

CHAOS PERSISTS IN LARGE-SCALE MULTI-AGENT LEARNING DESPITE ADAPTIVE LEARNING RATES

EMMANOUIL V. VLATAKIS-GKARAGKOUNIS[§], LAMPROS FLOKAS^{*}, AND GEORGIOS PILIOURAS^{◇,*}

ABSTRACT. Multi-agent learning is intrinsically harder, more unstable and unpredictable than single agent optimization. For this reason, numerous specialized heuristics and techniques have been designed towards the goal of achieving convergence to equilibria in self-play. One such celebrated approach is the use of dynamically adaptive learning rates. Although such techniques are known to allow for improved convergence guarantees in small games, it has been much harder to analyze them in more relevant settings with large populations of agents. These settings are particularly hard as recent work has established that learning with fixed rates will become chaotic given large enough populations [9, 17]. In this work, we show that chaos persists in large population congestion games despite using adaptive learning rates even for the ubiquitous Multiplicative Weight Updates algorithm, even in the presence of only two strategies. At a technical level, due to the non-autonomous nature of the system, our approach goes beyond conventional period-three techniques [29] by studying fundamental properties of the dynamics including invariant sets, volume expansion and turbulent sets. We complement our theoretical insights with experiments showcasing that slight variations to system parameters lead to a wide variety of unpredictable behaviors.

CONTENTS

1. Introduction	2
2. Problem setup and preliminaries	4
2.1. Dynamical systems & Li Yorke chaos	4
2.2. Linear congestion games	5
2.3. Multiplicative Weights Update	5
2.4. Dynamic Learning Rate Model	5
3. Refinements in Fixed Learning rates	6
4. Chaos in Uniformly Convergent Non-autonomous Dynamical Systems	7
5. Turbulent Sets and Chaos in MWU map	9
6. Conclusion	11
Acknowledgments	11
References	11
Appendix A. Omitted Proofs of Section 3	14
A.1. The order of $\{f_{\min}(a), x_{\max}(a), b, x_{\min}(a), f_{\max}(a)\}$ in the fixed learning rate regime.	14
A.2. Forward invariant, Perpetual & Absorbing sets for fixed learning rate.	15
A.3. Two period trajectories in MWU maps.	16
A.4. Volume Expansion & Instability of mixed equilibrium $x^* = b$	19
Appendix B. Omitted Proofs of Section 4	21
B.1. Forward & Absorbing Set in Dynamic Learning Rates.	21
B.2. Uniform Convergence of Césaro means & Learning Rate.	22
B.3. Uniform Convergence of varying step-size MWU map for uniform convergent update step rule.	23
Appendix C. Omitted Proofs of Section 5	26
C.1. Decaying Volume Turbulent sets in Fixed Learning Rate Regime	26
C.2. Tracking properties, make chaos explicit via symbolic dynamics	27
Appendix D. Bifurcation Plots	29

[§] UNIVERSITY OF CALIFORNIA, BERKELEY. EMVLATAKIS@BERKELEY.EDU

^{*} COLUMBIA UNIVERSITY. LAMFLOKAS@CS.COLUMBIA.EDU

[◇] SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN.

^{*} DEEPMIND. GEORGIOS.PILIOURAS@GMAIL.COM

Key words and phrases. Nash equilibrium; Li-Yorke Chaos; Adaptive MWU; Turbulence sets.

1. INTRODUCTION

Arguably one of the most thorny problems on the intersection of learning and games is the development of simple and practical algorithms that converge to Nash equilibria. The problem in its full generality is known to be intractable [18, 23], however, recent developments in Machine Learning have lead to a revived interest in the problem even in special classes of games. Unfortunately, the additional scrutiny has only helped crystallize the severity of the problem at hand through a diverse set of non-convergence, instability results even in well motivated special classes of games [2, 4–7, 19, 22, 24, 28, 31]. Worse yet, standard online learning meta-algorithms, such as Multiplicative Weights Updates (MWU), have been shown to exhibit chaos [9, 13–15, 17, 33], which has been shown to be rather common in more complex games, e.g., as we increase the number of agents [9, 17, 21, 37]. Given this proliferation of negative results, what other approaches are left to explore?

An interesting hint can be found in some of the earliest AI work on the subject. [42] established arguably one of the first non-convergence results in the area, showing that gradient dynamics do not suffice to achieve point-wise convergence even in the trivial case of two agent, two strategy games. On the positive side, when those dynamics are non-convergent, time-average convergence results can be established. Building up on their work, [12] showed that a modification of these standard dynamics where the agents can dynamically update their step-size based on payoff cues from their environment suffices to stabilize these dynamics in all 2×2 games. Informally, the specific heuristic, Win or Learn Fast (WoLF), has the agents increase their learning rate when they are "losing" in an effort to escape from a non-promising region of the state space. Unfortunately, WoLF requires that each agent has knowledge to a Nash equilibrium strategy for themselves, which means it can only be applied in small games. On the positive side, this first result has led to several other more similar heuristics, with promising results in small games [1, 8, 10, 11, 25, 26]. While staying in the realm of small games, such heuristics have been shown to robustly improve the stability of other popular learning heuristics such as Proximal Policy Optimization [35, 38]. On the negative side, even for slightly larger games most positive results are largely empirical in nature and effectively very little is known about the behavior of such techniques in games with many agents. These works raise our key motivating question:

Does the simple idea of judiciously increasing the learning rate to "escape" barren regions of the state space scale to games with a large number of agents? If not, what type of formal instability results can be established?

Our approach. We focus on one of the most well studied class of large population games, non-atomic congestion games [32, 36] where the agents learn using the ubiquitous MWU update rule [3, 20, 30]. Furthermore, given the recent works of [9, 17], we focus on games with exactly two strategies/routes and linear cost functions, since such settings are already sufficiently hard for dynamics. Specifically, given fixed, non-adaptive learning rates, MWU dynamics typically will bifurcate to chaotic behavior for a large enough population size. Critically, however, these previous results are based on the celebrated "Period three implies chaos" methodology pioneered by Li-Yorke [29], which is only applicable in autonomous, i.e., time-invariant systems. In contrast, understanding the emergent behavior in our case will require totally different techniques.

Our model will allow for a wide range of dynamically adaptive learning rates. In particular, instead of having only two learning rates, fast and slow, as, e.g., in [12], we will allow for a continuous range of learning rates. Moreover, the decision to increase/decrease the learning rate will be driven by a regret-like measure [39]. Intuitively, each agent compares the historical time-average performance of both routes available to them. If both routes appear very similar to each other then the agents favor large learning rates so as to reach sufficiently different configurations where they can hopefully find

informative payoff signals to exploit. On the other hand, if one route is significantly better than the other then the agents will take smaller but non-vanishing step sizes so as to take advantage of such opportunities without overcorrecting.

Our main result. We prove that our class of dynamics exhibits Li-Yorke chaotic behavior, implying the existence of an uncountable set of initial conditions that become “scrambled” by the dynamics (Theorem 4). Formally, for any two initial conditions from this set, the trajectories of the game dynamics come arbitrarily close to each other and move away from each other infinitely often.




Figure 1: An assortment of bifurcation diagrams for different adaptive learning rate schemes. Larger volatility (larger limit step-size) leads to more prominently unpredictable, chaotic behavior (dense limit-sets for large fraction of the parameter space).

Our techniques. Despite the intense recent interest in developing formal arguments about chaos in game settings [9, 17, 27, 34] proving chaos in non-autonomous dynamical systems is a challenging endeavor, as the departure from one-dimensional autonomous systems gives up from many established equivalences among different notions of chaos, such as volume expansion, topological entropy, and arbitrarily long periodic points.

One possible approach to addressing this challenge would be to first establish that the sequence of employed maps converges uniformly to a fixed chaotic map, and then transfer the resulting scrambled set of the limit system to the non-autonomous one.

Although this property is not true for our system, it suffices to use a slightly relaxed version of it. We establish uniform convergence guarantees for initialization sets that are confined to the interior

of the strategy space (Lemma 10), while tolerating the discrepancy between the varying-step and the corresponding MWU map in the limit. To confirm the existence of an uncountable scrambled set, we propose a novel connection between the behavior of MWU maps with varying learning rates and symbolic dynamics. Essentially, our proof strategy entails assigning every binary sequence to a different initialization such that scrambled binary strings correspond to scrambled initializations of the our non-autonomous dynamical system (Proof of Theorem 4).

The first part of our analysis (Section 3) involves enhancing previously existing results for the case of fixed learning rate dynamics. Specifically, we present an explicit construction of a *perpetual set* $\mathcal{F}(a)$, i.e., a forward invariant set where our dynamical system is surjective, $f_a(\mathcal{F}(a)) = \mathcal{F}(a)$, where f_a represents the MWU system with fixed learning rate a . (Lemma 3). Next we further show that, even in the presence of agent volatility due to their high learning rates, if their initial strategies lie within the interior of the strategy space, their strategies will eventually be absorbed into the perpetual set we have constructed (Lemma 4).

In the case of fixed learning rate, we are able to quantify an even stronger volume expansion result. Specifically, we show that the image of any arbitrarily small neighborhood of the mixed equilibrium of our game will converge exactly to the perpetual set we have constructed (Theorem 2).

Turning our attention to the adaptively dynamic learning rate (Section 4), we show that thanks to $\mathcal{F}(a)$, while learning rate may vary over iterations and initializations, there exists a set Δ –the closure of all perpetual sets $\mathcal{F}(a_n)$ –to which our dynamics eventually will be absorbed (Lemmas 5,6), provided the initialization set lies in the interior of solution space. Leveraging this property, we show that despite the volatility due to agents employing potentially high learning rates, our average regret-like concepts converge for nearly all initial conditions (Lemmas 7,8,9). It is noteworthy that this result requires a substantial degree of additional technicality in contrast to its analogue for the fixed rate regime [17].

Before developing our machinery for the symbolic dynamics reduction in Section 5, our proof strategy includes an additional step of demonstrating that our system does not collapse early in any fixed point or subspace of strategies that lacks volume expansion, as this would eliminate chaos. To achieve this, we rely on Theorem 3 which guarantees that almost all neighborhoods of mixed equilibrium will avoid this scenario by volume expanding and approximately covering $\mathcal{F}(\lim_{n \rightarrow \infty} a_n)$. Putting everything together our main result about chaotic limit behavior follows.

2. PROBLEM SETUP AND PRELIMINARIES

2.1. Dynamical systems & Li Yorke chaos. A discrete autonomous m -dimensional dynamical system is described by an equation $x_{n+1} = f(x_n)$ for a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^m$. In contrast a discrete m -dimensional non-autonomous dynamical system is described by an equation $x_{n+1} = f(x_n, n)$ for a function $f : \mathbb{R}^m \times \mathbb{N} \rightarrow \mathbb{R}^m$. In both of these cases we can view the $x_n(x_0)$ iterate as a function of the initialization x_0 . We thus term the sequence $\{x_n(x_0)\}_{n \in \mathbb{N}}$ the orbit of the dynamical system of x_0 .

We call two initializations p, q of a dynamical system scrambled if $\limsup_{n \rightarrow \infty} \|x_n(p) - x_n(q)\| > 0$ and $\liminf_{n \rightarrow \infty} \|x_n(p) - x_n(q)\| = 0$. Intuitively the orbits of p, q continuously alternate between moving apart and arbitrarily close respectively. A set S is scrambled if every pair $p, q \in S$ is scrambled. We are now ready to define the notion of Li Yorke chaos for both autonomous and non-autonomous dynamical systems.

Definition 1. An autonomous/non-autonomous dynamical system is called Li Yorke chaotic if there is an uncountably infinite set S that is scrambled.

A famous result for the case of one dimensional autonomous dynamical systems provides a simple and easily verified sufficient condition for chaos based on periodic orbits:

Theorem 1 ([29]). *Let J be an interval and $f : J \rightarrow J$ be a continuous map of a dynamical system. If f has a period-three orbit, then it is Li-Yorke chaotic.*

2.2. Linear congestion games. In this work we consider the case of non-atomic two strategy linear congestion games. In a non-atomic congestion game, there is a continuum of agents each of which controls an infinitesimal fraction of the total flow N . In each iteration, a fraction of agents x choose the first action and $1 - x$ the second one, suffering costs $c_1(x) = \gamma Nx$ and $c_2(1 - x) = \delta N(1 - x)$ respectively where $\gamma, \delta > 0$ are the linear cost coefficients of each action respectively.

2.3. Multiplicative Weights Update. The dynamics of the MWU meta-algorithm correspond to an autonomous dynamical system. Given the common learning rate of all agents $\eta > 0$, the MWU update is as follows

$$x_{n+1} = \frac{x_n e^{-\eta c_1(x_n)}}{x_n e^{-\eta c_1(x_n)} + (1 - x_n) e^{-\eta c_2(1-x_n)}} = \frac{x_n}{x_n + (1 - x_n) e^{-\eta(c_2(1-x_n) - c_1(x_n))}}$$

We can further simplify the MWU update rule by making the following changes of variables

$$a = \frac{N\eta}{\delta + \gamma} \quad b = \frac{\delta}{\delta + \gamma}$$

where a corresponds to a normalized learning rate that accounts for the scale of c_1 and c_2 and $(b, 1 - b)$ is the Nash equilibrium of the game. The simplified update rule can be expressed via a parametric map f over a and b that captures the whole class of games we are interested in

$$x_{n+1} = f(x_n, a, b) = \frac{x_n}{x_n + (1 - x_n) \exp(a(x_n - b))}. \quad (1)$$

When a, b are fixed we may skip them from the notation of MWU map and use just $f(x)$.

2.4. Dynamic Learning Rate Model. Here we extend the MWU meta-algorithm to adaptive learning rates η . The basis of our adapting dynamics is a notion of pseudo-regret that the agents suffer when choosing the first action over the second one. The definition captures the average difference of costs between the two actions weighted by the learning rate at each iteration η_i :

$$r_{n+1} = \frac{1}{n+1} \sum_{i=0}^n \eta_i [c_1(x_i) - c_2(1 - x_i)].$$

Based on this calculation of pseudo-regret, the agents may seek to adapt their learning rates η_i in order to balance between exploration and exploitation. When $|r_n|$ is big then it is clear than one action dominates over the other so the agents seek to move in the direction of improving costs but hopefully not too aggressively so as not to overshoot, akin to a gradient-like dynamics. In contrast, when $|r_n|$ is close to zero the average penalty of switching between actions is small since they suffer the same weighted cost in average. Thus there is an opportunity for the agents to increase their learning rate to encourage exploration while the cost differences are low. To hedge between exploration and exploitation as above, agents can pick appropriate function $h : \mathbb{R} \rightarrow \mathbb{R}^+$ that peaks at 0 and setting $\eta_n = h(r_n)$.

