Chapter 1

# Insights into the Impact of Social Networks on Evolutionary Games

**Anton Chechetka**
**Katia Sycara**
**Paul Scerri**
Carnegie Mellon University
{antonc, katia, pscerri}@cs.cmu.edu

In this chapter, we explore the use of evolutionary game theory (EGT) [14, 10, 9] to model the dynamics of adaptive opponent strategies for large population of players. In particular, we explore effects of information propagation through social networks in Evolutionary Games. The key underlying phenomenon that the information diffusion aims to capture is that reasoning about the experiences of acquaintances can dramatically impact the dynamics of a society. We present experimental results from agent-based simulations that show the impact of diffusion through social networks on the player strategies of an evolutionary game and the sensitivity of the dynamics to features of the social network.

## 1 Introduction

We use evolutionary game theory (EGT) [14, 7, 4] to model the dynamics of adaptive opponent strategies for large population of players. Previous EGT work has produced interesting, and sometimes counter-intuitive results about how populations of self-interested agents will evolve over time [6, 5].

In our model, at each stage of the game, boundedly rational players observe the strategies and payoffs of a subset of others and use this information to choose their strategies for the next stage of the interaction. Building on EGT, we introduce a model of interaction where, unlike the standard EGT setting, the basic stage game changes over time depending on the global state of the population (state here means the strategies chosen by the players). More precisely, each player has three strategies available (cooperate $C$, defect $D$ and do-nothing $N$) and the payoffs of the basic stage game are resampled when the proportion of the players playing $D$ crosses a certain threshold from above. This feature requires long-term reasoning by the players that is not needed in the standard EGT setting. A possible example of a similar real-world situation is a power struggle between different groups. When cooperation drops sufficiently and there are many defections, the situation turns to chaos. When order is restored, i.e., when cooperation resumes, the power structure and thus, the payoffs, will likely be different than before the chaos. The payoffs are kept constant while most of the players $C$ooperate (support the status quo) or do-$N$othing, but when enough players are unhappy and choose to $D$efect, the power balance breaks and radically different one may emerge afterwards.

The available strategies were chosen to abstractly capture and model violent uprisings in a society. Players playing $C$ cooperate with the current regime and receive reward when interacting with others playing $C$. If a player has a good position in a regime, it has large incentive to continue playing $C$. $D$ is a strategy played to change the payoffs over a long term, but at an unavoidable immediate cost. Intuitively, it resembles resorting to insurgency or other violent tactics to overthrow a regime. When many players play $D$, playing $C$ can lead to very low payoffs. For example, one can imagine a person trying to run a small business during a violent uprising. If these costs are too high, but the player has no incentive to change the regime, playing $N$ can limit payoffs – both negative and positive, until the situation stabilizes. Intuitively, this might correspond to going into hiding or temporarily leaving the conflicted area.

Similar to [9, 8], we investigate the spatial aspect of the interaction. Previous work has shown that spatial interaction can change which strategies are most effective, e.g., in [3] an interation lattice changed which strategies were most effective in an iterative prisoners dilemma game. In our model, the players are connected into a *social network*, through which the rewards are propagated[13, 11]. Thus the players can benefit (or suffer) indirectly depending on how well off their friends in the network are. We show empirically that the connectivity pattern of the network, as well as the amount of information available to the players, have significant influence on the outcome of the interaction. In particular, the presence of a dense scale-free network or small-worlds network led to far higher proportions of players playing $C$ than other social network types.

