Adaptive Dynamics

Game Theory · Module 7 — Evolutionary Game Theory

Node 7.6

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Evolutionary Game Theory **Node:** 7.6

Adaptive dynamics is the theory of evolution in **continuous strategy spaces**. Nodes 7.1-7.5 worked with discrete strategy sets: agents choose from a finite menu, and the population state is a distribution over that menu. Adaptive dynamics drops the finiteness assumption and studies what happens when strategy is a continuous variable – a real number, a vector, a function – and evolution proceeds by small mutations in strategy space. The key objects are: the **invasion fitness function** (whether a rare mutant with a slightly different strategy can invade the resident population), the **canonical equation of adaptive dynamics** (an ODE describing how the resident strategy evolves under successive invasions), and **evolutionary branching** (the phenomenon where a population splits into two coexisting types at a singular point). Adaptive dynamics is where evolutionary game theory makes contact with **dynamical systems theory** and, through the thesis's lens, with the **control-theoretic view of optimization** developed in VQA Modules 6-7. The canonical equation is a gradient-like flow on strategy space; evolutionary branching is a bifurcation; adaptive dynamics is, structurally, optimization by a population of controllers who cannot see the objective function directly but can measure local fitness gradients through mutation and selection.

Formal Definition

Setting. A population of agents, each characterized by a strategy $x \in \mathcal{X} \subseteq \mathbb{R}^d$ (the **trait space**). The population is monomorphic (all agents share the same trait x) except for rare mutants.

Invasion fitness. When a rare mutant with strategy y appears in a resident population at strategy x, the **invasion fitness** is

$$f(y; x) = \text{per-capita growth rate of mutant } y \text{ in the environment set by resident } x.$$

The mutant invades iff $f(y; x) > 0$. The resident is in ecological equilibrium: $f(x; x) = 0$ (zero growth at equilibrium).

Selection gradient. The local direction of selection at resident strategy x is

$$D(x) = \left.\frac{\partial f(y; x)}{\partial y}\right|_{y=x},$$

the gradient of invasion fitness with respect to the mutant strategy, evaluated at the resident. If $D(x) > 0$, mutants with slightly higher strategies invade; if $D(x) < 0$, mutants with slightly lower strategies invade.

Canonical equation of adaptive dynamics (Dieckmann-Law 1996). Under the assumptions that mutations are rare, small, and symmetric, the expected trajectory of the resident strategy is

$$\dot{x} = \mu(x) \cdot \sigma^2(x) \cdot N(x) \cdot D(x),$$

where $\mu(x)$ is the mutation rate, $\sigma^2(x)$ is the mutation variance, and $N(x)$ is the population size at equilibrium. Ignoring the scaling, the canonical equation is a **gradient ascent** on invasion fitness:

$$\dot{x} \propto D(x) = \nabla_y f(y;x)\big|_{y=x}.$$

Singular strategies. A strategy x^* is a **singular strategy** if $D(x^*) = 0$ – a critical point of the adaptive dynamics.

Classification of singular strategies. A singular x^* is:

- **Convergence stable (CSS)** if the adaptive dynamics converge to x^* from nearby initial conditions: $\frac{d}{dx}D(x)\big|_{x=x^*} < 0$.
- **Evolutionarily stable (ESS)** if no nearby mutant can invade the resident at x^*: $\frac{\partial^2 f(y;x^*)}{\partial y^2}\big|_{y=x^*} < 0$.
- **Evolutionary branching point** if x^* is convergence-stable but *not* ESS: the dynamics carry the population to x^*, but once there, mutants on both sides can invade, and the population splits into two coexisting lineages that diverge.

Intuition

Adaptive dynamics captures an important idea: **evolution is a controller that has no access to the objective function, only to local fitness gradients measured through mutation and selection**. The population "explores" strategy space by producing mutants with slightly different traits. Mutants that are fitter invade; mutants that are less fit perish. The net effect is that the population climbs the fitness landscape – but the landscape changes as the population moves, because the fitness function depends on the resident strategy. This is the key difference from optimization: in adaptive dynamics, the objective (invasion fitness) is a function of both the mutant strategy and the resident strategy, so the landscape shifts under the optimizer's feet.

