FRACTIONAL REPLICATOR DYNAMICS

Documents de travail GREDEG
GREDEG Working Papers Series

MAXIME MENUET

GREDEG WP No. 2026-05
https://ideas.repec.org/s/gre/wpaper.html

Fractional Replicator Dynamics

Maxime Menuet*

February 16, 2026

Abstract

This paper studies how the temporal structure of adjustment shapes evolutionary dynamics in symmetric games. We introduce a fractional replicator dynamic that modifies the classical replicator only through the time operator, replacing the ordinary derivative with a fractional derivative of order $\alpha \in (0, 1]$. This formulation preserves payoff monotonicity, feasibility, and the equilibrium set, while allowing past payoff differences to affect current behavior through long memory with power-law decay. We show that fractional time fundamentally alters local stability and equilibrium selection. While evolutionarily stable strategies remain locally asymptotically stable, equilibria that are unstable under the classical replicator can become locally stable when memory is sufficiently persistent, generating stability switching and purely temporal bifurcations without any change in payoffs or strategic interaction. Moreover, convergence toward stable equilibria becomes polynomial rather than exponential, implying slow adjustment even in simple games. As a consequence, equilibrium selection may fail to be completed over economically relevant horizons despite being guaranteed asymptotically. Finally, we provide a microfoundation based on standard payoff-monotone revision processes with heterogeneous and asynchronous revision opportunities. Fractional dynamics thus offer a parsimonious way to incorporate persistent memory into evolutionary models while preserving their core economic structure.

---

*University Côte d'Azur, GREDEG, CNRS, France maxime.menuet@univ-cotedazur.fr.

1 Introduction

Evolutionary game theory provides a canonical framework for analyzing how strategic behavior evolves in large populations. A central object of this framework is the replicator dynamics, a payoff-monotone selection rule under which strategies that perform above the population average expand, while underperforming strategies contract. Because it offers a robust link between individual incentives and aggregate outcomes, the replicator dynamics has become a benchmark model in economics, biology, and related fields (Taylor and Jonker, 1978; Weibull, 1997). A key implication of this approach is that evolutionary selection is entirely incentive-driven. Long-run outcomes are determined by the strategic structure of the game: dominated strategies are eliminated, evolutionarily stable strategies are locally attracting, and coordination games converge to stable conventions (Weibull, 1997; Hofbauer and Sigmund, 1998). The temporal structure of adjustment plays no role in determining local stability. Time may affect the speed of convergence, but not which states are selected. In this sense, standard evolutionary models treat time as neutral.

However, this neutrality is strong and empirically nontrivial. It follows directly from the memoryless formulation of the replicator as a first-order differential equation. Strategy shares respond only to contemporaneous payoff differences, and past interactions exert no direct influence on current adjustment. As a result, local stability is fully characterized by the spectrum of the Jacobian of the replicator vector field at a stationary state, and deviations decay at exponential rates.

A growing body of experimental and empirical evidence challenges this view. In many environments, adjustment is slow and highly persistent even when incentives are transparent and feedback is immediate. Dominated strategies often survive for extended periods (Erev and Roth, 1998), behavior exhibits strong serial dependence (Cheung and Friedman, 1997; Camerer and Hua Ho, 1999), and populations may remain close to unstable thresholds over experimentally relevant horizons (Young, 1993; Fréchette, 2012). In continuous-time Prisoner's Dilemma experiments, Friedman and Oprea (2012) show that cooperation declines gradually rather than exponentially, despite instantaneous feedback and fixed incentives. These evidences suggest that the way past payoff information is aggregated over time may shape evolutionary dynamics in ways not captured by memoryless adjustment.

This paper shows that time is not neutral in evolutionary dynamics. Long memory can alter local stability and slow adjustment from exponential to polynomial rates, so that equilibria may become stable while equilibrium selection remains incomplete over finite horizons, even in deterministic, payoff-monotone environments.

Our framework. To study how temporal aggregation affects evolutionary adjustment, we introduce a minimal modification of the classical replicator dynamics. We consider a unit-mass population repeatedly matched to play a symmetric normal-form game, where strategy shares evolve in a payoff-monotone manner: strategies earning above-average payoffs expand, while underperforming strategies contract. We preserve this framework unchanged and modify only the temporal operator governing adjustment. Specifically, we replace the ordinary time derivative in the replicator equation with a fractional (Caputo) derivative of order $\alpha \in (0, 1]$.

The resulting *fractional replicator dynamics* (FRD) introduce long memory into the adjustment process: current changes in strategy shares depend on a weighted history of past payoff differences. The parameter α governs the persistence of memory. When $\alpha = 1$, the model reduces to the classical replicator with memoryless, exponential adjustment. When $\alpha < 1$, past experiences continue to influence behavior over long horizons, with effects that decay gradually according to a power law. Economically, lower values of α correspond to slower adjustment of aggregate behavior, which may reflect persistence in beliefs, habits, or institutional practices. Importantly, this inertia does not reflect bounded rationality or non-standard preferences. At every instant, agents remain fully payoff-responsive; only the temporal aggregation of payoff information is altered.

Fractional dynamics provide a parsimonious representation of persistent adjustment. Unlike models with finite delays or higher-order dynamics, fractional derivatives capture memory with infinite temporal support and power-law decay. While standard in the analysis of systems with long memory and anomalous relaxation (Diethelm, 2010), this approach has not previously been integrated into evolutionary game theory in a way that preserves payoff monotonicity and equilibrium structure.

Our results. Our findings show that introducing long memory fundamentally reshapes evolutionary dynamics in ways that go beyond a mere slowdown of adjustment.

Our first result shows that introducing long memory preserves the core structure of evolutionary game dynamics. For all $\alpha \in (0, 1]$, population shares remain in the simplex and the set of stationary states coincides with that of the classical replicator. In particular, Nash equilibria remain rest points for all values of α, and no new steady states are created by memory. Thus, fractional adjustment leaves incentives and feasibility unchanged, allowing us to isolate the role of temporal aggregation. Any change in behavior arises solely from how past payoff information is accumulated over time, not from modifications of the strategic environment. In this sense, the fractional replicator captures a pure timing effect.

Our second result concerns local stability. Under the FRD, stability depends jointly on payoff incentives and on the memory parameter α. While every evolutionarily stable strategy (ESS) remains locally asymptotically stable for all α, equilibria that are unstable under the classical replicator may become locally stable when memory is sufficiently persistent. This gives rise to a *stability-switching phenomenon*: varying α can induce qualitative changes in stability without any change in payoffs or strategic interaction.

The intuition is that long memory modifies how populations respond to local payoff gradients. With memoryless adjustment ($\alpha = 1$), dynamics react only to contemporaneous incentives, so destabilizing local forces—such as rotational or expansive components—act without attenuation. Under fractional adjustment ($0 < \alpha < 1$), current motion reflects an aggregation of past payoff differences, which introduces inertia into the adjustment process, smoothing short-run responses and weakening destabilizing forces. As a result, deviations that would grow or cycle under classical dynamics may be damped when memory is sufficiently persistent.

This mechanism is transparent in cyclic environments such as Rock–Paper–Scissors games. In the exact zero-sum case, the interior Nash equilibrium is neutrally stable under the classical replicator, generating closed orbits. Any amount of fractional memory ($\alpha < 1$) breaks this knife-edge behavior by damping oscillations and inducing convergence. A similar logic applies in near zero-sum games. Small symmetric perturbations make the interior equilibrium unstable under the classical replicator, producing outward spirals. Under fractional adjustment, memory attenuates these destabilizing forces. When α falls below an endogenous threshold, temporal aggregation dominates and the interior equilibrium becomes locally asymptotically stable despite being unstable under memoryless dynamics.

Our third result characterizes adjustment paths and convergence rates under the FRD. While the classical replicator implies exponential convergence toward locally stable equilibria, we show that under FRD convergence follows a Mittag–Leffler law and is generically polynomial. In particular, deviations from a locally stable equilibrium decay at rate $t^{-\alpha}$, so convergence becomes arbitrarily slow as memory becomes more persistent.

These differences are quantitatively meaningful. In dominance games such as the Prisoner's Dilemma, the classical replicator predicts rapid extinction of the dominated strategy: in our simulations, cooperation falls below 5% within roughly five time units when $\alpha = 1$. Under fractional adjustment, convergence slows sharply. For $\alpha = 0.6$, the same threshold is reached only after several tens of periods; for $\alpha = 0.4$, cooperation remains above 10%

beyond $t = 50$; and for $\alpha = 0.2$, it still exceeds 10% at horizons around $t = 100$, despite cooperation being locally unstable.

The consequences are more pronounced in coordination games. Although the set of equilibria and their basins of attraction are unchanged, polynomial convergence implies that selection may fail to be completed within any finite, economically relevant horizon. Large sets of initial conditions remain in the interior of the state space for extended periods, lingering near unstable thresholds rather than converging to a convention. For a fixed horizon T, our simulations of the 2×2 coordination game show that more than 90% of trajectories reach an equilibrium neighborhood by time T when $\alpha \geq 0.8$, whereas only about half do so when $\alpha \approx 0.6$. When $\alpha \leq 0.4$, a large majority—often exceeding 70%—fail to select any equilibrium within the same horizon.

These results show that slow adjustment and persistent non-equilibrium behavior need not reflect noise, bounded rationality, or strategic indeterminacy. They can arise endogenously from the intrinsic dynamics of adaptation when past experience remains salient. Crucially, this does not require modifying payoff incentives or the revision protocol itself, since the replicator structure is preserved, and only the temporal operator governing adjustment is altered. Long-run equilibria therefore remain unchanged, yet equilibrium selection becomes a fundamentally finite-time phenomenon shaped by the temporal structure of adjustment. This perspective bridges the gap between sharp equilibrium predictions and the slow, history-dependent patterns commonly observed in data, and motivates a re-examination of evolutionary selection beyond purely asymptotic analysis.

Related literature. The present paper is related to a few strands of the literature.

First, our analysis relates to the literature on evolutionary game dynamics, which studies how strategic behavior evolves under payoff-driven selection. Within this literature, a variety of deterministic adjustment rules have been proposed, including the replicator dynamics (Taylor and Jonker, 1978), best-response dynamics (Gilboa and Matsui, 1991), Brown–von Neumann–Nash dynamics (Brown and Von Neumann, 1950), and Smith dynamics (Smith, 1984).[1] These models differ in their underlying revision protocols but share a common feature: local stability is determined by payoff incentives through the Jacobian of the underlying vector field at a stationary state. Time affects only the speed of convergence and is therefore neutral with respect to equilibrium selection. Our paper de-

---

[1] Among these models, the replicator has become the canonical benchmark due to its well-known microfoundation and analytical tractability, and its properties are well established (see Cressman and Tao, 2014): dominated strategies vanish, and evolutionarily stable strategies are locally asymptotically stable.

parts from this neutrality. While preserving the replicator's payoff-monotonicity, feasibility, and equilibrium structure, we show that the temporal structure of adjustment itself can affect local stability. Introducing memory through fractional adjustment can stabilize equilibria that are unstable under the classical replicator and render finite-time equilibrium selection non-degenerate, even in deterministic environments.[2]

Second, a large literature departs from the predictions of the replicator dynamics by introducing stochastic perturbations, mutations, or learning frictions. In stochastic evolutionary models, long-run outcomes are selected through rare shocks or mutations (Foster and Young, 1990; Kandori et al., 1993; Cabrales, 2000; Choudhury and Aydinyan, 2023), and persistence away from equilibrium is driven by exogenous noise. A related strand introduces inertia through learning rules. Reinforcement learning and belief-based learning models allow past experience to influence current behavior (Cheung and Friedman, 1997; Erev and Roth, 1998; Camerer and Hua Ho, 1999). While empirically successful, these approaches typically modify the decision rule itself and rely on additional behavioral parameters governing learning speed or forgetting. Finally, models with explicit delays, inertial terms, or higher-order dynamics show that temporal structure can affect stability and generate oscillations or slow convergence (see, e.g., Cressman, 2003; Gopalsamy, 2013; Obando et al., 2016; Mertikopoulos and Sandholm, 2016; Mertikopoulos et al., 2018). These models, however, rely on finite delays or augmented state variables and retain exponential discounting of the past.[3]

Our approach differs from all three strands. We introduce neither noise nor mutations, and we do not modify payoff incentives or choice rules. Adjustment remains fully payoff-monotone at every instant. Persistence arises instead from the way past payoff information is aggregated over time. Fractional dynamics capture long memory with infinite support, generating power-law decay rather than exponential forgetting. This allows slow adjustment and stability switching to emerge endogenously in a fully deterministic setting.

---

[2] This mechanism is complementary to recent work emphasizing the sensitivity of evolutionary predictions to the choice of adjustment dynamics. In particular, Hofbauer and Sandholm (2011) show that many standard results of evolutionary game theory are not robust across reasonable classes of deterministic dynamics. While their analysis focuses on how qualitative outcomes depend on the revision protocol, our paper highlights a distinct channel: holding the payoff-induced vector field fixed, changing the temporal law of adjustment through memory can fundamentally alter local stability.

[3] A related literature studies rational learning in games and establishes conditions under which behavior converges to Nash equilibrium. Bayesian and belief-based learning models show that, under appropriate informational assumptions, rational agents eventually learn to play Nash equilibria (Kalai and Lehrer, 1993; Hart and Mas-Colell, 2000). These results provide strong foundations for equilibrium selection in the long run. However, they are largely silent about the dynamics of adjustment over finite horizons, in particular about convergence speeds and the persistence of out-of-equilibrium behavior.

Finally, fractional differential equations provide a well-established mathematical framework for modeling systems with long memory (Diethelm and Ford, 2002; Deng et al., 2007). They are widely used in physics and engineering to capture persistent dynamics such as anomalous diffusion and viscoelastic behavior (Rivero et al., 2013), but have received little attention in economic dynamics. To the best of our knowledge, this paper provides the first integration of fractional time into continuous-time evolutionary game theory in a way that preserves payoff-based selection and standard equilibrium structure.

Outline. Section 2 presents the standard evolutionary framework and reviews the basic results of the classical replicator dynamics. Section 3 introduces the fractional replicator dynamics (FRD) and establishes its fundamental properties. Section 4 contains the main theoretical findings, organized around three theorems: a local stability criterion (Theorem 1), a memory-induced stability switching result (Theorem 2), and a characterization of Mittag–Leffler (polynomial) convergence (Theorem 3). Sections 5 and 6 provide numerical illustrations, focusing respectively on slow convergence and finite-time equilibrium selection, and on stability switching in 3×3 environments. Section 7 concludes. Appendix A collects the proofs of all propositions and theorems. The Online Appendix develops a microfoundation of the FRD. We show that the fractional replicator arises as the scaling limit of a standard payoff-monotone revision process when agents revise strategies at heterogeneous and asynchronous times with heavy-tailed waiting times. In this interpretation, the fractional order captures heterogeneity in revision frequencies rather than departures from rational or payoff-driven behavior.

2 The Evolutionary Framework

We consider a unit-mass population of agents who are randomly matched in pairs to play a symmetric normal-form game with a finite set of pure strategies $S = \{s_1, \dots, s_n\}$, $n \geq 2$. Time is continuous. Let $p_i(t)$ denote the proportion of the population playing strategy s_i at time t. The population state is represented by the vector $p(t) = (p_1(t), \dots, p_n(t)) \in \Delta^n$, where $\Delta^n := \left\{ p \in \mathbb{R}^n_+ : \sum_{i=1}^n p_i = 1 \right\}$ is the n-dimensional unit simplex.

