The Emergence of Cooperation in Evolutionary Game Theory

Anatol E. Wegner

Term Essay for Emergent States of Matter

Abstract: Cooperation is a fundamental aspect of biological systems. The evolution of complex organism, intelligence and complex societies all depend on cooperative behavior. Yet it is not clear how cooperation can evolve through natural selection. This essay will review game theoretical models of cooperation following a chronological approach with the final aim to explore their connection to physics.

Motivation and Historical Background:

Cooperation is a fundamental aspect of biological systems. The emergence of mutli-cellular organisms from uni-cellular ones, collective behavior in groups of animals, human societies are just a few of the countless examples of complex biological systems that depend on cooperative interactions. Yet it is not clear how cooperation can emerge in the face of exploitation. Cooperators make an investment for the common good while exploiters reap of the benefit avoiding costs, thus based on the evolutionary principle that the fittest survives or simply on rational behavior to generate the highest payoff one would expect exploitation to be dominant and cooperation to be rare in nature. Therefore the problem of finding mechanisms the can lead to the evolution of cooperation through natural selection has remained a core problem of biology since Darwin.

Darwin's theory provides a mechanism for evolution (natural selection) but also is explicitly competitive (the survival of the fittest) and based on the struggle for existence. One of the first to criticize the prevailing blindly competitive interpretation of evolution and draw attention to cooperation in evolution was P. Kropotkin in his classic "Mutual Aid- A Factor of Evolution" first published in 1890[1]. Still, until the 1960's collective phenomena drew little attention. Selection was mainly thought of as happening at the level of populations or even species. Cooperation was regarded as being adaptive. One of the first extensions of evolutionary theory accounting for cooperation and altruism was Hamilton's kinship theory [2] which took the gene-eye's view of evolution. This theory regarded the gene to be the fundamental unit of evolution that looks beyond its mortal bearer and seeks a potentially immortal set of its replicas to exist in a population related to it and explained altruism and cooperation as a mechanism for a gene to increase its frequency [3]. But still there are a lot of examples of cooperation, though rarely at the level of self sacrifice, between unrelated individuals and even different species in nature that this argument fails to explain.

Following the pioneering paper by Trivers, Axelrod used game theory to explain the emergence of cooperation based on a well studied game the Prisoners Dilemma. Since then game theoretical approaches in evolution have attracted huge interests and have been widely used to explain mechanisms through which cooperation can emerge and be maintained in different settings. The game theoretical approach is based on the theory first formulated by Morgenstern and Von Neumann [4] in the 1940's and applies methods previously developed in economics to evolutionary settings.

Game theory offers a mathematical framework for addressing these questions. In this paper we are going to review the most promising results of this approach and try to establish a link to physics.

Introduction

The Prisoners Dilemma

Game theoretical approaches to the evolution of cooperation are generally based on two player games called social dilemmas of which the most widely studied one is the prisoner's dilemma (PD). In game theory games are generally represented using payoff matrices. The PD game is explained in the table below.

Prisoner's Dilemma

	C	D
Payoff to C	$b - c$ (R)	$-c$ (S)
Payoff to D	b (T)	0 (P)

Table 1. The payoff matrix of the PD game. Another commonly used terminology is in terms of cost (c) and benefit (b), c represents the cost to cooperate whereas b is the benefit obtained from cooperating. If a one of the players cooperates and the other deflects the deflector get all of the benefit.

In the PD each player has two choices to cooperate C or to deflect D. If both cooperate each gets the payoff R (for reward), if one cooperates and the other deflects the cooperator gets the S (for sucker's payoff) and the deflector get T (for temptation), and if both deflect each gets P (for punishment). The prisoners dilemma is characterized by the payoff structure T>R>P>S. Therefore no matter what the opponent does the rational choice is D, but if both players deflect they do worse than they could have by both cooperating. Hence the dilemma.

The Iterated Prisoners Dilemma Game, Axelrod's-Tournaments [5]

If two individuals play the PD only once the only strategy that is a solution to the game is to deflect. This solution is stable in the sense of Nash equilibrium which can be defined as a state where none of the players can increase its gain by changing its strategy unilaterally. Therefore first extension used to study cooperation was the iterated prisoners dilemma in which players play a series of consecutive PD's. This provides greater room for cooperation. To see which strategies are effective Axelrod organized a series of round-robin computer tournaments. The rules of the tournament were:1)interactions are between pairs of players,2)Each player had two choices and choices were made simultaneously, 3)Payoffs were fixed,4)At each move of the game each player had access to the history of the game up to that move. Then in later tournaments an additional "shadow of the future was introduced" which corresponds to the game having a certain probability that it will end at the next round.

