Decision-Payoff Co-Evolutionary Dynamics in a Closed Loop for Population Games

Prof. Ming Cao

Institute of Engineering and Technology
University of Groningen
The Netherlands

An evolutionary approach to coordination of self-interested agents

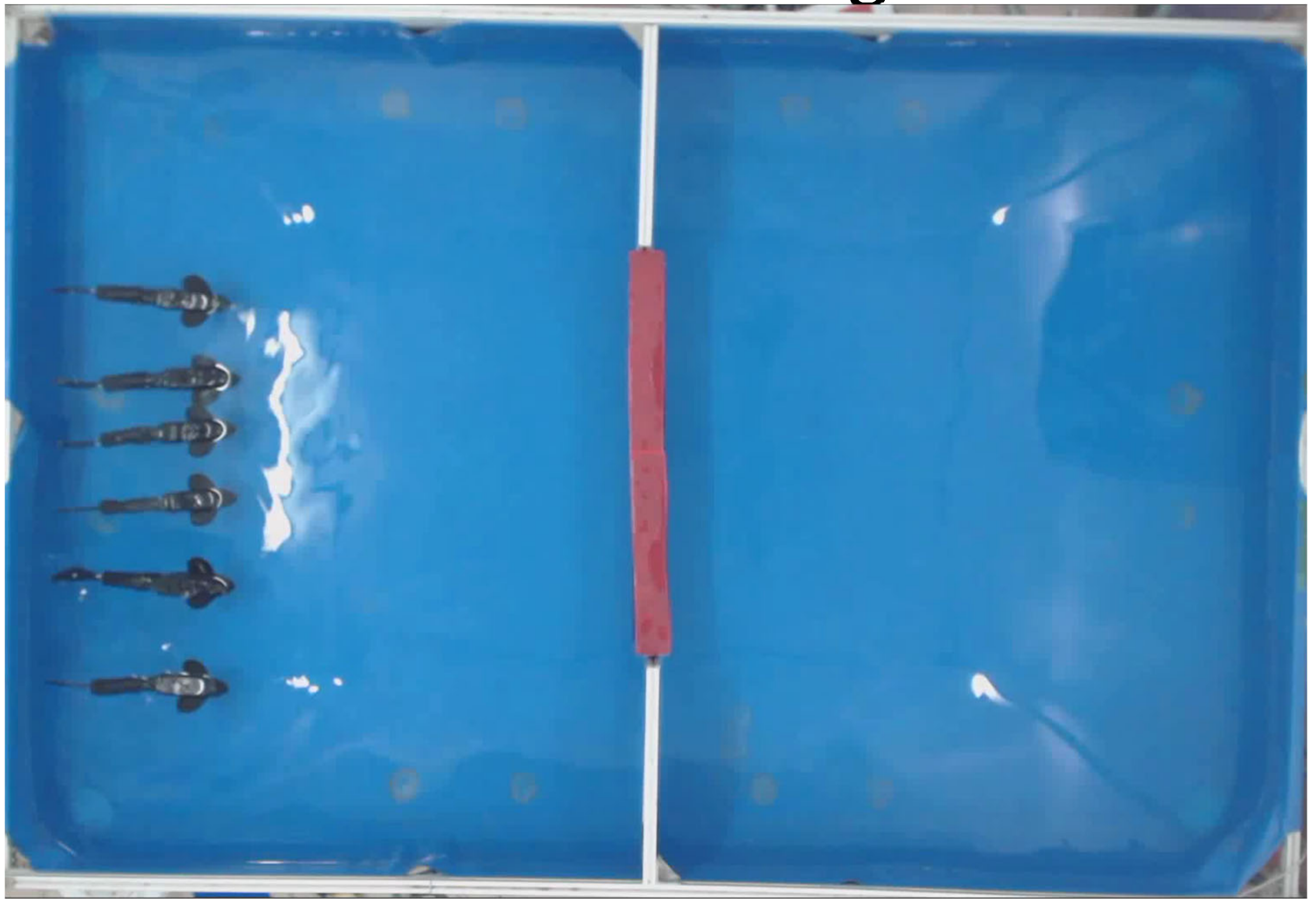

Carry out the task repeatedly; adjust strategies each time

- each time the task is taken as a group game
- new insight into how cooperation emerge as an evolutionary outcome

The paradox of cooperation



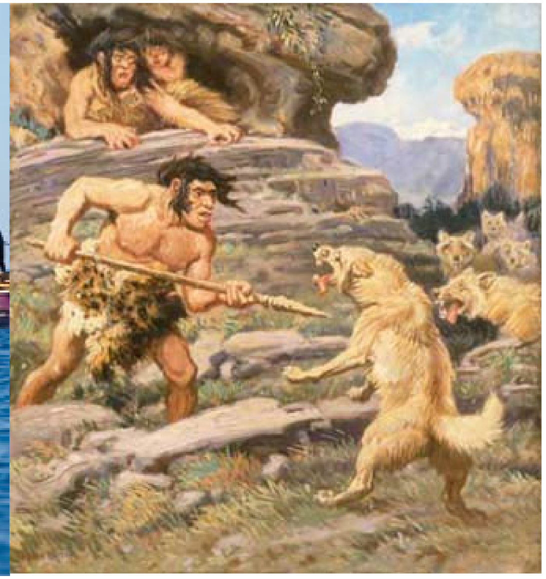

Natural selection is based on competition. How can it lead to cooperation?

Cooperation is often costly for the individual, while benefits are distributed over a collective

Cooperation (altruism) is an evolutionary puzzle!

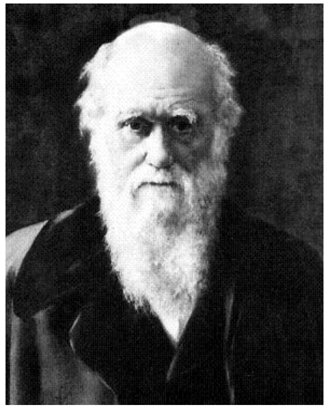

Charles Darwin
(1809-1882)

Mechanism for evolution of cooperation is a central topic

THE QUESTIONS

The Top 25
Essays by our news staff on 25 big questions facing science over the next quarter-century.

- What Is the Universe Made Of?
- What is the Biological Basis of Consciousness?
- Why Do Humans Have So Few Genes?
- To What Extent Are Genetic Variation and Personal Health Linked?
- Can the Laws of Physics Be Unified?
- How Much Can Human Life Span Be Extended?
- What Controls Organ Regeneration?
- How Can a Skin Cell Become a Nerve Cell?
- How Does a Single Somatic Cell Become a Whole Plant?
- How Does Earth's Interior Work?
- Are We Alone in the Universe?
- How and Where Did Life on Earth Arise?
- What Determines Species Diversity?
- What Genetic Changes Made Us Uniquely Human?
- How Are Memories Stored and Retrieved?
- How Did Cooperative Behavior Evolve?

