No-Regret Dynamics Do Not Eliminate Dominated Strategies

SUBMISSION 2233

We show that no-regret learning dynamics need not eliminate strictly dominated strategies. We construct a bimatrix game and show that, if both players follow (continuous or discrete time) projected gradient dynamics in this game, then there exists a periodic orbit that has a segment along which a strictly dominated strategy is played with positive probability. Moreover, a positive measure set of initial conditions has trajectories which converge to this periodic orbit.

Since the discrete-time projected gradient dynamic with step size $1/\sqrt{t}$ has the no-regret property, our results imply that no-regret learning can persistently select strictly dominated strategies. Therefore, the desirable properties of no-regret learning in single-agent settings do not extend to strategic interactions, even for simple dominance-based refinements.

Contents

Manuscript submitted for review to the 27th ACM Conference on Economics & Computation (EC'26).

1 Introduction

No-regret learning has emerged as a foundational framework for multi-agent reinforcement learning, offering an appealing alternative to traditional equilibrium-based approaches. The foundation traces back to work by [Blackwell, 1956] and [Hannan, 1957], but gained renewed prominence through the seminal contributions of [Foster and Vohra, 1997] and [Hart and Mas-Colell, 2000], who established fundamental connections between regret minimization and equilibrium concepts in game theory. This framework has since been extensively developed in both the game theory and computer science literatures (Cesa-Bianchi and Lugosi, 2006; Shoham and Leyton-Brown, 2009).

The appeal of no-regret learning rests on several compelling features. First, it requires only minimal rationality assumptions from agents: players need not know the game structure, the number of opponents, or their payoff functions. Instead, agents simply aim to minimize regret, that is, to perform as well as the best alternative strategy in hindsight. Second, the framework provides theoretical guarantees. When all players employ algorithms with vanishing external regret, their empirical distribution of play converges to the set of coarse correlated equilibria (CCE) [Foster and Vohra, 1997]. If players instead minimize swap regret, convergence is guaranteed to the set of correlated equilibria (CE) (Hart and Mas-Colell, 2000; Blum and Mansour, 2007). These convergence results hold for a broad class of no-regret algorithms including multiplicative weights, follow-the-regularized-leader, and gradient-based methods. Third, the framework provides a unifying lens for understanding learning dynamics across diverse settings, from online prediction to sequential decision-making in strategic environments [Cesa-Bianchi and Lugosi, 2006].

However, the convergence guarantee to CCE raises questions about whether such outcomes satisfy basic strategic rationality requirements. [Moulin and Vial, 1978] showed that certain CCE assign positive probability to strictly dominated strategies. [Viossat and Zapechelnyuk, 2013] strengthened this result, constructing games with CCE in which all players use strictly dominated strategies with positive probability in every outcome. These examples demonstrate that minimizing external regret does not guarantee basic forms of strategic rationality.

A natural question is whether specific learning algorithms avoid this issue in practice. [Hofbauer and Weibull, 1996] established that replicator dynamics eliminate strictly dominated strategies from almost all initial conditions. Given the known connections between certain no-regret algorithms and replicator dynamics, one might hope that gradient-based no-regret learning would similarly eliminate dominated strategies. [Hofbauer and Sandholm, 2011] investigated this question more generally, showing that dynamics satisfying certain properties fail to eliminate strictly dominated strategies in the "hypnodisc" game for a positive measure set of initial conditions. However, projected gradient dynamics do not satisfy the properties considered in that paper.

Our main contribution is to show that gradient descent with Euclidean (also called greedy or eager) projection does not eliminate dominated strategies, even when all players employ no-regret learning algorithms. We construct a simple bimatrix game (Matching Pennies with a Feeble Twin) and we prove that, if both players follow continuous-time or discrete-time projected gradient dynamics, the there exists a periodic orbit that has a segment along which a strictly dominated strategy is played with positive probability. Moreover, a positive measure set of initial conditions has trajectories which converge to this orbit. And, since the discrete-time projected gradient dynamic with step size $1/\sqrt{t}$ has the no-regret property, we conclude that no-regret learning can persistently select strictly dominated strategies.

Our result sheds new light on the properties of no-regret learning. In particular, the elimination of strictly dominated strategies and the no-regret guarantee are entirely distinct properties of algorithms. In gradient descent with lazy projection, players accumulate gradients and project when computing their next strategy, corresponding to the follow-the-regularized-leader framework. In gradient descent with Euclidean (also called greedy or eager) projection, players project immediately after each gradient step. While both methods achieve no-regret, their relationship to dominated strategy elimination differs: the former is known to eliminate strictly dominated strategies , while our results show that the latter does not.

2 Set Up

Let Σ_1, Σ_2 be finite action sets, and let $A, B \in \mathbb{R}^{|\Sigma_1| \times |\Sigma_2|}$ be the corresponding payoff matrices. The associated two-player bimatrix game is the strategic-form game $G = (\{1,2\}, \Sigma_1, \Sigma_2, u_1, u_2)$, where the payoff functions $u_1, u_2 : \Sigma_1 \times \Sigma_2 \to \mathbb{R}$ are defined by

$$u_1(i,j) = A_{ij} \qquad \text{and} \qquad u_2(i,j) = B_{ij}, \quad (i,j) \in \Sigma_1 \times \Sigma_2.$$

Thus, if the action profile (i,j) is played, player 1 receives payoff A_{ij} and player 2 receives payoff B_{ij}. Let $X = \Delta(\Sigma_1)$ and $Y = \Delta(\Sigma_2)$ denote the simplices of mixed strategies for players 1 and 2, respectively. By a slight, but standard, abuse of notation, we extend u_1 and u_2 to mixed strategies by bilinearity. For $x \in X$ and $y \in Y$, the expected utilities of players 1 and 2 are given by

$$u_1(x,y) = \sum_{i \in \Sigma_1} \sum_{j \in \Sigma_2} x_i \, A_{ij} \, y_j = x^\top A y,$$
$$u_2(x,y) = \sum_{i \in \Sigma_1} \sum_{j \in \Sigma_2} x_i \, B_{ij} \, y_j = x^\top B y.$$

Definition 2.1. A pair of mixed strategies $(x^*, y^*) \in X \times Y$ is a *Nash equilibrium* if no player can improve their expected payoff by unilaterally deviating. That is,

$$u_1(x^*, y^*) \geq u_1(x, y^*) \quad \forall x \in X, \qquad u_2(x^*, y^*) \geq u_2(x^*, y) \quad \forall y \in Y.$$

Definition 2.2. An action $\alpha \in \Sigma_1$ of player 1 is *strictly dominated by a pure action* $\alpha' \in \Sigma_1$ if

$$u_1(e_{\alpha'}, y) > u_1(e_\alpha, y) \quad \text{for all } y \in \Delta(\Sigma_2),$$

where e_α denotes the pure strategy assigning probability 1 to α. An analogue definition holds for player 2's actions[1].

2.1 Projected gradient algorithm

We now introduce the projected gradient dynamics.

The gradient of the strategies of player 1 and player 2 respectively are

$$\nabla_x u_1(x,y) = Ay, \qquad \nabla_y u_2(x,y) = B^\top x.$$

Namely, the gradient of player 1 is the vector of expected utilities of the pure actions in Σ_1 when played against y, Ay, and the gradient of player 2 is the vector of expected utilities of the pure actions in Σ_2 when played against x, $B^\top x$.

Given an initial mixed strategy profile (x^0, y^0) and a sequence of step sizes $(\gamma^t)_{t \in \mathbb{N}}$, the discrete projected gradient algorithm produces a sequence $((x^t, y^t))_{t \in \mathbb{N}}$ in $X \times Y$ in which each player updates in the direction of their utility gradient and then projects back onto the simplex. Concretely, we denote by Π_X and Π_Y the Euclidean projection onto the sets X and Y (respectively), for every t in $\mathbb{N}$,

$$x^{t+1} = \Pi_X\big(x^t + \gamma^t \nabla_x u_1(x^t, y^t)\big) = \Pi_X\big(x^t + \gamma^t A y^t\big),$$
$$y^{t+1} = \Pi_Y\big(y^t + \gamma^t \nabla_y u_2(x^t, y^t)\big) = \Pi_Y\big(y^t + \gamma^t B^\top x^t\big).$$

[1]A pure action can also be dominated by a mixed action, but our examples use the more restrictive definition.

2.2 Regret minimising and projected gradient

The no-regret property formalises the idea that, ex post, each player's cumulative payoff is asymptotically as good as if they had committed in advance to the best fixed pure action against the realised sequence of opponent play.

Definition 2.3. Let $(x^t)_{t=1}^T \subset X$ be the mixed strategies played by player 1 while the opponent's mixed strategies are $(y^t)_{t=1}^T \subset Y$. For each pure action $i \in \Sigma_1$ write $e_i \in X$ for the unit vector that places mass 1 on i. The regret of player 1 up to time T is

$$R_1(T) \coloneqq \max_{i \in \Sigma_1} \sum_{t=1}^{T} \big(u_1(e_i, y^t) - u_1(x^t, y^t)\big).$$

Player 1 is said to have no regret if $R_1(T)/T \to 0$ as $T \to \infty$. The analogous definition applies to player 2.

Projected-gradient learning is known to be no-regret under the canonical step size choice $\gamma_t = t^{-1/2}$ [Zinkevich, 2003]. An inspection of that argument shows it extends to step sizes of the form $\gamma_t = t^{-\alpha}$ for any $\alpha \in (0, 1)$. We record the property as a lemma and attribute it to Zinkevich; the minor adjustments are collected in Appendix B.

Lemma 2.4. *Let $((x^t, y^t))_{t\geq 1}$ be the sequence produced by the discrete projected gradient algorithm with step sizes $\gamma_t = t^{-\alpha}$ for some $\alpha \in (0, 1)$ and arbitrary initial $(x^0, y^0) \in X \times Y$. Then each player has vanishing average regret:*

$$\frac{R_k(T)}{T} \to 0 \quad \text{as } T \to \infty, \qquad k = 1, 2.$$

Projected gradient is therefore a no-regret algorithm for the usual polynomially-decaying step sizes of interest.

3 Results

The goal of this paper is to show that there exists a set of positive Lebesgue measure in $X \times Y$ such that, for any initial point in this set, the dynamics described above fail to eliminate certain strictly dominated strategies in the long run. In particular, we will show that, for a chosen game and a non-measure zero set of initial conditions, there exist a strictly dominated action that is played with a strictly positive probability infinitely often and is never permanently eliminated.

