Cooperation in Networked Populations of Selfish Adaptive Agents: Sensitivity to Learning Speed

Gulyás, László

Veröffentlichungsversion / Published Version
Zeitschriftenartikel / journal article

European Quarterly of Political Attitudes and Mentalities EQPAM
Vol.2, No.1, January 2013

ISSN 2285 – 4916
ISSN-L 2285 - 4916

Cooperation in Networked Populations of Selfish Adaptive Agents: Sensitivity to Learning Speed[1]

László Gulyás[1,2]

Submission for EWWW-PAM'2012: May 27th, 2012 Accepted for EWWW-PAM'2012: Jun 11th, 2012
Accepted for EQPAM: October 25th, 2012

Abstract

This paper investigates the evolution of cooperation in iterated Prisoner's Dilemma (IPD) games with individually learning agents, subject to the structure of the interaction network. In particular, we study how Tit-for-Tat or All-Defection comes to dominate the population on Watts-Strogatz networks, under varying learning speeds and average network path lengths. We find that the presence of a cooperative regime (where almost the entire population plays Tit-for-Tat) is dependent on the quickness of information spreading across the network. More precisely, cooperation hinges on the relation between individual adaptation speed and average path length in the interaction topology. Our results are in good agreement with previous works both on discrete choice dynamics on networks and in the evolution of cooperation literature.

1. Introduction

Selfishness is a naturally expected and accepted evolutionary reaction to a harsh environment. The wide-spread assumption of rational human decision-making is thus in contrast with the empirically often observed altruistic human behavior.

Corresponding Author: Dr. László Gulyás, Assistant Professor, Department of History and Philosophy of Sicence, Loránd Eötvös University, Budapest, Hungary
Address: 1117 Budapest, Pázmány Péter sétány 1/c, Hungary
e-mail: lgulyas@aitia.ai

European Quarterly of Political Attitudes and Mentalities - EQPAM, Volume2, No.1, January 2013, pp. 55-73.
ISSN 2285 – 4916
ISSN–L 2285 – 4916

[1] This paper has been previously published in the Proceedings of the 1st *Eastern European Exploratory Workshop on Political Attitudes and Mentalities,* EEEW-PAM'2012, September 2nd, 2012, University of Bucharest, Romania, in: Voinea C.F. (ed.), *"Political Attitudes and Mentalities. Old Theoretical Foundations – New Experimental and Prediction Challenges",* University of Bucharest-Ars Docendi Press, 2012, Bucharest, pp.117-134.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

The puzzle is how cooperation can emerge in a population of selfish, rational agents. Evolutionary game theory offers an effective framework to study puzzles of this kind. In evolutionary game theory agents play a game repeatedly, with a specific mechanism to update their strategies between successive rounds. (Weibull, 1995) (Gulyas and Platkowski, 2004).

In the context of this paper, similarly to most fields of the social sciences, cooperation is understood strategically: i.e., as the (missing) tendency of the agent to take unilateral advantage over its partners. The classic framework to explore such situations is the (Iterated) Prisoner's Dilemma game (IPD).

Following the series of works by Axlerod, this framework is now ubiquitous in many disciplines and forms the base of many models, studying a variety of problems. For example, see (Cederman, 1997) (Cohen, 1994) (Gowa and Mansfield, 1993) (Pruitt, 1998) (Szabo and Fath, 2006) Within this framework, the success of altruistic cooperation is often understood as the presence of the tit-for-tat (*TFT*) strategy, that cooperates on the first round and then reciprocates the moves of its partner, against the all-defective (*ALLD*) strategy that exploits its opponents in all rounds (i.e., never cooperates).

The role of interaction topology (i.e., network structure) was also studied in IPD settings. Axelrod, Riolo and Cohen (CAR) study how cooperation (c.f., trust) can evolve spontaneously in a population of selfish agents, using a networked IPD games framework with an evolutionary adaptation rule. (Axelrod, Riolo and Cohen, 2000) In the context of their evolutionary model, they found that context-preservation (i.e., a static network) and not network structure was the key for the survival (and domination) of cooperative behavior (i.e., agents playing *TFT*). (Cohen, Riolo and Axelrod, 2001) On the other hand, when network dynamics was allowed, *ALLD* agents took over the world. The independence of the emergence of cooperation result has a strong message in the context of cooperation and trust in society. However, it is in stark contrast with the growing body of results in information/behavior cascades, where discrete choice dynamics were found to be dependent on network topology, and especially on the *l* average path length in the network. (Gulyas and Dugundji, 2003) (Dugundji and Gulyas, 2006) (Herrero, 2002) (Nikoletopoulos et al, 2004) (Watts, 2002) (Klüver and Schmidt, 1999)

Therefore, in earlier works, we revisited the question of network dependence in IPD games with the aim to consolidate the previously discussed apparent contradiction. (Gulyas, forthcoming) Our hypothesis was that independence from the network topology was caused by the sharp, step-like nature of adaptation in the evolutionary setting, when agents switched from *ALLD* to *TFT* after a single round of underperformance. In order to investigate this hypothesis, we departed from the classic IPD framework applied in the CAR-model and designed a model using an individually inductive adaptation rule that 'smoothens out' the strategy change. This framework assumes a discrete setting with individual learning, instead of evolutionary adaptation. Studying this altered model on static Watts-Strogatz networks, we found that context-preservation is not always sufficient for achieving a cooperative regime. In particular, our numerical results suggested that the performance of the *TFT* strategy may depend on the (static) network structure, eliminating the contradiction above.

In the present paper we further explore this dependence. Especially, we focus on cooperation's dependence on the speed of learning, relative to network structure. We find that cooperation is dependent on the quickness of information spreading across the network, which depends on the relation of adaptation speed and average path length in the interaction topology.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

As a special case, these results yield the findings of (Axelrod, Riolo and Cohen, 2000) on cooperation's independence from (static) network topologies with the evolutionary adaptation rule. Also they are in good agreement with earlier results on discrete choice dynamics.

The paper is structured as follows. The next section introduces the classic evolutionary IPD framework, which is followed by the description of our modified, individually inductive IPD model. Section 4 discusses earlier results on the network dependence of these two families of models. Section 5 summarizes the main contributions of this paper: cooperation's dependence on the adaptation speed in the individually inductive framework. This is followed by a discussion of related works, while the last section concludes the paper.

1. The Classic Evolutionary IPD framework

In the Prisoner's Dilemma (PD) game, two agents face the same binary choice: each has the option to either cooperate (*C*) or to defect (*D*) – without knowing the choice of the partner. The payoff matrix of this game is shown on Table 1, with the condition that $T>R>P>S$ and $2R>S+T$.[3] That is, choosing *D* always has a higher expected payoff than choosing *C*, yet the combined payoff of the two agents is best when both of them cooperate. The deficiency of defection becomes more pronounced when the partners play consecutive games, as the difference $2R-(S+T)>0$ accumulates. This makes the number of games (iterations, n_i) important. Iterated Prisoner's Dilemma (IPD) games ($n_i>1$) are interesting if the agents have a memory of the last $n_m < n_i$ actions taken by their partner. (Schweitzer, Mach and Mühlenbein, 2005)

Table 1. The payoff matrix of the Prisoner's Dilemma game, T > R > P > S, 2R > S+T.

	Player 2's action			
Player 1's action	Cooperates		Defects	
Cooperates	R (reward)	R (reward)	S (sucker)	T (temptation)
Defects	T (temptation)	S (sucker)	P (punishment)	P (punishment)
	Player 1's Payoff	Player 2's Payoff	Player 1's Payoff	Player 2's Payoff

In the following, we restrict our analysis to one-step memory (n_m =1) strategies, in accordance with (Cohen, Riolo and Axelrod, 2001) and (Gulyas, forthcoming). In the deterministic case, such strategies can be described with a sequence of 3 actions: A_I, A_C, A_D. That is, the action taken initially (when no information about the partner is available), in response to a cooperative and a defective action, respectively. Among the 8 possible strategies, we focus on the four studied in (Gulyas, forthcoming), as listed on Table 2.[4]

Table 2. The 4 one-step memory strategies studied.

