

Hawks and Doves in an Artificial Dynamically Structured Society

E. Pestelacci and M. Tomassini

Information Systems Institute, HEC, University of Lausanne, Switzerland
{enea.pestelacci,marco.tomassini@unil.ch}

Abstract

Using a dynamical network model of society, we show that cooperation is the norm in the Hawks-Doves game when individuals are allowed to break ties with undesirable neighbors and to make new acquaintances in their extended neighborhood. This is an interesting result, as standard theory for mixing populations prescribes that a certain fraction of defectors must always exist at equilibrium. We discuss the empirical network structure reasons that allow cooperators to thrive in the population.

Introduction and Previous Work

Hawks-Doves, also known as Chicken, is a two-person, symmetric game with the following payoff bi-matrix:

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

In this matrix, D stands for strategy “hawk”, and C stands for strategy “dove”. Metaphorically, a hawkish behavior means a strategy of fighting, while a dove, when facing a confrontation, will always yield. R is the *reward* the two players receive if they both cooperate (C), P is the *punishment* for bilateral defection (D), and T is the *temptation*, i.e. the payoff that a player receives if it defects, while the other cooperates. In this case, the cooperator gets the *sucker’s payoff* S. The game has a structure similar to that of the *Prisoner’s Dilemma* (Axelrod, 1984). However, the ordering of payoffs for the Prisoner’s Dilemma is $T > R > P > S$ rendering defection the best rational individual choice, while in the Hawks-Doves game the ordering is $T > R > S > P$ thus making mutual defection, i.e. result (D,D), the worst possible outcome. Note that in game theory, as long as the above orderings are respected, the actual numerical payoff values do not matter (Vega-Redondo, 2003).

In contrast to the Prisoner’s Dilemma which has a unique Nash equilibrium that corresponds to both players defecting, the strategy pairs (C,D) and (D,C) are both Nash equilibria of the Hawks-Doves game in pure strategies, so the game is

antagonistic, and there is a third equilibrium in mixed strategies in which strategy D is played with probability p , and strategy C with probability $1 - p$, where $0 \leq p \leq 1$ depends on the actual payoff values. We recall that a Nash equilibrium is a combination of strategies (pure or mixed) of the different players such that any unilateral deviation by any agent from this combination can only decrease her expected payoff (Vega-Redondo, 2003).

As is the case for the Prisoner’s Dilemma (Axelrod, 1984; Lindgren and Nordahl, 1994), Hawks-Doves, for all its simplicity, appears to capture some important features of social interactions. In this sense, it applies in many situations in which “parading”, “retreating”, and “escalating” are common. One striking example of a situation that has been thought to lead to a Hawks-Doves dilemma is the Cuban missile crisis in 1962 (Poundstone, 1992). Other well known applications are found in the animal kingdom (Maynard Smith, 1982).

Considering now not just two players but rather a large, mixing population of identical players where randomly chosen pairs play a sequence of two-person games, *evolutionary game theory* (Hofbauer and Sigmund, 1998) prescribes that the only *Evolutionary Stable Strategy* (ESS) of the population is the mixed strategy, giving rise, at equilibrium, to a frequency of hawks in the population equal to p , the probability with which strategy hawk, i.e. D, would be played in a mixed strategy.

In the case of the Prisoner’s Dilemma, one finds a unique ESS with all the individuals defecting. However, Nowak and May (1992) showed that cooperation in the population is sustainable in the Prisoner’s Dilemma under certain conditions, provided that the network of the interactions between players has a lattice spatial structure. Killingback and Doebeli (1996) extended the spatial approach to the Hawks-Doves game and found that a planar lattice structure with only nearest-neighbor interactions may favor cooperation, i.e. the fraction of doves in the population is often higher than what is predicted by evolutionary game theory. In a more recent work however, Hauert and Doebeli (2004) were led to a different conclusion, namely that spatial struc-

