RIEMANNIAN GAME DYNAMICS

P. Mertikopoulos[1] W. H. Sandholm[2]

[1]CNRS – Laboratoire d'Informatique de Grenoble

[2]University of Wisconsin–Madison

GEL 2016 – Bagnères-de-Luchon, January 5, 2016

Outline

Game dynamics

Rationality concepts can often be justified as a the long-term outcome of simple dynamic processes (game dynamics). However, your mileage may vary:

- Replicator dynamics (Taylor and Jonker, 1978)
  - Positive correlation between payoffs and dynamics
  - Nash equilibria are fixed points – but so is every pure state
  - Dominated strategies are eliminated
  - Periodic orbits in 2-player zero-sum games
- Projection dynamics (Nagurney and Zhang, 1997)
  - Positive correlation between payoffs and dynamics
  - Nash $\Leftrightarrow$ fixed points
  - Dominated strategies may survive
  - Limit cycles and periodic orbits in 2-player zero-sum games

Game dynamics

Rationality concepts can often be justified as a the long-term outcome of simple dynamic processes (game dynamics). However, your mileage may vary:

- Replicator dynamics (Taylor and Jonker, 1978)
  - Positive correlation between payoffs and dynamics ✓
  - Nash equilibria are fixed points – but so is every pure state ✗
  - Dominated strategies are eliminated ✗
  - Periodic orbits in 2-player zero-sum games ✓✗
- Projection dynamics (Nagurney and Zhang, 1997)
  - Positive correlation between payoffs and dynamics ✓
  - Nash $\Leftrightarrow$ fixed points ✗
  - Dominated strategies may survive ✗
  - Limit cycles and periodic orbits in 2-player zero-sum games ✓✗

Game dynamics

Rationality concepts can often be justified as a the long-term outcome of simple dynamic processes (game dynamics). However, your mileage may vary:

- Replicator dynamics (Taylor and Jonker, 1978)
  - Positive correlation between payoffs and dynamics ✓
  - Nash equilibria are fixed points – but so is every pure state ✗
  - Dominated strategies are eliminated ✗
  - Periodic orbits in 2-player zero-sum games ✓✗
- Projection dynamics (Nagurney and Zhang, 1997)
  - Positive correlation between payoffs and dynamics ✓
  - Nash $\Leftrightarrow$ fixed points ✗
  - Dominated strategies may survive ✗
  - Limit cycles and periodic orbits in 2-player zero-sum games ✓✗

What this talk is about:

Provide a unifying framework for the foregoing dynamics and extend a range of rationality properties known for certain special cases.

Population games

Throughout this talk we will focus on *population games:*

- Continuum of *nonatomic players*
- Common (finite) set of *actions* $\mathcal{A}$
- *Population states* $x = (x_\alpha)_{\alpha \in \mathcal{A}}$ are defined by the proportions $x_\alpha \in [0,1]$ of players playing each action $\alpha \in \mathcal{A}$.
- Set of population states: $\mathcal{X} = \Delta(\mathcal{A})$.
- Individual action rewards determined by associated *payoff functions* $v_\alpha \colon \mathcal{X} \to \mathbb{R}$

Example (Random matching)

- A population of agents is randomly matched to play a 2-player symmetric normal form game with payoff matrix $A = (A_{\alpha\beta})_{\alpha,\beta=1}^{n}$.
- The payoff to α-strategists when the population state is $x \in \mathcal{X}$ is

$$v_\alpha(x) = \sum\nolimits_{\beta \in \mathcal{A}} A_{\alpha\beta} x_\beta$$

Evolutionary dynamics

Evolutionary dynamics are rules that assign to games a dynamical system of the form:

$$\dot{x} = V(x) \tag{ED}$$

where $V(x)$ is usually specified as a function of state/payoff pairs: $V(x) \equiv \tilde{V}(x, v(x))$.

Evolutionary dynamics

Evolutionary dynamics are rules that assign to games a dynamical system of the form:

$$\dot{x} = V(x) \tag{ED}$$

where $V(x)$ is usually specified as a function of state/payoff pairs: $V(x) \equiv \tilde{V}(x, v(x))$.

Evolutionary dynamics are *admissible* if and only if $x(t) \in \mathcal{X}$ for all time $\implies$ $V(x)$ must lie in the *tangent cone* of $\mathcal{X}$ at x:

$$V(x) \in \mathrm{TC}_{\mathcal{X}}(x) = \left\{ z \in \mathbb{R}^{\mathcal{A}} : \sum_\alpha z_\alpha = 0 \text{ and } z_\alpha \geq 0 \text{ if } x_\alpha = 0 \right\}$$

Evolutionary dynamics

Evolutionary dynamics are rules that assign to games a dynamical system of the form:

$$\dot{x} = V(x) \tag{ED}$$

where $V(x)$ is usually specified as a function of state/payoff pairs: $V(x) \equiv \tilde{V}(x, v(x))$.

Evolutionary dynamics are *admissible* if and only if $x(t) \in \mathcal{X}$ for all time $\implies$ $V(x)$ must lie in the *tangent cone* of $\mathcal{X}$ at x:

$$V(x) \in \mathrm{TC}_{\mathcal{X}}(x) = \left\{z \in \mathbb{R}^{\mathcal{A}} : \sum_{\alpha} z_\alpha = 0 \text{ and } z_\alpha \geq 0 \text{ if } x_\alpha = 0\right\}$$

In many cases, $V(x)$ actually lies in the *tangent space* to $\mathcal{X}$ at x, viz.:

$$V(x) \in \mathrm{T}_{\mathcal{X}}(x) \equiv \left\{z \in \mathbb{R}^{\mathcal{A}} : \sum_{\alpha} z_\alpha = 0\right\}$$

In that case, no new strategies can enter the fray (the support of $x(t)$ remains invariant)

Revision protocols

Microfoundations for (ED) are often provided by means of a *revision protocol* that specifies the (unconditional) switch rates $s_{\alpha\beta}$ at which α-strategists revise their strategy to β.

