Introduction to Evolutionary Game Theory

Marco Tomassini

University of Lausanne
Lausanne, Switzerland
marco.tomassini@unil.ch

http://www.sigevo.org/gecco-2014/

GECCO'14, July 12-16, 2014, Vancouver, BC, Canada.
ACM 978-1-4503-2881-4/14/07.

http://dx.doi.org/10.1145/2598394.2605363

UNIL | Université de Lausanne

Contents

UNIL | Université de Lausanne

Standard Game Theory

- A mathematical theory of decision under conflicting situations
- A player's decision depends on the other players' decisions and viceversa
- The theory postulates that the players are intelligent rational agents

UNIL | Université de Lausanne

Intelligent and Rational Agents

Such an agent (player) must be able to:

- Determine the set of possible actions
- Know how consequences are related to a given action
- Sort the consequences according to a value scale: utility or payoff
- Select the action that guarantees utility maximization

UNIL | Université de Lausanne

Rational Players

- Players are expected utility maximizers
- Players have common knowledge of each other rationality. A fact is common knowledge if every player knows it, every player knows that every player knows it ...
- Pre-play communication between players has no effect on the outcome: everything works as if players played the game simultaneously and independently (for normal form games)

UNIL | Université de Lausanne

The Subject of this Tutorial

here we only deal with complete information games in which each player is a selfish utility maximizer: Non-Cooperative Games

Other important subjects in game theory not treated here:

- Cooperative games in which players may form coalitions
- Incomplete information games (Bayesian Games)
- Iterated games: games played repeatedly between the same players with memory of the past interactions

UNIL | Université de Lausanne

Games and Social Dilemmas: Prisoners and Public Goods

Two strategies are available to the players: either make a donation of 5$ to the other player or make no donation (0$). If they both make a donation, the amount that each one receives is 5+5=10$ (10$ offered by a third party). If a player makes no donation and the other does, the first gets 15$ while the donor loses 5$. If neither donates, they both get 0$ (K. Sigmund).

Prisoner's Dilemma

	makes a donation	*makes no donation*
makes a donation	$10, 10$	$-5, \mathbf{15}$
makes no donation	$\mathbf{15}, -5$	$\mathbf{0}, \mathbf{0}$

UNIL | Université de Lausanne

Normal or Strategic Form

- Complete information game Γ:

$$\Gamma = (N, C_i, u_i), \ \forall i \in N,$$

where:

N is the finite set of players,

C_i is the ensemble of Strategies available to player i,

$C = \times_{j \in N} C_j$: the set of possible strategy profiles,

and $u_i : C \to \mathbb{R}$ is the utility of player i

UNIL | Université de Lausanne

Two-Person Games: Normal Form

For a two-person game (players A and B):

$C_A = \{a_1, a_2, \ldots, a_m\}$: set of A's strategies (lines)
$C_B = \{b_1, b_2, \ldots, b_n\}$: B's strategies (columns)

	b_1	b_2	$\ldots$	b_n
a_1	$g_{1,1}$	$g_{1,2}$	$\ldots$	$g_{1,n}$
a_2	$g_{2,1}$	$g_{2,2}$	$\ldots$	$g_{2,n}$
$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$
a_m	$g_{m,1}$	$g_{m,2}$	$\ldots$	$g_{m,n}$

$g_{i,j} = u_i, u_j$ are the utilities (payoffs) to A and to B when A plays strategy i and B plays strategy j

UNIL | Université de Lausanne

Two-Person Games: Extensive Form

- This is the most complete representation: it takes time into account
- The game is represented as a tree and, at each tree node, one particular player chooses an action
- A game in extensive form has a unique corresponding normal form representation but the reverse is not true in general

UNIL | Université de Lausanne

Two-Person Games: Extensive Form

The "Matching Pennies" Game

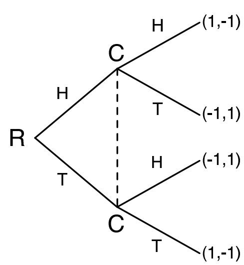

and its corresponding normal form:

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

UNIL | Université de Lausanne

Nash Equilibrium

The most important concept in standard game theory:

Theorem: Every finite game Γ in strategic form has at least one equilibrium in pure or mixed strategies

