

Evolution & Learning in Games

Econ 243B

Jean-Paul Carvalho

Lecture 1: Foundations of Evolution & Learning in Games I

Classical Game Theory

We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and preferable.

von Neumann and Morgenstern (1944, p. 4445)

Classical Game Theory

The classical approach to game theory provides:

- ▶ A mathematical description of a strategic interaction.
- ▶ A prediction of play based on equilibrium analysis:
 - ▶ the leading *solution concept* being Nash equilibrium.

A Coordination Game

Stag Hunt

	<i>H</i>	<i>S</i>
<i>H</i>	<u>2</u> , <u>2</u>	0, 0
<i>S</i>	0, 0	<u>3</u> , <u>3</u>

NE:

- ▶ $(H, H), (S, S)$
- ▶ $(\frac{1}{3}, \frac{2}{3})$ for both players.

An Anti-Coordination Game

Hawk Dove

	<i>Hawk</i>	<i>Dove</i>
<i>Hawk</i>	-2	<u>0</u>
<i>Dove</i>	<u>4</u>	0

NE:

- ▶ $(Hawk, Dove), (Dove, Hawk)$
- ▶ $(\frac{2}{3}, \frac{1}{3})$ for both players.

Classical Requirements for Nash Equilibrium Play

The traditional justification for NE as a “rational” outcome of a game requires three things.

1. Players are **rational**:

- ▶ They can formulate *strategies* that maximize their payoff given what everyone else is doing.

2. Players have **knowledge** of the game they are playing:

- ▶ They know the strategy set,
- ▶ They know the payoffs generated by each strategy profile.

3. Players have **equilibrium knowledge**:

- ▶ They can correctly anticipate what other players will do.

Classical Requirements for Nash Equilibrium Play

All of these requirements are problematic:

1. Players are **rational**:

- ▶ They can formulate *strategies* that maximize their payoff given what everyone else is doing. (**consider chess**)

2. Players have **knowledge** of the game they are playing:

- ▶ They know the strategy set, (**consider sports**)
- ▶ They know the payoffs generated by each strategy profile. (**consider traffic congestion**)

3. Players have **equilibrium knowledge**:

- ▶ They can correctly anticipate what other players will do. (**consider games with multiple equilibria such as Stag Hunt and Hawk Dove**)

The Problem of Equilibrium Knowledge

Is the following not reasonable?

- ▶ Player 1 chooses *stag* expecting player 2 to choose *stag*. But, player 2 chooses *hare*, expecting player 1 to choose *hare*.
- ▶ Player 1 chooses *hawk* expecting player 2 to choose *dove*. But, player 2 chooses *hawk*, expecting player 1 to choose *dove*.

These are both collections of **rationalizable strategies**:

- ▶ Each player's strategy is a best response to some *rationalizable belief* about the other player's strategy.

The Problem of Equilibrium Selection

Another problem is that of **multiple equilibria**.

- ▶ Nash theorem guarantees the existence of at least one NE in a large class of games.
- ▶ However, it does NOT in general make a *unique* prediction of play.
- ▶ Which NE is the most “plausible” prediction?
 - ▶ The equilibrium refinements program emerged to address this question.
 - ▶ Placed unreasonable burdens on the rationality of players.

An Evolutionary Approach

Partly in response to the shortcomings of the equilibrium refinements program, a new field emerged known as **evolutionary game theory**.

The evolutionary approach to the social sciences is based on:

- ▶ **boundedly rational**
- ▶ **populations** of agents,
- ▶ who may (or may not) **learn** or **evolve** their way into equilibrium,
- ▶ by **gradually revising**
- ▶ **simple, myopic rules** of behavior.

An Evolutionary Approach

We shall now take up the “mass-action” interpretation of equilibrium points... It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.

