Evolutionary Games

The Darwin Connection

Marco Tomassini

marco.tomassini@unil.ch

University of Lausanne, Switzerland

GECCO'07, July 7-11, 2007, London, England, United Kingdom.

ACM 978-1-59593-698-1/07/0007.

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

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
- Select the action that guarantees value maximization

Basic Terminology

- *N-person games*: games with N participants
- *Two-person games*: two participants, the case treated in these lectures
- *Cooperative* games: in which players can collaborate and form coalitions to mutual advantage
- *Non-Cooperative games*: players are egoistic utility maximizers, the case treated in these lectures
- *one-shot game*: a game that is played only once, the case treated here
- *repeated or iterated game*: a game that is played more than once between the same players

To Recap: Rational Players

- Players are expected utility maximizers
- 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
- This allows a rigorous mathematical treatment but does not necessarily correspond to actual decision-making processes

Extended Form of a Game

A simple two-person game: Each player puts a dollar in the pot. Player 1 then takes a card at random from a deck (red or black equally likely). Player 1 looks at the card privately, and decides whether to *raise* or *fold*. If he folds, he shows the card and wins a dollar if the card is red, otherwise player 2 wins. If player 1 raises then he puts another dollar in the pot. Now player two must decide whether to *meet* or *pass*. If she passes, then the game ends with 1 taking all the money. If she meets, then she must add a dollar too.

Extended Form of the Game

Player 1 now shows the card, and he takes all the money if the card is red, and player 2 takes the money if it is black

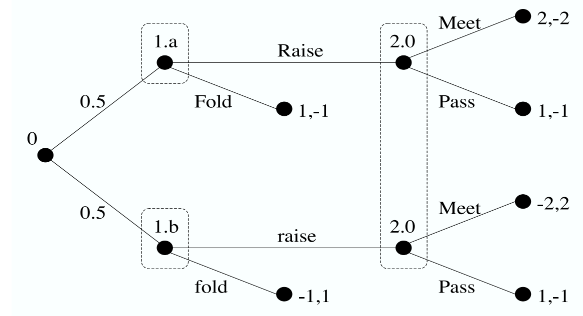

Normal or Strategic Form

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

where N is the set of players,

C_i is the ensemble of *Strategies* available to player i,

and $u_i : \times_{j \in N} C_j \to R$ is the utility of player i

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	$g_{m,1}$	$g_{m,2}$	$\ldots$	$g_{m,n}$

$g_{i,j}$ is the result of the game when A plays strategy i and B plays strategy j

Equivalence of Ext. and Normal Forms

The normal form of a game is a **static** model that ignores timing, i.e. the sequence of moves. It treats players as if they choose their strategies simultaneously.

Von Neumann and Morgenstern gave a construction whereby any finite game in extensive form Γ^e reduces uniquely to a game in normal form Γ

The Previous Game in Normal Form

	M	P
Rr	0,0	1,-1
Rf	0.5,-0.5	0,0
Fr	-0.5,0.5	1,-1
Ff	0,0	0,0

Strategies of player 1 (rows) = $\{Rr, Rf, Fr, Ff\}$

Rr: Raise on red card, raise on black card

Rf: Raise on red card, fold on black card

Fr: Fold on red card, raise on black

Ff: Fold on red card, fold on black

Strategies of player 2 (columns) = $\{M, P\}$; M= Meet, P=Pass

Nash Equilibrium

The most important concept in standard game theory:

"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

Randomized Strategies

A randomized (or mixed) strategy 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 $\sum_{c_i \in C_i} \sigma(c_i) = 1$

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

Some Classical Examples

Prisoner's Dilemma

Invented by A. Tucker in 1950 (RAND Corporation, Santa Monica, CA)

	C	D
C	(R,R)	(S,T)
D	(T,S)	(P,P)

C = cooperate, D = defect, R = Reward, T = Temptation, P = Punishment, S = Sucker's payoff
with the constraints:

T > R > P > S, R > (T + S)/2

Prisoner's Dilemma

An actual possible payoff bi-matrix for the game :

	C	D
C	(3,3)	$(0,4^*)$
D	$(^{+}4,0)$	$(^{+}1,1^*)$

A Nash equilibrium is a pair of strategies (pure or mixed) such that any other choice would be a less good reply for each player
$+$: best reply for player A
$*$: best reply for player B
The (unique) Nash equilibrium of the game is (D,D)

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) et (D,H)), and a third eq. in mixed strategies (play H with probability $1/3$ and D with $2/3$)

Can We Get out of the Dilemma?

Can cooperation emerge and remain stable? The answer is “yes” if the game is played an undefined number of times between players that have a memory of past encounters

Robert Axelrod organized world-famous PD computer tournaments; “The Evolution of Cooperation”, 1984

Problems with Nash Equilibria

- Many games have more than one equilibrium. How to choose, given that they are often non-equivalent? Equilibrium selection problem
- Some “equilibrium refinements” may eliminate “unstable” equilibria (*trembling hands*) but this does not always work
- The common knowledge of rationality is a very demanding assumption
- But: limited rationality concepts also have their problems

A Coordination Game

In the following game (sometimes called “battle of the sexes”):

	T	M
T	$^{+}3, 1^{*}$	$0, 0$
M	$0, 0$	$^{+}1, 3^{*}$

A prefers to go to the theater (T); B prefers to go to the movie (M). But both players would prefer to go out together, rather than separately.
(T,T) et (M,M) are Nash equilibria in pure strategies but they are not equivalent. There is a third Nash equilibrium in mixed strategies, which is inefficient because the players behave in a random and uncoordinated manner.

Evolutionary Game Theory

John Maynard Smith (1974-1982)
“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

Evolutionarily Stable Strategies (ESS)

ESS are strategies that cannot be invaded by a “mutant” strategy. In the Hawk-Dove game, neither H nor D are ESS, as they can be invaded by players playing the other strategy

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

Evolutionarily Stable Strategies

In a formal way, $C = \{1, 2, \ldots, k\}$ is the set of pure strategies, $\Delta(C_i) = \{x \in R^k : \sum_{i \in C} x_i = 1\}$ is the associated mixed strategy set, and the payoff to strategy $x \in \Delta(C_i)$ when played against "mutant" strategy $y \in \Delta(C_i)$ is $u(x, y)$

Let $\epsilon \in (0, 1)$ be the share of mutants in the population. Given that pairs of players are drawn from the population with uniform probability to play the game, the probability that a player will play y is ϵ, while the probability of playing x is $1 - \epsilon$; this is equivalent to playing the mixed strategy $w = \epsilon\, y + (1 - \epsilon)\, x$

Evolutionary Stability

the payoff of the "established" strategy x against w is thus $u(x, w)$ and that of the "mutant" strategy y is $u(y, w)$. Strategy x will be **evolutionarily stable** if

$$u[x, \epsilon y + (1 - \epsilon)x] > u[y, \epsilon y + (1 - \epsilon)x]$$

$\forall y \in \Delta(C_i), y \neq x$, and granted that the share of mutants ϵ is "sufficiently small"

Evolutionary Stability

Evolutionary stability can also be stated as follows:

- $u(y, x) \leq u(x, x) \;\; \forall y,$
- $u(y, x) = u(x, x) \Rightarrow u(y, y) < u(x, y) \;\; \forall y \neq x$

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

Prisoner's Dilemma Again

	C	D
C	(3,3)	(0,4^*)
D	($^+4$,0)	($^+1, 1^*$)

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 non-iterated case

Replicator Dynamics

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 does not include a mutation mechanism

Replicator Dynamics

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

$$\frac{dx_i}{dt} = x_i(u_i - \bar{u}),$$

where x_i is the strategy of player 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