In these tournaments tit for tat (TFT) emerged as a clear winner against even quite sophisticated strategies. TFT consists of cooperating on the first move and then repeating the opponent's previous move in the next round. The success of TFT was explained by it being the nice as it did not deflect first, provocable as it retaliates against deflection by deflection, forgiving as even after deflection it still was ready to respond to cooperation by cooperation and at last simple in the sense that it was easy to understand for the opponent. Once TFT was established as a model for cooperation based on reciprocity it was investigated further in ecological simulations based on population dynamics were in a initially mixed population of strategies the frequency of a strategy was changed proportional to its success in the previous round. In these, TFT also quickly became the most common strategy. The next question that arises is how a strategy like TFT can emerge and invade an initially uncooperative population and maintain itself in the face of reinvasion. People soon realized that this needed the population to have some kind of structure that assured cooperative players interacted more with themselves rather than with the average population. This corresponds to the formation of clusters.

Axelrod's TFT attracted a great deal of interest and subsequent research concentrated on changing the game environment and relaxing the assumptions on various aspects of the game. The effects of changing interactions, choices, payoffs, population structure, noise etc. were studied extensively. The possibilities to extend the approach are numerous so we refer the reader to the review articles [6], [7] and [8] for further details. Instead we would like to concentrate on the connection between the evolutionary game theory approach to cooperation and physics. For this we will, first, take a look at the effects of spatial structure on the emergence of cooperation. Second, we will see how increasing the number of players and giving them the choice to stay out of the game changes the game dynamics.

Mixed Populations and Replicator Dynamics

In a mixed population with no structure (every player interacts with the others randomly with equal probability) the population dynamics is given by the replicator equations. In a population with a fraction ρ of cooperators and 1- ρ of deflectors the average payoff is $P_C=\rho R+(1-\rho)S$ for cooperators and $P_D=\rho T+(1-\rho)P$ for deflectors[9]. Then the rate of change of ρ is given by:

$\dot{\rho}=\rho(P_C-\overline{P})=\rho(1-\rho)(P_C-P_D)$ as $P_D > P_C$, ρ converges to zero with time. Thus this shows that in well mixed populations cooperators are doomed to extinction. To overcome this dilemma one has to consider structured populations where players only interact with a restricted neighborhood.

PD on lattices

The above problem was already pointed out by Axelrod in his original paper. But what really increased the interest in games with structured populations, sometimes called games on grids, was the work of Nowak and May in 1992 [10]. We refer the reader to [11] for other examples of games on grids and a very accessible and interactive online tutorial on evolutionary game theory.

In the paper n^2 players play PD on a two dimensional n by n lattice. Each payer is either always cooperates of deflects. At each round players play the PD with their neighbors. Then the player on each site is replaced by the highest scoring player in the neighborhood including the player itself. The PD has a payoff matrix with R=1, T=b (b>1), S=P=0. As the dynamics is discrete one can identify a range for the parameter b (1.8< b < 2.0) for which the outcome of the system is chaotic. In this range of parameter values the system exhibits a variety of fascinating spatial patterns. Moreover Nowak and May demonstrated that for a very large set of initial conditions the fraction of cooperators approaches to 12log2-8=0.318. Some of the patterns and the plot of cooperator frequencies from the article are shown below.