Theorem 3.1. *There exists a finite bimatrix game and a non-measure zero set of initial conditions such that, , when both players update their strategies according to the discrete-time projected gradient dynamics with step size sequence $(\gamma_t)_{t\in\mathbb{N}}$ satisfying $\max_t \gamma_t < \delta$ for sufficiently small $\delta > 0$, there exists a strictly dominated action α for player one satisfying*

$$\limsup_{t\to\infty} x_\alpha^t > 0.$$

We conjecture that the phenomenon described in Theorem 3.1 is not isolated: in fact, we expect it to occur for an open set of bimatrix games. This is because our results relies upon several structural proporties that vary smoothly with the payoffs: the Lipschitz constant, an attracting neighbourhood, and the return time. Since we prove that in any dynamical system with these properties,

the periodic orbit that we construct is understood to a far higher level of numerical accuracy than the basin of attraction that it sits within, providing substantial margin for perturbation. The Lipschitz constant, basin of attraction, and return time all vary continuously with the payoff parameters. Since Theorem... guarantees that any game with an approximate orbit under our conditions, and these condition vary smoothy with the payoffs, admits an attracting orbit, continuity arguments suggest the phenomenon persists under small payoff perturbations.

Furthermore, by choosing an appropriate step size, we obtain:

Corollary 3.2. *No-regret dynamics are not guaranteed to eliminate strictly dominated strategies.*

We begin by introducing two examples, in simpler frameworks, that illustrate the possible mechanisms keeping a strictly dominated strategy in the support. Then we introduce our main example (which is more technically involved, and on which Theorem 3.1 is based), with two players independently using projected gradient learning, and with a strictly dominated strategy that does not disappear from the support.

3.1 Single learning player again a strategic opponent

Suppose a player, player 1, uses projected gradient with step-size $\frac{1}{\sqrt{t}}$, against an opponent that aims to keep a dominated strategy repeatedly appearing in the support of player 1's strategy (with a limsup probability bounded away from 0).

We show a simple example and give the intuition as to why in this case, the dominated strategy does not vanish. The full details are in Appendix C.

For this example, only the payoffs of player 1 are relevant, and they are given in the matrix below.

$$\begin{array}{c|cc} & L & R \\ \hline T & 1 & 0 \\ M & 0 & 1 \\ B & -d & 1-d \end{array} \tag{1}$$

For $d > 0$ Action B is strictly dominated by action A. We further assume $d < \frac{1}{2}$. The strategy of player 2 that guarantees that action B does not disappear relies on two simple observations: when player 2 plays L, the probability of T increases, and will increase to 1 if T is played sufficiently many times. Then, if player 2 switches and plays R, both the probabilities of action M and action B increase. Alternating between playing L many times and playing R many times, we show in Appendix C that the limsup of he probability of B is $\frac{1-2d}{2-d}$, which is bounded away from 0 for any $d \in (0, \frac{1}{2})$.

3.2 Single learning player playing against itself

A single population game, with the population adapting according to the continuous time projected gradient dynamics is considered in [Sandholm et al., 2008]. They consider the following game, which is a variation of Rock-Paper-Scissors with a feeble twin:

$$\begin{array}{c|cccc} & R & P & S & FS \\ \hline R & 0 & -3 & 2 & 2 \\ P & 2 & 0 & -3 & -3 \\ S & -3 & 2 & 0 & 0 \\ FS & -3-c & 2-c & -c & -c \end{array} \tag{2}$$

They show the survival of the dominated strategy by analysing numerically a closed orbit of the dynamic, for $c = \frac{1}{10}$, which includes a part where the dominated strategy appears . They also claim that the orbit is attracting, which implies that there is a non-zero measure of starting points from which the dynamic ends up in this orbit ([Sandholm et al., 2008], Theorem 5.2).

We prove that, in the single-population case, replacing the continuous dynamic with discrete dynamics with decreasing step-sizes (and in particular, with step sizes that ensure the dynamic has the no-regret property) keeps the survival of the dominated strategy, due to continuity arguments combined with the orbit being attracting.

Unfortunately, this analysis does not trivially extend to two-population games. Even if we limit the initial strategy to be symmetric, which means that we have a zero-measure set of starting points, we still cannot trivially obtain the survival of the dominated strategy. The reason is that the orbit being attracting is stated only for symmetric deviations from the orbit, while for two populations asymmetric deviations should be considered as well. Without the orbit being attracting, we cannot know that the discrete dynamics has the same behaviour as the continuous one.

Therefore, in the next section, we introduce an example designed for two-populations games, and where the dominated strategy survives.

3.3 Two independent players

The proof of Theorem 1 is by analysing the following bimatrix game:

$$\begin{array}{c|cc} & L & R \\ \hline T & (k,0) & (0,1) \\ M & (0,k) & (1,0) \\ B & (-d,1) & (1-d,0) \end{array} \tag{3}$$

where $k \geq 1$ and $d > 0$ is small. Specifically, we analyse the case $d = \frac{1}{5}$ and $k = 5$. Ignoring the Bottom action, the game is ordinally equivalent to classic Matching Pennies, but is cardinally skewed towards (Top, Left) for the row player and towards (Middle, Left) for the column player. The Bottom action for the row player is a "feeble twin" of Middle and is strictly dominated by it.

Figure 1 shows what simulations suggest is a cycle generated by the projected dynamics. The cycle consists of several segments, with varying supports, as detailed in Appendix E. Importantly, there are segments of the cycle for which x_3 is bounded away from zero and, therefore, along which the strictly dominated strategy of player 1 persists.

The ultimate aim of our analysis in the appendix is to *prove* that this cycle actually exists.

One notable feature of the cycle is that one segment (ζ_6 to ζ_1) is on the edge of the simplex: Player 1 plays a pure strategy while only Player 2's mixed strategy evolves. Since Player 2 has only two actions, the dynamics in this part of the cycle follow a line segment. This structure is particularly convenient because, once the trajectory enters this line segment, it reliably follows the same portion of the cycle in each iteration.

The proof considers the continuous-time approximation of the discrete-time gradient dynamics. In the continuous-time dynamics, even if we do not know the exact point at which the trajectory first reaches the line segment, it will always leave the segment at the same exit point for each cycle. Consequently, the trajectory repeatedly traverses the same cycle, which allows us to predict the sequence of segments and ensures the persistence of the dominated strategy throughout the cycle. As a consequence, there exists a set of initial conditions of positive measure from which the strictly dominated strategy persists under the continuous-time dynamics.

Once the continuous dynamic is clear, it is combined with [Dupuis and Nagurney, 1993] to show that the discrete-time projected gradient dynamic with step size $1/\sqrt{t}$ approximates the same cycle. In the discrete approximations we show that the dominated strategy is not eliminated. Furthermore, we show that there is a positive measure set of initial strategies that converges to this cycle. This will then yield Theorem 3.1.

References

Blackwell, D. (1956). An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8.

Blum, A. and Mansour, Y. (2007). From external to internal regret. *Journal of Machine Learning Research*, 8(6):1307–1324.

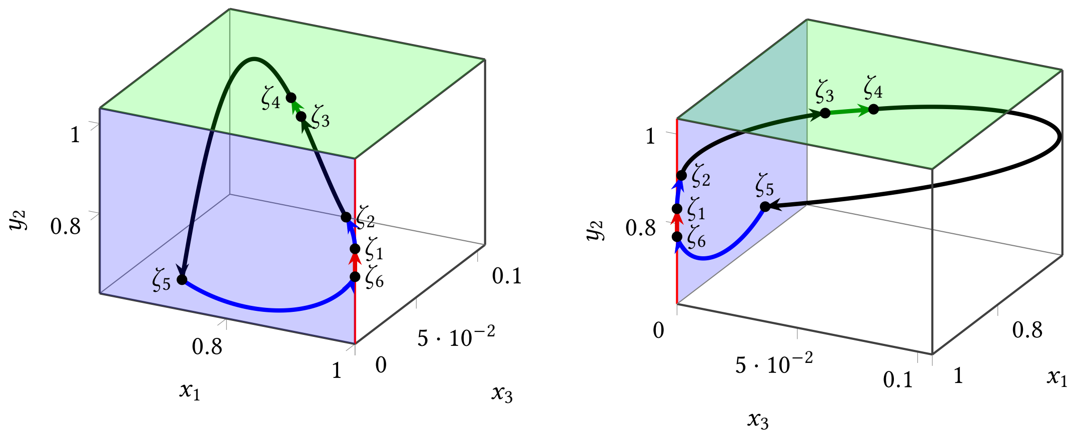

(a) View from the front

(b) View from the side

Fig. 1. The cycle.

Note. The path segments in blue are against the blue face $x_3 = 0$, and the path segment in green is against the green face $y_2 = 1$. The path segment in red is on the edge $x_1 = 1$. The path segments in black are in the relative interior.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games.* Cambridge University Press.

Dupuis, P. and Nagurney, A. (1993). Dynamical systems and variational inequalities. *Annals of Operations Research*, 44(1):7–42.

Foster, D. P. and Vohra, R. V. (1997). Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55.

Hannan, J. (1957). Approximation to Bayes risk in repeated play. In *Contributions to the Theory of Games*, volume 3, pages 97–139. Princeton University Press.

Hart, S. and Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150.

Hazan, E., Agarwal, A., and Kale, S. (2007). Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2):169–192.

Hofbauer, J. and Sandholm, W. H. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6(3):341–377.

Hofbauer, J. and Weibull, J. W. (1996). Evolutionary selection against dominated strategies. *Journal of Economic Theory*, 71(2):558–573.

Moulin, H. and Vial, J.-P. (1978). Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221.

Sandholm, W. H., Dokumacı, E., and Lahkar, R. (2008). The projection dynamic and the replicator dynamic. *Games and Economic Behavior*, 64(2):666–683.

Shoham, Y. and Leyton-Brown, K. (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* Cambridge University Press.

Viossat, Y. and Zapechelnyuk, A. (2013). No-regret dynamics and fictitious play. *Journal of Economic Theory*, 148(2):825–842.

Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936.

A Projected Gradient Dynamics

Given an initial mixed strategy profile (x^0, y^0) and a sequence of step sizes $(\gamma^i)_{i\in\mathbb{N}}$, the discrete projected gradient algorithm is the dynamical system $((x^i, y^i))_{i\in\mathbb{N}}$ on the set $X \times Y$ where for every

t in $\mathbb{N}$,

$$x^{t+1} = \Pi_X(x^t + \gamma^t Ay),$$
$$y^{t+1} = \Pi_Y(y^t + \gamma^t x^{t\top} B).$$

Alternatively, given an initial mixed strategy profile (x^0, y^0), the continuous projected gradient algorithm is the dynamical system on the set $X \times Y$ defined by the differential equation:

$$\dot{x} = \Pi_{T_X(x,y)}(Ay),$$
$$\dot{y} = \Pi_{T_Y(x,y)}(x^\top B),$$

where $T_S(x, y)$ denotes the tangent cone of a set S at (x, y).