Name (abbreviation)	Action sequence
Always cooperates (ALLC)	CCC
Tit-for-Tat (TFT)	CCD
Ant-Tit-for-Tat (ATFT)	DDC
Always defects (ALLD)	DDD

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

2. Networked Evolutionary Iterated Prisoner's Dilemma Games

The evolution of cooperation is studied in a 'society' A of more than 2 agents (N=|A|>2). In this context, agents playing IPD games have to be 'paired', i.e., the interaction topology of the society must be specified. The most general formulation of the topology is a graph (G=(A, E), where $E \subseteq A \times A$): the nodes are the agents and the links represent pairs playing IPD games. Each agent plays an IPD game with all of its $E(a)$ neighbors independently, but simultaneously. (For simplicity, we assume an undirected network.) Agents are assigned a random initial strategy drawn from a uniform distribution.

At the end of each round (i.e., after each agent has played with all of its neighbors), the agents have the option to change their strategies. In the evolutionary family of models, this is done by observing the profitability of the strategies applied by the agent's neighbors. (Axelrod, Riolo and Cohen, 2000) (Cohen, Riolo, Axelrod, 2001) Profitability is expressed as the average payoff per game achieved in the current round by the neighbor in question. Let PO_a^t denote the payoff collected, and n_a^t the number of IPD games played by agent a in round t. Then the profitability of agent a at time t is expressed as

$$\pi_a^t = \frac{PO_a^t}{n_a^t} \tag{1}$$

The evolutionary adaptation rule assumes that at the end of the round the agents copy the strategy of their most profitable neighbor. If σ_a^t stands for the strategy played by agent a in round t, then this adaptation rule can be expressed as

$$\sigma_a^{t+1} = \sigma^t_{\arg\max_{i \in E(a)} \pi_i^t} \tag{2}$$

If different strategies scored the same maximum value, the agent applies a random choice. Notice that this adaptation rule assumes that the agents have the means to *ex post* discover the strategies played by their opponents. However, this does not mean that they would have the same capability during the IPD game.

3. Individually Inductive IPD Framework

The individually inductive framework introduced in (Gulyas, forthcoming) is based on the classical approach of the previous section. However, for the sake of a new adaptation rule, basic model elements need to be altered. The agents of our model will continue to play one of the four 1-step memory (n_m=1) strategies as above in each round. (See Table 2.) However, they will probabilistically pick one of these strategies *at the beginning of each round* and play this fixed strategy *against all of their opponents*.[5] Let $p_a^t(\sigma)$ denote the probability that agent a plays strategy $\sigma \in \{ALLC, TFT, ATFT, ALLD\}$ at time t. Naturally,

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

$$\sum_{\sigma\in\{ALLC,TFT,ATFT,ALLD\}} p_a^t(\sigma) = 1 \qquad (3)$$

for all $a\in A$ and $t\geq 0$.

In order to approximate the random initial strategy configuration of the evolutionary adaptation model, agents are assigned a special initial strategy probability distribution. In this distribution, one strategy has a definite advantage ($p\geq 0.25$, set to 0.3 in our experiments) and the complementing probability is distributed equally among the 3 remaining strategies (see Table 3). The random initial configuration is created in such a way that an approximately equal number of agents will have a bias for each strategy.

Table 3. Initial distribution of strategy probabilities for the four strategy classes of the evolutionary adaptation model (p is a model parameter).

	$p_a^0(ALLC)$	$p_a^0(TFT)$	$p_a^0(ATFT)$	$p_a^0(ALLD)$
Former *ALLC*	p	$(1-p)/3$	$(1-p)/3$	$(1-p)/3$
Former *TFT*	$(1-p)/3$	p	$(1-p)/3$	$(1-p)/3$
Former *ATFT*	$(1-p)/3$	$(1-p)/3$	p	$(1-p)/3$
Former *ALLD*	$(1-p)/3$	$(1-p)/3$	$(1-p)/3$	P

At the end of the round agents have the option to adapt to their more successful neighbors, just like in the evolutionary case. However, instead of copying their opponents directly, the agents only change their probabilities for the four strategies. They will increase the weight of the strategy of their best performing neighbor with a constant c_A value. Then the weights of the strategies are renormalized, in order to create a probability distribution again.

That is, if $\sigma_M^t = \sigma^t_{\arg\max_{i\in E(a)}\pi_i^t}$ stands for the best strategy in the neighborhood, then the new strategy probability distribution of agent a will be the following[6]:

$$p_a^{t+1}(\sigma) = \begin{cases} \dfrac{p_a^t(\sigma)+c_A}{1+c_A} & ,\sigma=\sigma_M^t \\ \dfrac{p_a^t(\sigma)}{1+c_A} & ,\sigma\neq\sigma_M^t \end{cases} \qquad (4)$$

4. Cooperation's Sensitivity to Network Structure

In the following we focus on Watts-Strogatz networks, which start from a regular D-dimensional lattice with periodic boundary conditions and k-neighborhood (i.e., the agents are connected to their neighbors at most k-steps away on the lattice). Then with a certain p_{Rew} probability, each connection is 'rewired': i.e., it is replaced by a random link subject to the following constraints: neither self-loops nor multiple links between the same two nodes are allowed. (Watts, 1999) In particular, we study two-dimensional Watts-Strogatz networks (D=2) with 1-neighborhood (k=1). (Watts and Strogatz, 1998) The motivation for the latter is to keep the average number of neighbors (z=4) consistent with the interaction network topologies studied in the (Axelrod, Riolo and Cohen, 2000) and (Cohen, Riolo, Axelrod, 2001) papers. Furthermore, in accordance with the CAR studies, we fix the number of iterations at n_i=4.

Table 4. Virtual Experiment #1 varied the p_{Rew} rewiring parameter of Watts-Strogatz networks between 0 and 1 and the size of the network (N) between 100 and 10^4, creating 10 different (random) network instances for each value. The behavior of the evolutionary IPD framework was explored on each of these instances with 10 replications for T=5, R=3, P=1 and S=0.

		p_{Rew}										
		0.0	*0.1*	*0.2*	*0.3*	*0.4*	*0.5*	*0.6*	*0.7*	*0.8*	*0.9*	*1.0*
N	*100*	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	256	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	400	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	625	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	900	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	2500	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	3600	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	10000	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10

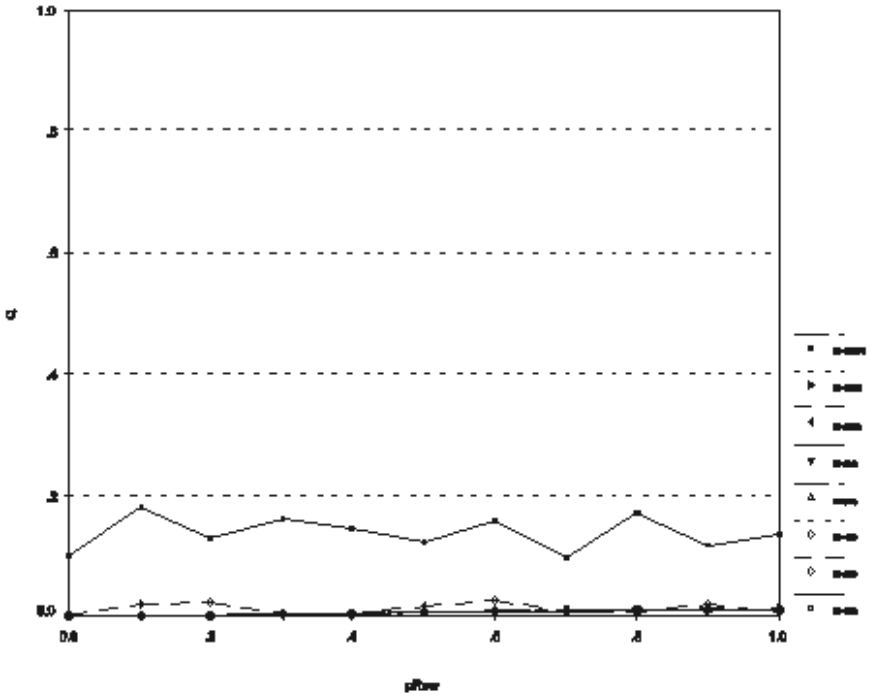

Figure 1.
Evolutionary adaptation's independence of network structure on Watts-Strogatz networks after 10^3 rounds. (Results from Virtual Experiment #1.) The points are averages for 10 network instances and 10 different runs for each network, given each combination of p_{Rew} and N. The values for the N=100 case show an averaging effect: all individual simulations converge to either 0 or 1.0.

