Replicator Dynamics

Game Theory · Module 7 — Evolutionary Game Theory

Node 7.1

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Evolutionary Game Theory **Node:** 7.1

Classical game theory, through Module 6, rests on a single image: rational agents who know the game, reason about it, and select strategies that are best responses to what other rational agents will select. Module 7 replaces that image with a different one. There are no agents doing reasoning. There is a population of carriers — organisms, firms, poker players — each hard-wired or culturally locked into a strategy, and the strategies that do well *reproduce*, while the strategies that do poorly *disappear*. Equilibria are not selected by introspection but by differential survival. Node 7.1 introduces the canonical dynamics of this picture: the **replicator equation** of Taylor and Jonker (1978), which translates the Darwinian slogan "fitness differences drive frequency changes" into a precise ODE on the simplex. Everything downstream in Module 7 — ESS, Hawk-Dove, population games, stochastic stability, adaptive dynamics — is built on this one dynamical system, or on a variant of it. And the most striking result of the node is that in spite of being derived without any reference to rationality, replicator dynamics recover Nash equilibrium as their rest points. The classical and evolutionary pictures meet at the same fixed set.

Formal Definition

Fix a finite symmetric two-player game with strategy set $S = \{1, 2, \dots, n\}$ and payoff matrix $A \in \mathbb{R}^{n\times n}$, where A_{ij} is the payoff to a player using strategy i against an opponent using strategy j. Consider an infinite population in which each agent is locked into a single pure strategy. Let $x_i(t)$ denote the **frequency** (relative proportion) of strategy i in the population at time t, with

$$x(t) = (x_1(t), \dots, x_n(t)) \in \Delta^{n-1}, \qquad \sum_i x_i = 1, \quad x_i \geq 0.$$

The **expected fitness** of strategy i against the current population mixture is

$$f_i(x) = (Ax)_i = \sum_j A_{ij} x_j,$$

and the **mean population fitness** is

$$\bar{f}(x) = x^\top A x = \sum_i x_i f_i(x).$$

The **replicator equation** (Taylor–Jonker, 1978) is the system of ODEs

$$\dot{x}_i = x_i\big(f_i(x) - \bar{f}(x)\big), \qquad i = 1, \dots, n.$$

Interpretation. A strategy grows in frequency if and only if its fitness exceeds the population average. A strategy below average shrinks. The multiplicative factor x_i enforces the boundary: a strategy at frequency zero stays at zero (it has no carriers to reproduce), so the faces of the simplex are invariant.

Theorem (invariance of the simplex). The simplex Δ^{n-1} is forward-invariant under the replicator equation: $x(0) \in \Delta^{n-1}$ implies $x(t) \in \Delta^{n-1}$ for all $t \geq 0$. Each face $\{x_i = 0\}$ and each vertex e_i is invariant.

Theorem (Taylor–Jonker folk theorem, 1978). The rest points of the replicator equation in $\operatorname{int} \Delta^{n-1}$ coincide with the set of symmetric Nash equilibria of $(A, A^\top)$ with full support. Every symmetric Nash equilibrium is a rest point. Every strict symmetric Nash equilibrium is an asymptotically stable rest point.

The replicator equation is nonlinear (quadratic in x) but has the remarkable property that it is coordinate-wise linear in the fitness difference $f_i - \bar{f}$. This makes it the gentlest possible "rationality-free" dynamic that still recovers Nash equilibria as its rest set.

Intuition

Imagine a population of strategies, each represented by some number of individuals. Every generation, each individual plays the game against a random opponent drawn from the population, accumulates a payoff, and this payoff determines its reproductive rate. High-payoff strategies leave more offspring (or, in cultural versions, are imitated more). Low-payoff strategies leave fewer. Over many generations, the frequencies x_i evolve, and the evolution is driven by the gap between each strategy's fitness and the population average. The replicator equation is the continuous-time limit of this logic.

Two structural features deserve emphasis.

First, the dynamic is **payoff-monotone** but not **best-response driven**. A strategy with slightly-above-average fitness grows slightly; a strategy with much-above-average fitness grows rapidly; but no strategy ever "jumps" to the best response in one step. This is what makes replicator dynamics a model of **gradual adaptation** through differential reproduction, not of rational recalculation. Best-response dynamics (Module 9) sit at the other extreme: agents instantly switch to the optimal action. Replicator dynamics are the slow, smooth, population-level version.

Second, the dynamic has **no memory and no foresight**. Agents do not plan for the future or remember the past. The current state $x(t)$ determines the current velocity $\dot{x}(t)$, and nothing else matters. This is a radical departure from the Bayesian rationality of Modules 1–4, but it captures something biologists and cultural theorists have long understood: *most* strategy change in real populations is not driven by explicit optimization. Firms do not solve Bellman equations to decide pricing; species do not compute ESSs to decide how aggressively to defend territory. What happens is that bad strategies leave fewer descendants — literal descendants in biology, metaphorical descendants (imitators, students, survivors) in economics and poker. Replicator dynamics is the minimal mathematical model of this fact.

