Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 6:
Deterministic Dynamics

Best Responses

- The **best response dynamic** is generated by agents always switching to their current best response.
- The set of rest points of the dynamic (where $\dot{x} = \mathbf{0}$) are the set of Nash equilibria:

$$NE(F) = \{x \in X : x_i > 0 \implies F_i(x) = \max_{j \in S} F_j(x)\}.$$

- The dynamic has some peculiar features: the best response correspondence is discontinuous (small changes in the state x can produce sharp changes in responses) and multivalued (there could be multiple best responses to a state).
- Differential inclusions—set-valued differential equations—are used.

The Best Response Dynamic

- The best response dynamic is given by the following differential inclusion:

$$\dot{x} \in M\big(F(x)\big) - x. \tag{1}$$

Definition. A Carathéodory solution to the differential inclusion $\dot{x} \in V(x)$ is a Lipschitz continuous trajectory $\{x_t\}_{t \geq 0}$ that satisfies $\dot{x}_t \in V(x_t)$ at all but a measure zero set of times in $[0, \infty)$.

Theorem. Fix a continuous population game F. Then for each $\xi \in X$, there exists a trajectory $\{x_t\}_{t \geq 0}$ with $x_0 = \xi$ that is a Carathéodory solution to the differential inclusion (1).

Solution Trajectories

- As we shall see, while solutions to the best response dynamic exist, the best response protocol is discontinuous so the solutions *need not be unique*; multiple solution trajectories can emanate from a single initial condition.

- Yet they can be quite simple.

- Let $\{x_t\}$ be a solution to (1) and suppose that the best response to state x_t is the pure strategy $i \in S$ at all times $t \in [0, T]$.

- Then during this interval, evolution is described by the affine differential equation:

$$\dot{x} = e_i - x.$$

Solution Trajectories

- Hence the state x moves directly toward vertex e_i of the set X, proceeding more slowly as the vertex is approached.
- This means that the state x_t lies on the segment containing x_0 and e_i throughout the interval $[0, T]$.
- Solving $\dot{x} = e_i - x$ we get the following explicit formula for x_t:

$$x_t = (1 - \mathsf{exp}^{-t})e_i + \mathsf{exp}^{-t} x_0 \qquad \text{for all } t \in [0, T].$$

Examples

Two-Strategy Coordination

- Let the strategy set be $S = \{U, D\}$ and the payoff matrix be:

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

- The game $F(x) = Ax$ has three Nash equilibria, two pure (e_U and e_D) and a mixed equilibrium $(x_U^*, x_D^*) = (\frac{2}{3}, \frac{1}{3})$.

Examples

Denote the state by $\chi = x_D$, so that $\chi^* = \frac{1}{3}$.

Then the best response-dynamic can be expressed as:

Examples

- From every initial condition except χ^*, there is a unique solution trajectory of the dynamic that converges to a pure Nash equilibrium:

$$\chi_0 < \chi^* \implies \chi_t = e^{-t}\chi_0. \tag{2}$$

$$\chi_0 > \chi^* \implies \chi_t = 1 - e^{-t}(1 - \chi_0). \tag{3}$$

Examples

- There are many solution trajectories from χ^*:
  - a stationary trajectory,
  - one that proceeds to $\chi = 0$ according to (2),
  - another that proceeds to $\chi = 1$ according to (3).
- Notice that solutions (2) and (3) quickly leave the vicinity of χ^*.
- In contrast, for Lipschitz continuous dynamics:
  1. solutions from all initial conditions are unique,
  2. solutions that start near a stationary point move very slowly near that point.

Imitative Dynamics

Recall that imitative dynamics are based on learning protocols of the form:

$$\rho_{ij}(\pi, x) = x_j r_{ij}(\pi, x),$$

where r_{ij} is a conditional imitation rate.

These generate a mean dynamic of the form:

$$\begin{aligned}
\dot{x}_i &= \sum_{j \in S} x_j \rho_{ji}\big(F(x), x\big) - x_i \sum_{j \in S} \rho_{ij}\big(F(x), x\big) \\
&= \sum_{j \in S} x_j x_i r_{ji}\big(F(x), x\big) - x_i \sum_{j \in S} x_j r_{ij}\big(F(x), x\big) \\
&= x_i \sum_{j \in S} x_j \big[r_{ji}\big(F(x), x\big) - r_{ij}\big(F(x), x\big)\big].
\end{aligned}$$

Imitative Dynamics

Definition. Suppose that the conditional imitation rates are Lipschitz continuous and that net conditional imitation rates are *monotone*, that is:

$$\pi_j \geq \pi_i \Leftrightarrow r_{kj}(\pi, x) - r_{jk}(\pi, x) \geq r_{ki}(\pi, x) - r_{ik}(\pi, x)$$

for all $i, j, k \in S$.