Science AAAS

AAAS.ORG | FEEDBACK | HELP | LIBRARIANS

All Science Journals | Enter Search T

UNIVERSITY OF GRONINGEN | ALERTS | ACCESS

NEWS | *SCIENCE* JOURNALS | CAREERS | MULTIMEDIA | COLLECTIONS

Subject Collections | Online Extras | *Science* Special Collections | Archived Collections | About Collections

Home > Collections > Online Extras > Special Issues 2005 > 1 July 2005: 125th Anniversary

About Our Sites

- About *Science*
- About *Science*NOW
- About *Science Signaling*
- About *Science Translational Medicine*
- About *Science* Careers
- About Multimedia
- About Collections

Alerts & Feeds

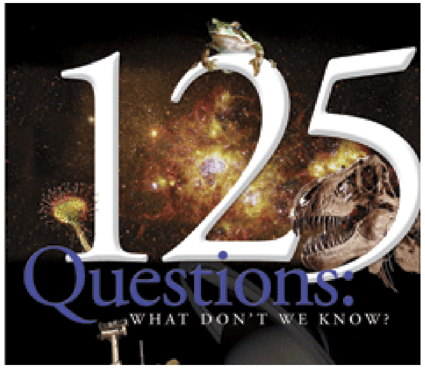

In a special collection of articles published beginning 1 July 2005, *Science* Magazine and its online companion sites celebrate the journal's 125th anniversary with a look forward -- at the most compelling puzzles and questions facing scientists today. A special, free news feature in *Science* explores 125 big questions that face scientific inquiry over the next quarter-century; accompanying the feature are several online extras including a reader's forum on the big questions. The Signal Transduction Knowledge Environment highlights some classic *Science* papers that have influenced the study of cell signaling. The Science of Aging Knowledge Environment looks at several important questions confronting researchers on aging. And *Science*'s Next Wave introduces us to four young scientists building their careers grappling with some of the very questions that *Science* has identified.

Jump to Features in This Special Collection:

- In *Science* Magazine
- In *Science*'s STKE
- In *Science*'s SAGE KE
- In *Science*'s Next Wave

How Did Cooperative Behavior Evolve?

REVIEW

Five Rules for the Evolution of Cooperation

Martin A. Nowak

Cooperation is needed for evolution to construct new levels of organization. Genomes, cells, multicellular organisms, social insects, and human society are all based on cooperation. Cooperation means that selfish replicators forgo some of their reproductive potential to help one another. But natural selection implies competition and therefore opposes cooperation unless a specific mechanism is at work. Here I discuss five mechanisms for the evolution of cooperation: kin selection, direct reciprocity, indirect reciprocity, network reciprocity, and group selection. For each mechanism, a simple rule is derived that specifies whether natural selection can lead to cooperation.

M. A. Nowak, *Science*, V314, 1560-1563, 2006

- Kin selection
- Direct reciprocity (“tit-for-tat”)
- Indirect reciprocity
- Network reciprocity
- Group selection

rich theory – limited predictive power…

Free-rider problem: the Prisoner's dilemma

	cooperate	defect
cooperate	$b-c$	$-c$
defect	b	0

$$b > c > 0$$

Mutual cooperation more profitable than mutual defection

But: under all circumstances defect is the dominating strategy

At group level: Public goods game

Everybody profits from public good, whether contributing or not

➡ Erosion of the public good

Evolutionary game theory: History and motivation

John Maynard Smith was interested in why so many animals engage in ritualized fighting ("The logic of animal conflict", *Nature*, 1973)

Evolutionary game theory (EGT) refers to the study of large populations of interacting agents, and how various behaviors and traits might evolve.

Differences from classical game theory

- Players = sub-populations, employing a common strategy
- Strategies = behaviors that learn to update or traits encoded in genes
- Payoffs = fitness, which determines update or reproductive rates

Key concept: The fitness of an individual must be evaluated in the context of the population in which it lives and interacts

Outline

- Evolutionary game model: replicator dynamics
- Two-population dynamics with environmental feedback
- Controlling evolutionary network dynamics using incentives

Dynamical system description for evolutionary games

Matrix game (symmetric two-player normal form with finitely many strategies)

Pure strategies: $e_1, e_2, \ldots, e_m$ unit vectors in $\mathbb{R}^m$

Mixed strategies: $p = (p_1, \ldots, p_m) \in \Delta^m$

where $\Delta^m = \{p | \sum_{i=1}^m p_i = 1\}$
$p_i \geq 0$

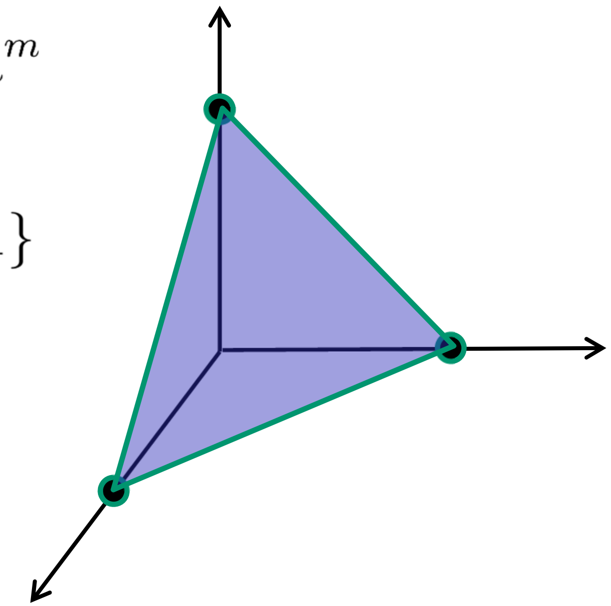

- Payoff: Individuals interact in a two-player game
- Let $\pi(e_i, e_j)$ be the payoff of e_i against e_j ,
  then the *m*-by-*m* payoff matrix *A* has entries $A_{ij} = \pi(e_i, e_j)$
- The payoff of *p* against *q* is $\pi(p, q) = \sum_{i,j=1}^m p_i \pi(e_i, e_j) q_j = p^T A q$

Dynamical system description for evolutionary games

Matrix game (symmetric two-player normal form with finitely many strategies)

Assumptions:

- Well-mixed large population (all-to-all network)
- Random pairwise interaction per unit time
- Payoff translate directly into fitness that determines reproductive rate
- Individuals use pure strategies

Key concept: The fitness of an individual must be evaluated in the context of the population in which it lives and interacts

Let $x_i(t)$ denote the share of those individuals using strategy e_i at time t, and $x = [x_1, x_2, \ldots, x_m]$ be the population vector. Then

- The payoff of p against q is $\pi(p,q) = \sum_{i,j=1}^{m} p_i \pi(e_i, e_j) q_j = p^T A q$

Replicator dynamics: $\dot{x}_i = x_i(\pi(e_i, x) - \pi(x, x))$

Evolutionary game dynamics

Replicator dynamics: $\dot{x}_i = x_i(\pi(e_i, x) - \pi(x, x))$

Evolutionary game dynamics

Replicator dynamics: $\dot{x}_i = x_i(\pi(e_i, x) - \pi(x, x))$

Central topics:

- Will the solution converge? If so, will it converge to the *Nash equilibrium* of the game?
- Which equilibrium is *evolutionarily stable*?
- Are the stable solutions also *optimal* maximizing some or all players' payoffs?

