Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 10.
Nonconvergence of Evolutionary Dynamics

Outline

- In this lecture, we shall further explore the conditions under which an evolutionary dynamic is locally stable.
- Once again we shall study the linear approximation to a (non-linear) evolutionary dynamic around a rest point x^*.
- We will be particularly interested in the conditions under which this linear dynamic does not converge to x^* from states in the neighborhood of x^*.
- This will essentially be a review of some standard results for systems of linear ordinary differential equations.

Outline

- We shall then:
  - Explore a method for studying the long-run behavior of nonconvergent dynamics,
  - Define a class of games in which nonconvergence is bound to occur, and
  - Introduce the concept of chaotic dynamics.

Linear Approximations around Rest Points

- The single-population dynamic $\dot{x} = V(x)$, which we shall refer to as (D), describes the evolution of the population state through the simplex X.

- Recall that near x^*, the dynamic (D) can typically be well approximated by the linear dynamic:

$$\dot{z} = DV(x^*)z, \qquad (L)$$

  in a neighborhood of the origin.

- Note that (L) is a dynamic on the tangent space TX.

Eigenvalues and Stability

- (L) approximates the motion of deviations from x^* following a small displacement z.
- To check for local stability of x^* under (D), we need to check whether the origin is stable under (L).
- The stability of the origin under (L) is completely determined by the eigenvalues of the Jacobian matrix $DV(x^*)$:
  - The origin is stable if all the eigenvalues have negative real part.
  - The origin is unstable if at least one eigenvalue has positive real part.
  - The origin is a saddle if some of each.

Eigenvalues and Stability

- Note that if $DV(x^*)$ is positive definite, i.e. $z'DV(x^*)z > 0$ for all nonzero $z \in TX$, then all eigenvalues have positive real part.

  —Therefore, x^* is unstable; all solutions that start near x^* are repelled.

- If $DV(x^*)$ is negative definite, i.e. $z'DV(x^*)z < 0$ for all nonzero $z \in TX$, then all eigenvalues have negative real part and the opposite is true.

Explicit Solutions and Stability

- To fully characterize the behavior of (L), we need to first reduce it to a simpler form.
- Consider an arbitrary matrix $A \in \mathbb{R}^{n \times n}$. We can find a matrix with the same eigenvalues as A, i.e. a matrix that is similar to A, but is easier to work with.
- The matrix A is **similar** to matrix $B \in \mathbb{R}^{n \times n}$ if there exists an invertible matrix $S \in \mathbb{C}^{n \times n}$, called a *similarity matrix*, such that:

$$B = S^{-1} A S.$$

Real Jordan Matrices

- We shall now find a simple class of matrices with the property that *every matrix is similar to a unique representative from this class*.

- A **real Jordan matrix** is block diagonal matrix whose diagonal blocks—*Jordan blocks*—are of the following four types:

$$J_1 = (\lambda); \qquad J_2 = \begin{pmatrix} a & b \\ -b & a \end{pmatrix};$$

$$J_3 = \begin{pmatrix} \lambda & 1 & 0 & 0 & 0 \\ 0 & \lambda & 1 & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & 0 & \lambda \end{pmatrix}; \quad J_4 = \begin{pmatrix} J_2 & I & 0 & 0 & 0 \\ 0 & J_2 & I & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & J_2 & I \\ 0 & 0 & 0 & 0 & J_2 \end{pmatrix}.$$

Real Jordan Matrices

Theorem 10.1. Every matrix $A \in \mathbb{R}^{n \times n}$ is similar to a real Jordan matrix $J = S^{-1}AS$. The latter is unique up to an ordering of the Jordan blocks.

- Each J_1 block corresponds to a real eigenvalue λ.
- Each J_2 block corresponds to a pair of complex eigenvalues $a \pm \mathrm{i}b$.
- Each J_3 block corresponds to a real eigenvalue with less than full geometric multiplicity (i.e. λ corresponds to at least two eigenvectors that are not linearly independent).
- Each J_4 block corresponds to a pair of complex eigenvalues with less than full geometric multiplicities.

Linear Dynamics on the Plane

There are three generic types of 2×2 matrices:

1. Diagonalizable matrices with two real eigenvalues—their real Jordan form is a diagonal matrix containing two J_1 blocks.
2. Diagonalizable matrices with two complex eigenvalues—their real Jordan form is a J_2 matrix.
3. Nondiagonalizable matrices with one real eigenvalue—their real Jordan form is a J_3 matrix.

Linear Dynamics on the Plane

1. When A has two real eigenvalues, $\dot{z} = Az$ and its solution from initial condition $z_0 = \xi$ are of the following form:

$$\dot{z} = Az = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}; \qquad z_t = \begin{pmatrix} \xi_1 e^{\lambda t} \\ \xi_2 e^{\mu t} \end{pmatrix}.$$

- If λ and μ are both negative, then the origin is a *stable node*,
- If both are positive, then the origin is an *unstable node*, and
- If the signs differ, then the origin is a *saddle*.