5. Evolutionary Adaptation

In order to assess the cooperation's dependence on the structure of Watts-Strogatz networks, we performed a virtual computational experiment, as summarized on Table 4. Figure 1

demonstrates that the evolution of cooperation outcome is independent of the network structure on Watts-Strogatz graphs, using the evolutionary adaptation method. On the horizontal axis p_{Rew} is varied from 0 to 1. For each value 10 different random instances of the appropriate Watts-Strogatz network class is created and each of them is tested with 10 different random initial configurations. After 10^3 iterations, we count the number of agents playing the non-cooperative *ALLD* strategy and normalize the result with the system size.

The vertical axis of Figure 1 shows the average of the values obtained from the 10×10 simulations, defined as

$$Q = \frac{\sum_{a \in A} \chi\left(\sigma_a^{1000} = ALLD\right)}{N} \quad (5)$$

where $\chi(\cdot)$ yields 1 for true statements and 0 otherwise.

Table 5. Virtual Experiment #2 varied the p_{Rew} rewiring parameter of Watts-Strogatz networks between 0 and 1 and the size of the network (N) between 100 and 10^4, creating 10 different (random) network instances for each value. The behavior of the individually inductive IPD framework was explored on each of these instances with 10 replications for p=0.3, c_A=0.1, T=5, R=3, P=1 and S=0.

		p_{Rew}										
		0.0	*0.1*	*0.2*	*0.3*	*0.4*	*0.5*	*0.6*	*0.7*	*0.8*	*0.9*	*1.0*
N	*100*	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	256	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	400	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	625	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	900	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	2500	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	3600	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	10000	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10

6. Individually Inductive Adaptation

Following Virtual Experiment #1, we performed a similar experiment with the individually inductive adaptation model, summarized on Table 5. We found that, in contrast to the case of the evolutionary adaptation rule, the aggregate outcome is highly dependent on network structure, in much the same way as with other discrete choices on networks. Cooperation (i.e. *TFT*) is successful when agents in the network are 'far apart', while on small worlds (i.e., when the average agent-to-agent distance l scales with $\log N$) *ALLD* takes over.

On Figure 2 we observe a phase shift from the *TFT* dominated large worlds to the *ALLD* dominated small-worlds as the p_{Rew} probability of rewiring in the Watts-Strogatz network increases. The Q ratio of *ALLD* domination scales as a $f(N)=\alpha N + \beta$ linear function of the system size. Figure 2b illustrates the collapse of l as a function of p_{Rew}.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

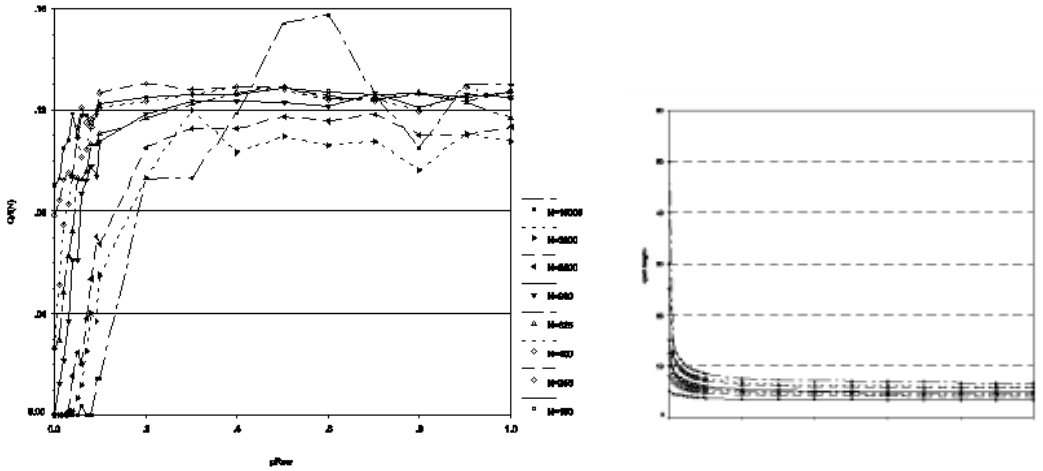

(a) (b)

Figure 2.
(a) Sensitivity to network structure with the individually inductive adaptation rule (***p***=0.3, c_A=0.1). (Results from Virtual Experiment #2.) Numerical results showing the *Q* dominance of *ALLD* outcomes after 10^3 rounds, scaling as a polinomial function of *N*, with $\alpha=5\cdot10^{-4}$, $\beta=7.8$ (see text). (b): Average path length (*l*) versus p_{Rew} in Watts-Strogatz networks. (Created by measuring the networks generated during Virtual Experiment #2. The points are averages of the 10 network instances for each combination of *N* and p_{Rew}.)

7. Cooperation's Sensitivity to Adaptation Speed

Figure 2 demonstrated cooperation's sensitivity to network structure, in the individually inductive adaptation framework, with the adaptation speed parameter (c_A) kept constant. In particular, the results show that the emergence of cooperation is dependent on the average path length in the interaction network. This suggests that cooperation is sensitive to the quickness at which information (strategy adaptation) can spread in the network, which is clearly also dependent on adaptation speed.

Virtual Experiment #3 was designed to confirm this hypothesis. (See Table 6.) As it is shown on Figure 3 (left column), the phase shift observed earlier (see Figure 2) appears for relatively low adaptation speeds only (i.e., $c_A \leq 0.4$). In particular, when the quickness of adaptation approximates that of the evolutionary adaptation rule (e.g., when $c_A > 1-p$) the model reproduces the strong network-independence result of the evolutionary adaptation rule. (Axelrod, Riolo and Cohen, 2000).

Table 6.
Virtual Experiment #3 varied the p_{Rew} rewiring parameter of Watts-Strogatz networks between 0 and 1; the speed of adaptation (c_A) parameter between 0.1 and 1.0; and the size of the network (*N*) over 400, 625 and 900. We created 10 different (random) network instances for each combination of *N* and p_{Rew}. The behavior of the individually inductive IPD framework was explored on each of these instances with all values of c_A with 10 replications for ***p***=0.3, *T*=5, *R*=3, *P*=1 and *S*=0.

N=400, 625, 900		c_A									
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
p_{Rew}	0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.1	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10

	0.2	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.3	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.4	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.5	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.6	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.7	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.8	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.9	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	1.0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10

Table 7.

Virtual Experiment #4 varied the p_{Rew} rewiring parameter of Watts-Strogatz networks between 0 and 1; the initial strategy probability bias (p) parameter between 0.25 and 0.35 and between 0.4 and 1.0; and the size of the network (N) over 400, 625 and 900. We created 10 different (random) network instances for each combination of N and p_{Rew}. The behavior of the individually inductive IPD framework was explored on each of these instances with all values of p with 10 replications for $\boldsymbol{c_A}$=0.3, T=5, R=3, P=1 and S=0.