"A survey on the analysis and control of evolutionary matrix games," J. R. Riehl, P. Ramazi and M. Cao. *Annual Reviews in Control*, 45(6), 87-106, 2018

Environmental feedback

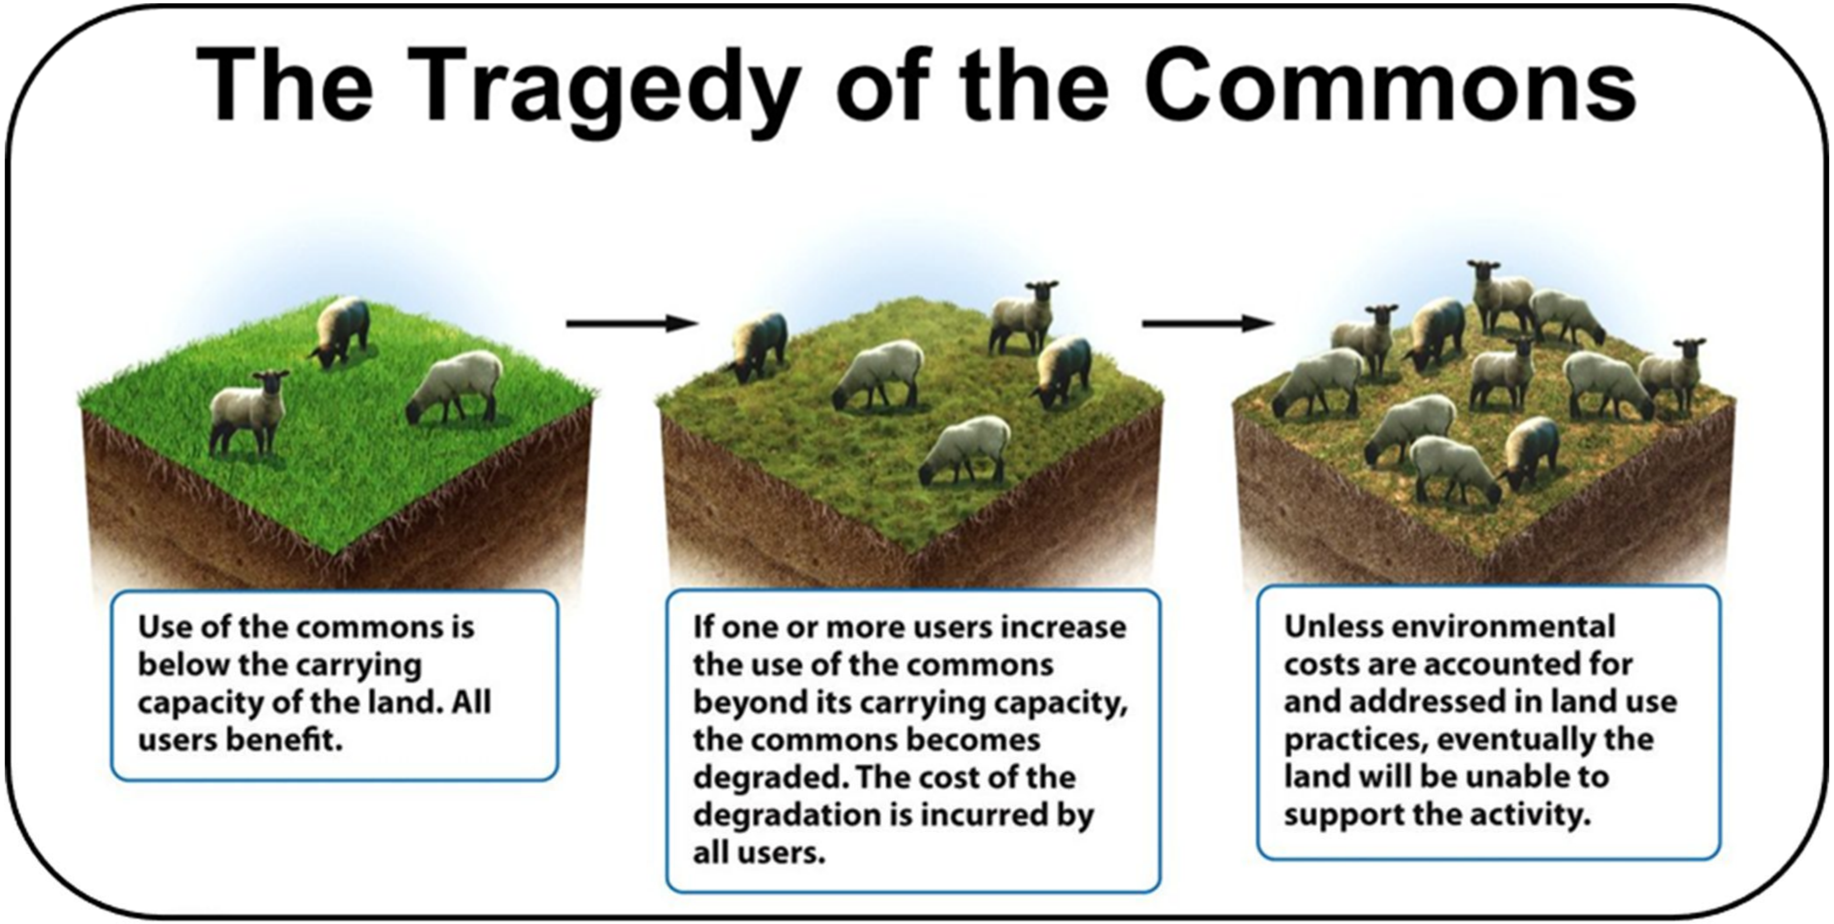

Environmental feedback

- Known:

Individuals' actions can reshape the surrounding environment.

