

Coordination Games on Small-Worlds: Artificial Agents Vs. Experiments

Enea Pestelacci¹, Marco Tomassini¹, and Alberto Antonioni¹

¹Information Systems Department, HEC, University of Lausanne, Switzerland
alberto.antonioni@unil.ch

Abstract

Effective coordination is a key social ingredient and social structure may be approximated by networks of contacts. Using Stag Hunt games, which provide socially efficient and inefficient equilibria, we compare our simulation results using artificial players and evolutionary game theory with laboratory experimental work with human subjects on small-world type networks and with theoretical results. The conclusion is that the apparently encouraging results obtained in the few human experiments in which the local interaction structure seems to promote efficient equilibria, is neither supported by simulation results nor by theoretical ones.

Introduction

Many types of conflicting interactions between agents in biology and society can be usefully described with the tools of Game Theory (Vega-Redondo, 2003). The Prisoner's Dilemma and the Hawk-Dove games are well known metaphor for representing the tension that appears in society when individual objectives are in conflict with socially desirable outcomes, and most of the vast research literature has focused on conflicting situations in order to uncover the mechanisms that could lead to cooperation instead of socially harmful interactions. However, in many important situations in society agents are not required to use aggressive strategies. In fact, many frequent social and economic activities simply require individuals to coordinate their actions on a common goal since in many cases the best course of action is to conform to the standard behavior. For example, if one is used to drive on the right side of the road and travels to a country where the norm is reversed, it pays off to follow the local norm. Games that express this extremely common kind of interactions are called *coordination games*.

Coordination games are apparently simple but they confront the players with multiple Nash equilibria (NE) and thus with the problem of how to choose among them. Evolutionary game theory (EGT) offers a dynamical view which is based on concepts of positively selecting fitter variants in the population, i.e. strategies that score best are more likely to survive and provides a justification for the appearance of stable states of the dynamics that represent solutions of the

game (Vega-Redondo, 2003).

For mathematical convenience, standard EGT is based on infinite mixing populations where pairs of individuals are drawn uniformly at random at each step and play the game. Correlations are absent by definition and the population has an homogeneous structure. However, everyday observation tells us that in animal and human societies, individuals usually tend to interact more often with some specified subset of partners; for instance, teenagers tend to adopt the fashions of their close friends group; closely connected groups usually follow the same religion, and so on. In short, social interaction is mediated by networks, in which vertices identify people, firms etc., and edges identify some kind of relation between the concerned vertices such as friendship, collaboration, economic exchange and so on. Thus, locality of interaction plays an important role. Recently, in the wake of a surge of activity in network research in many fields (Newman, 2003), the dynamical behavior of games on networks that are more likely to represent actual social interactions than regular grids has been investigated (see (Szabó and Fáth, 2007; Roca et al., 2009) for comprehensive recent reviews). These studies have been conducted on games of conflict such as the Prisoner's dilemma or the Hawk-Dove in most cases and have shown that there are network structures, such as scale-free and actual social networks that may favor the emergence of cooperation with respect to the fully mixing populations used in the theory (Santos et al., 2006; Luthi et al., 2008; Roca et al., 2009). Recently, some work has been done following this approach for games of the coordination type too to try to unravel the effect of structure on the population behavior, e.g. (Roca et al., 2009; Tomassini and Pestelacci, 2010).

Several analytically rigorous results are available for coordination games in well-mixed populations (Kandori et al., 1993), as well as populations with a simple local interaction structure such as rings and grids (Ellison, 1993; Morris, 2000). These results are very useful; however, while game theory has normative value, its prescriptions are not always reflected in the way people act when confronted with these situations. This has been made manifest by a host of results

of experiments with people (Camerer, 2003). Coordination games are no exception and also confront the theory with many puzzles. For coordination games on small-worlds and regular networks the recent laboratory experiments carried out in (Cassar, 2007) and in (My et al., 1999; Keser et al., 1998) are particularly relevant.

It has been argued that multi-agent learning simulations have the potential for greatly improving our knowledge of the game-theoretical interactions in artificial societies (Shoham et al., 2007). We also believe that numerical simulations, with their possibility of modeling many different situations, may shed light on the factors, both endogenous such as strategy update policy and exogenous, such as population structure, that have an influence on the game outcome. In this way, this can be a valuable tool to experiment in both the theoretical and experimental sides and to build a bridge between the two.

