1. Population games & evolutionary dynamics

Jonathan Newton, Kyoto University

Population games

Consider a situation in which

- There are many agents.
- There is no assumption of equilibrium.
- Agents use some behavioral rule in order to choose their strategies.
- The share of agents playing each strategy changes over time.

Population games

Define a population game

1. Continuum of mass 1 of agents.
2. Strategies $S = \{1, \ldots, n\}$.
3. Set of population states

$$X = \left\{ x \in \mathbb{R}^n : \sum_{i \in S} x_i = 1 \right\}.$$

4. Payoff function $\pi_i : X \to \mathbb{R}$ for each strategy $i \in S$.

Best response dynamics

The dynamics that immediately come to mind when we think of Nash equilibrium.

Best response dynamic

- Let time be continuous, $t \geq 0$.
- Let any given agent become active and update his strategy at a given Poisson rate (to avoid notation, let this equal 1).
- When an agent updates, he chooses a best response to the current population state.
- Strategy i is a best response to x if i solves $\max_{j \in S} \pi_j(x)$.

Best response dynamics

- Let $BR(x)$ be the set of best responses (including mixed best responses) to population state (mixed strategy) x,

$$BR(x) = \left\{ w \in X : w_i > 0 \implies i \in \max_{j \in S} \pi_j(x) \right\}.$$

Best response dynamics

- Let $BR(x)$ be the set of best responses (including mixed best responses) to population state (mixed strategy) x,

$$BR(x) = \left\{ w \in X : w_i > 0 \implies i \in \max_{j \in S} \pi_j(x) \right\}.$$

- Best response dynamics (BRD) satisfy

$$\dot{x} \in BR(x) - x \tag{1}$$

where $\dot{x}$ is the time derivative of x.

Best response dynamics

- Let $BR(x)$ be the set of best responses (including mixed best responses) to population state (mixed strategy) x,

$$BR(x) = \left\{ w \in X : w_i > 0 \implies i \in \max_{j \in S} \pi_j(x) \right\}.$$

- Best response dynamics (BRD) satisfy

$$\dot{x} \in BR(x) - x \tag{1}$$

  where $\dot{x}$ is the time derivative of x.
- Under a BRD, trajectory $\{x(t)\}_{t \geq 0}$ satisfies (1).

Best response dynamic

Coordination game

- Let payoffs be given by

$$\pi(x) = \Pi\, x, \qquad \Pi = \begin{pmatrix} 4 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

- That is, a coordination game

$$\pi_1(x) = 4\,x_1$$
$$\pi_2(x) = 3\,x_2$$
$$\pi_3(x) = 2\,x_3$$

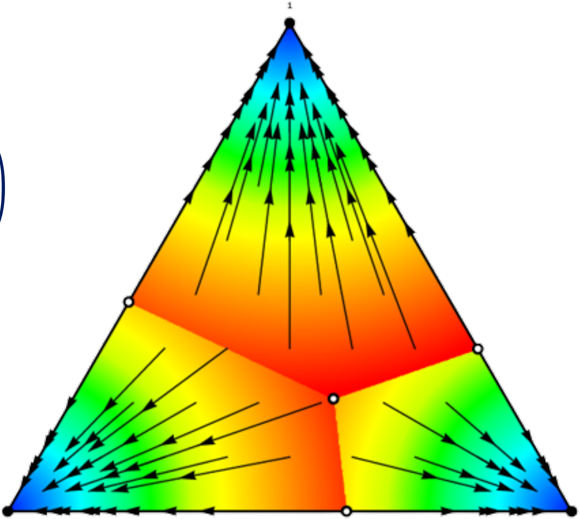

Best response dynamic

Coordination game

- Diagram illustrates trajectory of BRD.
- Colour indicates how fast or slow it is moving.
- Stable rest points (filled circles) correspond to strict Nash equilibria.
- Unstable rest points (unfilled circles).