Then the mean dynamic above is called an **imitative dynamic**.

Properties of Imitative Dynamics

- All imitative dynamics satisfy ***support invariance***.
- Since imitative dynamics are Lipschitz continuous, we know they also exhibit ***uniqueness*** and ***forward and backward invariance***:

Proposition. For every initial condition $\xi \in X$, an imitative dynamic admits a unique solution trajectory $\mathcal{T}_{(-\infty,\infty)} = \{x : (-\infty, \infty) \to X |\, x \text{ is continuous}\}$.

Example

Pure Coordination

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

The **replicator dynamic** is:

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

No closed-form solution to the initial value problem, but we can immediately deduce:

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

Rest Points of Imitative Dynamics

- Recall that:

$$NE(F) = \{x \in X : x_i > 0 \implies F_i(x) = \max_{j \in S} F_j(x)\}.$$

- The set of rest points of an imitative dynamic is the set of **restricted equilibria**:

$$RE(F) = \{x \in X : x_i > 0 \implies F_i(x) = \max_{j \in S : x_j > 0} F_j(x)\}.$$

These are the Nash equilibria of a restricted version of F in which only strategies in the support of x can be played.

Limit Sets

- More generally, the limiting behavior of deterministic dynamics can be characterized as follows.

- The ω-limit of trajectory $\{x_t\}_{t\geq 0}$ is the set of all points that the trajectory approaches arbitrarily closely infinitely often:

$$\omega(\{x_t\}) = \Big\{ y \in X : \text{there exists } \{t_k\}_{k=1}^{\infty} \text{ with } \lim_{k\to\infty} t_k = \infty \text{ such that } \lim_{k\to\infty} x_{t_k} = y \Big\}.$$

- If $\omega(\{x_t\}) = x^*$, a singleton, then x^* is a called an **absorbing state** (rest point).

- More generally, $\omega(\{x_t\})$ is called a **recurrence class** or **ω-limit set** of the dynamic.

Limit Sets

- For dynamics that admit a unique forward solution trajectory from each initial condition, $\omega(\xi)$ denotes the ω-limit set of the trajectory starting from state ξ.

- The set of all ω-limit points of all solution trajectories is:

$$\Omega(V_F) = \bigcup_{\xi \in X} \omega(\xi).$$

- The notion of recurrence (or the set of recurrence classes) of a deterministic dynamic is captured by $\Omega(V_F)$.

Analyzing Convergence

Example. **Random Matching in Rock-Paper-Scissors**

Suppose a win is worth $w > 0$, a loss $-l < 0$ and a draw 0. Then for $F(x) = Ax$:

$$A = \begin{pmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{pmatrix}.$$

When $w = l$, let us call the game (standard) *RPS*. When $w > l$, we call it *good RPS* and when $w < l$ we call it *bad RPS*.

For all cases, the unique Nash equilibrium of A is $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$, an interior equilibrium.

Best Response Dynamic in standard RPS

- One can construct a figure which appears to indicate that every solution trajectory converges to the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

- To prove this, we claim that along every solution trajectory $\{x_t\}$, whenever the best response is unique, we have:

$$\frac{d}{dt}\left(\max_{k \in S} F_k(x_t)\right) = -\max_{k \in S} F_k(x_t). \tag{4}$$

Note: Best response correspondence is not continuous, but max. payoff is continuous.

Example

- To establish the claim, let x_t be a state in which there is a unique optimal strategy, say Paper.
- At this state $\dot{x}_t = e_P - x_t$. Since $F_P(x) = w(x_R - x_S)$:

$$\begin{aligned} \frac{d}{dt} F_P(x_t) &= \nabla F_P(x_t)' \dot{x}_t \\ &= (w \quad 0 \quad -w)(e_P - x_t) \\ &= -w(x_R - x_S) \\ &= -F_P(x_t). \end{aligned} \tag{5}$$

Example

- Because any solution trajectory passes through states with multiple best responses at most a countable number of times (see Figure), (4) can be integrated with respect to time.

- This yields:

$$\max_{k \in S} F_k(x_t) = e^{-t} \max_{k \in S} F_k(x_0). \tag{6}$$

- In standard RPS, payoffs to each strategy are non-negative and equal zero only at the Nash equilibrium x^*.