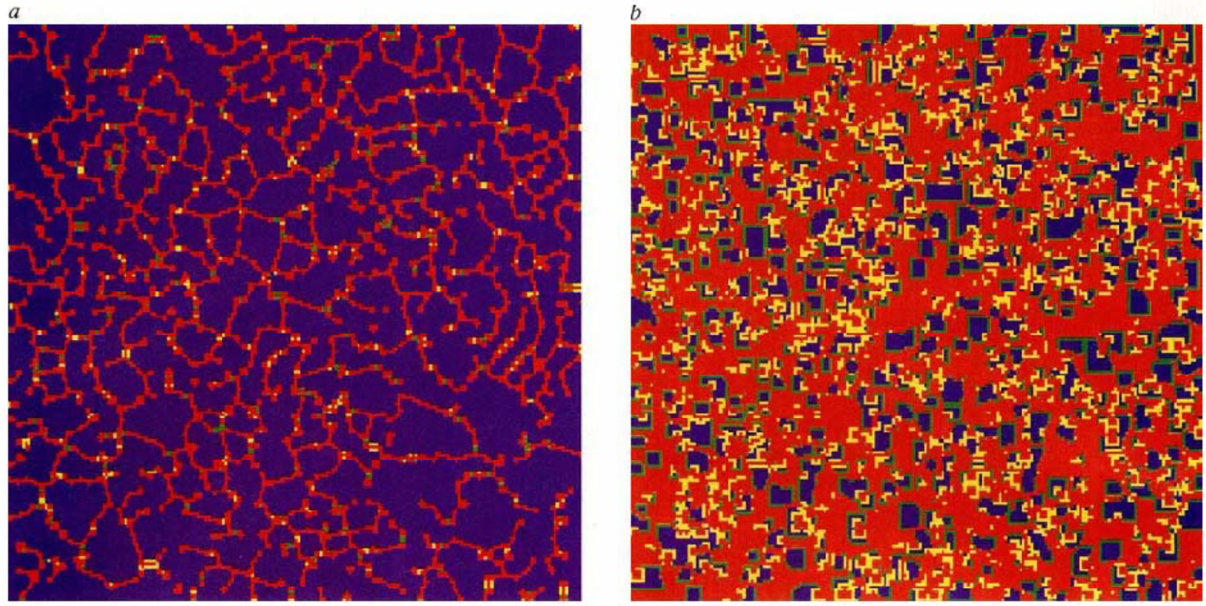

FIG. 1 The spatial Prisoners' Dilemma can generate a large variety of qualitatively different patterns, depending on the magnitude of the parameter, b, which represents the advantage for defectors. This figure shows two examples. Both simulations are performed on a 200×200 square lattice with fixed boundary conditions, and start with the same random initial configuration with 10% defectors (and 90% cooperators). The asymptotic pattern after 200 generations is shown. The colour coding is as follows: blue represents a cooperator (C) that was already a C in the preceding generation; red is a defector (D) following a D; yellow a D following a C; green a C following a D. a, An irregular, but static pattern (mainly of interlaced networks) emerges if $1.75 < b < 1.8$. The equilibrium frequency of C depends on the initial conditions, but is usually between 0.7 and 0.95. For lower b values (provided $b > \frac{5}{3}$), D persists as line fragments less connected than shown here, or as scattered small oscillators ('D-blinkers'). b, Spatial chaos characterizes the region $1.8 < b < 2$. The large proportion of yellow and green indicates many changes from one generation to the next. Here, as outlined in the text, 2×2 or bigger C clusters can invade D regions, and vice versa. C and D coexist indefinitely in a chaotically shifting balance, with the frequency of C being (almost) completely independent of the initial conditions at ~ 0.318.

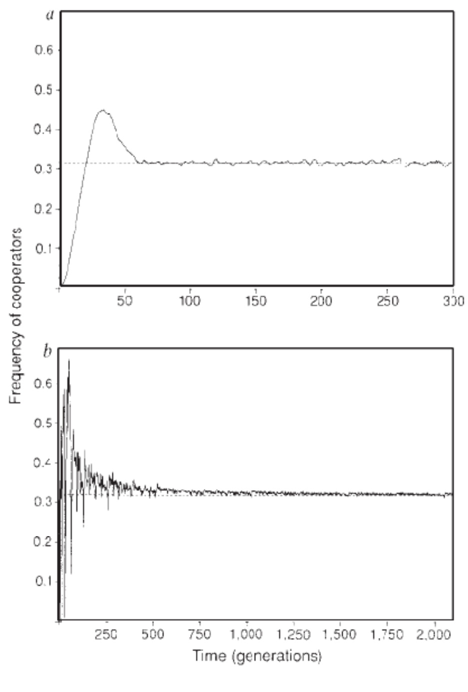

FIG. 2 The frequency of cooperators in simulations with random or symmetrical initial conditions, within the interesting region $1.8 < b < 2$. a, The frequency of cooperators, $f_C(t)$, for 300 generations, starting with a random initial configuration of $f_C(0) = 0.6$. The simulation is performed on a 400×400 square lattice with fixed boundary conditions, and each player interacts with 9 neighbours (including self). The dashed line represents $f_C = 12 \log 2 - 8$ (see b). b, The frequency of C within the dynamic fractal generated by a single D invading an infinite array of C.

PD on graphs and networks

Complex networks are widely used as models of social and biological interactions. Therefore the relation between well known network topologies and the evolution of cooperation is of great interest. In [12] a simple condition for the evolution of cooperation on networks and graphs is described in terms of cost to benefit ratio and the degree of the graph.