The paper is organized as follows. In the next section we present a brief introduction to the subject of coordination games. Then we describe the dynamical model, as well as the main results obtained in previous work. The following sections deal with the main theme of the present study, namely, the relationship between recent experimental results and our simulations. Finally, we present our conclusions.

Coordination Games

General two-person, two strategies coordination games have the normal form of Table 1. With $a > d$ and $b > c$, (α, α) and (β, β) are both Nash equilibria. Now, if we assume that $a > b$ and $(a - d) \leq (b - c)$ then (β, β) is the risk-dominant equilibrium, while (α, α) is the Pareto-dominant one (Harsanyi and Selten, 1988). This simply means that players get a higher payoff by coordinating on (α, α) but they risk less by using strategy β instead. There is also a third equilibrium in mixed strategies but it is evolutionarily unstable. A well known example of games of this

	α	β
α	a, a	c, d
β	d, c	b, b

Table 1: A general two-person, two strategies coordination game.

type are the so-called Stag-Hunt games (Skyrms, 2004). This class of games has been extensively studied analytically in an evolutionary setting (Kandori et al., 1993; Ellison, 1993) and by numerical simulation on several model network types (Skyrms, 2004; Luthi et al., 2008; Roca et al., 2009).

Evolutionary Games on Networks

The network of agents will be represented by an undirected graph $G(V, E)$, where the set of vertices V represents the

agents, while the set of edges (or links) E represents their symmetric interactions. The population size N is the cardinality of V . A neighbor of an agent i is any other agent j at distance one from i . The set of neighbors of i is called V_i and its cardinality is the degree k_i of vertex $i \in V$. The average degree of the network is called \bar{k} .

Strategy Revision Rules

Since we shall adopt an evolutionary approach, we must define the decision rules by which individuals will update their strategy during time. Let $\sigma_i \in \{\alpha, \beta\}$ be the current strategy of player i and let us call M the payoff matrix of the game, see Table 1. The quantity

$$\Pi_i(t) = \sum_{j \in V_i} \sigma_i(t) M \sigma_j^T(t)$$

is the accumulated payoff collected by agent i at time step t and $\sigma_i(t)$ is a vector giving the strategy profile at time t .

Here we shall describe two among the most commonly used strategy revision rules. These rules, although they are extremely simple, also make sense when human players are concerned, at least at a very low level of knowledge and information processing capabilities. The first rule is to switch to the strategy of the neighbor that has scored best in the last time step. This *imitation of the best* policy can be described in the following way: the strategy $\sigma_i(t)$ of individual i at time step t will be

$$\sigma_i(t) = \sigma_j(t - 1),$$

where

$$j \in \{V_i \cup i\} \text{ s.t. } \Pi_j = \max\{\Pi_k(t - 1)\}, \forall k \in \{V_i \cup i\}.$$

That is, individual i will adopt the strategy of the player with the highest payoff among its neighbors including itself. If there is a tie, the winner individual is chosen uniformly at random, but otherwise the rule is deterministic.

At a slightly higher sophistication level, a well known adaptive learning rule is *myopic best-response* (Young, 1998), also called best-reply, which embodies a primitive form of bounded rationality and for which rigorous results are known. In the local version of this model, time is discrete i.e. $t = 0, 1, 2, \dots$ and, at each time step, an agent has the opportunity of revising her current strategy with probability p . She does so by considering the current actions of her neighbors and switching to the action that would maximize her payoff if the neighbors would stick to their current choices. In other words, $\hat{\sigma}_i$ is a best response for player i if $\Pi_i(\hat{\sigma}_i(t)) > \Pi_i(\sigma_i(t)), \forall \sigma_i$. In case of a tie, agent i keeps its current strategy.

The model is thus completely local and an agent only needs to know her own current strategy, the game payoff matrix, who are her neighbors, and their current strategies. This

rule is called myopic because the agents only care about immediate payoff, they cannot see far into the future. In order to introduce some stochasticity, an agent can make a mistake with some small probability q . These small random effects are meant to capture various sources of uncertainty such as deliberate and involuntary decision errors. Deliberate errors might play the role of experimentation, and involuntary ones might be linked with insufficient familiarity with the game, for example. This dynamic will be called best-response with noise.