N=400, 625, 900		p								
		0.25	0.26	0.27	0.28	0.29	0.3	0.31	0.32	0.33
p_{Rew}	0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.1	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.2	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.3	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.4	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.5	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.6	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.7	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.8	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.9	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	1.0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
N=400, 625, 900		p								
		0.34	0.35	0.4	0.5	0.6	0.7	0.8	0.9	1.0
p_{Rew}	0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.1	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.2	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.3	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.4	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.5	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.6	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.7	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.8	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.9	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	1.0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

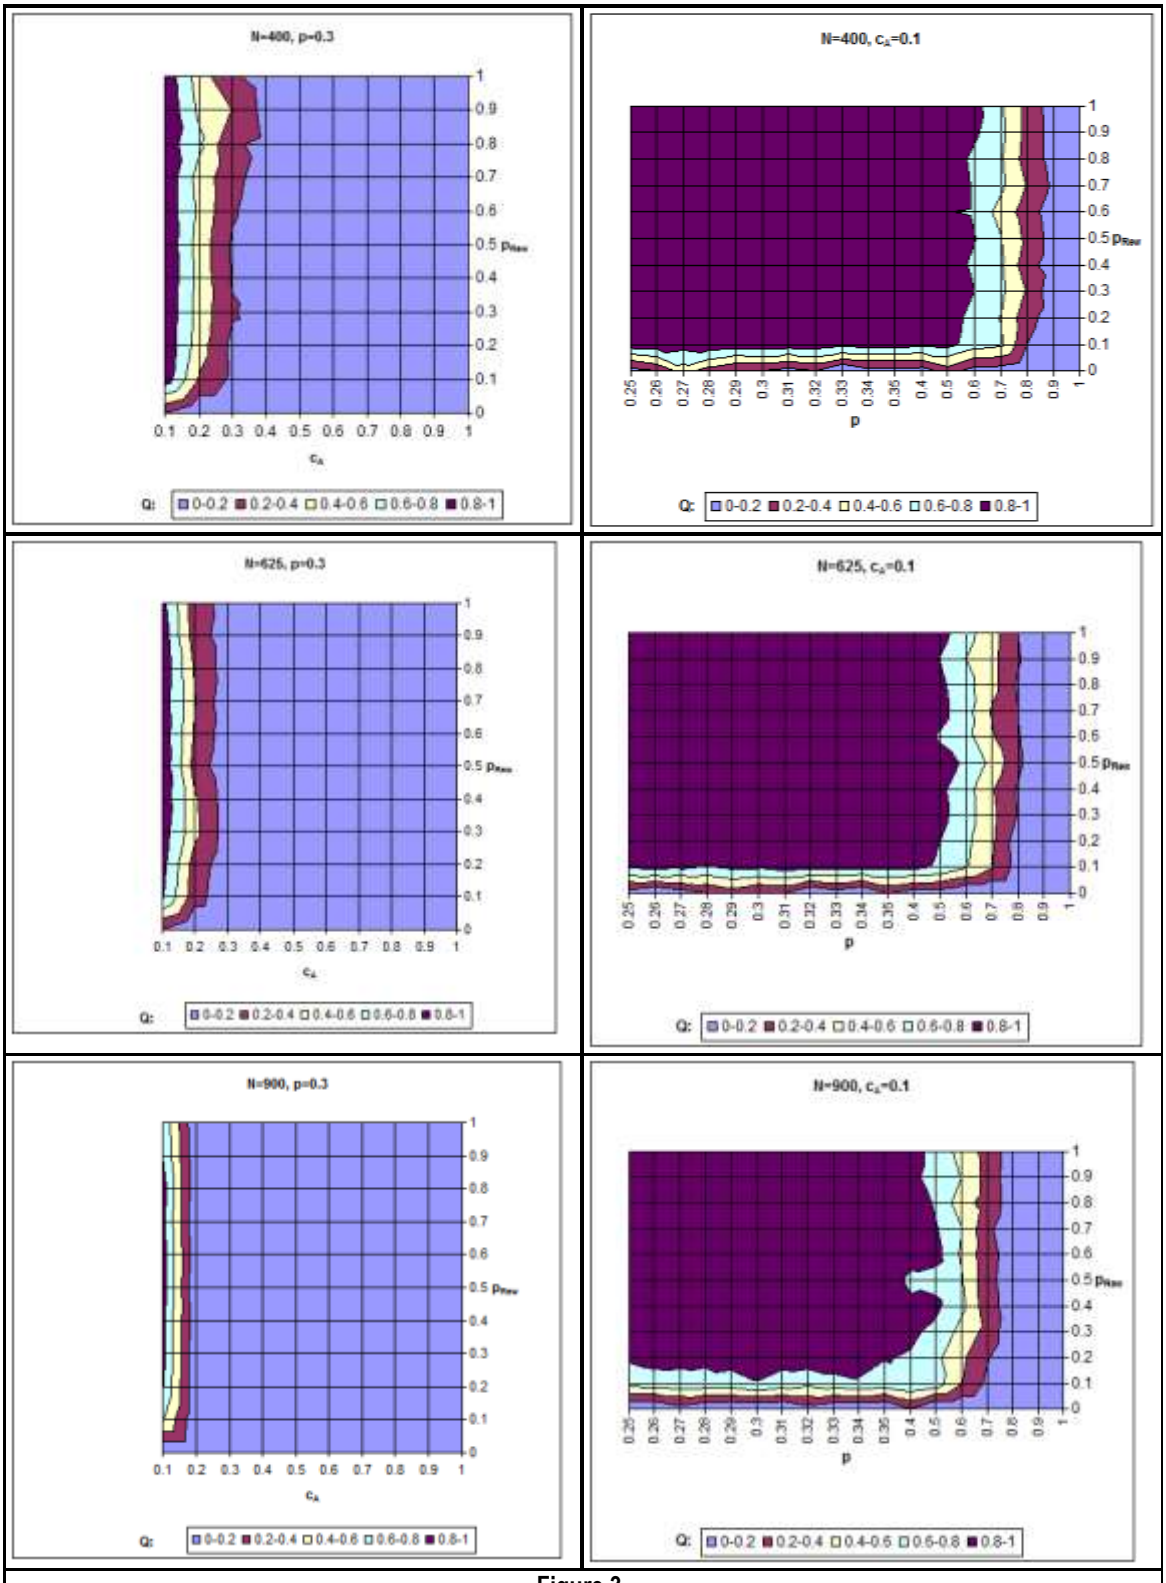

Figure 3.
Sensitivity to learning speed with the individually inductive adaptation rule. Results from Virtual Experiment #3 (left column) and Virtual Experiment #4 (right column), showing the *Q* dominance of *ALLD* outcomes after 10^3 rounds for

combinations of N, p_{Rew} and c_A or p. Each data point is an average over 10 replications per each of the 10 network instances per combinations of N and p_{Rew}.

Virtual Experiment #4 tested how these findings depend on the initial strategy probability bias (p) parameter, as summarized on Table 7. (The minimum p value tested was 0.2 as any value below 0.25 actually represents a negative bias for the 'initial strategy'.) The results of this virtual experiment are shown on Figure 3 (right panel) and they confirm the findings of Virtual Experiment #3. It is worth noting, however, that in case of both Virtual Experiment #3 and #4, the results have a pronounced dependence on system size (N).

Table 8.