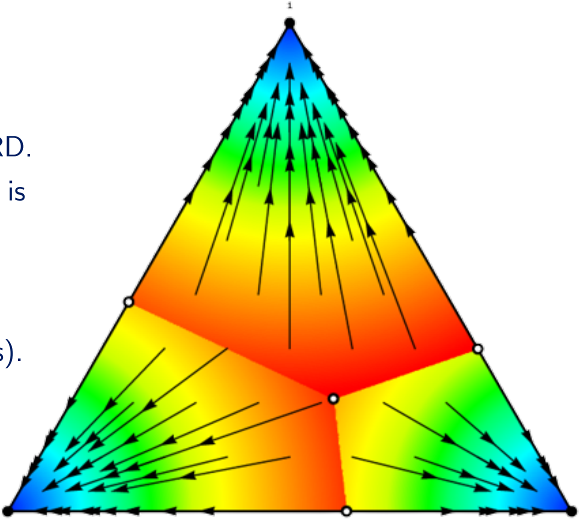

Similar to best response

BNN, Smith, Sample BRD

Average payoff at x is

$$\bar{\pi}(x) = \sum_{i \in S} x_i \, \pi_i(x)$$

- Agent with current strategy i chooses a strategy j at random.
  - Brown-von Neumann-Nash dynamic. If $\pi_j(x) > \bar{\pi}(x)$, then switch to j with probability proportional to $\pi_j(x) - \bar{\pi}(x)$.
  - Smith dynamic. If $\pi_j(x) > \pi_i(x)$, then switch to j with probability proportional to $\pi_j(x) - \pi_i(x)$.
- Sample BRD. Agent samples k actions from the population and best responds to the distribution given by the sample.

Logit dynamic

Noisy best response

- Logit dynamic satisfies

$$\dot{x} = M(x) - x, \quad \text{for} \quad M_i(x) = \frac{e^{\frac{1}{\eta}\pi_i(x)}}{\sum_{j \in S} e^{\frac{1}{\eta}\pi_j(x)}}$$

- As $\eta \to 0$, approaches a BRD.
- As $\eta \to \infty$, approaches uniform random choice.