The simulation represents a dynamical system in which time t is discrete, i.e. $t = 0, 1, \dots$. Let us call $\Sigma(t) = (\sigma_1(t), \dots, \sigma_N(t))$ the strategy profile at time t . For the imitation of the best and best response rules the evolution of $\Sigma(t)$ is deterministic. In the best response with noise case the resulting process is stochastic. It can be described by a Markov chain (Kandori et al., 1993) since the probability of strategy profile $\Sigma(t) = (\sigma_1(t), \dots, \sigma_N(t))$ at time step $t + 1$ only depends on the previous time step:

$$Pr(\Sigma(t + 1) | \Sigma(t), \Sigma(t - 1), \dots) = Pr(\Sigma(t + 1) | \Sigma(t)).$$

It is clear that more refined forms of learning, such as reinforcement learning could be used to represent the agents' decisions (Camerer, 2003). However, these more sophisticated approach do not have yet a firm theoretical basis and could not be compared with baseline dynamical models. This is the reason why, in the interest of simplicity, we stick with very simple basic protocol revision rules here.

Summary of Previous Simulation and Theoretical Results

This section summarizes previous numerical results on Stag Hunt games. Several populations topologies have been studied, including regular lattices (Skyrms, 2004; Roca et al., 2009), random graphs (Roca et al., 2009; Luthi et al., 2008), scale-free graphs (Luthi et al., 2008; Roca et al., 2009), model and actual social networks (Luthi et al., 2008) using several strategy update rules such as replicator dynamics, imitation of the best, and best response dynamics. In the average, for initially equidistributed strategies, at the steady state the population is monomorphic, with all individuals playing α or β . For all network types, the more efficient α strategy is enhanced with respect to what would happen in a mixing population. This is true for all update rules except best reply, for which the topology does not seem to play an important role (Roca et al., 2009). Social networks also favor the Pareto-efficient outcome in the average but the steady state population is often dimorphic, i.e. there is a mix of the two strategies. The reason why there can be mixed states in social networks has been attributed (Tomassini and Pestelacci, 2010) to the presence of *communities*. In fact, social networks can usually be partitioned into recognizable clusters (Newman, 2003); within these clusters strategies

may become dominant as in the pure coordination case just by chance. In other words, as soon as a strategy dominates in a given cluster, it is difficult to eradicate it from outside since other communities, being weakly connected, have little influence.

We now briefly comment on the relationship between the results of numerical simulations and well known theoretical results on Stag-Hunt games (for a recent review see (Weidenholzer, 2010)). These theoretical models are based on ergodic stochastic processes in very large well mixed populations and state that, when using best-response dynamics in random two-person encounters, and in the presence of a little amount of noise, both for well mixed populations as well as for populations structured as rings, the risk-dominant strategy should take over the population in the long run (Kandori et al., 1993; Ellison, 1993). But coordination seems to be sensitive to the exact type of revision protocol and dynamic. For example, Robson and Vega-Redondo (Robson and Vega-Redondo, 1996) found that the Pareto-dominant equilibrium is selected if players are immediately randomly rematched after each encounter.

Simulations results on networked populations indirectly confirm the above, i.e., at the steady state there is always either a single strategy, but not necessarily the risk-dominant one. However, owing to network reciprocity effects related to clustering, a mix of both strategies is also possible. In summary, it can be said that network effects tend to reinforce cooperation on the Pareto-dominant case, which is a socially appreciable effect. However, these results must be taken with a grain of salt. Numerically studies deal with finite, network-structured populations during a limited amount of time, while theoretical results have been established for large well mixed populations in the very long run. Thus, numerical results and theoretical predictions based on different assumptions do not necessarily agree with each other.

Discussion of Some Experimental Results on Coordination Games

In this section we comment on some experimental results on coordination games in the light of the conclusions that have been reached by numerical simulation and also with respect to theoretical results. There have been many experiments in the field and we cannot be exhaustive; however, the main conclusions are the following. When the analog of a (finite and generally small) well mixed population of players have been used, the general result is that polymorphic final states are rare, the initial state of play i.e. the strategy played at the first period is a good predictor of convergence, and the risk-dominant equilibrium is often reached in the laboratory, i.e. coordination failures emerge, although in some cases, especially in finitely repeated games and by varying the payoff structure, coordination on the Pareto-efficient equilibrium can also be achieved (Cooper et al., 1992; Bat-

talio et al., 2001). It has also been observed that the number of rounds, the size of the group, and the fact of playing repeatedly with the same player may influence the result, i.e. small groups, higher number of rounds, and repeated interactions have been shown to favor the Pareto-efficient outcome (Huyck et al., 1993, 1990). This lends support to the idea that human agents play the games using some imperfect decision rules that, nonetheless, may be similar to some variant of myopic best response, perhaps with a longer memory of past encounters instead of just one step behind. No doubt, human decision-making is a lot more complex, but simple learning rules should somehow evolve during these experiments.