John Nash, 1951; Nobel Prize in Economy in 1994

UNIL | Université de Lausanne

Randomized Strategies

A randomized (or mixed) strategy σ_i for player i is a probability distribution $\Delta(C_i)$ over the set of "pure" strategies C_i

$\sigma \in \times_{i \in N} \Delta(C_i)$ is a randomized strategy profile for each player and for each pure strategy $c_i \in C_i$

$\sigma(c_i)$ represents the probability that player i chooses c_i with $\sum_{c_i \in C_i} \sigma(c_i) = 1$

Mixed Strategies: The Unit Simplex

$\sigma_i \in \mathbb{R}^{m_i}$ belongs to the unit *simplex*. For $m = 2$ and $m = 3$ the unit simplex looks like this:

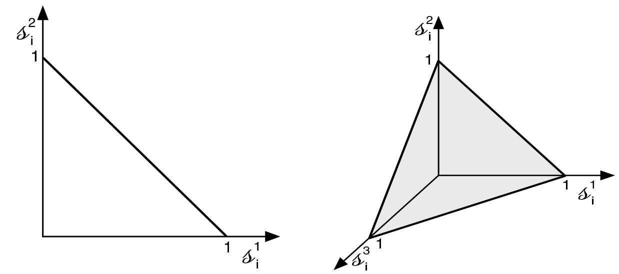

Nash Equilibrium

Calling $u_i(\sigma)$ the expected payoff that player i would get when the players choose their strategies independently according to the strategy profile σ, a *Nash equilibrium* is such that:

$$u_i(\sigma) \geq u_i(\sigma_{-i}, \tau_i), \quad \forall i \in N, \quad \forall \tau_i \in \Delta(C_i)$$

with (σ_{-i}, τ_i) a randomized strategy profile equal to σ except for the i-th component τ_i.

Thus a randomized strategy profile is a Nash equilibrium iff no player could increase her expected payoff by unilaterally deviating from this strategy profile

Coordination: Rousseau and Stag Hunt

Two people are needed to hunt a stag. A single hunter can hunt a rabbit alone

Stag Hunt

	Stag	*Rabbit*
Stag	**3, 3**	0, 2
Rabbit	2, 0	**1, 1**

(Stag,Stag) is the Pareto-optimal outcome

(Rabbit,Rabbit) is the risk-efficient outcome

Both are NE of the game (plus a third unstable NE in mixed strategies)

Heads and Tails

heads and tails

	H	T
H	$\mathbf{1}, -1$	$-1, \mathbf{1}$
T	$-1, \mathbf{1}$	$\mathbf{1}, -1$

There is no Nash equilibrium in pure strategies. There is only one Nash equilibrium in mixed strategies in which H and T are played with probability 1/2 each

Suppose the column player plays H with probability p and T with probability $1 - p$. Then:

$E_{row}[H] = p - (1 - p) = 2p - 1$
$E_{row}[T] = -p + (1 - p) = -2p + 1$
$E_{row}[H] = E_{row}[T] \implies 2p - 1 = -2p + 1 \implies p = 1/2$

And the same result is obtained for the column player

UNIL | Université de Lausanne

N-person Games: Public Goods

This kind of dilemma appears when more than two actors interact in team behaviors such as fishing, securing mutual defence, disposal of polluting wastes, trade unions, ...

The social dilemma appears because people cannot be excluded from enjoying the public good but they can choose not to play fair and try to free-ride on others

UNIL | Université de Lausanne

Tragedy of the Commons

- a grazing land open to all
- each farmer will keep as many cattle as it is possible without depleting the field too quickly
- "what is the utility to me of adding one more animal?" say it is 1
- the additional overgrazing is shared by all: it is small if only one farmer adds an animal
- if one more animal does the trick, why not adding another one?
- however, if they are rational, all the farmers will reason in the same way
- as a result, the common is ruined

Garret Hardin, "The tragedy of the commons", Science, 162, 1243-1248, 1968

UNIL | Université de Lausanne

N-person Games: Public Goods NE

N individuals where each individual can be a cooperator or a defector. Cooperators contribute a cost $c > 0$ each to the public good, defectors do not contribute.

