Chapter 16

Evolutionary Game Theory

Maynard Smith's export of game theory to biology replaced rationality with *selection*: strategies that earn higher payoffs (fitness) reproduce faster. The resulting theory illuminates animal conflict, the evolution of cooperation, social norms—and, read as a model of imitation and learning, the behavior of decidedly non-Darwinian agents like firms and traders.

16.1 Evolutionarily stable strategies

Consider a large population whose members are randomly matched in pairs to play a symmetric two-player game with payoff $u(a, b)$ to a player using a against b.

Definition 16.1: Evolutionarily stable strategy (Maynard Smith–Price 1973)

A (possibly mixed) strategy σ^* is an *evolutionarily stable strategy* (ESS) if for every mutant strategy $\sigma \neq \sigma^*$ there is an invasion barrier $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$:

$$u(\sigma^*,\, (1-\varepsilon)\sigma^* + \varepsilon\sigma) \;>\; u(\sigma,\, (1-\varepsilon)\sigma^* + \varepsilon\sigma).$$

Equivalently: (i) (σ^*, σ^*) is a symmetric Nash equilibrium, and (ii) if $u(\sigma, \sigma^*) = u(\sigma^*, \sigma^*)$ for a mutant $\sigma \neq \sigma^*$, then $u(\sigma^*, \sigma) > u(\sigma, \sigma)$.

Key Idea: Stability without rationality

ESS refines symmetric Nash equilibrium using no rationality at all: condition (i) says no mutant does strictly better against the incumbent population; condition (ii) says that neutral mutants (equally good against incumbents) must do strictly worse against *themselves*, so they cannot drift in. Equilibrium emerges from differential survival, not deliberation—answering the oldest objection to Nash equilibrium ("real agents don't compute fixed points").

Example 16.1: Hawk–Dove revisited

In Example 5.3 with $C > V$, the mixed strategy $\sigma^* =$ Hawk with probability V/C is the unique ESS: a population of pure Doves is invaded by Hawks (condition (i) fails), pure Hawks by Doves, and the mixed population resists all mutants. Prediction: ritualized limited conflict in nature, with escalation frequency governed by the value/cost ratio—one of theoretical biology's most celebrated explanations.

Example 16.2: ESS may fail to exist

Rock–Paper–Scissors has a unique symmetric equilibrium (uniform mixing), but it is *not* an ESS: any pure mutant earns exactly the equilibrium payoff against the incumbent and ties against itself, so condition (ii) fails with equality. Nature concurs: the side-blotched lizard *Uta stansburiana* has three male morphs (orange beats blue beats yellow beats orange) whose frequencies *cycle* over years instead of settling—exactly the dynamic behavior the missing ESS predicts.

16.2 Replicator dynamics

Definition 16.2: Replicator dynamics (Taylor–Jonker 1978)

Let $x_a(t)$ be the population share playing pure strategy a, and let $u(a, x)$ be its expected payoff against the population mix x. The *replicator equation* is

$$\dot{x}_a \;=\; x_a\big[\,u(a,x) - \bar{u}(x)\,\big], \qquad \bar{u}(x) = \textstyle\sum_b x_b\, u(b,x)\;:$$

strategies grow in proportion to their payoff advantage over the population average.

Theorem 16.1: Folk theorem of evolutionary dynamics

For the replicator dynamics: (i) every Nash equilibrium is a rest point; (ii) every stable rest point is a Nash equilibrium; (iii) every ESS is asymptotically stable (Hofbauer–Schuster–Sigmund); (iv) strictly dominated strategies are eliminated along interior trajectories. The converses generally fail: rest points need not be Nash (unstable vertices), and asymptotically stable states need not be ESS.

Example 16.3: Stag Hunt dynamics and basins of attraction

For the Stag Hunt (Example 3.2), the replicator dynamics on the share x of Stag players has three rest points: $x = 0$, $x = 1$, and the mixed equilibrium $x^* = \frac{1}{2}$ (with those payoffs, indifference at $4x = 2$, i.e. $x^* = 1/2$). Both pure equilibria are stable; the mixed one is the unstable watershed between their basins. Which convention evolves depends on initial conditions—and the *risk-dominant* equilibrium has the larger basin. Stochastic evolutionary models (Kandori–Mailath–Rob 1993; Young 1993) go further: with rare mutations, the population spends almost all its long-run time at the risk-dominant equilibrium. Evolution provides the equilibrium *selection* theory that rationality alone could not (Chapter 4).

16.3 The evolution of cooperation

How can cooperation evolve when defection is dominant in the one-shot Prisoner's Dilemma? The main mechanisms (Nowak's "five rules"), each translatable into a precise payoff condition:

1. **Direct reciprocity**: repeated interaction with the same partner (Chapter 9); cooperation viable when the continuation probability exceeds the cost–benefit ratio c/b.
2. **Indirect reciprocity**: reputation—help those who help others; viable when the probability your action is observed exceeds c/b. ("For direct reciprocity you need a face; for indirect reciprocity, a name.")

3. **Kin selection**: Hamilton's rule $rb > c$ (relatedness r).
4. **Network/spatial reciprocity**: cooperators forming clusters on graphs survive when b/c exceeds the average degree (Ohtsuki–Nowak).
5. **Group selection**: competition between groups can favor within-group cooperators.

Historical Note: From lizards to language

Evolutionary game theory now spans: animal signaling (the handicap principle as costly signaling, cf. Chapter 11); the evolution of fairness norms (ultimatum-game behavior, Example 6.4, as an evolved commitment device); microbial games (bacteriocin-producing *E. coli* strains playing rock–paper–scissors in petri dishes); cancer modeling (tumor subclones as competing strategies, with "evolutionary therapy" designed as a game against the tumor); and the evolution of language conventions (signaling games where meaning itself is the equilibrium, cf. the babbling discussion in Chapter 11).

16.4 Learning in games (the bridge to algorithms)

Replicator dynamics also describes *imitation* in human populations, and its discrete-time cousin is precisely the *multiplicative weights* algorithm: agents increase the probability of strategies in proportion (exponentially) to accumulated payoffs. This is no coincidence—evolution, learning, and optimization dynamics form one mathematical family, taken up in Chapter 17: no-regret learning converges (in time average) to coarse correlated equilibria in general games and to minimax play in zero-sum games, while last-iterate behavior may cycle or be chaotic (as the lizards already demonstrated).

Exercises

Exercise 16.1: ESS verification

For the symmetric game with payoffs $u(A,A) = 2$, $u(A,B) = 5$, $u(B,A) = 1$, $u(B,B) = 3$: find all symmetric Nash equilibria and determine which are ESS.

Exercise 16.2: Replicator phase line

Write the replicator equation for the Hawk–Dove game with $V = 2$, $C = 4$ as a single ODE for the Hawk share x. Find all rest points, determine their stability by the sign of $\dot{x}$, and sketch the phase line.

Exercise 16.3: Dominated strategies die

Prove claim (iv) of Theorem 16.1 for a strategy a strictly dominated by pure strategy b: show $\frac{d}{dt}\log(x_a/x_b) \le -\varepsilon < 0$ along interior replicator trajectories, so $x_a \to 0$.

Exercise 16.4: Invasion barriers

In the Stag Hunt of Example 3.2, a fraction ε of Stag-playing mutants enters a population of Hare players. Compute the payoff of each strategy as a function of ε and find the minimal mutant mass needed for Stag to spread. Interpret: why are "better" conventions hard to

ignite, and what real-world interventions (assurance contracts, coordinated switching dates) lower the barrier?