In the model each player occupies a vertex on the graph. Cooperators help neighboring cooperators at a cost c, so if a cooperator has i cooperating neighbors the payoff he receives is bi-ck . On the other hand deflectors don't help each other so they don't have any costs but they can benefit from cooperators so the payoff a deflector receives when it has j cooperating neighbors is bj. The updating is through a "death-birth" mechanism that is at each round a player is chosen at random and replaced by one of his neighbors with a probability that depends on the fitness of each neighbor. The model assumes weak selection which is commonly used in models which means that the fitness is not directly proportional to the payoff but rather 1-w+wP, where P is the payoff and w is a small parameter w<<1. Then the fixation probability for a cooperator is calculated through computer simulations. The fixation probability is the probability that a single cooperator turns an entire population from deflection to cooperation. If selection is neutral to cooperation the probability is 1/N (N being the size of the population). So if the fixation probability is greater the 1/N selection favors cooperation. The condition b/c>k (k=the average number of neighbors) for the evolution of cooperation initially derived using approximations is confirmed by the simulations. The results of the paper [12] are summarized in the figure below.

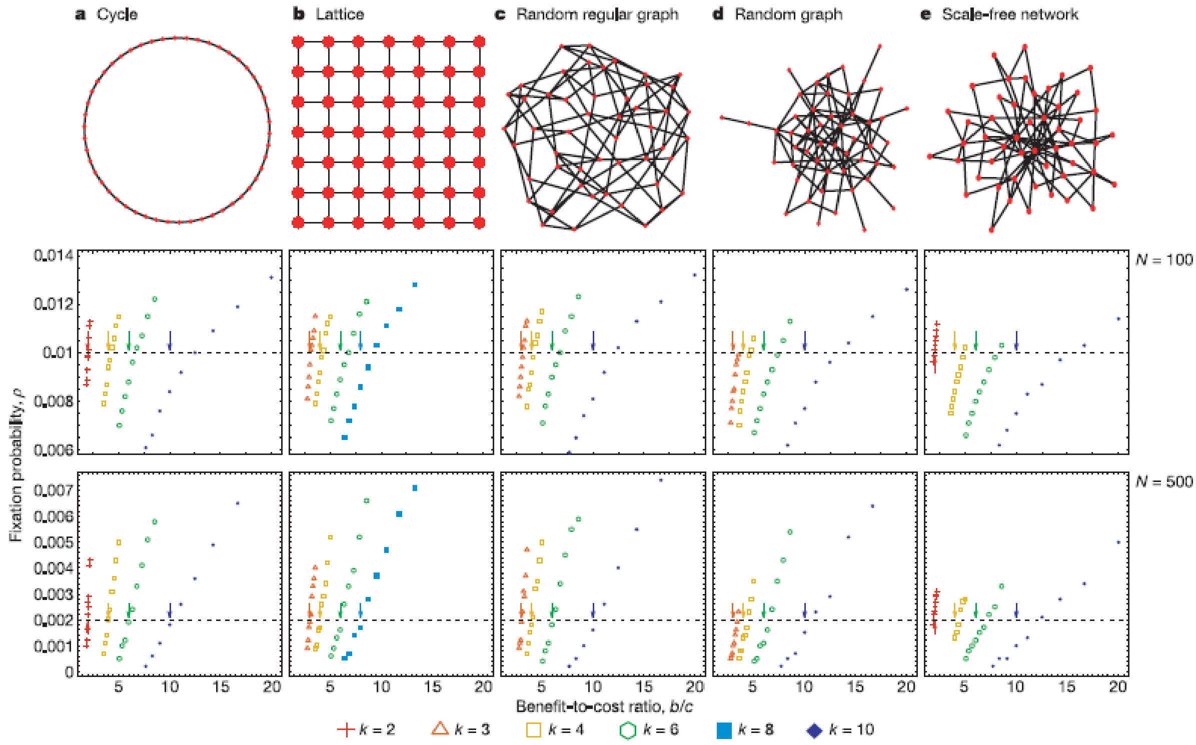

Figure 3 | The simple rule, $b/c > k$, is in good agreement with numerical simulations. The parameter k denotes the degree of the graph, which is given by the (average) number of neighbours per individual. The first row illustrates the type of graph for $k = 2$ (**a**) and $k = 4$ (**b**–**e**). The second and third rows show simulation data for population sizes $N = 100$ and $N = 500$. The fixation probability, ρ, of cooperators is determined by the fraction of runs where cooperators reached fixation out of 10^6 runs under weak selection, $w = 0.01$. Each type of graph is simulated for different (average) degrees ranging from $k = 2$ to $k = 10$. The arrows mark $b/c = k$. The dotted horizontal line indicates the fixation probability $1/N$ of neutral evolution. The data suggest that $b/c > k$ is necessary but not sufficient. The discrepancy is larger for non-regular graphs (**d**, **e**) with high average degree ($k = 10$). This is not surprising given that the derivation of the rule is for regular graphs and in the limit $N \gg k$. Note that the larger population size, $N = 500$, gives better agreement.