A more interesting situation from the point of view of the present work is the one in which some more specific population structure has been recreated in the laboratory setting. We are aware of three experiments of this type, the work of (Cassar, 2007), and the studies of (My et al., 1999) and of (Keser et al., 1998).

Keser et al. used a ring structure where each player has a neighbor on either side and a well mixed structure for comparison. Their conclusions were that in the ring the preferred equilibrium is the risk-dominant one, while the payoff-dominant equilibrium was the more frequent result in the globally communicating population. This is in qualitative agreement with the theoretical predictions of (Ellison, 1993) for the ring and (Robson and Vega-Redondo, 1996) for the mixing case.

My et al. performed a comparative experimental study of Stag Hunt games with three different payoff matrices on mixing and structured populations. The population with local structure was composed by a circle of eight people where each player only interacts with her immediate right and left neighbors. They find that the first period modal choice of strategy, which is the payoff dominant one, plays a major role in the final outcome. In the global population case, the steady state generally lies in the same basin of attraction as the initial state. This result, which is commonly observed in many laboratory experiments, does not agree with the theoretical results of (Kandori et al., 1993) which predict that all the probability at stochastic equilibrium be placed on the risk-dominant state. However, we have to bear in mind that the latter have been established for stochastic processes in the very long run taking place in large populations and neither of these conditions can be satisfied in a laboratory setting.

For the ring structure, the convergence to the risk-dominant outcome is more frequent than in the well mixed case, especially when the payoff matrix values are such that the Pareto-superior basin shrinks. However, still often times the system converges to the Pareto-dominant state, which disagrees with the theoretical predictions of (Ellison, 1993) based on noisy best reply dynamics. By examining the detailed history of play, the experimenters have found that, while in the

global population subjects on average play myopic best response, in the ring with local structure a kind of imitation rule fits the data better than best reply. This is in qualitative agreement with the very extensive numerical studies of (Roca et al., 2009), where the simple strategy of imitating the individual having the best payoff in the neighborhood is the one that best promotes cooperation in the Stag Hunt played on several classes of networks.

The experimental Study of Cassar and its Numerical Simulation

The study of Cassar (Cassar, 2007) is the most interesting one from the standpoint of the present paper as it investigates network structures that are more realistic than the ring and the two-dimensional lattice, although the ring is also used in the experiments for comparison. One particular Stag Hunt payoff matrix is used in (Cassar, 2007) with the following payoff values (see Table 1): $a = 5$, $b = 1$, $c = -1$, $d = 4$. With this choice the frequency α of stag players at the (unstable) mixed equilibrium would be $2/3$, since at this point the expected value playing strategy α , $E[\alpha]$, is equal to the expected value $E[\beta]$, which implies $5p - (1-p) = 4p + 1 - p$, i.e. $p = 2/3$, where p is the probability with which α is played in the mixed strategy or, equivalently, the α fraction in the population. This leads to a basin of attraction for the payoff-dominant strategy α that is half the size of the corresponding basin for the risk-dominant strategy.