The link to Nash equilibrium is the surprise of the theory. We wrote down a dynamical system that has no mention of rationality, best response, or equilibrium. Yet its rest points are exactly the symmetric Nash equilibria. This is not a coincidence and it is not trivial to prove — it falls out of the identity $\dot{x}_i = x_i(f_i - \bar{f})$: a rest point in the interior has $f_i = \bar{f}$ for all i with $x_i > 0$, which is precisely the indifference condition of a Nash equilibrium. So the machinery of rational choice and the machinery of differential reproduction converge on the same object.

Structural Role

7.1 is the gateway node of Module 7 and the structural bridge between the "reasoning" picture (Modules 1–6) and the "selection" picture (Modules 7–9):

- **7.2 Evolutionarily stable strategies** refines the rest-point characterization by asking which Nash equilibria are actually stable under replicator dynamics (not just fixed points).
- **7.3 Hawk-Dove and asymmetric contests** is the canonical biological example, where the replicator equation is solved explicitly and the interior rest point is an ESS.
- **7.4 Population games** generalizes the ODE to continuum-of-agent games where each agent controls a negligible share and fitness is a general function of the population distribution.
- **7.5 Stochastic evolution and mutation** adds a noise term, producing processes like the Moran process and Kandori–Mailath–Rob stochastic stability.
- **7.6 Adaptive dynamics** zooms into the "mutant invasion" logic and derives the canonical equation governing slow evolution of continuous traits.

The replicator equation also exports to Module 9 (learning dynamics): it is formally equivalent to **exponential weights / multiplicative weights update** in continuous time, which gives the classical Lyapunov connection between no-regret learning and Nash equilibrium. This is one of the deepest cross-links in the course: a biological dynamic written down in 1978 turns out to be the continuous-time limit of an online-learning algorithm written down in the 1990s. The structural reason is that both describe "how to update a distribution over options when each option has been getting a fitness/utility signal."

Mathematical Development

Derivation from population dynamics

Let $N_i(t)$ denote the number of individuals using strategy i, and let $N(t) = \sum_i N_i(t)$ be the total population. Suppose N_i grows at a rate proportional to its fitness:

$$\dot{N}_i = f_i(x)N_i, \qquad x_i = N_i/N.$$

Then $\dot{N} = \sum_i f_i(x)N_i = N\bar{f}(x)$. For the frequency $x_i = N_i/N$:

$$\dot{x}_i = \frac{\dot{N}_i N - N_i \dot{N}}{N^2} = \frac{f_i N_i}{N} - \frac{N_i \bar{f}}{N} = x_i(f_i - \bar{f}).$$

The replicator equation is the natural dynamic on frequencies induced by the exponential-growth dynamic on absolute numbers. The nonlinearity (multiplication by x_i) comes from the change of variables from counts to frequencies.

Invariance and boundary behavior

Simplex invariance. Summing the replicator equation,

$$\sum_i \dot{x}_i = \sum_i x_i f_i - \bar{f}\sum_i x_i = \bar{f} - \bar{f} = 0,$$

so $\sum_i x_i$ is conserved. If $x(0) \in \Delta^{n-1}$ then $x(t) \in \Delta^{n-1}$ for all t.

Face invariance. If $x_i = 0$ then $\dot{x}_i = 0$: once extinct a strategy stays extinct. This is the single most important structural property of the replicator: it cannot "re-invent" a strategy that has died out. Stochastic dynamics (7.5) repair this by adding a mutation term.

Linearity invariance

Proposition. Adding a column-constant vector to the payoff matrix does not change the replicator dynamic.

Proof. Let $A' = A + \mathbf{1}c^\top$. Then $f_i'(x) = f_i(x) + c^\top x$, so $f_i' - \bar{f}' = f_i - \bar{f}$. The replicator equation is unchanged. $\square$

The dynamic depends only on *differences* within each row, a positive-affine invariance inherited from the utility invariance of 1.1.

Rest points and Nash equilibrium

Theorem. A state $x^* \in \operatorname{int} \Delta^{n-1}$ is a rest point of the replicator equation iff $f_i(x^*) = \bar{f}(x^*)$ for all i, iff x^* is a symmetric Nash equilibrium of $(A, A^\top)$ with full support.