## 2 Game Details

We consider a finite population $X$ of players. At each stage all the players are randomly matched in triples to play the basic stage game. Each player thus participates in every stage. Each player has three strategies available:

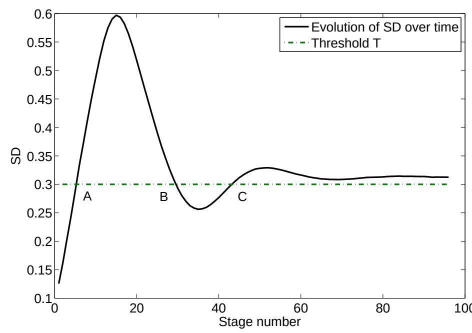


**Figure 1**: An example trace of an individual run of the system. $x$-axis is the stage number ("'time step"), $y$-axis is the proportion $SD$ of the population playing $D$. The level of threshold $T$ is also plotted for a reference.

cooperate ($C$), defect ($D$) and do-nothing ($N$) (one can interpret these choices as participating in democratic process, resorting to insurgency and minimizing interactions with the outer world correspondingly). The payoff $p_i(k)$ of the stage $k$ game to player $x_i$ is

| | | 0 opponents play $D$ | 1 or 2 opponents play $D$ |
|---|---|---|---|
| $x_i$'s strategy | $C$ | $cc_i - \#_i(N)$[1] | $cd$[2] |
| | $D$ | $dc$ | $dd$ |
| | $N$ | $n$ | |

where $cc_i - 2 > n > dc > dd > cd$. Note that the payoff matrices for different players can only differ in the value of $cc_i$. All the other payoffs are constant across the population.

Denote $SD(k)$ the proportion of the population that defected during stage $k$:

$$SD(k) = \frac{\text{number of players that played } D \text{ during stage } k}{|X|},$$

Before the start of the first stage, $cc_i$ are sampled uniformly from an interval $[CC_{min}, CC_{max}]$. If during stage $k^*$ the series $SD(k)$ crosses a fixed threshold[3] $T \in (0,1)$ from above, i.e.

$$SD(k^* - 1) > T \text{ and } SD(k^*) < T,$$

then all $cc_i$ are resampled. Otherwise they stay the same as for previous stage. For example, in an individual run plotted in Fig. 1 the values of $cc_i$ would be resampled only at point $B$.

[1] $\#_i(N)$ means the number of $i$'s opponents playing $N$.

[2] Here is a simple rule for distinguishing between these 4 variables: the first letter corresponds to $x_i$'s strategy, the second letter is $c$ if both of the $x_i$'s opponents play $C$ and $d$ otherwise. For example, $cd$ is the payoff of playing $C$ given that at least one of the opponents plays $D$.

[3] See the end of this section for the interpretation of this threshold.

One can interpret the above interaction as a power struggle: if the proportion of players supporting status quo (i.e. cooperating or doing nothing) is high enough, the payoffs for each individual players do not change. When enough players defect, the system "falls into chaos" and after it emerges back from this state, a new power balance is formed and the payoffs change correspondingly. Threshold $T$ in this interpretation is the minimum number of defectors that brings the system into chaos.

## 2.1 Impact of social networks

A social network for finite population $X$ is an undirected graph $< X, E >$. Two players $x_i$ and $x_j$ are neighbors in the network if and only if $(x_i, x_j) \in E$. We investigate the effect of reward sharing in social networks. After each stage $k$ every player $x_i$ obtains in addition to its own payoff $p_i$ a shared payoff $ps_i$:

$$ps_i(k) = \alpha \sum_{x_j \in \text{neighbors}(x_i)} p_j(k),$$

where $\alpha \in [0, 1]$ is a parameter of the system.

Notice that this does not incur payoff redistribution: the shared payoff is not subtracted from payoffs of the players that cause it. One can interpret this phenomenon as players being more happy when their friends are happy.

#### Social network type

The *small-world property* of the network means that the average distance between two nodes in the network is small. It has been shown [12] that regular non-small-world networks, such as grids, may be transformed to small-world ones by changing only a small fractions of edges. We followed the algorithm from [12] to generate the networks with probability 0.1 of rewiring any edge of the regular structure.

In scale-free networks [2] the number of neighbors of a vertex is distributed according to a scale-free power law, therefore few highly-connected vertices dominate the connectivity. Many real-world networks possess the small-worlds and/or scale-free properties [2, 12].

The impact of both small-worlds and scale-free networks are explored below.