The following proposition characterises these dynamics in a way that accounts for the simplex constraints of the strategy spaces.

Proposition A.1. *Let $(x, y) \in X \times Y$ be a mixed-strategy profile. For player $i \in \{1, 2\}$, define the* active set $S_i(x_i) \subseteq \Sigma_i$ *as the smallest superset of* $\operatorname{supp}(x_i)$ *that maximises*

$$\frac{1}{|S_i(x_i)|} \sum_{\beta \in S_i(x_i)} u_i(\beta, x_{-i}),$$

where x_{-i} denotes the current strategy of the other player.

Then, at time t, the discrete-time update is given by, for each action $\alpha \in \Sigma_i$,

$$x_{i,\alpha}^{t+1} = \begin{cases} x_{i,\alpha}^t + \gamma^t \left(u_i(\alpha, x_{-i}^t) - \frac{1}{|S_i(x_i^t)|} \sum_{\beta \in S_i(x_i^t)} u_i(\beta, x_{-i}^t) \right), & \text{if } \alpha \in S_i(x_i^t), \\ x_{i,\alpha}^t, & \text{if } \alpha \notin S_i(x_i^t), \end{cases}$$

where $(\gamma^t)_{t \in \mathbb{N}}$ is a sequence of step sizes.

Furthermore, the continuous-time differential equation is:

$$\dot{x}_{i,\alpha} = \begin{cases} u_i(\alpha, x_{-i}) - \frac{1}{|S_i(x_i)|} \sum_{\beta \in S_i(x_i)} u_i(\beta, x_{-i}), & \text{if } \alpha \in S_i(x_i), \\ 0, & \text{otherwise.} \end{cases}$$

B Step Size Condition for Projected Gradient to be No Regret

Theorem 1 in [Zinkevich, 2003] shows that the discrete-time projected gradient dynamic has the no-regret property for step size $t^{-\frac{1}{2}}$.

The notations are such that the step-size is η_t, $||F||$ is the diameter of the feasible set (for us it is a simplex, so 1), and $||\nabla_c||$ is a bound on the gradient (so some game-dependent constant). The computation of the proof uses the general notation η_t until the final steps. In the final steps we have a bound of the accumulated regret over T periods being:

$$R_G(T) \leq ||F||^2 \frac{1}{2\eta_T} + \frac{||\nabla_C||^2}{2} \sum_{t=1}^{T} \eta_t.$$

Let $\eta_t = t^{-\alpha}$.

We have for $\alpha \in (0, 1)$

$$\sum_{t=1}^{T} \eta_t = \sum_{t=1}^{T} t^{-\alpha} \leq 1 + \int_{t=1}^{T} t^{-\alpha} dt = 1 + \frac{T^{1-\alpha}}{1-\alpha}$$

Therefore, $\frac{\sum_{t=1}^{T}\eta_t}{T} = \frac{1+\frac{T^{1-\alpha}}{1-\alpha}}{T}$, that vanishes at T grows,

The first expression divided by T gives $\frac{||F||^2\frac{1}{2\eta_T}}{T}$ which vanishes as T increases, as $\frac{T^\alpha}{T}$ does.

Observe that the last inequality does not hold for $\alpha = 1$. We are not aware of a result showing that for $\alpha = 1$ projected gradient is no-regret. For $\alpha = 1$ [Hazan et al., 2007] prove that projected gradient is no regret when the reward is strongly convex, which does not apply in our model, as the rewards are linear.

C A single player against a strategic opponent example

Playing $y = 1$ from full support will increase the probabilities of the last two actions and decrease the probability of the first, until the first strategy goes out of the support.

If we begin at a point where the probability of the first action is close to 1 (and the last two are close to 0), then the sum of step-sizes until the first action goes out of the support is approximately $x = \frac{2}{2-d}$, and then the probability of the last action is $\frac{1-2d}{2-d}$, which is bounded away from zero for any $d < \frac{1}{2}$.

Then the derivative changes to:

$$\dot{x}_i = \begin{pmatrix} \frac{2}{3} - \frac{4}{3}y + \frac{d}{3} \\ -\frac{1}{3} + \frac{2}{3}y + \frac{d}{3} \\ -\frac{1}{3} + \frac{2}{3}y - \frac{2}{3}d \end{pmatrix} = \begin{pmatrix} -\frac{2}{3} + \frac{d}{3} \\ \frac{1}{3} + \frac{d}{3} \\ \frac{1}{3} - \frac{2}{3}d \end{pmatrix} \tag{4}$$

Playing $y = 1$ after the first action goes out of the support. After the first action goes out of the support, if player 2 keeps playing $y = 1$, it will not go back to the support. In this case, the constant we subtract from the gradient (so that the sum is 1) can change the average only of the two actions that are in the support. Then, the derivative changes to:

$$\dot{x}_i = \begin{pmatrix} 0 - \frac{2-d}{2} \\ 1 - \frac{2-d}{2} \\ 1 - d - \frac{2-d}{2} \end{pmatrix} = \begin{pmatrix} -1 + \frac{d}{2} \\ +\frac{d}{2} \\ -\frac{d}{2} \end{pmatrix} \tag{5}$$

Eventually, the third action will disappear as well, and as long as player 2 keeps playing $y = 1$, the only action in the support will be the 2nd action.

Suppose that once only the second action is left in the support, player 2 switches to playing repeatedly $y = 0$, will result in the following derivative:

$$\dot{x}_i = \begin{pmatrix} \frac{2}{3} + \frac{d}{3} \\ -\frac{1}{3} + \frac{d}{3} \\ -\frac{1}{3} - \frac{2}{3}d \end{pmatrix} \tag{6}$$

This will (immediately) get the first action in the support, so the support is now the first and second actions. The third action will remain outside the support.

Playing repeatedly $y = 0$, when the first and second action are in the support, will lead to the following derivatives (subtracting the average of the first two actions).

$$\dot{x}_i = \begin{pmatrix} 1 - \frac{1}{2} \\ 0 - \frac{1}{2} \\ -d - \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \\ -\frac{1}{2} \\ -d - \frac{2}{2} \end{pmatrix} \tag{7}$$

This will go on until only the first action remains.

Suppose that after only the first action remains in the support, player 2 switches again to playing repeatedly $y = 1$. This will lead to the derivatives.

$$\dot{x}_i = \begin{pmatrix} -\frac{2}{3} + \frac{d}{3} \\ \frac{1}{3} + \frac{d}{3} \\ \frac{1}{3} - \frac{2}{3}d \end{pmatrix} \tag{8}$$

This will get both of the last actions into the support, back to playing $y = 1$ in full support, where we started.

We conclude that by player 2 switching between playing $y = 0$ for a large enough number of periods and playing $y = 1$ for large enough number of periods, the probability of the last action will have a limsup of at least $\frac{1-2d}{2-d}$.

D Continuous Time Analysis

We now establish rigorously that the counterexample presented in Section 3.3 exhibits a periodic orbit under continuous-time projected gradient dynamics. Recall the game:

	L	R
T	$(k, 0)$	$(0, 1)$
M	$(0, k)$	$(1, 0)$
B	$(-d, 1)$	$(1-d, 0)$

(9)

for $d = \frac{1}{5}$ and $k = 5$. Figure 1 depicts the trajectory suggested by numerical computations. The objective of this appendix is twofold: first, to verify that such a cycle exists for the continuous-time projected gradient dynamics, and second, to demonstrate that an open set of initial conditions converges to this cycle, establishing that it attracts a positive-measure basin.

Our approach relies on constructing an approximate periodic orbit through precise numerical computation, then proving that this approximate orbit lies sufficiently close to a true trajectory. To formalise this, we introduce the following definition.

Definition D.1. Let ϕ be a flow on a metric space (X, d). Given $\varepsilon > 0, T > 0$, and $x, y \in X$, an (ε, T)-chain from x to y with respect to ϕ and d is a pair of finite sequences $x = x_1, \dots, x_{n-1}, x_n = y \in X$ and $\tau_1, \dots, \tau_{n-1} \in [T, \infty)$, denoted together by $(x_1, \dots, x_n;\ \tau_1, \dots, \tau_{n-1})$ such that

$$d\big(\phi^{\tau_i}(x_i), x_{i+1}\big) < \varepsilon,$$

for $i = 1, 2, \dots, n-1$.

Intuitively, an (ε, T)-chain describes a sequence of points that approximately follows the flow ϕ, with each consecutive pair of points connected by a forward evolution of the dynamics up to an error of at most ε. The constraint $\tau_i \geq T$ prevents the accumulation of arbitrarily many small corrections over vanishingly short time intervals, which could otherwise produce sequences bearing little resemblance to genuine trajectories. Such chains provide a bridge between numerical approximations and rigorous dynamical analysis: by establishing that computed points form an (ε, T)-chain with sufficiently small ε and moderate T, we can leverage continuity of the flow to deduce the existence of nearby true orbits.

We now present the numerically computed trajectory that forms our approximate periodic orbit.

Proposition D.2. *Fix the parameter values $k = 5$ and $d = \frac{1}{5}$. Define the following points and times as:*

$$
\begin{aligned}
t_1 &= 0,\\
x_1 &= (1,0,0), \quad y_1 = \left(\tfrac{1}{6}, \tfrac{5}{6}\right)\\[4pt]
t_2 &= 0.1371722820224960,\\
x_2 &= \big(0.9860858564985280,\ 0.0139141435014720,\ 0\big), \quad y_2 = (0.1,\ 0.9)\\[4pt]
t_3 &= 0.4687739703009671,\\
x_3 &= \big(0.8401172405785831,\ 0.1200586202892916,\ 0.0398241391321253\big), \quad y_3 = (0,\ 1)\\[4pt]
t_4 &= 0.5402025417295385,\\
x_4 &= \big(0.7972600977214402,\ 0.1486300488607202,\ 0.0541098534178396\big), \quad y_4 = (0,\ 1)\\[4pt]
t_5 &= 1.4137524726528787,\\
x_5 &= \big(0.7307698183724514,\ 0.2692301816275486,\ 0\big),\\
y_5 &= (0.3120795975837672,\ 0.6879204024162328)\\[4pt]
t_6 &= 2.0228034995770646,\\
x_6 &= (1,0,0), \quad y_6 = (0.2290078031758348,\ 0.7709921968241652)\\[4pt]
t_7 &= 2.1474857725954009,\\
x_7 &= (1,0,0), \quad y_7 = \left(\tfrac{1}{6}, \tfrac{5}{6}\right).
\end{aligned}
$$

Let $\zeta_i = (x_i, y_i)$ for $i = 1, \ldots, 7$, and define $\tau_i = t_{i+1} - t_i$ for $i = 1, \ldots, 6$. Then $(\zeta_1, \ldots, \zeta_7;\ \tau_1, \ldots, \tau_6)$ is an $(10^{-16}, 1/20)$-chain from ζ_1 to ζ_7 with respect to the flow induced by the continuous-time projected gradient dynamics under the Euclidean metric.