Proof sketch. Rest point $\Leftrightarrow x_i(f_i - \bar{f}) = 0$ for all i. In the interior $x_i > 0$, so $f_i = \bar{f}$ for all i — the classical indifference condition of mixed Nash equilibrium. The vertices e_i are always rest points (monomorphic populations). □

Lyapunov function for interior equilibria

Suppose $x^* \in \operatorname{int} \Delta^{n-1}$ is an interior Nash equilibrium. Define the relative entropy

$$H(x^*, x) = \sum_i x_i^* \log \frac{x_i^*}{x_i}.$$

Theorem. Along solutions of the replicator equation,

$$\frac{d}{dt} H(x^*, x) = -\big((x^* - x)^\top Ax\big).$$

If x^* is an ESS (see 7.2), then $(x^* - x)^\top Ax > 0$ for all x near x^*, so H is strictly decreasing: x^* is asymptotically stable with H as Lyapunov function.

This is the central stability result connecting replicator dynamics to evolutionary stability. The relative entropy is the natural "distance" on the simplex, and ESS is exactly the condition under which it decreases along trajectories.

Worked Example

Rock-Paper-Scissors

Symmetric two-player game, $S = \{R, P, S\}$, payoff matrix

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

The unique symmetric Nash equilibrium is $x^* = (1/3, 1/3, 1/3)$.

Rest point. Compute $Ax^* = (0, 0, 0)^\top$: every row sums to zero. So $f_i(x^*) = 0$ for all i and $\bar{f}(x^*) = 0$. Thus $\dot{x}_i = 0$ at x^*.

Stability. Linearize at x^*. The relevant Jacobian (on the tangent subspace to the simplex) has purely imaginary eigenvalues $\pm i/\sqrt{3}$. The linearization is a center — not asymptotically stable and not unstable.

Check a candidate Lyapunov: $H(x^*, x) = -(1/3) \sum_i \log x_i - \log 3$. Along replicator trajectories,

$$\dot{H} = -(1/3) \sum_i \dot{x}_i / x_i = -(1/3) \sum_i (f_i - \bar{f}) = 0,$$

since $\sum_i f_i = 0$ by the antisymmetry of A. So H is **conserved** along trajectories. **Rock-Paper-Scissors replicator dynamics has closed orbits** encircling x^* — the population cycles forever.

Starting at $x = (0.7, 0.2, 0.1)$, rock dominates, so paper grows (paper beats rock) and scissors decays. Later paper dominates and scissors grows. Then scissors dominates and rock grows. The system traces a closed loop around x^*. Nash equilibrium need not be an attractor: it is only a rest point.

Comparison: a stable perturbation

Replace A with $A - \epsilon I$ for small $\epsilon > 0$. Now trajectories spiral *inward* and x^* becomes asymptotically stable. The relative entropy H is strictly decreasing. This is the generic picture for a stable mixed ESS.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Population frequencies $x_i(t)$	Proportion of regulars using playing style i at time t
Payoff matrix A	Expected win rate (bb/100) of style i against style j
Strategy (pure)	A locked-in style: LAG, TAG, nit, solver-monkey, exploit-hunter
Fitness $f_i(x)$	Expected win rate of style i against the current population mix
Mean fitness $\bar{f}(x)$	Average win rate at the tables (near zero, or negative from rake)
Replicator equation	Style frequencies change over time driven by differential success
Rest point (interior)	Metagame equilibrium: every style has equal win rate against the mix

The mapping is not metaphorical. Poker populations literally evolve by differential reproduction: players whose styles lose money drop out (financial extinction); players whose styles make money remain and are imitated by new entrants (cultural reproduction).

Concrete situation: the evolution of the poker meta

In 2006 the dominant winning style was LAG (loose-aggressive). By 2012 it was tight-aggressive “regs” grinding 6-max. By 2016 it was solver-trained GTO. By 2022 it was exploit-aware GTO with explicit population reads. Each transition looks, at the population level, exactly like a replicator trajectory: a style with above-average fitness grew until it saturated the population; a counter-style (with higher fitness against the now-dominant population) emerged and grew; the cycle continued.

Take the 2006–2012 transition. In 2006, the population was predominantly loose recreational players, and LAG had high fitness against them. As LAGs grew in frequency, the population became more aggressive, and LAG-vs-LAG was lower in fitness than LAG-vs-recreational. Tight-aggressive style, which exploited LAG aggression, now had the highest fitness. TAG grew, and the population cycled. This is not a metaphor — it is a replicator equation in action, with A given by expected win-rate matrices and $x(t)$ given by observed style frequencies over time.

Replicator dynamics predicts that in a game with cyclic dominance (style A beats B, B beats C, C beats A) the population cycles forever, not converges. This is what is observed in the poker meta-game over decades, with recurring oscillation between aggressive and tight styles.