Virtual Experiment #5 varied the initial strategy probability bias (p) parameter between 0.2 and 1.0; the speed of adaptation (c_A) parameter between 0.1 and 1.0; the size of the network (N) over 400, 625, 900 and 2500; while the p_{Rew} rewiring parameter of Watts-Strogatz networks took two different values: 0.2 and 0.9. We created 10 different (random) network instances for each combination of N and p_{Rew}. The behavior of the individually inductive IPD framework was explored on each of these instances with all combinations of p and c_A, with 10 replications for T=5, R=3, P=1 and S=0.

p_{Rew}=0.2, 0.9		c_A									
N=400, 625, 900, 2500		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
p	0.2	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.3	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.4	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.5	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.6	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.7	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.8	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	0.9	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10
	1.0	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10	10x10

In order to scrutinize our results so far, and to further explore the dependence of cooperation's success on adaptation speed, we designed Virtual Experiment #5 to vary both adaptation speed (c_A) and initial strategy probability bias (p) for various combinations of N and p_{Rew}. The values for p_{Rew} were selected to generate networks with the 'small-world' property, i.e., with low average path lengths (l). (See Table 8.) The results shown on Figure 4 confirm our earlier findings. In addition, they demonstrate that a sufficiently high initial bias (p) may also contribute to the emergence of cooperation. Notice that this is independent of the biased strategy in question, as the initial population consists of a roughly equal number of agents inclined for all the 4 studied strategies.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