The common fund grows by an amount $b > c$ for each contribution and at the end the contributions are shared equally among the N players; we assume $c > b/N$.
if all contribute, the total fund will be Nb and each individual gets a payoff of $Nb/N - c = b - c > 0$.

But now free riders enter the picture: if N_c is the number of those who are contributing, a player gets $b(N_c + 1)/N = bN_c/N + b/N - c$ if she contributes, while she gets bN_c/N if she doesn't; but since $c > b/N$, it is rational for her not to contribute.

Thus, although it would be socially efficient for everybody to contribute, the NE dictates that nobody contributes and the public good is not provided, or the common resource is depleted.

UNIL | Université de Lausanne

Problems with Nash Equilibria

- Many games have more than one equilibrium. How to choose, given that they are often non-equivalent from a social or economical point of view?
- Full rationality and its common knowledge among players is a very demanding assumption
- The theory lacks a dynamical approach: how do agents reach these equilibria?
- ...

Evolutionary Game Theory

John Maynard Smith (1920-2004)
"Evolution and the Theory of Games", Cambridge University Press, 1982

Selection and reproduction of the fittest is the key idea

Evolutionary Processes

An evolutionary system must possess the following fundamental elements:

- A population of individuals
- a source of variation that provides diversity through, say, mutations and recombinations of genetic material, and
- a selection mechanism that favors fitter variants over others that are less adapted to the current environment

Evolutionary Games

- A very large population of players
- Randomly paired individuals play the game and are replaced for the next run
- Players have no identity, they are anonymous
- A player is "programmed" to play a given strategy
- A player need not be intelligent and rational

Evolutionary Stability

Let x be the incumbent strategy, i.e. the strategy currently played by all individuals in the population and y be the mutant strategy

One way of defining evolutionary stability is as follows:

- $u(x,x) > u(y,x) \ \ \forall y,$
- $u(x,x) = u(y,x) \Rightarrow u(x,y) > u(y,y) \ \ \forall y \neq x$

i.e., x is evolutionarily stable if either x is a strict best reply to any y, or it is as good against itself as any other mutant, and x is a better reply to any mutant y than y is to itself.

The important conclusion is that $\Delta^{ESS} \subset \Delta^{NE}$, which means that some Nash equilibria may not be an ESS

UNIL | Université de Lausanne

Hawks and Doves

Also known as the "chicken" (J. Dean movie) or the "snowdrift" game. It's a metaphor for "arm races" and other "bullying" games, also in the animal kingdom

	H	D
H	-2,-2	**2,0**
D	**0,2**	1,1

In this game there are two Nash equilibria in pure strategies ((H,D) and (D,H)), and a third eq. in mixed strategies (play H with probability 1/3 and D with 2/3)

UNIL | Université de Lausanne

Evolutionarily Stable Strategies in the Hawk-Dove Game

In the Hawk-Dove game, neither H nor D are ESS, as they can be invaded by players playing the other strategy since:

$u(H,H) < u(D,H)$ and $u(D,D) < u(H,D)$

	H	D
H	-2,-2	**2,0**
D	**0, 2**	1,1

The only ESS is the mixed strategy equilibrium: this corresponds to a population that stabilises itself with a proportion of 1/3 hawks and 2/3 doves

The evolutionary approach can thus (but not always) reduce the number of Nash equilibria

UNIL | Université de Lausanne

Prisoner's Dilemma Again

	C	D
C	3,3	0,**4**
D	**4**,0	**1,1**

Since $u(D,D) > u(C,D)$, D is evolutionarily stable

Conversely, $u(C,C) < u(D,C)$ and thus C is not an ESS

The unique Nash equilibrium (D,D) is also the only ESS. Thus, evolutionary game theory does not help to better understand these situations. At least in the one-shot case

UNIL | Université de Lausanne

Replicator Dynamics

We need an explicitly dynamic foundation for the static evolutionary stability concept

We have seen that evolutionary processes have two basic elements:

- a source of **mutation** that provides diversity, and
- a **selection** mechanism that favors fitter variants

Evolutionary stability emphasizes the role of mutations, while replicator dynamics focuses on selection and, in its basic form, does not include a mutation mechanism

UNIL | Université de Lausanne

Replicator Dynamics

Given an initial distribution of strategies among the agents, the ensuing strategy frequency evolution in the population is dictated by a system of linear differential equations of the type:

$$\frac{dp_i}{dt} = \dot{p}_i = p_i(u_i - \bar{u}), \quad i = 1, \dots, m$$

where p_i is the frequency of strategy i, u_i is its expected payoff, and $\bar{u}$ is the mean payoff of the population.