Appendix E provides the proof of this proposition, where we verify each segment of the trajectory analytically.

The (ε, T)-chain constructed in Proposition D.2 exhibits several features crucial to our main result. First, the numerical values satisfy $\zeta_7 = \zeta_1$, indicating that the trajectory returns to its starting point with period approximately $t_7 \approx 2.147$. The precise periodicity is established rigorously in Proposition D.3 below.

Second, and most importantly, the dominated action B (corresponding to x_3) is played with strictly positive probability during segments 3 and 4 of the trajectory. Specifically, at time $t_3 \approx 0.469$, we have $x_3(t_3) \approx 0.040 > 0$, and this coordinate remains positive until time $t_5 \approx 1.414$. This demonstrates that dominated strategies persist along the cycle rather than being eliminated by the learning dynamics, contradicting standard folk wisdom about strategy elimination in gradient-based learning.

We now establish not only that this approximate orbit corresponds to a genuine periodic trajectory, but also that it attracts a neighbourhood of positive measure.

Proposition D.3. *Let $(\zeta_k)_{k=1}^{7}$ and $(t_k)_{k=1}^{7}$ denote the coordinates and hitting times computed in Proposition D.2. Define the set*

$$C := \{(x(t), y(t)) : t \in [0, \infty)\}, \tag{10}$$

where $(x(t), y(t))$ solves the continuous-time projected gradient ODE with initial condition $((1,0,0), (\frac{1}{6}, \frac{5}{6})) = \zeta_1$. Then, C is a cycle of period T', where $|T' - t_7| < 0.12$. Additionally, there exists a neighbourhood $\mathcal{N}$ of ζ_6 with positive Lebesgue measure such that every trajectory starting in $\mathcal{N}$ is contained in C for all times greater than T'.

The proof of Proposition D.3 is given in Appendix F.

E Dynamics over each segment

In this appendix, we prove the following proposition:

PROPOSITION D.2. *Fix the parameter values $k = 5$ and $d = \frac{1}{5}$. Define the following points and times as:*

$$
\begin{aligned}
t_1 &= 0, \\
x_1 &= (1, 0, 0), \quad y_1 = \left(\tfrac{1}{6}, \tfrac{5}{6}\right) \\
t_2 &= 0.1371722820224960, \\
x_2 &= \big(0.9860858564985280,\ 0.0139141435014720,\ 0\big), \quad y_2 = (0.1,\ 0.9) \\
t_3 &= 0.4687739703009671, \\
x_3 &= \big(0.8401172405785831,\ 0.1200586202892916,\ 0.0398241391321253\big), \quad y_3 = (0,\ 1) \\
t_4 &= 0.5402025417295385, \\
x_4 &= \big(0.7972600977214402,\ 0.1486300488607202,\ 0.0541098534178396\big), \quad y_4 = (0,\ 1) \\
t_5 &= 1.4137524726528787, \\
x_5 &= \big(0.7307698183724514,\ 0.2692301816275486,\ 0\big), \\
y_5 &= (0.3120795975837672,\ 0.6879204024162328) \\
t_6 &= 2.0228034995770646, \\
x_6 &= (1, 0, 0), \quad y_6 = (0.2290078031758348,\ 0.7709921968241652) \\
t_7 &= 2.1474857725954009, \\
x_7 &= (1, 0, 0), \quad y_7 = \left(\tfrac{1}{6}, \tfrac{5}{6}\right).
\end{aligned}
$$

Let $\zeta_i = (x_i, y_i)$ for $i = 1, \dots, 7$, and define $\tau_i = t_{i+1} - t_i$ for $i = 1, \dots, 6$. Then $(\zeta_1, \dots, \zeta_7;\ \tau_1, \dots, \tau_6)$ is an $(10^{-16}, 1/20)$-chain from ζ_1 to ζ_7 with respect to the flow induced by the continuous-time projected gradient dynamics under the Euclidean metric.

PROOF. We verify that the sequence $(\zeta_1, \dots, \zeta_7;\ \tau_1, \dots, \tau_6)$, where $\tau_i = t_{i+1} - t_i$, forms an $(10^{-16}, 1/20)$-chain for the continuous-time projected gradient dynamics. Our strategy is to show that each segment of the trajectory lies in a region where the dynamics admit explicit solutions, and that all numerical values are accurate to at least 16 decimal places. Direct calculation confirms that $\tau_i \geq 1/20$ for all $i \in \{1, \dots, 6\}$.

We analyse the chain by considering six trajectory segments, where segment i connects ζ_{i-1} to ζ_i for $i = 2, \dots, 7$. Throughout, we use Proposition A.1 to determine the applicable dynamical equations in each region.

Segment 2: From ζ_1 to ζ_2. In this segment, both players have interior mixed strategies with $x_1, x_2, y_1, y_2 \in (0, 1)$ and $x_3 = 0$. By Proposition A.1, the projected gradient dynamics reduce to

$$
\begin{aligned}
\dot{x}_1 &= \tfrac{1}{2}\big(k - (1 + k)y_2\big), \\
\dot{y}_2 &= \tfrac{1}{2}\big((1 + k)x_1 - k\big).
\end{aligned}
$$

This linear system has a unique rest point at $\left(\frac{k}{1+k}, \frac{k}{1+k}\right) = \left(\frac{5}{6}, \frac{5}{6}\right)$. The trajectories in the (x_1, y_2)-plane are circles centred at this point, traversed anticlockwise. Starting from $\zeta_1 = \big((1, 0, 0), (\frac{1}{6}, \frac{5}{6})\big)$, the initial point $(x_1, y_2) = (1, \frac{5}{6})$ lies at distance $\frac{1}{6}$ from the centre, so the trajectory follows a circle of radius $r_2 = \frac{1}{1+k} = \frac{1}{6}$.

The transition to segment 3 occurs when action x_3 begins to receive positive probability, which by Proposition A.1 requires $y_2 \geq \frac{k+2d}{1+k} = \frac{5.4}{6} = 0.9$. Thus, ζ_2 is determined by intersecting the circle with the horizontal line $y_2 = 0.9$:

$$\left(x_1 - \tfrac{5}{6}\right)^2 + \left(0.9 - \tfrac{5}{6}\right)^2 = \left(\tfrac{1}{6}\right)^2.$$

Since $\left(0.9 - \frac{5}{6}\right)^2 = \left(\frac{1}{15}\right)^2$, we obtain

$$x_1 = \tfrac{5}{6} + \sqrt{\tfrac{1}{36} - \tfrac{1}{225}}.$$

This yields $x_1 \approx 0.9860858564985280$, and thus $x_2 = 1 - x_1 \approx 0.0139141435014720$. The time $t_2 - t_1$ is computed by integrating the angular velocity around the circle, giving $\tau_1 = t_2 \approx 0.1371722820224960$. These values match ζ_2 as stated.

Segment 3: From ζ_2 to ζ_3. In this segment, all strategies are played with positive probability, so the dynamics evolve in the interior. For $k = 5$ and $d = \frac{1}{5}$, the solution takes the form

$$\begin{aligned}
S(t) &= c_1 - dt,\\
R(t) &= c_5 \cos(c_0 t) + c_6 \sin(c_0 t) + c_3 + c_4 t,\\
x_1(t) &= \tfrac{1}{3}\big(S(t) + 2R(t)\big),\\
x_3(t) &= \tfrac{1}{3}\big(S(t) - R(t)\big),\\
y_2(t) &= \tfrac{1}{1+k}\big(k + d - c_4 + c_5 c_0 \sin(c_0 t) - c_6 c_0 \cos(c_0 t)\big),
\end{aligned}$$

where the constants are determined by the initial conditions at ζ_2:

$$\begin{aligned}
c_0 &= \sqrt{\tfrac{(k+1)(k+3)}{6}} = 2\sqrt{2},\\
c_1 &= \tfrac{k+\sqrt{1-4d^2}}{k+1} = \tfrac{25+\sqrt{91}}{30},\\
c_3 &= \tfrac{k(k+3-2\sqrt{1-4d^2})}{(k+3)(k+1)} = \tfrac{20-\sqrt{91}}{24},\\
c_4 &= \tfrac{2kd}{k+3} = \tfrac{1}{4},\\
c_5 &= \tfrac{3\sqrt{1-4d^2}}{k+3} = \tfrac{3\sqrt{91}}{40},\\
c_6 &= -\tfrac{3d(k+1)}{(k+3)c_0} = -\tfrac{9\sqrt{2}}{80}.
\end{aligned}$$

The trajectory exits this region at time t_3 when $y_2(t_3) = 1$. Solving this equation yields

$$t^\star := t_3 - t_2 = \tfrac{1}{2\sqrt{2}}\left[\arcsin\left(\tfrac{7}{\sqrt{51}}\right) - \arctan\left(\tfrac{3}{\sqrt{42}}\right)\right] \approx 0.3316016882784711.$$

Evaluating the trajectory at this time gives $x_1(t_3) \approx 0.8401172405785831$ and $x_3(t_3) \approx 0.0398241391321253$, which together with $x_2(t_3) = 1 - x_1(t_3) - x_3(t_3)$ and $y_2(t_3) = 1$ (so $y_1(t_3) = 0$) recovers ζ_3 as stated.

Segment 4: From ζ_3 to ζ_4. Here we have $x_1, x_3 \in (0,1)$, $x_2 = 0$, and $y_2 = 1$ (so $y_1 = 0$). The dynamics simplify to

$$\begin{aligned}
\dot{x}_1 &= \tfrac{1}{3}(d-2) = -\tfrac{3}{5},\\
\dot{x}_3 &= \tfrac{1}{3}(1-2d) = \tfrac{1}{5}.
\end{aligned}$$

Integrating from ζ_3 gives $x_1(t) = x_1(t_3) - \frac{3}{5}(t - t_3)$ and $x_3(t) = x_3(t_3) + \frac{1}{5}(t - t_3)$.