ture does not seem to favor cooperation in the Hawks-Doves game. Further studies (Tomassini et al., 2006) extended the structured population approach to other graph structures representing small worlds. Small-world networks are produced by randomly rewiring a few links in an otherwise regular lattice such as a ring or a grid (Watts and Strogatz, 1998). These “shortcuts”, as they are called, give rise to graphs that have short path lengths between any two nodes in the average as in random graphs, but in contrast to the latter, also have a great deal of local structure as conventionally measured by the *clustering coefficient*¹. These structures are much more typical of the networks that have been analyzed in technology, society, and biology than regular lattices or random graphs (Newman, 2003). In Tomassini et al. (2006) it was found that cooperation may be either enhanced or inhibited in small-world networks depending on the gain-to-cost ratio $r = R/(R - P)$, and on the strategy update rule. But Watts–Strogatz small worlds and scale-free networks, although more realistic than lattices or random graphs, are not faithful representation of typical social networks. Santos and Pacheco (2005) and Santos et al. (2006b) extended the study of the Hawk-Doves game to scale-free networks, i.e. to networks having a power-law distribution of the connectivity degree (Newman, 2003). They found that cooperation is remarkably enhanced in them with respect to previously described population structures through the existence of highly connected cooperator hubs. However, pure static scale-free networks are not found among the typical socioeconomic networks that have been studied (Amaral et al., 2000; Newman, 2001, 2003). Using real and model static social networks, Luthi et al. (2008) also found that cooperation is enhanced, although to a lesser degree, thanks to the existence of tight clusters of cooperators that reinforce each other. Static networks having a resemblance with actual social networks are a good starting point; however, the static approach ignores fluctuations and non-equilibrium phenomena. Instead, real social networks are dynamical, i.e. nodes may join the network forming new links, and old nodes may leave it as social actors come and go. Furthermore, new links between agents already in the network may also form or be dismissed. Thus, the motivation of the present work is to study the co-evolution of strategy and network structure and to investigate under which conditions cooperative behavior may emerge and be stable in the Hawks-Doves game. A related goal is to study the topological structures of the emergent networks and their relationships with the strategic choices of the agents. Some previous work has been done

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 measures the amount of “cliquishness” of the neighborhood of node i and it 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).

on evolutionary games on dynamic networks (Zimmermann and Eguíluz, 2005; Luthi et al., 2006; Santos et al., 2006a). The only one citing the Hawks-Doves game is (Santos et al., 2006a) but our model differs in several important respects and we obtain new results on the structure of the cooperating clusters.

The paper is organized as follows. In the next section we present our dynamical models. This is followed by an exhaustive numerical study of the game’s parameter space. After that we describe and discuss the statistical structure of the emerging networks and finally we give our conclusions.

Model and Dynamics

Our model is strictly local. No player uses information other than the strength of the links with its neighbors and the knowledge of her own payoff and, indirectly, the payoffs of her immediate neighbors. Moreover, as the model is an evolutionary one, no rationality, in the sense of game theory, is needed (Vega-Redondo, 2003). Players just adapt their behavior such that they imitate more successful strategies in their environment with higher probability. Furthermore, they are able to locally assess the worth of an interaction and possibly dismiss a relationship that does not pay off enough. The model and its dynamics are described in detail in the following sections.

Network and Interaction Structure. The network of agents is 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 such that there is an edge $\{ij\} \in E$. 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 will be called \bar{k} . Although there is formally a single undirected link between a player i and another player $j \in V_i$, we shall maintain two links: one going from i to j and another one in the reverse direction. Each link has a weight or “force” f_{ij} (respectively f_{ji}). This weight, say f_{ij} , represents in an indirect way the “trust” player i attributes to player j . This weight may take any value in $[0, 1]$ and its variation is dictated by the payoff earned by i in each encounter with j , as explained below.

The idea behind the introduction of the forces f_{ij} is loosely inspired by the potentiation/depotentiation of connections between neurons in neural networks, an effect known as the *Hebb rule* (Hebb, 1949). In our context, it can be seen as a kind of “memory” of previous encounters. However, it must be distinguished from the memory used in iterated games, in which players “remember” a certain number of previous moves and can thus conform their future strategy on the analysis of those past encounters (Vega-Redondo, 2003). Our interactions are strictly one-shot, i.e. players “forget” the results of previous rounds and cannot recognize

previous partners and their possible playing patterns. However, a certain amount of past history is implicitly contained in the numbers f_{ij} and this information may be used by an agent when it will come to decide whether or not an interaction should be dismissed (see below).