Thus, strategies that are better than the average will increase their share in the population, while inferior strategies will decrease in time

UNIL | Université de Lausanne

Replicator Dynamics Example

Consider the following symmetric 2x2 game payoff matrix:

$$A = \begin{pmatrix} 1 & 3 \\ 2 & 0 \end{pmatrix}$$

Let p_1 and p_2 be the frequencies of strategy 1 and 2 in the population, so that the strategy profile at time t is $\mathbf{p}^T = (p_1 \;\; p_2)$ (we use the symbol p_i instead of x_i for the frequency of strategy i)

The expected utility of a player using strategy 1 and 2 are, respectively: $E[1, \mathbf{p}] = p_1 + 3p_2$, and $E[2, \mathbf{p}] = 2p_1$

The average fitness of the population is
$E[\mathbf{p}, \mathbf{p}] = (p_1 \;\; p_2)^T A \begin{pmatrix} p_1 \\ p_2 \end{pmatrix} = p_1^2 + 5p_1p_2$

UNIL | Université de Lausanne

Replicator Dynamics Example Continued

We can now write the replicator equations:

$$\begin{aligned} \frac{dp_1(t)}{dt} &= p_1([p_1 + 3p_2 - (p_1^2 + 5p_1p_2)] \\ \frac{dp_2(t)}{dt} &= p_2([2p_1 - (p_1^2 + 5p_1p_2)] \end{aligned}$$

whose fixed points are given by $dp_i(t)/dt = 0$, which implies:

$p_1([p_1 + 3p_2 - (p_1^2 + 5p_1p_2)] = 0$ and $p_2([2p_1 - (p_1^2 + 5p_1p_2)] = 0$

from these, one can check that the solutions are:

$(p_1 = 0, p_2 = 1)$, $(p_1 = 1, p_2 = 0)$, and $(p_1 = 3/4, p_2 = 1/4)$

UNIL | Université de Lausanne

Replicator Dynamics Example Ctnd.

The first two solutions are trivial; they say that if only one strategy is nothing will change. They are represented by the vertices of the simplex.

The non-trivial solution $P = (p_1 = 3/4, p_2 = 1/4)$ is a stable attractor: no matter where the system starts inside the square, the dynamics will be attracted to the point P.

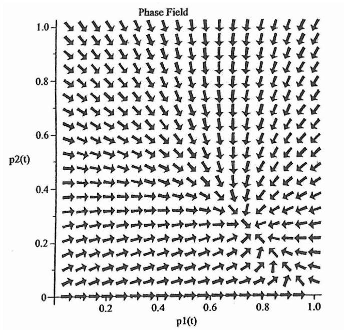

UNIL | Université de Lausanne

Replicator Dynamics Example Cntd.

Since $p_1 = 1 - p_2$, let's call $p = p_1, p_2 = 1 - p$; the system is actually one-dimensional:

$$\frac{dp(t)}{dt} = p(t)\,(1 - p(t))\,(3 - 4p(t))$$

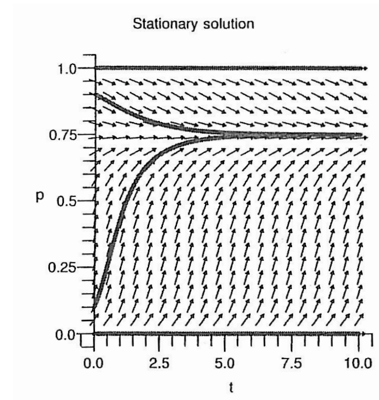

Figures redrawn from E. N. Barron, "Game Theory"

UNIL | Université de Lausanne

Replicator Dynamics Main Results

- under replicator (or imitation) dynamics only selection is active; there are no mutations
- players are only programmed to play pure strategies
- a mixed strategy is now represented by the equilibrium share of pure strategies in the population
- stationary stable states of the replicator dynamics may correspond to ESS

UNIL | Université de Lausanne

Some Problems with EGT

- An ESS is a refinement of NE but An ESS may not exist for some games
- RD may have equilibria that are not NE
- Fitness is well adapted to the animal kingdom; what does it mean is socio-economic situations?