- Unknown

Does the environment affect individuals' decisions? If so, how?

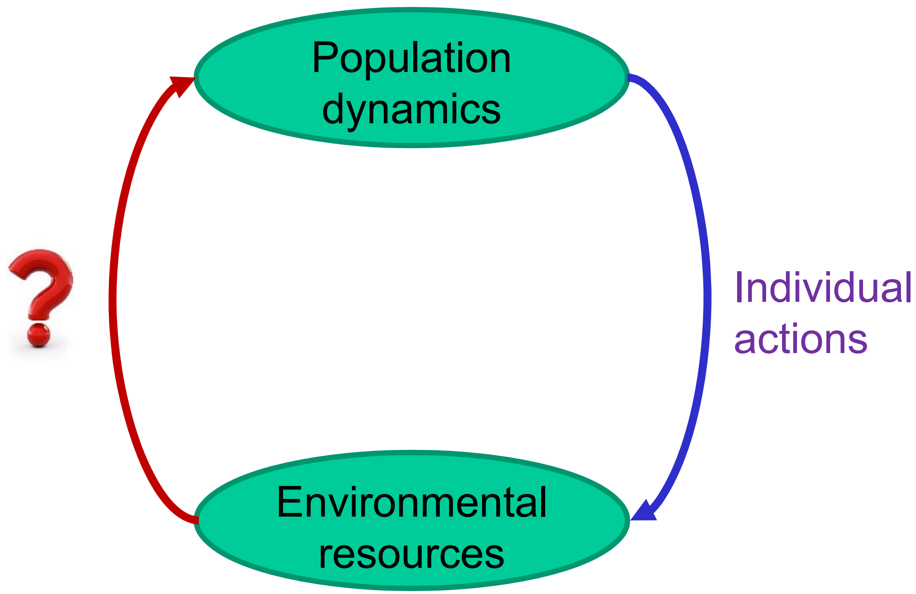

Recent report:

- Changes in the environment modulate the payoffs, leading to environment-dependent payoff matrices

$$\begin{array}{c} \\ C \\ D \end{array} \begin{array}{c} \begin{array}{cc} C & D \end{array} \\ \begin{pmatrix} a_{11}(r) & a_{12}(r) \\ a_{21}(r) & a_{22}(r) \end{pmatrix} \end{array}$$

Replicator dynamics with game-environment feedback for a single population [Weitz, 2016, PNAS].

Co-Evolutionary game dynamics

For multiple populations of individuals, how will the population dynamics *co-evolve* with the environmental dynamics?

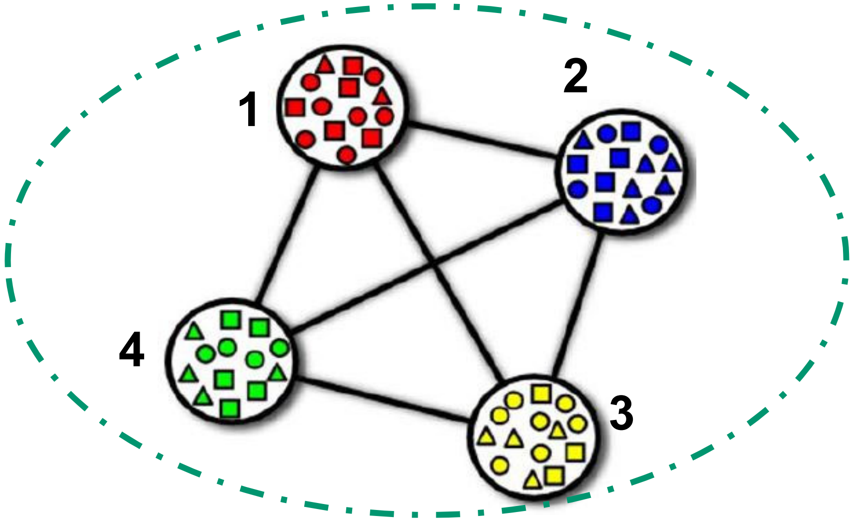

environment

This results in two layers of coevolution:

1. coevolution of multiple populations of individuals;
2. coevolution of populations and the environment.

Outline

- Evolutionary game model: replicator dynamics
- Two-population dynamics with environmental feedback
- Controlling evolutionary network dynamics using incentives

Two-population replicator dynamics

- Replicator dynamics for the matrix game with strategies $\{e_1, ..., e_m\}$

$$\dot{x}_i = x_i(\pi(e_i, x) - \pi(x, x)),$$

- Two-population replicator dynamics

$$\dot{p}_i^k = p_i^k[U_i^k(\mathbf{p}) - \bar{U}^k(\mathbf{p})], \qquad k = 1, 2$$

where p_i^k is the proportion of individuals in population k using s_i.

- We use x and y to denote the states of the two populations.

Payoff matrices with environmental feedback

Individuals from population 1 interact with individuals from population 2, and vise versa.

Dynamic payoff matrices

$$A(r)_{12} = \begin{bmatrix} a_{11}r + b_{11} & \dots & a_{1m}r + b_{1m} \\ \vdots & \ddots & \vdots \\ a_{m1}r + b_{m1} & \dots & a_{mm}r + b_{mm} \end{bmatrix}$$

$$A(r)_{21} = \begin{bmatrix} c_{11}r + d_{11} & \dots & c_{1m}r + d_{1m} \\ \vdots & \ddots & \vdots \\ c_{m1}r + d_{m1} & \dots & c_{mm}r + d_{mm} \end{bmatrix}$$

Dynamics of the environment

$$\dot{r} = r(1-r)h(x,y)$$

where $h(x,y)$ denotes the impact of population states on the environment, which may enhance or degrade resources.

Co-evolutionary game dynamics model

Consider

$$h(x,y) = \sum_{i\in\mathcal{S}} \mu_i x_i - \sum x_{-i} + \sum_{j\in\mathcal{S}} \rho_j y_j - \sum y_{-j}$$

with $\mu_i > 0,\ \rho_j > 0$ representing the ratios of enhancement to degradation.

Replicator dynamics with environmental feedback:

$$\Sigma : \begin{cases} \dot{x}_i = x_i[(A(r)_{12}\mathbf{y})_i - \mathbf{x}^T A(r)_{12}\mathbf{y}] \\ \dot{y}_j = y_j[(A(r)_{21}\mathbf{x})_j - \mathbf{y}^T A(r)_{21}\mathbf{x}] \\ \dot{r} = r(1-r)[\sum_{i,j\in\mathcal{S}}(\mu_i x_i + \rho_j y_j) - \sum(x_{-i} + y_{-j})], \end{cases}$$

where $(x_i, y_j, r) \in \Omega :\ \Delta^{m-1} \times \Delta^{m-1} \times I_{[0,1]}$.