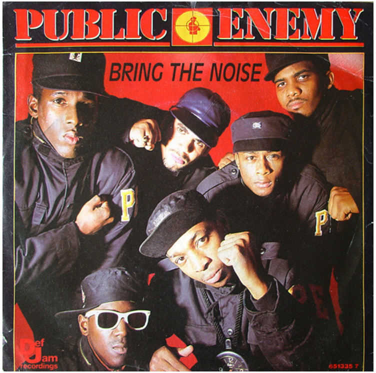

Logit dynamic

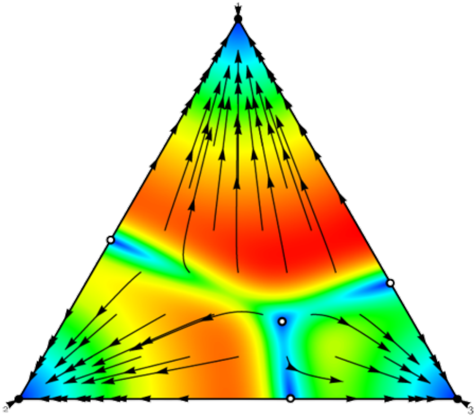

$\eta = 0.2$

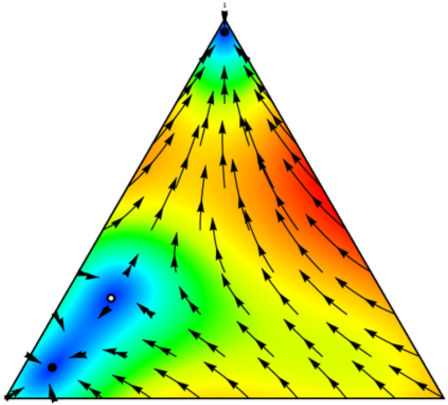

$\eta = 0.95$

Logit dynamic

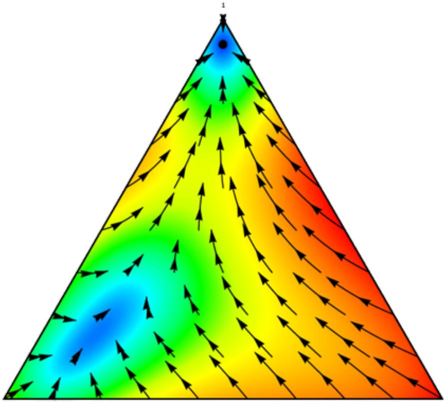

$\eta = 1.1$

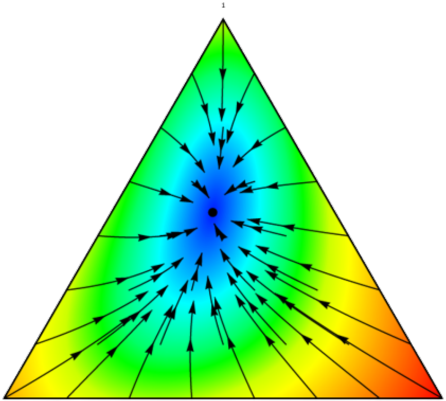

$\eta = 2$

Riemannian game dynamics

A class of dynamics that includes the replicator dynamic.

Riemannian game dynamics
Definition

- Tangent space (possible directions) $\mathbb{R}^S_0 := \{z \in \mathbb{R}^S_+ : \sum_{i \in S} z_i = 0\}$.

Riemannian game dynamics

Definition

- Tangent space (possible directions) $\mathbb{R}_0^S := \{z \in \mathbb{R}_+^S : \sum_{i \in S} z_i = 0\}$.
- Continuously differentiable cost of motion

$$C : \mathbb{R}_0^S \times X^{int} \to (0, \infty)$$

such that $C(\cdot, x)$ is a $+$ve definite quadratic form.

Riemannian game dynamics

Definition

- Tangent space (possible directions) $\mathbb{R}^S_0 := \{z \in \mathbb{R}^S_+ : \sum_{i \in S} z_i = 0\}$.
- Continuously differentiable cost of motion

$$C : \mathbb{R}^S_0 \times X^{int} \to (0, \infty)$$

  such that $C(\cdot, x)$ is a +ve definite quadratic form.
- Riemannian game dynamics defined by

$$\dot{x} = \arg\max_{z \in \mathbb{R}^S_0} \left(\sum_{i \in S} \pi_i(x)\, z_i - C(z, x) \right).$$

Riemannian game dynamics

Definition

- Tangent space (possible directions) $\mathbb{R}^S_0 := \{z \in \mathbb{R}^S_+ : \sum_{i \in S} z_i = 0\}$.
- Continuously differentiable cost of motion

$$C : \mathbb{R}^S_0 \times X^{int} \to (0, \infty)$$

  such that $C(\cdot, x)$ is a $+$ve definite quadratic form.
- Riemannian game dynamics defined by

$$\dot{x} = \arg\max_{z \in \mathbb{R}^S_0} \left(\sum_{i \in S} \pi_i(x)\, z_i - C(z, x) \right).$$

- $C(z, x)$ corresponds to a Riemannian metric on X, hence the name.

Riemannian game dynamics

Replicator dynamic

- Letting $C(z, x) = \frac{1}{2} \sum_{i \in S} \frac{z_i^2}{x_i}$,
- Solving we obtain the replicator dynamics

$$\dot{x}_i = x_i \Big(\pi_i(x) - \bar{\pi}(x) \Big).$$

- Difference between growth rates $\frac{\dot{x}_i}{x_i}$ and $\frac{\dot{x}_j}{x_j}$ is proportional to payoff difference $\pi_i(x) - \pi_j(x)$.

Positive correlation (PC)

All dynamics so far have $\dot{x}$ as a function of x,

$$\dot{x} = V_\pi(x).$$

Definition (Positive correlation)

$V_\pi(x) \neq 0$ implies that $V_\pi(x)' \pi(x) > 0$.

- PC requires that whenever a population is not at rest, the covariance between strategies' growth rates and payoffs is positive.
- Satisfied by Riemannian (including replicator), BRD, BNN, Smith dynamics.