UNIL | Université de Lausanne

Non-Existence of an ESS

Rock–Scissors–Paper game

	R	S	P
R	0,0	**1**,-1	-1,**1**
S	-1,**1**	0,0	**1**,-1
P	**1**,-1	-1,**1**	0, 0

There is a unique NE $\sigma^* = (1/3, 1/3, 1/3)$.

Take $\sigma' = (1, 0, 0)$, then $\sigma^*.A\,\sigma' = \sigma'.A\,\sigma' = 0$, while $\sigma'.A\,\sigma^* = \sigma^*.A\,\sigma^* = 0$, which violates one of the conditions for σ^* to be an ESS

UNIL | Université de Lausanne

Relationships Between RD, ESSd, and NE

- Every NE is a fixed point of the RD (and thus, since $\Delta^{ESS} \subset \Delta^{NE}$, all ESS are among the RD fixed points)
- But the RD may have attractors that are not NE; for example, all the vertices of the strategy simplex are fixed points of the RD, but not necessarily NE
- There can also be nonmonomorphic fixed points that do not correspond to NE (see Weibull's book)
- In conclusion, instead of providing a refinement of NE, the RD gives an **extension**

UNIL | Université de Lausanne

But We Can Do Better

We must supplement the notion of stationarity of the dynamics with additional criteria of robustness and stability:

Asymptotic and Lyapunov Stability (Qualitative):

- Any path that starts close enough to the equilibrium converges to it (AS)
- Any path that starts sufficiently close to the equilibrium remains arbitrarily close to it (Lyapunov stability)

UNIL | Université de Lausanne

Rock-Scissors-Paper Again

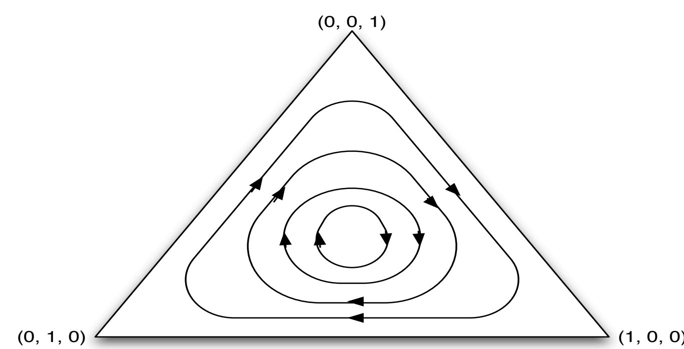

The unique NE and fixed point of the RD is not asymptotically stable but it is Lyapunov-stable: small perturbations will have small effects

UNIL | Université de Lausanne

What is Fitness in Social Contexts?

- EGT has been conceived for biological contexts in which Darwinian selection has a well established meaning; the mutation mechanisms in society are unlikely to be independent of selection
- In socio-economic settings there are additional mechanisms such as trial and error, more sophisticated learning
- Pure imitation leads to equations that are indistinguishable from standard RD; however, imitation is an extremely inflexible mechanism for social agents
- Other extensions are based on multiple population models, i.e. players may be of different types such as buyers and sellers

Where Does Cooperation Come From?

- Cooperation in a PD situation is possible, at least to some extent, as clearly shown by field behavior observation and experimental games in the laboratory
- Mechanisms that have been shown to favor the emergence of cooperation in human and animal societies (M. Nowak)
  - Direct Reciprocity
  - Indirect Reciprocity
  - Kin Selection
  - Group Selection
  - Network Selection

Different Mechanisms That May Promote Cooperation

- Direct Reciprocity : works when there are repeated encounters between the same players and they have memory of the past. The strategy Tit-For-Tat in repeated games is a well known example.
- Indirect Reciprocity : Also works if the game is iterated. My behavior depends on what you have done to me and to others. Punishment and reputation play a key role.
- Kin Selection : the games occur between individuals that are genetically related. Altruistic acts are more likely the stricter the kin relationship is.