The transition to segment 5 occurs when $\dot{y}_2$ becomes non-positive, which by Proposition A.1 happens when

$$(1+k)x_1 + (k-1)x_3 = k.$$

Substituting the explicit formulas and using $k = 5$, $d = \frac{1}{5}$:

$$6\big(x_1(t_3) - \tfrac{3}{5}(t - t_3)\big) + 4\big(x_3(t_3) + \tfrac{1}{5}(t - t_3)\big) = 5.$$

This simplifies to $-\frac{14}{5}(t - t_3) = 5 - 6x_1(t_3) - 4x_3(t_3)$. Remarkably, substituting the numerical values of $x_1(t_3)$ and $x_3(t_3)$ yields

$$t^\dagger := t_4 - t_3 = \tfrac{1}{14} \approx 0.0714285714285714$$

to very high precision (beyond 50 decimal places). Consequently,

$$\begin{aligned} x_1(t_4) &= x_1(t_3) - \tfrac{3}{5} \cdot \tfrac{1}{14} = x_1(t_3) - \tfrac{3}{70} \approx 0.7972600977214402, \\ x_3(t_4) &= x_3(t_3) + \tfrac{1}{5} \cdot \tfrac{1}{14} = x_3(t_3) + \tfrac{1}{70} \approx 0.0541098534178396, \end{aligned}$$

giving ζ_4 as stated.

Segment 5: From ζ_4 to ζ_5. The dynamics return to the interior. Evaluating the constants at ζ_4 with $k = 5$ and $d = \frac{1}{5}$ yields

$$\begin{aligned} c_0 &= 2\sqrt{2} \approx 2.8284271247461903, \\ c_1 &\approx 0.9054798045571194, \\ c_2 &\approx 0.7431502443036007, \\ c_3 &\approx 0.7431502443036008, \\ c_4 &= \tfrac{1}{4}, \\ c_5 &\approx 0 \quad (\text{numerically } |c_5| < 10^{-15}), \\ c_6 &= -\tfrac{21\sqrt{2}}{80} \approx -0.3712310601229374. \end{aligned}$$

The trajectory leaves this region at time t_5 when $x_3(t_5) = 0$. Solving numerically gives

$$t^\ddagger := t_5 - t_4 \approx 0.8735499309233403.$$

At this time, we obtain $x_1(t_5) \approx 0.7307698183724514$ and $y_2(t_5) \approx 0.6879204024162328$, with $x_2(t_5) \approx 0.2692301816275486$ and $y_1(t_5) \approx 0.3120795975837672$, recovering ζ_5 as stated.

Segment 6: From ζ_5 to ζ_6. We return to the boundary with $x_3 = 0$ and interior strategies $x_1, y_2 \in (0, 1)$. The dynamics again follow

$$\begin{aligned} \dot{x}_1 &= \tfrac{1}{2}\big(k - (1 + k)y_2\big), \\ \dot{y}_2 &= \tfrac{1}{2}\big((1 + k)x_1 - k\big), \end{aligned}$$

so (x_1, y_2) moves anticlockwise along a circle centred at $\big(\frac{5}{6}, \frac{5}{6}\big)$. The radius of this circle is the Euclidean distance from ζ_5 to the centre:

$$r_6 = \sqrt{\big(x_1(t_5) - \tfrac{5}{6}\big)^2 + \big(y_2(t_5) - \tfrac{5}{6}\big)^2} \approx 0.1779443595032799.$$

The transition to segment 7 occurs when $x_1 = 1$. To find the corresponding value of y_2, we solve

$$\big(1 - \tfrac{5}{6}\big)^2 + \big(y_2 - \tfrac{5}{6}\big)^2 = r_6^2.$$

This gives $y_2 - \frac{5}{6} = \pm\sqrt{r_6^2 - \frac{1}{36}}$. Since the trajectory moves anticlockwise from a point with $y_2(t_5) < \frac{5}{6}$, we take the negative solution, yielding $y_2(t_6) \approx 0.7709921968241652$. Thus $y_1(t_6) = 1 - y_2(t_6) \approx 0.2290078031758348$, which together with $x_6 = (1, 0, 0)$, which is correct to 10^{-16}, gives ζ_6 as stated.

Segment 7: From ζ_6 to ζ_7. In this final segment, we have $x_1 = 1$ (so $x_2 = x_3 = 0$) and $y_1, y_2 \in (0, 1)$. By Proposition A.1, the dynamics reduce to

$$\dot{y}_2 = \tfrac{1}{2}\big((1+k) - k\big) = \tfrac{1}{2}.$$

Thus y_2 increases linearly from $y_2(t_6) \approx 0.7709921968241652$ to $y_2(t_7) = \frac{5}{6}$. The time required is

$$\tau_6 = t_7 - t_6 = 2\left(\tfrac{5}{6} - y_2(t_6)\right) \approx 0.1246822730183363,$$

giving $t_7 \approx 2.1474857725954009$ and $\zeta_7 = \big((1,0,0), (\tfrac{1}{6}, \tfrac{5}{6})\big) = \zeta_1$ as stated.

Since all intermediate points and times have been verified to the stated precision, this completes the proof. □

F Proof of Proposition D.3

In the following, we prove Proposition D.3. Our strategy consists of three steps. First, we establish that a neighbourhood of ζ_6 forms a basin of attraction for a line segment $\mathcal{L}$ along which the dynamics are one-dimensional (Lemma F.1). Second, we bound the Lipschitz constant of the vector field to control error propagation through the cycle (Lemma F.3). Finally, we combine these results with the numerical computations from Proposition D.2 to verify that trajectories starting from ζ_1 return to the basin of attraction, thereby establishing both periodicity and the existence of an attracting neighbourhood.

Lemma F.1. *Recall the critical point $\zeta_6 = ((1,0,0), (1-y_2^*, y_2^*))$, where $y_2^* = 0.7709921968241652$. Define the line segment*

$$\mathcal{L} := \left\{(1,0,0) \times (1-y_2, y_2) : y_2 \in \left[y_2^*, \tfrac{k-d}{k+1}\right]\right\}.$$

Let $A := k - (k+1)y_2(0)$ and $\mathcal{T} := \frac{2}{k+1}(A-d)$. If the initial condition $(x(0), y(0))$ satisfies

(1) $x_1(0) \in \left(1 - \dfrac{A^2 - d^2}{2(k+1)}, 1\right]$, *and*

(2) $y_2(0) \in \left(\dfrac{2}{1-d^2} y_2^* - \dfrac{(1+d^2)(k-d)}{(k+1)(1-d^2)}, \dfrac{k-d}{k+1}\right)$,

then $(x(\mathcal{T}), y(\mathcal{T})) \in \mathcal{L}$. Furthermore, defining the first hitting times

$$T_x := \inf\{t \geq 0 : x_1(t) = 1\}, \qquad T_y := \inf\{t \geq 0 : y_2(t) = \tfrac{k-d}{k+1}\},$$

we have $T_x < \mathcal{T} \leq T_y$.

Proof. Let the initial condition be

$$x(0) = (x_1(0), x_2(0), x_3(0)), \qquad y(0) = (1 - y_2(0), y_2(0)),$$

satisfying conditions (1) and (2). We structure the proof in three parts: first, we verify that these conditions place the initial point in an appropriate region; second, we prove that the trajectory reaches the edge $\{x_1 = 1\}$ before y_2 exceeds $\frac{k-d}{k+1}$; third, we show that upon reaching this edge, y_2 has already surpassed y_2^*, which places the trajectory in $\mathcal{L}$.

We first verify the initial conditions. From condition (2), we have $y_2(0) < \frac{k-d}{k+1}$, which implies $A > d$, hence $1 - \frac{A^2-d^2}{2(k+1)} < 1$ as required. By using the lower bound on $y_2(0)$, one can show that $A < 0.5626$. Therefore,

$$x_1(0) > 1 - \frac{A^2 - d^2}{2(k+1)} > 0.975 > \frac{k}{k+1}.$$

We define the first hitting times

$$T_x := \inf\{t \geq 0 : x_1(t) = 1\}, \qquad T_y := \inf\{t \geq 0 : y_2(t) = \tfrac{k-d}{k+1}\},$$

and prove that $T_x < T_y$, which shows that the trajectory reaches the edge $\{x_1 = 1\}$ before y_2 can exceed $\frac{k-d}{k+1}$. If $x_1(0) = 1$, then $T_x = 0$ and the result follows trivially. Henceforth we assume $x_1(0) < 1$.

The dynamics in the relevant regions are:

$$y_2'(t) \in \begin{cases} \frac{1}{2} & \text{on the edge } x_1 = 1 \\ \frac{1}{2}((k+1)x_1(t) - k) & \text{on the face } x_3 = 0 \\ \frac{1}{2}((k+1)x_1(t) + (k-1)x_3(t) - k) & \text{in the relative interior} \\ x_1(t) - \frac{1}{2} & \text{on the face } x_2 = 0\text{, provided } y_2(t) \leq \frac{k-d}{k+1} \end{cases}$$

and

$$x_1'(t) \in \begin{cases} 0 & \text{on the edge } x_1 = 1 \\ \frac{1}{2}(k - (k+1)y_2(t)) & \text{on the face } x_3 = 0 \\ \frac{1}{3}(2k + d - 2(k+1)y_2(t)) & \text{in the relative interior} \\ \frac{1}{2}(k + d - (k+1)y_2(t)) & \text{on the face } x_2 = 0\text{, provided } y_2(t) \leq \frac{k-d}{k+1} \end{cases}$$

For $t \in [0, T_y]$, we have $y_2(t) \leq \frac{k-d}{k+1}$, so the maximal growth rate for y_2 is $\frac{1}{2}$ (attained on the edge $x_1 = 1$). Therefore,

$$y_2(t) \leq y_2(0) + \tfrac{1}{2}t.$$

Since $y_2(T_y) = \frac{k-d}{k+1}$, we obtain

$$\frac{k-d}{k+1} \leq y_2(0) + \tfrac{1}{2}T_y,$$

and consequently

$$T_y \geq 2(\tfrac{k-d}{k+1} - y_2(0)) = \frac{2}{k+1}(A - d).$$

Define $\mathcal{T} := \frac{2}{k+1}(A-d)$. We will next find a lower bound for $x_1(t)$ for $t \in [0, T_y]$. For $t \in [0, T_y]$, we have that

$$0 < \tfrac{1}{2}(k - (k+1)y_2(t)) \leq \tfrac{1}{2}(k + d - (k+1)y_2(t)) \leq \tfrac{1}{3}(2k + d - 2(k+1)y_2(t)).$$