We also define a quantity s_i called *satisfaction* of an agent i which is the sum of all the weights of the links between i and its neighbors V_i divided by the total number of links of that node k_i :

$$s_i = \frac{\sum_{j \in V_i} f_{ij}}{k_i}.$$

We clearly have $0 \leq s_i \leq 1$.

Initialization. The constant size of the networks during the simulations is $N = 1000$. The initial graph is generated randomly with a mean degree $\bar{k} = 10$ which is of the order of those actually found in many social networks; see, for instance, (Newman, 2003). Players are distributed uniformly at random over the graph vertices with 50% cooperators. Forces between any pair of neighboring players are initialized at 0.5.

We use a parameter q which is a real number in $[0, 1]$ and it represents the frequency with which an agent wishes to dismiss a link with one of its neighbors. The higher q , the faster the link reorganization in the network. This parameter has a role analogous to the “time scale” parameter of (Santos et al., 2006a) and it controls the speed at which topological changes occur in the network. All the agents have the same value of q . It is an important consideration, as social networks may structurally evolve at widely different speeds, depending on the kind of interaction between agents. For example, e-mail networks change their structure at a faster pace than, say, scientific collaboration networks.

Update Timing. Usually, agents systems such as the present one, are updated synchronously (Nowak and May, 1992; Santos and Pacheco, 2005; Zimmermann and Eguíluz, 2005). However, strictly speaking, simultaneous update is physically unfeasible as it would require a global clock, while real extended systems in biology and society in general have to take into account finite signal propagation speed. Simultaneity may cause some artificial effects in the dynamics which are not observed in real systems (Huberman and Glance, 1993; Luthi et al., 2006). On the other hand, updating a randomly chosen agent at a time also seems a rather arbitrary extreme case that is not likely to represent reality very accurately. We have thus chosen to update our population in a partially synchronous manner. In practice, we define a fraction $f = n/N$ (with $N = an, a \in \mathbb{N}$) and, at each simulated discrete time step, we update only $n \leq N$ agents randomly chosen with replacement. This is called a *microstep*. After N/n microsteps a whole population update, i.e. a *macrostep* will have taken place. With $n = N$

we recover the fully synchronous update, while $n = 1$ gives the extreme case of the fully asynchronous update. In this work we use $f = 0.01$.

Strategy and Link Dynamics

Here we describe in detail how individual strategies, links, and link weights are updated. Once a given node i is chosen to be activated, i.e. it belongs to the fraction f of nodes that are to be updated in a given microstep, i goes through the following steps:

- if the degree of agent i , $k_i = 0$ then player i is an isolated node. In this case a link with strength 0.5 is created from i to a player j chosen uniformly at random among the other $N - 1$ players in the network.
- otherwise,
 - either agent i updates its strategy according to a local *replicator dynamics* rule with probability $1 - q$ or, with probability q , agent i may delete a link with a given neighbor j and creates a new 0.5 force link with another node k ;
 - the forces between i and its neighbors V_i are updated

Let us now describe each step in more detail.

Strategy Evolution. We use a local version of replicator dynamics (RD) as described in (Luthi et al., 2008). The local dynamics of a player i only depends on its own strategy and on the strategies of the k_i players in its neighborhood V_i . Let us call π_{ij} the payoff player i receives when interacting with neighbor j . This payoff is defined as

$$\pi_{ij} = \sigma_i(t) M \sigma_j^T(t),$$

where M is the payoff matrix of the game and $\sigma_i(t)$ and $\sigma_j(t)$ are the strategies played by i and j at time t . The quantity