A noticeable property of these results is their similarity with Hamilton's rule of kin selection which states that selection can favor cooperation if rb>c. Where r is the degree of relatedness b is the benefit and c the cost. Thus one can identify 1/k with r.

The Public Goods Game and Loners

Another extension of the PD is to increase the number of players. This generalization is called the public goods game (PGG). In this game n players have the option to invest in a common pool which is then increased by a certain factor α and consequently divided equally among the participants regardless of their individual contribution. The case where everybody deflects is the only Nash equilibrium of the game. Therefore, as in the PD, the rational thing to do is to deflect but this means that the group will forego the possible benefit of the game ([7] box 2, [13]).

One of the solutions offered to this dilemma is to make participation voluntary, creating another option L (for loner). Loners rely on a small but steady income σ. This results in rock-paper-scissors type of hierarchy between the strategies. If everybody cooperates it pays off to deflect, if everybody deflects it is better to leave the game and if there are many loners the effective number of players decreases which favors cooperation. As a result different strategies coexist with oscillating frequencies. The PGG has very rich dynamical properties and can lead to interesting phase transitions and the formation of rich spatiotemporal patterns.

Physical Models-Phase Transitions

After the spatial PD became popular there was an effort to apply methods of statistical and condensed matter physics to evolutionary game theory [14],[15],[16]. These models in general heavily depend on Monte Carlo simulations. The models consider players distributed on a certain lattice, and assign a payoff dependent transition probability between two strategies x and y that is given by:

$$W[s(\mathbf{y}) \rightarrow s(\mathbf{x})] = \frac{1}{1 + \exp\{[P(\mathbf{x}) - P(\mathbf{y}) + \tau]/K\}} \qquad \text{Eq. 1.}$$

Here P(x) and P(y) are the payoffs of the strategies in the previous round and τ is a parameter that represents the cost of changing strategies. K has an interpretation as a measure of the noise in the system that allows for irrational strategy changes which don't maximize payoffs and also as measure of variations in the payoffs. The payoffs are rescaled such that *R*=1, *T*=1+*r*, *S*=-*r*, and *P*=0, where *r*=*c*/(*b*-*c*) denotes the ratio of the costs of cooperation to the net benefits of cooperation. Note that this is differs from the cost to benefit ratio in the previous part. With these rules established cites are updated in an asynchronous fashion in a Monte Carlo simulation starting from random initial conditions. At each step a two neighboring cites are selected at random and x adopts the strategy y with the probability W described in Eq. 1. Then the stationary state is characterized by averaging over the sampling time. The Monte Carlo simulations allow the determination of critical exponents near the critical value r_c which characterizes the extinction threshold for cooperators. Moreover using the critical exponent of the phase transition from the cooperative to the non-cooperative state on a square lattice can be classified to be in the directed percolation universality class [15].

The figures below show the results of the simulations (τ=0) for a square lattice and various network topologies [15].

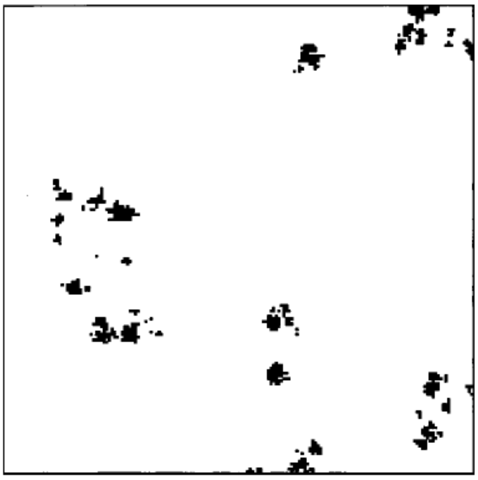

Fig. 4. Typical distribution of cooperators (black) in a sea of defectors (white) on a square lattice for $r=0.0211$ and $\kappa=0.1$, just below the extinction threshold $r_c=0.021\,12$ (2) of cooperators. Note that the distribution is essentially independent of the initial lattice configuration. However, in finite systems the frequency of cooperators should not be too low, so as to avoid accidental extinctions while approaching the stationary state.