# 3 Player reasoning

## 3.1 Information available to players

Before describing the player reasoning algorithm one has to define what information is available to the player, i.e. define an observation model. We assume that the players are aware of the overall behavior of the game, but may not be aware of the true values of parameters, such as the proportion $SD(k)$ of the population that played $D$ at stage $k$. The players only observe the actions of their opponents for the given stage, as opposed to observing the whole population. Therefore, the observations available to $x_i$ after stage $k$ are its payoff $p_i(k)$, shared payoff $ps_i(k)$, and proportion $SC_i^{obs}(k), SD_i^{obs}(k), SN_i^{obs}(k) \in \{0, 0.5, 1\}$ of its direct opponents playing $C$, $D$ and $N$ during the $k^{th}$ stage.

Note that the information about the global properties of social network connectivity, such as density or whether the network is small-worlds or scale free, is not available to players. Therefore, this global information is not used in the reasoning algorithm.

## 3.2 The reasoning algorithm

It is easy to see that for any triple of players, a single-stage game has 2 Nash equilibria in pure strategies: everybody cooperating and everybody defecting. The cooperative equilibrium Pareto-dominates the "all-defect" equilibrium. Therefore, if the "all-cooperate" payoffs $cc_i$ were always held constant across the stages, one would expect a population of rational players to always play $C$. However, the payoffs are resampled once the proportion of players playing $C$ drops below $T$ and then grows above $T$ again. This provides an incentive, for the players which happened to receive relatively low values of $cc_i$, to play $D$ for some period of time in order to try and cause the resampling of payoffs. On the other hand, if a significant share of the players play $D$, some of the players may decide to play $N$, which guarantees a fixed payoff and provides an opportunity to "wait until the violence ends".

A natural way for a player to choose a strategy for the next stage is to compare the (approximate) cumulative future expected payoffs resulting from different strategies. Denote $EP_i(X)$ the approximate cumulative future expected payoff for player $i$ and strategy $X$. Let $SX_i(k)$ be $i$'s estimate of the share of population playing $X$ on time step $k$. Then the action selection for step $k+1$ is as follows. If $SD_i(k) > T$, player $i$ chooses action $\arg\max_{X=C,N} EP_i(X)$. Otherwise it chooses $\arg\max_{X=C,D,N} EP_i(X)$. The reason for treating situation $SD_i(k) > T$ specially is that once the share of defectors reaches the threshold, reducing the share of players below $T$ is in common interest of all the players, and the approximate computations of expected utilities do not always capture this feature.

The previous paragraph assumed $EP_i(X)$ to be known. We now turn to their approximate computation.

First consider $EP_i(D)$. The only incentive for a player $i$ to play $D$ is to try to bring the system into chaos in hope that, when the system emerges from chaos, the resampled all-cooperate payoff $cc_i$ for that player will be higher then it is now. Denote $TTC_i$ the $i$'s estimate of the number of stages that it will need to play $D$ before the share of those playing $D$ is higher than $T$, $TC_i$ - estimate of the number of stages that the system will spend above the threshold and finally,

$TS_i$ the length of the following "stability period". Then

$$\begin{aligned} EP_i(D) &\approx (TTC_i + TC_i)E\left[p_i(D)\right] + TS_iE[cc_i^{new}] \\ &= TTR_iE[p_i(D)] + TS_i\frac{CC_{min} + CC_{max}}{2}, \end{aligned} \tag{1}$$

where $TTR_i \equiv TTC_i + TC_i$ is "time to resampling" and

$$E[p_i(D)] = P(\#_i(D) = 0)dc + P(\#_i(D) > 0)dd.$$

Expected payoff for action $C$ over the time period is approximated as

$$\begin{aligned} EP(C) \approx\; & TS_i(p_i(C) + ps_i) + TTC_iE[p_i(c)] \\ & + TC_i(P(\#_i(D) > 0)cd + P(\#_i(D) = 0)(p_i(C) + ps_i)), \end{aligned} \tag{2}$$

where $P(\#_i(D) > 0) = 1 - (1 - T)^2$ and

$$\begin{aligned} E[p_i(C)] =\; & P(\#_i(C) = 2)cc_i + P(\#_i(C) = 1, \#_i(N) = 1)(cc_i - 1) \\ & + P(\#_i(N) = 2)(cc_i - 2) + P(\#_i(D) > 0)cd \end{aligned}$$

(note that the probabilities here sum to one).