This dynamics of adapting η_i to the pseudo-regret can be implemented as a 2-dimensional non-autonomous dynamical system with x_n and r_n as its state variables:

$$x_{n+1} = f\left(x_n, \frac{Ng(r_n)}{\gamma + \delta}, b\right) \quad r_{n+1} = \frac{nr_n + h(r_n)(c_1(x_n) - c_2(1 - x_n))}{n+1}. \quad (2)$$

The system is non-autonomous in the 2-dimensional space as the update rule of r_n depends on n . We can always view d -dimensional non-autonomous systems as autonomous systems in $d + 1$ dimensions by treating n as an additional dimension. However, this does not really simplify things as [Theorem 1](#) does not apply beyond a single dimension.

For simplicity, we rewrite the dynamics in terms of the normalized learning rates a_i . Let a_i be the normalized learning rates corresponding to η_i based on [Eq. \(1\)](#). We can then rewrite r_{n+1} as

$$r_{n+1} = \frac{1}{n+1} \sum_{i=0}^n a_i [x_i - b].$$

Similarly, given that η_i and α_i are equivalent up to game specific scale parameters, we can always find a g such that $a_n = g(r_n)$. Instead of working in terms of the cumbersome dynamical system update rule Eq. (2), we will be performing our analysis in terms of these simpler equations

$$x_{n+1} = f(x_n, a_n, b) \quad a_{n+1} = g\left(\frac{\sum_{i=0}^n a_i(x_i - b)}{n+1}\right).$$

In the following sections we will make some assumptions on g . We will chose g to be a continuous function bounded in $[a_{\min}, a_{\max}] \subset \mathbb{R}^+$. We will also assume that the initialization a_0 is also a constant in $[a_{\min}, a_{\max}]$.

Throughout this work we will think of x_n and a_n as functions of the initialization x_0, a_0 . Thus, the complete notation of the n -th iterate would be $x_n(x_0, a_0)$ and $a_n(x_0, a_0)$, however, when it does not hinder understanding we will drop the explicit description of all the dependencies.

3. REFINEMENTS IN FIXED LEARNING RATES




Figure 2: The order of $f_{\min}(a)$, $x_{\max}(a)$, b , $x_{\min}(a)$, $f_{\max}(a)$ in x -axis.

We start our analysis by providing a refined analysis in the regime of fixed learning rate as in [17]. In the following we will examine conditions of such systems under the assumption of $a > 4$. Let us study the local minima and maxima of f :

$$x_{\min/\max}(a) = \frac{1}{2} \pm \sqrt{\frac{1}{4} - \frac{1}{a}}$$

Correspondingly we will denote $f_{\max}(a) = f(x_{\max}(a), a, b)$ and $f_{\min}(a) = f(x_{\min}(a), a, b)$. Our first observation is connected with the order of these values for high enough learning rate:

Lemma 1. *For every $b \in (0, 1)$, there is a a_b such that for all $a > a_b$*

$$f_{\min}(a) < x_{\max}(a) < b < x_{\min}(a) < f_{\max}(a)$$

and $f_{\min}(a)$ is decreasing and $f_{\max}(a)$ is increasing.

Additionally, by construction, we can show that the interval $\mathcal{F}(a) = [f_{\min}(a), f_{\max}(a)]$ consists a forward invariant set for our dynamical system. In other words, if the MWU map with fixed learning rate a starts in a state that is within I , it will remain in that set for all future times.

Lemma 2. *For every $b \in (0, 1)$ there is a s_b such that $\mathcal{F}(a)$ is forward invariant for all $a > s_b$, i.e.*

$$x \in \mathcal{F}(a) \Rightarrow f(x, a, b) \in \mathcal{F}(a).$$

We can also prove that f is surjective on $\mathcal{F}(a)$ for high enough learning rates.

Lemma 3. *For $a > a_b$, f is surjective on $\mathcal{F}(a)$, i.e. $f(\mathcal{F}(a), a, b) = \mathcal{F}(a)$.*

When $\mathcal{F}(a)$ is both surjective and forward invariant as is the case for $a > a_b, s_b$, we will call $\mathcal{F}(a)$ a *perpetual* set. Actually, $\mathcal{F}(a)$ is also an absorbing set for all $x \in (0, 1)$, in the following sense:

Lemma 4. *Let $a > s_b$, $x_0 \in [\gamma, \zeta] \Rightarrow \exists n_0 : \forall n \geq n_0$ s.t. $x_n \in \mathcal{F}(a)$, for every $[\gamma, \zeta] \subset (0, 1)$.*

We now turn to the chaotic properties of f . Given that period-three orbits do not carry over to the dynamic learning rate case, we choose to study an alternative property of chaotic maps, namely volume expansion. In autonomous maps, the existence of period-three orbits together with [Theorem 1](#) imply that there is an initialization set whose volume under the dynamics expands quickly. Intuitively, the more volume this initialization set covers the more chaotic the dynamic is. While [Theorem 1](#) guarantees the existence of volume expansion, it does not quantify how much volume it eventually covers. We prove that any interval around b is sufficient to eventually cover $\mathcal{F}(a)$:

Theorem 2. *For every $b \in (0, 1)$, there is a v_b such that for all $a > v_b$ and any interval $[\gamma, \delta] : \{b\} \subset [\gamma, \delta] \subset (0, 1)$, it holds that*

$$\exists n_0 : \forall n \geq n_0 \quad f^n([\gamma, \delta], a, b) = \mathcal{F}(a).$$

The proof of this theorem, which notably has not been established in any prior work, relies on a novel argument that is based on the monotonicity of f on $\mathcal{D} = [x_{\max}(a), x_{\min}(a)]$, the lack of period-2 trajectories in $\mathcal{D}$ and the instability of fixed point $x = b$.

The key takeaway for the next section is that we can focus our efforts on analyzing the behavior of the non autonomous system in the interior of $(0, 1)$ with special attention to neighborhoods of b that alone can exhibit chaotic behavior. This is especially important because as we will see our notion of pseudo-regret converges uniformly in closed intervals in the interior of $(0, 1)$ but not on $[0, 1]$.

4. CHAOS IN UNIFORMLY CONVERGENT NON-AUTONOMOUS DYNAMICAL SYSTEMS

Turning our attention to the dynamic learning rate setting, we will refer to $x_n(x_0)$ as the n -th iterate given the initialization x_0 , and $a_n(x_0)$ the learning rate at the same iteration, respectively.

Leveraging the existence of the perpetual set $\mathcal{F}(a)$ in the fixed rate case, we can construct a forward invariant absorbing set. This set will be crucial in proving the Li-Yorke chaotic behavior for the non-autonomous case. Intuitively, although the learning rate is varying both among different initializations and iterations, there exists a set which corresponds to the closure of all perpetual sets $\{\mathcal{F}(a_n)\}$, to which our dynamics are always absorbed. This motivates the definition of set Δ as

$$\Delta = \left[\min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a), \max_{a \in [a_{\min}, a_{\max}]} f_{\max}(a) \right]$$

where $a_{\min}$ and $a_{\max}$ are the $\min_{x_0 \in [0, 1]} g(x_0)$ and $\max_{x_0 \in [0, 1]} g(x_0)$ correspondingly.

Lemma 5 (Forward Invariance Property). *For all $a_{\min} > s_b$ we get that Δ is forward invariant, i.e.*

$$x_n(x_0) \in \Delta \implies x_{n+1}(x_0) \in \Delta.$$

More interestingly, we show that eventually, any subinterval of $(0, 1)$ will be absorbed within Δ .

Lemma 6 (Absorption Property). *For all $a_{\min} > s_b$ and $[\gamma, \zeta] \subset (0, 1)$, there is an $n_0 \geq 0$ so that*

$$\forall n \geq n_0 \quad x_n([\gamma, \zeta]) \subseteq \Delta.$$

The key observation for the proof of the aforementioned lemmas is the fact that at any iteration $x_n(x_0)$ will always come closer to the perpetual absorbing set $\mathcal{F}(a_n(x_0))$ and thus to its closure Δ .

An immediate consequence of this lemma is the following corollary which will be dominant element for the uniform convergence both of our proposed *pseudo-regret* notion and the Césaro mean of the iterations. Analytically, for a given subinterval $[\gamma, \zeta]$ of $(0, 1)$, we can always choose some $\delta > 0$ such

that $\Delta \cup [\gamma, \zeta] \subseteq (\delta, 1 - \delta)$. Again such choice of δ is possible because the attraction to Δ is not eventual phenomenon – at each iteration, $\text{dist}(x_n(x_0), \Delta)$ will be either zero or always decreasing. This argumentation is expressed by the following statement:

Corollary 1. *Let $a_{\min} > s_b$, then for every $[\gamma, \zeta] \subset (0, 1)$, there is an $\delta > 0$ so that*

$$\forall n \geq 0 \quad x_n([\gamma, \zeta]) \subseteq (\delta, 1 - \delta).$$

Having established Corollary 1, we are now in a position to present our claims regarding the average convergence for any initialization in $(0, 1)$. Notably, while the individual iterates of the system, as we shall demonstrate in the ensuing section, exhibit chaotic characteristics, the iterate averages display stabilization over time.

Corollary 1 indicates that for any $x_0 \in (0, 1)$, $x_n(x_0)$ will remain bounded away from the endpoints 0, 1. We now seek to connect how the boundness away from 0, 1 affects the average iterate of a trajectory. By induction on $x_{n+1} = f(x_n, a_n, b)$ we can prove the following

$$x_n = \frac{x_0}{x_0 + (1 - x_0) \exp\left(\sum_{i=0}^{n-1} a_i(x_i - b)\right)}. \quad (3)$$

Observe that if we choose a $x_0 \in (0, 1)$, then x_n is bounded away from 0, 1 if and only if $\sum_{i=0}^{n-1} a_i(x_i - b)$ is bounded. Specifically we can prove:

Lemma 7. *Let $a_{\min} > s_b$ and $x_0 \in (0, 1)$, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n a_i(x_0)(x_i(x_0) - b) = 0.$$

Upon examination of the proof, a salient conclusion is that as x_0 moves farther away from the absorbing set Δ and closer to the endpoints 0, 1, the aforementioned convergence rate becomes increasingly slow, and fails in the case that $x_0 \in \{0, 1\}$, precluding a blanket result for the closed interval $[0, 1]$. Conversely, if we restrict our attention to an arbitrary subinterval $[\gamma, \zeta]$, bounded away from the endpoints, we can always exploit the minimum convergence rate of this interval to derive a uniform convergence bound for the learning rate

Lemma 8. *Let $a_{\min} > s_b$. The sequence of functions $a_n(\cdot)$ is converging uniformly to the constant function $a^*(\cdot) = g(0)$ in every interval $[\gamma, \zeta] \subset (0, 1)$.*

Having established that $a_n(x_0)$ converges to $a^* = g(0)$ for any $x_0 \in (0, 1)$, we are able to strengthen Lemma 7, demonstrating that its unweighted version of the Césaro mean (i.e. the average iterate in the limit of infinite time) also converges to b .

Lemma 9. *Let $a_{\min} > s_b$ and $x_0 \in (0, 1)$, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n x_i(x_0) = b.$$

It is noteworthy that, despite this being equivalent to the guarantee provided in prior work, e.g. [17], for the fixed learning rate case, the machinery developed here for the adaptive rate regime has been more complex and involved.

The next lemma will be essential in our proofs of chaos and examines the relationship between $(x_n(x_0), a_n(x_0))$ and the MWU map with a fixed learning rate $f(x_n(x_0), a^*, b)$.

Lemma 10. *Let $a_{\min} > s_b$. For every $k, \epsilon > 0$ and $[\gamma, \zeta] \subset (0, 1)$*

$$\exists n_0 : \forall n \geq n_0 \quad \max_{x_0 \in [\gamma, \zeta]} |x_{n+k}(x_0) - f^k(x_n(x_0), a^*, b)| \leq \epsilon.$$

We can intuitively think of this result as a form of uniform convergence result for the sequence of MWU maps with varying learning rates to the fixed learning rate map. We establish this result by showing that for small changes in the learning rate, i.e., $a \in (a_n(x_0) - \tilde{\epsilon}, a_n(x_0) + \tilde{\epsilon})$, the sensitivity of $f(x_n(x_0), a, b)$ is independent of the choice of x_0 . It is thus not sufficient to prove that $f(x, a, b)$ is continuous in a . Instead we use the Mean Value Theorem to argue that $f(x, a, b)$ is Lipschitz continuous in a with a Lipschitz constant independent of x . It is crucial to note that this uniform convergence result holds only in intervals $[\gamma, \zeta] \subset (0, 1)$ because only then does the learning rate a_n converge uniformly to $a(\cdot) = a^*$.

We now prove our first volume expansion lemma for the dynamic learning rate case, namely that neighborhoods of b within $\mathcal{D}(a_{\min})$ eventually expand to at least $\mathcal{F}(a_{\min})$. The intuition behind this result is based on the monotonicity property of $\mathcal{F}(a)$. Specifically, for high enough learning rates $a_1 \leq a_2$, we have that $\mathcal{F}(a_1) \subseteq \mathcal{F}(a_2)$.

Lemma 11. *For every $b \in (0, 1)$, there is an z_b such that for all $a_{\min} > z_b$ it holds that for every $[\gamma, \delta]$ such that $\{b\} \subset [\gamma, \delta] \subseteq \mathcal{D}(a_{\min}) = [x_{\max}(a_{\min}), x_{\min}(a_{\min})]$*

$$\exists n_0 : \forall n \geq n_0 \quad x_n([\gamma, \delta]) \supseteq \mathcal{F}(a_{\min}).$$

Remark 1. It is worth mentioning that the series of Lemmas 6, 5, 7 and 9, 11 of this section –pertaining to forward invariance, absorption, volume expansion, and convergence of Césaro means –hold actually regardless of the choice of the update rule and are of independent interest.

For the case of the fixed learning rate, Theorem 2 makes a more refined prediction compared to Lemma 11 as the former provides an equality. In the dynamic learning rate case there is a gap between the eventual image upper bound, Δ and the volume expansion lower bound $\mathcal{F}(a_{\min})$. In order to close this gap we will make use of the fact that the learning rate converges to $g(0)$.