General Form of the Paradigmatic Games

	C	D
C	R,R	S,T
D	T,S	P,P

Payoffs: R=Reward, T=Temptation, P=Punishment, S=Sucker

$T > R > P > S$: Prisoner's Dilemma (PD)

$T > R > S > P$: Hawk-Dove (HD)

$R > T > P > S$: Stag Hunt (SH)

$R > S > T > P$ or $R > T > S > P$: Harmony Games (H)

The Games' Phase Space

There is an infinite number of games respecting the previous payoff constraints. It has become standard in the field to restrict oneself to a 2-D configuration space by taking $R = 1$, $P = 0$, $0 \leq T \leq 2$, and $-1 \leq S \leq 1$.

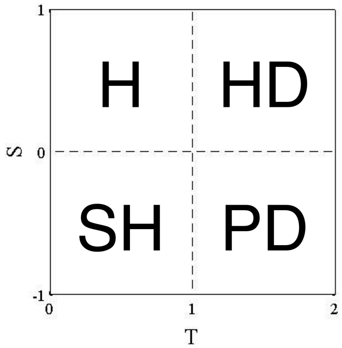

UNIL | Université de Lausanne

Replicator Dynamics Diagrams in Well Mixed Populations

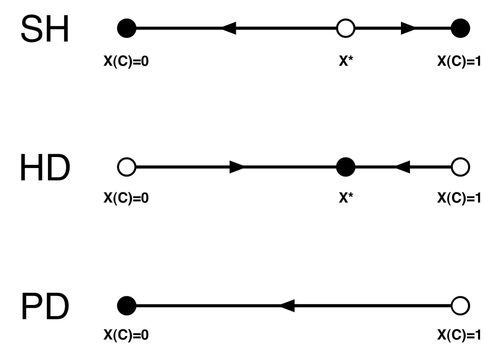

UNIL | Université de Lausanne

Theoretical Results of RD in Well Mixed Populations for the Whole Game Space

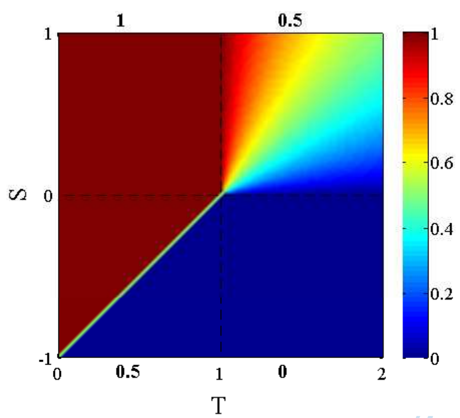

UNIL | Université de Lausanne

Structured Populations: Networks of Agents

The first approaches considered 2-D grid (Left image) structured populations (Nowak & May)