This is exactly the structure of the thesis's "separation of objective from dynamics" principle from VQA Module 6. The objective is $f(y;x)$ – invasion fitness – and it depends on the current state x. The dynamics are the canonical equation $\dot{x} \propto D(x)$. The optimizer (selection) does not see f directly; it only sees $D(x)$, the local gradient. The separation is clean: the dynamics are gradient-like, but the gradient is of a state-dependent objective.

Evolutionary branching is the adaptive-dynamics analogue of a bifurcation in dynamical systems: a smooth parameter change (the population reaching a singular point) causes a qualitative change in the dynamics (from monomorphism to dimorphism). The branching phenomenon is uniquely biological – classical optimization does not have it, because classical objectives do not depend on the optimizer's position in the way that invasion fitness depends on the resident.

The broader lesson is that adaptive dynamics is a concrete example of a "controller with incomplete information about the plant": the population is the controller, the ecological environment is the plant, and the controller has access only to local fitness gradients, not to the global fitness landscape. This connects directly to the HOPSO-as-controller framework of VQA Module 10.7 and to the learning-in-games framework of Module 9.

Structural Role

7.6 is the closing node of Module 7, connecting evolutionary dynamics to the broader dynamical-systems and control-theoretic themes that will be developed in Modules 9 and 10.

- **7.1 (Replicator dynamics)** is the discrete-strategy special case; adaptive dynamics extends the replicator to continuous traits.
- **7.4 (Population games)** provides the population-level framework; adaptive dynamics adds the trait-space dimension.

- **7.5 (Stochastic evolution)** studies mutations in discrete strategy spaces; adaptive dynamics studies mutations in continuous ones.
- **9.5 (Convergence to equilibrium)** will use the selection-gradient framework to study when learning dynamics converge.
- **9.6 (RL in games)** will use the "controller with local fitness gradient" view.
- **VQA 6-7 (Optimization dynamics, control-theoretic view)** is the thesis's own framework that the canonical equation parallels.

The thesis voice enters here: the canonical equation $\dot{x} \propto D(x)$ is a dynamical system on strategy space, and the population is a controller that navigates this space using only local information. The separation of the invasion fitness (objective) from the canonical equation (dynamics) is the same architectural principle that the thesis argues for in VQA optimizer design.

Mathematical Development

Derivation of the canonical equation

Start with a monomorphic resident population at strategy x. A rare mutant appears with strategy $y = x + \epsilon\xi$, where ξ is a random mutation direction and ϵ is small. The mutant's invasion fitness is

$$f(x + \epsilon\xi; x) \approx \epsilon\xi \cdot D(x) + O(\epsilon^2),$$

using $f(x; x) = 0$. The mutant invades with probability approximately proportional to $\max(0, f(y; x))$.

Averaging over many rare-mutation events (mutations arrive at rate μ, have variance σ^2, and the population size at equilibrium is N):

$$\dot{x} = \mu\sigma^2 N \cdot D(x).$$

This assumes: (i) mutations are rare (so the population is always monomorphic between invasions), (ii) mutations are small (so the first-order expansion is valid), and (iii) establishment probabilities are proportional to fitness advantage (the branching-process approximation).

Pairwise invasibility plots (PIPs)

A **PIP** is a graphical tool: the (x, y)-plane where x is the resident strategy and y is the mutant strategy, with regions where $f(y; x) > 0$ (invasion possible) shaded. The diagonal $y = x$ is always the zero contour. The selection gradient is the slope of the invasion-positive region at the diagonal.

From the PIP, one can read off: direction of selection (which side of the diagonal is fitter), singular strategies (where the zero contour crosses the diagonal tangentially), and ESS vs branching (whether the zero contour curves toward or away from the diagonal at singular points).

Evolutionary branching: the canonical example

Consider a model where organisms compete for a resource along a niche axis $x \in \mathbb{R}$. The carrying capacity is $K(x)$, peaked at some optimal trait value. The competition kernel $\alpha(x, y)$ describes how much a resident at x is affected by a competitor at y. Classic assumption: K is broader than α (there is more niche space than competition covers).

The invasion fitness is

$$f(y; x) = r\left(1 - \frac{\alpha(y, x)N^*(x)}{K(y)}\right),$$

where $N^*(x) = K(x)/\alpha(x, x)$ is the equilibrium population size.