Linear Dynamics on the Plane

2. When A has two complex eigenvalues $a \pm \mathrm{i}b$, then:

$$\dot{z} = Az = \begin{pmatrix} a & b \\ -b & a \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}; \quad z_t = \begin{pmatrix} \xi_1 e^{at} \cos bt + \xi_2 e^{at} \sin bt \\ \xi_1 e^{at} \sin bt + \xi_2 e^{at} \cos bt \end{pmatrix}.$$

The stability of the origin is determined by the real part of the eigenvalues:

- If $a < 0$, then the origin is a *stable spiral*,
- If $a > 0$, then the origin is an *unstable spiral*, and
- If $a = 0$, then the origin is a *center*, with each solution following a closed orbit around the origin.

Linear Dynamics on the Plane

3. When A has lone eigenvalue λ:

$$\dot{z} = Az = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}; \qquad z_t = \begin{pmatrix} \xi_1 e^{\lambda t} + \xi_2 t e^{\lambda t} \\ \xi_2 e^{\lambda t} \end{pmatrix}.$$

The origin is:

- *stable* if $\lambda < 0$,
- *unstable* if $\lambda > 0$.

Characterizing Long-Run Behavior of Nonconvergent Dynamics

- This is often an impossible task.
- But as we have seen, the replicator dynamic in certain contexts has useful *conservative properties* which allow us to characterize its long-run behavior even when the dynamic does not converge.
- In particular, we have seen that in null stable games, all interior solutions of the replicator dynamic preserve the value of the strict Lyapunov function:

$$h_{x^*}(x) = \sum_{i \in S} x_i^* \log \frac{x_i^*}{x_i}.$$

Characterizing Long-Run Behavior of Nonconvergent Dynamics

- We know that standard RPS is a null stable game (good RPS is strictly stable and bad RPS is unstable).
- Let x^* be the unique Nash equilibrium $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.
- Then:

$$\begin{aligned} h_{x^*}(x) &= \sum_{i \in S} \tfrac{1}{3} \log \tfrac{1/3}{x_i} \\ &= \tfrac{1}{3} \sum_{i \in S} \left[\log(1/3) - \log(x_i) \right] \\ &= \log(1/3) - \tfrac{1}{3} \sum_{i \in S} \log(x_i) \\ &= \log(1/3) - \tfrac{1}{3} \log\left(x_1 x_2 x_3\right). \end{aligned} \tag{1}$$

Characterizing Long-Run Behavior of Nonconvergent Dynamics

- Therefore, if every solution trajectory preserves $h_{x^*}(x)$, then it preserves $x_1x_2x_3$ (an affine transformation of $h_{x^*}(x)$).

- That is, the level sets of $x_1x_2x_3$ form closed orbits around x^*.

Convergence of Time Averages

- Even if the process itself does not converge, the average population share over time for each strategy $i \in S$ could converge to its Nash equilibrium share.

- Let the average value of the state over the time interval $[0, t]$ be:

$$\bar{x}_t = \frac{1}{t} \int_0^t x_s ds.$$

- One can show that in standard RPS under the replicator dynamic, $\{\bar{x}_t\}_{t \geq 0}$ converges to the Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ as t approaches infinity:

$$\lim_{t \to \infty} |\bar{x}_t - x^*| = 0.$$

Games with Nonconvergent Dynamics

- We have focussed on RPS in much of our discussion so far, but we can generalize these insights to a broader class of games in which convergence can fail, called **circulant games** of which RPS is a member.

- The matrix $A \in \mathbb{R}^{n \times n}$ is called a *circulant matrix* if it is of the form:

$$A = \begin{pmatrix} a_1 & a_2 & \cdots & a_{n-1} & a_n \\ a_n & a_1 & a_2 & \dots & a_{n-1} \\ \ddots & \ddots & \ddots & \ddots & \ddots \\ a_3 & \cdots & a_n & a_1 & a_2 \\ a_2 & a_3 & \cdots & a_n & a_1 \end{pmatrix}.$$

Circulant Games

- When A is a payoff matrix for a symmetric normal form game, then A is a *circulant game*.
- The barycenter $x^* = \frac{1}{n}\mathbf{1}$ is always in the set of Nash equilibria of such games.
- RPS is a circulant game with $n = 3$, $a_1 = 0$, $a_2 = -\ell$, and $a_3 = w$.

Chaotic Evolutionary Dynamics

- The ω limit sets we have focused on are fairly simple, mainly rest points and closed orbits of a dynamic.
- In one-dimensional systems, all continuous-time dynamics converge to equilibrium.
- In two dimensional systems rest points, closed orbits, chains of rest points and connecting orbits exhaust the possibilities.

Chaotic Evolutionary Dynamics

- For flows in three or more dimensions, however, ω-limit sets can be complicated sets known as **chaotic (or strange) attractors**.

- In addition, chaotic dynamics are defined by sensitive dependence on initial conditions:

  —solution trajectories starting from nearby points on the attractor move apart at an exponential rate.