Interactions are anonymous and symmetric. The game is represented by a real payoff matrix $A = (a_{ij})_{i,j=1}^n \in \mathbb{R}^{n \times n}$, where a_{ij} is the payoff received by a player using strategy s_i against an opponent using s_j. We assume throughout that payoffs are bounded, $\max_{i,j} |a_{ij}| < \infty$. Given a population state $p \in \Delta^n$, the expected payoff to strategy

s_i is $\pi_i(p) = (Ap)_i = \sum_{j=1}^{n} a_{ij} p_j$. For any two mixed strategies $p, q \in \Delta^n$, we write $\pi(p,q) = p^\top Aq$ for the expected payoff to a player using mix p against an opponent playing q. The average payoff in the population is $\bar{\pi}(p) = \pi(p,p) = p^\top Ap = \sum_{i=1}^{n} p_i \, \pi_i(p)$.

Classical Replicator Dynamics The standard evolutionary model assumes that strategies earning above-average payoffs increase their population share over time. This leads to the well-known replicator dynamics:[4]

$$\dot{p}_i(t) = f_i(p(t)) := p_i(t)\big[\pi_i(p(t)) - \bar{\pi}(p(t))\big], \qquad p_i(0) \in [0,1],\ i = 1, \ldots, n. \tag{1}$$

The term $p_i(t)$ reflects the fact that only strategies with positive mass can change. The bracketed term measures the excess payoff of strategy s_i relative to the population average. If strategy s_i yields above-average payoff, its share grows ($\dot{p}_i > 0$); if it performs below average, its share shrinks ($\dot{p}_i < 0$). Only payoff differences matter, and growth is payoff-monotonic.

We collect here a set of standard properties that will be used throughout the paper.

Proposition 1.

(i) If $p(0) \in \Delta^n$, then $p(t) \in \Delta^n$ for all $t \geq 0$.

(ii) Every Nash equilibrium is a fixed point of the replicator dynamics (1).

At a Nash equilibrium, every strategy in use earns exactly the average payoff, so the vector field vanishes. Strategies with zero mass cannot suddenly appear, because the multiplicative factor p_i forces their growth rate to zero even if they would earn higher payoffs. This explains both the invariance of faces of the simplex and why every Nash equilibrium is a rest point of the dynamic. However, a Nash equilibrium need not be dynamically stable. A stronger notion is that of evolutionary stability, which captures robustness to invasion by rare mutants (Smith and Price, 1973).

Definition 1. *A state $p^\star \in \Delta^n$ is an* evolutionarily stable strategy (ESS) *if:*

(i) $p^\star$ is a Nash equilibrium;

(ii) $\forall q \neq p^\star$, $\pi(q, p^\star) = \pi(p^\star, p^\star) \Rightarrow \pi(p^\star, q) > \pi(q, q)$.

---

[4] The replicator dynamic admits several well-established microfoundations, including imitation of more successful strategies, fitness-proportional reproduction, and revision protocols with payoff-monotonic choice (Börgers and Sarin, 1997; Taylor and Jonker, 1978; Hofbauer and Sigmund, 1998; Benaïm and Weibull, 2003; Sandholm, 2010).

Evolutionary stability is thus a refinement of Nash equilibrium based on invasion fitness. Condition (ii) requires that if a mutant population q earns the same payoff against $p^\star$ as residents do against themselves, then residents must earn strictly more against mutants than mutants earn against themselves. Under payoff-monotonic dynamics such as the replicator equation, this notion of stability has a well-known dynamic implication. In particular, every strict ESS is locally asymptotically stable under the classical replicator dynamics (see Weibull, 1997, Proposition. 3.10).

Proposition 2. *Every strict ESS is a locally asymptotically stable fixed point of* (1).

Intuitively, if no nearby mutant type can profitably invade, then any small deviation decreases in frequency and the population trajectory returns to $p^\star$. In symmetric 2×2 games, there is only one invasion direction, so evolutionary stability and local asymptotic stability coincide. In higher dimensions, local stability does not necessarily imply ESS, because invasion may occur along directions that are neutral at first order but profitable at higher order.

Near a stable fixed point, deviations decay exponentially over time:

Proposition 3 (Exponential convergence)**.** *Let $p^\star$ be a locally asymptotically stable fixed point. Assume that $p^\star$ is hyperbolic on the tangent space $T\Delta^n = \{x \in \mathbb{R}^n : \mathbf{1}^\top x = 0\}$. Then there exist constants $C, c > 0$ and a neighborhood $\mathcal{N}(p^\star)$ such that, for all $p_0 \in \mathcal{N}(p^\star)$, $\|p(t) - p^\star\| \leq Ce^{-ct}$, $t \geq 0$. Moreover, a linearization yields*

$$p(t) - p^\star = e^{J_T t}(p_0 - p^\star) + o\left(e^{-ct}\right),$$

where $J_T = Df(p^\star)$ is the Jacobian of the replicator vector field restricted to $T\Delta^n$.

Under the classical replicator dynamics, once a trajectory enters a neighborhood of a locally stable equilibrium, convergence is exponential. This exponential clock implies rapid adjustment and memoryless behavior: at each instant, dynamics respond only to current payoff differences, while the influence of past interactions vanishes at an exponential rate.

From a methodological standpoint, this feature delivers sharp analytical results and clean equilibrium selection but sits uneasily with a broad range of empirical regularities. Experimental and field evidence often document adjustment paths characterized by prolonged plateaus, slow elimination of dominated strategies, strong path dependence, or heavy-tailed transients, as discussed in the Introduction.

The central claim of this paper is that these discrepancies do not originate in the selection mechanism itself. The standard payoff-based term $f_i(p) = p_i(\pi_i(p) - \bar{\pi}(p))$ provides a well-founded description of incentive-driven selection. Rather, the tension with evidence arises from the implicit assumption governing the timing of adjustment. In the classical replicator, this timing is memoryless: past payoff differences are discounted instantaneously, and behavior responds exclusively to contemporaneous incentives.

The next section relaxes this assumption while leaving the selection mechanism unchanged. We modify the adjustment clock by introducing a fractional time operator that allows past payoff differences to influence current dynamics through persistent memory. This generalization preserves the economic logic of payoff-based selection, while generating adjustment paths that are slower and history-dependent.

3 Fractional Replicator Dynamics (FRD)

A natural way to introduce long memory is to allow adjustment to depend on a history buffer. When agents interact, past payoff differences and past strategy revisions continue to affect current behavior. This section formalizes this idea by replacing the time derivative ($\dot{p}_i$) with a fractional (Caputo) operator.

For an absolutely continuous function $x : [0, \infty) \to \mathbb{R}$ and $\alpha \in (0, 1]$, the Caputo derivative is defined by

$$ {}^{C}D_t^{\alpha} x(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t \frac{\dot{x}(\tau)}{(t-\tau)^{\alpha}} \, d\tau, \tag{2} $$

where $\Gamma(\cdot)$ denotes the Euler gamma function.

When $\alpha = 1$, this reduces to the standard derivative, ${}^{C}D_t^{1} x(t) = \dot{x}(t)$. When $\alpha < 1$, past changes $\dot{x}(\tau)$ enter current adjustment with a slowly decaying kernel $(t-\tau)^{-\alpha}$, generating power-law memory. We define the fractional replicator dynamics by replacing $\dot{p}(t)$ in the classical replicator (1) with the Caputo derivative (2).

Definition 2 (Fractional replicator dynamics)**.** *For $\alpha \in (0, 1]$, the fractional replicator dynamics (FRD) are given by*

$$ {}^{C}D_t^{\alpha} p_i(t) = f_i(p(t)) := p_i(t)\big[\pi_i(p(t)) - \bar{\pi}(p(t))\big], \qquad p_i(0) \in [0, 1],\ i = 1, \ldots, n. \tag{3} $$

The fractional order α provides a structural way to model how past experiences affect current behavior. When $\alpha = 1$, (3) coincides with the classical (memoryless) replicator (1). When $\alpha < 1$, the rate of change at time t depends on the entire past history $\{p(\tau) : 0 \leq \tau \leq t\}$, not only on the current payoff difference.

An equivalent and often useful formulation is the Volterra representation;

Proposition 4 (Volterra representation). *Let $0 < \alpha < 1$ and $T > 0$. A trajectory $p(\cdot) \in C([0,T];\Delta^n)$ solves the Caputo system (3) on $[0,T]$ if and only if it solves*

$$p_i(t) = p_i(0) + \frac{1}{\Gamma(\alpha)} \int_0^t (t-\tau)^{\alpha-1}\, p_i(\tau) \big[\pi_i(p(\tau)) - \bar{\pi}(p(\tau))\big]\, d\tau. \tag{4}$$

This representation makes the non-Markovian nature of the FRD explicit: the current state reflects a weighted history of past payoff differences, with a power-law kernel.[5]

Basic properties Despite being non-Markovian, the fractional replicator preserves the core geometric and strategic structure of the classical model.

Proposition 5 (Feasibility). *Let $p(0) \in \Delta^n$ and $0 < \alpha \leq 1$. Then the solution $p(t)$ to the FRD satisfies $p(t) \in \Delta^n$ for all $t \geq 0$.*

The replicator drift is tangent to the simplex, and the Caputo operator does not create or destroy mass. Strategies with zero mass cannot decrease further, and the population shares always sum to one. Memory does not violate the geometry of the state space.

Proposition 6 (Fixed points). *A state $p^\star \in \Delta^n$ is stationary under the FRD if and only if it is stationary under the classical replicator. In particular, every Nash equilibrium is a fixed point for all $\alpha \in (0,1]$.*

Introducing fractional memory does not modify the payoff structure nor the instantaneous vector field $p_i(\pi_i(p) - \bar{\pi}(p))$. What changes is only the temporal operator applied to this vector. Therefore, the fractionary differential equation still vanishes at the same states $p^\star$ as in the standard case. Consequently, as any Nash equilibrium is a fixed point for $\alpha = 1$ (Proposition 1), it follows that it is when $\alpha \in (0,1)$. Even when agents place weight on the entire history of past payoffs (as encoded by the Caputo operator), strategies that were mutually best responses in the classical sense still form rest points of the evolutionary process.

---

[5] In fractional analysis, two standard formulations are commonly used: the Caputo derivative, as adopted in (3), and the Riemann–Liouville derivative. Proposition 4 shows that, for the replicator vector field $f(p)$ and $\alpha \in (0,1]$, these formulations are equivalent in the sense that they have the same Volterra representation.

In our FRD, the fractional order α governs the persistence of memory in the adjustment process. Slow, history-dependent adjustment is a common feature of adaptive systems. In evolutionary biology, delayed or age-structured reproduction generates much slower convergence than memoryless dynamics (Tuljapurkar, 2013; Metz and Diekmann, 2014). In neuroscience, learning and forgetting are often characterized by persistent responses rather than rapid decay (Henry et al., 2010). These findings point to a general link between temporal heterogeneity and slow adjustment.

In economics, α admits a natural interpretation as a measure of belief inertia or cultural persistence. Standard learning models typically assume exponential discounting of past experience, yet experimental evidence documents slower updating and strong serial dependence in behavior (Cheung and Friedman, 1997; Camerer and Hua Ho, 1999; Erev and Roth, 1998). Related forms of persistence arise in models of habit formation, reinforcement learning, and experience-based updating (Fudenberg and Levine, 1998), as well as in the literature on path dependence and institutional persistence (North, 1990; Greif, 1994; Acemoglu and Robinson, 2005).

Within our framework, these mechanisms are captured parsimoniously by a single parameter. Lower values of α correspond to stronger historical dependence and slower adjustment, while $\alpha = 1$ recovers memoryless dynamics. Importantly, varying α leaves equilibrium payoffs and asymptotic basins unchanged and affects only the temporal speed of adjustment. By separating strategic structure from the timing of adaptation, the fractional replicator provides a simple way to reconcile the evolutionary game theory with the empirical observation that social norms, beliefs, and institutions often evolve slowly and remain shaped by history.[6]

4 Main Theorems

This section establishes the core analytical properties of the fractional replicator dynamic (FRD). While the preceding section focused on feasibility and equilibrium invariance, we now examine the two respects in which fractional time fundamentally modifies the behavior of the classical replicator dynamic.

The first concerns *local stability*. In the classical case, local stability of a stationary point is determined by the signs of the real parts of the eigenvalues of the Jacobian matrix

---

[6] The Online Appendix provides a microfoundation by linking α to aggregation over heterogeneous and asynchronous revision times.

of the replicator vector field, evaluated at that point. Under fractional adjustment, this criterion no longer applies. Stability is instead governed by an angular condition on the eigenvalues of the same linearized system, where the admissible region depends explicitly on the memory parameter α. As α varies, eigenvalues may cross the stability boundary even though the payoff structure and the underlying vector field remain unchanged. Fractional time thus induces a bifurcation in stability driven purely by memory.

The second dimension concerns *the speed of convergence*. Even when local stability is preserved, fractional adjustment replaces exponential convergence by a Mittag–Leffler (power-law) decay. Consequently, asymptotic equilibrium selection becomes sharply decoupled from the rate at which behavior adjusts in finite time, with potentially long transients despite local stability.

Notation. Throughout, $f(p) = (f_1(p), \dots, f_n(p))$ denotes the classical replicator vector field and let $J := Df(p^\star)$ be its Jacobian at a stationary state $p^\star \in \Delta^n$. The fractional replicator dynamics (3) is quasi-positive, meaning that $p_i = 0 \implies f_i(p) \geq 0$, for all i, and satisfies the mass conservation property $\sum_{i=1}^n f_i(p) = 0$. Consequently, the simplex Δ^n is forward invariant (Proposition 5). For any stationary state $p^\star \in \Delta^n$, let $S(p^\star) := \{i \in \{1, \dots, n\} : \; p_i^\star > 0\}$ denote its support. We define the smallest face of the simplex containing $p^\star$ by

$$F(p^\star) := \{p \in \Delta^n : \; p_i = 0 \text{ for all } i \notin S(p^\star)\}.$$

By quasi-positivity, this face is positively invariant under the FRD, and $p^\star$ belongs to the relative interior of $F(p^\star)$.[7] Since the simplex is characterized by the affine constraint $\sum_i p_i = 1$, local perturbations of $p^\star$ within the invariant face $F(p^\star)$ evolve on the tangent space

$$T_{p^\star}F(p^\star) = \big\{y \in \mathbb{R}^n : \; \mathbf{1}^\top y = 0, \; y_i = 0 \text{ for all } i \notin S(p^\star)\big\}.$$

We define the tangent Jacobian by $J_T := Df(p^\star)\big|_{T_{p^\star}F(p^\star)}$, and write $\sigma(J_T)$ for the spectrum of J_T. All local stability and convergence results in this section are understood relative to the invariant face $F(p^\star)$. Within this face, the local behavior of the FRD is governed by the fractional linearization ${}^{C}D_t^{\alpha} y(t) = J_T\, y(t)$, which allows us to treat interior and boundary stationary states in a unified framework.

---

[7] We formally prove this claim in Lemma 2 in Appendix. For a face $F \subset \Delta^n$, its relative interior consists of all points in F whose coordinates corresponding to active strategies are strictly positive.

4.1 Local Stability and the Role of Memory

Under the classical replicator dynamic (1), local stability is governed entirely by the payoff structure through the Jacobian of the replicator vector field. A stationary state $p^\star$ is locally asymptotically stable iff all eigenvalues of the Jacobian have strictly negative real parts. In this setup, time plays no independent role as stability is a strategic property. Fractional adjustment breaks this logic. Local stability is no longer determined by the signs of real parts of eigenvalues, but by their complex arguments through an angular criterion. As a result, the same payoff structure can generate qualitatively different local dynamics depending solely on how strongly the past is remembered. The criterion is due to Matignon (1996).