Finally, expected payoff for $N$ over the same time interval is

$$EP(H) = (TTC_i + TC_i + TS_i)n.$$

One can see that a player only expects to get the shared payoff in case of all-cooperative outcomes.

In our model, time of stability $TS_i$ and time in chaos $TC_i$ are system constants that do not differ across the population.

The belief $SX_i(k)$ about the proportion of players playing $X$ at stage $k$ is maintained by each player individually. After each stage each player learns about the strategies of its opponents for that stage. $SX_i$ is then updated according to

$$SX_i(k+1) = \gamma SX_i^{obs}(k+1) + (1-\gamma)SX_i(k) \tag{3}$$

where $\gamma \in (0, 1]$ is learning rate. Each player also maintains $\delta SX_i(k)$, an estimate of

$$\delta SX(k) \equiv SX(k) - SX(k-1),$$

using an expression analogous to Eq. 3 to update it. In the expressions (1-2) $P(\#_i(X))$ are approximated straightforwardly using $SX_i$, for example

$$P(\#_i(C) = 2) \approx SC_i^2(k)$$

Having $SX_i$ and $\delta SX_i$, each player can estimate $TTC_i$ using a linear approximation. For $SD_i < T$, we have ($TTC$ is a system-wide constant)

$$TTC_i = \begin{cases} TTC, & \delta SD_i \leq 0 \\ \frac{T - SD_i}{\delta SD_i}, & \delta SD_i > 0 \end{cases}$$

For $SD_i \geq T$, $TTC_i = 0$.

## 4 Experimental results

In our experiments the population size was fixed to 1000 players. The numerical values of payoff constants were

$$dc = -1, \quad dd = -3, \quad cd = -5, \quad CC_{min} = 3, \quad CC_{max} = 10$$

Estimated time of stability was fixed to $TS_i = TS = 50$ stages, "chaos threshold" $T = 0.3$. Initial player-specific values were $SC_i(0) = 1, \quad \delta SC_i(0) = -0.02$. For each set of specific parameter values the results were averaged over 500 runs. Unless otherwise noted, the players were connected via a scale-free network with average density of 8.

We were primarily interest in how different parameters of the model affect the evolution of proportion of players playing $C$ over time. On all graphs $x$-axis denotes the stage of the interaction, $y$-axis denotes $SC$, $SD$ and $SN$. In previous work[1], we presented results for the case where action $N$ was not available to the players. In each of the figures below we contrast the results when $N$ is and is not available to the players.

Note that because the plotted results are averages over 100 runs. Averages provide more meaningful information about the influence of the parameters values on the system, than do individual runs which can vary distinctly from run to run. Most parameter values allow the $SC$ to fall below $T$ on some occasions, but what varies is how often this occurs, how rapidly changes happen and how quickly cooperation resumes. These effects are more clearly seen on graphs of averages than many individual runs super-imposed on a single graph. Notice that the fact that the value of $SD$ on the plots rarely rises above $T$ does not mean that payoffs are almost never resampled - individual runs have much more variance and resampling happens quite often. It simply means that on average $SD$ is below $T$.

Figure 2 shows the baseline configuration, with 2(a) showing the case where $N$ is available and 2(b) showing the case where it is not. In both cases, early in the game many players choose $D$ to either try to change the payoffs or protect against losses. When $N$ is available to the players, many choose this action in response to others playing $D$. Eventually this discourages the use of $D$ an equilibrium settles in. While the initial dynamics in both cases are similar, notice that over time the proportion of $C$ is far higher in the case where $N$ is available than when it is not. This may indicate that if players are able to avoid spasms of violence without getting hurt, the outcome for all will be better.