- Then (6) implies that the maximal payoff across strategies $k \in S$ falls over time converging to zero as t approaches infinity; this occurs as x_t converges to the Nash equilibrium x^*.

Lyapunov Functions

The most common method for proving global convergence in dynamical systems is by constructing a **strict Lyapunov function**:

- A scalar-valued function.
- The value of the function changes monotonically along every solution trajectory.
- No general method of constructing Lyapunov function, but know for many games + dynamics.
- The Lyapunov function allows us to (partially) characterize the evolution of play without requiring explicit solutions to the differential equation (or inclusion).

Definition. The C^1 function $L : X \to \mathbb{R}$ is a (decreasing) strict Lyapunov function for the differential equation $\dot{x} = V_F(x)$ if along any solution trajectory $\dot{L}(x) \leq 0$, with equality only at rest points of V_F.

Stable Games

RPS (standard or good) is an example of a stable game.

Definition. The population game $F : X \to \mathbb{R}^n$ is a **stable game** if:

$$(y - x)'\big(F(y) - F(x)\big) \leq 0 \qquad \text{for all } x, y \in X \tag{7}$$

If the inequality in condition (7) holds strictly whenever $x \neq y$, F is a strictly stable game (e.g. good RPS), whereas if this inequality always binds, F is a null stable game (e.g. standard RPS).

Proposition. If F is a null stable game, $NE(F)$ is a convex set. If F is a strictly stable game, $NE(F)$ is a singleton.

Best Response Dynamics in Stable Games

Theorem. Let F be a C^1 stable game, and let $\dot{x} \in V_F(x)$ be the best response dynamic for F. Define the Lipschitz continuous function $G : X \to \mathbb{R}_+$ by:

$$G(x) = max_{i \in S}[F_i(x) - \overline{F}(x)],$$

which is non-negative and satisfies $G^{-1}(0) = NE(F)$.

Moreover, if $\{x_t\}_{t \geq 0}$ is a solution to V_F then $\dot{G}(x_t) \leq -G(x_t)$ for almost all $t \geq 0$.

Hence G is a Lyapunov function for $V_F(x)$ and $\lim_{t \to \infty} x_t \in NE(F)$.

Replicator Dynamic in Stable Games

- For convergence of the replicator dynamics, we need to restrict attention to a subset of all initial conditions $\xi \in X$, because if ξ places no mass on a strategy in the support of a Nash equilibrium x^*, then the dynamic cannot converge to x^* from ξ.

- Let the support of x be $S(x) = \{i \in S : x_i > 0\}$. Then $X_y = \{x \in X : S(y) \subseteq S(x)\}$ is the set of states in X whose supports contain the support of y.

Replicator Dynamics in Stable Games

- The Lyapunov function (in the single-population case) is $h_y : X_y \to \mathbb{R}$ where:

$$h_y(x) = \sum_{i \in S(y)} y_i \log \frac{y_i}{x_i}.$$

h_y is known as the relative entropy of y given x.

Replicator Dynamics in Stable Games

Theorem. Let F be a strictly stable game with unique Nash equilibrium x^*, and let $\dot{x} = V_F(x)$ be the replicator dynamic for F.

Then h_{x^*} is non-negative and $\dot{h}_{x^*}(x) \leq 0$, with equality only when $x = x^*$.

Therefore, $\lim_{t \to \infty} x_t = x^*$.

Characterizing Long-Run Behavior of Nonconvergent Dynamics

- This is often an impossible task.
- But 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, 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:

$$\begin{aligned} h_{x^*}(x) &= \sum_{i \in S} \frac{1}{3} \log \frac{1/3}{x_i} \\ &= \frac{1}{3} \sum_{i \in S} \left[\log(1/3) - \log(x_i) \right] \\ &= \log(1/3) - \frac{1}{3} \sum_{i \in S} \log(x_i) \\ &= \log(1/3) - \frac{1}{3} \log \left(x_1 x_2 x_3 \right). \end{aligned} \tag{8}$$

Characterizing Long-Run Behavior of Nonconvergent Dynamics

- Therefore, if every solution trajectory preserves $h_{x^*}(x)$, then it preserves $x_1 x_2 x_3$ (an affine transformation of $h_{x^*}(x)$).
- Hence all interior trajectories preserve volume, $x_1 x_2 x_3$.
- That is, the level sets of $x_1 x_2 x_3$ 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.