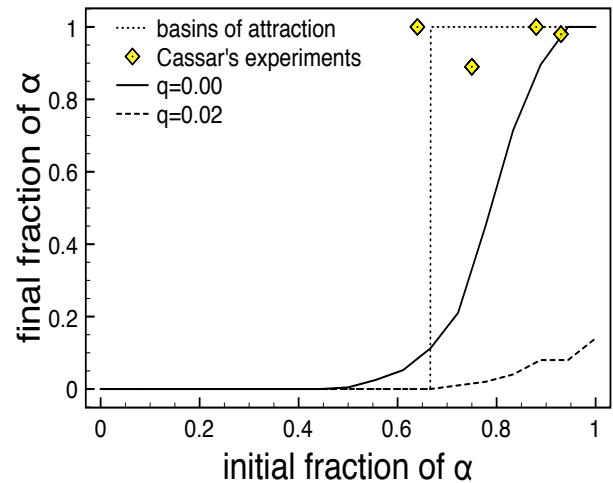


Figure 1: Final average ratio of α -players as a function of their initial ratio in small-world networks of size $N = 18$ and $\bar{k} = 4$. With noise (dashed curve) the system converges almost always to the risk-dominant steady state. Without noise (continuous curve) the payoff-dominant steady state is often reached when the initial ratio is in the corresponding basin of attraction. The dotted line marks the theoretical Nash equilibria and their basins of attraction. The small squares represent the results of Cassar's experiments.

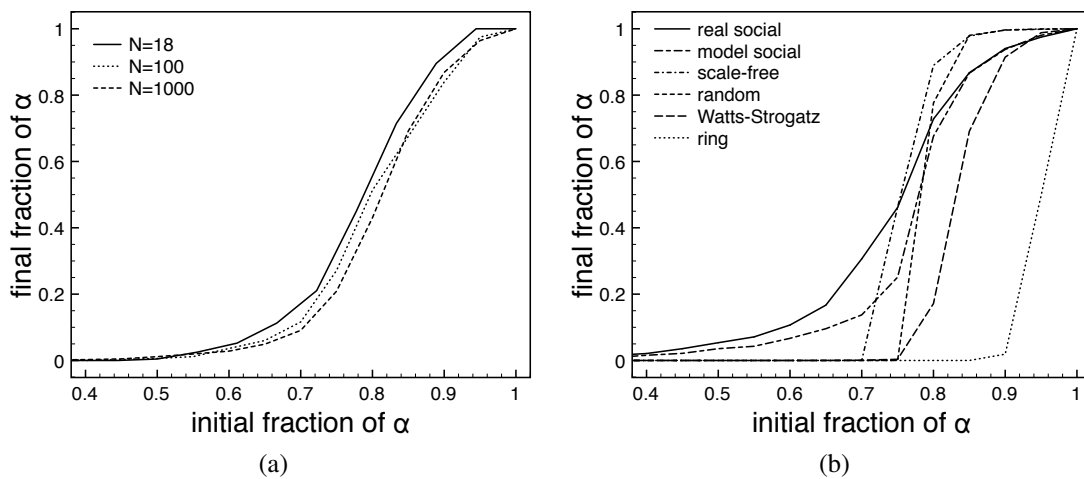


Figure 2: Final ratio of α -players as a function of their initial ratio. (a): Watts–Strogatz small-world networks scaling; $\bar{k} = 4$. (b): Other network topologies with size $N = 1000$ and $\bar{k} = 6$. The horizontal scale starts at $\alpha = 0.4$. Results are averages over 50 runs for each network class using myopic best response as a strategy update rule.

Summarizing Cassar’s experimental settings, groups of 18 subjects were used, virtually connected with a local network according to three types of graph topology: ring, random, and Watts–Strogatz small world (Watts and Strogatz, 1998). Watts–Strogatz graphs are constructed starting from a regular lattice of low degree and rewiring each link in turn with some small probability to a node chosen uniformly at random. Thanks to the formation of shortcuts between distant parts of the ring the clustering coefficient remains high, while the path lengths are dramatically shortened. Although the resulting graphs are poor representations of actual social networks, some statistical quantities are qualitatively correctly reproduced (Watts and Strogatz, 1998; Newman, 2003).

The degree of each node k was exactly 4 for the ring, while it was $\bar{k} = 4$ on the average for random and small-world networks. Of course, a single realization of the ring was used, while three different realizations of each of the other two topologies were generated.

Cassar’s results can be summarized as follows. In all three networks the Pareto-dominant equilibrium was the preferred result, with a significant advantage for the small-world networks in terms of coordination on the efficient outcome. Likewise, the ring was more favorable than the random graph. The frequency of choice of the Pareto-dominant outcome on the small-world graphs is unusually high, about 95%. Thus, the qualitative conclusion is that rings, and especially small-world networks, are favorable topological structures for coordination on the socially efficient outcome. This is in contrast with theoretical results on rings using noisy best response dynamics (Ellison, 1993) while there are no theoretical results on Watts–Strogatz small worlds to compare with. However, from the extensive numerical work

of (Roca et al., 2009) it appears that several different graph structures do favor the payoff-dominant equilibrium in the population for the Stag Hunt for most of the strategy update rules tried, but not for best reply dynamics. In light of the above, Cassar’s results seem to us less compelling than they would appear at first sight.