Building upon the full range of the developed machinery of this section, we will strengthen the above volume expansion result by showing that $a_{\min}$ could be actually substituted by any $\tilde{a} = a^* - \epsilon$ for any sufficient small $\epsilon > 0$ such that $a^* - \epsilon > a_{\min} > z_b$. Our first observation is that thanks to the uniform convergence of $a_n \rightarrow a^*$ in $[\gamma, \delta] \subset (0, 1)$ there exists a $n^\dagger$ such that

$$\forall n > n^\dagger \quad \forall x \in [\gamma, \delta] \quad a_n(x) > a^* - \epsilon = \tilde{a}$$

This strengthens the previous volume expansion result to at least the set $\mathcal{F}(\tilde{a})$ for any $\tilde{a}$ close to a^* .

Theorem 3. *For $a_{\min} > z_b$ and for any sufficient small $\epsilon > 0$ such that $a^* - \epsilon > a_{\min} > z_b$ we have that for all $[\gamma, \delta]$ such that $\{b\} \subset [\gamma, \delta] \subset (0, 1)$, it holds that*

$$\exists n_0 : \forall n \geq n_0 \quad x_n([\gamma, \delta]) \supseteq \mathcal{F}(a^* - \epsilon).$$

In a completely similar fashion we can show that the absorbing set Δ can also be refined. For a given interval $[\gamma, \delta] \subset (0, 1)$, as the learning rates converge to a^* , the trajectories will tend to be absorbed by $\mathcal{F}(a^*)$. Thus the dynamic learning rate volume expansion behavior matches fixed learning rate case in the long run.

5. TURBULENT SETS AND CHAOS IN MWU MAP

The roadmap of this section is our construction of the turbulent sets and their connection with the symbolic dynamics in order to prove Li-Yorke chaos in our non-autonomous dynamical system. It should be noticed that our novel approach is a major departure from the standard techniques of 3-period orbit arguments that have been extensively used in the case of fixed learning rates.

We start with some useful definition for our reduction.

Definition 2 ([40]). A continuous map $f : I \rightarrow I$ is called *turbulent* if there exist compact subintervals J, K with at most one common point such that $J \cup K \subseteq f(J) \cap J(K)$. Additionally, the map is called *strictly turbulent* if the subintervals can be chosen to be disjoint. Finally, the corresponding set J, K are called turbulent sets.




Figure 3: Illustration of the inductive construction of V_a^k sets. In order to construct an exponential decaying sequence, we choose as $V_a^1 = \arg \min_{Z \in \{Z_1, Z_2\}} \text{diam}(Z) \leq \text{diam}(V_a^0 = K_a)$.

Delving into the proof of [29], it becomes clear that the chaotic map $f(x, a, b)$ has a periodic orbit of period 3 that exists within the interior of $\mathcal{F}(a)$. We take advantage of this property to demonstrate that $f^2(\cdot, a, b)$ is in fact a strictly turbulent map.

Lemma 12. *For every $b \in (0, 1)$, there is a u_b such that for all $a > u_b$, there exist closed and disjoint intervals K_a and J_a in the interior of $\mathcal{F}(a)$ such that $f^2(K_a, a, b)$ and $f^2(J_a, a, b)$ are neighborhoods of $K_a \cup J_a$.*

We begin by examining the properties of a fixed learning rate MWU map and its ability to create exponential decaying volume (length) turbulent sets. Our analysis centers around the key observation that since $f^2(K_a)$ covers $K_a \cup J_a$, then there exist necessarily, by continuity of f^2 , at least two distinct subintervals, Z_1, Z_2 , within K_a whose f^2 -image is precisely K_a and J_a , respectively. By repeatedly applying this principle, we demonstrate that the f^4 -image of these subintervals cover again $K_a \cup J_a$, and through induction, we can extend this property to higher compositions of f .

Lemma 13. *For every $a > u_b$, there exist closed intervals $V_a^k \subseteq K_a$ and $U_a^k \subseteq J_a$ such that*

$$\lim_{k \rightarrow \infty} \text{diam}(V_a^k) = \lim_{k \rightarrow \infty} \text{diam}(U_a^k) = 0$$

and for every $k \geq 0$ it holds that $f^{2k+2}(V_a^k, a, b)$ and $f^{2k+2}(U_a^k, a, b)$ are neighborhoods of $K_a \cup J_a$.

The following lemma plays an essential role in our symbolic dynamic proof of the Li-Yorke chaotic behavior. It allows us to construct a scrambled set of initial conditions through a set of abstract symbolic orbits.

On a technical level, we utilize a range of machinery developed in this paper to prove this lemma. Specifically, we use Theorem 3 to firstly describe turbulent sets $J_{a^*} \cup K_{a^*}$ that lie within the interior of $\mathcal{F}(a^*)$ and, for small enough ϵ , within $\mathcal{F}(a^* - \epsilon)$, and secondly to ensure that the non-autonomous dynamical system covers $J_{a^*} \cup K_{a^*}$. Furthermore, to extend the implication of Lemma 13 for the turbulent map $f^2(\cdot, a^*, b)$ to the non-autonomous system x_n map for the sets J_{a^*}, K_{a^*} , we employ the uniform convergence guarantee provided by Lemma 10, which controls the discrepancy between $x_{n+2}(S)$ and $f^2(x_n(S), a^*, b)$ for any subinterval S of J_{a^*}, K_{a^*} .

Lemma 14 (Tracking Lemma). *If $b \in (0, 1) \setminus \{\frac{1}{2}\}$, there exists a d_b such that if $a_{\min} > d_b$, we can construct an increasing sequence n_i with the following properties. For every sequence of intervals A_i with $A_i = V_{a^*}^i$ or $A_i = U_{a^*}^i$, there exists a $x_0 \in [0, 1]$ such that for all $i \geq 0$ it holds that $x_{n_i}(x_0) \in A_i$.*

It is important to note that Lemma 14 ensures the ability to construct the same sequence of n_i for any distinct sequence of A_i . Our approach to demonstrate the Li-Yorke chaotic behavior through symbolic dynamics is summarized in the following high-level steps:

- Assume that we can construct an uncountable set $\mathcal{S}$ of scrambled infinite length binary sequences, i.e., for every pair of sequences $\sigma, \tau \in \mathcal{S}$ there exist (i) an infinite length subsequence $(k_i)_{i \in \mathbb{N}}$ where the two sequences differ, i.e., $\sigma_{k_i} \neq \tau_{k_i}$ and (ii) an infinite length subsequence $(\ell_i)_{i \in \mathbb{N}}$ where the two sequences are equal, i.e., $\sigma_{\ell_i} = \tau_{\ell_i}$.
- For each element of $\sigma \in \mathcal{S}$ we construct a sequence of sets $\mathcal{A}_\sigma = (A_i)_{i \in \mathbb{N}}$ as follows: If the k -th place element is 0 we use $V_{a^*}^k$, whereas if it is 1 we pick $U_{a^*}^k$. We now apply Lemma 14 for each of the sequence of sets to get a corresponding initialization x_0^σ . We call this set of initializations Q .
- Since every pair of strings σ, τ is scrambled, we know that we can construct two infinite subsequences $(\mu_i)_{i \in \mathbb{N}}, (\nu_i)_{i \in \mathbb{N}}$ such that $\{x_{\mu_i}(x_0^\sigma)\}, \{x_{\mu_i}(x_0^\tau)\}$ belong to the same turbulent sets and $\{x_{\nu_i}(x_0^\sigma)\}, \{x_{\nu_i}(x_0^\tau)\}$ belong to the disjoint turbulent ones. Therefore, we can show that for every pair $\sigma, \tau \in Q$: (i) $\liminf \|\{x_n(x_0^\sigma) - x_n(x_0^\tau)\}\| = 0$ (ii) $\limsup \|\{x_n(x_0^\sigma) - x_n(x_0^\tau)\}\| > 0$

Formalizing the outlined proof sketch, our final result follows:

Theorem 4. *If $b \in (0, 1) \setminus \{\frac{1}{2}\}$, there exists a d_b such that if $a_{\min} > d_b$, the dynamics of Equation 2 are Li-Yorke chaotic.*

In this work we have focused on the dynamics of Eq. (2). But our proof strategy can be readily generalized to any rule $a_n(\cdot)$ as long as it uniformly converges to a sufficiently high constant rate.

Corollary 2. *Let $a_n(\cdot)$ be a sequence of maps uniformly converging in $[0, 1]$ to $a^* > d_b$. Then the resulting dynamic learning rate system is Li-Yorke chaotic.*

6. CONCLUSION

We have formally analyzed and established chaotic behavior for a class of multi-agent learning systems with a heuristically updated, variable learning rate. At the technical crux of all prior formal analysis of Li-Yorke chaos in games (e.g., [9, 16, 17, 27, 34]) lied the celebrated methodology based on period three orbits [29], which is only applicable in autonomous, i.e., time-invariant systems. In contrast, we had to delve deeper into the geometry and structural properties of these dynamics, which itself evolve with time showing that formal analysis of chaos is still possible. This opens the possibility of extending prior results to more realistic time-varying models.

ACKNOWLEDGMENTS

Emmanouil V. Vlatakis-Gkaragkounis is grateful for financial support by the Post-Doctoral FODSI-Simons Fellowship, Pancretan Association of America and Simons Collaboration on Algorithms and Geometry and Onassis Doctoral Fellowship. This research/project is also supported in part by the National Research Foundation, Singapore and DSO National Laboratories under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), NRF 2018 Fellowship NRF-NRFF2018-07, NRF2019-NRF-ANR095 ALIAS grant, grant PIESGP-AI-2020-01, AME Programmatic Fund (Grant No.A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) and Provost's Chair Professorship grant RGEPPV2101

REFERENCES

- [1] Abdallah, S. and Lesser, V. A multiagent reinforcement learning algorithm with non-linear dynamics. *Journal of Artificial Intelligence Research*, 33:521–549, 2008.
- [2] Andrade, G. P., Frongillo, R., and Piliouras, G. Learning in matrix games can be arbitrarily complex. In Belkin, M. and Kpotufe, S. (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 159–185. PMLR, 15–19 Aug 2021. URL <https://proceedings.mlr.press/v134/andrade21a.html>.

- [3] 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.
- [4] Bailey, J. P. and Piliouras, G. Multiplicative weights update in zero-sum games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 321–338. ACM, 2018.
- [5] 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.
- [6] Bailey, J. P., Gidel, G., and Piliouras, G. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *Conference on Learning Theory*, pp. 391–407. PMLR, 2020.
- [7] Balduzzi, D., Racanière, S., Martens, J., Foerster, J. N., Tuyls, K., and Graepel, T. The mechanics of n-player differentiable games. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 363–372. PMLR, 2018. URL <http://proceedings.mlr.press/v80/balduzzi18a.html>.
- [8] Banerjee, B. and Peng, J. Adaptive policy gradient in multiagent learning. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pp. 686–692, 2003.
- [9] Bielawski, J., Chotibut, T., Falniowski, F., Kosiorowski, G., Misiurewicz, M., and Piliouras, G. Follow-the-regularized-leader routes to chaos in routing games. In *ICML*, 2021.
- [10] Bloembergen, D., Tuyls, K., Hennes, D., and Kaisers, M. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research*, 53:659–697, 2015.
- [11] Bowling, M. Convergence and no-regret in multiagent learning. *Advances in neural information processing systems*, 17, 2004.
- [12] Bowling, M. and Veloso, M. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136(2):215–250, 2002.
- [13] Cheung, Y. K. and Piliouras, G. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *Conference on Learning Theory*, pp. 807–834. PMLR, 2019.
- [14] Cheung, Y. K. and Piliouras, G. Chaos, extremism and optimism: Volume analysis of learning in games. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [15] Cheung, Y. K. and Tao, Y. Chaos of learning beyond zero-sum and coordination via game decompositions. In *International Conference on Learning Representations (ICLR)*, 2021.
- [16] Cheung, Y. K., Leonardos, S., and Piliouras, G. Learning in markets: Greed leads to chaos but following the price is right. 2021.
- [17] Chotibut, T., Falniowski, F., Misiurewicz, M., and Piliouras, G. The route to chaos in routing games: When is price of anarchy too optimistic? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [18] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. The complexity of computing a nash equilibrium. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’06, pp. 71–78, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595931341. doi: 10.1145/1132516.1132527. URL <https://doi.org/10.1145/1132516.1132527>.
- [19] Flokas, L., Vlatakis-Gkaragkounis, E.-V., Lianas, T., Mertikopoulos, P., and Piliouras, G. No-regret learning and mixed nash equilibria: They do not mix. In *NeurIPS*, 2020.
- [20] Freund, Y. and Schapire, R. E. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- [21] 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.
- [22] Giannou, A., Vlatakis-Gkaragkounis, E. V., and Mertikopoulos, P. Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information. In Belkin, M.

- and Kpotufe, S. (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 2147–2148. PMLR, 15–19 Aug 2021. URL <https://proceedings.mlr.press/v134/giannou21a.html>.
- [23] Hart, S. and Mas-Colell, A. Uncoupled dynamics do not lead to nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
- [24] Hsieh, Y.-P., Mertikopoulos, P., and Cevher, V. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. In *International Conference on Machine Learning*, pp. 4337–4348. PMLR, 2021.
- [25] Kaisers, M., Tuyls, K., Parsons, S., and Thuijsman, F. An evolutionary model of multi-agent learning with a varying exploration rate. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 1255–1256, 2009.
- [26] Leonardos, S. and Piliouras, G. Exploration-exploitation in multi-agent learning: Catastrophe theory meets game theory. *Artificial Intelligence*, 304:103653, 2022.
- [27] Leonardos, S., Monnot, B., Reijsbergen, D., Skoulakis, E., and Piliouras, G. Dynamical analysis of the eip-1559 ethereum fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pp. 114–126, 2021.
- [28] Letcher, A. On the impossibility of global convergence in multi-loss optimization, 2021.
- [29] Li, T.-Y. and Yorke, J. A. Period three implies chaos. *The American Mathematical Monthly*, 82(10):985–992, 1975.
- [30] Littlestone, N. and Warmuth, M. K. The weighted majority algorithm. *Inf. Comput.*, 108(2): 212–261, February 1994. ISSN 0890-5401. doi: 10.1006/inco.1994.1009. URL <http://dx.doi.org/10.1006/inco.1994.1009>.
- [31] Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pp. 2703–2717, USA, 2018. Society for Industrial and Applied Mathematics. ISBN 9781611975031.
- [32] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007. ISBN 0521872820.
- [33] Palaiojanos, 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.
- [34] Piliouras, G. and Yu, F.-Y. Multi-agent performative prediction: From global stability and optimality to chaos. *arXiv preprint arXiv:2201.10483*, 2022.
- [35] Ratcliffe, D. S., Hofmann, K., and Devlin, S. Win or learn fast proximal policy optimisation. In *2019 IEEE Conference on Games (CoG)*, pp. 1–4. IEEE, 2019.
- [36] Roughgarden, T. and Tardos, É. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.
- [37] Sanders, J. B. T., Farmer, J. D., and Galla, T. The prevalence of chaotic dynamics in games with many players. *Scientific reports*, 8(1):1–13, 2018.
- [38] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [39] Shalev-Shwartz, S. et al. Online learning and online convex optimization. *Foundations and trends in Machine Learning*, 4(2):107–194, 2011.
- [40] Shi, Y. and Yu, P. Study on chaos induced by turbulent maps in noncompact sets. *Chaos, Solitons & Fractals*, 28(5):1165–1180, 2006.
- [41] Shub, M. *Global Stability of Dynamical Systems*. Springer-Verlag, 1987.
- [42] Singh, S., Kearns, M. J., and Mansour, Y. Nash convergence of gradient dynamics in general-sum games. In *UAI*, pp. 541–548. Citeseer, 2000.