Figure 3 shows the impact of setting the network density to 2, 4, 8 and 16. In general, the higher the average network degree, the more players played $C$ and the more quickly players stopped playing $D$. For the less dense networks, players often chose $D$ early on, but in the most dense network, the lure of shared rewards was too high for players to have incentive to try to move the system towards chaos. In the less dense networks, the availability of the $N$ action allowed the system to move toward all playing $C$, but as in the baseline case, without the $N$ action, some level of $SD$ persisted. When the average network density was 4,

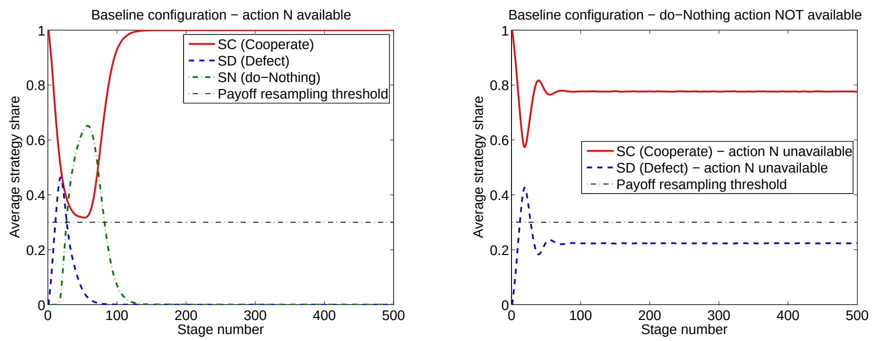


(a) (b)

**Figure 2**: Baseline configuration (scale-free network with density 8) with available action $N$ (a) and with $N$ not available (b).

the system moved back towards $SC = 1$ faster than when the network density was 2. This result may indicate that dense social networks are critical to stable societies.

Figure 4 shows what happens when there is no sharing across the social network. The sharp early peak in $SD$ is similar to the sparse network shown above. This is one of the few cases where the availability of the $N$ action leads to a lower $SC$ over the course of the game. However, the option to play $N$ is extensively used and $SD$ is reduced to 0. Over an extended period of time, $SC$ does rise to 1, but $N$ dominates for a long time.

If the type of the network is set to small-worlds instead of scale-free (with the average of 4 neighbors), $SC$ stays very close to 1 regardless of the availability of $N$ to the players (there is no plot for this case, because the results are so trivial). This remarkable relative stability is likely due to the very even sharing of reward across all members of the team, reducing the possibility of a cascade towards chaos. This result may suggest that human societies that have a more scale-free nature will be more likely to descend into chaos.

Figure 5 shows the result as the learning rate is set to 0.05, 0.1, 0.4 and 0.8. Smaller learning rate means that the players are reluctant to change their estimates of the parameters; the closer the learning rate to 1, the more importance is attributed to the most recent observations.

Several interesting effects occur due to the learning rate. Firstly, an intermediate learning rate induces an oscillation in behavior with increasing and decreasing $SD$. Higher or lower learning rates induce different behavior. A high learning rate quickly settles the population down to playing $C$, because the players are better able to estimate future rewards which are maximized by a stable society. A low learning rate eventually allows a stable society but not before a large $SD$ has occurred. Interestingly, none of these effects were observed when the $N$ action was not available to the players. With learning eventual behavior (except for the intermediate learning rate) $SC$ was higher when $N$ was available.

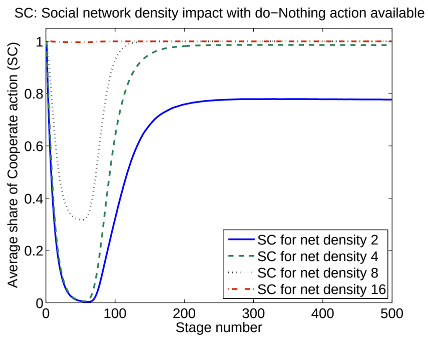


(a)