Theorem 1 (Local stability under FRD)**.** *Let $p^\star \in \Delta^n$ be a stationary state of the FRD. Then $p^\star$ is locally asymptotically stable, relative to the invariant face $F(p^\star)$, if and only if every eigenvalue $\lambda \in \sigma(J_T)$ satisfies $|\arg(\lambda)| > \alpha\pi/2$.*

For $\alpha = 1$, the stability region coincides with the left half-plane, and the condition reduces to the classical requirement $\Re(\lambda) < 0$. When $\alpha < 1$, the stability region contracts to the complement of a cone around the positive real axis with opening angle $\alpha\pi$, reflecting the altered temporal aggregation of payoff differences induced by fractional adjustment.

To interpret this condition, consider the linearized dynamics around $p^\star$. Each eigenvalue $\lambda \in \sigma(J_T)$ corresponds to a local adjustment mode. Eigenvalues with small arguments are associated with predominantly non-oscillatory modes, in which deviations from the stationary state evolve mainly through monotonic expansion or contraction. By contrast, eigenvalues with arguments close to $\pi/2$ correspond to oscillatory modes, characterized by persistent local cycling with limited net drift.

Fractional adjustment affects these modes asymmetrically. Because the fractional operator aggregates past payoff differences using a slowly decaying kernel, the effective adjustment rate depends on the temporal structure of deviations. Oscillatory modes, which repeatedly revisit nearby states, are more strongly dampened by memory accumulation than monotonic modes. As a result, eigenvalues with a substantial imaginary component can satisfy the fractional stability condition even when they would generate neutral cycles or weak instability under the classical replicator.

Hence fractional memory does not uniformly slow down adjustment. Rather, it selectively alters the balance between drift and oscillation, transforming certain cyclic forces into dissipative dynamics. This mechanism explains why payoff structures that generate persistent cycling under memoryless adjustment may converge under fractional time.

Because stability now depends explicitly on the temporal law of adjustment, it is useful to index dynamic stability by α.

Definition 3 (α-stability)**.** *A stationary state $p^\star \in \Delta^n$ is said to be α*-stable *if it is locally asymptotically stable under the FRD with Caputo derivative order $\alpha \in (0,1]$, relative to the smallest invariant face $F(p^\star)$.*

By construction, α-stability coincides with standard local asymptotic stability when $\alpha = 1$. For $\alpha < 1$, α-stability refines refines classical stability by making explicit that dynamic viability is a joint property of (i) the payoff structure, which determines the local direction and rotation of motion (encoded in J_T), and (ii) the temporal aggregation of payoff information, which governs how strongly past deviations continue to shape current adjustment (encoded by α).

This definition formalizes the idea that stability is not purely strategic: A stationary state may fail to be reached under memoryless adjustment and yet become locally attracting once history-dependent adjustment is introduced.[8]

Proposition 7. *If $p^\star$ is a strict ESS in the classical sense, then $p^\star$ is α-stable for all $\alpha \in (0,1]$.*

Fractional memory does not destabilize evolutionarily robust equilibria. Fractional time preserves the canonical link between ESS and local attraction. However, the converse does not hold in general. In games with three or more strategies, there exist stationary states that are not ESS and are unstable under the classical replicator ($\alpha = 1$), yet become α-stable once memory is sufficiently strong. This reflects a distinction between evolutionary stability and dynamic stability under memory-dependent adjustment.

Consequently, α acts as a temporal selection parameter. It determines which stationary states are dynamically viable once adjustment frictions and memory are taken into account. ESS remain robust for all α, but additional stationary states may become locally attracting as memory strengthens. Stability is therefore not purely strategic, but jointly determined by incentives and temporal structure.

4.2 Memory-Induced Stability Switching

We now isolate a general mechanism through which fractional memory reverses local stability in multidimensional games.

---

[8] Note that, by definition, α-stability is always understood relative to the smallest invariant face of the simplex containing the stationary state. This convention allows us to treat interior and boundary equilibria in a unified definition.

Theorem 2 (Memory-induced stability switching)**.** *Let* $p^\star \in \Delta^n$ *be a stationary state of the FRD, and let* J_T *denote the tangent Jacobian at* $p^\star$*. Assume that* $\vartheta := \min_{\lambda \in \sigma(J_T)} |\arg(\lambda)| \in (0, \pi/2)$*. Then:*

(i) $p^\star$ *is locally unstable under the classical replicator* $(\alpha = 1)$*;*

(ii) there exists a threshold $\bar{\alpha} := 2\vartheta/\pi \in (0, 1)$ *such that* $p^\star$ *is* α*-stable for all* $\alpha \in (0, \bar{\alpha})$*.*

This result follows directly from the local stability criterion established in Theorem 1. Under the FRD, local stability is characterized by an angular condition on the spectrum of the tangent Jacobian: a stationary state is α-stable if and only if every eigenvalue λ satisfies $|\arg(\lambda)| > \alpha\pi/2$. In contrast with the classical replicator, stability is therefore not determined solely by the sign of the real parts of the eigenvalues, but by their position in the complex plane.

When $\alpha = 1$, as we have seen, FRD coincides with the classical replicator and the angular condition reduces to $|\arg(\lambda)| > \pi/2$, which is equivalent to the standard requirement $\Re(\lambda) < 0$ for all $\lambda \in \sigma(J_T)$. In the present theorem, the spectral assumption $\vartheta := \min_{\lambda \in \sigma(J_T)} |\arg(\lambda)| \in (0, \pi/2)$ implies that this classical condition fails: at least one eigenvalue lies outside the open left half-plane, so the stationary state is locally unstable under the classical replicator dynamics. For $0 < \alpha < 1$, however, the stability condition becomes less restrictive. Since the critical angle $\alpha\pi/2$ decreases with α, eigenvalues that are destabilizing in the classical case may satisfy the fractional stability condition when memory is sufficiently strong. Defining the threshold $\bar{\alpha} = 2\vartheta/\pi$, it follows that for all $\alpha \in (0, \bar{\alpha})$ the angular condition holds for every eigenvalue of J_T, and the stationary state becomes locally asymptotically stable. Hence, memory can stabilize a stationary state that is unstable in the memoryless dynamics. Importantly, this stability reversal occurs without any change in payoffs, best-response behavior, or stationary points: it is driven entirely by memory, which alters the temporal structure of adjustment rather than the strategic environment. The mechanism is therefore a *purely temporal bifurcation*. [9]

Corollary 1 (Monotonicity in α)**.** *Fix a stationary state* $p^\star$ *and its spectrum* $\sigma(J_T)$*. If* $p^\star$ *is* α_0*-stable for some* $\alpha_0 \in (0, 1]$*, then it is* α*-stable for every* $\alpha \in (0, \alpha_0]$*.*

---

[9] Note that this result highlights a key dimensionality requirement. When $n = 2$, the Jacobian admits only real eigenvalues, so the angular stability criterion collapses to the classical sign test and stability cannot depend on α. By contrast, for $n \geq 3$, complex eigenvalues arise generically in economically relevant environments with cyclic or non-transitive interactions. In such settings, memory becomes a genuine stability parameter: varying α alone can turn an unstable stationary state into a stable one.

Corollary 1 characterizes how local stability varies with the degree of memory. By Theorem 1, local stability requires that all eigenvalues of J_T lie outside a sector of angle $\alpha\pi$ centered on the positive real axis. As α decreases, this sector contracts monotonically, making the spectral condition for stability weakly easier to satisfy. Consequently, the set of stationary states that are locally asymptotically stable under FRD expands as memory persistence increases.

A direct implication is that memory cannot eliminate local stability. If a stationary state is (locally) stable at some α_0, it remains stable for all smaller α. Conversely, a stationary state that is unstable under the classical replicator dynamics ($\alpha = 1$) may become locally stable for sufficiently small α. Theorem 2 provides a sharp characterization of this stability reversal when instability is driven by complex eigenvalues.

Economically, this monotonicity implies that memory affects adjustment dynamics in a one-sided and purely stabilizing manner. Varying α modifies the temporal aggregation of past payoffs without altering incentives, payoffs, or best responses. Memory therefore attenuates payoff-induced rotational or locally expansive forces in the dynamics but cannot generate new sources of local instability.

4.3 Speed of Convergence: Mittag–Leffler Relaxation

We now turn to our second main result, which concerns the speed at which stable stationary states are approached. While the previous section showed that memory can alter local stability properties, fractional adjustment also fundamentally changes the dynamics of convergence.

Theorem 3 (Mittag–Leffler convergence)**.** *Let $\alpha \in (0,1)$ and let $p^\star$ be an α-stable stationary state of FRD. Then there exists a neighborhood $\mathcal{U}$ of $p^\star$ such that, for every $p(0) \in \mathcal{U}$,*

$$p(t) - p^\star \;=\; E_\alpha\big(J_T t^\alpha\big)\,(p(0) - p^\star) \;+\; o(\|p(0) - p^\star\|) \quad \text{as } \|p(0) - p^\star\| \to 0,$$

where $E_\alpha(z) = \sum_{k=0}^{\infty} z^k/\Gamma(\alpha k + 1)$ denotes the Mittag–Leffler function.

Moreover, the convergence toward $p^\star$ is polynomial: there exists $K > 0$ such that, for all sufficiently large t, we have $\|p(t) - p^\star\| \le K\, t^{-\alpha}$.

When $\alpha = 1$, $E_1(J_T t) = e^{J_T t}$ and convergence is exponential. When $\alpha < 1$, the “exponential clock” is replaced by a heavy-tailed relaxation law. This difference is not a mere rescaling of time: polynomial decay implies that transients are long-lived and that early deviations retain influence over economically relevant horizons. In particular, the half-life of a deviation is not proportional to a constant rate c^{-1} (as under exponential decay) but depends on the level of the deviation itself, a hallmark of power-law relaxation.

Economically, Mittag–Leffler convergence provides a deterministic mechanism for persistent off-equilibrium behavior. Even when selection is asymptotically unambiguous, dominated strategies and unstable thresholds may remain observable for long periods because the approach to the attractor is slow in a precise, power-law sense. This is the basis for our finite-time indeterminacy results.

This section have analytically demonstrated two distinct channels through which long memory affects evolutionary dynamics. The first channel is *quantitative*. Fractional adjustment preserves the equilibrium set and the local direction of selection whenever stability holds, but replaces exponential convergence with Mittag–Leffler (polynomial) relaxation (Theorem 3). The second channel is *qualitative* and inherently multidimensional. In games with $n \geq 3$, fractional memory can alter local stability when the payoff-induced Jacobian admits complex eigenvalues (Theorem 2).

The remainder of the paper illustrates these two mechanisms. Section 5 studies slow convergence and finite-time equilibrium selection in Prisoner's Dilemma and coordination games. Section 6 illustrates memory-induced stability switching in three-strategy games with cyclic interactions.

5 Finite-Time Indeterminacy and Equilibrium Selection

This section illustrates the convergence result of Theorem 3 and shows that the change in the clock has first-order implications for what an economist can observe in finite time. Even when the equilibrium set and asymptotic basins of attraction are unchanged, the polynomial rate of convergence induced by long memory generates large regions of the state space that remain away from any equilibrium neighborhood over empirically relevant horizons. As a result, equilibrium selection becomes history-dependent.

We illustrate these implications in two classes of environments. We first study 2×2 games, including a strict dominance environment (Prisoner's Dilemma) and a coordination game, to show how fractional memory slows convergence without altering equilibrium selection asymptotically. We then turn to a symmetric 3×3 coordination game, where finite-time indeterminacy becomes particularly transparent on the simplex.

5.1 2×2 games: dominance and coordination under fractional time

We focus on the two-strategy case ($n = 2$) with the payoff matrix

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}.$$

where $p(t) \in [0, 1]$ is the population share playing s_1 (so that $1 - p(t)$ plays s_2). The payoff difference between s_1 and s_2 is then an affine function and can be written as $\Delta(p) := \pi_1(p) - \pi_2(p) = Ap + B$, where $A := (a_{11} - a_{21}) + (a_{22} - a_{12}) \neq 0$ and $B := a_{12} - a_{22}$. The induced replicator drift therefore takes the form

$$f(p) = p(1-p)\,\Delta(p) = p(1-p)(Ap + B).$$

Prisoner's Dilemma. The Prisoner's Dilemma provides the simplest case of strict dominance. Strategy s_2 (Defect) strictly dominates s_1 (Cooperate), so the payoff difference $\Delta(p) = Ap + B$ is negative for all $p \in (0, 1)$ (e.g., when $A < 0$ and $B < 0$).

In the two-strategy case ($n = 2$), Nash equilibria and ESS coincide. It follows from Corollary 6 that $p^\star = 0$ is the unique Nash equilibrium and the unique ESS. Moreover, $p^\star = 0$ is asymptotically selected under both the classical and fractional replicator. Applying Theorem 3, the fraction of cooperation solves locally $p(t) = E_\alpha(\lambda t^\alpha)\, p_0$, $\lambda = f'(0) < 0$, so that, for large t,

$$p(t) \sim \frac{p_0}{|\lambda|\,\Gamma(1-\alpha)}\, t^{-\alpha}.$$

Hence cooperation still vanishes in the long run, but the speed of extinction depends critically on α. When $\alpha = 1$, convergence is exponential; but when $\alpha < 1$, extinction follows a heavy-tailed power law, which implies that substantial levels of cooperation may persist for a long time. We illustrate this feature starting from $p_0 = 0.5$, with $A = B = -1$.[10] Figure 1 shows $p(t)$ for $\alpha \in \{1, 0.8, 0.6, 0.4, 0.2\}$. With $\alpha = 1$, cooperation collapses almost instantaneously (after roughly $t = 5$). In contrast, as long as $\alpha < 1$, the trajectory decays much more slowly. For $\alpha = 0.2$, the share of cooperation is still greater than 10% after $t \approx 100$. Although cooperation remains doomed in the very long run, it may survive

---

[10] All simulations of the FRD are performed in MATLAB using the `fde12` routine of Diethelm et al. (2004), which implements a predictor-corrector scheme for Caputo fractional differential equations. For $\alpha = 1$, the results are cross-checked using a standard ODE solver (`ode45`), which confirms that the fractional code accurately reproduces the classical replicator dynamic. Results are robust to smaller time steps and alternative discretizations. The codes are available upon request.

in short(medium)-run for an order of magnitude longer when memory is long-lived.

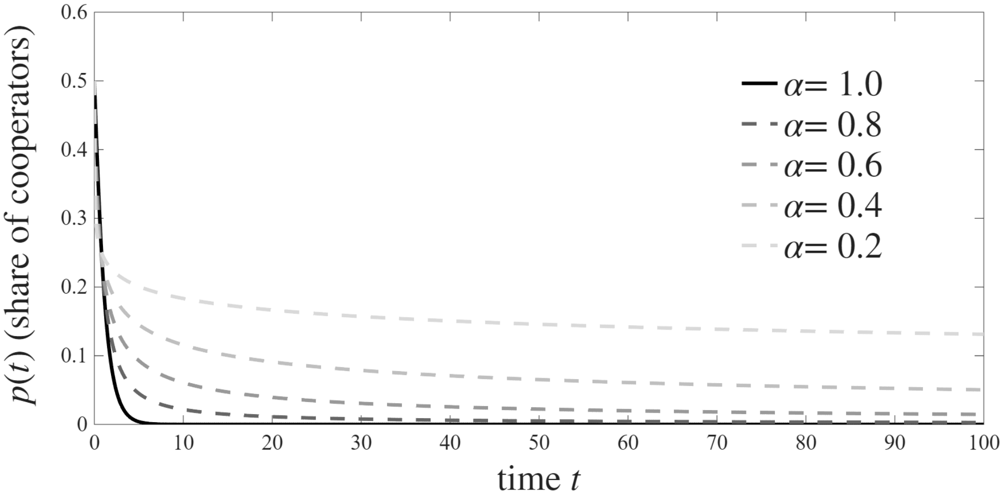

Figure 1: Trajectories $p(t)$ under the fractional replicator for several values of α.