APPENDIX A. OMITTED PROOFS OF SECTION 3

In this section we will focus on the case where the dynamical system has a fixed learning rate.

$$x_{n+1} = f(x_n, a, b) = \frac{x_n}{x_n + (1 - x_n) \exp(a(x_n - b))}$$

The derivative of f for a fixed a and b is

$$f'(x, a, b) = \frac{(ax^2 - ax + 1) \exp(a(x - b))}{(x + (1 - x) \exp(a(x - b)))^2}$$

Critical/stationary points of f are solutions of $ax^2 - ax + 1 = 0$. By taking the determinant, we get than for $0 < a \leq 4$, there is no solution so f is increasing and no chaos can exist, instead the system converges to equilibrium. We will thus require $a > 4$ to enable chaotic behaviour in the system. Let us study the local minima and maxima of f :

$$x_{\max}(a) = \frac{1}{2} - \sqrt{\frac{1}{4} - \frac{1}{a}} \quad x_{\min}(a) = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{1}{a}}$$

A.1. The order of $\{f_{\min}(a), x_{\max}(a), b, x_{\min}(a), f_{\max}(a)\}$ in the fixed learning rate regime.

The following facts establish the preliminary necessary observation to prove structural Lemma 1. More precisely, we can show the following straightforward facts:

Fact A.1. *For every $b \in (0, 1)$, there is a $s_b > 4$ such that for all $a > s_b$*

$$x_{\max}(a) < b < x_{\min}(a).$$

Proof. We have $x_{\max}(a)$ is decreasing and $\lim_{a \rightarrow \infty} x_{\max}(a) = 0$. Symmetrically we have $x_{\min}(a)$ is increasing and $\lim_{a \rightarrow \infty} x_{\min}(a) = 1$. The fact follows immediately. ■

Fact A.2. *If $a > 4$, then $f(x_{\min}(a), a, b) = f_{\min}(a) < f(x_{\max}(a), a, b) = f_{\max}(a)$.*

Proof. Since the function is decreasing after the local maximum and there is no stationary point until the local minimum, the local minimum has smaller value. ■

Fact A.3. *If $a > s_b$, then $f_{\min}(a) < b < f_{\max}(a)$.*

Proof. For $a > s_b$ we know that $f_{\min}(a) < f(b, a, b) = b$ since f is decreasing in $[b, x_{\min}(a)]$. Symmetrically, $f_{\max}(a) > f(b, a, b) = b$ since again f is decreasing in $[x_{\max}(a), b]$. ■

Having presented the aforementioned intuitive facts, we are ready to prove Lemma 1

Lemma A.1. *[Restate Lemma 1] For every $b \in (0, 1)$, there is a a_b such that for all $a > a_b$*

$$f_{\min}(a) < x_{\max}(a) < b < x_{\min}(a) < f_{\max}(a)$$

and $f_{\min}(a)$ is decreasing and $f_{\max}(a)$ is increasing.

Proof. $f_{\min}(a) < x_{\max}(a)$ is equivalent to

$$x_{\min}(a) - x_{\min}(a)x_{\max}(a) - x_{\max}^2(a) \exp[\alpha(x_{\min}(a) - b)] < 0$$

Observe that the first two terms are bounded and that for the third term we have

$$\lim_{a \rightarrow \infty} x_{\max}^2(a) \exp[\alpha(x_{\min}(a) - b)] = \infty$$

So we can pick an a_b large enough so that $a > a_b$ implies that the inequality holds. The case of $x_{\min}(a) < f_{\max}(a)$ is symmetric. Moving on to the monotonicity of $f_{\min}(a)$, to avoid overloading notation let us call x_1 the x argument of f and x_2 its a argument

$$\frac{\partial f_{\min}(a)}{\partial a} = \frac{\partial f(x_{\min}(a), a, b)}{\partial x_1} \frac{\partial x_{\min}(a)}{\partial a} + \frac{\partial f(x_{\min}(a), a, b)}{\partial x_2}$$

Observe that since $x_{\min}(a)$ is a local minimum, we have that

$$\frac{\partial f(x_{\min}(a), a, b)}{\partial x_1} = 0$$

Additionally we can pick a_b large enough so that $a > a_b$ implies $x_{\min}(a) > b$. Let us write $f(x_{\min}(a), a, b)$

$$f(x_{\min}(a), a, b) = \frac{x_{\min}(a)}{x_{\min}(a) + (1 - x_{\min}(a)) \exp(a(x_{\min}(a) - b))}.$$

Treating $x_{\min}(a)$ as a constant independent of a , the exponential in the denominator is an increasing function of a so $f(x_{\min}(a), a, b)$ is decreasing with respect to a . This makes

$$\forall a > a_b \quad \frac{\partial f(x_{\min}(a), a, b)}{\partial x_2} < 0$$

As a result we have that

$$\forall a > a_b \quad \frac{\partial f_{\min}(a)}{\partial a} < 0 \tag{A.1}$$

and $f_{\min}(a)$ is decreasing. The case of $f_{\max}(a)$ is symmetric. We take the maximum of all the required a_b to get the result. ■

A.2. Forward invariant, Perpetual & Absorbing sets for fixed learning rate. Having settled the order among $\{f_{\min}(a), x_{\max}(a), b, x_{\min}(a), f_{\max}(a)\}$ for high enough learning rates, we are ready to prove the (i) forward invariant, (ii) perpetual and (iii) absorbing property of $\mathcal{F}(a) = [f_{\min}(a), f_{\max}(a)]$.

For the sake of readability, we recall first the formal definitions of these properties:

- (1) **A forward invariant set:** a set of states such that if the system starts in any state in the set, it will remain in the set for all future time.
- (2) **A perpetual set:** a special case of a forward invariant set, whose image consists itself.
- (3) **An (global/local) absorbing set:** a forward invariant set that also includes (globally/locally) all possible future states of the system.

Lemma A.2. *[Restated Lemma 2] For every $b \in (0, 1)$ there is a s_b such that $\mathcal{F}(a)$ is forward invariant for all $a > s_b$, i.e.*

$$x \in \mathcal{F}(a) \Rightarrow f(x, a, b) \in \mathcal{F}(a).$$

Proof. By continuity of f , we only need to consider four points to determine the image of $[f_{\min}(a), f_{\max}(a)]$: $f_{\min}(a)$, $f_{\max}(a)$ as well as $x_{\max}(a)$, $x_{\min}(a)$. Since $f(x, a, b) \geq x$ in $(0, b)$ and $f_{\max}(a)$ is the maximum in this interval, we know that

$$f_{\max}(a) \geq f(f_{\min}(a), a, b) \geq f_{\min}(a)$$

Since $f(x, a, b) \leq x$ in $(b, 1)$ and $f_{\min}(a)$ is the minimum in this interval, we know that

$$f_{\max}(a) \geq f(f_{\max}(a), a, b) \geq f_{\min}(a)$$

Of course the images of the local optima $x_{\max}(a)$, $x_{\min}(a)$ trivially belong to $[f_{\min}(a), f_{\max}(a)]$. ■

But even points that are outside of this interval are monotonically attracted to it without overshooting. We have the following lemma that shows this monotonic attracting property

Lemma A.3. *Let $a > s_b$, then*

$$\begin{aligned} x \in (0, f_{\min}(a)) &\implies f(x, a, b) \in (x, f_{\max}(a)] \\ x \in (f_{\max}(a), 1) &\implies f(x, a, b) \in [f_{\min}(a), x). \end{aligned}$$

Proof. We will prove the first one, the second one is entirely symmetric. If $x < f_{\min}(a)$ then $x < b$ and thus $f(x, a, b) > x$ and $f(x, a, b) \leq f_{\max}(a)$. The first implication follows immediately. ■

Consequently, we have the following lemma:

Lemma A.4 (Restated Lemma 3). *For $a > a_b$, $\mathcal{F}(a)$ is surjective, i.e. $f(\mathcal{F}(a), a, b) = \mathcal{F}(a)$.*

Proof. We already know that the right hand side is a super set of the left hand side by Lemma A.2. But by Lemma A.1 we have that $x_{\max}(a)$ and $x_{\min}(a)$ belong to $[f_{\min}(a), f_{\max}(a)]$. So the left hand is a super set of the left hand side. Thus the result follows. $\blacksquare$

More generally we can prove the following absorbing condition:

Lemma A.5 (Restated Lemma 4). *Let $a > s_b$, $x_0 \in [\gamma, \zeta] \Rightarrow \exists n_0 : \forall n \geq n_0$ s.t. $x_n \in \mathcal{F}(a)$, for every $[\gamma, \zeta] \subset (0, 1)$.*

Proof. If $x_0 \in [f_{\min}(a), f_{\max}(a)]$ then the results follows trivially for $n_0 = 0$. Let us take the case

$$x_0 \in [\gamma, f_{\min}(a)) = \Delta_0$$

We know that in Δ_0 we have that $f(x, a, b) > x$. We can define a uniform bound on their difference

$$\min_{x \in \Delta_0} [f(x, a, b) - x] = \xi > 0$$

We now know that

$$x_1 \in [\gamma + \xi, f_{\max}(a))$$

If $x_1 \in \Delta$ then the result follows trivially for $n_0 = 1$. Otherwise we have that

$$x_1 \in [\gamma + \xi, f_{\min}(a)) \subset \Delta_0$$

Applying recursively, either there is a n_0 such that $x_{n_0} \in \Delta$ and the theorem follows trivially or for all n we have that x_n stays in a subset of Δ_0

$$x_n \in [\gamma + n\xi, f_{\min}(a)) \subset \Delta_0$$

But this is impossible since there is a $n_0 > 0$ such that

$$\forall n \geq n_0 : \gamma + n\xi > f_{\min}(a)$$

We are now left with the symmetric case of

$$x_0 \in (f_{\max}(a), \zeta]$$

which we can handle just like above. $\blacksquare$

A.3. Two period trajectories in MWU maps. We start with a fundamental observation from Calculus of continuous injective function

Claim 1. *Let g be a continuous decreasing function on a closed interval I and that there exists a $x_0 \in I$ such that $\forall n : g^n(x_0) \in I$. Then either $g^n(x_0)$ converges to a fixed point x^* or to a 2-period trajectory.*

Proof. Let us take $g^{n_j}(x_0)$ a subsequence that converges to $\liminf_{n \rightarrow \infty} g^n(x_0)$.

$$\limsup_{n \rightarrow \infty} g^n(x_0) \geq \lim_{j \rightarrow \infty} g^{n_j+1}(x_0) = g(\liminf_{n \rightarrow \infty} g^n(x_0))$$

Now let us take $g^{n_k}(x_0)$ a subsequence that converges to $\limsup_{n \rightarrow \infty} g^n(x_0)$. Since g is decreasing and thus invertible in I we know that $g^{n_k-1}(x_0)$ also converges

$$\limsup_{n \rightarrow \infty} g^n(x_0) = g(\lim_{k \rightarrow \infty} g^{n_k-1}(x_0)) \leq g(\liminf_{n \rightarrow \infty} g^n(x_0))$$

The two steps clearly imply that

$$\limsup_{n \rightarrow \infty} g^n(x_0) = g(\liminf_{n \rightarrow \infty} g^n(x_0))$$

Symmetrically, with the same arguments we have that

$$\liminf_{n \rightarrow \infty} g^n(x_0) = g(\limsup_{n \rightarrow \infty} g^n(x_0))$$

Thus, if we denote $x_1 = g(\limsup_{n \rightarrow \infty} g^n(x_0)) = \liminf_{n \rightarrow \infty} g^n(x_0)$ and $x_2 = g(\liminf_{n \rightarrow \infty} g^n(x_0)) = \limsup_{n \rightarrow \infty} g^n(x_0)$ then either $x_1 = x_2 = x^*$, which consists a fixed point for map g $-g(x^*) = x^*$, or $x_1 \neq x_2$ consist a 2-period trajectory $\{x_1 = g(x_2), x_2 = g(x_1)\}$. $\blacksquare$

Interestingly, if we restrict our attention to the interval $\mathcal{D} = [x_{\max}(a), x_{\min}(a)]$, we can show that there is no 2-period trajectory:

Lemma A.6. *For every $b \in (0, 1)$, there is a ℓ_b such that for all $a > \ell_b$, there is no period two trajectory for which both endpoints belong to $[x_{\max}(a), x_{\min}(a)]$.*

Proof. Endpoints of period two trajectories satisfy the equation $f^2(x, a, b) = x$

$$f^2(x, a, b) = \frac{x}{x + (1-x)\exp(a(x + f(x, a, b) - 2b))} = x$$

Ignoring $x = 0$ and $x = 1$, which are merely fixed points, this is equivalent to

$$x + f(x, a, b) = 2b$$

After some manipulation, the formula above is equivalent to

$$\gamma_{a,b}(x) = (2b - x - 1)x + (2b - x)(1 - x)\exp(a(x - b)) = 0$$

We take the first and second and third derivative of this function

$$\begin{aligned}\gamma'_{a,b}(x) &= 2b - 2x - 1 + \exp(a(x - b))[(2x - 2b - 1) + a(2b - x)(1 - x)] \\ \gamma''_{a,b}(x) &= -2 + \exp(a(x - b)) [a^2(2b - x)(1 - x) + 2a(2x - 2b - 1) + 2] \\ \gamma'''_{a,b}(x) &= a \exp(a(x - b)) [a^2(2b - x)(1 - x) + 3a(2x - 2b - 1) + 6]\end{aligned}$$

Let us define the following finite quantity

$$\mu_b = \frac{\max_{x \in [0, b]} 3(2b + 1 - 2x)}{\min_{x \in [0, b]} (2b - x)(1 - x)}$$

For $a > \mu_b$, we have that $\gamma'''_{a,b}(x) > 0$ in $[0, b]$. Moving on to $\gamma''_{a,b}(x)$, it is increasing in $[0, b]$ and

$$\gamma''_{a,b}(0) = -2 + \exp(-ab) [2ba^2 - 2a(2b - 1) + 2] \quad \gamma''_{a,b}(b) = a^2b(1 - b) - 2a$$

Clearly we have

$$\lim_{a \rightarrow \infty} \gamma''_{a,b}(0) = -2 \quad \lim_{a \rightarrow \infty} \gamma''_{a,b}(b) = \infty$$

We can thus pick an $\kappa_b > \mu_b$ such that for all $a > \kappa_b$

$$\gamma''_{a,b}(0) < 0 \quad \gamma''_{a,b}(b) > 0.$$

Thus $\gamma''_{a,b}(x)$ has exactly one root in $[0, b]$ given its monotonicity. Moving on to $\gamma'_{a,b}(x)$, it starts of as decreasing and moves to increasing in $[0, b]$ with

$$\begin{aligned}\gamma'_{a,b}(0) &= 2b - 1 + \exp(-ab) [(-2b - 1) + 2ab] \\ \gamma'_{a,b}(b) &= -2 + ab(1 - b)\end{aligned}$$

To continue our analysis we will study the following cases: $b < \frac{1}{2}$, $b = \frac{1}{2}$, $b > \frac{1}{2}$.