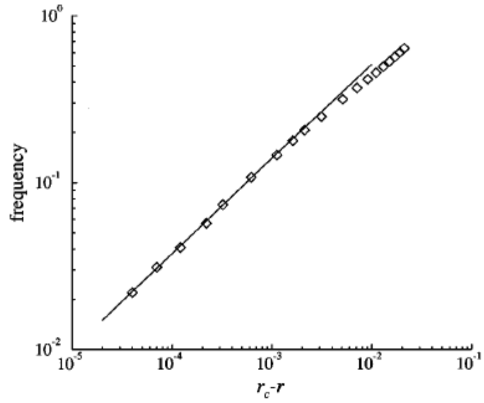

Fig. 5 Log-log plot of the average fraction of cooperators ρ as a function of the distance to the extinction threshold r_c-r. The solid line shows that in the vicinity of r_c, the power law $\approx (r_c-r)^\beta$ perfectly fits the data with $r_c=0.021\,12$ (2) and $\beta=0.57$ (3). The system size was increased as r approaches r_c from $N=1.6\times10^5$ to $N=10^6$.

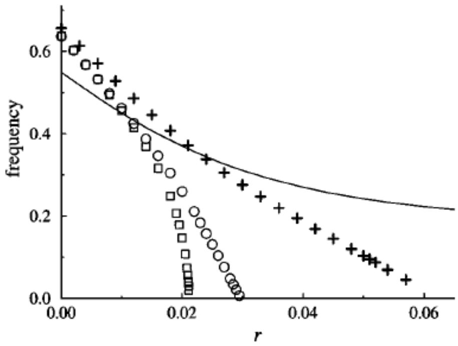

Fig. 6. Fraction of cooperators ρ as a function of r for different population structures: square lattice (□), random regular graph (+), and regular small world networks (○) for $Q=0.03$, $\kappa=0.1$, and $N=1.6\times10^5-10^6$. For increasing r, the spatial correlations result in a critical transition on the square lattice (see Fig. 3), whereas on random regular graph and small world networks the lack of correlations lead to a linear decrease in cooperation, that is, a mean-field type transition. The data referring to homogeneous D states (cooperators go extinct and defectors reach fixation) is omitted. The pair approximation (solid line) correctly predicts the trend, but significantly overestimates the benefits of population structures

Maybe we should mention the pair approximation at this point which is a widely applied analytical method in evolutionary game theory. The pair approximation is mean field theory based on tracking the frequency of pairs of strategies. But in general the pair approximation fails to predict the frequencies accurately and is mainly used as a tool to predict trends in strategy frequencies. We refer the interested reader to the appendix of [15] for further details.

Another problem this method has been applied to is the voluntary PGG[16]. The extension of the model is straightforward. In this variant a PGG is played with nearest neighbors on a square lattice, i.e. a PPG with 5 players. The transition probability between strategies is again given by Eq.1. The payoffs for a round with $n_C + n_D + n_L = 5$ cooperating, deflecting and loner players is given by:

$$P(\mathbf{x}) = \begin{cases} \frac{r n_c}{n_c + n_d} - 1 & \text{if } s(\mathbf{x}) = C, \\ \frac{r n_c}{n_c + n_d} & \text{if } s(\mathbf{x}) = D, \\ \sigma & \text{if } s(\mathbf{x}) = L, \end{cases}$$

Where r is the multiplication factor of the common good and σ is the loner payoff.

This can also be generalized to lattices with arbitrary structure. Using the above described Monte Carlo method the frequencies of the strategies and their sample trajectories as a function of r can be calculated. The interesting aspect of the voluntary PGG is that it has two phase transitions corresponding to the disappearance of loners and cooperators. The critical exponents of these transitions can also be calculated. Moreover the cyclic rock-paper-scissors hierarchy can lead to the formation of various patterns[18]. The figures 7, 8 and 9 show some of these results from [16].

These results show that methods and concepts of statistical physics and condensed matter physics can be provide new insights to evolutionary game theory.