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 correspond to ESS

(see Weibull's book for details)

Simple Replicator Dynamics Analysis

Let's consider **symmetric games** ($R^T = C$):

	C1	C2
R1	a_1, a_1	a_2, a_3
R2	a_3, a_2	a_4, a_4

Assume that $P\{R1\} = p, P\{R2\} = 1 - p$ (and, because of symmetry, $P\{C1\} = q = p, P\{C2\} = 1 - q = 1 - p$). Thus:

$$E[R1] = E[C1] = pa_1 + (1-p)a_2, E[R2] = E[C2] = pa_3 + (1-p)a_4$$

Simple Replicator Dynamics Analysis

Let's call δ the net gain (or loss) arising from choosing the first strategy $R1$ over the second $R2$:

$$\delta = E[R1] - E[R2] = pa_1 + (1-p)a_2 - [pa_3 + (1-p)a_4],$$

which is a straight line $A + Bp$, with

$$A = (a_2 - a_4), \quad B = (a_1 + a_4) - (a_2 + a_3)$$

The replicator dynamics will favour the relatively more successful strategy. We thus study the behavior of $\Delta\delta/\Delta p = f(p)$

Replicator Dynamics: Hawks-Doves

	H	D
H	-2,-2	$^+2, 0^*$
D	$^+0, 2^*$	1,1

$A = a_2 - a_4 = 2 - 1 = 1, \;\; B = (a_1 + a_4) - (a_2 + a_3) = (-2+1) - (2+0) = -3$

$\delta > 0 \Longrightarrow p > -A/B$ (first strategy more successful)
$\delta < 0 \Longrightarrow p < -A/B$ (second strategy more successful)
$\delta = 0 \Longrightarrow p = p^* - A/B = (-1)/(-3) = 1/3$ (Nash equilibrium)

Replicator Dynamics: Hawks-Doves

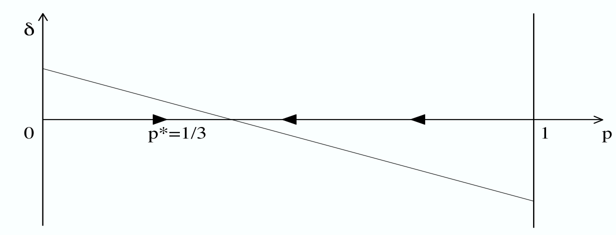

$\delta(0) = A = 1 > 0, B < 0:$ Positive intercept, negative slope

evolutionary equilibrium = nash equilibrium

Replicator Dynamics: Stag Hunt

	R	S
R	$^+1, 1^-$	2,0
S	0,2	$^+3, 3^-$

$A = a_2 - a_4 = -1, \;\; B = (a_1 + a_4) - (a_2 + a_3) = (1+3) - (2+0) = 2$

$\delta > 0 \Longrightarrow p > -A/B$ (first strategy more successful)
$\delta < 0 \Longrightarrow p < -A/B$ (second strategy more successful)
$\delta = 0 \Longrightarrow p = p^* - A/B = 1/2$ (Nash equilibrium)