Case: $b < \frac{1}{2}$. For the first case

$$\lim_{a \rightarrow \infty} \gamma'_{a,b}(0) = 2b - 1 < 0 \quad \lim_{a \rightarrow \infty} \gamma'_{a,b}(b) = \infty$$

Thus we can pick a $\nu_b > \kappa_b$ such that for $a > \nu_b$, $\gamma'_{a,b}(0) < 0$ and $\gamma'_{a,b}(b) > 0$. Since $\gamma'_{a,b}(x)$ starts decreasing and moves to increasing, it has exactly one root. Moving on to $\gamma_{a,b}(x)$ we have

$$\gamma_{a,b}(0) = 2b \exp(-ab) > 0 \quad \gamma_{a,b}(b) = 0$$

Given that $\gamma_{a,b}(x)$ starts decreasing and moves to increasing in $[0, b]$, it can have up to two roots in $[0, b]$, one of which is b . We can observe that

$$\lim_{a \rightarrow \infty} \gamma_{a,b} \left(\frac{b}{2} \right) = \left(\frac{3b}{2} - 1 \right) \left(\frac{b}{2} \right) < 0.$$

We can pick a $\xi_b > \nu_b$ such that for all $a > \xi_b$ it holds that $\gamma_{a,b} \left(\frac{b}{2} \right) < 0$. For $a > \xi_b$ we have that $\gamma_{a,b}(x)$ has exactly one root in $[0, b)$ that is located in $[0, \frac{b}{2}]$.

Case: $b = \frac{1}{2}$. For the second case we have that

$$\gamma'_{a,b}(0) = \exp \left(-\frac{a}{2} \right) (a - 2) \quad \lim_{a \rightarrow \infty} \gamma'_{a,b} \left(\frac{b}{2} \right) = b - 1 < 0 \quad \lim_{a \rightarrow \infty} \gamma'_{a,b}(b) = \infty$$

So we can pick a $\nu_b > \kappa_b$ such that for $a > \nu_b$, $\gamma'_{a,b}(0) > 0$ and $\gamma'_{a,b} \left(\frac{b}{2} \right) < 0$ and $\gamma'_{a,b}(b) > 0$. Since $\gamma'_{a,b}(x)$ starts decreasing and moves to increasing, it has exactly two roots. Moving on to $\gamma_{a,b}(x)$ we have

$$\gamma_{a,b}(0) = 2b \exp(-ab) > 0 \quad \gamma_{a,b}(b) = 0$$

In $[0, b]$, we have that $\gamma_{a,b}(x)$ starts increasing and positive, then switches to decreasing and then to increasing. In the first section it cannot have any root. In the second section it can have at most one root and in the third section it has exactly one root b . Just like above we can observe that

$$\lim_{a \rightarrow \infty} \gamma_{a,b} \left(\frac{b}{2} \right) = \left(\frac{3}{4} - 1 \right) \left(\frac{1}{4} \right) < 0.$$

Following the same steps as above, we can pick a ξ_b such that for $a > \xi_b$ we have that $\gamma_{a,b}(x)$ has exactly one root in $[0, b)$ that is located in $[0, \frac{b}{2}]$.

Case: $b > \frac{1}{2}$. For the last case we have that

$$\lim_{a \rightarrow \infty} \gamma'_{a,b}(0) = 2b - 1 > 0 \quad \lim_{a \rightarrow \infty} \gamma'_{a,b} \left(\frac{b}{2} \right) = b - 1 < 0 \quad \lim_{a \rightarrow \infty} \gamma'_{a,b}(b) = \infty$$

Just like before we can pick a $\nu_b > \kappa_b$ such that for $a > \nu_b$, $\gamma'_{a,b}(0) > 0$ and $\gamma'_{a,b} \left(\frac{b}{2} \right) < 0$ and $\gamma'_{a,b}(b) > 0$. Since $\gamma'_{a,b}(x)$ starts decreasing and moves to increasing, it has exactly two roots. Moving on to $\gamma_{a,b}(x)$

$$\gamma_{a,b}(0) = 2b \exp(-ab) > 0 \quad \gamma_{a,b}(b) = 0$$

In $[0, b]$, we have that $\gamma_{a,b}(x)$ starts increasing and positive, then switches to decreasing and then to increasing. In the first section it cannot have any root. In the second section it can have at most one root and in the third section it has exactly one root b . We can observe that

$$\gamma_{a,b} \left(\frac{3b-1}{2} \right) = \left(\frac{b-1}{2} \right) \left(\frac{3b-1}{2} \right) < 0.$$

Following the same steps as above, we can pick a ξ_b such that for $a > \xi_b$ we have that $\gamma_{a,b}(x)$ has exactly one root in $[0, b)$ that is located in $[0, \frac{3b-1}{2}]$.

Solution pairs. In all cases, we have identified an ξ_b such that for $a > \xi_b$ we have that $\gamma_{a,b}(x)$ has exactly one root in $[0, b)$. If $\gamma_{a,b}(x) = 0$, then we know that $\gamma_{a,b}(f(x, a, b)) = 0$. Since $x + f(x, a, b) = 2b$, we know that $f(x, a, b)$ needs to be in $(b, 2b]$. Symmetrically, any root of $\gamma_{a,b}(x) = 0$ with $x > b$ can only form a periodic trajectory with an $f(x, a, b) = 2b - x < b$ that is also satisfies $\gamma_{a,b}(f(x, a, b)) = 0$. But since there is only one root in $[0, b)$ and points cannot participate in multiple periodic trajectories, we have a unique solution for $\gamma_{a,b}(x) = 0$ in $(b, 1)$. Let us define $x_l(a) < b < x_r(a)$ the points of the unique two-periodic trajectory as functions of a . These functions are bounded and thus they need to have at least one limit point. We can use the following equations

$$\lim_{a \rightarrow \infty} \gamma_{a,b}(x_l(a)) = \lim_{a \rightarrow \infty} \gamma_{a,b}(x_r(a)) = 0$$

to derive the properties of these limit points. In all three cases above, we proved that the solution $x_l(a)$ is bounded away from b for $a > \xi_b$. As a result all limit points must satisfy

$$(2b - x - 1)x = 0 \implies x = 0 \text{ or } x = 2b - 1.$$

Similarly, since $x_l(a) + x_r(a) = 2b$ and $x_l(a)$ is bounded away from b , it must be the case that $x_r(a)$ is bounded away from b . As such all limit points of $x_r(a)$ must satisfy

$$(2b - x)(1 - x) \implies x = 2b \text{ or } x = 1.$$

Observe that the limit points of $x_l(a)$ and $x_r(a)$ must come in pairs that sum to $2b$ as $x_l(a) + x_r(a) = 2b$. We will once again do a case by case study. For $b < \frac{1}{2}$, the only viable pair is $(0, 2b)$ since $2b - 1 < 0$. For $b = \frac{1}{2}$, there is only one pair $(0, 1)$. For the case of $b > \frac{1}{2}$, the only viable pair is $(2b - 1, 1)$ since $2b > 1$. In all of the cases, the limit points of $x_r(a)$ and $x_l(a)$ are unique and thus $x_r(a)$ and $x_l(a)$ converge.

Convergence rate. We are now ready to argue why at least one of $x_r(a)$ and $x_l(a)$ do not belong in $[x_{\max}(a), x_{\min}(a)]$ for sufficiently large a . Once again we will do a case by case analysis on b . For $b < \frac{1}{2}$, we will argue that there is a ℓ_b such that for $a > \ell_b$ we have that $x_l(a) < x_{\max}(a)$. We know that $x_l(a) = f(x_r(a), a, b)$. Thus $x_l(a) < x_{\max}(a)$ is equivalent to

$$x_r(a) - x_{\max}(a)x_r(a) - x_{\max}(a)(1 - x_r(a)) \exp[\alpha(x_r(a) - b)] < 0$$

Observe that the first two terms are bounded but for the third term we have

$$\lim_{a \rightarrow \infty} x_{\max}(a)(1 - x_r(a)) \exp[\alpha(x_r(a) - b)] = \infty$$

given that $\lim_{a \rightarrow \infty} x_r(a) = 2b$ and the exponential goes to ∞ much faster than $x_{\max}(a)$ goes to 0. Thus we can choose a ℓ_b such that for $a > \ell_b$ we have $x_l(a) < x_{\max}(a)$. For the case of $b > \frac{1}{2}$, we use the same arguments to prove that there is an ℓ_b such that for $a > \ell_b$ we have $x_r(a) > x_{\min}(a)$. For the case of $b = \frac{1}{2}$, we will argue that there is a ℓ_b such that for $a > \ell_b$ we have that $x_l(a) < x_{\max}(a)$. We need study the convergence rate of $x_r(a)$ to 1. We have that

$$\gamma_{a, \frac{1}{2}}(x_r(a)) = 0 \implies x_r(a) = (1 - x_r(a)) \exp\left[\frac{\alpha}{2}\left(x_r(a) - \frac{1}{2}\right)\right]$$

We can apply the same argument as in the $b < \frac{1}{2}$ case and use the equation above to prove that

$$\lim_{a \rightarrow \infty} x_{\max}(a)(1 - x_r(a)) \exp\left[\alpha\left(x_r(a) - \frac{1}{2}\right)\right] = \lim_{a \rightarrow \infty} x_{\max}(a)x_r(a) \exp\left[\frac{\alpha}{2}\left(x_r(a) - \frac{1}{2}\right)\right] = \infty$$

because $\lim_{a \rightarrow \infty} x_r(a) = 1$ and the exponential goes to ∞ much faster than $x_{\max}(a)$ goes to 0. The resulting ℓ_b in all cases satisfy the requirements of the theorem. $\blacksquare$

A.4. Volume Expansion & Instability of mixed equilibrium $x^* = b$.

Lemma A.7. *For every $b \in (0, 1)$, there is a k_b such that for $a > k_b$ it holds that there exists some neighborhood $\mathcal{N}_\delta = (b - \delta, b + \delta)$, such that almost all initializations from $\mathcal{N}_\delta$ do not converge to $x^* = b$.*

Proof. We can write down $f'(b, a, b)$ as

$$f'(b, a, b) = ab^2 - ab + 1$$

Since $\lim_{a \rightarrow \infty} f'(b, a, b) = -\infty$, thus there is a k_b such that for $a > k_b$ it holds that there exists some neighborhood $\mathcal{N}_\delta = (b - \delta, b + \delta)$ where $|f'(x, a, b)| > 1 \forall x \in \mathcal{N}_\delta$. Leveraging Unstable Manifold Theorem (See [41]), the result immediately follows. $\blacksquare$

We are now ready to prove that any neighborhood of $x^* = b$ lemma

Lemma A.8. *For every $b \in (0, 1)$, there is a z_b such that for all $a > z_b$.*

$$\{b\} \subset [\gamma, \delta] \subseteq [x_{\max}(a), x_{\min}(a)] \implies \exists n_0 : \quad \forall n \geq n_0 \quad f^n([\gamma, \delta], a, b) = [f_{\min}(a), f_{\max}(a)]$$

Proof. Let us take the following interval

$$\begin{aligned} R &= f([\gamma, b], a, b) \cap [b, \delta] \\ &= [f(b, a, b), f(\gamma, a, b)] \cap [b, \delta] \\ &= [b, f(\gamma, a, b)] \cap [b, \delta] \\ &= [b, \min\{f(\gamma, a, b), \delta\}] \end{aligned}$$

Let us pick an $x_0 \in R$. Let us assume that

$$x_{\max}(a) \leq f^n(x_0, a, b) \leq x_{\min}(a)$$

By Lemma A.7 and picking $a > k_b$ we can always choose a $x_0 \in R$ such that $f^n(x_0, a, b)$ does not converge to b . Since b is the unique fixed point of $f(x, a, b)$ in $[x_{\max}(a), x_{\min}(a)]$, we have that $f^n(x_0, a, b)$ does not converge at all and as a result

$$\liminf_{n \rightarrow \infty} f^n(x_0, a, b) \neq \limsup_{n \rightarrow \infty} f^n(x_0, a, b)$$

By Claim 1 on f , which is decreasing on $[x_{\max}(a), x_{\min}(a)]$, we have that the points above form a two-period trajectory of $f(x, a, b)$. By Lemma A.6, if we pick ℓ_b such that for $a > \ell_b$ no period two trajectory can exist inside $[x_{\max}(a), x_{\min}(a)]$ yielding a contradiction. As a result

$$\exists n^* : f^{n^*}(x_0, a, b) \notin [x_{\max}(a), x_{\min}(a)]$$

We will first study the case of $f^{n^*}(x_0, a, b) > x_{\min}(a)$. For this case, we have

$$f^{n^*}([\gamma, \delta], a, b) \supseteq f^{n^*}([b, x_0], a, b) \supseteq [b, x_{\min}(a)].$$

In the next iteration

$$f^{n^*+1}([\gamma, \delta], a, b) \supseteq [f_{\min}(a), b].$$

Observe though that because $x_0 \in R$, we can pick a $x_1 \in [\gamma, b]$ such that $f(x_1, a, b) = x_0$. We then have

$$f^{n^*+1}(x_1, a, b) = f^{n^*}(x_0, a, b) > x_{\min}(a)$$

This yields

$$f^{n^*+1}([\gamma, \delta], a, b) \supseteq [f_{\min}(a), x_{\min}(a)]$$

For $a > a_b$ by Lemma 1 $f_{\min}(a) < x_{\max}(a)$ and by Lemma 2 we have that $[f_{\min}(a), f_{\max}(a)]$ is forward invariant. As a result we have the following

$$[f_{\min}(a), f_{\max}(a)] \supseteq f^{n^*+2}([\gamma, \delta], a, b) \supseteq [f_{\min}(a), f_{\max}(a)]$$

By Lemma A.4 we know that $[f_{\min}(a), f_{\max}(a)]$ is perpetual. So we can choose $n_0 = n^* + 2$ to fulfill the requirements of the theorem. The case of $f^{n^*}(x_0, a, b) < x_{\max}(a)$ is symmetric. $\blacksquare$

We are ready now to prove the main volume expansion claims of the section.