What this understanding buys you

The meta is not an agreement, it is a dynamical system. When people talk about “the meta right now,” they are describing the current state $x(t)$ of a replicator equation. The meta is not chosen; it is the current frequency distribution. It will change driven by fitness differentials. You cannot fight the meta any more than you can fight gravity.

Your edge is a rate, not a state. Your fitness as a player is $f_i(x)$ — your expected win rate *against the current population*. A style that wins today may lose tomorrow because the population changed, not because the style did. Most player "declines" are actually meta shifts that moved the fitness differential against a frozen style.

The "solver wins" claim is a replicator theorem. GTO tends to win over decades because it is the rest point that is stable against invasion (7.2). But in pure cyclic structures, even GTO only produces *neutral* stability — the population oscillates around it.

Cross-Links

- **7.2 Evolutionarily stable strategies** — asking which rest points are Lyapunov-stable.
- **7.3 Hawk-Dove** — the canonical example with an explicit solvable replicator equation.
- **7.4 Population games** — the generalization to continuum populations and general fitness functionals.
- **9.2 Fictitious play and best-response dynamics** — the "rational" cousin of replicator.
- **9.4 Multiplicative weights update** — the discrete-time online-learning algorithm whose continuous limit is the replicator.
- **2.1 Nash equilibrium** — the classical concept that replicator dynamics recover as rest points.
- **Statistics 4.x KL divergence** — the Lyapunov function for replicator dynamics is the relative entropy.
- **Physics 3.x Hamiltonian systems** — Rock-Paper-Scissors replicator is a conservative system with closed orbits.

Structural Compression

Invariant: The rest points of the replicator equation $\dot{x}_i = x_i(f_i(x) - \bar{f}(x))$ in the interior of the simplex coincide with the symmetric Nash equilibria of the underlying game; strict Nash equilibria are asymptotically stable; cyclic-dominance games produce closed orbits around the interior mixed Nash; and relative entropy to an ESS is a Lyapunov function for its stability.

Mechanism: Frequencies grow in proportion to the gap between strategy fitness and mean fitness; nonlinearity from the change of variables from counts to frequencies; invariance of the simplex and of each face; linearity invariance under column-additive payoff transformations; the rest-point condition reduces to the indifference condition of mixed Nash.

Cross-System Echo: The replicator equation is formally identical to the continuous-time limit of **multiplicative weights update** in online learning (9.4), where "fitness" is per-expert reward and "frequency" is probability mass on each expert. It is also the Kolmogorov forward equation for a Moran-like birth-death process in the infinite-population limit (7.5). And it is the gradient flow of expected payoff in the Shahshahani metric on the simplex — a Riemannian geometry in which relative entropy is distance. Evolutionary, learning-theoretic, and geometric interpretations are structurally the same object, which is why replicator dynamics keeps reappearing across fields.

Exercises

1. Derive the replicator equation from the exponential population dynamic $\dot{N}_i = f_i(x)N_i$ by computing $\dot{x}_i$ where $x_i = N_i/N$.

2. For the two-strategy symmetric game with payoff matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, write out the replicator equation in the single variable $x = x_1 \in [0, 1]$, find all rest points, and classify them (stable / unstable) as a function of a, b, c, d.
3. For Rock-Paper-Scissors, verify by direct computation that $H(x^*, x) = -(1/3)\sum_i \log x_i - \log 3$ is conserved along replicator trajectories. Numerically integrate from a generic initial condition and plot the orbit in the simplex.
4. Show that if we replace A with $A + \mathbf{1}c^\top$ (column-constant shift), the replicator dynamic is unchanged. Conclude that the dynamic depends only on "row differences."
5. Prove that every strict symmetric Nash equilibrium of $(A, A^\top)$ is an asymptotically stable rest point of the replicator equation. (Hint: for a strict equilibrium $x^* = e_i$, show $f_i > f_j$ for all $j \neq i$ in a neighborhood of e_i.)
6. Consider the replicator in a game where strategy 3 is strictly dominated by strategy 1. Show $x_3(t) \to 0$ as $t \to \infty$ from any interior initial condition. This is the "iterated elimination of strictly dominated strategies" theorem for replicator dynamics.
7. **Argue, structurally,** why the replicator equation recovers Nash equilibria as its rest points despite making no reference to rationality or best response. What is it about the "fitness difference drives frequency change" mechanism that produces the same fixed set as conscious optimization? Connect to the online-learning interpretation (multiplicative weights update) and to the gradient-flow interpretation on the Shahshahani metric.

Game Theory · Module 7 — Evolutionary Game Theory · Node 7.1
Concepts: replicator-dynamics, dynamical-systems, nash-equilibrium, matrix-representation, expectation-value
Prerequisites: 1.1, 1.2, 2.1, 2.2
Enables: 7.2, 7.3, 7.4, 7.5, 7.6

hbar.university — own.your.cognition