John Nash (PhD thesis, p. 21)

Out-of-Equilibrium Dynamics

By specifying an explicit dynamic process of adjustment to equilibrium, analyses of evolution and learning in games enable us to address the following questions:

1. Are equilibria stable in the face of random shocks? (**local stability**)
2. Does evolution/learning lead to an equilibrium from any initial state? (**global convergence and nonconvergence**)
3. Which equilibria, if any, does an evolutionary dynamic lead to? (**equilibrium selection**)

Evolutionary Game Theory: The Biological Approach

- ▶ Game theory was initially developed by mathematicians and economists.
- ▶ Evolutionary biologists adapted these techniques/concepts in developing evolutionary game theory—see for e.g. the pioneering work of British biologist John Maynard Smith. EGT was later imported back into economics.
- ▶ Owing to this intellectual history, and because social scientific approaches share some deep similarities with the biological approach, we shall start by reviewing the basic biological approach to evolution.

The Biological Approach

Ingredients:

1. Inheritance:

- ▶ Players are *programmed* with a strategy. (Players are essentially strategies.)

2. Selection:

- ▶ Strategies that do better, given what everyone else is doing, proliferate.
- ▶ In particular, payoffs are interpreted as *reproduction rates* of strategies.
- ▶ Extends Darwin's notion of survival of the fittest from an exogenous environment to an interactive setting.

3. Mutation:

- ▶ Equilibrium states can be perturbed by random shocks.
- ▶ To be *stable*, an equilibrium must be resistant to invasion by "mutant strategies".

The Strategic Context

Players

- ▶ The population consists of a continuum of agents.

Strategies

- ▶ The set of (pure) strategies is $S = \{1, \dots, n\}$, with typical members i, j and s .
- ▶ The mass of agents programmed with strategy i is m_i , where $\sum_{i=1}^n m_i = m$.
- ▶ Let $x_i = \frac{m_i}{m}$ denote the proportion of players choosing strategy $i \in S$.

The Strategic Context

Population States

- ▶ The set of population states (or strategy distributions) is $X = \{x \in \mathbb{R}_+^n : \sum_{i \in S} x_i = 1\}$.
- ▶ X is the unit simplex in \mathbb{R}^n .
- ▶ The set of vertices of X are the pure population states—those in which all agents choose the same strategy.
- ▶ These are the standard basis vectors in \mathbb{R}^n :

$$e_1 = (1, 0, 0, \dots), e_2 = (0, 1, 0, \dots), e_3 = (0, 0, 1, \dots), \dots$$

Payoffs

- ▶ A *continuous* payoff function $F : X \rightarrow \mathbb{R}^n$ assigns to each population state a vector of payoffs, consisting of a real number for each strategy.
- ▶ $F_i : X \rightarrow \mathbb{R}$ denotes the payoff function for strategy i .

Some Useful Facts

Consider the expected payoff to strategy i if i is matched with another strategy drawn uniformly at random from the population to play the following *two-player* game:

	1	2	...	n
i	$u(i, 1)$	$u(i, 2)$...	$u(i, n)$

The expected payoff to strategy i in state x is:

$$\begin{aligned} F_i(x) &= x_1 u(i, 1) + x_2 u(i, 2) \dots + x_n u(i, n) \\ &= \sum_{j=1}^n x_j u(i, j) \\ &= \sum_{j=1}^n x_j F_i(e_j). \end{aligned}$$

Some Useful Facts

The **average payoff** in the population is:

$$\begin{aligned}\bar{F}(x) &= x_1F_1(x) + x_2F_2(x) \dots + x_nF_n(x) \\ &= \sum_{i=1}^n x_iF_i(x).\end{aligned}$$

Note: this is the same as the payoff from playing the mixed strategy x against itself.

The Replicator Dynamic

- ▶ Suppose that payoffs represent *fitness* (rates of reproduction) and reproduction takes place in continuous time.
- ▶ This yields a continuous-time evolutionary dynamic called the **replicator dynamic** (Taylor and Jonker 1978).
- ▶ The replicators here are pure strategies that are copied without error from parent to child.
 - ▶ As the population state x changes, so do the payoffs and thereby the fitness of each strategy.
- ▶ The replicator dynamic is formalized as a (deterministic) system of ordinary differential equations without mutation.
- ▶ We will (at first) analyze the role of mutations indirectly by examining the local stability of various limit states of the replicator dynamic.