Specific payoff matrices

Consider the following asymmetric payoff matrices

$$A(r) = (1-r)\begin{bmatrix} T_1 & P_1 \\ R_1 & S_1 \end{bmatrix} + r\begin{bmatrix} R_1 & S_1 \\ T_1 & P_1 \end{bmatrix}$$

$$B(r) = (1-r)\begin{bmatrix} T_2 & P_2 \\ R_2 & S_2 \end{bmatrix} + r\begin{bmatrix} R_2 & S_2 \\ T_2 & P_2 \end{bmatrix}$$

with $P_1 > S_1,\ T_1 > R_1;\ P_2 > S_2,\ T_2 > R_2$.

Each matrix has an embedded symmetry to ensure that mutual cooperation is a Nash equilibrium when $r = 0$ and mutual defection is a Nash equilibrium when $r = 1$.

Then the dynamics become

$$\Sigma_1 : \begin{cases} \dot{x} = x(1-x)[\delta_{PS_1} + (\delta_{TR_1} - \delta_{PS_1})y](1-2r) \\ \dot{y} = y(1-y)[\delta_{PS_2} + (\delta_{TR_2} - \delta_{PS_2})x](1-2r) \\ \dot{r} = r(1-r)[(1+\theta_1)x + (1+\theta_2)y - 2] \end{cases}$$

with $\delta_{PS_i} = P_i - S_i > 0,\ \delta_{TR_i} = T_i - R_i > 0$ and θ_k is the enhancement to degradation ratio in population *k*.

Observation from simple computations

- Invariant cubic domain $I^3_{[0,1]} = [0,1]^3$
- eight isolated fixed points and one set of interior fixed points

$$\left\{(x,y,r) : (1+\theta_1)x + (1+\theta_2)y = 2, r = \frac{1}{2}\right\}$$

- The eight corner fixed points are unstable;
- The eigenvalues at the interior equilibrium are

$$\lambda_1 = 0,\ \lambda_{2,3} = \pm\sqrt{K}i, K > 0$$

- simulations

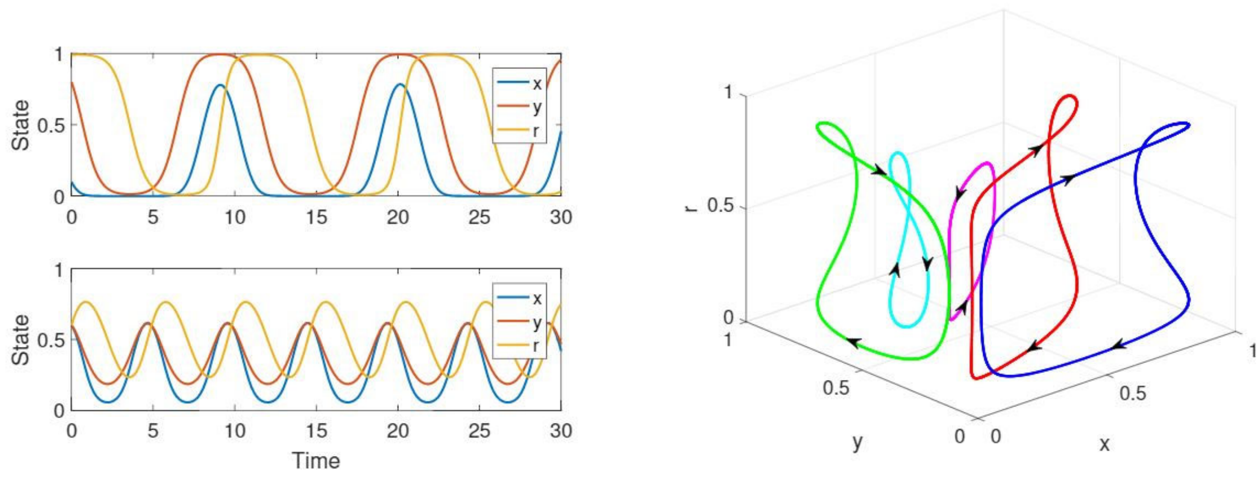

(a) Time Evolution

(b) Periodic Orbits

Main result to prove

Theorem

The two-population co-evolutionary game dynamics have infinitely many periodic orbits in the interior of $I^3_{[0,1]} = [0,1]^3$.

Viewpoint 1: Reversible system

- Reversible system

A dynamical system is said to be *reversible* if there is an involution *G* in its phase space which reverses the direction of time, i.e. the dynamics are invariant under a change in the sign of time.

- Periodic Orbits

An orbit (not a fixed point) is *periodic and symmetric* with respect to *G* if and only if the orbit intersects $\mathrm{Fix}(G)$ at precisely two points.

For system

$$\Sigma_1 : \begin{cases} \dot{x} = x(1-x)[\delta_{PS_1} + (\delta_{TR_1} - \delta_{PS_1})y](1-2r) \\ \dot{y} = y(1-y)[\delta_{PS_2} + (\delta_{TR_2} - \delta_{PS_2})x](1-2r) \\ \dot{r} = r(1-r)[(1+\theta_1)x + (1+\theta_2)y - 2] \end{cases}$$

one can find

$$G : x \to x, y \to y, r \to 1-r$$

$$Fix(G) : \quad \{r = 1/2\}$$

such that

$$\Sigma_1 \overset{G}{\Rightarrow} \tilde{\Sigma}_1 \overset{G}{\Rightarrow} \Sigma_1$$