At the singular strategy x^* (the trait maximizing carrying capacity), the selection gradient vanishes. If the competition kernel is narrower than the carrying capacity curve, the second derivative of invasion fitness is

positive at x^* – meaning x^* is not ESS. The population converges to x^* (convergence stability) but then branches: mutants on both sides of x^* can invade, and the population splits into two lineages that evolve away from each other. This is **evolutionary branching** – the onset of speciation in the adaptive-dynamics framework.

Adaptive dynamics as gradient flow

When the invasion fitness derives from a symmetric payoff function $\pi(x, y) = \pi(y, x)$, the selection gradient simplifies. For potential games, the canonical equation becomes exact gradient ascent on the potential. For non-potential games, the dynamics are gradient-*like* but not true gradient – there may be limit cycles or chaos. This mirrors the VQA setting: HOPSO's dynamics are gradient-like (include a gradient term) but not pure gradient (include momentum, swarm attraction, and noise).

Connection to replicator dynamics

In the limit of small mutations in a finite-strategy replicator, the adaptive dynamics on a one-dimensional trait space recovers a discretized version of the canonical equation. The replicator dynamics describe what happens when mutations are not rare (many strategies coexist); adaptive dynamics describe what happens when mutations are rare (the population is nearly monomorphic and evolves by successive invasions). The two are complementary regimes of the same evolutionary process.

Worked Example

Adaptive dynamics in a symmetric contest game

Two contestants each choose a level of aggression $x, y \in [0, \infty)$. Payoff to the mutant:

$$\pi(y, x) = \frac{y}{y + x} \cdot V - y \cdot c,$$

where V is the prize value and c is the per-unit cost of aggression.

Selection gradient.

$$D(x) = \left.\frac{\partial \pi(y, x)}{\partial y}\right|_{y=x} = \frac{x}{(2x)^2} V - c = \frac{V}{4x} - c.$$

Singular strategy. $D(x^*) = 0$: $x^* = V/(4c)$.

Convergence stability. $D'(x) = -V/(4x^2) < 0$ for all $x > 0$. Convergence-stable.

ESS check. $\partial^2 \pi / \partial y^2 \big|_{y=x=x^*} = -V/(4(x^*)^2) < 0$. ESS confirmed.

Numerical example. $V = 8, c = 1$. Singular strategy: $x^* = 2$. The population evolves toward aggression level 2. Payoff at equilibrium: $\pi(2, 2) = 4 - 2 = 2$.

Interpretation as a control problem. The population is "optimizing" by climbing the selection gradient. The singular point x^* is where the gradient vanishes. The dynamics are a natural-selection implementation of gradient descent on the contest cost. But the objective $\pi(x, x) = V/2 - cx$ is not the quantity being maximized – what selection maximizes is invasion fitness $f(y; x)$ at $y = x$, which has the same critical point but different curvature. This mismatch is the reason convergence stability and ESS are distinct conditions.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Continuous strategy x	A player's aggression parameter: 3-bet frequency, c-bet size, bluff-to-value ratio
Invasion fitness $f(y; x)$	The EV of a slightly different strategy y against the current population playing x
Selection gradient $D(x)$	The direction in which a small strategy tweak improves EV against the current field
Canonical equation	The trajectory of the meta's aggregate strategy over time
Singular strategy x^*	The GTO equilibrium or the meta's fixed point
Evolutionary branching	The meta splitting into two coexisting archetypes (e.g., nits and LAGs both being viable)

Concrete situation: the c-bet frequency meta

Consider the evolution of continuation-bet frequency on the flop as a continuous variable $x \in [0, 1]$. Early online poker: the meta was to c-bet nearly always ($x \approx 0.8$). This was a local ESS: deviating slightly was unprofitable because opponents folded too much. But the meta attracted exploiters who started check-raising against the high c-bet frequency, and the selection gradient reversed. The meta evolved toward a singular strategy around $x^* \approx 0.55$, where the selection gradient vanished. At this point, both "c-bet more in some spots" and "c-bet less in some spots" could invade – an evolutionary branching event. The result was a split meta: players who c-bet aggressively in certain board textures and passively in others, coexisting with players who take different branches.