$$\hat{\Pi}_i(t) = \sum_{j \in V_i} \pi_{ij}(t)$$

is the rescaled accumulated payoff (Luthi et al., 2008) collected by player i at time step t . The rule according to which agents update their strategies is the conventional RD in which strategies that do better than the average increase their share in the population, while those that fare worse than average decrease. To update the strategy of player i , another player j is drawn at random from the neighborhood V_i . It is assumed that the probability of switching strategy is a function ϕ of the payoff difference; ϕ is required to be monotonic increasing; here it has been taken linear (Hofbauer and Sigmund, 1998). Strategy σ_i is replaced by σ_j with probability

$$p_i = \phi(\hat{\Pi}_j - \hat{\Pi}_i). \quad (1)$$

The major differences with standard RD is that two-person encounters between players are only possible among neighbors, instead of being drawn from the whole population, and the latter is finite in our case. Other commonly used strategy update rules include imitating the best in the neighborhood (Nowak and May, 1992; Zimmermann and Eguíluz, 2005), or replicating in proportion to the payoff (Hauert and Doebeli, 2004; Tomassini et al., 2006).

Link Evolution. The active agent i , which has $k_i \neq 0$ neighbors will, with probability q , attempt to dismiss an interaction with one of its neighbors in the following way. Player i will look at its satisfaction s_i . The higher s_i , the more satisfied the player, since a high satisfaction is a consequence of successful strategic interactions with the neighbors. Thus, the natural tendency is to try to dismiss a link when s_i is low. This is simulated by drawing a uniform pseudo-random number $r \in [0, 1]$ and breaking a link when $r \geq s_i$. Assuming that the decision is taken to cut a link, which one, among the possible k_i , should be chosen? Our solution only relies on the strength of the relevant links. First a neighbor j is chosen with probability proportional to $1 - f_{ij}$, i.e. the stronger the link, the less likely it is that it will be selected. This intuitively corresponds to i 's observation that it is preferable to dismiss an interaction with a neighbor j that has contributed little to i 's payoff over several rounds of play. However, in our system dismissing a link is not free: j may "object" to the decision. The intuitive idea is that, in real social situations, it is seldom possible to take unilateral decisions: often there is a cost associated, and we represent this hidden cost by a probability $1 - (f_{ij} + f_{ji})/2$ with which j may refuse to be cut away. In other words, the link is less likely to be deleted if j appreciates i , i.e. when f_{ji} is high. If the link is not cut there is no further attempt during the current microstep update.

Assuming that the $\{ij\}$ link is finally cut, how is a new link to be formed? The solution adopted here is inspired by the observation that, in social settings, links are usually created more easily between people who have a mutual acquaintance than those who do not. First, a neighbor k is chosen in $V_i \setminus \{j\}$ with probability proportional to f_{ik} , thus favoring neighbors i trusts. Next, k in turn chooses player l in his neighborhood V_k using the same principle, i.e. with probability proportional to f_{kl} . If i and l are not connected, a link $\{il\}$ is created, otherwise the process is repeated in V_l . Again, if the selected node, say m , is not connected to i , a new link $\{im\}$ is established. If this also fails, a new link between i and a randomly chosen node is created. In all cases the new link is initialized with a strength of 0.5 in both directions. This rewiring process is schematically depicted in Fig. 1 for the case in which a link can be successfully established between players i and l thanks to their mutual acquaintance k .

At this point, we would like to stress several important dif-




Figure 1: Illustration of the rewiring of link $\{ij\}$ to $\{il\}$. Agent k is chosen to introduce player l to i (see text).

ferences with previous work in which links can be dismissed in evolutionary games on networks. In (Zimmermann and Eguíluz, 2005), only links between defectors are allowed to be cut unilaterally and the study is restricted to the Prisoner's Dilemma. Instead, in our case, any link has a finite probability to be abandoned, even a profitable link between cooperators if it is recent, although links that are more stable, i.e. have high strengths, are less likely to be rewired. This smoother situation is made possible thanks to our bilateral view of a link which is completely different from the undirected choice made in (Zimmermann and Eguíluz, 2005). It also allows for a moderate amount of "noise" in the system, which could reflect to a certain extent the uncertainties present in the system.