Theorem A.1 (Restated Theorem 2). *For every $b \in (0, 1)$, there is a v_b such that for all $a > v_b$ and any interval $[\gamma, \delta] : \{b\} \subset [\gamma, \delta] \subset (0, 1)$, it holds that*

$$\exists n_0 : \forall n \geq n_0 \quad f^n([\gamma, \delta], a, b) = \mathcal{F}(a).$$

Proof. We can apply Lemma A.8 on $s = [\gamma, \delta] \cap [x_{\max}(a), x_{\min}(a)] \supset \{b\}$. Then there is a n_1^* such that

$$\forall n \geq n_1^* \quad f^n([\gamma, \delta], a, b) \supseteq f^n(s, a, b) = [f_{\min}(a), f_{\max}(a)]$$

But by Lemma 4, for $a > r_b$ there is a n_2^* such that

$$\forall n \geq n_2^* \quad [f_{\min}(a), f_{\max}(a)] \supseteq f^n([\gamma, \delta], a, b)$$

Taking $v_b = \max\{z_b, r_b\}$ and $n_0 = \max\{n_1^*, n_2^*\}$ satisfies the theorem requirements. $\blacksquare$

APPENDIX B. OMITTED PROOFS OF SECTION 4

B.1. Forward & Absorbing Set in Dynamic Learning Rates. In this first part of this section we prove the forward invariance and absorption property of set

$$\Delta = \left[\min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a), \max_{a \in [a_{\min}, a_{\max}]} f_{\max}(a) \right].$$

Lemma B.1 (Restated Lemma 5). *For all $a_{\min} > s_b$, then for every n it holds that Δ is forward invariant, i.e.*

$$x_n(x_0) \in \Delta \implies x_{n+1}(x_0) \in \Delta.$$

Proof. We have three cases to consider. The first one

$$x_n(x_0) \in [f_{\min}(a_n(x_0)), f_{\max}(a_n(x_0))] \implies x_{n+1}(x_0) \in [f_{\min}(a_n(x_0)), f_{\max}(a_n(x_0))]$$

and the result follows trivially because of the perpetual set $\mathcal{F}(a_n(x_0))$. For the second case we get

$$x_n(x_0) \in \left[\min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a), f_{\min}(a_n(x_0)) \right) \implies x_{n+1}(x_0) \in \left[\min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a), f_{\max}(a_n(x_0)) \right]$$

and thus the result follows for this one as well. The last case is

$$x_n(x_0) \in \left(f_{\max}(a_n(x_0)), \max_{a \in [a_{\min}, a_{\max}]} f_{\max}(a) \right] \implies x_{n+1}(x_0) \in \left[f_{\min}(a_n(x_0)), \max_{a \in [a_{\min}, a_{\max}]} f_{\max}(a) \right]$$

Clearly the result holds for all cases. $\blacksquare$

Lemma B.2 (Restated Lemma 6). *For all $a_{\min} > s_b$ and $[\gamma, \zeta] \subset (0, 1)$, there is an $n_0 \geq 0$ so that*

$$\forall n \geq n_0 \quad x_n([\gamma, \zeta]) \subseteq \Delta.$$

Proof. Let us define the three following sets

$$U_0 = \left[\gamma, \min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a) \right) \quad \Delta_0 = \Delta \cap [\gamma, \zeta] \quad V_0 = \left(\max_{a \in [a_{\min}, a_{\max}]} f_{\max}(a), \zeta \right]$$

It suffices to prove the theorem for each of them (if they are non-empty) because we can pick

$$n_0 = \max\{n_0(U_0), n_0(\Delta_0), n_0(V_0)\}$$

(excluding any empty sets) to satisfy the theorem. Based on Lemma 5, it is clear that $n_0(\Delta_0) = 0$. Moving on to the case of U_0 , we know that in U_0 we have that $f(x, a, b) > x$. We can define a uniform bound on their difference

$$\min_{x \in U_0} \min_{a \in [a_{\min}, a_{\max}]} [f(x, a, b) - x] = \xi > 0$$

We now know that

$$x_1(U_0) \in \left[\gamma + \xi, \max_{a \in [a_{\min}, a_{\max}]} f_{\max}(a) \right)$$

If $x_1(U_0) \subset \Delta$ then the result follows trivially for $n_0 = 1$. Otherwise we have that

$$x_1(U_0) \in \left[\gamma + \xi, \min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a) \right) \subset U_0$$

Applying recursively, either there is a n_0 such that $x_{n_0}(U_0) \subset \Delta$ and the theorem holds for U_0 or for all n we have that $x_n(U_0)$ stays in a subset of U_0

$$x_n(U_0) \in \left[\gamma + n\xi, \min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a) \right) \subset U_0$$

But this is impossible since there is a $n_0 > 0$ such that

$$\forall n \geq n_0 : \gamma + n\xi > \min_{a \in [a_{\min}, a_{\max}]} f_{\min}(a)$$

The case of V_0 can be handled symmetrically. $\blacksquare$

B.2. Uniform Convergence of Césaro means & Learning Rate. In this section, we prove the uniform convergence of learning rate, pseudo-regret and the average iteration while $n \rightarrow \infty$. We start with the asymptotic behavior of the introduced notion of pseudo-regret

Lemma B.3 (Restated Lemma 7). *Let $a_{\min} > s_b$ and $x_0 \in (0, 1)$, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n a_i(x_0)(x_i(x_0) - b) = 0.$$

Proof. Observe that $x_0 \in (0, 1)$, so by Lemma 1 there is a $\delta > 0$ such that $\delta < x_n < 1 - \delta$. Thus we have

$$\frac{x_0}{1-\delta} < x_0 + (1-x_0) \exp \left(\sum_{i=0}^{n-1} a_i(x_0)(x_i(x_0) - b) \right) < \frac{x_0}{\delta}$$

so we have that

$$\delta^2 < \frac{x_0 \delta}{1-\delta} < (1-x_0) \exp \left(\sum_{i=0}^{n-1} a_i(x_0)(x_i(x_0) - b) \right) < \frac{x_0(1-\delta)}{\delta} < \frac{1}{\delta}$$

Clearly we have that

$$\delta^2 < \exp \left(\sum_{i=0}^{n-1} a_i(x_0)(x_i(x_0) - b) \right) < \frac{1}{\delta^2}$$

By taking the logarithm and dividing by n we have

$$\frac{\ln(\delta^2)}{n} < \frac{1}{n} \sum_{i=0}^{n-1} a_i(x_0)(x_i(x_0) - b) < -\frac{\ln(\delta^2)}{n}$$

By taking the limit, the result follows easily. ■

We proceed now to the uniform convergence of the adaptively changing learning rate

Lemma B.4 (Restated Lemma 8). *Let $a_{\min} > s_b$. The sequence of functions $a_n(\cdot)$ is converging uniformly to the constant function $a^*(\cdot) = g(0)$ in every interval $[\gamma, \zeta] \subset (0, 1)$.*

Proof. By Lemma 1 $x_0 \in [\gamma, \zeta]$ there is a $\delta > 0$ such $\delta < x_n < 1 - \delta$. Following the steps of Lemma 7

$$\frac{\ln(\delta^2)}{n} < \frac{1}{n} \sum_{i=0}^{n-1} a_i(x_0)(x_i(x_0) - b) < -\frac{\ln(\delta^2)}{n}$$

Clearly this implies that for any $n > 0$

$$a_n(x_0) = g \left(\frac{1}{n} \sum_{i=0}^{n-1} a_i(x_0)(x_i(x_0) - b) \right) \in g \left(\left[\frac{\ln(\delta^2)}{n}, -\frac{\ln(\delta^2)}{n} \right] \right)$$

Since g is continuous, for every $\epsilon > 0$ there is a n_0 such that $\forall n \geq n_0$ we have that

$$|a_n(x_0) - g(0)| \leq \epsilon.$$

Uniform convergence follows immediately. ■

Next, we can prove the following uniform convergence result for the Césaro mean of the iterations of our non-autonomous dynamical system

Lemma B.5 (Restated Lemma 9). *Let $a_{\min} > s_b$ and $x_0 \in (0, 1)$, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n x_i(x_0) = b.$$

Proof. Let us define $\rho_n = a_n(x_0) - a^*$. Obviously $\lim_{t \rightarrow \infty} \rho_n = 0$. Then applying Lemma 7 we have

$$0 = \lim_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n a_i(x_0)(x_i(x_0) - b) = \lim_{n \rightarrow \infty} \left[\frac{a^*}{n+1} \sum_{i=0}^n (x_i(x_0) - b) + \frac{1}{n+1} \sum_{i=0}^n \rho_i(x_i(x_0) - b) \right]$$

It is easy to prove the following statement:

Fact B.1. *Let $\lim_{n \rightarrow \infty} \gamma_n \rightarrow \gamma$. Then $\lim_{n \rightarrow \infty} \frac{\sum_{i=0}^n \gamma_i}{n+1} = \gamma$.*

Clearly we have that $\lim_{n \rightarrow \infty} \rho_n(x_n(x_0) - b) = 0$ since $x_n(x_0)$ is bounded. And thus we obviously have

$$0 = \lim_{n \rightarrow \infty} \left[\frac{a^*}{n+1} \sum_{i=0}^n (x_i(x_0) - b) \right]$$

Given that a^* is positive, the theorem follows. $\blacksquare$

B.3. Uniform Convergence of varying step-size MWU map for uniform convergent update step rule.

Lemma B.6 (Restated Lemma 10). *Let $a_{\min} > r_b$. For every $k, \epsilon > 0$ and $[\gamma, \zeta] \subset (0, 1)$*

$$\exists n_0 : \forall n \geq n_0 \quad \max_{x_0 \in [\gamma, \zeta]} |x_{n+k}(x_0) - f^k(x_n(x_0), a^*, b)| \leq \epsilon.$$

Proof. For the sake of readability, we will present the case $k = 2$ and with a completely similar fashion we can prove the result for any iterate k .

Lemma B.7. *Let $a_{\min} > s_b$. For every $\epsilon > 0$ and $[\gamma, \zeta] \subset (0, 1)$*

$$\exists n_0 : \forall n \geq n_0 \quad \max_{x_0 \in [\gamma, \zeta]} |x_{n+2}(x_0) - f^2(x_n(x_0), a^*, b)| \leq \epsilon.$$

Proof. By Corollary 1 there is a $\delta > 0$ such that for any $x_0 \in [\gamma, \zeta]$, we have that $x_n(x_0) \in (\delta, 1 - \delta)$. Let us define the following function that takes an x and applies f with learning rates a_1 and a_2

$$q(x, a_1, a_2, b) = f(f(x, a_1, b), a_2, b)$$

With this definition in mind we have that

$$D(x_0) = x_{n+2}(x_0) - f^2(x_n(x_0), a^*, b) = q(x_n(x_0), a_n(x_0), a_{n+1}(x_0), b) - q(x_n(x_0), a^*, a^*, b)$$

By the mean value theorem we have that there exists an $c \in (0, 1)$ such that

$$(a_1, a_2) = c(a_n(x_0), a_{n+1}(x_0)) + (1 - c)(a^*, a^*)$$

and also

$$D(x_0) = \frac{\partial q(x_n(x_0), a_1, a_2, b)}{\partial a_1} (a_n(x_0) - a^*) + \frac{\partial q(x_n(x_0), a_1, a_2, b)}{\partial a_2} (a_{n+1}(x_0) - a^*)$$

The following value is finite

$$\phi = \max_{\substack{x \in (\delta, 1-\delta), \\ a_1, a_2 \in (a_{\min}, a_{\max})}} \sqrt{\left(\frac{\partial q(x, a_1, a_2, b)}{\partial a_1} \right)^2 + \left(\frac{\partial q(x, a_1, a_2, b)}{\partial a_2} \right)^2}.$$

Then we have

$$|x_{n+2}(x_0) - f^2(x_n(x_0), a^*, b)| \leq \phi \sqrt{(a_n(x_0) - a^*)^2 + (a_{n+1}(x_0) - a^*)^2}$$

By the uniform convergence of $a_n(x_0)$ and $a_{n+1}(x_0)$ to a^* , the result follows immediately. $\blacksquare$

Lemma B.8 (Restated Lemma 11). *For every $b \in (0, 1)$, there is an z_b such that for all $a_{\min} > z_b$ it holds that for every $[\gamma, \delta]$ such that $\{b\} \subset [\gamma, \delta] \subseteq \mathcal{D}(a_{\min}) = [x_{\max}(a_{\min}), x_{\min}(a_{\min})]$*

$$\exists n_0 : \forall n \geq n_0 \quad x_n([\gamma, \delta]) \supseteq \mathcal{F}(a_{\min}).$$