Therefore,

$$x_1'(t) \geq \tfrac{1}{2}\big(k - (k+1)y_2(t)\big).$$

and thus, using the upper bound $y_2(t) \leq y_2(0) + \frac{1}{2}t$ in the right-hand side yields the differential inequality

$$x_1'(t) \geq \tfrac{1}{2}\big(k - (k+1)\big(y_2(0) + \tfrac{1}{2}t\big)\big) = \tfrac{1}{2}A - \tfrac{1}{4}(k+1)t.$$

Integrating from 0 to $t \in [0, T_y]$, for $t \leq T_x$, gives the lower bound:

$$\begin{aligned} x_1(t) &\geq x_1(0) + \int_0^t \Big(\tfrac{1}{2}A - \tfrac{1}{4}(k+1)s\Big)ds \\ &= x_1(0) + \tfrac{1}{2}At - \tfrac{1}{8}(k+1)t^2. \end{aligned}$$

Suppose, for contradiction, that $\mathcal{T} \le T_x$. Hence,

$$\begin{aligned} x_1(\mathcal{T}) &\ge x_1(0) + \tfrac{1}{2}A\left(\frac{2}{k+1}(A-d)\right) - \tfrac{1}{8}(k+1)\left(\frac{2}{k+1}(A-d)\right)^2 \\ &= x_1(0) + \frac{A(A-d)}{k+1} - \frac{(A-d)^2}{2(k+1)} \\ &= x_1(0) + \frac{2A(A-d)-(A-d)^2}{2(k+1)} \\ &= x_1(0) + \frac{A^2-d^2}{2(k+1)} \end{aligned}$$

By condition (1), we have $x_1(0) > 1 - \frac{A^2-d^2}{2(k+1)}$. Therefore $x_1(\mathcal{T}) > 1$, which contradicts $x_1(t) \in [0,1]$ for all t. Therefore, our assumption $\mathcal{T} \le T_x$ must be false, so $T_x < \mathcal{T}$. Thus,

$$T_x < \mathcal{T} \le T_y. \tag{11}$$

Having established that the trajectory reaches $x_1 = 1$ by time $\mathcal{T}$, we now show that $y_2(\mathcal{T}) \ge y_2^*$, which implies $(x(\mathcal{T}), y(\mathcal{T})) \in \mathcal{L}$.

From the initial conditions assumptions, we have

$$\begin{aligned} x_1(0) &> 1 - \frac{A^2-d^2}{2(k+1)} \\ &= 1 - \frac{(k-(k+1)y_2(0))^2-d^2}{2(k+1)} \\ &> 1 - \frac{1-d^2}{2(k+1)}. \end{aligned}$$

For $t \in [0,\mathcal{T}]$, the minimum growth rate for y_2 is $\frac{1}{2}((k+1)x_1(t)-k)$, which is bounded below by

$$\begin{aligned} y_2'(t) &\ge \tfrac{1}{2}((k+1)x_1(0)-k) \\ &> \tfrac{1}{2}\left((k+1)\left(1-\frac{1-d^2}{2(k+1)}\right)-k\right) \\ &= \frac{1+d^2}{4}. \end{aligned}$$

Integrating from 0 to $\mathcal{T}$,

$$\begin{aligned} y_2(\mathcal{T}) &\ge y_2(0) + \frac{2}{k+1}(A-d)\cdot\frac{1+d^2}{4} \\ &= y_2(0) + \frac{1+d^2}{2(k+1)}\big(k-(k+1)y_2(0)-d\big) \\ &= y_2(0) + \frac{1+d^2}{2}\left(\frac{k-d}{k+1}-y_2(0)\right) \\ &= \frac{1-d^2}{2}y_2(0) + \frac{1+d^2}{2}\cdot\frac{k-d}{k+1}. \end{aligned}$$

By condition (2), we have $y_2(\mathcal{T}) \ge y_2^*$. Hence, $(x(\mathcal{T}), y(\mathcal{T})) \in \mathcal{L}$. □

We then have the following corollary:

Corollary F.2. *For $k = 5$ and $d = \frac{1}{5}$, let $\zeta_6 = ((1,0,0),(1-y_2^*, y_2^*))$ with $y_2^* = 0.7709921968241652$, and define $\tau_6 := 2(\frac{k}{k+1} - y_2^*)$ as the travel time along $\mathcal{L}$ from ζ_6 to $\zeta_1 = ((1,0,0),(\frac{1}{k+1}, \frac{k}{k+1}))$.*

Then, for $\varepsilon_0 = 10^{-5/2}$, the trajectory starting from any initial condition $(x(0), y(0))$ satisfying $\|(x(0), y(0)) - \zeta_6\| < \varepsilon_0$, enters the line segment $\mathcal{L}$ in finite time bounded by $\mathcal{T}_{\max} \leq 0.12$, and subsequently reaches ζ_1 by total time $t_{(x(0),y(0))} \in (\tau_6 - 2\varepsilon_0, \tau_6 + 0.12)$.

Proof. We verify that the ball of radius $\varepsilon_0 = 10^{-5/2}$ around ζ_6 satisfies the conditions of Lemma F.1, then compute the maximum entry time.

Condition (2) of Lemma F.1 requires

$$y_2 \in \left(\frac{2}{1-d^2} y_2^* - \frac{(1+d^2)(k-d)}{(k+1)(1-d^2)}, \frac{k-d}{k+1}\right).$$

Hence, it suffices to have $y_2 \in (0.7396, 0.8)$. This condition is certainly satisfied if $|y_2 - y_2^*| < 0.003$.

Condition (1) of Lemma F.1 requires

$$x_1 \in \left(1 - \frac{(k-(k+1)y_2^2 - d^2}{2(k+1)}, 1\right].$$

If $|y_2 - y_2^*| < 0.003$, then it suffices to have

$$x_1 \in \left(1 - \frac{(k-(k+1)(y_2^* + 0.003))^2 - d^2}{2(k+1)}, 1\right].$$

For this, we can take $|x_1 - 1| < 0.0072$.

The square of the Euclidean distance in $X \times Y$ is given by

$$\begin{aligned}
\|(x,y) - \zeta_6\|^2 &= \|x - (1,0,0)\|^2 + \|y - (1-y_2^*, y_2^*)\|^2 \\
&= (x_1 - 1)^2 + x_2^2 + x_3^2 + (y_1 - (1-y_2^*))^2 + (y_2 - y_2^*)^2 \\
&= 2x_2^2 + 2x_2x_3 + 2x_3^2 + 2(y_2 - y_2^*)^2.
\end{aligned}$$

Note that $2x_2^2 + 2x_2x_3 + 2x_3^2 = \frac{3}{2}(x_2+x_3)^2 + \frac{1}{2}(x_2-x_3)^2$. Therefore, $2x_2^2 + 2x_2x_3 + 2x_3^2 \geq \frac{3}{2}(x_2+x_3)^2 = \frac{3}{2}(1-x_1)^2$.

Hence, if $\|(x,y) - \zeta_6\|^2 < 10^{-5}$, then $2(y_2 - y_2^*)^2 < 10^{-5}$ and $\frac{3}{2}(1-x_1)^2 < 10^{-5}$. Thus, if $\|(x,y) - \zeta_6\|^2 < 10^{-5}$, then $|y_2 - y_2^*| < 0.003$ and $|x_1 - 1| < 0.0072$. As both conditions (1) and (2) are satisfied, we conclude that every initial condition in $B(\zeta_6, \varepsilon_0)$ lies in the local basin of Lemma F.1, and the trajectory enters $\mathcal{L}$ in finite time.

By Lemma F.1, the trajectory enters $\mathcal{L}$ by time $\mathcal{T} = \frac{2}{k+1}(k - (k+1)y_2 - d)$. This expression is decreasing in y_2, so its maximum over $B(\zeta_6, \varepsilon_0)$ is

$$\mathcal{T}_{\max} \leq \frac{2}{k+1}(k - (k+1)(y_2^* - 0.003) - d) \leq 0.12.$$

We now establish bounds on the total time to reach ζ_1. For the upper bound, observe that at time $\mathcal{T} \leq \mathcal{T}_{\max}$, the trajectory has entered $\mathcal{L}$ at some point $((1,0,0),(1-y_2(\mathcal{T}), y_2(\mathcal{T})))$ with $y_2(\mathcal{T}) \geq y_2^*$. From this point, the trajectory reaches ζ_1 in additional time at most $2(\frac{k}{k+1} - y_2^*) = \tau_6$. Therefore, $t_{(x(0),y(0))} \leq \mathcal{T}_{\max} + \tau_6 \leq \tau_6 + 0.12$.

For the lower bound, we note that to reach ζ_1, the y_2 coordinate must increase from $y_2(0)$ to $\frac{k}{k+1}$. Since $y_2' \leq \frac{1}{2}$ throughout, we have

$$t_{(x(0),y(0))} \geq 2\left(\frac{k}{k+1} - y_2(0)\right) \geq 2\left(\frac{k}{k+1} - (y_2^* + \varepsilon_0)\right) = \tau_6 - 2\varepsilon_0.$$

Therefore, $t_{(x(0),y(0))} \in (\tau_6 - 2\varepsilon_0, \tau_6 + 0.12)$. □

Having established that trajectories near ζ_6 enter the line segment $\mathcal{L}$, we now turn to proving that there exists a cycle that is approximated by the (ε, T)-chain in Proposition D.2. To make this argument rigorous, we require a bound on the Lipschitz constant of the vector field, which will allow us to control the propagation of numerical errors when tracking the trajectory through each segment.

LEMMA F.3. *Fix $k = 5$ and $d = 1/5$. The projected gradient vector field on $X \times Y$ has Lipschitz constant $K \leq 2\sqrt{5}$.*

PROOF. We compute the Lipschitz constant by examining the Jacobian of the vector field on each segment of the cycle. Since the dynamics are piecewise smooth, the Lipschitz constant is bounded by the supremum of the Jacobian operator norms across all segments.

Using the parameterisation (x_1, x_3, y_2) (where $x_2 = 1 - x_1 - x_3$ and $y_1 = 1 - y_2$), we compute the Jacobian matrix J of the vector field $(\dot{x}_1, \dot{x}_3, \dot{y}_2)$ on each segment. Then, for each segment, the operator norm of J is $\|J\| = \sqrt{\lambda_{\max}(J^\top J)}$.