Revision protocols

Microfoundations for (ED) are often provided by means of a *revision protocol* that specifies the (unconditional) switch rates $s_{\alpha\beta}$ at which α-strategists revise their strategy to β.

Such protocols $s_{\alpha\beta}$ are usually written as

$$s_{\alpha\beta} = p_\alpha r_{\alpha\beta}$$

where

- p_α is the rate at which α-strategists receive a revision opportunity.
- $r_{\alpha\beta}$ is the (*conditional*) switch rate from α to β (i.e. once an α-strategist has received a revision opportunity).

Revision protocols

Microfoundations for (ED) are often provided by means of a *revision protocol* that specifies the (unconditional) switch rates $s_{\alpha\beta}$ at which α-strategists revise their strategy to β.

Such protocols $s_{\alpha\beta}$ are usually written as

$$s_{\alpha\beta} = p_\alpha r_{\alpha\beta}$$

where

- p_α is the rate at which α-strategists receive a revision opportunity.
- $r_{\alpha\beta}$ is the (*conditional*) switch rate from α to β (i.e. once an α-strategist has received a revision opportunity).

The mean dynamics generated by a revision protocol are then:

$$\dot{x}_\alpha = \sum\nolimits_\beta \left[s_{\beta\alpha}(x, v(x)) - s_{\alpha\beta}(x, v(x)) \right], \qquad \text{(MD)}$$

i.e the rate of change of x_α is the inflow/outflow difference from/to strategy α.

Exampel 1: the replicator dynamics

The *replicator dynamics* (Taylor and Jonker, 1978) are defined as

$$\dot{x}_\alpha = x_\alpha \left[v_\alpha(x) - \sum_\beta x_\beta v_\beta(x) \right] \tag{RD}$$

i.e. the *per capita* growth rate of a strategy is the difference between its individual payoff and the mean population payoff.

Revision protocols that give rise to (RD) (up to constant payoff shifts):

- Imitation of success: $s_{\alpha\beta}(x, v) = x_\alpha x_\beta v_\beta$
- Pure imitation driven by dissatisfaction: $s_{\alpha\beta}(x, v) = -x_\alpha x_\beta v_\alpha$
- Pairwise proportional imitation: $s_{\alpha\beta}(x, v) = x_\alpha x_\beta [v_\beta - v_\alpha]_+$

Example 2: projection dynamics

The *projection dynamics* (Nagurney and Zhang, 1978; Lahkar and Sandholm, 2008) are defined on the interior of $\mathcal{X}$ as:

$$\dot{x}_\alpha = v_\alpha(x) - \frac{1}{|\mathcal{A}|}\sum_{\beta\in\mathcal{A}} v_\beta(x) \tag{RD}$$

i.e. the (absolute) growth rate of a strategy is given by the difference between its individual payoff and the (unweighted) average strategy payoff.

Revision protocols that give rise to (RD) (up to constant payoff shifts):

- Imitation of success driven by insecurity: $s_{\alpha\beta}(x,v) = \frac{1}{|\mathcal{A}|} v_\beta$
- Pure imitation driven by dissatisfaction and insecurity: $s_{\alpha\beta}(x,v) = -\frac{1}{|\mathcal{A}|} v_\alpha$
- Pairwise proportional imitation driven by insecurity: $s_{\alpha\beta}(x,v) = \frac{1}{|\mathcal{A}|}[v_\beta - v_\alpha]_+$

Similar to the replicator dynamics, but agents now quicker to abandon strategies that are used by few others.

Geometric underpinnings

- Under the projection dynamics, $\dot{x}$ is the closest point projection of $v(x)$ to the tangent cone of $\mathcal{X}$ at x – but what does "closest" mean?
- In potential games, the replicator dynamics ascend the game's potential function, and this ascent is the steepest possible under a non-Euclidean metric
- In potential games, payoffs are potential gradients, so the replicator dynamics also represent a "closest point projection" – but under a non-Euclidean metric

Geometric underpinnings

- Under the projection dynamics, $\dot{x}$ is the closest point projection of $v(x)$ to the tangent cone of $\mathcal{X}$ at x – but what does "closest" mean?
- In potential games, the replicator dynamics ascend the game's potential function, and this ascent is the steepest possible under a non-Euclidean metric
- In potential games, payoffs are potential gradients, so the replicator dynamics also represent a "closest point projection" – but under a non-Euclidean metric

Objective of the talk

Study evolutionary dynamics whose vectors of motion $V(x)$ are aligned to payoffs $v(x)$ to the "greatest possible" extent.

Outline

Background

Geometric Preliminaries

The Dynamics

Analysis

Scalar products and metrics

The notion of "closeness" presupposes a notion of "distance". Given a vector space W, this is often defined via a *scalar product*, i.e. a bilinear pairing $\langle\cdot,\cdot\rangle$ that is:

1. Symmetric: $\langle w, w'\rangle = \langle w', w\rangle$ for all $w, w' \in W$
2. Positive-definite: $\langle w, w\rangle \geq 0$ for all $w \in W$, with equality iff $w = 0$.

Scalar products and metrics

The notion of "closeness" presupposes a notion of "distance". Given a vector space W, this is often defined via a *scalar product*, i.e. a bilinear pairing $\langle\cdot,\cdot\rangle$ that is:

1. Symmetric: $\langle w, w'\rangle = \langle w', w\rangle$ for all $w, w' \in W$
2. Positive-definite: $\langle w, w\rangle \geq 0$ for all $w \in W$, with equality iff $w = 0$.