To obtain more insight into the matter, we decided to simulate the game behavior on an ensemble of computer-generated networks of the same size $N = 18$ as those used by Cassar with best response dynamics. We are aware of the limitations of the comparison: artificial agents are not the same thing as rational or semi-rational humans in the laboratory and time scales are vastly different since only a limited number of runs can be effectively tested in experiments. Nevertheless, we think that the exercise is worthwhile and can shed some light into the question. One important thing to note is that in (Cassar, 2007) the first period move in most cases is the payoff-dominant one, which might be due to psychological reasons in human subjects and is frequently observed in experiments (see also (Battalio et al., 2001; My et al., 1999)). In order to explore the whole spectrum thus avoiding such initial bias in the simulations, we have studied several different initial proportions between 0 and 1. Small-world instances were generated anew for each run and each computed point is the average of 50 runs. We have used a fully asynchronous update scheme in which a randomly selected agent is chosen for update with replacement at each discrete time step. To detect steady states of the dynamics we first let the system evolve for a transient period of $5000 \times N \simeq 5 \times 10^6$ time steps. After a quasi-equilibrium state is reached past the transient, averages are calculated during $500 \times N$ additional time steps. A steady state has always been reached in all simulations performed within the

prescribed amount of time, for most of them well before the limit.

As an update rule we used both myopic best reply as well as best reply with a small amount of mutation $q = 0.02$. Figure 1 reports the average results of 50 runs for each case. As prescribed by theory (Kandori et al., 1993; Ellison, 1993) and confirmed by simulations, the noisy dynamics leads essentially to risk-dominant outcomes. On the other hand, with deterministic best response dynamics, the results are that, in general, the system reaches at steady state the basin corresponding to its initial strategy proportion, with a slight advantage for the risk-dominant equilibrium, also in qualitative agreement with the expected theoretical results. Focusing more specifically on the average initial conditions that arose in Cassar's experiment, i.e. with a proportion of α of about 0.7, one sees that at this point the amount of cooperation found in the simulations is much lower, about 0.30 instead of full or almost full cooperation found in the few laboratory experiments shown as small squares. Again, note that the above results are for automata playing mechanically a deterministic myopic best response. Instead, Cassar's results have been obtained with human players; nevertheless, the difference is striking.

Cassar tried to relate her results to some statistical topological features of the networks. Her main suggestion was that the higher the clustering coefficient¹, the higher the probability of players choosing the Pareto-superior strategy. Unfortunately, given the small size $N = 18$ of such networks and only three network realizations each for random and Watts–Strogatz, all sampled quantities such as the degree distribution function $p(k)$ and mean clustering coefficient \bar{C} are too noisy to be statistically significant. For example, for a random graph, the clustering coefficient C asymptotically tends to 0 as $N \rightarrow \infty$. However, for small N clustering remains high in random graphs, which is actually the case for the values reported by Cassar. Thus, it is difficult to relate \bar{C} with the game dynamics for such small networks. With those caveats in mind, in order to get an idea as to the effect on the dynamics of scaling-up the network, we report in Fig. 2 (a) the results on graphs of size $N = 100$ and $N = 1000$, together with those for $N = 18$. It is apparent that, apart from smoothing the finite-size fluctuations, scaling-up the graph has only the effect of shifting the inset of cooperation on the payoff-dominant outcome a bit further to the right. In Fig. 2 (b) we report the fraction of population coordinating on the payoff-dominant strategy α as a function of the initial proportion of α -strategists in various network types of size $N = 1000$ for the payoff values used in Cassar

¹The clustering coefficient C_i of a node i is defined as $C_i = 2E_i/k_i(k_i - 1)$, where E_i is the number of edges in the neighborhood of i . Thus C_i characterizes the extent to which nodes adjacent to node i are connected to each other. The clustering coefficient of the graph is simply the average over all nodes: $C = \frac{1}{N} \sum_{i=1}^N C_i$ (Newman, 2003).