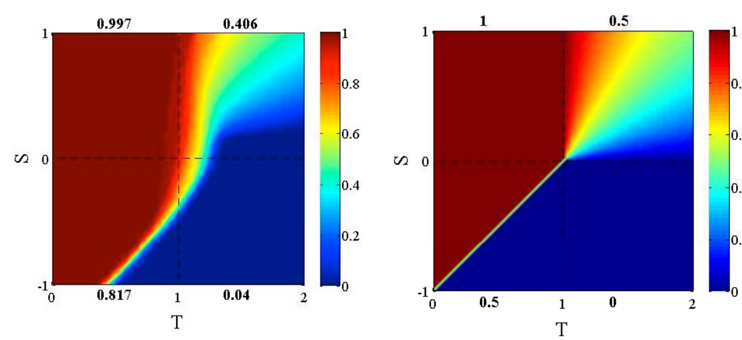

UNIL | Université de Lausanne

Stag Hunt Evolution on Grids

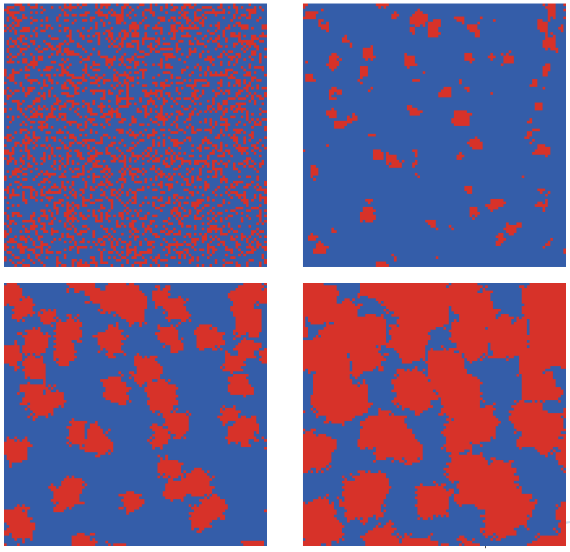

Social Networks

Model social networks may help cooperation to evolve even for the PD

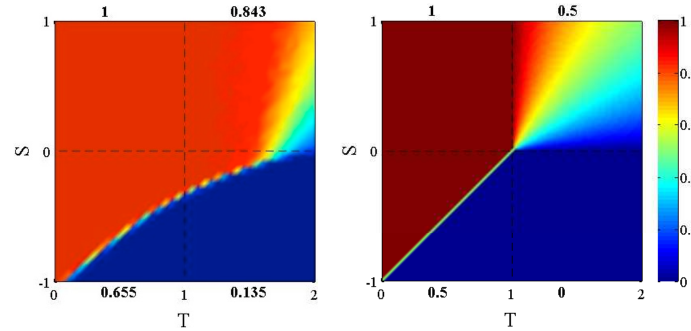

These results are encouraging as these nets are much closer to actual social networks in terms of degree distribution, clustering, assortativity and community structure

UNIL | Université de Lausanne

Dynamical Networks: Fluctuations May Suppress Cooperation

If links are subject to random noise, e.g. links are rewired at random with a certain frequency ω, keeping the degree distribution constant, cooperation becomes harder and harder with increasing ω

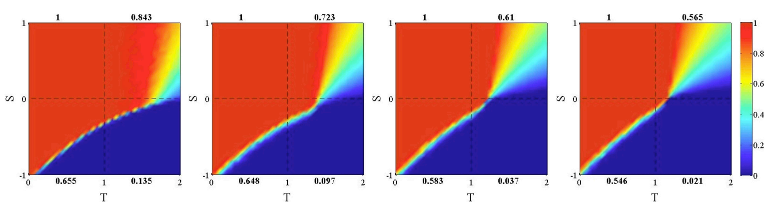

$\omega = 0$ $\omega = 0.001$ $\omega = 0.01$ $\omega = 2 \times 0.01$

UNIL | Université de Lausanne

Dynamical Networks: Purposeful Rewiring

- we use fixed size networks in which strategies and links may co-evolve
- strategy dynamics is as before; link dynamics is based on a concept of link strength being equated with a degree of player's satisfaction
- link strength evolves through a reinforcement learning process: if the interaction is beneficial in terms of payoff the corresponding link strength increases, otherwise it is reduced
- the relative speeds of strategy update and link rewiring are controlled by a parameter $0 \leq q \leq 1$ such that the whole region between rather stable and rather "fluid" networks may be simulated

UNIL | Université de Lausanne

Network Effects: Summing Up

- There appears to be a generally positive effect of network structure on cooperation and coordination, independently of other possible mechanisms that foster cooperation
- Network clustering and degree heterogeneity have the largest impact. Clustering being more important in the SH game and degree heterogeneity in HD. PD is the hardest one but it is significantly influenced too
- In social networks clustering and community structure are the key to understand the promotion of cooperation/efficient coordination

UNIL | Université de Lausanne

Network Effects Ctd.

- When link dynamics or mobility is taken into account, the topological mechanisms that help promote cooperation and coordination are perturbed and the situation tends to the well mixed case for increasing noise levels
- However, if link cutting and rewiring is done purposefully and strategically, perhaps as a result of a reinforcement learning mechanism, then cooperation and Pareto-efficient coordination are again achievable to a large extent
- The last point has been recently demonstrated through laboratory experiments with human participants

UNIL | Université de Lausanne

Where to go from here?

Good general books, both for standard game theory and evolutionary GT:

F. Vega-Redondo: Economics and the Theory of Games, Cambridge, 2003

E. N. Barron: Game theory: An Introduction, J. Wiley, 2013

More advanced books on Evolutionary Game Theory:

J.Weibull, Evolutionary Game Theory, MIT Press, Cambridge MA, 1995.

J. Hofbauer and K. Sigmund: Evolutionary Games and Population Dynamics, Cambridge, 1998.

UNIL | Université de Lausanne