Proof. Our first observation is that for two $a_1 > a_2 > 0$

$$\begin{aligned} 1 > x > b &\implies f(x, a_1, b) < f(x, a_2, b) \\ 0 < x < b &\implies f(x, a_1, b) > f(x, a_2, b) \\ x = b &\implies f(x, a_1, b) = f(x, a_2, b) = b \end{aligned} \tag{B.1}$$

Let us define the following subset of $[\gamma, \delta]$

$$[\gamma', \delta'] = [\gamma, \min\{f(\gamma, a_{\min}, b), \delta\}]$$

Using the same arguments as in [Lemma A.8](#) we know that there is a $x' \in (b, \delta']$ and a minimal n^* such that

$$f^{n^*}(x', a_{\min}, b) \notin [x_{\max}(a_{\min}), x_{\min}(a_{\min})]$$

We are going to assume $f^{n^*}(x', a_{\min}, b) > x_{\min}(a_{\min})$. The case of $f^{n^*}(x', a_{\min}, b) < x_{\max}(a_{\min})$ is entirely symmetric. Since n^* is minimal we have that

$$\forall 0 \leq n < n^* \quad f^n(x', a_{\min}, b) \in [x_{\max}(a_{\min}), x_{\min}(a_{\min})]$$

As a result we have that

$$\begin{aligned} \forall 0 \leq 2n < n^* \quad f^{2n}([b, x'], a_{\min}, b) &= [b, f^{2n}(x', a_{\min}, b)] \\ \forall 0 < 2n + 1 < n^* \quad f^{2n+1}([b, x'], a_{\min}, b) &= [f^{2n+1}(x', a_{\min}, b), b] \end{aligned}$$

Using the above equations as well as [Equation \(B.1\)](#) we can recursively prove that

$$\forall 0 \leq n < n^* \quad x_n([b, x']) \supseteq f^n([b, x'], a_{\min}, b)$$

We can deduce that

$$x_{n^*}([b, x']) \supseteq [b, f^{n^*}(x', a_{\min}, b)] \supseteq [b, x_{\min}(a_{\min})]$$

With one more iteration we have

$$x_{n^*+1}([b, x']) \supseteq [f_{\min}(a_{\min}), b]$$

We can also pick a $x'' \in [\gamma', b]$ such that $f(x'', a_{\min}, b) = x'$. By construction, we know that $n^* + 1$ is the first iteration such that

$$f^{n^*+1}(x'', a_{\min}, b) \notin [x_{\max}(a_{\min}), x_{\min}(a_{\min})].$$

By similar arguments as above we can prove that

$$x_{n^*+1}([x'', b]) \supseteq [b, f^{n^*+1}(x'', a_{\min}, b)] = [b, f^{n^*}(x', a_{\min}, b)] \supseteq [b, x_{\min}(a_{\min})]$$

As a result we now have

$$x_{n^*+1}([x'', x']) \supseteq [f_{\min}(a_{\min}), x_{\min}(a_{\min})] \supseteq [x_{\max}(a_{\min}), x_{\min}(a_{\min})]$$

It follows directly that

$$x_{n^*+2}([\gamma, \delta]) \supseteq x_{n^*+2}([x'', x']) \supseteq [f_{\min}(a_{\min}), f_{\max}(a_{\min})] \supseteq [x_{\max}(a_{\min}), x_{\min}(a_{\min})]$$

Applying the above step recursively we get

$$\forall n \geq n^* + 2 \quad x_n([\gamma, \delta]) \supseteq [f_{\min}(a_{\min}), f_{\max}(a_{\min})]$$

so $n_0 = n^* + 2$ satisfies the requirements of the theorem. ■

Theorem B.1 (Restated Theorem 3). *For $a_{\min} > z_b$ and for any sufficient small $\epsilon > 0$ such that $a^* - \epsilon > a_{\min} > z_b$ we have that for all $[\gamma, \delta]$ such that $\{b\} \subset [\gamma, \delta] \subset (0, 1)$, it holds that*

$$\exists n_0 : \forall n \geq n_0 \quad x_n([\gamma, \delta]) \supseteq \mathcal{F}(a^* - \epsilon).$$

Proof. We will imitate the proof strategy of the above lemma by showing that $a_{\min}$ could be substitute by any $\tilde{a} = a^* - \epsilon$ for any sufficient small $\epsilon > 0$ such that $a^* - \epsilon > a_{\min} > z_b$. Our first observation is that thanks to the uniform convergence of $a_n \rightarrow a^*$ in $[\gamma, \delta] \subset [0, 1]$ there exists a $n^\dagger$ such that

$$\forall n \geq n^\dagger \quad a_n([\gamma, \delta]) \subseteq (a^* - \epsilon, a^* + \epsilon)$$

which implies that there exists a $n^\dagger$ such that

$$\forall n > n^\dagger \quad \forall x \in [\gamma, \delta] \quad a_n(x) > a^* - \epsilon = \tilde{a} \quad (\text{B.2})$$

Again notice that for any $a_n(x) > \tilde{a} > 0$

$$\begin{aligned} 1 > x > b &\implies f(x, a_n(x), b) < f(x, \tilde{a}, b) \\ 0 < x < b &\implies f(x, a_n(x), b) > f(x, \tilde{a}, b) \\ x = b &\implies f(x, a_n(x), b) = f(x, \tilde{a}, b) = b \end{aligned} \quad (\text{B.3})$$

Let us define the following intervals

$$\begin{aligned} [r, q] &= x_{n^\dagger}([\gamma, \delta]) \cap [x_{\max}(\tilde{a}), x_{\min}(\tilde{a})] \supset \{b\} \\ [\gamma', \delta'] &= [r, \min\{f(r, \tilde{a}, b), q\}] \end{aligned}$$

Using the same arguments as in [Lemma A.8](#) we know that there is a $x' \in (b, \delta']$ and a minimal n^* such that

$$f^{n^*}(x', \tilde{a}, b) \notin [x_{\max}(\tilde{a}), x_{\min}(\tilde{a})]$$

We are going to assume $f^{n^*}(x', \tilde{a}, b) > x_{\min}(\tilde{a})$. The case of $f^{n^*}(x', \tilde{a}, b) < x_{\max}(\tilde{a})$ is entirely symmetric. Since n^* is minimal we have that

$$\forall 0 \leq n < n^* \quad f^n(x', \tilde{a}, b) \in [x_{\max}(\tilde{a}), x_{\min}(\tilde{a})]$$

As a result we have that

$$\begin{aligned} \forall 0 \leq 2n < n^* \quad f^{2n}([b, x'], \tilde{a}, b) &= [b, f^{2n}(x', \tilde{a}, b)] \\ \forall 0 < 2n + 1 < n^* \quad f^{2n+1}([b, x'], \tilde{a}, b) &= [f^{2n+1}(x', \tilde{a}, b), b] \end{aligned}$$

Using the above equations as well as [Equations \(B.2\) and \(B.3\)](#) we can recursively prove that

$$\forall 0 \leq n < n^* \quad x_{n+n^\dagger}([b, x']) \supseteq f^n([b, x'], \tilde{a}, b)$$

We can deduce that

$$x_{n^\dagger+n^*}([b, x']) \supseteq [b, f^{n^*}(x', \tilde{a}, b)] \supseteq [b, x_{\min}(\tilde{a})]$$

With one more iteration we have

$$x_{n^\dagger+n^*+1}([b, x']) \supseteq [f_{\min}(\tilde{a}), b]$$

We can also pick a $x'' \in [\gamma', b]$ such that $f(x'', \tilde{a}, b) = x'$. By construction, we know that $n^* + 1$ is the first iteration such that

$$f^{n^\dagger+n^*+1}(x'', \tilde{a}, b) \notin [x_{\max}(\tilde{a}), x_{\min}(\tilde{a})].$$

By similar arguments as above we can prove that

$$x_{n^\dagger+n^*+1}([x'', b]) \supseteq [b, f^{n^\dagger+n^*+1}(x'', \tilde{a}, b)] = [b, f^{n^\dagger+n^*}(x', \tilde{a}, b)] \supseteq [b, x_{\min}(\tilde{a})]$$

As a result we now have

$$x_{n^\dagger+n^*+1}([x'', x']) \supseteq [f_{\min}(\tilde{a}), x_{\min}(\tilde{a})] \supseteq [x_{\max}(\tilde{a}), x_{\min}(\tilde{a})]$$

It follows directly that

$$x_{n^\dagger+n^*+2}([\gamma, \delta]) \supseteq x_{n^\dagger+n^*+2}([x'', x']) \supseteq [f_{\min}(\tilde{a}), f_{\max}(\tilde{a})] \supseteq [x_{\max}(\tilde{a}), x_{\min}(\tilde{a})]$$

Applying the above step recursively we get

$$\forall n \geq n^\dagger + n^* + 2 \quad x_n([\gamma, \delta]) \supseteq [f_{\min}(\tilde{a}), f_{\max}(\tilde{a})]$$

so $n_0 = n^\dagger + n^* + 2$ satisfies the requirements of the theorem. ■

APPENDIX C. OMITTED PROOFS OF SECTION 5

We start this appendix by recalling the proof of 3-period focusing on the location the 3-period orbit and the $\mathcal{F}(a)$.

Theorem C.1. *If $b \in (0, 1) \setminus \{\frac{1}{2}\}$, then there exists u_b such that for all $a > u_b$ it holds that the map $f(x, a, b)$ has periodic orbit of period 3 in the interior of $\mathcal{F}(a) = [f_{\min}(a), f_{\max}(a)]$.*

Proof. If $f(x, a, b) > x$ is equivalent to $x < b$ and $f^3(x, a, b) < x$ is equivalent to $x + f(a, b, x) + f^2(a, b, x) > 3b$. Assume that $0 < b < 1/2$. Then $3b - 1 < b$, so we can take $x > 0$ such that $3b - 1 < x < b$. Then $f(x, a, b) > x$. Moreover, $\exp(a(x - b))$ goes to 0 as a goes to infinity, so

$$\lim_{a \rightarrow +\infty} f(a, b, x) = \lim_{a \rightarrow +\infty} \frac{x}{x + (1 - x) \exp(a(x - b))} = 1$$

Thus, since $3b - x < 1$, there exists u_b such that for all $a > u_b$ then $f(x, a, b) > 3b - x$, so $x + f(x, a, b) + f^2(x, a, b) > 3b$. Hence, if $a > u_b$ then $f^3(x, a, b) < x < f(x, a, b)$. Now, from the main theorem in [29] it follows that if $a > u_b$ then f has a periodic point of period 3. Observe that the chosen x depends only on b and not on a . Since

$$\lim_{a \rightarrow \infty} f_{\min}(a) = 0 \quad \lim_{a \rightarrow \infty} f_{\max}(a) = 1$$

we can pick u_b large enough such that for $a > u_b$ $x \in [f_{\min}(a), f_{\max}(a)]$. Picking $u_b > a_b$ we also have that $f(x, a, b)$ and $f^2(x, a, b)$ belong to $[f_{\min}(a), f_{\max}(a)]$ as well for $a > u_b$. The period 3 constructed by [29] has thus all its points in the interior of $[f_{\min}(a), f_{\max}(a)]$. The case of $1/2 < b < 1$ is symmetric because $f(x, a, b) = f(1 - x, a, 1 - b)$. ■

C.1. Decaying Volume Turbulent sets in Fixed Learning Rate Regime.

Lemma C.1 (Restated Lemma 12). *For every $a > u_b$, there exist closed and disjoint intervals K_a and J_a in the interior of $[f_{\min}(a), f_{\max}(a)]$ such that $f^2(K_a, a, b)$ and $f^2(J_a, a, b)$ are neighborhoods of $K_a \cup J_a$.*

Proof. The orbit of period 3 has the form $f^2(x, a, b) < x < f(x, a, b)$ or its mirror image (See Theorem 1 in [29]). Without loss of generality, assume it has the form above. Then we can choose,

- (1) d between x and $f(x, a, b)$, so that $f(d, a, b) = x$ and hence $d < f^2(d, a, b)$.
- (2) z between $f(x)$ and x , so close that $f^2(z, a, b) > d$.
- (3) q between z and x , so close to x that $f^2(z, a, b) < z$.
- (4) c between x and d , so close to x that $f^2(c, a, b) < z$.

Then, $J_a = [z, q]$ and $K_a = [c, d]$ are disjoint and

$$[z, d] \subset [f^2(q, a, b), f^2(z, a, b)] \subset f^2(J_a), [z, d] \subset [f^2(c, a, b), f^2(d, a, b)] \subset f^2(K_a)$$

Observe that $f^2(K_a, a, b)$ and $f^2(J_a, a, b)$ are supersets of K_a and J_a that share no endpoints with K_a and J_a so they are neighborhoods of $K_a \cup J_a$. Also since $x, f(x, a, b)$ and $f^2(x, a, b)$ are in the interior of the interval $[f_{\min}(a), f_{\max}(a)]$, we have that J_a and K_a have this property as well. ■

Moreover we can prove the following claim for the sets K_a, J_a

Lemma C.2 (Restated Lemma 13). *For every $a > s_b$, there exist closed intervals $V_a^k \subseteq K_a$ and $U_a^k \subseteq J_a$ such that*

$$\lim_{k \rightarrow \infty} \text{diam}(V_a^k) = \lim_{k \rightarrow \infty} \text{diam}(U_a^k) = 0$$

and for every $k \geq 0$ it holds that $f^{2k+2}(V_a^k, a, b)$ and $f^{2k+2}(U_a^k, a, b)$ are neighborhoods of $K_a \cup J_a$.