Notes: Prisoner's Dilemma with parameters $A = B = -1$ and $p(0) = 0.5$.

To quantify the slowdown induced by memory, we define the time-to-threshold $T_\varepsilon(\alpha) := \inf\{t \geq 0 : \; p(t) \leq \varepsilon\}$, the first time at which the cooperative share falls below a tolerance level ε. Figure 2 plots $T_\varepsilon(\alpha)$ for $\varepsilon \in \{10\%, 5\%, 1\%\}$ as a function of the memory parameter α. Two patterns emerge. First, $T_\varepsilon(\alpha)$ is strictly decreasing in α and grows rapidly as α decreases, consistent with the Mittag–Leffler scaling $T_\varepsilon(\alpha) \approx \big(p_0/(\varepsilon|\lambda|\Gamma(1-\alpha))\big)^{1/\alpha}$. Second, tighter thresholds require substantially more time for any given α, yielding the ordering $T_{10\%}(\alpha) < T_{5\%}(\alpha) < T_{1\%}(\alpha)$ for all α. Quantitatively, reaching 10% cooperation takes about 70 periods at $\alpha = 0.3$, compared with roughly 9 periods at $\alpha = 0.5$ and only a few periods when α is close to 1. The slowdown is more pronounced for stricter thresholds: for $\varepsilon = 5\%$, we find $T_{5\%}(0.45) \approx 75$ and $T_{5\%}(0.6) \approx 10$, while for $\varepsilon = 1\%$ even moderate memory generates long persistence ($T_{1\%}(0.65) \approx 90$ and $T_{1\%}(0.9) \approx 10$).

Finally, Figure 3 examines the tail behaviour of the trajectories by plotting $\log p(t)$ against $\log t$. This transformation allows us to clearly distinguish exponential decay from power-law decay. The black curve ($\alpha = 1$) bends downward, as expected from an exponential law. By contrast, for all $\alpha < 1$, the trajectories align with almost straight lines in log-log scale. Each of these lines has an asymptotic slope approximately equal to $-\alpha$, as predicted by the Mittag-Leffler asymptotics $p(t) \sim t^{-\alpha}$. The dotted lines plotted in the background confirm this behaviour: they represent the theoretical reference $t^{-\alpha}$ and closely overlap with the numerical trajectories for sufficiently large t.

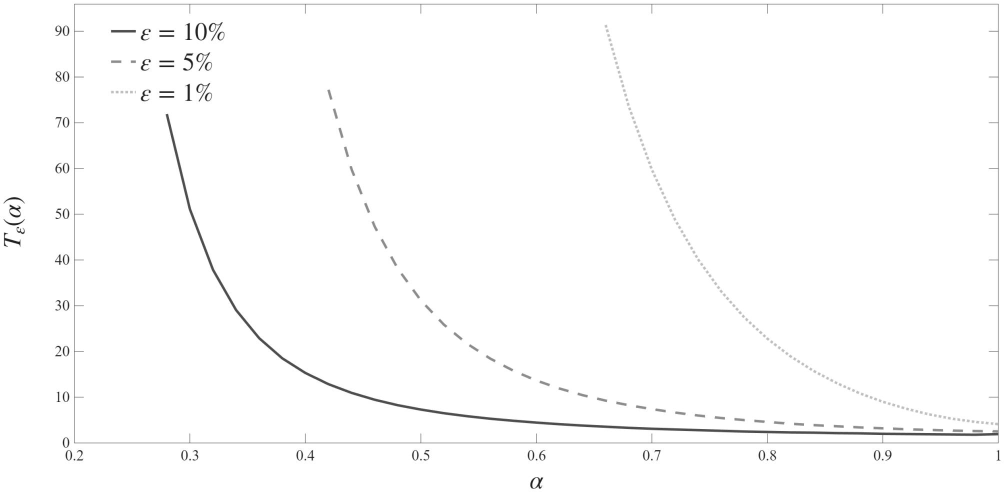

Figure 2: Time-to-threshold $T_\varepsilon(\alpha)$ under the fractional replicator.

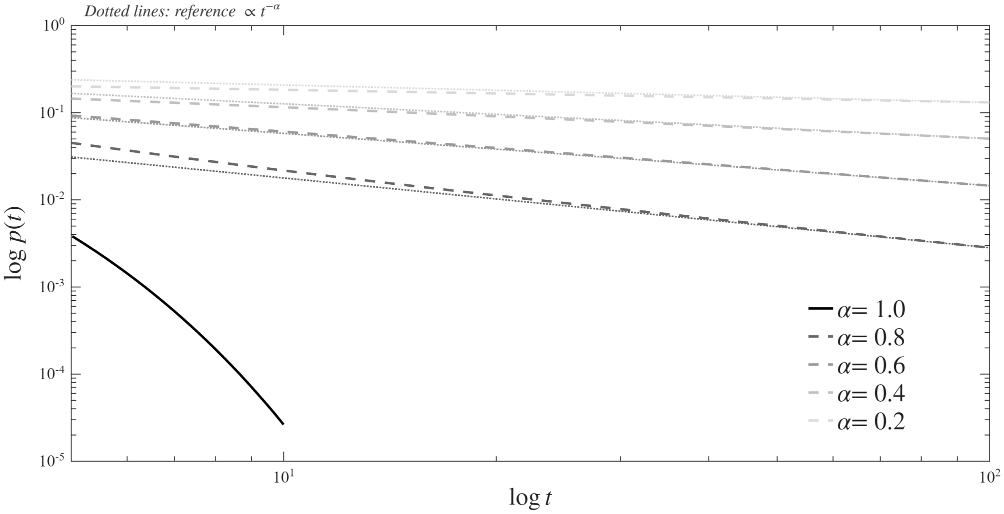

Figure 3: Log–log plot of $\log p(t)$ vs. $\log t$ under the fractional replicator.

Notes: Log–log plot illustrating power-law decay when $\alpha < 1$. The slope is approximately $-\alpha$, consistent with Mittag–Leffler asymptotics.

Coordination and finite-time non-degenerate selection. In a coordination game, the payoff difference $\Delta(p) = Ap + B$ crosses zero once in the interior of the unit interval. Assuming $A > 0$ and $B < 0$, this unique crossing occurs at $p^\dagger := -B/A \in (0,1)$. The induced replicator dynamics therefore admit three stationary states: the boundary points $p = 0$ and $p = 1$, which are locally asymptotically stable, and the interior point $p^\dagger$, which is locally unstable.

Under the classical replicator, the threshold partitions the state space sharply: initial conditions with $p(0) > p^\dagger$ converge exponentially to 1, and those with $p(0) < p^\dagger$ converge exponentially to 0. Hence, in any realistic time frame, nearly every population reaches one of the two conventions. Under the fractional dynamic, however, the set of equilibria and their basins remain the same, but convergence slows dramatically. As a consequence, finite-time equilibrium selection is no longer unique—both conventions may be observed with strictly positive frequency within relevant horizons.

To formalize this, define for $k \in \{0,1\}$ the hitting time $T_\varepsilon^{(k)}(\alpha) = \inf\{t \geq 0 : \ |p(t) - k| \leq \varepsilon\}$. Although long-run selection is unaffected by α, Theorem 3 implies that near a stable corner with slope $\lambda_k = f'(k) < 0$,

$$T_\varepsilon^{(k)}(\alpha) \ \approx \ \left(\frac{|p_0 - k|}{\varepsilon \, |\lambda_k| \, \Gamma(1-\alpha)} \right)^{1/\alpha},$$

so $T_\varepsilon^{(k)}(\alpha)$ explodes as α declines.

We perform numerical simulations with $A = 1$ and $B = -0.6$ (so $p^\dagger = 0.6$). Figure 4 reports the hitting times $T_\varepsilon^{(0)}(\alpha)$ and $T_\varepsilon^{(1)}(\alpha)$. When $\alpha \simeq 1$, both equilibria are reached within 10 periods; at $\alpha = 0.6$, the same paths require 40–50 periods; at $\alpha = 0.5$, more than 70; and once $\alpha < 0.4$, many trajectories do not reach any ε-neighborhood by $T = 100$.

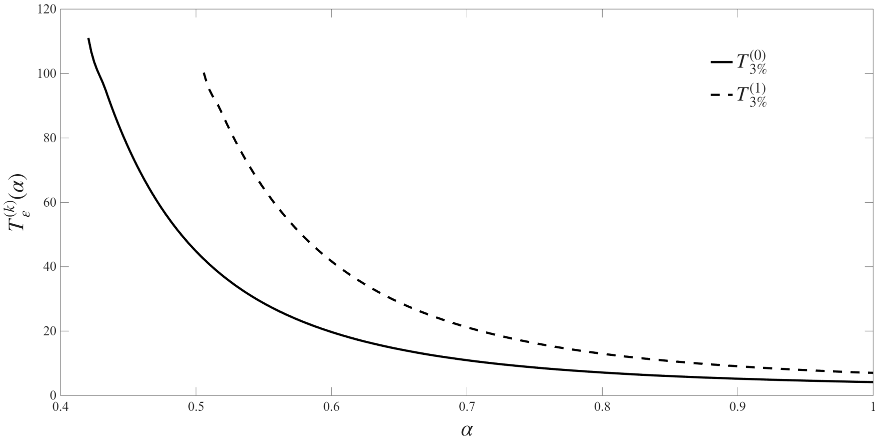

Figure 4: Hitting times $T_\varepsilon^{(0)}(\alpha)$ and $T_\varepsilon^{(1)}(\alpha)$ in a 2×2 coordination game.

Notes: Coordination game with $A = 1$ and $B = -0.6$ (so $p^\dagger = -B/A = 0.6$). The tolerance is $\varepsilon = 0.05$. $T_\varepsilon^{(k)}(\alpha)$ denotes the first time at which the trajectory enters an ε-neighborhood of equilibrium $k \in \{0,1\}$.

To assess which convention is effectively selected over finite horizons, we conduct Monte Carlo simulations with initial conditions drawn uniformly from $[0,1]$, that is, $p_0 \sim$

Unif$[0,1]$. We fix the time horizon at $T = 60$. This horizon is long enough for coordination to occur under the classical replicator dynamics, yet short enough to remain economically and behaviorally relevant—for instance, the duration of an experimental session or a small number of generations in a cultural evolution setting.

For each simulated trajectory, we record whether the state reaches an ε-neighborhood of either boundary equilibrium $s_0 = 0$ or $s_1 = 1$ first, and whether this occurs before the horizon T. Let $T_\varepsilon^{(i)}$ denote the first hitting time of the ε-neighborhood of s_i. This procedure induces the following finite-time selection probabilities:

$$\mathbb{P}_T(s_1 \mid \alpha) = \Pr\big(T_\varepsilon^{(1)} < T_\varepsilon^{(0)} \leq T\big), \qquad \mathbb{P}_T(s_0 \mid \alpha) = \Pr\big(T_\varepsilon^{(0)} < T_\varepsilon^{(1)} \leq T\big).$$

Equivalently, these probabilities correspond to the Lebesgue measures of the sets of initial conditions whose trajectories reach s_1 or s_0 first within the finite horizon.

Importantly, these probabilities need not sum to one. For sufficiently small α, a non-negligible set of trajectories converges too slowly to reach either equilibrium before time T and remains in the interior of the state space. We denote the measure of this set by $|\mathcal{I}_T(\alpha,\varepsilon)| = \Pr\big(\min\{T_\varepsilon^{(0)}, T_\varepsilon^{(1)}\} > T\big)$, and refer to $\mathcal{I}_T(\alpha,\varepsilon)$ as the *indecision region*.

These three outcomes—convergence to s_0, convergence to s_1, or no convergence—partition the unit interval of initial conditions. Figure 5 visualizes this partition in the (α, p_0) plane. Each point is colored according to the outcome of the trajectory: light for s_1, dark for s_0, and gray for trajectories that have not reached either equilibrium by $T = 60$ (here $\varepsilon = 0.02$). When memory is short ($\alpha \geq 0.7$), the picture resembles the classical replicator: the unstable threshold $p^\dagger = 0.6$ sharply separates the two basins, and virtually every initial condition converges to a corner. As α falls below 0.7, a gray band emerges around $p^\dagger$ where convergence does not occur within the horizon. This indecision region expands quickly: it represents 40–60% of initial conditions at $\alpha \approx 0.5$, and over 80% for $\alpha \approx 0.3$. For α between 0.1 and 0.2, more than 90% of initial conditions remain outside any equilibrium neighbourhood at $T = 60$.

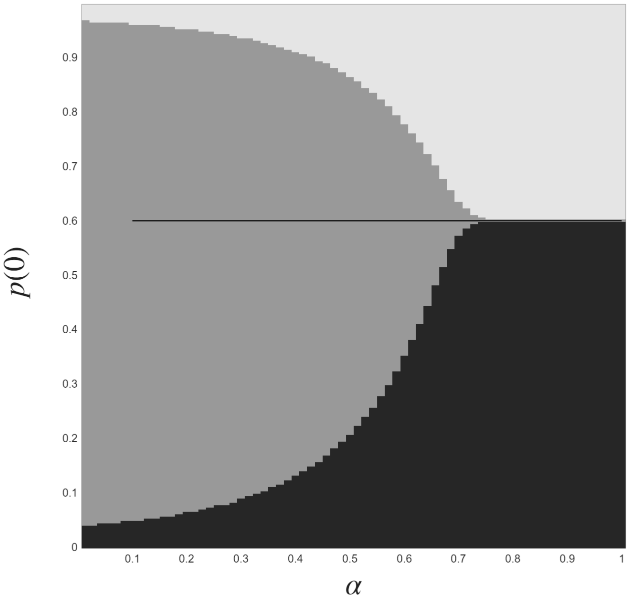

Figure 5: Finite-time selection map in the (α, p_0) plane.

Notes: Selection map computed over a grid of initial conditions $p_0 \in [0, 1]$ for horizon $T = 60$ and tolerance $\varepsilon = 0.02$. Light areas correspond to finite-time selection of s_1, dark areas to selection of s_0, and gray areas to unresolved trajectories that do not reach any ε-neighborhood by time T.

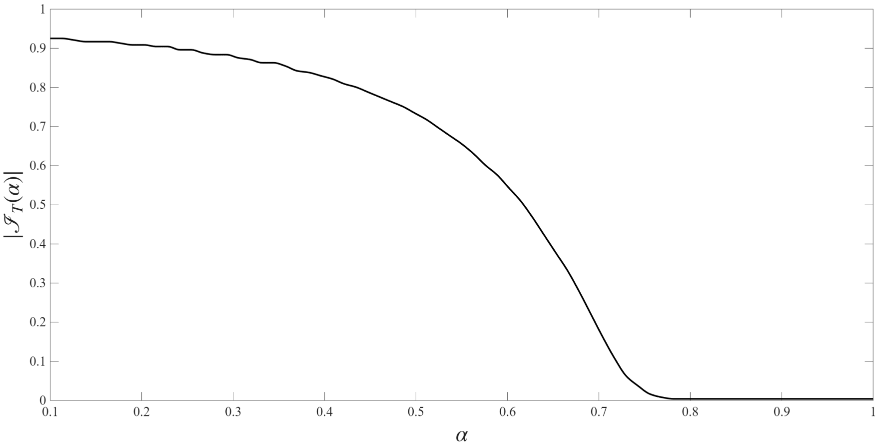

Figure 6: Width of the indecision region $|\mathcal{I}_T(\alpha)|$.

Notes: The indecision region $\mathcal{I}_T(\alpha)$ consists of initial conditions p_0 that do not reach any ε-neighborhood of a pure equilibrium by time T. Parameters: $A = 1$, $B = -0.6$, $\varepsilon = 0.05$, $T = 60$.