The Replicator Dynamic

- ▶ Let the rate of growth of strategy i be:

$$\frac{\dot{m}_i}{m_i} = [\beta - \delta + F_i(x)],$$

where β and δ are “background” birth and death rates (which are independent of payoffs).

- ▶ This is the interpretation of payoffs as fitness (reproduction rates) in biological models of evolution.

Derivation

What is the rate of growth of x_i , the *population share* of strategy i ?

By definition:

$$\begin{aligned}x_i &= \frac{m_i}{m} \\ \ln(x_i) &= \ln(m_i) - \ln(m) \\ \frac{\dot{x}_i}{x_i} &= \frac{\dot{m}_i}{m_i} - \frac{\dot{m}}{m} \\ &= [\beta - \delta + F_i(x)] - \sum_{j=1}^n x_j [\beta - \delta + F_j(x)] \\ &= F_i(x) - \bar{F}(x).\end{aligned}$$

That is, the growth rate of a strategy equals the excess of its payoff over the average payoff.

Some Properties of the Replicator Dynamic

The following results are immediate:

- ▶ Those subpopulations that are associated with better than average payoffs grow and *vice versa*.
- ▶ The subpopulations associated with pure best replies to the current population state $x \in X$ have the highest growth rate.
- ▶ $\dot{x}_i = x_i[F_i(x) - \bar{F}(x)]$, so that if $m_i = 0$ at T , then $m_i = 0$ for all $t > T$.

Relative Growth Rates

The ratio of any two population shares x_i and x_j increases (resp. decreases) over time if strategy i earns a higher (resp. lower) payoff than strategy j .

$$\begin{aligned}\frac{d}{dt} \left[\frac{x_i}{x_j} \right] &= \frac{\dot{x}_i x_j - x_i \dot{x}_j}{x_j x_j} \\ &= \frac{\dot{x}_i}{x_j} - \frac{\dot{x}_j x_i}{x_j x_j} \\ &= \frac{x_i}{x_j} \left[\frac{\dot{x}_i}{x_i} - \frac{\dot{x}_j}{x_j} \right] \\ &= \frac{x_i}{x_j} \left[F_i(x) - \bar{F}(x) - (F_j(x) - \bar{F}(x)) \right] \\ &= \frac{x_i}{x_j} \left[F_i(x) - F_j(x) \right].\end{aligned}$$

Invariance under Payoff Transformations

Suppose the payoff function $F_i(x)$ is replaced by a positive affine transformation:

$$G_i(x) = \alpha + \gamma F_i(x).$$

EXERCISE: Show that the replicator dynamic is invariant to such a change, modulo a change of timescale.

In particular, show that:

$$\frac{\dot{x}_i}{x_i} = \gamma[F_i(x) - \bar{F}(x)].$$

Example: Pure Coordination

Pure Coordination

	1	2
1	1 1	0 0
2	0 0	2 2

$$\frac{\dot{x}_1}{x_1} = (1 - x_1)(3x_1 - 2).$$

Therefore, $\frac{\dot{x}_1}{x_1} > 0$ iff $3x_1 > 2$ or $x_1 > \frac{2}{3}$.

Example: Impure Coordination

Stag Hunt

	1	2
1	2, 2	0, 0
2	0, 2	3, 3

$$\frac{\dot{x}_1}{x_1} = (1 - x_1)(3x_1 - 1).$$

Therefore, $\frac{\dot{x}_1}{x_1} > 0$ iff $3x_1 > 1$ or $x_1 > \frac{1}{3}$.

Example: Anti-Coordination

Hawk Dove

	1	2
1	-2	0
2	4	0

$$\frac{\dot{x}_1}{x_1} = (1 - x_1)(4 - 6x_1).$$

Therefore, $\frac{\dot{x}_1}{x_1} > 0$ iff $6x_1 < 4$ or $x_1 < \frac{2}{3}$.