Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 4:
Evolution in Games

The Evolutionary Approach

Evolutionary game theory is the study of:

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

Population Games

- Number of agents is large,
- Individual agents are small,
- Anonymous interaction,
- The number of ‘roles’ is finite,
  - each agent is a member of one of a finite number of populations.
  - members of a population have identical strategy sets and payoff functions.
- Payoffs are continuous in the population state (sometimes require continuous differentiability C^1).

Population Games

Players

- The population is a set of agents (possibly a continuum).

Strategies

- The set of (pure) strategies is $S = \{1, \ldots, n\}$, with typical members i, j and s.
- The mass of agents choosing 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 $i \in S$.

Population Games

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 $\mathbb{R}^n$:

$$e_1 = (1, 0, 0, \ldots), e_2 = (0, 1, 0, \ldots), e_3 = (0, 0, 1, \ldots), \quad \ldots$$

Payoffs

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

Equivalence to Mixed Strategies

Consider random matching to play a *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:

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

which depends *linearly* on the population state.

Note: There are many contexts in which agents' payoffs depend 'directly' on the strategies of all other players.

Nash Equilibria of Population Games

x^* is a Nash equilibrium of the population game if

$$(x^* - x)' F(x^*) \geq 0 \text{ for all } x \in X.$$

- **Monomorphic equilibria:** $x^* = e_i$.
- **Polymorphic equilibria:** $x^* \neq e_i$ for some $i \in S$; requires $F_i(x^*) = F_j(x^*) \geq F_k(x^*))$ for all i, j in support of x^* and k not in the support of x^*.

Theorem. Every population game with a continuum of agents admits at least one Nash equilibrium.

—Proved in the usual way using Kakutani's fixed point theorem.

Average Population Payoffs

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

$$\begin{aligned}\overline{F}(x) &= x_1 F_1(x) + x_2 F_2(x) \ldots + x_n F_n(x) \\ &= \sum_{i=1}^{n} x_i F_i(x).\end{aligned}$$

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

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 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.

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 in strategy i's population share x_i?

By definition:

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

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.
- Support invariance: $\dot{x}_i = x_i[F_i(x) - \overline{F}(x)]$, so that if $m_i = 0$ at T, then $m_i = 0$ for all 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.

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

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:

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

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	2, 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, -2	4, 0
2	0, 4	0, 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 $x_1 < \frac{2}{3}$.

Example: Prisoners' Dilemma

PD

	1	2
1	3, 3	0, 5
2	5, 0	1, 1

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

Therefore, $\frac{\dot{x}_1}{x_1} < 0$ for all $x_1 < 1$.

The Iterated Prisoners' Dilemma

- When matched, two players engage in a series of PD games.
- The engagement ends after the current round with probability $\delta < \frac{1}{2}$. We call this the *stopping probability*.
- Consider a population in which three strategies are present:
  - C—always cooperate,
  - D—always defect,
  - T—tit-for-tat, i.e. start by cooperating, thenceforth cooperate in period t if partner cooperated in $t - 1$.

Expected Payoffs Within Each Pairing

	C	D	T
C	$\frac{3}{\delta}$	0	$\frac{3}{\delta}$
D	$\frac{5}{\delta}$	$\frac{1}{\delta}$	$4 + \frac{1}{\delta}$
T	$\frac{3}{\delta}$	$\frac{1}{\delta} - 1$	$\frac{3}{\delta}$

Note:

Payoff from playing T against D is $0 + (1-\delta)\frac{1}{\delta} = \frac{1}{\delta} - 1$.

Payoff from playing D against T is $5 + (1-\delta)\frac{1}{\delta} = 4 + \frac{1}{\delta}$.

Expected Payoffs Over All Pairings

$$\begin{aligned}
F_C(x) &= (x_C + x_T)\tfrac{3}{\delta} \\
F_D(x) &= x_C\tfrac{5}{\delta} + x_D\tfrac{1}{\delta} + x_T\left(4 + \tfrac{1}{\delta}\right) \\
F_T(x) &= (x_C + x_T)\tfrac{3}{\delta} + x_D(\tfrac{1}{\delta} - 1)
\end{aligned}$$

Replicator Dynamics

$$\frac{d}{dt}\left[\frac{x_T}{x_C}\right] = \frac{x_T}{x_C}\big(F_T(x) - F_C(x)\big)$$
$$= \frac{x_T}{x_C}\left[x_D\left(\tfrac{1}{\delta} - 1\right)\right],$$

which is positive because $\delta < \frac{1}{2}$.

$$\frac{d}{dt}\left[\frac{x_T}{x_D}\right] = \frac{x_T}{x_D}\big(F_T(x) - F_D(x)\big)$$
$$= \frac{x_T}{x_D}\left[-x_C\tfrac{2}{\delta} - x_D + \underbrace{x_T\left(\tfrac{2}{\delta} - 4\right)}_{>0}\right],$$

which is positive for x_T sufficiently large.

Vector Field