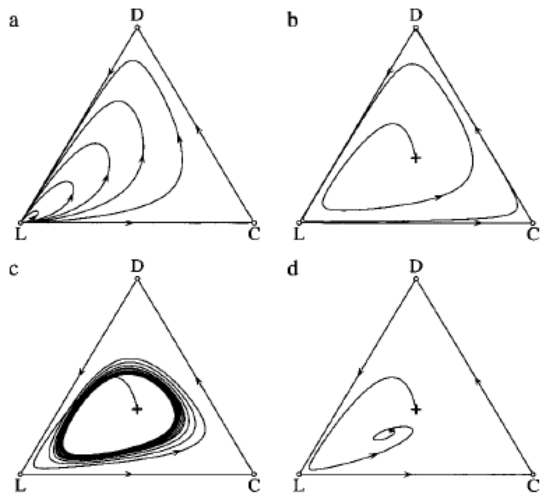

Fig. 7. Sample trajectories of the evolution of the frequencies of cooperators, defectors, and loners in the voluntary prisoner's dilemma for different population structures. The boundary of the simplex S_3 consists of a heteroclinic cycle which reflects the cyclic dominance of the three strategies. (a) In well-mixed populations the system relaxes into homogenous states of all loners. (b) For RRG the trajectories spiral outward and eventually end in one of the three absorbing states, but usually they end in the loner corner as in (a). (c) Regular small world networks networks ($Q=0.03$) substantially change this outcome and reveal an asymptotically stable limit cycle leading to persistent global oscillations of the three strategies. (d) On square lattices the system evolves toward a stable stationary state with all three strategies coexisting [$\rho_D=0.229$ (1), $\rho_C=0.269$ (1), and $\rho_L=0.502$ (1)]. All simulations [(b)–(d)] were done for the for $r=0.4$, $\sigma=0.3$, $\kappa=0.1$, and $N=10^6$. Note that for these parameters defectors invariably reach fixation (cooperators go extinct) in the absence of the loners. The simulations in (b)–(d) have random initial configurations with identical concentrations of all three strategies (marked by +).

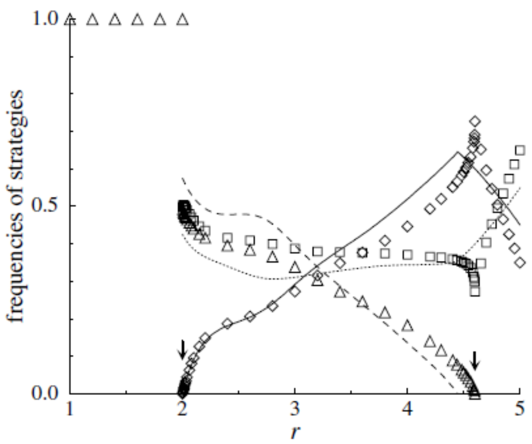

FIG. 8. Frequencies of cooperators C (open squares), defectors D (open diamonds), and loners L (open triangles) as a function of r for $\sigma = 1$ and $\tau = K = 0.1$. The results of pair approximation are shown as dotted (C), solid (D), and dashed (L) lines. The values of r_D and r_L are indicated by arrows.

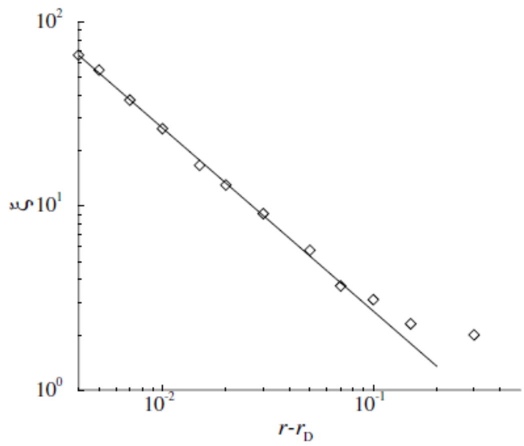

FIG. 9. Log-log plot of the correlation length ξ vs $r - r_D$ for the three strategies (parameters as in Fig. 8). The solid line denotes the fitted power law with an exponent -0.99.

Other Mechanisms-Mobility

A recently discovered factor that can increase cooperation is mobility, which can be describe as the ability of players to migrate to locations with more promising conditions. A recently published article [17] demonstrates that a combination of mobility and imitation of successful strategies can result in dominantly cooperating population even in an initially purely deflecting one. Moreover mobility has been shown to promote bio-diversity in rock-paper-scissors type cyclic competitions [18]. This article [18] also shows how such competition can lead to spiraling spatial patterns under suitable conditions.