Lyapunov & potential functions

Important methods for analyzing convergence.

Lyapunov functions

- A common way to analyze a dynamic is to find a Lyapunov function.
- Such functions are monotonic along solution trajectories.
- Attain maxima (or minima) at rest points of the dynamic.

Definition (Lyapunov function)

A continuously differentiable function $L : X \to \mathbb{R}$ is an (increasing) strict Lyapunov function for $\dot{x} = V_\pi(x)$ if $\dot{L}(x) \equiv \nabla L(x)' V_\pi(x) \geq 0$ for all $x \in X$, with equality only at rest points of V_π.

Potential games

Define Φ to normalize vectors while retaining information on differences.

$$\Phi = \begin{pmatrix} 1-\frac{1}{n} & -\frac{1}{n} & \cdots & -\frac{1}{n} \\ -\frac{1}{n} & 1-\frac{1}{n} & \cdots & -\frac{1}{n} \\ \vdots & \ddots & \ddots & \ddots \\ -\frac{1}{n} & -\frac{1}{n} & \cdots & 1-\frac{1}{n} \end{pmatrix}$$

Definition (Potential game)

Let $\pi : X \to \mathbb{R}^n$ be a population game. π is a potential game if it admits a continuously differentiable potential function $f : X \to \mathbb{R}$ such that

$$\nabla f(x) = \Phi\, \pi(x) \quad \text{for all } x \in X.$$

Evolution & Social Science, Brisbane 2023

Climbing potential

Examples of potential games:

- Normal form games composed of a common interest term plus an externality term.
- Cournot competition (f equals total surplus).
- Congestion games.

Theorem (Climbing potential)

Let π be a potential game with potential function f, and let $\dot{x} = V_\pi(x)$ satisfy (PC). Then f is a strict Lyapunov function for V_π.

PC and potential

Replicator for coordination (potential game) and Rock-Paper-Scissors (not a potential game)