Case One: Segments S_7 and S_4. On segment S_7, we have $x_1 = 1, x_3 = 0$, and the dynamics reduce to $\dot{y}_2 = 1/2$. On segment S_4, we have $y_2 = 1$ and the dynamics are $\dot{x}_1 = \frac{1}{3}(d-2)$ and $\dot{x}_3 = \frac{1}{3}(1-2d)$. In both cases, the Jacobian is $J = 0$, and so $\|J\| = 0$.

Case Two: Segments S_2 and S_6. In these segments, the dynamics are

$$\dot{x}_1 = \frac{1}{2}(k - (1+k)y_2), \quad \dot{x}_3 = 0, \quad \dot{y}_2 = \frac{1}{2}((1+k)x_1 - k). \tag{12}$$

Taking partial derivatives with respect to (x_1, x_3, y_2), the Jacobian is

$$J = \begin{bmatrix} 0 & 0 & -\frac{1+k}{2} \\ 0 & 0 & 0 \\ \frac{1+k}{2} & 0 & 0 \end{bmatrix}. \tag{13}$$

Substituting $k = 5$ gives $\frac{1+k}{2} = 3$, we have:

$$J^\top J = \begin{bmatrix} 0 & 0 & 3 \\ 0 & 0 & 0 \\ -3 & 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 & -3 \\ 0 & 0 & 0 \\ 3 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 9 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 9 \end{bmatrix}. \tag{14}$$

The eigenvalues of $J^\top J$ are $\{9, 0, 9\}$, so $\|J\| = \sqrt{9} = 3$.

Case Three: Segments S_3 and S_5. When $(x_1, x_3, y_2) \in (0,1)^3$, the dynamics are

$$\dot{x}_1 = \frac{1}{3}(2k + d - 2(1+k)y_2), \tag{15}$$

$$\dot{x}_3 = \frac{1}{3}((1+k)y_2 - 2d - k), \tag{16}$$

$$\dot{y}_2 = \frac{1}{2}((1+k)x_1 + (k-1)x_3 - k). \tag{17}$$

With $k = 5$ and $d = 1/5$, the Jacobian is:

$$J = \begin{bmatrix} 0 & 0 & -4 \\ 0 & 0 & 2 \\ 3 & 2 & 0 \end{bmatrix}. \tag{18}$$

We compute

$$J^\top J = \begin{bmatrix} 0 & 0 & 3 \\ 0 & 0 & 2 \\ -4 & 2 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 & -4 \\ 0 & 0 & 2 \\ 3 & 2 & 0 \end{bmatrix} = \begin{bmatrix} 9 & 6 & 0 \\ 6 & 4 & 0 \\ 0 & 0 & 20 \end{bmatrix}. \tag{19}$$

Since $J^\top J$ is block diagonal, its eigenvalues are the eigenvalues of the 2×2 block $\left[\begin{smallmatrix} 9 & 6 \\ 6 & 4 \end{smallmatrix}\right]$ together with the $(3, 3)$ entry. For the 2×2 block, the characteristic polynomial is

$$\det \begin{bmatrix} 9-\lambda & 6 \\ 6 & 4-\lambda \end{bmatrix} = (9-\lambda)(4-\lambda) - 36 = \lambda^2 - 13\lambda = \lambda(\lambda - 13), \tag{20}$$

giving eigenvalues 0 and 13. Combined with the eigenvalue 20 from the $(3, 3)$ entry, the eigenvalues of $J^\top J$ are $\{0, 13, 20\}$. Therefore,

$$\|J\| = \sqrt{20} = 2\sqrt{5} \approx 4.472. \tag{21}$$

Taking the maximum across all segments, we obtain $K = \max\{0, 3, 2\sqrt{5}\} = 2\sqrt{5}$. □

We will also require the following result, which establishes that an approximate closed chain in a Lipschitz flow yields a periodic orbit that attracts all trajectories in a neighbourhood, provided the exponentially-accumulated chain error is smaller than this neighbourhood.

Theorem F.4. *Let ϕ be a flow on a metric space (X, d) generated by a Lipschitz vector field with constant $K > 0$. Let $\varepsilon, T > 0$ and $(\zeta_1, \ldots, \zeta_n; \tau_1, \ldots, \tau_{n-1})$ be an (ε, T)-chain with $\zeta_n = \zeta_1$. Suppose that there exists $\varepsilon' > 0$ and a neighbourhood $\mathcal{N}$ of ζ_{n-1} such that:*

(1) *$\mathcal{N}$ is open with* $\text{radius}(\mathcal{N}) := \sup\{r > 0 : B_r(\zeta_{n-1}) \subseteq \mathcal{N}\}$;

(2) *For every $x_0 \in \mathcal{N}$, there exists a unique time $t_{x_0} \in (\tau_{n-1} - \varepsilon', \tau_{n-1} + \varepsilon')$ such that $\phi^{t_{x_0}}(x_0) = \zeta_1$.*

Let $t_k := \sum_{i=1}^{k-1} \tau_i$ for $k = 2, \ldots, n$ with $t_1 := 0$ denote the cumulative times along the chain, and define

$$E := \varepsilon \sum_{k=2}^{n-1} e^{K(t_{n-1} - t_k)}.$$

If $E <$ $\text{radius}(\mathcal{N})$*, then:*

(1) *The flow admits a periodic orbit C passing through ζ_1;*

(2) *The period T' of C satisfies $|T' - t_n| < \varepsilon'$;*

(3) *For every $x \in \mathcal{N}$, there exists $t_x \leq \tau_{n-1} + \varepsilon'$ such that $\phi^{t_x}(x) \in C$, and $\phi^t(x) \in C$ for all $t \geq t_x$.*

The proof of this is found in Appendix H

With the above result an Lipschitz constant established, we can now bound the accumulated error as a trajectory flows through the sequence of segments $S_2, \ldots, S_7$. This enables us to prove that trajectories starting from ζ_1 return sufficiently close to ζ_6 to re-enter the basin of attraction established in Lemma F.1, thereby completing the cycle.

Proposition D.3. *Let $(\zeta_k)_{k=1}^7$ and $(t_k)_{k=1}^7$ denote the coordinates and hitting times computed in Proposition D.2. Define the set*

$$C := \{(x(t), y(t)) : t \in [0, \infty)\}, \tag{10}$$

where $(x(t), y(t))$ solves the continuous-time projected gradient ODE with initial condition $((1, 0, 0), (\frac{1}{6}, \frac{5}{6})) = \zeta_1$. Then, C is a cycle of period T', where $|T' - t_7| < 0.12$. Additionally, there exists a neighbourhood $\mathcal{N}$ of ζ_6 with positive Lebesgue measure such that every trajectory starting in $\mathcal{N}$ is contained in C for all times greater than T'.

Proof. By Proposition D.2, the sequence $(\zeta_1, \ldots, \zeta_7; \tau_1, \ldots, \tau_6)$ forms a $(10^{-16}, 1/20)$-chain with $\zeta_7 = \zeta_1$, where $\tau_k := t_{k+1} - t_k$ for $k = 1, \ldots, 6$. Furthermore, by Lemma F.3, the projected gradient flow ϕ on $X \times Y$ is generated by a Lipschitz vector field with constant $K = 2\sqrt{5}$.

Define $\varepsilon := 10^{-16}$ and let

$$E := \varepsilon \left(\sum_{j=2}^{6} e^{K(t_6 - t_j)} \right).$$

Using the times from Proposition D.2 and $K = 2\sqrt{5}$, we compute $E \le 7 \times 10^{-13}$. From Corollary F.2, there exists $\varepsilon_0 = 10^{-5/2}$ such that the trajectory starting from any initial condition $(x(0), y(0))$ satisfying $\|(x(0), y(0)) - \zeta_6\| < \varepsilon_0$ enters the line segment $\mathcal{L}$ in finite time bounded by $\mathcal{T}_{\max} \le 0.12$, and subsequently reaches ζ_1 by total time $t_{(x(0),y(0))} \in (\tau_6 - 2\varepsilon_0, \tau_6 + 0.12)$. Furthermore, as $\tau_6 > 0.12$, we have $t_{(x_0,y_0)} \in (0, 2\tau_6)$.

We now apply Theorem H.1. As $(0, 2\tau_6) = (\tau_6 - \tau_6, \tau_6 + \tau_6)$, the above establishes conditions (i) and (ii) with $\varepsilon' = \tau_6$. Since $E < 7 \times 10^{-13} < 10^{-5/2} = \text{radius}(\mathcal{N})$, we can therefore conclude:

(1) The flow admits a periodic orbit C passing through ζ_1.
(2) The period T' of C satisfies $|T' - t_7| < \tau_6$.
(3) For every $(x, y) \in \mathcal{N}$, there exists $t_{(x,y)} \le 2\tau_6$ such that $\phi^{t_{(x,y)}}(x, y) \in C$, and $\phi^t(x, y) \in C$ for all $t \ge t_{(x,y)}$.

As $\tau_6 < 0.125$ and $2\tau_6 < t_7 - \tau_6 < T'$, we obtain the desired result. □

G Proof of Theorem 3.1

We finally return our attention to proving Theorem 3.1.

Proof. Let C be the cycle and $\mathcal{N}$ be the neighbourhood of ζ_6 defined in Proposition D.3. Therefore, $\mathcal{N}$ has positive Lebesgue measure and every continuous time trajectory starting in $\mathcal{N}$ is contained in C by time T'.

Since $\mathcal{N}$ is an open neighbourhood of ζ_6, pick $r \in (0, 0.1)$ such that the closed ball $\overline{B(\zeta_6, r)} \subset \mathcal{N}$. Choose $z^0 = (x^0, y^0) \in \overline{B(\zeta_6, r)}$. Denote $z^n = (x^n, y^n)$ for $n \in \mathbb{N}$ as the solution in discrete time starting from z^0 and $z(t) = (x(t), y(t))$ for $t \ge 0$ as the solution in continuous time starting from z^0. Define $\tau_0 = 0$, $\tau_n = \sum_{i=0}^{n-1} \gamma^i$ and the piecewise constant interpolation $Z^t = z^n$ for $t \in [\tau_n, \tau_{n+1})$.

We will construct recursively a sequence of return times $(t_k)_{k \ge 1}$ with the properties

(i) $t_1 \le 2T'$ and $t_{k+1} - t_k \le 2T'$ for every $k \ge 1$,
(ii) $Z^{t_k} \in \overline{B(\zeta_6, r/2)}$ for every $k \ge 1$,
(iii) $\sup_{u \in [0,2T']} \|Z^{t_k+u} - z^{(Z^{t_k})}(u)\| < \frac{r}{2}$ for every $k \ge 1$,

where $z^{(Z^{t_k})}(\cdot)$ denotes the continuous solution started at $z^{(Z^{t_k})}(0) = Z^{t_k}$.