Some Empirical Examples

Predator inspection in fish: In many species of fish individuals separate from their shoal and approach predators for inspection. This behavior was studied extensively by Milinski in stickleback fish [19]. Using a system of mirrors, Milinski provided the fish approaching a live predator with either a simulated cooperating companion or a simulated defecting one. The results were consistent with tit for tat. Though, it is still argued whether the game is a PD or the less stringent version of the PD, the snowdrift game.

Yeast : Certain kind of yeast produce an enzyme to hydrolyze sucrose, the enzyme is used in common and thus there might exist cells which produce enzyme while some others profit from the common resource without contributing.

Cooperation in humans: Actually humans are very promising subjects for the experimental study of cooperation. Because repeated and controlled experiments are possible. In general this is not the case for other social animals whose natural behavior can't be observed under controlled conditions. Cooperative behavior in human has been the subject of extensive research [20]. Maybe the most relevant experiments to our discussion are public goods games. Interestingly human subjects don not follow rational reasoning in experimental public goods games and cooperate much more than expected [21]. This has led people to question the rationality assumptions in economics.

For a more detailed list of empirical examples see [7] box 6.

Conclusion

Evolutionary game theory is a very diverse subject and a still very active field of research. There are a huge number of mechanisms which can promote cooperation, in this essay only a fraction of them were explored. The aim was to introduce the subject in a consistent way with the outlook to establish a link to physics.

The theory of evolution is one of the most important intellectual achievements of mankind and dominates our way of thinking. The results explored in this essay show that even simple models of cooperation can lead to complex and diverse outcomes. This shows that mathematical models are crucial to advance our understanding of evolution.

References:

1. *Kropotkin P., Mutual Aid: A Factor of Evolution (1998 paperback ed.). London: Freedom Press. ISBN 0-900384-36-0.*
2. *Hamilton, W. D. (1963). "The evolution of altruistic behavior". American Naturalist 97: 354–356.*
3. *R. Dawkins, The Selfish Gene (Oxford Univ. Press, Oxford,1976).*
4. *J. von Neumann and O. Morgenstern, Theory of Games and Economic Behaviour ,Princeton U. P, Princeton, 1944*
5. *Axelrod, R. & Hamilton, W.D. (1981). The evolution of cooperation. Science 211, 1390–1396.*
6. *Axelrod, R. & Dion.D. (1988). The further evolution of cooperation. Science 242, 1385–1390.*
7. *Doebeli M. and Hauert C. , (2005) Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game Ecology Letters,8: 748–766.*
8. *Nowak, M. A. & Sigmund, K. (2004) Evolutionary Dynamics of biological games. Science 303, 793–799*
9. *J. Hofbauer, K. Sigmund, Evolutionary Games and Population Dynamics (Cambridge Univ. Press, Cambridge, 1998).*
10. *M. A. Nowak and R. M. May,(1992) ''Evolutionary games and spatial chaos'' Nature 359, 826–829*
11. *http://www.univie.ac.at/virtuallabs/*
12. *Ohtsuki H.,Hauert C.,Lieberman H and Nowak M. A.(2006)A simple rule for the evolution of cooperation on graphs and social networks Nature 441, 502-505*
13. *Hauert, C., Haiden, N. & Sigmund, K. (2004). The dynamics of public goods. DCDS B 4, 575–587.*
14. *Szabo´ , G. & To˜ke, C. (1998). Evolutionary prisoner's dilemma game on a square lattice. Phys. Rev. E 58, 69–73.*
15. *Szabo´ , G. & Hauert, C. (2002). Phase transitions and volunteering in spatial public goods games. Phys. Rev. Lett. 89, 118101.*
16. *C. Hauert and G. Szabo, (2005) Game Theory and Physics, Am. J. Phys. 73, 405.*
17. *Helbing, D., Yu, W., 2009. The outbreak of cooperation among success-driven individuals under noisy conditions. Proc. Natl.Acad. Sci. USA 106, 3680–3685.*
18. *Reichenbach T., Mobilia M. and Frey E. (2007) Mobility promotes and jeopardizes biodiversity in rock-paper-scissors games,Nature 448, 1046-1049*
19. *Milinski, M. (1987). Tit for tat in sticklebacks and the evolution of cooperation. Nature 325, 433–435*
20. *The Handbook of Experimental Economics, edited by J. H.Kagel and A. E. Roth (Princeton University Press, Princeton, 1995).*
21. *Fehr, E. & Gachter, S. (2002). Altruistic punishment in humans.Nature 415, 137–140.*