(2007). It can be seen that the clustering coefficient does not seem to play an important role on the population behavior. In fact, rings and Watts–Strogatz small-world graphs which both have high clustering values lead to the lowest amount of payoff dominance. On the other hand, both model and real social networks, which also have high clustering, show more coordination on the payoff-dominant strategy for α below the theoretical 0.66 value, as well as slightly diminished value in the region above this value. The explanation for this behavior is related to the community structure that these networks possess (Tomassini and Pstelacci, 2010). In fact, very often at steady state the population is polymorphic, with a minority of clusters in which α dominates below 0.66 and a minority of clusters of agents playing β above this limit. Table 2 illustrates the above by giving the mean clustering coefficient \bar{C} and the modularity Q^2 of the irregular network types for $N = 1000$. The modularity values have been computed with Newman's and Girvan's divisive algorithm based on betweenness Newman and Girvan (2004).

In conclusion, these numerical experiments confirm that the key factor to promote cooperation in networks of agents playing coordination games according to best response when risk-dominance should theoretically prevail, is the network community structure, not the clustering coefficient. Conversely, this same community structure makes it possible for a fraction of β -strategists to survive in clusters when payoff-dominance should prevail.




Figure 3: Final average fraction of α -players as a function of their initial fraction in the population in small-world networks of size $N = 18$, $N = 100$, and $N = 1000$ with $\bar{k} = 4$. Agent strategy update rule is by imitation of the best.

²According to Newman (Newman, 2006), where quantitative definitions are given, modularity is proportional to the number of edges falling within clusters minus the expected number in an equivalent network with edges placed at random.

	Ring	Small-World	Random	Scale-Free	Model Social	Real Social
\bar{C}	0.6	0.44	0.006	0.03	0.57	0.69
Q	-	-	0.31	0.30	0.66	0.69

Table 2: Average clustering coefficient \bar{C} and modularity Q for various network types of size $N = 1000$ and $\bar{k} = 6$. The values are averages over 20 independent graph realizations. A - sign means that Q is not meaningful.

Since multi-agent simulations are cheap, while laboratory experiments demand a lot of time and resources, we have also simulated the same system assuming that the agents play unconditional imitation of the best in their neighborhood, instead of playing best response. Imitation of the best is a primitive strategy for humans, but it could be used in the absence of more refined reasoning tools, as in the experiment of (My et al., 1999). After all, such imitative behavior is very common in the stock market. The results for different initial shares of α and for three network sizes are shown in Fig. 3, and should be compared with Figs. 1 and 2 (a). The notable feature is that the fraction of population playing α is strongly enhanced with respect to the simulations using best reply as a strategy update rule. This is in agreement with the numerical findings of (Roca et al., 2009) where it is shown that unconditional imitation of the best gives rise to the highest amount of efficient coordination on all network types tested. Indeed, Cassar’s experimental observations would be much closer to the results using imitation of the best than to those updating with best reply, as can be seen by comparing Figs. 1 and 3. However, in Cassar’s experiment, neighbors’ payoffs were not made known to the players and thus they could not employ a decision rule based on payoff differences. Indeed, Cassar’s analysis of the subjects’ behaviors favored rules based on myopic best reply and inertia, which means that after having chosen a strategy, a player may keep it for some time.

Clearly, a delicate point is the actual decision rule, or rules, humans do use during these experiments. While the simulated protocol revision rules used in simulations are extremely simple and homogeneous in the agents population, this is probably not the case with human players. Certainly, some amount of more sophisticated learning is at work which is not fully represented in the basic rules, as explained in Camerer’s book (Camerer, 2003), for example. For this reason, we think that it is extremely useful to validate statistical learning models arising from the experiments. These could then in turn guide and pave the way for better and more realistic strategy revision rules.

Summary and Conclusions

In this work we have studied general coordination games on complex networks by numerical simulation and we have compared the results with those of the few experimental studies that have been performed on structured populations.