Replicator Dynamics: Stag Hunt

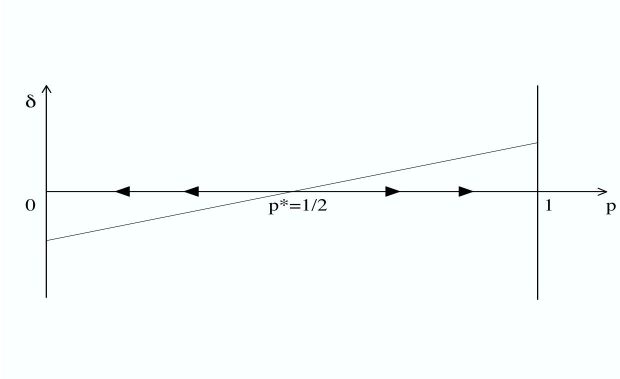

$\delta(0) = A = -1 < 0, B > 0$: Negative intercept, positive slope

Evolutionary equilibria are $p = 1$ and $p = 0$

Replicator Dynamics: PD

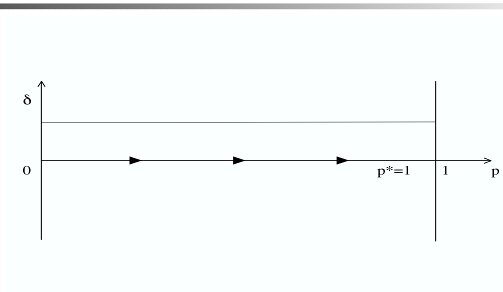

$[C] = 3p + 0 \times (1-p) = 3p; E[D] = 4p + 1 \times (1-p) = 3p + 1;$

$E[D] - E[C] = 1$, independent of p

Defection always predominates; it is the dominant strategy, the unique Nash equilibrium, and the unique evolutionary equilibrium

Spatially Structured Populations

Standard evolutionary game theory is valid for large, “mixing” populations in which any two players have the same probability of being selected to play the game

In 1992, Nowak et May showed by extensive numerical simulation that a regular structure with local interactions only can support a certain amount of cooperation, thanks to the formation of robust local cooperator clusters

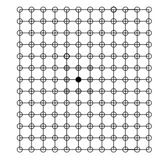

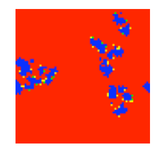

Relational Graphs

Neither random populations nor regularly structured ones are faithful descriptions of the relational structures that one finds in society

Regular Lattices ? Random Graphs

Does cooperation emerge in social networks?

Some Graph Statistics

- The **average path length** $\langle L \rangle$ of a graph G is the mean of all the shortest paths from all vertices to all other vertices
- The **clustering coefficient** C of G is (informally) the likelihood, averaged over all nodes in G, that nodes that are connected to a given node are also connected between them.
- The **average degree** $\langle k \rangle$ of G is the average number of neighbors of each node
- The **degree distribution** function $P(k)$ of G is the probability that a given node has exactly k neighbors

Small-World Networks I

- *Small-world graphs* are networks in which the average path length is short ($\langle L \rangle = O(logN)$, where N is the number of vertices). Thus, one can travel from any vertex to any other vertex in comparatively few steps, even in large graphs.
- This is also the case for standard random graphs. However, small worlds with the same number of vertices have a larger clustering coefficient. In other words, while random graphs are homogeneous in the average ($C = p = \langle k \rangle / N$), small-world networks have more local structure.

Small-World Networks II

- two important kinds of small-world networks are the *Watts-Strogatz* model and the *Scale-Free* model
- the latter is much more typical of real networks, while the former is an algorithmic construction that can be used in *artificial systems*, where there are no hard constraints on the topology

Small-World Networks III

The Scale-Free Model

Scale-Free graphs are also small-world (high clustering, low average path length) but they are characterized by a degree distribution function $P(k)$ of the power-law form: $P(k) = c\,k^{-\gamma}$, with c and γ positive constants, whereas random graphs and, to some extent, the Watts-Strogatz small-world model have a binomial $P(k)$

This form has been found in many real-world networks such as the Internet, the WWW, some biological networks, citation and collaboration networks and several others

Degree Distribution Functions

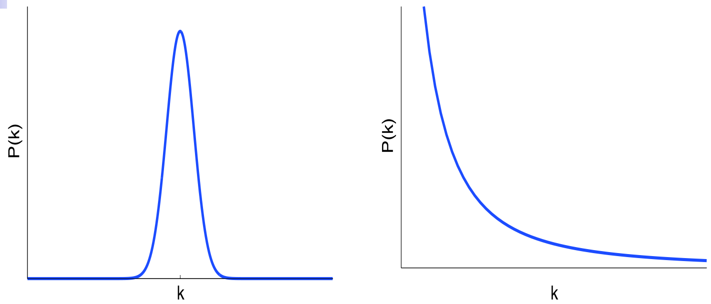

“Poissonian” DDF; Power law DDF.