In (Santos et al., 2006a), links can be cut by an unsatisfied player, where the concept of satisfaction is different from ours, and simply means that a cooperator or a defector will wish to break a link with a defector but there is no analogous of our "negotiation" process as the concept of link strength is absent. In (Luthi et al., 2006) links are cut according to a threshold decision rule and are rewired randomly anywhere in the network.

Updating the Link Strengths. Once the chosen agents have gone through their strategy or link update steps, the strengths of the links are updated accordingly in the following way:

$$f_{ij}(t+1) = f_{ij}(t) + \frac{\pi_{ij} - \bar{\pi}_{ij}}{k_i(\pi_{max} - \pi_{min})},$$

where π_{ij} is the payoff of i when interacting with j , $\bar{\pi}_{ij}$ is the payoff earned by i playing with j , if j were to play his other strategy, and π_{max} (π_{min}) is the maximal (minimal) possible payoff obtainable in a single interaction. This update is performed in both directions, i.e. both f_{ij} and f_{ji} are updated $\forall j \in V_i$.

Numerical Simulations

Simulation Parameters. We simulated the Hawks-Doves game with the dynamics described above exploring the entire game space by limiting our study to the variation of only

two game parameters. We set $R = 1$ and $P = 0$ and the two parameters are $1 \leq T \leq 2$ and $0 \leq S \leq 1$. Setting $R = 1$ and $P = 0$ determines the range of S (since $T > R > S > P$) and gives an upper bound of 2 for T , due to the $2R > T + S$ constraint, which ensures that mutual cooperation is preferred over an equal probability of unilateral cooperation and defection. Note however, that the only valid value pairs of (T, S) are those that satisfy the latter constraint.

We simulated networks of size $N = 1000$, randomly generated with an average degree $\bar{k} = 10$ and randomly initialized with 50% cooperators and 50% defectors. In all cases, the parameters are varied between their two bounds in steps of 0.1. For each set of values, we carry out 50 runs of at most 10000 macrosteps each, using a fresh graph realization in each run. After an initial transient period, the system is considered to have reached a pseudo-equilibrium strategy state when the strategy of the agents (C or D) does not change over 150 further macrosteps, which means 15×10^4 individual updates. We speak of pseudo-equilibria or steady states and not of true evolutionary equilibria because there is no analog of equilibrium conditions in the dynamical systems sense.

Cooperation and Stability. Cooperation results in contour plot form are shown in Fig. 2. We remark that, as observed in other structured populations, cooperation is achieved in almost the whole configuration space. Thus, the added degree of freedom represented by the possibility of refusing a partner and choosing a new one does indeed help to find player's arrangements that help cooperation. When considering the dependence on the fluidity parameter q , one sees in Fig. 2 that the higher q , the higher the cooperation level, although the differences are small, since full cooperation prevails already at $q = 0.2$. This was a somewhat expected result, since being able to break ties more often clearly gives cooperators more possibilities for finding and keeping fellow cooperators to interact with. The same effect has been previously observed in (Santos et al., 2006a) with the use of a different model both for strategy evolution and tie breaking. Thus the finding is robust and relatively independent of the other details of the models.

Compared with the level of cooperation observed in simulations in static networks, we can say that results are consistently better for co-evolving networks. For all values of q (Fig. 2) there is significantly more cooperation than what was found in model and real social networks (Luthi et al., 2008) where the same local replicator dynamics was used but with the constraints imposed by the invariant network structure. A comparable high cooperation level has only been found in static scale-free networks (Santos et al., 2006b), which is theoretically interesting, but those topologies are unlikely models for social networks, which often show fat-tailed degree distribution functions but not pure

power-laws (see, for instance, (Amaral et al., 2000; Newman, 2001)). As a further indication of the latter, we shall see later that, indeed, emerging networks do not have a power-law degree distribution.