The *norm* of $w \in W$ is then defined as

$$\|w\| = \langle w, w\rangle^{1/2}$$

or, in components w.r.t. a basis $\{e_\alpha\}_{\alpha=1}^n$ of W:

$$\|w\| = \left(\sum_\alpha \sum_\beta g_{\alpha\beta} w_\alpha w_\beta\right)^{1/2}$$

where the so-called *metric elements* $g_{\alpha\beta}$ of $\langle\cdot,\cdot\rangle$ are defined as:

$$g_{\alpha\beta} = \langle e_\alpha, e_\beta\rangle$$

Notation: $\langle\cdot,\cdot\rangle$ will often be identified with its *metric tensor* $g = (g_{\alpha\beta})_{\alpha,\beta=1}^n$

Riemannian metrics

Scalar products can be position-dependent, leading to the notion of a *Riemannian metric :*

Definition

A *Riemannian metric* on an open subset U of W is a continuous assignment of a metric tensor $g(x)$ to each point $x \in U$. In words, a Riemannian metric on U prescribes a way of measuring angles and lengths at each point $x \in U$, so it determines the geometry of U.

Riemannian metrics

Scalar products can be position-dependent, leading to the notion of a *Riemannian metric :*

Definition

A *Riemannian metric* on an open subset U of W is a continuous assignment of a metric tensor $g(x)$ to each point $x \in U$. In words, a Riemannian metric on U prescribes a way of measuring angles and lengths at each point $x \in U$, so it determines the geometry of U.

Examples

1. The *Euclidean metric* on $U = \mathbb{R}^{\mathcal{A}}$ is

$$g_{\alpha\beta}(x) = \delta_{\alpha\beta}$$

leading to the familiar expressions $\langle w, w' \rangle_x = \sum_\alpha w_\alpha w'_\alpha$ and $\|w\|_x = \left(\sum_\alpha w_\alpha^2\right)^{1/2}$

2. The *Shahshahani metric* on $U = \mathbb{R}^{\mathcal{A}}_{>0}$ is given by

$$g_{\alpha\beta}(x) = \delta_{\alpha\beta}/x_\beta$$

leading to $\langle w, w' \rangle = \sum_\alpha w_\alpha w'_\alpha / x_\alpha$ and $\|w\|_x^2 = \left(\sum_\alpha w_\alpha^2 / x_\alpha\right)^{1/2}$

Deformations under different metrics I

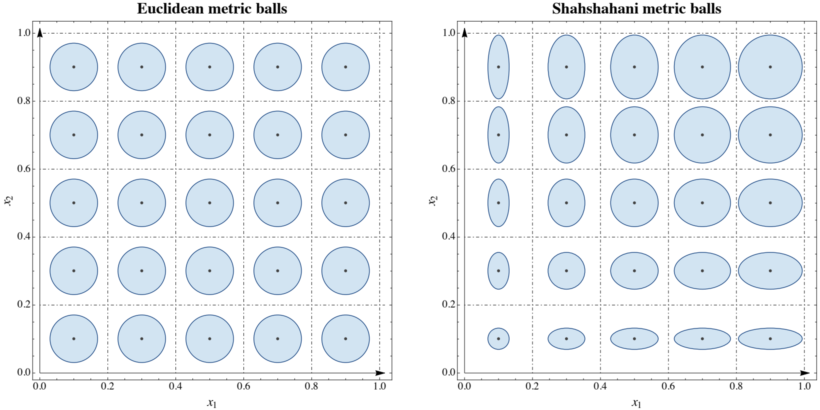

Unit balls on the positive orthant of $\mathbb{R}^2$ under the Euclidean and Shahshahani metrics

Deformations under different metrics II

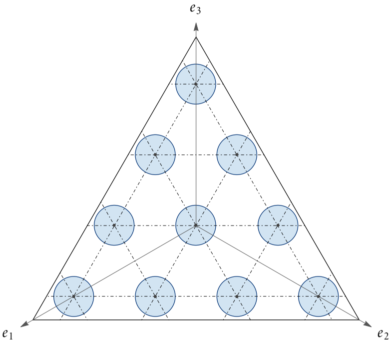

Unit balls on the unit simplex of $\mathbb{R}^2$ under the Euclidean and Shahshahani metrics

Duality: payoffs as covectors

The scalar product $\langle w, w' \rangle$ between two vectors $w, w' \in W$ is not to be confused with the action $\langle \omega | w \rangle \equiv \omega(w)$ of a linear functional $\omega \in W^*$ and a vector $w \in W$. In components:

$$\langle \omega | w \rangle = \sum_\alpha \omega_\alpha w_\alpha$$

but this is *not* a scalar product (ω and w live in different spaces)

Terminology: because of duality, linear functionals $\omega \in W^*$ are called covectors

Duality: payoffs as covectors

The scalar product $\langle w, w' \rangle$ between two vectors $w, w' \in W$ is not to be confused with the action $\langle \omega | w \rangle \equiv \omega(w)$ of a linear functional $\omega \in W^*$ and a vector $w \in W$. In components:

$$\langle \omega | w \rangle = \sum_\alpha \omega_\alpha w_\alpha$$

but this is *not* a scalar product (ω and w live in different spaces)

Terminology: because of duality, linear functionals $\omega \in W^*$ are called covectors

Payoff "vectors" naturally act as linear functionals $\implies$ should be treated as covectors

Examples

1. *Nash equilibrium:* $\sum_\alpha v_\alpha(x^*)(x_\alpha - x^*_\alpha) \leq 0$ is a requirement on the action of $v(x^*)$ on the displacement vector $x - x^*$
2. *Potential games:* the derivative $\nabla f = (\partial_1 f, \dots, \partial_n f)$ of a function f naturally acts on vectors to yield the directional derivative $\nabla_z f(x) = \langle \nabla f(x) | z \rangle$
   $\implies v(x) = \nabla f(x)$ behaves as a covector
3. *Positive correlation:* $\langle v(x) | V(x) \rangle \geq 0$ is a requirement on the action of $v(x)$ on the vector of motion $V(x)$

Primal equivalents: payoffs as vectors

Given a Riemannian metric, there is a canonical way to translate covectors into vectors:

Definition

The *primal equivalent* of a covector $\omega \in W^*$ is the unique vector $\omega^\sharp \in W$ such that

$$\langle \omega | w \rangle = \langle \omega^\sharp, w \rangle \quad \text{for all } w \in W.$$

In components, $\omega^\sharp_\alpha = \sum_\beta g^\sharp_{\alpha\beta}(x)\omega_\beta$ where $g^\sharp(x) \equiv g^{-1}(x)$ is the inverse matrix of $g(x)$

Examples

1. In the Euclidean case, $g_{\alpha\beta} = \delta_{\alpha\beta} = g^\sharp_{\alpha\beta}$ so we have

$$\omega^\sharp_\alpha = \omega_\alpha$$

(Source of the standard confusion between scalar products and duality pairings)

2. In the Shahshahani case, $g_{\alpha\beta} = \delta_{\alpha\beta}/x_\beta$, so $g^\sharp_{\alpha\beta} = x_\alpha \delta_{\alpha\beta}$, and we obtain

$$\omega^\sharp_\alpha = x_\alpha \omega_\alpha$$

Extensions to boundary points

The Shahshahani metric explodes at the boundary of $\mathcal{X}$, so care must be taken there:

Definition

Let $\mathcal{K}$ be a closed convex subset of $\mathbb{R}^{\mathcal{A}}$ with nonempty interior. We say that a Riemannian metric g on $\mathcal{K}^\circ$ is *extendable* to $\mathcal{K}$ if the associated sharp operator $g^\sharp$ admits a (necessarily unique) continuous extension to $\mathcal{K}$ such that

$$\mathrm{T}_{\mathcal{K}}(x) \subseteq \operatorname{im} g^\sharp(x) \quad \text{for all } x \in \mathcal{K}.$$

- The space $\operatorname{im} g^\sharp(x)$ is called the *domain* of g at x and is denoted as $\operatorname{dom} g(x)$.
- If $\operatorname{dom} g(x) = \mathbb{R}^{\mathcal{A}}$, we say that g is *full-rank extendable*; instead, if $\operatorname{dom} g(x) = \mathrm{T}_{\mathcal{K}}(x)$, we say that g is *minimal-rank extendable*.

Intuition: the extension of g must metrize (at least) all tangent vectors to $\mathcal{K}$ at x.

Extensions to boundary points, cont'd

Why this complicated definition?

Proposition

If g *is extendable, there exists a unique scalar product* $\langle\cdot,\cdot\rangle_x$ *on* $\operatorname{dom} g(x)$ *such that* $\langle w|w'\rangle_{x_k} \to \langle w|w'\rangle_x$ *for all* $w, w' \in \operatorname{dom} g(x)$ *and for all sequences* x_k *in* $\mathcal{K}^\circ$ *with* $x_k \to x$.

In particular, if $\mathcal{K} = \mathbb{R}_+^{\mathcal{A}}$ *and* g *is minimal-rank extendable, we have* $g^\sharp_{\alpha\beta}(x) = 0$ *whenever* $x_\alpha = 0$ *or* $x_\beta = 0$.

Extensions to boundary points, cont'd

Why this complicated definition?

Proposition

If g is extendable, there exists a unique scalar product $\langle\cdot,\cdot\rangle_x$ on $\operatorname{dom} g(x)$ such that $\langle w|w'\rangle_{x_k} \to \langle w|w'\rangle_x$ for all $w, w' \in \operatorname{dom} g(x)$ and for all sequences x_k in $\mathcal{K}^\circ$ with $x_k \to x$.

In particular, if $\mathcal{K} = \mathbb{R}_+^{\mathcal{A}}$ and g is minimal-rank extendable, we have $g^\sharp_{\alpha\beta}(x) = 0$ whenever $x_\alpha = 0$ or $x_\beta = 0$.

Examples

Let $\mathcal{K}$ denote the closed positive orthant of $\mathbb{R}^{\mathcal{A}}$ (= generalized population states). Then:

1. The Euclidean metric $g_{\alpha\beta}(x) = \delta_{\alpha\beta}$ is full-rank extendable (by default).
2. The Shahshahani metric $g_{\alpha\beta}(x) = \delta_{\alpha\beta}/x_\beta$ has $g^\sharp_{\alpha\beta}(x) = x_\alpha\delta_{\alpha\beta}$, so $\operatorname{dom} g(x) = \mathrm{T}_{\mathcal{K}}(x)$

   $\Longrightarrow$ The Shahshahani metric is minimal-rank extendable and

$$\langle w, w'\rangle_x = \sum_{\alpha\in\operatorname{supp}(x)} w_\alpha w'_\alpha/x_\alpha \quad \text{for all } w, w' \in \mathrm{T}_{\mathcal{K}}(x)$$

Tangent projections

Main idea: follow tangent projections of payoffs on $\mathcal{X}$

Definition

Let g be an extendable Riemannian metric on $\mathcal{K} = \mathbb{R}^n_+$. Then, the *tangent projection* of $w \in \operatorname{dom} g(x)$ to $\mathcal{X}$ at $x \in \mathcal{X}$ is

$$\Pi_x(w) = \arg\min\{\|w - z\|_x : z \in \mathrm{TC}_{\mathcal{X}}(x) \cap \operatorname{dom} g(x)\}$$

Tangent projections

Main idea: follow tangent projections of payoffs on $\mathcal{X}$

Definition

Let g be an extendable Riemannian metric on $\mathcal{K} = \mathbb{R}^n_+$. Then, the *tangent projection* of $w \in \operatorname{dom} g(x)$ to $\mathcal{X}$ at $x \in \mathcal{X}$ is

$$\Pi_x(w) = \arg\min\{\|w - z\|_x : z \in \mathrm{TC}_{\mathcal{X}}(x) \cap \operatorname{dom} g(x)\}$$

Remark: minimization is over all tangent vectors $\mathrm{TC}_{\mathcal{X}}(x)$ that are metrizable; this need not be the entire tangent cone if g is not full-rank extendable:

Definition

We call this set the *g-admissible tangent cone* of $\mathcal{X}$ at x:

$$\mathrm{Adm}_g(x) = \mathrm{TC}_{\mathcal{X}}(x) \cap \operatorname{dom} g(x)$$

Examples of g-admissible cones

Admissible cones under the Euclidean and Shahshahani metrics.