This branching is exactly the adaptive-dynamics prediction: convergence to a singular point, followed by diversification because the singular point is CSS but not ESS.

What this understanding buys you

The meta evolves by gradient ascent on a moving landscape. When you study the meta's evolution over time – increasing aggression, then decreasing, then stabilizing – you are watching the canonical equation play out. The landscape moves because opponent strategies change in response to yours.

Evolutionary branching explains why diverse styles coexist. If the meta were a simple ESS, only one style would survive. The fact that multiple viable styles coexist suggests the meta is near a branching point where multiple niches are available.

Cross-Links

- **7.1 Replicator dynamics** – the discrete-strategy parent framework.
- **7.4 Population games** – the population-level model.
- **7.5 Stochastic evolution** – mutations in discrete spaces; adaptive dynamics extends to continuous.
- **9.5 Convergence to equilibrium** – learning dynamics as the strategic analogue.
- **9.6 RL in games** – learning as a controller navigating fitness landscapes.
- **VQA 6 (Optimization dynamics)** – the canonical equation parallels the optimizer-as-controller framework.
- **VQA 10.7 (HOPSO)** – a structured controller navigating a landscape with local gradient information.
- **Biological systems – speciation** – evolutionary branching as the onset of speciation.

Structural Compression

Invariant: Adaptive dynamics describes evolution in continuous strategy spaces via the canonical equation $\dot{x} \propto D(x)$, a gradient-like flow driven by the selection gradient of invasion fitness. Singular strategies where $D = 0$ can be convergence-stable, evolutionarily stable, both, or neither; the distinction between CSS and ESS creates the possibility of evolutionary branching.

Mechanism: Rare small mutations generate variation; selection retains fitter mutants; the net effect is gradient ascent on the invasion fitness landscape. The landscape is state-dependent (it depends on the resident strategy), so the dynamics are not simple gradient flow but a coupled system where the optimizer and the objective co-evolve.

Cross-System Echo: Adaptive dynamics is the natural-selection implementation of the thesis's "optimizer as controller" principle. The canonical equation is a first-order controller navigating strategy space using only local gradient information from a state-dependent objective – exactly the structure of gradient-descent optimizers applied to parameterized quantum circuits in VQA. Evolutionary branching is the analogue of a bifurcation in a controlled dynamical system. The connection to HOPSO is direct: HOPSO's population of particles exploring a noisy landscape with local gradient feedback is a computational analogue of the adaptive-dynamics population exploring trait space via mutation and selection.

Exercises

1. For the symmetric contest game $\pi(y, x) = yV/(y + x) - yc$ with $V = 10, c = 2$, compute the singular strategy, verify convergence stability and ESS, and sketch the PIP.
2. Consider a resource-competition model where $K(x) = \exp(-x^2/2)$ and $\alpha(x, y) = \exp(-(x - y)^2/(2\sigma_\alpha^2))$. For what values of σ_α does the singular strategy exhibit evolutionary branching?
3. Derive the canonical equation from the assumption of rare small symmetric mutations and a branching-process establishment probability. Identify each assumption explicitly.
4. For a symmetric two-player game with payoff $\pi(x, y)$, carefully distinguish the convergence-stability condition (about the dynamics near x^*) from the ESS condition (about the local curvature of invasion fitness). Show by example that they can disagree.
5. In the Hawk-Dove game with continuous aggression x, let the payoff include a quadratic fighting cost: $\pi(y, x) = yV/(y + x) - cy - \beta y^2$. How does the quadratic cost change the singular strategy and its classification?
6. Construct a symmetric game where the singular strategy is CSS but not ESS (a branching point). Describe what happens qualitatively after branching.
7. **Argue, structurally,** why adaptive dynamics is the natural-selection implementation of the "optimizer as controller" principle from VQA Modules 6-7. Identify the state, the dynamics, the objective, and the feedback loop in both settings. What does evolutionary branching correspond to in the optimization context?

Game Theory · Module 7 — Evolutionary Game Theory · Node 7.6
Concepts: dynamical-systems, gradient-descent, bifurcation, stability, adaptive-systems
Prerequisites: 7.1, 7.4, 7.5
Enables: 9.5, 9.6

hbar.university — own.your.cognition