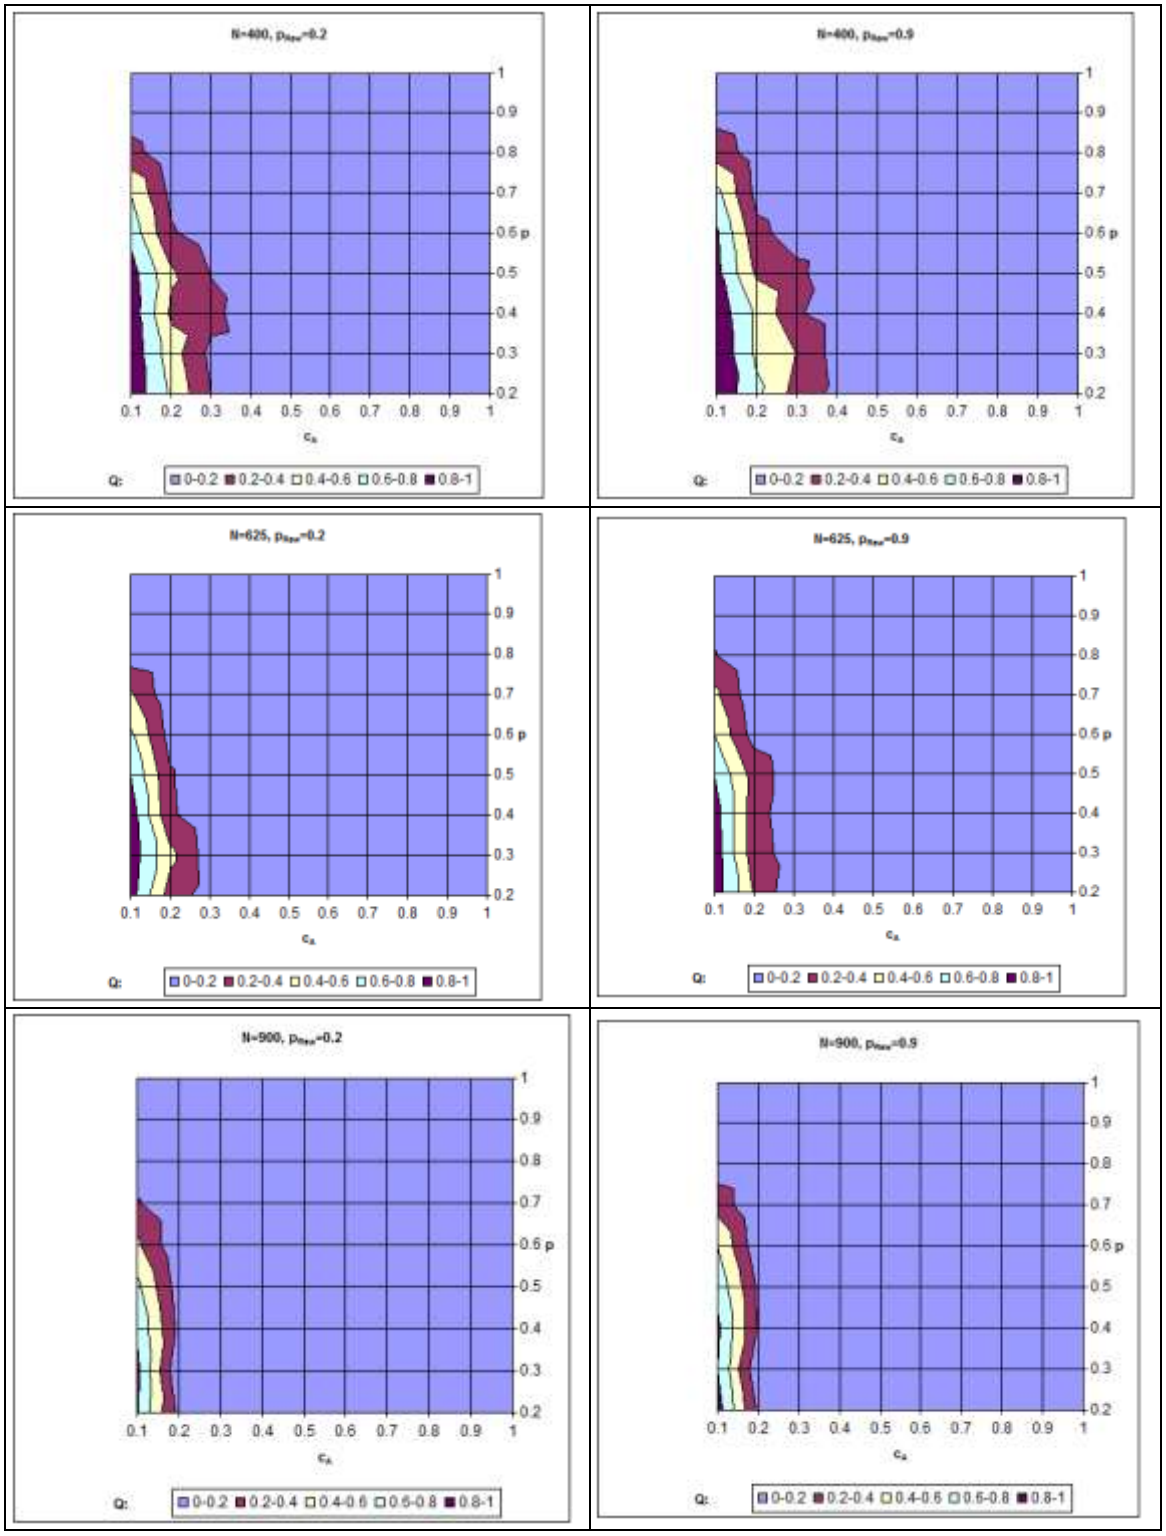

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

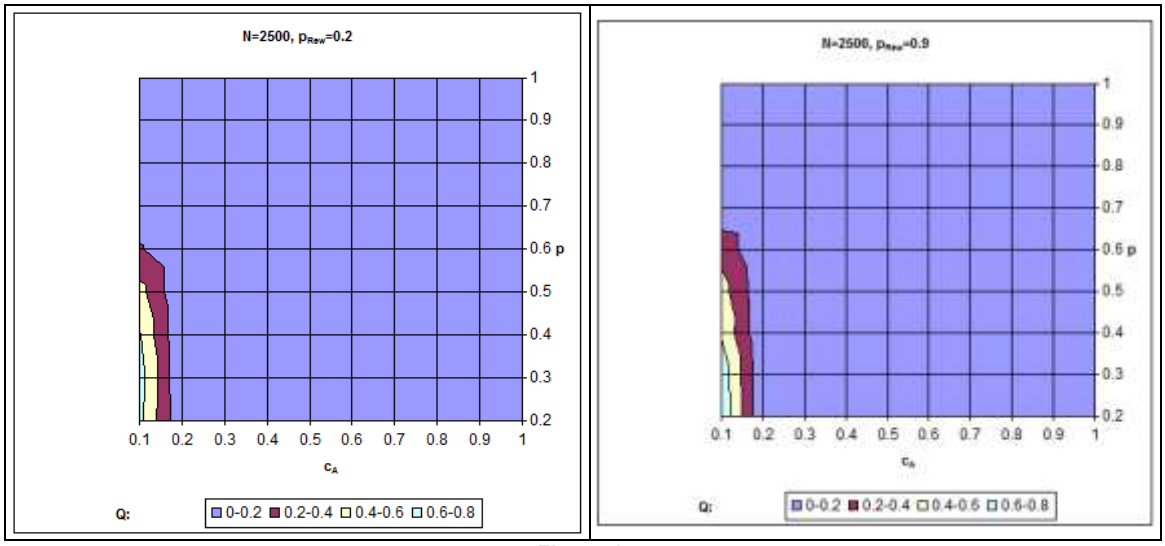

Figure 4.

Sensitivity to learning speed. Results from Virtual Experiment #5, showing the *Q* dominance of *ALLD* outcomes after 10^3 rounds. Each data point is an average over 10 replications per each of the 10 network instances per combinations of *N* and p_{Rew}.

The results of Virtual Experiments #3-#5 demonstrate that the emergence of cooperation does not only depend on the network structure as seen on Figure 2 (and in (Gulyas, forthcoming)), but also on adaptation speed. In particular, it depends both on the c_A and the p parameter of the individually inductive adaptation framework. These two parameters control, in essence, the number of reinforcement steps a strategy (σ) needs before its selection probability $p_a^t(\sigma)$ approaches 1. On the other hand, parameter p_{Rew} of Watts-Strogatz networks control the number of rewirings and consequently, the l average path length in the network. (Watts and Strogatz, 1998) (Newman, 2000) (Newman, 2003) These two observations imply that the emergence of cooperation may ultimately depend on how fast *TFT* can 'communicate' its success. That is, whether information travels fast enough via repeated adaptations of the agents to make the presence of fellow cooperators felt before agents facing defeat from *ALLD* partners 'freeze' into a non-cooperative strategy themselves.

Having a look at the time-evolution of single runs with long and short average path lengths, respectively, backs this intuition. As demonstrated by Figure 5, both runs start out in the same way: *ALLD* gaining ground. However, in a network with short path lengths, we observe a sudden come-back of cooperators.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

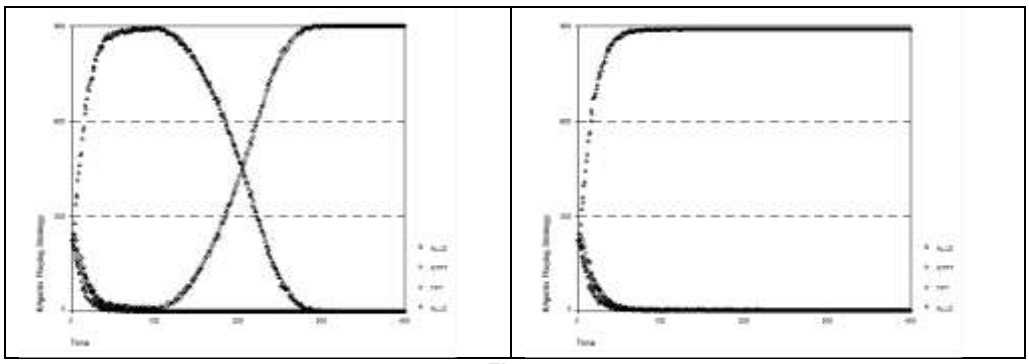

Figure 5.
Trajectories of two example runs from Virtual Experiment #2. The charts plot the time-evolution of the distribution of strategies played by the agents for p_{Rew} =0 (left) and p_{Rew}=1 (right). The other parameters were N=900, c_A=0.1 and p=0.3.

A rough analysis can further back our interpretation of the results. Initial cooperators (agents that are most likely to play *TFT* in the first round) have an initial $p_a^0(\sigma)=p$. In order to become a stable cooperator, this value needs to approximate 1. This is done via the observation of the success of *TFT* in the neighborhood (played either by one of the neighbors or by the agent itself). Each time this happens, $p_a^t(\sigma)$ is increased by c_A and then renormalized (divided by $1+c_A$). That is, omitting the normalization, an initial cooperator needs roughly $(1-p)/c_A$ reinforcements to become a stable cooperator. On the other hand, we know that in a network with the small-world property (i.e., when we observe a cooperative outcome on Figure 2 and Figure 3), the l average path length scales with the logarithm of network size. That is, the average number of hops that it takes to reach one agent from another via the links of IPD games (and observations of success performed during the adaptation phase) is around $ln(N)$. Therefore, our earlier intuition can be reformulated as the hypothesis that cooperation emerges when

$$\frac{1-p}{c_A} > \ln(N) \qquad (6)$$

Figure 6 shows the difference $\ln(N) - (1-\gamma \cdot p)/ \mu \cdot c_A$ in charts similar to that of Figure 4, for N=400, 625, 900 and 2500. We believe that the panels of Figure 6 are in good agreement with those on Figure 4. The match is not perfect, obviously, but that is to be expected, given the simplifications made during our analysis.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