Figure 6 plots the associated measure $|\mathcal{I}_T(\alpha)|$ as a function of α. At the same time, conditional on reaching an equilibrium before T, the probability of selecting s_1 versus s_0 stabilizes around a 60/40 split for $\alpha \in [0.5, 0.7]$ (Figure 5). These ratios should not be interpreted as steady-state frequencies: they arise because the drift toward s_1 is slightly stronger near that corner ($|\lambda_0| = 0.6$ versus $|\lambda_1| = 0.4$), and because finite-time selection does not coincide with long-run basins when convergence is slow.

These results point to a general property of fractional dynamics: finite-time equilibrium selection is *non-degenerate*. Even in a simple 2×2 coordination game, both conventions may appear with positive probability and a large share of initial conditions may remain undecided. The following proposition establishes this analytically.

Proposition 8. *Consider a* 2×2 *coordination game under the FRD with* $\alpha \in (0, 1)$. *Fix* $\varepsilon \in (0, 1/4)$ *and* $T > 0$. *Suppose* p_0 *has density* g *on* $[0, 1]$ *satisfying* $g(p) \geq m > 0$. *Let* $\lambda_k := f'(k)$ *denote the local slope at the boundary equilibrium* $k \in \{0, 1\}$ *and set* $\Lambda := \min\{|\lambda_0|, |\lambda_1|\}$. *Then there exists* $\eta > 0$ *such that, whenever*

$$\left(\frac{\eta}{\varepsilon \Lambda \Gamma(1-\alpha)} \right)^{1/\alpha} \leq T \leq \left(\frac{2\eta}{\varepsilon \Lambda \Gamma(1-\alpha)} \right)^{1/\alpha},$$

both selection probabilities are strictly positive: $\mathbb{P}_T(s_0) \geq m\,\eta$, *and* $\mathbb{P}_T(s_1) \geq m\,\eta$. *Hence, finite-time equilibrium selection is non-degenerate.*

The proposition formalizes the idea that, under fractional dynamics, equilibrium selection is generically non-degenerate over economically relevant time horizons. There exists an intermediate range of horizons T—neither too short nor asymptotically large—over which both conventions are selected with strictly positive probability. Trajectories starting sufficiently close to the unstable threshold $p^\dagger$ drift only slowly toward either corner. Because convergence is polynomial, these trajectories do not reach an equilibrium neighborhood within time T, while trajectories slightly farther away already do. As a result, a non-negligible mass of initial conditions selects s_0, another selects s_1, and a third remains transiently uncommitted. This phenomenon is not specific to one-dimensional state spaces. Moving to a multidimensional environment sharpens both its geometric interpretation and its empirical relevance.

5.2 Three-Strategy Coordination and Basin Observability

We consider the minimal multidimensional extension of the coordination environment, namely a symmetric 3×3 game with three competing conventions. Let $\{s_1, s_2, s_3\}$ denote the strategy set and consider the payoff matrix

$$A(\delta) = \begin{pmatrix} 1 & -\delta & -\delta \\ -\delta & 1 & -\delta \\ -\delta & -\delta & 1 \end{pmatrix}, \tag{5}$$

where $\delta > 0$ measures the coordination penalty. Matching on the same convention yields payoff 1, whereas mismatching yields $-\delta$. Each vertex of the simplex is a strict Nash equilibrium and an ESS. By contrast, the symmetric interior point $p^\star = (1/3, 1/3, 1/3)$ is not a stationary state of the classical replicator, since a decrease in the share of any strategy raises its relative payoff and induces drift toward the boundary.

Geometrically, this setting generalizes our 2×2 coordination game. While the 2×2 case features two stable equilibria separated by a single unstable threshold, the 3×3 case unfolds on the two-dimensional simplex, which is partitioned into three basins of attraction separated by separatrices. This geometry allows a direct visualization of equilibrium selection.

Under the classical replicator ($\alpha = 1$), the simplex is cleanly partitioned into three basins, and trajectories converge exponentially fast to one of the corners. Equilibrium selection is therefore completed within moderate time horizons. Under the FDR ($0 < \alpha < 1$), the payoff-induced vector field and the asymptotic basins are unchanged. Each corner remains locally stable and almost every initial condition converges to one of them in the long run. What changes is the temporal law of adjustment: the time required to enter a neighborhood of any equilibrium can be arbitrarily large, creating a wedge between asymptotic basin structure and finite-time observability. Under fractional adjustment, basins persist, but equilibrium selection need not be visible over economically relevant horizons.

Figure 7 illustrates this wedge between asymptotic basin structure and finite-time observability for two values of the coordination penalty, $\delta = 0.3$ (top panels) and $\delta = 0.6$ (bottom panels), under the classical replicator ($\alpha = 1$, left) and the fractional replicator ($\alpha = 0.5$, right).

When miscoordination is mild ($\delta = 0.3$), equilibrium selection is complete within horizon T under the classical replicator. Under fractional adjustment, by contrast, a non-negligible region around the center of the simplex remains outside any ε-neighborhood of the vertices over the same horizon, despite asymptotic convergence. When miscoordination is more severe ($\delta = 0.6$), payoff-induced drift away from the separatrices strengthens and finite-time selection accelerates, yet the delaying effect of memory remains visible.

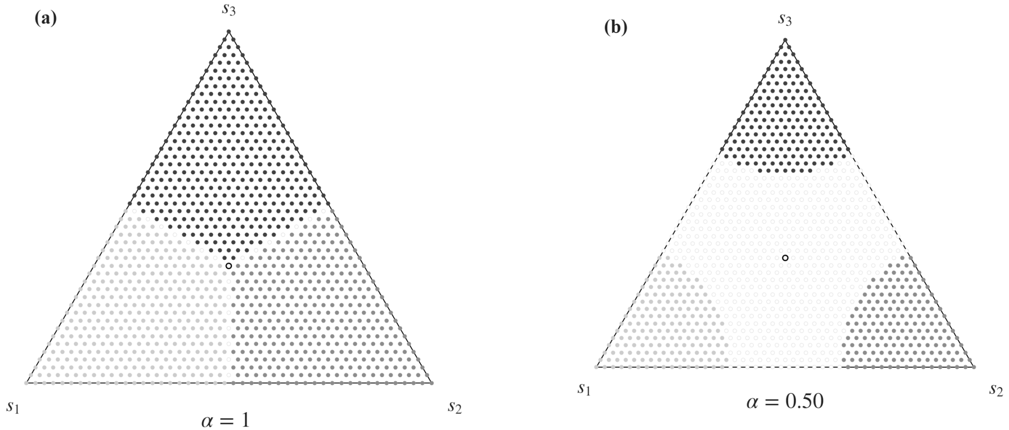

(a) $\delta = 0.3$

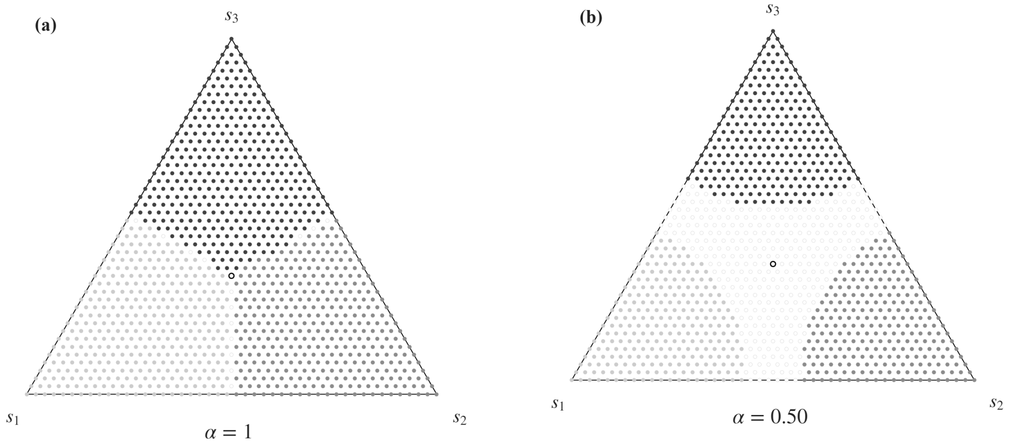

(b) $\delta = 0.6$

Figure 7: Finite-horizon basin classification in the three-strategy coordination game.

Notes: For each δ, the left panel reports the classical replicator ($\alpha = 1$) and the right panel the fractional replicator ($\alpha = 0.5$). Gray dots indicate initial conditions that have not reached any ε-neighborhood of a corner by time T. We consider $\varepsilon = 0.03$ and $T = 60$.

Our simulations confirm that fractional memory leaves the asymptotic basin geometry unchanged but substantially reduces its finite-time visibility. Over economically relevant horizons, the slowing induced by memory can dominate payoff effects, so sharp asymptotic predictions coexist with persistent finite-time indeterminacy.

Methodologically, fractional adjustment then provides a deterministic mechanism through which non-equilibrium or dominated configurations may persist over long but finite horizons, without the need of stochastic perturbations, mutations, or random switching. Apparent coordination failures or prolonged intermediate outcomes can therefore arise from the temporal structure of adjustment rather than from weak incentives or payoff indifference.

6 Stability Switching in a Three-Strategy Game

This section illustrates the stability-switching result established in Theorem 2 within a minimal three-strategy environment. The objective is show how fractional memory induces a genuine temporal bifurcation in a canonical evolutionary game.

The mechanism underlying the bifurcation is the interaction between rotational forces in the local payoff-induced dynamics and history-dependent adjustment. When the Jacobian of the replicator vector field admits complex eigenvalues, local stability is no longer determined solely by payoff primitives. Instead, it also depends on the memory parameter α. As a consequence, the same stationary Nash equilibrium may be neutrally stable, unstable, or locally attracting depending only on the temporal law of adjustment.

We proceed in three steps. We first consider the exact zero-sum case, in which classical replicator dynamics generate neutral cycles and the stability threshold is degenerate: any amount of memory collapses these cycles into a single attractor. We then introduce a small symmetric perturbation, which yields an interior stability threshold, say $\bar{\alpha}(\varepsilon) \in (0, 1)$. Finally, we interpret this threshold as a purely temporal bifurcation, in line with Theorem 2.

Exact zero-sum games (degenerate threshold). We first consider the limiting case of exact zero-sum interactions, namely the payoff matrix A is such that $A + A^\top = 0$. Suppose that such a game admits an interior Nash equilibrium $p^\star \in \mathrm{int}(\Delta^n)$. In this case, the Jacobian of the replicator vector field restricted to the tangent space J_T is antisymmetric. As a consequence, all its nonzero eigenvalues are purely imaginary, $\lambda = \pm i\omega$, implying $|\arg(\lambda)| = \pi/2$. Theorem 2 then yields the following corollary.

Corollary 2 (Exact zero-sum case)**.** *In symmetric zero-sum games with an interior Nash equilibrium $p^\star$, the stability threshold satisfies $\bar{\alpha} = 1$. The equilibrium is neutrally stable under the classical replicator dynamics ($\alpha = 1$) but α-stable for every $\alpha < 1$.*

As long as $\alpha < 1$, any amount of long memory collapses the continuum of neutral cycles generated by the classical replicator into a single attracting trajectory. In exact zero-sum games, the interior Nash equilibrium is surrounded by closed orbits under standard dynamics because local motion is purely rotational. Introducing fractional adjustment alters only the temporal law of motion, but this is sufficient to dampen the rotational component and generate convergence. As a result, a mixed Nash equilibrium that fails to be evolutionarily stable in the sense of Maynard Smith becomes locally asymptotically stable under fractional replicator dynamics.

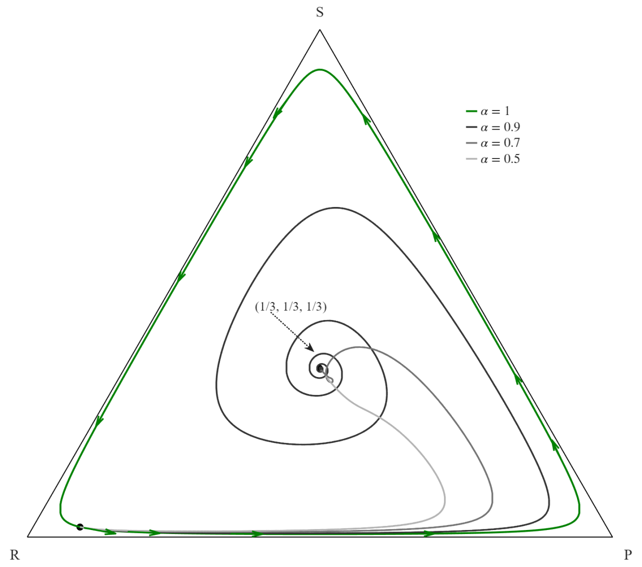

Figure 8: Replicator and fractional replicator dynamics in Rock–Paper–Scissors (RPS)

Notes: Exact zero-sum RPS. All trajectories start from $p(0) = (0.90, 0.08, 0.02)$ over $T = 500$.

Figure 8 illustrates this mechanism in the exact zero-sum Rock–Paper–Scissors game, namely

$$A = A_0 = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix},$$

With $\alpha = 1$, trajectories remain on closed orbits around the interior Nash equilibrium $p^\star = (1/3, 1/3, 1/3)$, reflecting the purely rotational local dynamics. By contrast, under fractional adjustment ($\alpha < 1$), trajectories exhibit damped oscillations and converge to $p^\star$. The figure reports several such trajectories for $\alpha \in \{0.9, 0.7, 0.5\}$, illustrating how even an arbitrarily small amount of memory suffices to eliminate cycling and induce local stability.

Near zero-sum games (interior thresholds). We now relax the exact zero-sum assumption and consider small symmetric perturbations of the Rock–Paper–Scissors game. This allows us to move from the degenerate case in which any amount of memory stabilizes the interior equilibrium to a setting in which stability depends on an interior threshold $\bar{\alpha} \in (0, 1)$.

Specifically, consider the family of payoff matrices $A_\eta = A_0 - \eta B$, where $\eta > 0$, A_0 is the antisymmetric Rock–Paper–Scissors matrix and

$$B = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

For η sufficiently small, the interior state $p^\star = (1/3, 1/3, 1/3)$ remains a Nash equilibrium. However, the introduction of a symmetric component generically shifts the eigenvalues of the tangent Jacobian away from the imaginary axis.

Under the classical replicator dynamics ($\alpha = 1$), this perturbation can turn the interior equilibrium into a repelling focus, so that trajectories spiral outward toward the boundary even though $p^\star$ remains Nash. Under fractional adjustment, by contrast, stability is governed by the spectral angle of the eigenvalues. Since $|\arg(\lambda_\eta)| < \pi/2$ for $\eta > 0$, the associated stability threshold

$$\bar{\alpha}(\eta) \;=\; \frac{2|\arg(\lambda_\eta)|}{\pi}$$

lies strictly inside $(0, 1)$ whenever the classical equilibrium is unstable.

Corollary 3 (Near zero-sum stabilization)**.** *Consider the near zero-sum family* $A_\eta = A_0 - \eta B$*, and suppose that for some* $\eta > 0$ *the interior Nash equilibrium* $p^\star = (1/3, 1/3, 1/3)$ *is locally unstable under the classical replicator dynamics (*$\alpha = 1$*). Then there exists a threshold* $\bar{\alpha}(\eta) \in (0, 1)$ *such that* $p^\star$ *is* α*-stable for all* $\alpha < \bar{\alpha}(\eta)$*.*

This result identifies a genuine stability-switching region in the memory parameter. For large α, the interior equilibrium is dynamically repelling, whereas for sufficiently small α, long memory stabilizes the same equilibrium, holding payoffs fixed.

Figure 9 illustrates this stability-switching mechanism for $\eta = 0.05$. Panel (a) compares the classical replicator dynamics ($\alpha = 1$) with fractional adjustment for $\alpha > \bar{\alpha}(0.05)$: in both cases trajectories display persistent cycling or outward drift. Panel (b) holds the payoff matrix and initial condition fixed but sets $\alpha < \bar{\alpha}(0.05)$, in which case trajectories exhibit damped oscillations converging to $p^\star$.