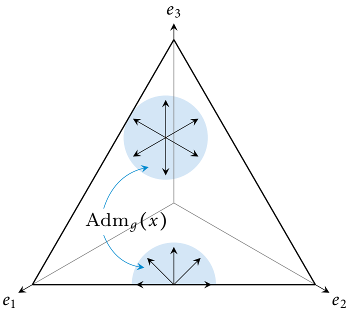

(e) Euclidean boundary behavior.

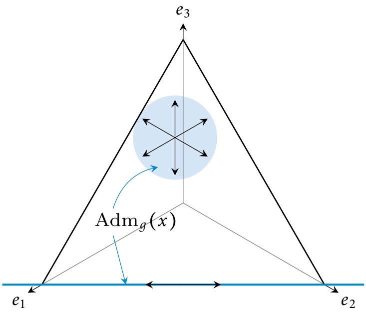

(f) Shahshahani boundary behavior.

Coordinate expressions

Case 1: computing tangent projections is straightforward if x is interior:

1. Note that the vector $n(x) = \mathbf{1}^\sharp$ is normal to $\mathcal{X}$ at x:

$$\langle n(x), z \rangle = \langle \mathbf{1} | z \rangle = \sum_\alpha z_\alpha = 0 \quad \text{for all } z \in \mathrm{T}_{\mathcal{X}}(x)$$

2. By subtracting the component of w along $n(x)$, we have

$$\Pi_x(w) = w - \operatorname{proj}_{n(x)} w = w - \frac{\langle n(x), w \rangle_x}{\| n(x) \|_x^2} n(x)$$

or, in coordinates:

$$(\Pi_x(w))_\alpha = w_\alpha - \frac{\sum_\gamma w_\gamma}{\sum_\gamma n_\gamma(x)} n_\alpha(x)$$

where $n_\alpha(x) = \sum_\beta g^\sharp_{\alpha\beta}(x)$ is the α-th component of $n(x)$.

Case 2: construction similar if $\operatorname{dom} g(x) = \mathbb{R}^{\operatorname{supp}(x)}$

Case 3: if $\operatorname{dom} g(x) \nsupseteq \mathbb{R}^{\operatorname{supp}(x)}$, must solve a convex problem.

Examples of unit normals

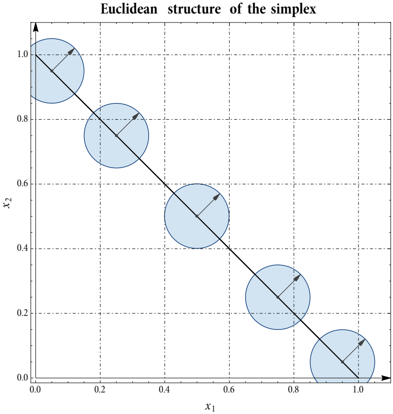

(g) Euclidean unit normals

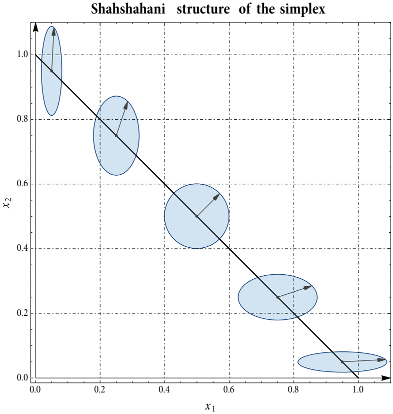

(h) Shahshahani unit normals

Examples of tangent projections

1. *Euclidean projections:*
   Given $w \in \mathbb{R}^n$, its tangent projection to $\mathcal{X}$ at x is

$$(\Pi_x(w))_\alpha = w_\alpha - |\mathcal{A}(x)|^{-1} \sum_{\gamma \in \mathcal{A}(x)} w_\gamma,$$

   where $\mathcal{A}(x)$ is a certain subset of $\mathcal{A}$ that contains $\mathrm{supp}(x)$ – details in Lahkar and Sandholm, 2008.

2. *Shahshahani projections:*
   Given $w \in \mathbb{R}^{\mathrm{supp}(x)}$, its tangent projection to $\mathcal{X}$ at x is

$$(\Pi_x(w))_\alpha = w_\alpha - x_\alpha \sum_\gamma w_\gamma$$

NB: Euclidean and Shahshahani projections do not have the same domain.

Examples of tangent projections

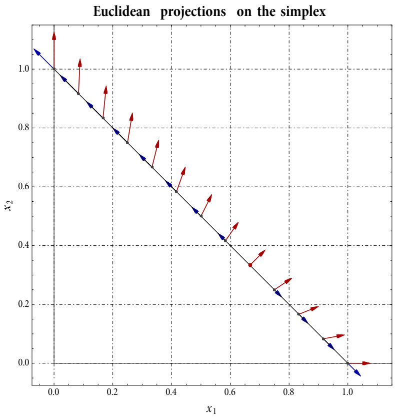

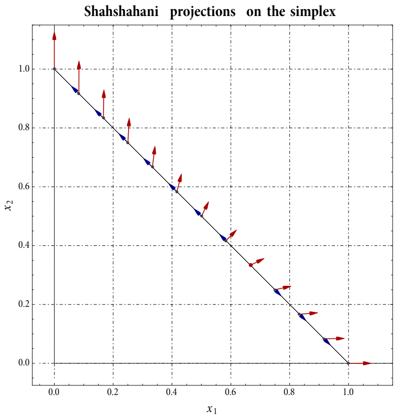

Outline

Definition and coordinate expressions

Recall the main idea: vectors of motion are aligned to payoffs as closely as possible

This gives the class of *Riemannian game dynamics*

$$\dot{x} = \Pi_x\left(v^\sharp(x)\right) \tag{Rie}$$

where $v^\sharp(x)$ is the primal equivalent of the payoff covector $v(x)$ at population state x and $\Pi_x(v^\sharp(x))$ is its tangent projection to $\mathcal{X}$ at x.