Viewpoint 1: Reversible system

Divide $I^3_{(0,1)} = (0,1)^3$ into four regions by the two planes

$$\{r = 1/2\}$$

$$\{(1+\theta_1)x + (1+\theta_2)y = 2\}$$

$$\Omega_1 : \{\frac{1}{2} < r < 1, (1+\theta_1)x + (1+\theta_2)y > 2\}$$

$$\Omega_2 : \{\frac{1}{2} < r < 1, (1+\theta_1)x + (1+\theta_2)y < 2\}$$

$$\Omega_3 : \{0 < r < \frac{1}{2}, (1+\theta_1)x + (1+\theta_2)y < 2\}$$

$$\Omega_4 : \{0 < r < \frac{1}{2}, (1+\theta_1)x + (1+\theta_2)y > 2\}$$

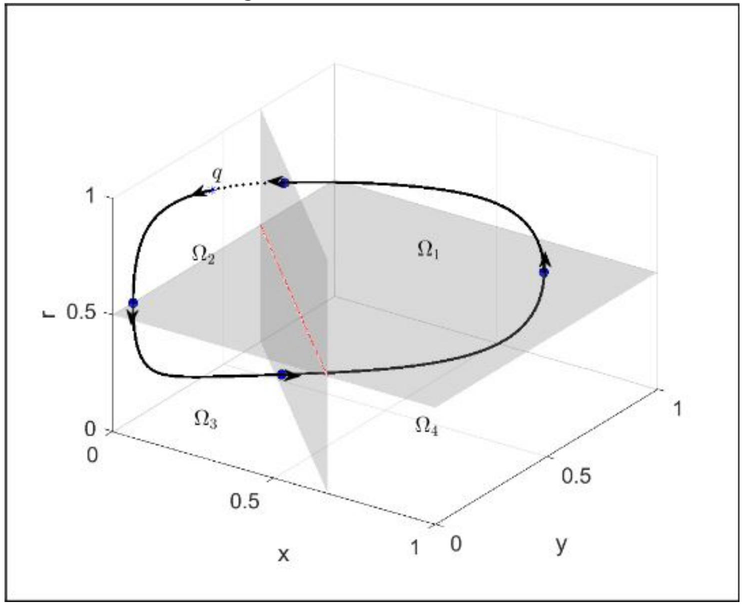

A typical trajectory

Consider an arbitrary trajectory starting from point *q*;

it goes across the plane $\{r = 1/2\}$ and into Ω_3;

crosses plane $\{(1+\theta_1)x + (1+\theta_2)y = 2\}$

and enters Ω_4;

then crosses plane $\{r = 1/2\}$ again;

returns to the starting point and forms a periodic orbit.

Viewpoint 2: Hamiltonian theory

For system Σ_1 , apply a change of variable

$$a = \varphi_2(y) - \varphi_1(x)$$

where $\varphi_1(x) = \ln(x) - \frac{\delta_{TR_1}}{\delta_{PS_1}} \ln(1-x);\ \varphi_1(y) = \ln(y) - \frac{\delta_{TR_2}}{\delta_{PS_2}} \ln(1-y)$

Apply globally real-analytic diffeomorphism

$$\phi : (x, y, r) \in I^3_{(0,1)} \to (z, a, r) \in I^2_{(0,1)} \times \mathbf{R}$$

arrive at

$$\Sigma_2 : \begin{cases} \Sigma_a : \begin{cases} \dot{z} = z(1-z)[\delta_{PS_1} + (\delta_{TR_1} - \delta_{PS_1})\varphi_2^{-1}(a + \varphi_1(z))](1-2r) \\ \dot{r} = r(1-r)[(1+\theta_1)z + (1+\theta_2)\varphi_2^{-1}(a+\varphi_1(z)) - 2] \end{cases} \\ \dot{a} = 0 \end{cases}$$

Viewpoint 2: Hamiltonian theory

Constant of Motion

For each a, there exists such a Hamiltonian function for system Σ_a

$$H_a(z,r) = H_{z,a}(z) + H_r(r)$$

$$H_r(r) = \delta_{PS_1}(\ln(r) + \ln(1-r))$$

$$H_{z,a}(z) = \int_{1/2}^{z} -\frac{(1+\theta_1)\tau + (1+\theta_2)\varphi_2^{-1}(a+\varphi_1(\tau)) - 2}{\tau(1-\tau)(\delta_{PS_1} + (\delta_{TR_1} - \delta_{PS_1})\varphi_2^{-1}(a+\varphi_1(\tau)))} d\tau$$

satisfying

$$\dot{H}_a(z,r) = \frac{\partial H_a}{\partial z}\dot{z} + \frac{\partial H_a}{\partial r}\dot{r} = 0$$

Level Sets and Closed Orbits

$$\Sigma_2 : \begin{cases} \Sigma_a : \begin{cases} \dot{z} = z(1-z)[\delta_{PS_1} + (\delta_{TR_1} - \delta_{PS_1})\varphi_2^{-1}(a+\varphi_1(z))](1-2r) \\ \dot{r} = r(1-r)[(1+\theta_1)z + (1+\theta_2)\varphi_2^{-1}(a+\varphi_1(z)) - 2] \end{cases} \\ \dot{a} = 0 \end{cases}$$

Viewpoint 2: Hamiltonian theory

Constant of Motion

For each a, there exists such a Hamiltonian function for system Σ_a

$$H_a(z,r) = H_{z,a}(z) + H_r(r)$$

$$H_r(r) = \delta_{PS_1}(\ln(r) + \ln(1-r))$$

$$H_{z,a}(z) = \int_{1/2}^{z} -\frac{(1+\theta_1)\tau + (1+\theta_2)\varphi_2^{-1}(a+\varphi_1(\tau)) - 2}{\tau(1-\tau)(\delta_{PS_1} + (\delta_{TR_1} - \delta_{PS_1})\varphi_2^{-1}(a+\varphi_1(\tau)))} d\tau$$

satisfying

$$\dot{H}_a(z,r) = \frac{\partial H_a}{\partial z}\dot{z} + \frac{\partial H_a}{\partial r}\dot{r} = 0$$

Level Sets and Closed Orbits

Each level set of the constant of motion is a periodic orbit except for the unique maximum.

Go back from Σ_2 to Σ_1

Since ϕ is a globally real-analytic diffeomorphism, the above result is applicable to the original coordinates.

Viewpoint 3: Volume preserving

For more general payoff matrices and multiple strategies, the periodic orbits may still exist.

The generalized system

$$\Sigma : \begin{cases} \dot{x}_i = x_i[(A(r)_{12}\mathbf{y})_i - \mathbf{x}^T A(r)_{12}\mathbf{y}] \\ \dot{y}_j = y_j[(A(r)_{21}\mathbf{x})_j - \mathbf{y}^T A(r)_{21}\mathbf{x}] \\ \dot{r} = r(1-r)[\sum_{i,j\in S_1}(\mu_i x_i + \rho_j y_j) - \sum(x_{-i} + y_{-j})], \end{cases}$$

with $(x_i, y_j, r) \in \Omega : \ \Delta^{m-1} \times \Delta^{m-1} \times I_{[0,1]}$.

Consider separately fixed points on the boundary and in the interior .