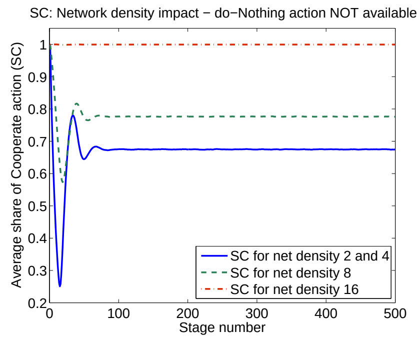


(b)

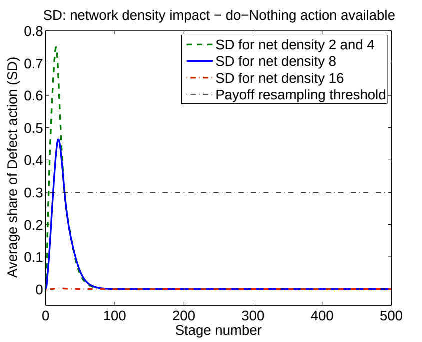


(c)

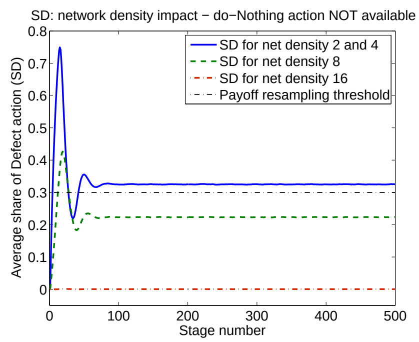


(d)

**Figure 3**: Impact of network density on the players' strategies. In the top row, the share of players playing *cooperate*, in the bottom - *defect*. On the left, the action $N$ is available to the players, on the right - not available.

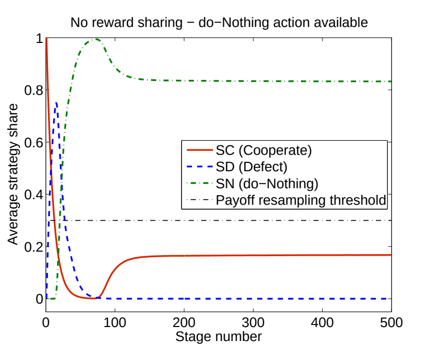


(a)

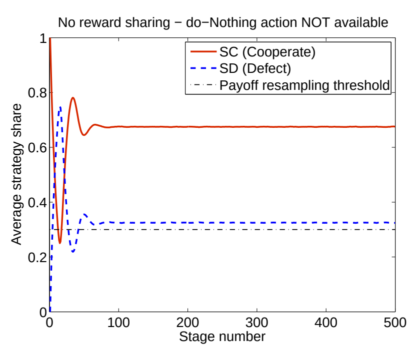


(b)

**Figure 4**: Results with reward sharing disabled with available action $N$ (1.4(a)) and with $N$ not available (1.4(b)).