In particular, when x is interior or g is minimal-rank extendable, we get (up to a change in time):

$$\dot{x}_\alpha = v^\sharp_\alpha(x)\sum\nolimits_\gamma n_\gamma(x) - n_\alpha(x)\sum\nolimits_\gamma v^\sharp_\gamma(x)$$

where $v^\sharp_\alpha(x) = \sum_\beta g^\sharp_{\alpha\beta}(x)$ and $n_\alpha(x) = \sum_\beta g^\sharp_{\alpha\beta}(x)$.

Examples

1. *The replicator dynamics:*
   If g is the Shahshahani metric, $g^\sharp_{\alpha\beta}(x) = \delta_{\alpha\beta} x_\beta$, so $v^\sharp_\alpha(x) = x_\alpha v_\alpha(x)$ and $n_\alpha(x) = x_\alpha$ $\Longrightarrow$

$$\dot{x}_\alpha = x_\alpha \left[v_\alpha(x) - \sum\nolimits_\beta x_\beta v_\beta(x) \right] \tag{RD}$$

2. *The projection dynamics:*
   If g is the Euclidean metric, $g_{\alpha\beta} = \delta_{\alpha\beta}$, $v^\sharp_\alpha = v_\alpha$ and $n_\alpha(x) = 1$. Ultimately, we get:

$$\dot{x}_\alpha = \begin{cases} v_\alpha(x) - |\mathcal{A}(x)|^{-1} \sum_{\beta \in \mathcal{A}(x)} v_\beta(x) & \text{if } \alpha \in \mathcal{A}(x), \\ 0 & \text{otherwise,} \end{cases} \tag{PD}$$

   where $\mathcal{A}(x)$ is a subset of $\mathcal{A}$ containing $\mathrm{supp}(x)$ – details in Lahkar and Sandholm (2008).

Examples, cont'd

3. *Separable metrics:*

Let g be of the form $g_{\alpha\beta}(x) = \delta_{\alpha\beta}/\phi(x_\beta)$ with $\phi\colon [0,\infty) \to [0,\infty)$ a *weighting function*. We then get the dynamics:

$$\dot{x}_\alpha = \phi(x_\alpha)\left[v_\alpha(x)\sum\nolimits_{\beta\in\mathcal{A}}\phi(x_\beta) - \sum\nolimits_{\beta\in\mathcal{A}}\phi(x_\beta)v_\beta(x)\right]$$

Weighted imitation of success: α-strategists receive revision opportunities with rate $\phi(x_\alpha)$ and switch to β with (conditional) rate $\phi(x_\beta)v_\alpha$

Examples, cont'd

4. *Hessian Riemmanian metrics:*

   Let $h\colon \mathbb{R}_+^{\mathcal{A}} \to \mathbb{R}$ be a smooth function s.t. $\mathrm{Hess}(h(x)) \succ 0$ for all $x \succ 0$. Then, the *Hessian Riemannian metric* induced by g is

$$g_{\alpha\beta}(x) = \frac{\partial^2 h}{\partial x_\alpha \partial x_\beta}$$

   and the induced dynamics take the form:

$$\dot{x}_\alpha = \sum_{\beta \in \mathrm{supp}(x_k)} \left[g^\sharp_{\alpha\beta}(x) - \frac{1}{G^\sharp(x)} g^\sharp_\alpha(x) g^\sharp_\beta(x) \right] v_\beta(x), \tag{HRD}$$

   where $g^\sharp_{\alpha\beta}(x) = \mathrm{Hess}(h(x))^{-1}_{\alpha\beta}$, $g^\sharp_\alpha = \sum_\beta g^\sharp_{\alpha\beta}$, and $G^\sharp(x) = \sum_\alpha g^\sharp_\alpha(x)$.

   Antecedent: closely related to a class of reinforcement learning dynamics studied by M. and Sandholm (2015).

Phase portraits

Phase portraits for a Rock-Paper-Scissors game under different metrics:

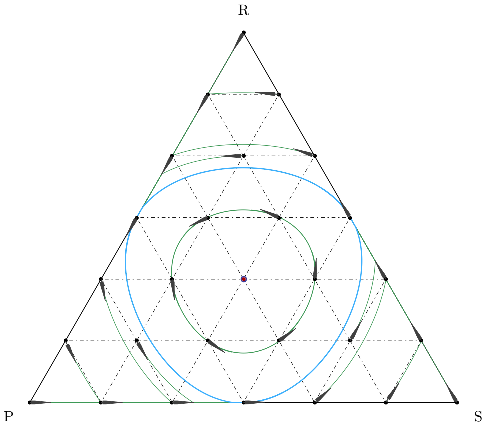

$$g_{\alpha\beta} = \delta_{\alpha\beta} \exp(x_\alpha)$$

Phase portraits

Phase portraits for a Rock-Paper-Scissors game under different metrics:

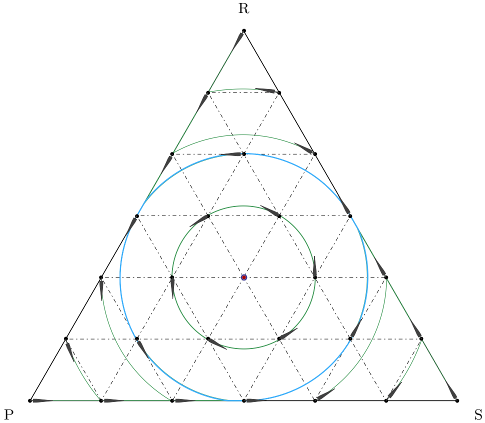

$$g_{\alpha\beta} = \delta_{\alpha\beta}$$

Phase portraits

Phase portraits for a Rock-Paper-Scissors game under different metrics:

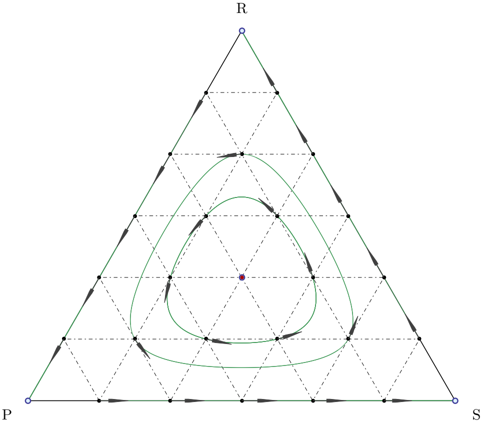

$$g_{\alpha\beta} = \delta_{\alpha\beta}/x_\beta$$

Phase portraits

Phase portraits for a Rock-Paper-Scissors game under different metrics:

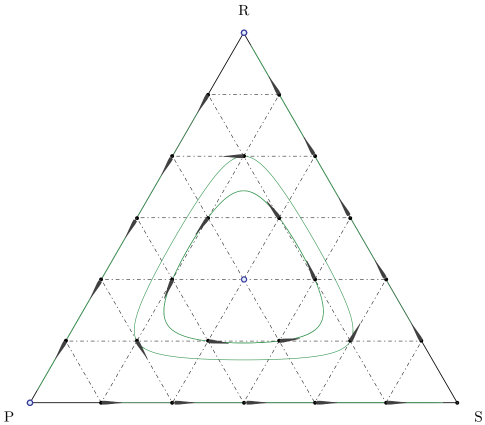

$$g_{\alpha\beta} = \delta_{\alpha\beta}/x_\alpha^{3/2}$$

Outline

Well-posedness

The first fundamental question is global existence and uniqueness of solutions. This is true if g is Lipschitz continuous:

Proposition

Let g be a Lipschitz extendable Riemannian metric on $\mathbb{R}^{\mathcal{A}}$. Then, for every $x_0 \in \mathcal{X}$, the induced dynamics admit a unique forward solution $x(t)$ with $x(0) = x_0$.

In particular, if the underlying Riemannian metric is minimal-rank extendable, all solution orbits are smooth and the relative interior of each face of $\mathcal{X}$ is forward invariant: $x_\alpha(t) = 0$ for some $t > 0$ if and only if $x_\alpha(0) = 0$.

Idea of proof.

- Existence is trivial.
- Main difficulty for uniqueness is the discontinuity of the motion field at the boundary of $\mathcal{X}$.
- The Cojocaru and Jonker theory for projected systems also fails (applies to Hilbert spaces, not Hilbert manifolds).
- Use ad hoc estimates to show that the system is 1-sided Lipschitz near discontinuity points and apply the differential inclusion results of Kunze (2000). □

Further results

Proposition

Riemannian game dynamics satisfy the positive correlation condition:

$$\langle v(x) | V(x) \rangle \geq 0,$$

with equality if and only if $V(x) = 0$.

Natural: motion field approximates payoff field to the greatest possible extent

Further results

Proposition

Riemannian game dynamics satisfy the positive correlation condition:

$$\langle v(x) | V(x) \rangle \geq 0,$$

with equality if and only if $V(x) = 0$.

Natural: motion field approximates payoff field to the greatest possible extent

However, the dynamics exhibit an important dichotomy w.r.t. equilibrium play:

Proposition

- *If g is full-rank extendable, $x^* \in \mathcal{X}$ is stationary if and only if it is a Nash equilibrium.*
- *If g is minimal-rank extendable, $x^* \in \mathcal{X}$ is stationary if and only if it is a restricted Nash equilibrium (i.e. an equilibrium of a restriction of the game to a subface of $\mathcal{X}$).*

(Explains the difference between the replicator dynamics and the projection dynamics)

Some observations

Some observations before moving further:

- Dominated strategies may survive under the (Euclidean) projection dynamics (Sandholm et al., 2008).
- Survival under the projection dynamics is due to discontinuities at the boundary; however, dominated strategies can also survive under smooth Riemannian dynamics.

Not all metrics lead to rational outcomes $\Longrightarrow$ Must refine the class of metrics under study

Distance-like functions

A key element of our analysis is the notion of a distance-like function:

Definition

Let g be an extendable Riemannian metric on $\mathcal{K}$ and let $p \in \mathcal{K}$. A smooth function $D_p : \mathcal{K}^\circ \to \mathbb{R}$ is called *distance-like* w.r.t. p if:

1. D_p is positive-definite: $D_p(x) \geq 0$ for all $x \in \mathcal{K}^\circ$ and $D_p(x_n) \to 0$ for a sequence $x_k \in \mathcal{K}^\circ$ if and only if $x_k \to p$.
2. D_p is radial:

$$\operatorname{grad} D_p(x) = \sigma(x)(x - p)$$

   for some smooth positive function $\sigma : \mathcal{K}^\circ \to (0, \infty)$.

If g admits distance-like functions, we say that g is *integrable.*

Examples

1. *The Euclidean case:* $D_p(x) = \frac{1}{2} \sum_\beta (x_\beta - p_\beta)^2$ is distance-like.
2. *The Shahshahani case:* the Kullback-Leibler divergence $D_p(x) = \sum_\alpha p_\alpha \log p_\alpha / x_\alpha$ is distance-like

Elimination of Dominated Strategies

A strategy $p \in \mathcal{X}$ is *dominated* by p' if $\langle v(x) | p - p' \rangle < 0$ for all $x \in \mathcal{X}$.