Actual distributions would be discrete (finite graphs), and would show a maximum degree (cutoff for the power law)

Scale-Free Network Construction

They have been found in nature and society, thus they must be the result of some dynamical process.
They can also be generated algorithmically in several ways for example:

- The Barábasi–Albert dynamical *preferential attachment method*
- the static *configuration* model, and
- several other ways (see Newman and Barábasi–Albert for details)

A Scale-Free Network

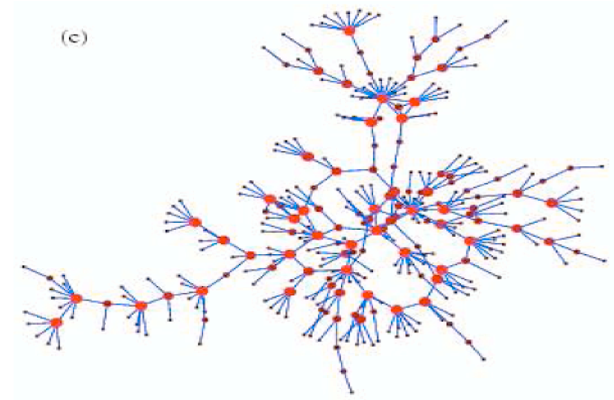

Social Networks

Social networks are often of the small-world type but usually do not belong to the pure scale-free family. The following is an example of an school acquaintance network, with cultural communities outlined:

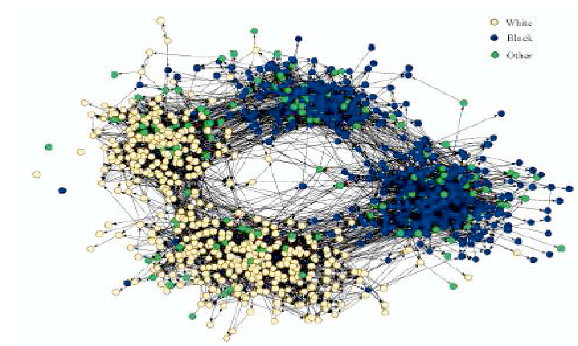

Cooperation on Scale-Free Graphs

On pure scale-free networks a high level of cooperation can be attained, both for the PD as well as the HD games:

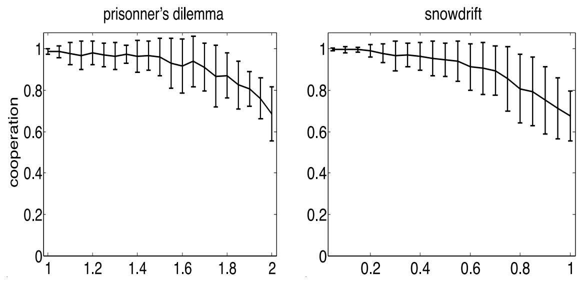

Fraction of cooperators in SF BA networks of size 10^4 with average degree $\bar{k} = 4$ using a discrete analogue of replicator dynamics. Averages over 50 runs

Cooperation on Scale-Free Graphs

But social networks have a less extreme (fat-tailed) degree distribution. So hubs cannot play such a fundamental role. When averaging a player's payoff over the number of links it has, results are strikingly different:

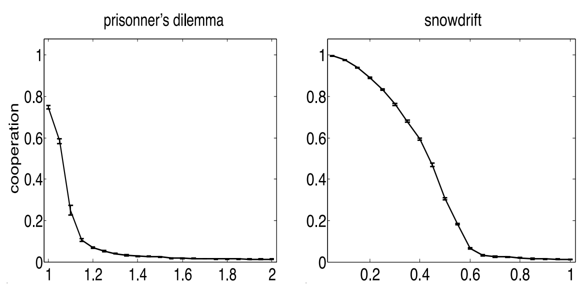

Stability of Cooperation on SFG

After a transient period, equilibrium is attained and it appears to be rather stable with respect to small rate strategy errors, although a fraction of defection avalanches can propagate through the system (errors appear in the highly connected players; errors in the strategy of low-connected players are almost unnoticeable)