Proof. We will prove the lemma by induction. Choosing $V_a^0 = K_a$ we know by Lemma 12 that $f^2(V_a^0, a, b)$ is a neighborhood of $K_a \cup J_a$. Now let us assume that we have a $V_a^k \subseteq K_a$ such that $f^{2k+2}(V_a^k, a, b)$ is a neighborhood of $K_a \cup J_a$. Since K_a and J_a are disjoint closed intervals, there are two disjoint closed intervals Z_1 and Z_2 of V_a^k such that $f^{2k+2}(Z_1, a, b)$ is a neighborhood of K_a and $f^{2k+2}(Z_2, a, b)$ is a neighborhood of J_a . As a result $f^{2k+4}(Z_1, a, b)$ and $f^{2k+4}(Z_2, a, b)$ are neighborhoods of $K_a \cup J_a$. Given that Z_1 and Z_2 are disjoint intervals of V_a^k it must be the case that

$$\min\{\text{diam}(Z_1), \text{diam}(Z_2)\} \leq \text{diam}(V_a^k)/2.$$

We pick the interval with the smallest diameter as V_a^{k+1} and thus

$$0 \leq \text{diam}(V_a^k) \leq 2^{-k} \text{diam}(K_a) \implies \lim_{k \rightarrow \infty} \text{diam}(V_a^k) = 0.$$

Choosing $U_a^0 = J_a$ we can follow the same arguments for the rest of the U_a^i . ■

C.2. Tracking properties, make chaos explicit via symbolic dynamics. In this section, we will demonstrate how to construct a scrambled set of initial conditions via a scrambled set of abstract symbolic orbits. Symbolic dynamics is a mathematical method used in dynamical systems theory to study the long-term behavior of a system. It involves representing the states of a system as a sequence of symbols, typically taken from a finite alphabet—in our case $\{0, 1\}^*$ —. In our proof, these symbols are chosen based on position of the trajectory of a point in strategy space. The resulting sequence of symbols is called a symbolic orbit, and the study of these orbits will provide the necessary insight to establish the long-term chaotic behavior of MWU.

Below we present the tracking lemma that translates a binary sequence to the trajectory of an initial condition inside the decaying sequence of turbulent sets $V_{a^*}^i, U_{a^*}^i$:

Lemma C.3 (Restated Lemma 14). *If $b \in (0, 1) \setminus \{\frac{1}{2}\}$, there exists a d_b such that if $a_{\min} > d_b$, we can construct an increasing sequence n_i with the following properties. For every sequence of intervals A_i with $A_i = V_{a^*}^i$ or $A_i = U_{a^*}^i$, there exists a $x_0 \in [0, 1]$ such that for all $i \geq 0$ it holds that $x_{n_i}(x_0) \in A_i$.*

Proof. By continuity of $f_{\min}$ and $f_{\max}$ we know that

$$\lim_{\epsilon \rightarrow 0} f_{\min}(a^* - \epsilon) = f_{\min}(a^*) \quad \lim_{\epsilon \rightarrow 0} f_{\max}(a^* - \epsilon) = f_{\max}(a^*)$$

By Lemma 12 we know that K_{a^*} and J_{a^*} are in the interior of $[f_{\min}(a^*), f_{\max}(a^*)]$. As a result there is a sufficiently small $\epsilon > 0$ such that both of the following properties hold

$$\begin{aligned} a^* &> a^* - \epsilon > a_{\min} \\ [f_{\min}(a^* - \epsilon), f_{\max}(a^* - \epsilon)] &\supseteq K_{a^*} \cup J_{a^*} \end{aligned}$$

Let us pick any interval $[\gamma, \delta]$ such that $\{b\} \subset [\gamma, \delta] \subset (0, 1)$. By Theorem B.1 and for the aforementioned ϵ , we know that there is a n^* such that

$$\forall n \geq n^*: \quad x_n([\gamma, \delta]) \supseteq [f_{\min}(a^* - \epsilon), f_{\max}(a^* - \epsilon)] \supseteq K_{a^*} \cup J_{a^*}$$

We also know that for all $i \geq 0$, $f^{2i+2}(V_{a^*}^i, a^*, b)$ and $f^{2i+2}(U_{a^*}^i, a^*, b)$ are neighborhoods of $K_{a^*} \cup J_{a^*}$. Thus for each $i \geq 0$ there must be an $\epsilon_i > 0$ such that $f^{2i+2}(V_{a^*}^i, a^*, b)$ and $f^{2i+2}(U_{a^*}^i, a^*, b)$ are ϵ_i neighborhoods of $K_{a^*} \cup J_{a^*}$. By Lemma 10, there is a sequence of m_i such that for all $i \geq 0$

$$\forall n \geq m_i \quad \max_{x_0 \in [\gamma, \delta]} |x_{n+2i+2}(x_0) - f^{2i+2}(x_n(x_0), a^*, b)| \leq \epsilon_i$$

We are now ready to construct the sequence n_i . We choose $n_0 = \max\{m_0, n^*\}$. For n_i with $i \geq 1$, we choose the minimum number with the following properties

$$\begin{aligned} (n_i - n_{i-1} - 2i) \mod 2 &= 0 \\ n_i &\geq n_{i-1} + 2i \quad n_i \geq m_i \end{aligned}$$

It now remains to construct the required $x_0 \in [0, 1]$ for each potential sequence of A_i . Since $x_{i_0}([\gamma, \delta])$ is a superset of the disjoint closed intervals $K_{a^*} = V_{a^*}^0$ and $J_{a^*} = U_{a^*}^0$ we know that there are closed minimal intervals I_0 and I_1 such that

$$x_{n_0}(I_0) = V_{a^*}^0 \quad x_{n_0}(I_1) = U_{a^*}^0$$

Since $n_0 \geq m_0$ we know that

$$\max_{x_0 \in I_0} |x_{n_0+2}(x_0) - f^2(x_{n_0}(x_0), a^*, b)| \leq \epsilon_0$$

Given that $x_{n_0}(I_0) = V_{a^*}^0$ and $f^2(V_{a^*}^0, a^*, b)$ is an ϵ_0 neighborhood of $K_{a^*} \cup J_{a^*}$, we can infer that

$$x_{n_0+2}(I_0) \supseteq K_{a^*} \cup J_{a^*}$$

With a similar analysis we can prove that

$$x_{n_0+2}(I_1) \supseteq K_{a^*} \cup J_{a^*}$$

Following the same steps repeatedly we can prove for any $k > 0$ that

$$x_{n_0+2k}(I_0) \supseteq K_{a^*} \cup J_{a^*} \quad x_{n_0+2k}(I_1) \supseteq K_{a^*} \cup J_{a^*}$$

Since $n_1 = n_0 + 2k$ for some $k > 0$ we can directly infer that

$$x_{n_1}(I_0) \supseteq K_{a^*} \cup J_{a^*} \quad x_{n_1}(I_1) \supseteq K_{a^*} \cup J_{a^*}$$

Hence we can construct minimal closed intervals $I_{00}, I_{01} \subset I_0$ and $I_{10}, I_{11} \subset I_1$ such that

$$x_{n_1}(I_{00}) = V_{a^*}^1 \quad x_{n_1}(I_{01}) = U_{a^*}^1 \quad x_{n_1}(I_{10}) = V_{a^*}^1 \quad x_{n_1}(I_{11}) = U_{a^*}^1$$

By induction, for any binary sequence c of length i we can construct two minimal closed intervals $I_{c_0 \dots c_{i-1} 0}$ and $I_{c_0 \dots c_{i-1} 1}$ subsets of $I_{c_0 \dots c_{i-1}}$ such that

$$x_{n_i}(I_{c_0 \dots c_{i-1} 0}) = V_{a^*}^i \quad x_{n_i}(I_{c_0 \dots c_{i-1} 1}) = U_{a^*}^i$$

Now for any sequence A_i such that $A_i = V_{a^*}^i$ or $A_i = U_{a^*}^i$ we can construct the corresponding sequence c_i that has $c_i = 0$ when $A_i = V_{a^*}^i$ and $c_i = 1$ when $A_i = U_{a^*}^i$. Let us define

$$I_c = \bigcap_{i=1}^{\infty} I_{c_0 \dots c_i}$$

By Cantor's intersection theorem, I_c is non empty. Any $x_0 \in I_c$ satisfies the requirements of the lemma

$$\forall i \geq 0 \quad x_0 \in I_{c_1 \dots c_i} \implies \forall i \geq 0 \quad x_{n_i}(x_0) \in A_i$$

■

Theorem C.2 (Restated Theorem 4). *If $b \in (0, 1) \setminus \{\frac{1}{2}\}$, there exists a d_b such that if $a_{\min} > d_b$, the dynamics of Equation 2 are Li-Yorke chaotic.*

Proof. We first prove that there is an uncountable set S of infinite length binary sequences with the following property: For every pair of sequences $i, j \in S$ there is an infinite length subsequence where the two sequences differ, i.e., one of i, j is 0 and the other is 1.

We first define the equivalence relation R over infinite length binary sequence such that two binary sequences are equivalent if and only if they differ in finitely many places. The relation is clearly reflexive, a sequence differs with itself in 0 places, it is by definition symmetric and it is transitive, if a sequence i differs in k places with sequence j and j differs in m places with a sequence t then i and t differ in up to $m + k$ places which is also finite. Thus we can partition all binary sequences in equivalence classes where all pairs of all elements in the same class differ in finite places.

We now prove that each equivalence class has countably infinite binary sequences. Let us pick an element σ of the equivalence class. For each member of the equivalence class w we can construct a finite subset of $\mathbb{N}$ by picking the indices where $s_t \neq w_t$. Inversely, for each finite subset of $\mathbb{N}$ we can construct a member of the equivalence by flipping the corresponding indices of σ . Thus there is a bijection between the members of the equivalence class and the finite subsets of $\mathbb{N}$. Because the finite subsets of $\mathbb{N}$ are countably infinite, so is each equivalence class.

This allows us to prove that the relation R has uncountably infinite number of equivalence classes. We proceed by contradiction. If we had countably finite equivalence classes each having countably many elements then all of the binary sequences would be countable which is false.

To construct S we only need to pick one element from each equivalence class. Since each pair belongs in a different equivalent class then they differ in infinite number of places which forms an infinite length subsequence. The set S is uncountable because there are uncountably many equivalence classes.

We now prove that there is an uncountable set S' of infinite length binary sequences that has the following two properties: First, for every pair of sequences $i, j \in S'$ there is an infinite length subsequence where the two sequences differ, i.e., one of i, j is 0 and the other is 1. Second, for every pair of sequences $i, j \in S'$ there is an infinite length subsequence where the two sequences are equal, i.e., they are either both 0 or 1.

The construction works as follows: For each element of $i \in S$ we construct a new sequence. In its even places the sequence is 0 and in its odd places we use the elements of i . Clearly by construction S' remains uncountable and satisfies the first property as we can pick the subsequences from the odd places. The second property also holds because all sequences have the same elements in the even places.

For each element of S' we construct a sequence of sets T as follows: If the k th place element is 0 we use $V_{a^*}^k$, whereas if it is 1 we pick $U_{a^*}^k$. We now apply Lemma 14 for each of the sequence of sets to get a corresponding initialization x_0 and a subsequence of iteration indices n_i . We call this set of initializations Q . By the construction in Lemma 14 all initializations use the same subsequence n_i .

Let us pick two initializations $\mu, \nu \in Q$. By the construction of T and S' we have the following: There is an infinite subsequence of n_i , which we call h_t , where $x_{n_{h_t}}(\mu) \in V_{a^*}^{h_t}$ and $x_{n_{h_t}}(\nu) \in U_{a^*}^{h_t}$ or vice versa. Because $V_{a^*}^{h_t}$ and $U_{a^*}^{h_t}$ are disjoint we have that there is an infinite subsequence where the trajectories of μ and ν are bounded away from each other. In other words, we have that

$$\limsup_{n \rightarrow \infty} |x_n(\mu) - x_n(\nu)| > 0$$

Again by construction of T and S' we also have that there is an infinite subsequence of n_i , which we call f_t such that either $x_{n_{f_t}}(\mu) \in V_{a^*}^{f_t}$ and $x_{n_{f_t}}(\nu) \in V_{a^*}^{f_t}$ or $x_{n_{f_t}}(\mu) \in U_{a^*}^{f_t}$ and $x_{n_{f_t}}(\nu) \in U_{a^*}^{f_t}$. Thus we have that

$$\forall t \geq 0 : |x_{n_{f_t}}(\mu) - x_{n_{f_t}}(\nu)| \leq \max\{\text{diam}(U_{a^*}^{f_t}), \text{diam}(V_{a^*}^{f_t})\}$$

By Lemma 13, we have that the right hand side converges to 0. As a result

$$\lim_{t \rightarrow \infty} |x_{n_{f_t}}(\mu) - x_{n_{f_t}}(\nu)| = 0$$

which directly implies that

$$\liminf_{n \rightarrow \infty} |x_n(\mu) - x_n(\nu)| = 0.$$

Technically our dynamics are expressed in terms of two state variables x_n and r_n so distances in orbits need to account for both dimensions. Because $r_n \rightarrow 0$ for all initializations $x_0 \in (0, 1)$, in all cases above we have $\lim_{n \rightarrow \infty} |r_n(\mu) - r_n(\nu)| = 0$ so the distances in the limit are not affected by the r_n dimension.

Picking $\{(x, r) : x \in Q \text{ and } r = 0\}$ as our scrambled set, which is equivalent to choosing $a_0 = g(0)$ for all $x \in Q$, we show that our dynamic learning rate dynamics are Li-Yorke chaotic. ■

APPENDIX D. BIFURCATION PLOTS

In this section, we showcase a series of bifurcation diagrams illustrating the adaptive scheme's limit behavior in response to varying constraints on a_∞ . These diagrams display the emergence of periodic points as a function of the proportion of users favoring the first link in the equilibrium state of the system (parameter b). A notable observation from these bifurcation plots is their symmetric nature. For instance, we anticipate analogous behavior when the equilibrium state accommodates




Figure 4: The cobweb diagram above displays the trajectory of iterates for the time-varying learning rate dynamics in (2), considering various equilibrium choices. The patterns formed by these paths reveal insights into the iterative behavior of the dynamics. Fixed points are identified by the intersection of the $y = x$ diagonal line and the limit function $\text{MWU}(x, \lim_{n \rightarrow \infty} a_n, \text{Equilibrium})$, with spirals converging towards these points. Period-2 orbits create rectangles, while higher-period cycles generate increasingly complex closed loops. Chaotic orbits, conversely, occupy an area, indicating an infinite sequence of non-repeating values. This phenomenon intensifies as we move further from the boundaries.

80% of users favoring the first link in the congestion game and when it accommodates 20%. This symmetry implies an inherent balance in the system's response to changes in user preferences.






Figure 5: When the limit behavior a_∞ is small the only system attractors are equilibria (roughly for $b < 0.2$ and $b > 0.8$) and period-two cycles.



(a) When the limit behavior a_∞ is significantly larger, which signifies a more volatile, highly fluctuating system, we see the emergence of more complex attractors, i.e., periodic orbits of large period, chaos.



(b) Even more aggressive schemes imply the creation of complex attractors even when the equilibrium b is much closer to the boundary i.e. to 0 or 1, in comparison with Fig. 6a.

Figure 6: Increased Limit behavior a_∞