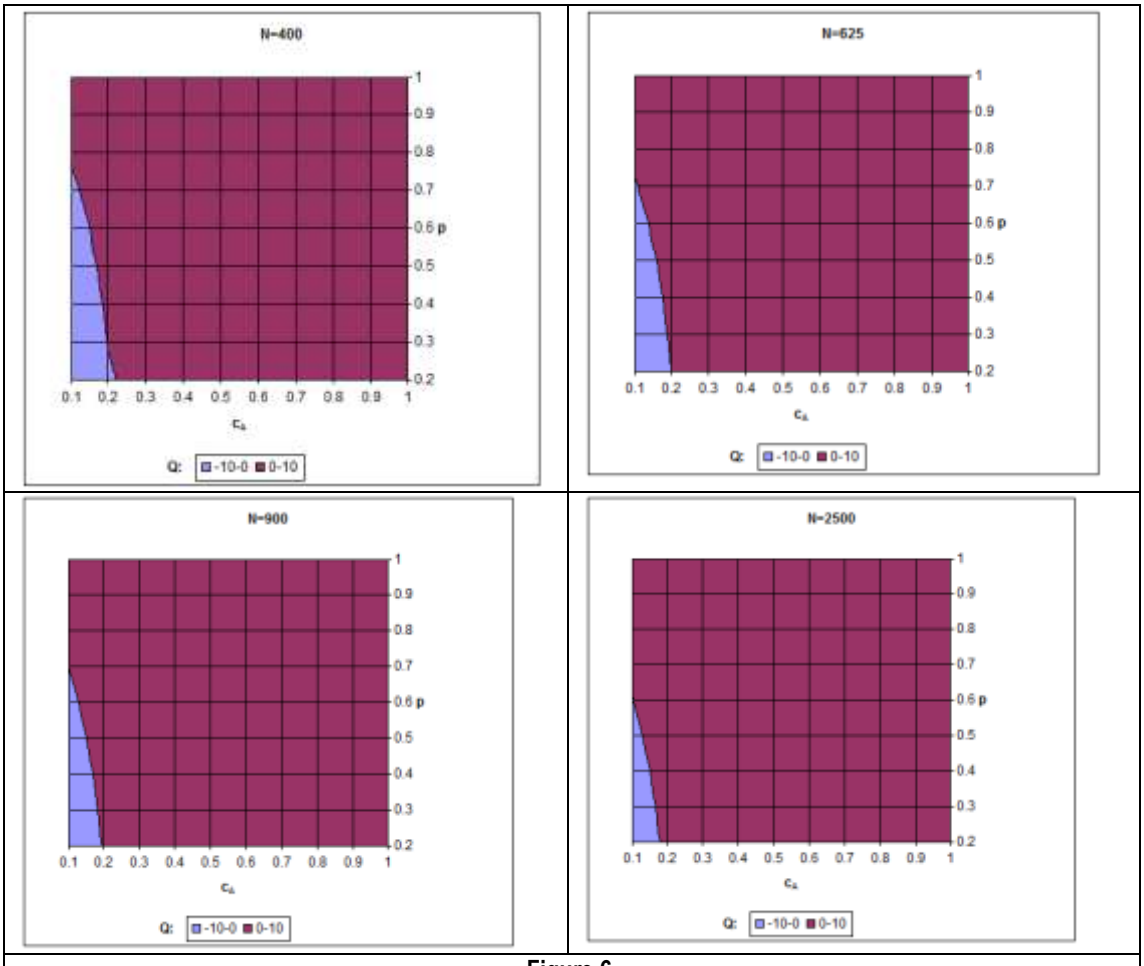

Figure 6.
Analytic results for the difference $\ln(N) - (1-\gamma p)/\mu c_A$, in case of γ=0.8 and μ=0.65, for N=400, 625, 900 and 2500; c_A values between 0.1 and 1.0 and values of p between 0.2 and 1.0.

Our approximate analysis offers two additional insights as well. First that the size of the non-cooperative regime (i.e., when *ALLD* comes to dominate) is indeed dependent on the N size of the population, because the l average path length also depends on it, albeit logarithmically. The second insight is that apparently, even for large networks, there are always combinations of c_A and p, where cooperation disappears. This is because

$$\lim_{c_A \to 0} \frac{1-p}{c_A} = \infty \qquad (7)$$

Therefore, for any fixed N and p, there is always a critical c_A^*, below which *ALLD* dominates the landscape. However, this also implies that the question of adaptation speed is only relevant in populations of moderate size, given that the interaction topology network has the small-world property. If the interaction topology is a 'large-world', that implies that the l average path length is greater than $\ln(N)$. Therefore, cooperation is even more common or, in other words, the critical adaptation speed below which *ALLD* wins is even lower than for small-worlds.

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

8. Related Works

The puzzle of altruism in a selfish, rational society is decades old. One of the first publications to discuss it in detail full of insight was (Axelrod, 1984). Since then many versions of the basic question have appeared together with many answers to address them. Axlerod's original answer was that altruism is due to the 'shadow of future', i.e., the rationally foreseen repetition of encounters, where selfish behavior dictates to measure immediate gains against future losses. Yet, many times this reciprocity argument fails to work due to various reasons, among which one obvious is that it assumes repeated encounters *with the same partner*. (Axelrod, Riolo and Cohen, 2000). This led to the development of theories on *indirect* reciprocity and the options for its spontaneous emergence. In such cases, agent A helps agent B and gets help from agent C. This framework naturally leads to models of reputation building, morality judgment and complex social interactions. (Nowak and Sigmund, 2005) However, at a more general level, explanations operating with indirect reciprocity exchange the assumption of *fixed partners* for the assumption of *identifiable partners*. That is because none of the above concepts can work without being able to tell cooperators apart from defectors. A generalization of this idea is today a busy and fruitful area of research: that of *tag-based systems*. (Holland 1993) (Hales and Edmonds, 2003). Tags are arbitrary, but recognizable labels assigned to agents that may serve to identify desirable and undesirable partners. Tags are therefore a means to establish patterns of interactions among agents. (Axelrod, Riolo and Cohen, 2000) (Riolo, 1997) And interaction patterns are of extreme importance, as noted by (Holland, 1995), because they "can have a strong effect not only on the success of individual agents, but also on the performance of the system as a whole". (Axelrod, Riolo and Cohen, 2000).

In this paper, we have refrained from allowing the agents to define interaction patterns by and for themselves. Instead, we focused on externally specified interaction topologies. Also, despite recent interest in dynamic networks, we limited our study to static interaction graphs, in order to focus on the context dependency issue emphasized by (Cohen, Riolo and Axelrod, 2001), who found that, as discussed earlier, the emergence of a cooperative regime is not dependent on the exact network structure, but on the stability of it. These results are also in good agreement with (Schweitzer and Mach, in preparation) and (Schweitzer, Mach, and Mühlenbein, 2005).

In another line of research (Santos, Pacheco and Lenaerts, 2006) found that heterogeneity in the number of connections the agents have may play a significant role in promoting of cooperation. (Szabo and Fath, 2006) further specified this finding by pointing out that it is subject to the additivity of the utility model (i.e., whether the payoffs are averaged over the number of games an agent plays). Similar results were discussed by Watts, in relation to his α- and β-graphs. (Watts, 1999). Even though the models discussed in this paper use non-additive utilities, we took the conservative approach and studied networks with near-homogenous degree distributions and with similar values of z average node degrees.

The focus of the present paper was to explore how the speed of adaptation contributes to the emergence of cooperation on Watts-Strogatz networks. While studying various methods of learning and adaptation is far from new in evolutionary game theory, these studies, to our knowledge, have not been linked to works exploring the effect of the interaction topology. (Fudenberg and Levine, 1998) (Axelrod, Riolo and Cohen, 2000) Interestingly, however, the issue

EQPAM Vol.2 No.1 January 2013
ISSN 2285 – 4916
ISSN-L 2285 - 4916

of adaptation speed was recently raised in the context of changing one's *tag* in tag systems. Hales argued that the quickness of tag change facilitates cooperation, and even went as far as arguing that perhaps this is a necessary condition for tag systems to produce high levels of cooperation (Hales, 2005).

9. Conclusions

The rationality of human decision-making is a wide-spread and natural assumption. However, it is in apparent contrast with the richness of altruistic behavior often observed at various levels of society. The puzzle of how cooperation may spontaneously emerge in a population of selfish, rational agents is old and yet, it continues to raise novel answers to novel questions. (Axelrod,1984) (Nowak and Sigmund, 2005) In this paper we have joined these series of works by studying an individually inductive version of an evolutionary IPD games framework. Our work was originally motivated by an apparent contradiction between the emergence of cooperation literature and results on discrete choice dynamics on social influence networks. In an earlier publication we resolved this issue by 'smoothing' the sharpness of strategy adaptation. (Gulyas, forthcoming) That result, however, prompted the question of how, in detail, the cooperative outcome depends on adaptation speed. This is what we have explored in this paper in the case of 'small-world' social network topologies introduced by Watts and Strogatz (Watts and Strogatz, 1998).

We have found that the presence of a cooperative regime (where almost the entire population plays Tit-for-Tat) is dependent on the quickness of information spreading across the network. More precisely, cooperation hinges on the relation between individual adaptation speed and average path length in the interaction topology. Our results are in good agreement with previous works both on discrete choice dynamics on networks and in the evolution of cooperation literature.

Our results suggest that there are two essential roads to cooperation on static networks: one is a 'large world' (when the average pair-to-pair distance is long), the other is quick adaptation. And since real social networks are known to be small-worlds, the common occurrence of altruism points to the effectiveness of human learning (Milgram, 1967; Watts, 1999; Newman, 2003).

In interpreting their results, Axelrod, Riolo and Cohen quote Putnam's definition of social capital that is defined in terms of the "connections among individuals – social networks and the norms of reciprocity and trustworthiness that arise from them." (Axelrod, Riolo and Cohen, 2000) (Putnam, 2000) They then argue that their finding on cooperation's independence of network structure should calm skeptics, as it suggests that social relationships may survive, even in the era of Internet and fast electronic communication. Following this line of thought, our results bring more ambiguous news. Shorter distances may indeed eliminate cooperation if adaptation is not quick enough. However, in the particular model studied here, 'quick enough' is not very demanding. At least, not in networks of realistic size.