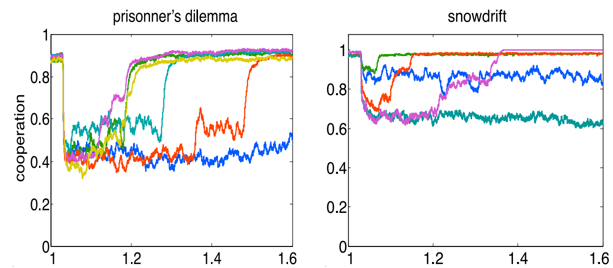

Evolving Networks

Actual social network structures are dynamic, not static. In fact, actors can join or leave the network and make new acquaintances, or abandon old ones at any time. This means that one should study **evolving networks**

In these networks, the rules of the game not only influence the player's strategy but also its propension to make new links or to cut them

The network **self-organises** according to these collective individual choices

Evolving Networks

For the PD game, strategy switching rules are the same as before, but now the agent i can also decide to cut an unsatisfying relationship with a neighbor; this can be either

- a D-D link, and/or
- a C-D link (seen from the cooperator's end)
- obviously, a C-C link is good for both players and they will tend to keep it
- the severed link can be re-established either with a random player, or with a player in the neighborhood of i

Evolving Networks

The system collective behavior at equilibrium will depend on the particular link-changing strategy adopted:

- if only D-D links are dismissed, cooperators chain will form that are rather stable against small perturbations
- if both D-D and C-D links are redirected, cooperation can still survive with cooperators having a tendency to cluster together; the resulting network has an exponential or Poisson degree distribution
- if the link redirection is biased toward a player in the immediate or second-order neighborhood, then cooperation tend to be stable at relatively high levels, and the emerging network tends to have a structure similar to that of other social networks

Some Conclusions and Ideas

- After more that 50 years, game theory has still not settled down!
- There have been many successful applications but also a lot of criticism of standard game theory and of its founding assumptions
- Evolutionary game theory has provided a good framework and some answers to the previous criticisms but not to all of them
- There is a tendency toward models of games where the players can progressively learn much as in standard machine learning

But...

Game theory is one of the few frameworks that is both mathematically rigorous and that can be applied to complex socio-economical, engineering, and biological problems when there are conflicting requirements

In this sense, it is more general than, say, multiobjective and constrained optimization as it is based on the aggregate result of individual decisions, and not on an externally assumed objective function

To Know More

Games General

R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge MA, 1991.

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

R. Axelrod. *The Evolution of Cooperation*. Basic Books, Inc., New York, 1984.

Networks General

R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.

M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

D. J. Watts. *Small worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, Princeton NJ, 1999.

Static Structured Populations

M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, 1992.

F.C. Santos and J.M. Pacheco. Scale-Free Networks Provide a Unifying Framework for the Emergence of Cooperation. *Phys. Rev. Lett.*, 95, 098104, 2005.

M. Tomassini and L. Luthi and M. Giacobini. Hawks and Doves on Small-World Networks. *Phys. Rev. E*, 73,016132, 2006.

Evolving Network Populations

M. G. Zimmermann and V. Eguíluz. Coevolution of Dynamical States and Interactions in Dynamic Networks. *Phys. Rev. E*, 69,065102, 2004.

L. Luthi and M. Giacobini and M. Tomassini. A Minimal Information Prisoner's Dilemma on Evolving Networks. *Artificial Life X*, MIT Press, Cambridge, MA, 438-444, 2006.