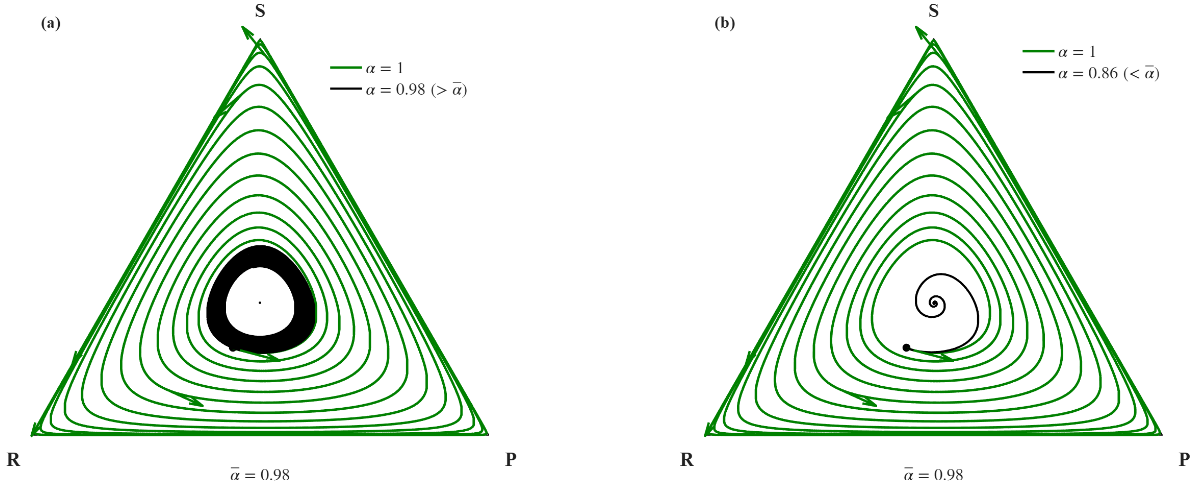

Figure 9: Memory-induced stability switching in a near zero-sum three-strategy game

Notes: Near zero-sum RPS with payoffs $A_\eta = A_0 - \varepsilon B$, $\eta = 0.05$, where A_0 is the antisymmetric RPS matrix and B has off-diagonal entries equal to one. Green: $\alpha = 1$. Black: fractional adjustment with $\alpha > \bar{\alpha}$ (panel a) and $\alpha < \bar{\alpha}$ (panel b). Common initial condition across trajectories.

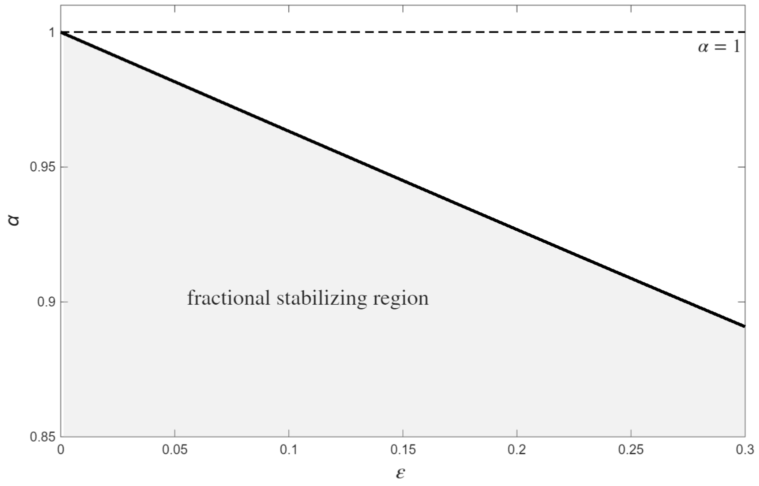

Figure 10: Fractional stability threshold in near zero-sum games

Notes: Stability threshold $\bar{\alpha}(\eta) = 2|\arg(\lambda_\eta)|/\pi$ for $A_\eta = A_0 - \eta B$. Shaded: (η, α) for which $p^\star$ is unstable at $\alpha = 1$ but α-stable under fractional adjustment. Dashed line: $\alpha = 1$.

To assess robustness, Figure 10 plots the analytical threshold $\bar{\alpha}(\eta)$ across η. The curve partitions the (η, α)-plane into an unstable region (above) and an α-stable region (below). The shaded area highlights the *fractional stabilizing region*—parameter values for which $p^\star$ is unstable under the classical replicator ($\alpha = 1$) yet becomes locally asymptotically stable once memory is sufficiently strong.

This section illustrates that local stability is not determined by incentives alone, but also by the temporal law of adjustment. When the payoff-induced vector field features significant rotation, memory dampens oscillatory components and restores dissipation. In this sense, time becomes a genuine bifurcation parameter governing dynamic stability. This kind od bifurcation arises in deterministic setting and does not rely on stochastic shocks, mutations, or random perturbations. It follows solely from the interaction between rotational incentives and history-dependent adjustment.

The mechanism is intrinsically multidimensional and cannot arise in two-strategy games, where the Jacobian has only real eigenvalues and stability reduces to the classical scalar condition. Fractional memory therefore expands the set of equilibria that are dynamically selected in the long run. Interior Nash equilibria that are not evolutionarily stable under the classical replicator may nevertheless become locally attracting once agents exhibit sufficiently persistent memory.

7 Concluding remarks

This paper studies how the temporal structure of adjustment shapes evolutionary selection. We introduce a fractional extension of the replicator dynamics that preserves payoff-based selection while allowing past payoff information to influence current behavior through long memory. The resulting fractional replicator dynamics leaves the equilibrium set and feasibility properties unchanged, but alters local stability, convergence rates, and equilibrium selection over finite horizons.

Our analysis yields three main insights. First, introducing long memory does not modify incentives or equilibrium structure: Nash equilibria remain stationary, and the simplex remains forward invariant for all fractional orders. Second, memory can affect stability. Equilibria that are unstable under the classical, memoryless replicator may become locally asymptotically stable once past experiences remain sufficiently salient, giving rise to memory-induced stability switching. Third, even when stability is unchanged, memory slows adjustment. Convergence becomes polynomial rather than exponential, creating a wedge between asymptotic basin geometry and finite-time selection.

These results challenge the implicit neutrality of time in standard evolutionary models. They show that sharp asymptotic predictions can coexist with persistent non-equilibrium behavior over economically relevant horizons, even in fully deterministic environments. Slow convergence, prolonged coexistence of strategies, or delayed coordination need not reflect noise, bounded rationality, or weak incentives. They can arise endogenously from the intrinsic dynamics of adjustment when past payoff information decays slowly.

Methodologically, our approach isolates a new channel through which evolutionary dynamics can be enriched without altering the selection mechanism itself. By modifying only the temporal operator, the fractional replicator preserves the microeconomic logic of payoff-monotone revision while allowing time to become a determinant of stability and selection. This perspective complements existing work emphasizing the role of revision protocols, by showing that even holding the payoff-induced vector field fixed, the law of temporal aggregation can matter.

Finally, our results point to several directions for future research. Fractional adjustment could be combined with stochastic perturbations to study the interaction between memory and noise. The framework may also be extended to other evolutionary or learning dynamics, such as best-response or reinforcement-based models. More generally, fractional analysis provides a tractable way to study persistence and history dependence in evolutionary systems, and its systematic integration into economic theory remains largely unexplored.

References

Acemoglu, D. and J. A. Robinson (2005). *Economic origins of dictatorship and democracy*. Cambridge university press.

Benaïm, M. and J. W. Weibull (2003). Deterministic approximation of stochastic evolution in games. *Econometrica 71*(3), 873–903.

Börgers, T. and R. Sarin (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory 77*(1), 1–14.

Brown, G. W. and J. Von Neumann (1950). Solutions of games by differential equations. In *Contributions to the Theory of Games*, pp. 73–79. Princeton University Press.

Cabrales, A. (2000). Stochastic replicator dynamics. *International Economic Review 41*(2), 451–481.

Camerer, C. and T. Hua Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica 67*(4), 827–874.

Cheung, Y.-W. and D. Friedman (1997). Individual learning in normal form games: Some laboratory results. *Games and economic behavior 19*(1), 46–76.

Choudhury, K. D. and T. Aydinyan (2023). Stochastic replicator dynamics: A theoretical analysis and an experimental assessment. *Games and Economic Behavior 142*, 851–865.

Cong, N. D., T. S. Doan, et al. (2016). Linearized asymptotic stability for fractional differential equations. *Electronic Journal of Qualitative Theory of Differential Equations 39*, 1–13.

Cressman, R. (2003). *Evolutionary dynamics and extensive form games*, Volume 5. MIT Press.

Cressman, R. and Y. Tao (2014). The replicator equation and other game dynamics. *Proceedings of the National Academy of Sciences 111*(supplement_3), 10810–10817.

Deng, W., C. Li, and J. Lü (2007). Stability analysis of linear fractional differential system with multiple time delays. *Nonlinear Dynamics 48*(4), 409–416.

Diethelm, K. (2010). *The Analysis of Fractional Differential Equations*. Springer.

Diethelm, K. and N. J. Ford (2002). Analysis of fractional differential equations. *Journal of Mathematical Analysis and Applications 265*(2), 229–248.

Diethelm, K., N. J. Ford, and A. D. Freed (2004). Detailed error analysis for a fractional adams method. *Numerical Algorithms 36*(1), 31–52.

Erev, I. and A. E. Roth (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Eeview*, 848–881.

Foster, D. and P. Young (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology 38*(2), 219–232.

Fréchette, G. R. (2012). Session-effects in the laboratory. *Experimental Economics 15*(3), 485–498.

Friedman, D. and R. Oprea (2012). A continuous dilemma. *American Economic Review 102*(1), 337–363.

Fudenberg, D. and D. K. Levine (1998). *The theory of learning in games*, Volume 2. MIT press.

Gilboa, I. and A. Matsui (1991). Social stability and equilibrium. *Econometrica*, 859–867.

Gopalsamy, K. (2013). *Stability and oscillations in delay differential equations of population dynamics*, Volume 74. Springer Science & Business Media.

Greif, A. (1994). Cultural beliefs and the organization of society: A historical and theoretical reflection on collectivist and individualist societies. *Journal of Political Economy 102*(5), 912–950.

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

Henry, B. I., T. Langlands, and P. Straka (2010). Fractional fokker-planck equations for subdiffusion with space-and time-dependent forces. *Physical Review Letters 105*(17), 170602.

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

Hofbauer, J. and K. Sigmund (1998). *Evolutionary games and population dynamics*. Cambridge university press.

Kalai, E. and E. Lehrer (1993). Rational learning leads to nash equilibrium. *Econometrica*, 1019–1045.

Kandori, M., G. J. Mailath, and R. Rob (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 29–56.

Matignon, D. (1996). Stability results for fractional differential equations with applications to control processing. In *Computational engineering in systems applications*, Volume 2, pp. 963–968. Lille, France.

Mertikopoulos, P., C. Papadimitriou, and G. Piliouras (2018). Cycles in adversarial regularized learning. In *Proceedings of the twenty-ninth annual ACM-SIAM symposium on discrete algorithms*, pp. 2703–2717. SIAM.

Mertikopoulos, P. and W. H. Sandholm (2016). Learning in games via reinforcement and regularization. *Mathematics of Operations Research 41*(4), 1297–1324.

Metz, J. A. and O. Diekmann (2014). *The dynamics of physiologically structured populations*, Volume 68. Springer.

North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge university press.

Obando, G., J. I. Poveda, and N. Quijano (2016). Replicator dynamics under perturbations and time delays. *Mathematics of Control, Signals, and Systems 28*(3), 20.

Rivero, M., S. V. Rogosin, J. A. Tenreiro Machado, and J. J. Trujillo (2013). Stability of fractional order systems. *Mathematical Problems in Engineering 2013*(1), 356215.

Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. MIT press.

Smith, J. M. and G. R. Price (1973). The logic of animal conflict. *Nature 246*(5427), 15–18.

Smith, M. J. (1984). The stability of a dynamic model of traffic assignment—an application of a method of lyapunov. *Transportation science 18*(3), 245–252.

Taylor, P. D. and L. B. Jonker (1978). Evolutionary stable strategies and game dynamics. *Mathematical biosciences 40*(1-2), 145–156.

Tuljapurkar, S. (2013). *Population dynamics in variable environments*, Volume 85. Springer Science & Business Media.

Weibull, J. W. (1997). *Evolutionary game theory*. MIT press.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 57–84.

A Mathematical Proofs

Preliminary. Throughout, let $f : \Delta^n \to \mathbb{R}^n$ denote the replicator vector field associated with the bounded payoff matrix $A = (a_{ij})$. For any population state $p \in \Delta^n$, individual payoffs are given by $\pi(p) = Ap$ and average payoff by $\bar{\pi}(p) = p^\top Ap$. The replicator dynamics are therefore defined componentwise as

$$f_i(p) = p_i\big(\pi_i(p) - \bar{\pi}(p)\big), \qquad i = 1, \ldots, n.$$

Proof of proposition 1. (i) Invariance of the simplex. Summing over i yields $\sum_{i=1}^n f_i(p) = \sum_{i=1}^n p_i \pi_i(p) - \bar{\pi}(p) \sum_{i=1}^n p_i = \bar{\pi}(p) - \bar{\pi}(p) \cdot 1 = 0$, for every $p \in \Delta^n$. Hence $\frac{d}{dt} \sum_i p_i(t) = 0$, so $\sum_i p_i(t) = \sum_i p_i(0) = 1$ for all t. Nonnegativity follows from boundary tangency: If $p_i = 0$ then $f_i(p) = 0 \cdot \big(\pi_i(p) - \bar{\pi}(p)\big) = 0$, so the vector field has no outward component through the face $\{p \in \Delta^n : p_i = 0\}$. Therefore every face of the simplex is forward invariant, and $p(t) \in \Delta^n$ for all $t \geq 0$ whenever $p(0) \in \Delta^n$.

(ii) Stationarity and Nash equilibria. See Weibull (1997), Proposition 3.5, for a detailed proof. If $p^\star$ is a symmetric Nash equilibrium, then every strategy used with positive probability yields the equilibrium payoff: $\pi_i(p^\star) = \bar{\pi}(p^\star)$ for all i such that $p_i^\star > 0$. Therefore, for each i, $f_i(p^\star) = p_i^\star\big(\pi_i(p^\star) - \bar{\pi}(p^\star)\big) = 0$, and for indices with $p_i^\star = 0$ we also have $f_i(p^\star) = 0$ by the factor $p_i^\star$. Hence $f(p^\star) = 0$, i.e., $p^\star$ is a stationary state. □

Proof of Proposition 2. See Weibull (1997), Proposition 3.10, for a detailed proof. The result is standard for payoff-monotone dynamics. A canonical argument constructs the divergence $V(p) = \sum_{i:p_i^\star > 0} p_i^\star \log\left(\frac{p_i^\star}{p_i}\right)$ as a strict local Lyapunov function. One then shows that $\dot{V}(p) = \pi(p,p) - \pi(p^\star, p) < 0$ in a neighborhood of the ESS $p^\star$, which establishes local asymptotic stability. □

Proof of proposition 3. Let $p^\star$ be a locally asymptotically stable fixed point and assume that it is hyperbolic on the tangent space, i.e. that every eigenvalue of the restricted Jacobian $J_T := Df(p^\star)\big|_{T\Delta^n}$, $T\Delta^n = \{x \in \mathbb{R}^n : \mathbf{1}^\top x = 0\}$, has strictly negative real part.