Acknowledgements

The partial support of the GVOP-3.2.2-2004.07-005/3.0 (ELTE Informatics Cooperative Research and Education Center) and the *OMFB-01076/2006* (*Mecenatúra*) grants of the Hungarian Government is gratefully acknowledged. Earlier versions of this paper were presented at EXYSTENCE Topical Workshop on "*Trust-Based Networks and Robustness in Organizations*", ETH Zurich, 2006 and at the International Conference on Social Networks at Collegium Budapest, Hungary, 2006.

References

1. Axelrod, R. (1984), Evolution of Cooperation, NY: Basic Books.
2. Axelrod, R., Riolo, R. L., Cohen, M. D. (2000), "Beyond Geography: Cooperation with persistent links in the absence of clustered neighborhoods", Personality and Social Psychology Review, 6, 341-346.
3. Cederman, L.-E. (1997): Emergent Actors in World Politics, Princeton University Press, Princeton, NJ, USA, p. 290.
4. Cohen, M. D., Riolo, R. L., Axelrod, R. (2001), „The Role of Social Structure in the Maintenance of Cooperative Regimes", Rationality and Society, 13, pp. 5-32.
5. Cohen, Y. (1994): Radicals, Reformers and Reactionaires: The Prisoner's Dilemma and the Collapse of Democracy in Latin America, The University of Chicago Press, Chicago, IL
6. Dugundji, E. R., Gulyas, L. (2006), "Socio-Dynamic Discrete Choice on Networks in Space: Impacts of Agent Heterogeneity on Emergent Outcomes", Workshop on Agent-based models of market dynamics and consumer behaviour, Institute of Advanced Studies, University of Surrey, UK
7. Fudenberg, D., Levine, D. K. (1998), Theory of Learning in Games, Cambridge, MA, MIT Press.
8. Gowa, J., Mansfield, E. D. (1993), "Power Politics and International Trade.", American Political Science Review, Vol. 87, No. 2 (Jun., 1993), pp. 408-420.
9. Gulyas, L. (forthcoming), "Cooperation in Networked Populations of Selfish Adaptive Agents: Sensitivity to Network Structure", Physica A
10. Gulyas, L., Dugundji, E. R. (2003), "Discrete Choice on Networks: An Agent-Based Approach", North American Association for Computational Social and Organizational Science Conference (NAACSOS 2003), Omni William Penn, Pittsburgh, PA
11. Gulyas, L., Platkowski, T. (2004): On Evolutionary 3-Person Prisoner's Dilemma Games on 2-D Lattice, In P.M.A. Sloot, B. Chopard, and A.G. Hoekstra (Eds.), Proceedings of the Sixth International conference on Cellular Automata for Research and Industry (ACRI 2004), LNCS 3305, Springer-Verlag, Berlin-Heidelberg, pp. 831-840.
12. Hales, D. (2005), "Change your tags fast! : A necessary condition for cooperation?", Proceedings of the Workshop on Multi-Agents and Multi-Agent-Based Simulation (MABS 2004), LNAI 3415, Springer, Germany
13. Hales, D. and Edmonds, B. (2003), "Evolving Social Rationality for MAS using 'Tags'", In Rosenschein, J. S., et al. (Eds.) Proceedings of the 2nd International Conference on Autonomous Agents and Multiagent Systems, Melbourne, July 2003 (AAMAS03), ACM Press, pp. 497-503.
14. Herrero, C. P. (2002), "Ising model in small-world networks", Physical Review E, Vol. 65., pp. 066110
15. Holland, J. H. (1995), Hidden Order: How Adaptation Builds Complexity, Reading, MA, Addison Wesley
16. Holland, J. H. (1993), "The Effect of Labels (Tags) on Social Interactions", Santa Fe Institute Working Paper 93-10-064. Santa Fe, NM.
17. Klüver, J., Schmidt, J. (1999), "Topology, Metric and Dynamics of Social Systems", Journal of Artificial Societies and Social Simulation, Vol. 2, No.3, http://www.soc.surrey.ac.uk/JASSS/2/3/7.html
18. Milgram, S. (1967), "The Small World Problem", Psychology Today, Vol. 1 (May), pp. 62-67.
19. Newman, M. (2000), "Models of the Small World: A Review", J. Stat. Phys., 101, pp. 819-841.
20. Newman, M. (2003), "Structure and Function of Complex Networks", SIAM Review, 45(2):167-256.
21. Nikoletopoulos, T., Coolen, A. C. C., Perez Castillo, I., Skantzos, N. S., Hatchett, J. P. L., Wemmenhove, B. (2004), "Replicated transfer matrix analysis of Ising spin models on 'small world' lattices", J. Phys. A: Math. Gen. 37, pp. 6455-6475.
22. Nowak, M. A., Sigmund, K. (2005), "Evolution of Indirect Reciprocity", Nature, Vol. 437, 27 October 2005, pp. 1291-1298.

23. Putnam, R. D. (2000), Bowling Alone: The Collapse and Revival of American Community, New York, Simon and Schuster.
24. Pruitt, D. G. (1998), "Social Conflict", In D.T. Gilbert, S.T. Fiske and G. Lindzey (Eds.), Handbook of Social Psychology, 4th Edition. Boston, McGraw-Hill
25. Riolo, R. L. (1997), "The Effects and Evolution of Tag-Mediated Selection of Partners in Populations Playing the Iterated Prisoner's Dilemma", In T. Back (Ed.), Proceedings of the 7th International Conference on Genetic Algorithms (ICGA97), San Francisco, CA, Morgan Kaufmann, pp. 378-385.
26. Santos, F. C., Pacheco, J. M., Lenaerts, T. (2006), "Evolutionary dynamics of social dilemmas in structured heterogeneous populations", Proceedings of the National Academy of Sciences, Vol. 103, No. 9, pp. 3490-3494
27. Schweitzer, F., Mach, R., and Mühlenbein, H (2005), „Agents with heterogeneous strategies interacting in a spatial IPD", in T. Lux, S. Reitz, and E. Samanidou (Eds.) Nonlinear Dynamics and Heterogeneous Interacting Agents, Lecture Notes in Economics and Mathematical Systems, Vol. 550, Springer Berlin, pp. 87-102.
28. Schweitzer, F., Mach, R. (in preparation), "Adaptation of Cooperation: Random vs. Local Interaction in a Spatial Agent Model", Personal Communications.
29. Szabo, G., Fath, G. (2006), "Evolutionary games on graphs", electronic source, arXiv.org:cond-mat/0607344
30. Watts, D. (1999), Small-worlds: The dynamics of Networks between order and randomness, Princeton University Press
31. Watts, D. (2002), "A Simple Model of Global Cascades on Random Networks", Proceedings of the National Academy of Sciences, Vol. 99, no. 9, April 03, 2002, 5766-5771.
32. Watts, D.J. and Strogatz, S.H. (1998), "Collective dynamics of 'small-world' networks", Nature (393)
33. Weibull, J. W. (1995), Evolutionary game theory, MIT Press.

Footnotes

1 Department of History and Philosophy of Sicence, Loránd Eötvös University, 1117 Budapest, Pázmány Péter sétány 1/c, Hungary}

2 AITIA International Inc., 1039 Budapest, Czetz János utca 48-50, Hungary

3 Henceforth, we adopt the following standard values: T=5, R=3, P=1, S=0.

4 This is a discretized version of the model discussed by Cohen et al. Their results apply to a wider set of strategies than discussed here.

5 Therefore, the proposed scheme is *not* a mixed strategy.

6 Notice that this adaptation rule has a nice convergence property when subjected to repeated reinforcement.