The construction is by induction.

Base Case:

By Theorem 3 in [Dupuis and Nagurney, 1993], there exists $\delta > 0$ such that if $\max_t \gamma^t < \delta$ the interpolated process satisfies

$$\sup_{t \in [0,2T']} \|Z^t - z(t)\| < \tfrac{r}{2}.$$

As $\overline{B(\zeta_6, r)}$ is compact, take the minimum such $\delta > 0$ over initial conditions in $\overline{B(\zeta_6, r)}$. By Proposition D.3, we have $z(T') \in C$. Therefore, as the period of the cycle is T', we have that there exists $t' \in [0, 2T']$ such that $z(t') = \zeta_6$. Hence, there exists a time $t_1 \in [0, 2T']$ such that $Z^{t_1} \in \overline{B(\zeta_6, r/2)}$.

By Theorem 3 in [Dupuis and Nagurney, 1993], by our choice of $\delta > 0$ (uniform over initial conditions in the compact set $\overline{B(\zeta_6, r)}$) we have,

$$\sup_{u \in [0,2T']} \|Z^{t_1+u} - z^{(Z^{t_1})}(u)\| < \frac{r}{2}.$$

Inductive step: Assume t_k satisfies the inductive assumptions. By Proposition D.3 the continuous trajectory $z^{(Z^{t_k})}(\cdot)$ enters the cycle C by time T' and thus hits ζ_6 at some time $s_k \in [0, 2T']$. By (iii) we have

$$\|Z^{t_k+s_k} - \zeta_6\| \le \|Z^{t_k+s_k} - z^{(Z^{t_k})}(s_k)\| + \|z^{(Z^{t_k})}(s_k) - \zeta_6\| < \frac{r}{2}.$$

Set $t_{k+1} := t_k + s_k$. Then $t_{k+1} - t_k = s_k \leq 2T'$ and $Z^{t_{k+1}} \in \overline{B(\zeta_6, r/2)}$. Moreover, the same uniform approximation guarantee (Theorem 3 applied to the shifted sequence starting at the update index corresponding to t_{k+1}) yields (iii) for $k+1$ as well. This completes the induction and produces the desired sequence $(t_k)_{k\geq 1}$.

As $Z^{t_k} \in \overline{B(\zeta_6, r/2)}$ for every $k \geq 1$, we have $Z^{t_k} \in \mathcal{N}$. Therefore, $z^{(Z^{t_k})}(t_k + u) \in C$ for $u \in [T', 2T']$ by Proposition D.3. By (iii) the interpolated trajectory Z is at distance at most $r/2$ from $z^{(Z^{t_k})}(t_k + u)$ at these times. Therefore, Z^t is within $r/2$ of C for periods of time T' infinitely often. Hence, as the cycle C is of period T', this implies Z^t passes within $r/2$ of every point on C infinitely often.

Finally, because C contains a segment on which the dominated action has mass at least 0.054 and $r < 0.1$, the uniform $r/2$–closeness implies that the discrete iterates place strictly positive mass on that action infinitely often.

Therefore, the discrete iterates $z^n = (x^n, y^n)$ satisfy $x_3^n > 0$ infinitely often, which implies

$$\limsup_{n\to\infty} x_3^n > 0.$$

Since action B (corresponding to x_3) is strictly dominated, this completes the proof. □

H Periodic Orbits and (ε, T)-chains

Theorem H.1. *Let ϕ be a flow on a metric space (X, d) generated by a Lipschitz vector field with constant $K > 0$. Let $\varepsilon, T > 0$ and $(\zeta_1, \ldots, \zeta_n; \tau_1, \ldots, \tau_{n-1})$ be an (ε, T)-chain with $\zeta_n = \zeta_1$. Suppose that there exists $\varepsilon' > 0$ and a neighbourhood $\mathcal{N}$ of ζ_{n-1} such that:*

(1) *$\mathcal{N}$ is open with* $\mathrm{radius}(\mathcal{N}) := \sup\{r > 0 : B_r(\zeta_{n-1}) \subseteq \mathcal{N}\}$;
(2) *For every $x_0 \in \mathcal{N}$, there exists a unique time $t_{x_0} \in (\tau_{n-1} - \varepsilon', \tau_{n-1} + \varepsilon')$ such that $\phi^{t_{x_0}}(x_0) = \zeta_1$.*

Let $t_k := \sum_{i=1}^{k-1} \tau_i$ for $k = 2, \ldots, n$ with $t_1 := 0$ denote the cumulative times along the chain, and define

$$E := \varepsilon \sum_{k=2}^{n-1} e^{K(t_{n-1} - t_k)}.$$

If $E <$ radius$(\mathcal{N})$, then:

(1) *The flow admits a periodic orbit C passing through ζ_1;*
(2) *The period T' of C satisfies $|T' - t_n| < \varepsilon'$;*
(3) *For every $x \in \mathcal{N}$, there exists $t_x \leq \tau_{n-1} + \varepsilon'$ such that $\phi^{t_x}(x) \in C$, and $\phi^t(x) \in C$ for all $t \geq t_x$.*

Proof. Since ϕ is generated by a Lipschitz vector field with constant $K > 0$, the flow satisfies the standard estimate for all $x, y \in X$ and $t \geq 0$:

$$d(\phi^t(x), \phi^t(y)) \leq e^{Kt} d(x, y).$$

We inductively bound $d_k := d(\phi^{t_k}(\zeta_1), \zeta_k)$ for $k = 1, \ldots, n-1$. At $k = 1$, we have $d_1 = d(\phi^0(\zeta_1), \zeta_1) = 0$.

We inductively bound $d(\phi^{t_k}(\zeta_1), \zeta_k)$ for $k = 1, \ldots, n-1$. At $k = 1$, we have $d(\phi^0(\zeta_1), \zeta_1) = 0$. For $k \geq 2$, using the triangle inequality, the definition of an (ε, T)-chain and the Lipschitz property:

$$\begin{aligned}
d_k &= d(\phi^{t_k}(\zeta_1), \zeta_k) \\
&= d(\phi^{\tau_{k-1}}(\phi^{t_{k-1}}(\zeta_1)), \zeta_k) \\
&\leq d(\phi^{\tau_{k-1}}(\phi^{t_{k-1}}(\zeta_1)), \phi^{\tau_{k-1}}(\zeta_{k-1})) + d(\phi^{\tau_{k-1}}(\zeta_{k-1}), \zeta_k) \\
&< e^{K\tau_{k-1}} \cdot d(\phi^{t_{k-1}}(\zeta_1), \zeta_{k-1}) + \varepsilon \\
&= e^{K\tau_{k-1}} d_{k-1} + \varepsilon.
\end{aligned}$$

Unrolling this recurrence from $k = 2$ to $k = n - 1$:

$$\begin{aligned}
d_2 &< \varepsilon, \\
d_3 &< e^{K\tau_2}\varepsilon + \varepsilon = \varepsilon(1 + e^{K\tau_2}), \\
d_4 &< e^{K\tau_3}(e^{K\tau_2}\varepsilon + \varepsilon) + \varepsilon = \varepsilon(1 + e^{K\tau_3} + e^{K(\tau_2+\tau_3)}), \\
&\vdots \\
d_{n-1} &< \varepsilon\Big(1 + e^{K\tau_{n-2}} + e^{K(\tau_{n-3}+\tau_{n-2})} + \cdots + e^{K(\tau_2+\cdots+\tau_{n-2})}\Big).
\end{aligned}$$

As $t_k := \sum_{i=1}^{k-1} \tau_i$ for $k = 2, \ldots, n$ with $t_1 := 0$, we have that $\tau_i = t_{i+1} - t_i$ for $i = 1, \ldots, n-1$. Therefore,

$$\begin{aligned}
d(\phi^{t_{n-1}}(\zeta_1), \zeta_{n-1}) &< \varepsilon\Big(1 + e^{K(t_{n-1}-t_{n-2})} + e^{K(t_{n-1}-t_{n-3})} + \cdots + e^{K(t_{n-1}-t_2)}\Big) \\
&= \varepsilon\left(1 + \sum_{j=2}^{n-2} e^{K(t_{n-1}-t_j)}\right) \\
&= \varepsilon \sum_{j=2}^{n-1} e^{K(t_{n-1}-t_j)} = E.
\end{aligned}$$

Since $E < \text{radius}(\mathcal{N})$ and $d(\phi^{t_{n-1}}(\zeta_1), \zeta_{n-1}) < E$ with $\zeta_{n-1} \in \mathcal{N}$, we have $\phi^{t_{n-1}}(\zeta_1) \in \mathcal{N}$. Therefore, the trajectory starting from $\phi^{t_{n-1}}(\zeta_1)$ reaches ζ_1 at time $t_{\phi^{t_{n-1}}(\zeta_1)} \in (\tau_{n-1} - \varepsilon', \tau_{n-1} + \varepsilon')$. That is, $\phi^{t_{n-1}+t_{\phi^{t_{n-1}}(\zeta_1)}}(\zeta_1) = \zeta_1$. Therefore, the trajectory starting from ζ_1 returns to ζ_1 at time $T' = t_{n-1} + t_{\phi^{t_{n-1}}(\zeta_1)}$, and the set

$$C := \{\phi^t(\zeta_1) : t \in [0, T']\}$$

is a periodic orbit of period T'. We have $T' = t_{n-1} + t_{\phi^{t_{n-1}}(\zeta_1)}$ and $t_n = t_{n-1} + \tau_{n-1}$. Therefore,

$$\begin{aligned}
|T' - t_n| &= |t_{n-1} + t_{\phi^{t_{n-1}}(\zeta_1)} - (t_{n-1} + \tau_{n-1})| \\
&= |t_{\phi^{t_{n-1}}(\zeta_1)} - \tau_{n-1}| \\
&< \varepsilon'.
\end{aligned}$$

To establish conclusion (3), let $x \in \mathcal{N}$ be arbitrary. Then, the trajectory starting from x reaches ζ_1 at time $t_x \in (\tau_{n-1} - \varepsilon', \tau_{n-1} + \varepsilon')$, that is, $\phi^{t_x}(x) = \zeta_1 \in C$. This shows that the trajectory from x reaches the periodic orbit C by time $t_x \leq \tau_{n-1} + \varepsilon'$. Once on C, the trajectory remains on C for all future times by periodicity. □