Since f is C^1 on Δ^n, the dynamics restricted to the invariant simplex define a C^1 local flow on the manifold Δ^n. By standard linearization for hyperbolic equilibria of C^1 vector fields, $p^\star$ is locally exponentially stable: there exist constants $C, c > 0$ and a neighborhood $\mathcal{N}(p^\star)$ such that for all $p_0 \in \mathcal{N}(p^\star)$, $\|p(t) - p^\star\| \leq Ce^{-ct}$, $t \geq 0$. Moreover, writing $x(t) = p(t) - p^\star \in T\Delta^n$, one has, as $t \to \infty$, $x(t) = e^{J_T t}x(0) + o\big(e^{-ct}\big)$, which is equivalent to the stated expansion $p(t) - p^\star = e^{J_T t}(p_0 - p^\star) + o\big(e^{-ct}\big)$. □

Proof of Proposition 4. This result follows Diethelm (2010), lemma 6.2. Fix $T > 0$ and $0 < \alpha < 1$. We prove the equivalence componentwise and therefore omit the index i for notational simplicity.

Preliminaries. For a locally integrable function $g : [0, T] \to \mathbb{R}$, define the Riemann–Liouville fractional integral of order $\alpha \in (0, 1)$ by

$$(I^{\alpha}g)(t) = \frac{1}{\Gamma(\alpha)} \int_0^t (t - \tau)^{\alpha-1} g(\tau)\, d\tau, \qquad t \in [0, T].$$

For an absolutely continuous function $x : [0, T] \to \mathbb{R}$, the Caputo fractional derivative of order $\alpha \in (0, 1)$ is given by

$${}^{C}D_t^{\alpha} x(t) = (I^{1-\alpha}\dot{x})(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t (t - \tau)^{-\alpha} \dot{x}(\tau)\, d\tau.$$

A fundamental identity linking these operators is

$$(I^{\alpha\,C}D^{\alpha}x)(t) = x(t) - x(0), \qquad t \in [0, T], \tag{6}$$

which holds for all absolutely continuous functions x. Equivalently,

$$x(t) = x(0) + \frac{1}{\Gamma(\alpha)} \int_0^t (t - \tau)^{\alpha-1}\, {}^{C}D_\tau^{\alpha} x(\tau)\, d\tau. \tag{7}$$

(⇒) Caputo implies Volterra formulation. Assume that $p(\cdot)$ satisfies the Caputo system (3) on $[0, T]$, i.e. ${}^{C}D_t^{\alpha} p(t) = f(p(t))$, $t \in [0, T]$. Applying (7) to $x(t) = p(t)$ and substituting ${}^{C}D_\tau^{\alpha} p(\tau) = f(p(\tau))$ yields the Volterra integral equation (4).

(⇐) Volterra formulation implies Caputo. Conversely, assume that $p(\cdot) \in C([0, T]; \Delta^n)$ satisfies

$$p(t) = p(0) + \frac{1}{\Gamma(\alpha)} \int_0^t (t - \tau)^{\alpha-1} f(p(\tau))\, d\tau, \qquad t \in [0, T].$$

In operator form, this can be written as $p(t) - p(0) = (I^{\alpha} f(p(\cdot)))(t)$. Since f is continuous on $[0, T]$, the right-hand side is absolutely continuous and the Caputo derivative is well-defined. Applying ${}^{C}D_t^{\alpha}$ to both sides and using the left-inverse property of the Caputo derivative: ${}^{C}D_t^{\alpha}(I^{\alpha}g)(t) = g(t)$, for all continuous g, we obtain: ${}^{C}D_t^{\alpha}\big(p(t) - p(0)\big) = f(p(t))$. Since the Caputo derivative of a constant vanishes, ${}^{C}D_t^{\alpha}(p(0)) = 0$, it follows that ${}^{C}D_t^{\alpha} p(t) = f(p(t))$, $t \in [0, T]$, which is precisely the Caputo formulation (3). □

Lemma 1. *For each $i \in \{1, \ldots, n\}$:*

(i) *(*Quasi-positivity*) For every $p \in \Delta^n$, $p_i = 0 \Rightarrow f_i(p) = 0$.*

(ii) *(Componentwise Lipschitz on Δ^n) $f_i : \Delta^n \to \mathbb{R}$ is Lipschitz (related to the ℓ^1 norm).*

Proof of lemma 1. (i) If $p_i = 0$, then $f_i(p) = p_i[(Ap)_i - p^\top Ap] = 0$.

(ii) Let $\| \cdot \|_1$ be the ℓ^1 norm and set $\|A\|_1 := \max_j \sum_{i=1}^n |a_{ij}|$. For any $p, q \in \Delta^n$, we have the standard bounds:

$$\begin{aligned}
&\|Ap - Aq\|_1 \leq \|A\|_1 \, \|p - q\|_1, \\
&|(p^\top Ap) - (q^\top Aq)| \leq \|A\|_1 \, \|p - q\|_1 + \|A\|_1 \, \|p - q\|_1 = 2\|A\|_1 \, \|p - q\|_1,
\end{aligned}$$

where we used $\|p\|_1 = \|q\|_1 = 1$ and $p^\top Ap - q^\top Aq = (p - q)^\top Ap + q^\top A(p - q)$.

Now, for each component,

$$\begin{aligned}
|f_i(p) - f_i(q)| &= \Big| p_i\big((Ap)_i - p^\top Ap\big) - q_i\big((Aq)_i - q^\top Aq\big) \Big| \\
&\leq |p_i - q_i| \cdot \Big| (Ap)_i - p^\top Ap \Big| \; + \; q_i \cdot \Big| (Ap - Aq)_i - \big(p^\top Ap - q^\top Aq\big) \Big|.
\end{aligned}$$

On Δ^n, payoffs are bounded, namely for all $r \in \Delta^n$, $|(Ar)_i| \leq \|A\|_1 \|r\|_1 = \|A\|_1$, and $|r^\top Ar| \leq \|A\|_1 \|r\|_1 \|r\|_1 = \|A\|_1$. Hence $\big|(Ar)_i - r^\top Ar\big| \leq 2\|A\|_1$. Consequently, for all i,

$$|f_i(p) - f_i(q)| \leq 2\|A\|_1 \, |p_i - q_i| + 3\|A\|_1 \, \|p - q\|_1.$$

Since $|p_i - q_i| \leq \|p - q\|_1$, the previous bound implies, for each i, $|f_i(p) - f_i(q)| \leq 5\|A\|_1 \, \|p - q\|_1$, so f_i is Lipschitz on Δ^n. Finally,

$$\|f(p) - f(q)\|_1 = \sum_{i=1}^n |f_i(p) - f_i(q)| \leq \sum_{i=1}^n 5\|A\|_1 \, \|p - q\|_1 = 5n\|A\|_1 \, \|p - q\|_1,$$

so the vector f is Lipschitz on Δ^n as well. □

Proof of proposition 5.

(a) Conservation of mass. By Proposition 4,

$$p_i(t) = p_i(0) + \frac{1}{\Gamma(\alpha)} \int_0^t (t - \tau)^{\alpha - 1} f_i(p(\tau)) \, d\tau.$$

Summing over i and using $\sum_i f_i(p) = 0$ yields $\sum_i p_i(t) = \sum_i p_i(0) = 1$ for all t.

(b) Positivity. Fix $i \in \{1, \ldots, n\}$. Define $f_i(p) = p_i \, g_i(p)$, where $g_i(p) := (Ap)_i - p^\top Ap$. Since g_i is continuous on the compact set Δ^n, it is bounded: there exists $M > 0$ such that

$|g_i(p)| \leq M$ for all $p \in \Delta^n$.

Assume by contradiction that p_i becomes negative at some time. Since p_i is continuous and $p_i(0) \geq 0$, there exists a first time $t^\star := \inf\{t \geq 0 : \ p_i(t) < 0\} \in (0, \infty)$, such that $p_i(t) \geq 0$ for all $t \in [0, t^\star]$, and $p_i(t^\star) = 0$.

Evaluating the Volterra representation at $t = t^\star$ yields

$$0 = p_i(t^\star) = p_i(0) + \frac{1}{\Gamma(\alpha)} \int_0^{t^\star} (t^\star - \tau)^{\alpha-1} \, p_i(\tau) g_i(p(\tau)) \, d\tau.$$

Taking absolute values and using the bound on g_i, we obtain

$$\left| \int_0^{t^\star} (t^\star - \tau)^{\alpha-1} \, p_i(\tau) g_i(p(\tau)) \, d\tau \right| \leq M \int_0^{t^\star} (t^\star - \tau)^{\alpha-1} \, p_i(\tau) \, d\tau. \tag{8}$$

We now control the right-hand side. Since p_i is continuous on $[0, t^\star]$ and satisfies $p_i(t^\star) = 0$, for every $\varepsilon > 0$ there exists $\delta \in (0, t^\star)$ such that $p_i(\tau) \leq \varepsilon$ for all $\tau \in [t^\star - \delta, t^\star]$. Decomposing the integral accordingly gives

$$\begin{aligned}\int_0^{t^\star} (t^\star - \tau)^{\alpha-1} p_i(\tau) \, d\tau &= \int_0^{t^\star - \delta} (t^\star - \tau)^{\alpha-1} p_i(\tau) \, d\tau + \int_{t^\star - \delta}^{t^\star} (t^\star - \tau)^{\alpha-1} p_i(\tau) \, d\tau \\ &\leq \Big(\sup_{t \in [0, t^\star]} p_i(t) \Big) \int_0^{t^\star - \delta} (t^\star - \tau)^{\alpha-1} \, d\tau + \varepsilon \int_{t^\star - \delta}^{t^\star} (t^\star - \tau)^{\alpha-1} \, d\tau.\end{aligned}$$

Both integrals are finite (they can be computed explicitly). In particular, the second term can be made arbitrarily small by choosing ε sufficiently small. It follows that the right-hand side of (8) can be made arbitrarily small.

Consequently, it follows that the Volterra integral term can be made arbitrarily small. Hence the equality $p_i(t^\star) = 0$ forces $p_i(0) = 0$. This contradicts the assumption $p_i(0) > 0$. Therefore, if $p_i(0) > 0$, no such first time $t^\star$ can exist, and thus $p_i(t) \geq 0$ for all $t \geq 0$.

If instead $p_i(0) = 0$, observe that the identically zero function $p_i(t) \equiv 0$ satisfies the i-th Volterra equation associated with the FRD. By Lemma 1.ii, this Volterra problem admits a unique solution for the given initial condition. It follows that $p_i(t) \equiv 0$ for all $t \geq 0$.

Since the index i was arbitrary, we conclude that $p(t) \in \mathbb{R}^n_+$ for all $t \geq 0$. Combined with conservation of mass, this implies that $p(t) \in \Delta^n$ for all $t \geq 0$. □

Proof of proposition 6.

(i) FRD-stationary ⇒ fixed point. If $p^\star$ is stationary under the FRD, then $p(t) \equiv p^\star$ is a

solution. Since the Caputo derivative of a constant is zero, ${}^{C}D_t^{\alpha} p_i(t) \equiv 0$, and by (3) we get $f_i(p^\star) = 0$ for all i. Hence $p^\star$ is a fixed point of the classical replicator.

(ii) Fixed point ⇒ FRD-stationary. Conversely, if $f(p^\star) = 0$, then in the Volterra representation,

$$p_i(t) = p_i(0) + \frac{1}{\Gamma(\alpha)} \int_0^t (t-\tau)^{\alpha-1} f_i(p^\star)\, d\tau = p_i^\star,$$

so $p(t) \equiv p^\star$ is a solution of the FRD.

(iii) Nash equilibria. If $p^\star$ is a Nash equilibrium, then $\pi_i(p^\star) = \bar{\pi}(p^\star)$ for all i in the support of $p^\star$, while $p_i^\star = 0$ outside the support. Therefore $f_i(p^\star) = p_i^\star[\pi_i(p^\star) - \bar{\pi}(p^\star)] = 0$ for all i, hence $p^\star$ is stationary under the FRD by (ii).

□

Lemma 2. *Let $p^\star \in \Delta^n$ and define its support $S(p^\star) := \{i \in \{1, \ldots, n\} : \; p_i^\star > 0\}$. Let*

$$F(p^\star) := \{p \in \Delta^n : \; p_i = 0 \textit{ for all } i \notin S(p^\star)\} \tag{9}$$

be the smallest face of the simplex containing $p^\star$. Then $F(p^\star)$ is forward invariant under the FRD. Moreover, $p^\star$ belongs to the relative interior of $F(p^\star)$.

Proof of lemma 2. We first show that $p^\star \in \mathrm{relint}(F(p^\star))$. By definition of $S(p^\star)$, we have $p_i^\star > 0$ for all $i \in S(p^\star)$ and $p_i^\star = 0$ for all $i \notin S(p^\star)$. The affine hull of $F(p^\star)$ is given by

$$\mathrm{aff}(F(p^\star)) = \{p \in \mathbb{R}^n : \; p_i = 0 \text{ for all } i \notin S(p^\star), \; \sum_{i \in S(p^\star)} p_i = 1\}.$$

Since $p_i^\star > 0$ for all $i \in S(p^\star)$, there exists $\varepsilon > 0$ such that every $p \in \mathrm{aff}(F(p^\star))$ satisfying $|p_i - p_i^\star| < \varepsilon$ for all $i \in S(p^\star)$ belongs to $F(p^\star)$. Hence $p^\star$ is an interior point of $F(p^\star)$ relative to its affine hull, that is, $p^\star \in \mathrm{relint}(F(p^\star))$.

We now prove forward invariance. Let $p(\cdot)$ be a solution of the FRD with initial condition $p(0) \in F(p^\star)$. Fix an index $i \notin S(p^\star)$. Then $p_i(0) = 0$ and we claim that $p_i(t) = 0$ for all $t \geq 0$. Using the Volterra representation of the Caputo system (Proposition 4) we have

$$p_i(t) = p_i(0) + \frac{1}{\Gamma(\alpha)} \int_0^t (t-\tau)^{\alpha-1} f_i(p(\tau))\, d\tau = \frac{1}{\Gamma(\alpha)} \int_0^t (t-\tau)^{\alpha-1} f_i(p(\tau))\, d\tau.$$

By quasi-positivity (Lemma 1(ii)), $f_i(p(\tau)) \geq 0$ whenever $p_i(\tau) = 0$. Since the kernel $(t - \tau)^{\alpha-1}$ is nonnegative and $p_i(0) = 0$, it follows that $p_i(t) \geq 0$ for all $t \geq 0$. Moreover, if $p_i(t)$ were strictly positive at some time, the integral representation would require $f_i(p(\tau)) > 0$

on a set of positive measure, which contradicts quasi-positivity at the boundary $p_i = 0$. Hence $p_i(t)$ cannot leave zero and must satisfy $p_i(t) = 0$ for all $t \geq 0$. Since the argument holds for every $i \notin S(p^\star)$, we conclude that $p(t) \in F(p^\star)$ for all $t \geq 0$. $\square$

Proof of Theorem 1. By Proposition 5 and Lemma 1(ii), our Caputo-based FRD is well-defined on Δ^n, since solutions exist uniquely and remain in the simplex.

Step 1 (reduction to the invariant face and tangent coordinates). Let $p^\star \in \Delta^n$ be a stationary state and let $F(p^\star)$ denote the smallest invariant face of the simplex containing $p^\star$ (see Eq. 9). By Lemma 2, $F(p^\star)$ is forward invariant and $p^\star \in \operatorname{relint}(F(p^\star))$.