Viewpoint 3: Volume preserving

Introduce $\tilde{r} = 1 - r$

$$\Sigma' = \begin{cases} \dot{x}_i = x_i[(A(r)_{12}\mathbf{y})_i - \mathbf{x}^T A(r)_{12}\mathbf{y}] \\ \dot{y}_j = y_j[(A(r)_{21}\mathbf{x})_j - \mathbf{y}^T A(r)_{21}\mathbf{x}] \\ \begin{cases} \dot{r} = r\tilde{r}[\sum_{i,j\in S_1}(\mu_i x_i + \rho_j y_j) - \sum(x_{-i} + y_{-j})] \\ \dot{\tilde{r}} = r\tilde{r}[\sum_{i,j\in S_1}(\sum(x_{-i} + y_{-j}) - \mu_i x_i + \rho_j y_j)] \end{cases} \end{cases}$$

the state space becomes $\hat{\Omega}:\ \Delta^{m-1} \times \Delta^{m-1} \times \Delta^1$

Adjust by multiplying a positive function $f = \frac{1}{r\tilde{r}\prod_i^m x_i \prod_j^m y_j}$, $\text{int}\ \hat{\Omega} \to \mathbf{R}_+$

yields

$$\begin{cases} \dot{x}_i = \frac{x_i}{r(1-r)\prod_i^m x_i \prod_j^m y_j}[(A(r)_{12}\mathbf{y})_i - \mathbf{x}^T A(r)_{12}\mathbf{y}] \\ \dot{y}_j = \frac{y_j}{r(1-r)\prod_i^m x_i \prod_j^m y_j}[(A(r)_{21}\mathbf{x})_j - \mathbf{y}^T A(r)_{21}\mathbf{x}] \\ \dot{r} = \frac{1}{\prod_i^m x_i \prod_j^m y_j}[\sum_{i,j\in S_1}(\mu_i x_i + \rho_j y_j) - \sum(x_{-i} + y_{-j})] \\ \dot{\tilde{r}} = \frac{1}{\prod_i^m x_i \prod_j^m y_j}[\sum(x_{-i} + y_{-j}) - \sum_{i,j\in S_1}(\mu_i x_i + \rho_j y_j)] \end{cases}$$

Record compactly

$$\dot{X} = \underbrace{\Pi(X)\overbrace{\text{diag}(X)F(X)}^{R(X)}}_{\xi(X)}$$

Viewpoint 3: Volume preserving

[Liouville theorem]

Consider $\dot{x} = \xi(x)$ defined on the open set U, if $G \subset U$ has volume V, then the volume $V(t)$ of $G(t)$ satisfies

$$\frac{d}{dt}V(t) = \int_{G(t)} \mathrm{tr}(D\xi(x))d(x_1, \ldots, x_n) = \int_{G(t)} \mathrm{div}\ \xi(x)d(x_1, \ldots, x_n).$$

Orthogonal projection matrix: $\hat{\Omega} \to T\hat{\Omega}$

$$\varPhi = \mathrm{diag}(\Phi^m, \Phi^m, \Phi^2), \quad \Phi^m = I - \frac{1}{m}\mathbf{1}_m^T\mathbf{1}_m$$

Computation of the Derivative and Divergence

$$\begin{aligned} D\xi(X) &= \Pi(X)DR(X) + R(X)\nabla\tilde{\Pi}(X)^T \\ &= \left(\Pi(X)D\tilde{R}(X) + R(X)\nabla\tilde{\Pi}(X)^T\right)\varPhi \end{aligned}$$

$$\Downarrow$$

$$\mathrm{div}\xi(X) \equiv 0$$

The volume is incompressible in $\mathrm{int}\ \hat{\Omega}$, i.e., volume-preserving.

Outline

- Evolutionary game model: replicator dynamics
- Two-population dynamics with environmental feedback
- Controlling evolutionary network dynamics using incentives

Agents' updating rule

Replicator dynamics at the population level in *continuous time* correspond to the **imitation updating rule** at the agent level in *discrete time* where agents maximize payoff against current neighbor actions.

Imitation: agents copy the action of the highest earning neighbor.

$\mathcal{S}_i^M(t)$ = set of strategies earning the maximum payoff in the neighborhood of agent i:

$$\mathcal{S}_i^M(t) = \left\{ x_j(t) \,\middle|\, y_j(t) = \max_{k \in \mathcal{N}_i \cup \{i\}} y_k(t) \right\}.$$

$$x_i(t+1) = \begin{cases} A & \mathcal{S}_i^M(t) = \{A\} \\ B & \mathcal{S}_i^M(t) = \{B\} \\ x_i(t) & \mathcal{S}_i^M(t) = \{A, B\} \end{cases}$$

Incentive-based control of A-coordinating networks

Suppose we can offer an incentive r for taking a particular action.

$$\begin{array}{c} \\ A \\ B \end{array} \begin{array}{c} \begin{array}{cc} A & \quad B \end{array} \\ \begin{pmatrix} a+r & b+r \\ c & d \end{pmatrix} \end{array}, \qquad a, b, c, d, r \in \mathbb{R}$$

How much would it cost to have all agents converge to A?

Cases:

- Uniform incentives
- Targeted incentives
- Targeted incentives subject to a budget constraint

Uniform incentive-based control

All agents receive the same incentive

$$\begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} a_i + r_0 & b_i + r_0 \\ c_i & d_i \end{pmatrix}, \end{array} \qquad a_i, b_i, c_i, d_i \in \mathbb{R}$$

Find the minimum value of the uniform incentive such that the entire network converges to *A*?

- ***A*-coordinating**: any agent who updates to Strategy *A* would also do so if some agents currently playing B were instead playing *A*
- ***A*-monotone**: Offering incentives to play *A* will never lead to an agent to switch away from *A*
- **Uniquely-convergent**: Offering incentives leads to a unique equilibrium

Theorem: Every network of *A*-coordinating agents is *A*-monotone and uniquely convergent.

Uniform incentive-based control

Proposition:
One can construct a finite set ***R*** that contains *r**

Because of the *A*-monotone property, one can carry out the binary search:

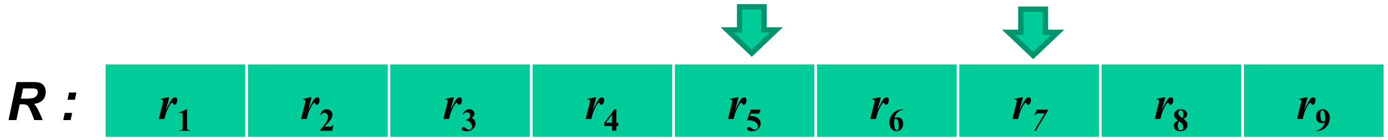

Theorem:
Within finite steps, binary search solves the uniform reward problem

Targeted incentive-based control

Suppose it's possible to offer different rewards to individual agents:

$$\begin{array}{c} \\ A \\ B \end{array}\begin{array}{c} \begin{array}{cc} A & B \end{array} \\ \begin{pmatrix} a_i + r_i & b_i + r_i \\ c_i & d_i \end{pmatrix} \end{array}, \quad a_i, b_i, c_i, d_i \in \mathbb{R}, \;\; r_i \in \mathbb{R}_{\geq 0}$$

Problem 1: Find $\mathbf{r} = (r_1, \ldots, r_n)$ that minimizes $\sum_i \mathbf{r}$ such that the entire network converges to A.