$$\Pi = \begin{pmatrix} 4 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

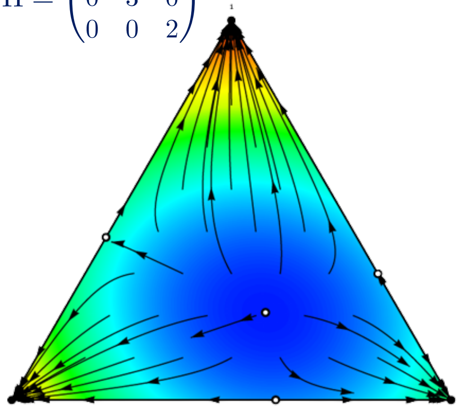

$$L(x) = f(x) = \frac{1}{2}\bar{\pi}(x)$$

$$\Pi = \begin{pmatrix} 0 & -3 & 1 \\ 1 & 0 & -3 \\ -3 & 1 & 0 \end{pmatrix}$$

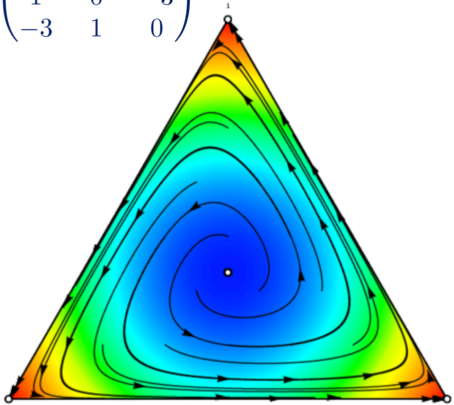

Instead use $L(x) = -x_1 x_2 x_3$

Logit does not satisfy PC

However, by adding an entropy term to potential, we get a Lyapunov function

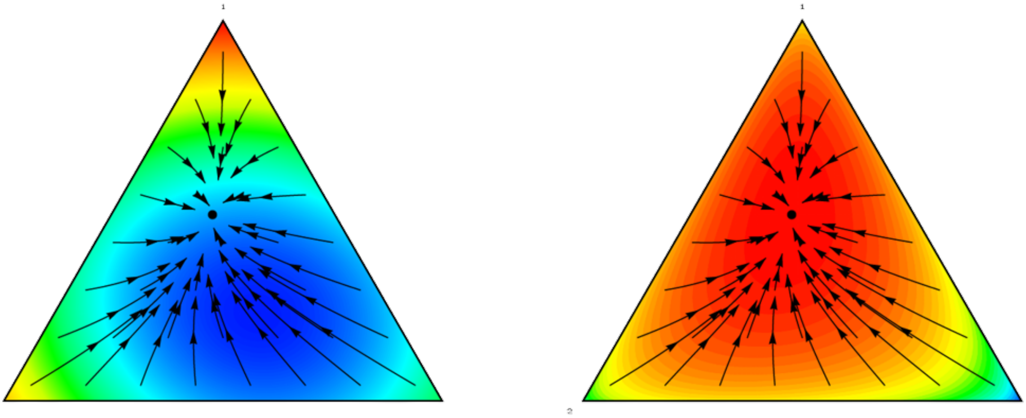

$L(x) \neq f(x) = \frac{1}{2}\bar{\pi}(x)$

$L(x) = f(x) - \eta \sum_{i \in S} x_i \log x_i$

Reinforcement, synthesis & higher order dynamics

Some fascinating connections.

Reinforcement learning

- For each $i \in S$, score $y_i(t)$ measures the performance of strategy i up to time t.

Evolution & Social Science, Brisbane 2023

Reinforcement learning

- For each $i \in S$, score $y_i(t)$ measures the performance of strategy i up to time t.
- Strategies updated via logit, but with scores (instead of payoffs),

$$x_i(t) = \frac{e^{\frac{1}{\eta} y_i(t)}}{\sum_{j \in S} e^{\frac{1}{\eta} y_j(t)}} \qquad \left[\text{all agents simultaneously}\right]$$

Reinforcement learning

- For each $i \in S$, score $y_i(t)$ measures the performance of strategy i up to time t.
- Strategies updated via logit, but with scores (instead of payoffs),

$$x_i(t) = \frac{e^{\frac{1}{\eta} y_i(t)}}{\sum_{j \in S} e^{\frac{1}{\eta} y_j(t)}} \qquad \text{[all agents simultaneously]}$$

- n-th order dynamics defined by how the score is updated according to payoffs.

$$\frac{d^n}{dt^n} y_i(t) = \pi_i\big(x(t)\big)$$

Reinforcement to replicator

- Taking logs

$$\log x_i(t) - \log x_j(t) = \frac{1}{\eta}\big(y_i(t) - y_j(t)\big)$$

- Differentiating

$$\frac{\dot{x}_i}{x_i} - \frac{\dot{x}_j}{x_j} = \frac{1}{\eta}\left(\frac{d}{dt}y_i(t) - \frac{d}{dt}y_j(t)\right)$$

- So, $\dot{y}_i = \pi_i(x)$ (i.e. $n = 1$) gives the replicator dynamic (!)

$$\frac{\dot{x}_i}{x_i} - \frac{\dot{x}_j}{x_j} = \frac{1}{\eta}\Big(\pi_i\big(x(t)\big) - \pi_j\big(x(t)\big)\Big)$$

Reinforcement to replicator

- For general n,

$$\frac{d^n}{dt^n} x_i = x_i \frac{1}{\eta} \Big(\pi_i(x(t)) - \bar{\pi}(x(t)) \Big)$$
$$+ \text{ terms independent of payoffs}$$

- For example, when $n = 2$, the acceleration of the dynamic is proportional to payoffs.
- By building n-th order replicator dynamics via reinforcement and logit, we ensure that x remains within X.

Wrapping up

- We have taken a quick look at some dynamics and methods that can be used to study population games.
- There are other dynamics that we have not had time to address here.
- For example, projection dynamics, tempered best response dynamics, imitation via pairwise comparison, best experienced payoff dynamics, completely uncoupled dynamics, regret testing, trial and error learning.
- There are also other classes of games for which there exist general results, for example contractive games and supermodular games.

newton@kier.kyoto-u.ac.jp

For references, see reading list.