Let $T := T_{p^\star}F(p^\star)$ be the tangent space to $F(p^\star)$ at $p^\star$. Choose a matrix $B \in \mathbb{R}^{n\times d}$, where $d = \dim T$, with full column rank such that $\operatorname{Im}(B) = T$, and let $L \in \mathbb{R}^{d\times n}$ be a left inverse of B, i.e. $L \times B = I_d$. Define local coordinates $z := L(p - p^\star)$, equivalently $p = p^\star + Bz$. Since $p^\star \in \operatorname{relint}(F(p^\star))$, there exists $r > 0$ such that $p^\star + Bz \in F(p^\star) \subset \Delta^n$ for all $\|z\| < r$. This parametrizes a neighborhood of $p^\star$ within the invariant face.

By linearity of the Caputo derivative and since $p^\star$ and B are constant, we have

$$ {}^{C}D_t^{\alpha} p(t) = {}^{C}D_t^{\alpha}(p^\star + Bz(t)) = B\, {}^{C}D_t^{\alpha} z(t). $$

Substituting into ${}^{C}D_t^{\alpha} p(t) = f(p(t))$ and multiplying by L yields the reduced Caputo system in $\mathbb{R}^d$:

$$ {}^{C}D_t^{\alpha} z(t) = L\, f(p^\star + Bz(t)). \tag{10} $$

Step 2 (linearization and tangent Jacobian). Since f is C^1 near $p^\star$ and $f(p^\star) = 0$, Taylor expansion gives: $f(p^\star + Bz) = Df(p^\star)\, Bz + r(Bz)$, with $\|r(Bz)\|/\|z\| \to 0$ as $z \to 0$.

Inserting this into (10) yields

$$ {}^{C}D_t^{\alpha} z(t) = \mathcal{A}z(t) + \rho(z(t)), \tag{11} $$

where $\mathcal{A} := L\, Df(p^\star)\, B$, and $\rho(z) := L\, r(Bz) = o(\|z\|)$ as $z \to 0$.

Since $F(p^\star)$ is invariant and $\sum_i f_i \equiv 0$, the Jacobian $Df(p^\star)$ maps T into itself, so the tangent Jacobian $J_T := Df(p^\star)\big|_T$ is well-defined. Moreover, for all z, we have $J_T(Bz) = Df(p^\star)(Bz) = B(\mathcal{A}z)$, which implies the intertwining relation $J_T \circ B = B \circ \mathcal{A}$. Hence $\mathcal{A}$ is a matrix representation of J_T, and in particular $\sigma(\mathcal{A}) = \sigma(J_T)$.

Step 3 (Matignon linearization criterion). By the linearization principle for Caputo systems applied to (11), the equilibrium $z = 0$ is locally asymptotically stable if and only if

$$|\arg(\lambda)| > \frac{\alpha\pi}{2} \qquad \text{for all } \lambda \in \sigma(\mathcal{A}). \tag{12}$$

This criterion was first established in Matignon (1996) and subsequently refined in Cong et al. (2016), Theorem 3.1, and Diethelm (2010), Theorem 7.20.

Using $\sigma(\mathcal{A}) = \sigma(J_T)$ and the local equivalence $p = p^\star + Bz$, condition (12) is equivalent to α-stability of $p^\star$ under the FDR (in the sense of Definition 3). □

Proof of Theorem 2. Let $p^\star \in \Delta^n$ be a stationary state of the FDR By Proposition 6, stationary states of FRD coincide with those of the classical replicator, so $p^\star$ is also stationary for $\alpha = 1$. Local stability is governed by the tangent Jacobian J_T restricted to $T_{p^\star}F(p^\star)$, as established in Theorem 1.

(i) Instability for $\alpha = 1$. When $\alpha = 1$, FRD reduces to the classical replicator ODE. By Theorem 1 specialized to $\alpha = 1$, local asymptotic stability requires $|\arg(\mu)| > \pi/2$ for all $\mu \in \sigma(J_T)$, which is equivalent to $\Re(\mu) < 0$. By assumption, $\vartheta := \min_{\mu\in\sigma(J_T)} |\arg(\mu)| < \pi/2$, so this condition is violated. Hence $p^\star$ is locally unstable under the classical replicator dynamics.

(ii) Stability for sufficiently small α. By Theorem 1, $p^\star$ is α-stable if and only if

$$|\arg(\mu)| > \frac{\alpha\pi}{2} \qquad \text{for all } \mu \in \sigma(J_T). \tag{13}$$

Define the spectral angle margin $\vartheta := \min_{\mu\in\sigma(J_T)} |\arg(\mu)| \in (0, \pi/2)$. Set $\bar{\alpha} := 2\vartheta/\pi \in (0, 1)$. For any $\alpha \in (0, \bar{\alpha})$, we have $\alpha\pi/2 < \vartheta$, and therefore $|\arg(\mu)| \geq \vartheta > \alpha\pi/2$ for all $\mu \in \sigma(J_T)$. Thus condition (13) holds, and $p^\star$ is locally stable for all $\alpha \in (0, \bar{\alpha})$. □

Proof of corollary 1. By α_0-stability and Theorem 1, all $\lambda \in \sigma(J_T)$ satisfy $|\arg(\lambda)| > \alpha_0\pi/2$. For any $\alpha \in (0, \alpha_0]$ one has $\alpha\pi/2 \leq \alpha_0\pi/2$, so the same inequality implies, i.e. $|\arg(\lambda)| > \alpha\pi/2, \forall\lambda \in \sigma(J_T)$. The conclusion follows again from Theorem 1. □

Proof of Theorem 3. Fix $\alpha \in (0, 1)$ and let $p^\star$ be an α-stable stationary state. We use the same local parametrization of the invariant face and the same tangent coordinates as in the proof of Theorem 1. In particular, there exist matrices B and L (with $L \times B = I$) such that, for p close to $p^\star$ within the invariant face, $p = p^\star + Bz$, $z = L(p - p^\star)$, and the FRD is locally equivalent to the reduced Caputo system

$${}^{C}D_t^\alpha z(t) = \mathcal{A}z(t) + \rho(z(t)), \qquad z(0) = z_0 := L(p(0) - p^\star), \tag{14}$$

where $\mathcal{A} := L\,Df(p^\star)\,B$ and $\rho(z) = o(\|z\|)$ as $z \to 0$. Moreover, the intertwining relation $J_T \circ B = B \circ \mathcal{A}$ holds, hence $\sigma(\mathcal{A}) = \sigma(J_T)$.

Since $p^\star$ is α-stable, Theorem 1 implies that $|\arg(\lambda)| > \alpha\pi/2$ for all $\lambda \in \sigma(J_T)$, and hence the same sector condition holds for $\sigma(\mathcal{A})$. By the linearization principle for Caputo systems (see Matignon, 1996; Cong et al., 2016; Diethelm, 2010), the equilibrium $z = 0$ is asymptotically stable and the solution admits the variation-of-constants representation

$$z(t) = E_\alpha(\mathcal{A}t^\alpha)\,z_0 + \int_0^t (t-\tau)^{\alpha-1} E_{\alpha,\alpha}\big(\mathcal{A}(t-\tau)^\alpha\big)\,\rho(z(\tau))\,\frac{d\tau}{\Gamma(\alpha)},$$

where $E_{\alpha,\beta}$ is the two-parameter Mittag–Leffler function $E_{\alpha,\beta}(s) := \sum_{k=0}^\infty s^k/\Gamma(\alpha k + \beta)$.

Standard resolvent estimates for Mittag–Leffler functions under the sector condition (see Diethelm, 2010, Theorems 7.3 and 7.2) imply that the integral term is $o(\|z_0\|)$ as $\|z_0\| \to 0$, uniformly in $t \geq 0$. Therefore, we can write $z(t) = E_\alpha(\mathcal{A}t^\alpha)\,z_0 + o(\|z_0\|)$. Returning to $p(t) = p^\star + Bz(t)$ and using $E_\alpha(J_T t^\alpha)B = BE_\alpha(\mathcal{A}t^\alpha)$ yields

$$p(t) - p^\star = E_\alpha(J_T t^\alpha)\,(p(0) - p^\star) + o(\|p(0) - p^\star\|),$$

which proves the first claim. Finally, under the same sector condition, the matrix Mittag–Leffler function satisfies the polynomial decay estimate $\|E_\alpha(\mathcal{A}t^\alpha)\| \leq C\,t^{-\alpha}$, for all sufficiently large t and for some $C > 0$ (Diethelm, 2010, Theorem 7.3). This implies $\|p(t) - p^\star\| \leq K\,t^{-\alpha}$ for large t. □

Proof of proposition 8. We consider the 2×2 coordination game where the two corners $k \in \{0, 1\}$ are locally asymptotically stable for the classical replicator, and therefore α-stable for FRD (Proposition 7 or directly by Theorem 1, since $\lambda_k := f'(k) < 0$).

Fix $k \in \{0, 1\}$ and define the deviation $x(t) := p(t) - k$. The dynamics near k can be written as a scalar Caputo equation ${}^C D_t^\alpha x(t) = \lambda_k x(t) + o(x(t))$.

Step 1 (local Mittag–Leffler approximation). By Theorem 3, there exists a neighborhood $\mathcal{U}_k$ of k such that, for every initial condition $p_0 \in \mathcal{U}_k$,

$$x(t) = E_\alpha(\lambda_k t^\alpha)\,x(0) + o(|x(0)|), \tag{15}$$

Moreover, since $\lambda_k < 0$, the Mittag–Leffler function satisfies the asymptotic bound (see

Diethelm, 2010, Thm 7.3): there exists $t_k > 0$ such that for all $t \geq t_k$,

$$0 \leq E_\alpha(\lambda_k t^\alpha) \leq \frac{1}{|\lambda_k|\,\Gamma(1-\alpha)}\, t^{-\alpha}\,(1+o(1))\,. \tag{16}$$

Hence, shrinking $\mathcal{U}_k$ if necessary, there exists $\eta_k > 0$ such that for all $|x(0)| \leq \eta_k$ and all $t \geq t_k$,

$$|x(t)| \leq \frac{2\,|x(0)|}{|\lambda_k|\,\Gamma(1-\alpha)}\, t^{-\alpha}. \tag{17}$$

Step 2 (uniform time-to-threshold bound). Define $\Lambda := \min\{|\lambda_0|, |\lambda_1|\} > 0, \eta := \min\{\eta_0, \eta_1\}$, and $\bar{t} := \max\{t_0, t_1\}$. Fix $k \in \{0,1\}$ and assume $|p_0 - k| \leq \eta$, so that (17) applies. Let $T \geq \bar{t}$ satisfy

$$\left(\frac{\eta}{\varepsilon\Lambda\Gamma(1-\alpha)}\right)^{1/\alpha} \leq T \leq \left(\frac{2\eta}{\varepsilon\Lambda\Gamma(1-\alpha)}\right)^{1/\alpha}.$$

Then, using $|x(0)| \leq \eta$ and $|\lambda_k| \geq \Lambda$, inequality (17) yields

$$|p(T) - k| = |x(T)| \leq \frac{2\eta}{\Lambda\,\Gamma(1-\alpha)}\, T^{-\alpha} \leq \varepsilon.$$

Therefore, any initial condition within distance η of k reaches the ε-neighborhood of k by time T, namely $p_0 \in [0,\eta] \implies T_\varepsilon^{(0)} \leq T$, and $p_0 \in [1-\eta, 1] \implies T_\varepsilon^{(1)} \leq T$.

Step 3 (positive probability mass). By assumption, the initial density satisfies $g(p) \geq m > 0$ on $[0,1]$; hence, $\Pr(p_0 \in [0,\eta]) \geq m\,\eta$, and $\Pr(p_0 \in [1-\eta, 1]) \geq m\,\eta$. On the event $\{p_0 \in [0,\eta]\}$, the trajectory reaches the ε-neighborhood of 0 by time T, so selection of s_0 by time T occurs with probability at least $m\eta$. A symmetric argument applies to s_1. Therefore, $\mathbb{P}_T(s_0) \geq m\,\eta$, and $\mathbb{P}_T(s_1) \geq m\,\eta$. ☐

Proof of Corollary 2. Assume that $A + A^\top = 0$ and let $p^\star \in \Delta^n$ be a Nash equilibrium. By Proposition 6, stationary states of the classical replicator and of the FRD coincide. Hence $p^\star$ is stationary under FRD for all $\alpha \in (0,1]$.

Since A is antisymmetric, one has $p^\top A p = 0$ for all $p \in \Delta^n$, and the replicator vector field simplifies to $f_i(p) = p_i(Ap)_i$. Linearizing at $p^\star$ and restricting to the tangent space of the invariant face $F(p^\star)$ yields an antisymmetric Jacobian J_T. As a consequence, all nonzero eigenvalues of J_T are purely imaginary, i.e. $\lambda = \pm i\omega$ with $\omega \neq 0$, and therefore satisfy $|\arg(\lambda)| = \pi/2$.

When $\alpha = 1$, the local stability condition of Theorem 1 requires $|\arg(\lambda)| > \pi/2$ for all $\lambda \in \sigma(J_T)$. Since this strict inequality fails for the nonzero eigenvalues, the equilibrium $p^\star$

is not locally asymptotically stable under the classical replicator dynamics. Because J_T has no eigenvalue with positive real part, the equilibrium is neutrally stable.

For any $\alpha \in (0,1)$, the fractional stability condition becomes $|\arg(\lambda)| > \alpha\pi/2$. Since $\alpha\pi/2 < \pi/2$, this condition is satisfied by all nonzero eigenvalues of J_T. Therefore, $p^\star$ is locally asymptotically stable under FRD for every $\alpha < 1$, i.e. it is α-stable for all $\alpha \in (0,1)$. Finally, the stability threshold in Theorem 2 is $\bar{\alpha} = 2\vartheta/\pi$, where $\vartheta = \min_{\lambda\in\sigma(J_T)} |\arg(\lambda)| = \pi/2$ in the zero-sum case. Hence $\bar{\alpha} = 1$. $\square$

Proof of Corollary 3. For η small enough, the interior Nash equilibrium $p^\star = (1/3, 1/3, 1/3)$ persists and, by Proposition 6, is stationary under the FRD for all $\alpha \in (0,1]$. Let $J_T(\eta) := Df_\eta(p^\star)|_{T_{p^\star}F(p^\star)}$ denote the tangent Jacobian at $p^\star$. By assumption, $p^\star$ is locally unstable under the classical replicator ($\alpha = 1$). By Theorem 1 specialized to $\alpha = 1$, this implies that the classical stability condition fails, i.e. $\min_{\lambda\in\sigma(J_T(\eta))} |\arg(\lambda)| < \pi/2$. Define the spectral angle margin $\vartheta(\eta) := \min_{\lambda\in\sigma(J_T(\eta))} |\arg(\lambda)| \in (0, \pi/2)$, and $\bar{\alpha}(\eta) := 2\vartheta(\eta)/\pi$, then $\bar{\alpha}(\eta) \in (0,1)$. For any $\alpha < \bar{\alpha}(\eta)$, we have $\alpha\pi/2 < \vartheta(\eta)$, and hence $|\arg(\lambda)| > \alpha\pi/2$ for all $\lambda \in \sigma(J_T(\eta))$. Applying Theorem 1 yields that $p^\star$ is locally asymptotically stable under the fractional replicator dynamics of order α. $\square$

DOCUMENTS DE TRAVAIL GREDEG PARUS EN 2026
GREDEG Working Papers Released in 2026

2026-01	Estelle Malavolti & Frédéric Marty *L'équilibre économique des aéroports secondaires européens à l'épreuve des dynamiques de concurrence fiscale*
2026-02	Charlie Joyez *A New Index of Export-Import Proximity: Conceptual Foundations and Global Patterns*
2026-03	Patrice Bougette & Frédéric Marty *The SCP Paradigm Revisited: What Structuralism Really Contributed to U.S. Antitrust*
2026-04	Imen Bouhlel, Nathalie Lazaric & Paolo Zeppini *Competitive Diffusion and Sustainability Transitions: The Case of Plastics Recycling Technologies*
2026-05	Maxime Menuet *Fractional Replicator Dynamics*