For general coordination games of the Stag Hunt type there is a tension between payoff-dominance and risk-dominance and thus it is of interest to know whether there exist population topology conditions that might favor the socially efficient, Pareto-superior outcome. We have simulated a particular, yet representative, coordination game on several classes of complex networks in order to compare the results with the laboratory experiments of Cassar (Cassar, 2007). This experiment with human beings is, to our knowledge, the only one to date which employs complex network structures resembling, at least from some statistical point of view, real social networks. Our results suggest that Cassar’s claims on the role of Watts–Strogatz small-world networks, and especially their clustering coefficient, on the predominance of payoff-dominant outcomes are inconclusive and are essentially due to favorable average initial conditions. These, in turn, seem to be a bias that is almost always present in such experiments and which may well be due to human psychological propensities, something that cannot be reproduced by the artificial agents used in the simulations but which can be easily simulated by generating the corresponding initial conditions. The numerical work also show that an important source of promotion of the efficient outcome is due to the community structure present in some networks for reasonable-sized networks, i.e. with a size of at least one hundred nodes. This, however, cannot be directly related to the experimental studies as the size of the populations used in the latter have been too small till now for any meaningful partition into clusters. In conclusion, we suggest that further laboratory work on a larger scale, such as those reported in Grujic et al. (2010) should be performed to elucidate the role that complex networks of contacts may have on the emergence of efficient coordination patterns when human agents are considered. In conclusion, we think that, although numerical multi-agent simulations cannot be directly compared with heterogeneous and possibly complex human decision rules, they are a useful guide for planning and interpreting laboratory experiments and social dynamics in general.

Acknowledgements

A. Antonioni and M. Tomassini are grateful to the Swiss National Science Foundation for financial support under contract number 200021-132802/1.

References

- Battalio, R., Samuelson, L., and Huyck, J. V. (2001). Optimization incentive and coordination failure in laboratory stag hunt games. *Econometrica*, 61:989–1018.
- Camerer, C. F. (2003). *Behavioral Game Theory*. Princeton University Press, Princeton, NJ.
- Cassar, A. (2007). Coordination and cooperation in local, random and small world networks: Experimental evidence. *Games and Economic Behavior*, 58:209–230.
- Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *Quarterly Journal of Economics*, 107:739–771.
- Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61:1047–1071.
- Grujic, J., Fosco, C., Araujo, L., Cuesta, J. A., and Sánchez, A. (2010). Social experiments in the mesoscale: humans playing a spatial Prisoner's Dilemma. *PLoS ONE*, 5(11):e13749.
- Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Huyck, J. B. V., Battalio, R. C., and Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *Amer. Econ. Rev.*, 80:234–249.
- Huyck, J. B. V., Battalio, R. C., and Beil, R. O. (1993). Asset markets as an equilibrium selection mechanism: Coordination failure, game form auctions, and tacit communication. *Games Econ. Behav.*, 5:485–504.
- Kandori, M., Mailath, G., and Rob, R. (1993). Learning, mutation, and long-run equilibria in games. *Econometrica*, 61:29–56.
- Keser, C., K-M-Erhart, and Berninghaus, S. (1998). Coordination and local interaction: experimental evidence. *Economics Letters*, 59:269–275.
- Luthi, L., Pestelacci, E., and Tomassini, M. (2008). Cooperation and community structure in social networks. *Physica A*, 387:955–966.
- Morris, S. (2000). Contagion. *Review of Economic Studies*, 67:57–78.
- My, K. B., Willinger, M., and Ziegelmeyer, A. (1999). Global versus local interaction in coordination games: an experimental investigation. Technical Report 9923, Working papers of BETA. ULP, Strasbourg.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45:167–256.
- Newman, M. E. J. (2006). Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA*, 103:8577–8582.
- Newman, M. E. J. and Girvan, M. (2004). Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113.
- Robson, A. J. and Vega-Redondo, F. (1996). Efficient equilibrium selection in evolutionary games with random matching. *J. Econ. Theory*, 70:65–92.
- Roca, C. P., Cuesta, J. A., and Sánchez, A. (2009). Evolutionary game theory: temporal and spatial effects beyond replicator dynamics. *Physics of Life Reviews*, 6:208–249.
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. (2006). Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl. Acad. Sci. USA*, 103:3490–3494.
- Shoham, Y., Powers, R., and Granager, T. (2007). If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171:365–377.
- Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Cambridge, UK.
- Szabó, G. and Fáth, G. (2007). Evolutionary games on graphs. *Physics Reports*, 446:97–216.
- Tomassini, M. and Pestelacci, E. (2010). Evolution of coordination in social networks: A numerical study. *Int. J. Mod. Phys. C*, 21(10):1277–1296.
- Vega-Redondo, F. (2003). *Economics and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393:440–442.
- Weidenholzer, S. (2010). Coordination games and local interactions: a survey of the game-theoretic literature. *Games*, 1:551–585.
- Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.