Elimination of Dominated Strategies

A strategy $p \in \mathcal{X}$ is *dominated* by p' if $\langle v(x) | p - p' \rangle < 0$ for all $x \in \mathcal{X}$.

A strategy $p \in \mathcal{X}$ becomes extinct along $x(t)$ if

$$\min\{x_\alpha(t) : \alpha \in \operatorname{supp}(p)\} \to 0 \quad \text{as } t \to \infty$$

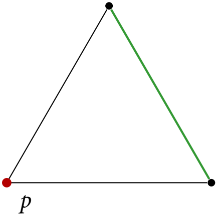

Elimination of Dominated Strategies

A strategy $p \in \mathcal{X}$ is *dominated* by p' if $\langle v(x) | p - p' \rangle < 0$ for all $x \in \mathcal{X}$.

A strategy $p \in \mathcal{X}$ becomes extinct along $x(t)$ if

$$\min\{x_\alpha(t) : \alpha \in \mathrm{supp}(p)\} \to 0 \quad \text{as } t \to \infty$$

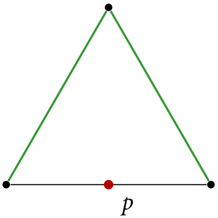

Elimination of Dominated Strategies

A strategy $p \in \mathcal{X}$ is *dominated* by p' if $\langle v(x) | p - p' \rangle < 0$ for all $x \in \mathcal{X}$.

A strategy $p \in \mathcal{X}$ becomes extinct along $x(t)$ if

$$\min\{x_\alpha(t) : \alpha \in \operatorname{supp}(p)\} \to 0 \quad \text{as } t \to \infty$$

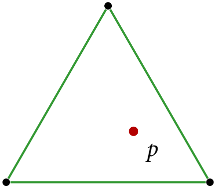

Elimination of Dominated Strategies

A strategy $p \in \mathcal{X}$ is *dominated* by p' if $\langle v(x) | p - p' \rangle < 0$ for all $x \in \mathcal{X}$.

A strategy $p \in \mathcal{X}$ becomes extinct along $x(t)$ if

$$\min\{x_\alpha(t) : \alpha \in \operatorname{supp}(p)\} \to 0 \quad \text{as } t \to \infty$$

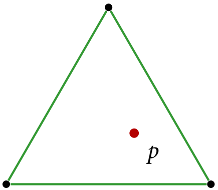

Theorem

If g is minimal-rank extendable and integrable, dominated strategies become extinct along every interior solution of the induced Riemannian dynamics.

Stability and convergence to equilibrium play

Proposition

Suppose that g is minimal-rank extendable and integrable, and let $x^* \in \mathcal{X}$. If a) $x(t) \to x^*$ for some interior orbit; or b) x^* is Lyapunov stable, then x^* is a Nash equilibrium.

Stability and convergence to equilibrium play

Proposition

Suppose that g is minimal-rank extendable and integrable, and let $x^* \in \mathcal{X}$. If a) $x(t) \to x^*$ for some interior orbit; or b) x^* is Lyapunov stable, then x^* is a Nash equilibrium.

Theorem

Suppose that g is minimal-rank extendable and integrable, and let $x^* \in \mathcal{X}$. If x^* is evolutionarily stable, then it is asymptotically stable under the induced Riemannian dynamics; and if it is globally evolutionarily stable, its basin of attraction includes the interior of $\mathcal{X}$.

(Converse direction: asymptotically stable states are isolated, perfect equilibria)

Permanence

An important evolutionary question is whether all strategies remain above a given threshold in the long run (if not, small fluctuations could lead to extinction of otherwise healthy types):

Definition (Hofbauer and Sigmund, 1998)

A dynamical system on $\mathcal{X}$ is called *permanent* if there exists some $\delta > 0$ such that $\liminf_{t\to\infty} x_\alpha(t) \geq \delta$ for all $\alpha \in \mathcal{A}$ and for all interior conditions $x(0) \in \mathcal{X}^\circ$.

Permanence

An important evolutionary question is whether all strategies remain above a given threshold in the long run (if not, small fluctuations could lead to extinction of otherwise healthy types):

Definition (Hofbauer and Sigmund, 1998)

A dynamical system on $\mathcal{X}$ is called *permanent* if there exists some $\delta > 0$ such that $\liminf_{t\to\infty} x_\alpha(t) \geq \delta$ for all $\alpha \in \mathcal{A}$ and for all interior conditions $x(0) \in \mathcal{X}^\circ$.

Theorem

Suppose that g is minimal-rank extendable and integrable, and assume there is a population state $p \in \mathcal{X}^\circ$ such that

$$\langle v(x^*) | p - x^* \rangle > 0 \tag{P}$$

for all restricted equilibria x^*. Then, the Riemannian game dynamics induced by g are permanent.

Permanence

An important evolutionary question is whether all strategies remain above a given threshold in the long run (if not, small fluctuations could lead to extinction of otherwise healthy types):

Definition (Hofbauer and Sigmund, 1998)

A dynamical system on $\mathcal{X}$ is called *permanent* if there exists some $\delta > 0$ such that $\liminf_{t\to\infty} x_\alpha(t) \geq \delta$ for all $\alpha \in \mathcal{A}$ and for all interior conditions $x(0) \in \mathcal{X}^\circ$.

Theorem

Suppose that g is minimal-rank extendable and integrable, and assume there is a population state $p \in \mathcal{X}^\circ$ such that

$$\langle v(x^*) | p - x^* \rangle > 0 \tag{P}$$

for all restricted equilibria x^*. Then, the Riemannian game dynamics induced by g are permanent.

Consequence: If (**??**) holds, the time-average $\bar{x}(t) = t^{-1} \int_0^t x(s)\, ds$ of the population's state converges to an interior Nash equilibrium.

Questions?