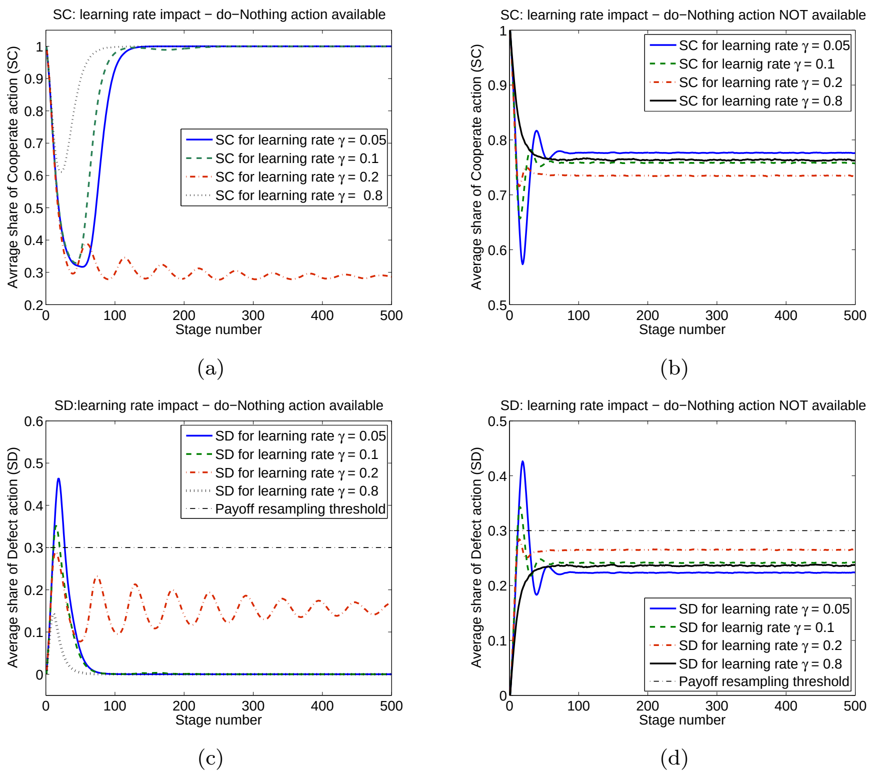


**Figure 5**: Impact of learning rate on the players' strategies. In the top row, the share of players playing *cooperate*, in the bottom - *defect*. On the left, the action $N$ is available to the players, on the right - not available.

# 5 Conclusions and future work

This paper presented an evolutionary game with players connected into a social network, sharing payoffs with their neighbors in that network. If individual players reason that increased long term payoffs might be higher if the whole society can be forced into chaos, they will accept significant short term costs and risk, to bring that situation about. The key conclusion from this game is that a society of *rational* agents who will all gain if they all play cooperative strategies can easily be induced to play strategies that are guaranteed to lead to a negative payoff.

Our experiments show that the existence and nature of a social network makes a dramatic difference to the evolution and conclusion of the game. Very dense networks or small worlds networks had far higher proportions of players playing cooperative strategies than when there is a sparse scale-free network. This result has implications for all EGT where interaction occurs between players, but only simple social networks are used. It is possible that such results will change if different interaction networks are used.

## Bibliography


[1] *Evolutionary Games and Social Networks in Adversary Reasoning* (Boston, MA, ).

[2] Barabási, Albert-László, and Réka Albert, "Emergence of scaling in random networks", *Science* **286** (1999).

[3] Brauchli, Kurt, Timothy Killingback, and Michael Doebeli, "Evolution of cooperation in spatially structured populations", *Journal of Theoretical Biology* (1999).

[4] Cabrales, Antonio, "Stochastic replicator dynamics", *International Economic Review* **41**, 2 (2000).

[5] d'Artigues, Agnès, and Thierry Vignolo, "Why global integration may lead to terrorism: An evolutionary theory of mimetic rivalry", *Economics Bulletin* **6**, 11 (2003).

[6] Frey, Bruno S., and Simon Luechinger, "Terrorism: Deterrence may backfire", *European Journal of Political Economy* **20**, 2 (2002).

[7] Hofbauer, Josef, and Karl Sigmund, "Evolutionary game dynamics", *Bulletin of the American Mathematical Society* **40**, 4 (2003).

[8] Killingback, Timothy, and Michael Doebeli, "Spatial evolutionary game theory: Hawks and doves revisited", *Proceedings of The Royal Society (Biological Sciences)* (1996).

[9] Nowak, Martin, and Robert May, "The spatial dilemmas of evolution", *International Journal of Bifurcation and Chaos* **3** (1993).

[10] Taylor, Peter, and Leo Jonker, "Evolutionary stable strategies and game dynamics", *Mathematical Biosciences* **40** (1978).

[11] Travers, J., and S. Milgram, "An experimental study of the small world problem", *Sociometry* **32** (1969), 425–443.

[12] Watts, Duncan, and Steven Strogatz, "Collective dynamics of small-world networks", *Nature* **393** (1998).

[13] Watts, Duncan J., Peter Sheridan Dodds, and M. E. J. Newman, "Identity and search in social networks", *Science* **296**, 5571 (2002), 1302–1305.

[14] Weibull, Jörgen, *Evolutionary Game Theory*, MIT Press (1995).