Problem 2 (budget constraint): Find $\mathbf{r}$ that maximizes the number of agents who converge to A subject to $\sum_i \mathbf{r} \leq \rho$.

Targeted incentive-based control

Computationally complex to solve exactly (conjectured to be NP)

We can compute the incentive $\check{r}_i$ needed such that at least one A-neighbor will switch to B

$$\check{r}_i = \max_{j \in \mathcal{N}_i^B} \max_{k \in \mathcal{N}_j^B} y_k - y_i,$$

where $\mathcal{N}_i^B := \{j \in \mathcal{N}_i \cup \{i\} : x_j = B\}$.

Algorithm: Iteratively choose agents to switch until the desired equilibrium is reached or the budget limit is exceeded.

Targeted incentive-based control

How should we choose these agents?

Several possibilities: max degree, min required incentives, etc.

Approach: Iteratively maximize a benefit-to-cost ratio

Benefit = # of agents who switch to *A*, cost = incentive

$$\max_i \frac{\Delta\Phi(x)^\alpha}{\check{r}_i^\beta}, \text{ where } \Delta\Phi(x) = \Phi(x(t_2)) - \Phi(x(t_1)),$$

$$\Phi(x) = \sum_{i=1}^{n} n_i^A(x), \qquad \alpha \text{ and } \beta \text{ are design parameters.}$$

Simulation results: Uniform vs. Targeted incentives

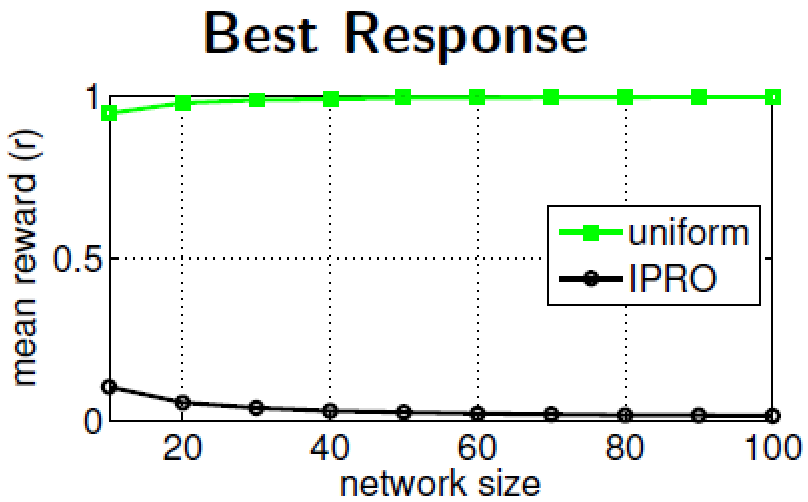

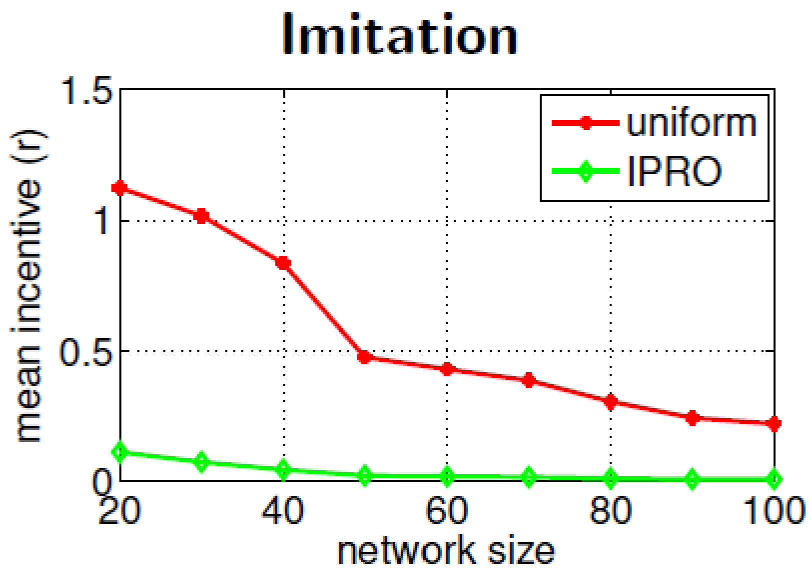

Outline

- Evolutionary game model: replicator dynamics
- Two-population dynamics with environmental feedback
- Controlling evolutionary network dynamics using incentives

Conclusions and outlook

Conclusions:

1) Use a generalized co-evolutionary model to study population game dynamics with environmental feedback;
2) Analysis of the model under specific payoff matrices;
3) Different methods to prove the existence of periodic orbits;
4) Volume-preserving property is disclosed;
5) Incentives help to drive coordinating networks towards desired equilibria
6) Targeted incentives are more efficient than uniform incentives

Outlook:

- Extensions to different types of games
- A general framework for controlling evolutionary games using environmental feedback

My research team on this topic

Former post docs and PhD students

James Riehl

Pouria Ramazi

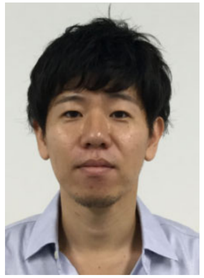

Yu Kawano

Current post docs and PhD students

Alain Govaert

Luke Gong

Some selected recent publications from my group on related topics

"Incentive-based control of asynchronous best-response dynamics on binary decision networks," J. R. Riehl, P. Ramazi, and M. Cao*. IEEE Trans. on Control of Network Systems*, 2019

"Evolutionary dynamics of two communities under environmental feedback," Y. Kawano, L. Gong, B. D. O. Anderson, and M. Cao*. IEEE Control Systems Letters*, special issue on Control and Network Theory for Biological Systems, 2019

"Evolutionary game dynamics of two interacting populations under in a co-evolving environment," L. Gong, J. Gao, and M. Cao*. Proc. IEEE CDC*, 2018

"A survey on the analysis and control of evolutionary matrix games," J. R. Riehl, P. Ramazi and M. Cao*. Annual Reviews in Control*, 45(6), 87-106, 2018

"Asynchronous decision-making dynamics under best-response update rule in finite heterogeneous populations," P. Ramazi and M. Cao*. IEEE Transactions on Automatic Control*, 63(3), 742-751, 2018.

"Networks of conforming or nonconforming individuals tend to reach satisfactory decisions," P. Ramazi, J. R. Riehl, and M. Cao*. Proceedings of the National Academy of Science of USA (PNAS)*, 113(46), pp12985-12990, 2016

--- The end ---