The above considerations are all the more interesting when one observes that the standard RD result is that the only asymptotically stable state for the game is a polymorphic population in which there is a fraction α of doves and a fraction $1 - \alpha$ of hawks, with α depending on the actual numerical payoff matrix values. To see the positive influence of making and breaking ties we can compare our results with what is prescribed by the standard RD solution. Referring to the payoff table of the Introduction section, let's assume that the column player plays C with probability α and D with probability $1 - \alpha$. In this case, the expected payoffs of the row player are:

$$E_r[C] = \alpha R + (1 - \alpha)S$$

and

$$E_r[D] = \alpha T + (1 - \alpha)P$$

The row player is indifferent to the choice of α when $E_r[C] = E_r[D]$. Solving for α gives:

$$\alpha = \frac{P - S}{R - S - T + P}. \quad (2)$$

Since the game is symmetric, the result for the column player is the same and $(\alpha C, (1 - \alpha)D)$ is a NE in mixed strategies. We have numerically solved the equation for all the sampled points in the game's parameter space. Let us now use the following payoff values in order to bring them within the explored game space (remember that NEs are invariant w.r.t. such an affine transformation):

	C	D
C	(1, 1)	(2/3, 4/3)
D	(4/3, 2/3)	(0, 0)

Substituting in 2 gives $\alpha = 2/3$, i.e. the dynamically stable polymorphic population should be composed by about 2/3 cooperators and 1/3 defectors. Now, if one looks at Fig. 2 at the points where $S = 2/3$ and $T = 4/3$, one can see that the point, and the region around it, is one of full cooperation instead. Even within the limits of the approximations caused by the finite population size and the local dynamics, the non-homogeneous graph structure and an increased level of tie rewiring has allowed the cooperation to be greatly enhanced with respect to the theoretical predictions of standard RD.

Structure of the Emerging Networks

In this section we present a statistical analysis of the global and local properties of the networks that emerge when the pseudo-equilibrium states of the dynamics are attained. First, the mean degree \bar{k} increases only slightly and tends




Figure 2: Average cooperation values for the Hawks-Doves game when the steady-state has been reached. Results are the average of 50 runs.




Figure 3: Average values of the clustering coefficient over 50 runs.

to stabilize around $\bar{k} = 11$. Next, let us consider first the clustering coefficient \mathcal{C} , which was previously defined. Random graphs are locally homogeneous in the average and for them \mathcal{C} is simply equal to the probability of having an edge between any pair of nodes independently. In contrast, real networks have local structures and thus higher values of \mathcal{C} . Fig. 3 gives the average clustering coefficient $\bar{\mathcal{C}} = \frac{1}{50} \sum_{i=1}^{50} \mathcal{C}$ for each sampled point in the Hawks-Doves configuration space, where 50 is the number of network realizations used for each simulation. The networks self-organize through dismissal of partners and choice of new ones and they acquire local structure, since the clustering coefficients are higher than that of the random graph with the same number of edges and nodes, which is $\bar{k}/N = 10/1000 = 0.01$. This effect was expected, since the model favors relinking with closer neighbors rather than arbitrary individuals. The clustering tends to increase with q (i.e. from left to right in Fig. 3).

The *degree distribution function* (DDF) $p(k)$ of a graph represents the probability that a randomly chosen node has degree k . Random graphs are characterized by DDF of Poissonian form $p(k) = \bar{k}^k e^{-\bar{k}} / k!$, while social and technological real networks often show long tails to the right, i.e. there are nodes that have an unusually large number of neighbors (Newman, 2003). In some extreme cases the DDF has a power-law form $p(k) \propto k^{-\gamma}$; the tail is particularly ex-




Figure 4: Empirical cumulative degree distribution functions for three different values of the temptation T . A Poissonian and an exponential distribution are also plotted for comparison. Distributions are discrete, the continuous lines are only a guide for the eye. Lin-log scales.

tended and there is no characteristic degree. The *cumulative degree distribution function* (CDDF) is just the probability that the degree is greater than or equal to k and has the advantage of being less noisy for high degrees. Fig. 4 shows the CDDFs for the Hawks-Doves for three values of T , and $q = 0.5$. A Poisson and an exponential distribution are also shown for comparison. The Poisson curve actually represents the initial degree distribution of the (random) popula-




Figure 5: Empirical cumulative degree distribution functions for three different values of the parameter T . Log-log scales.

tion graph. The distributions are far from the Poissonian that would apply if the networks would remain essentially random. However, they are also far from the power-law type, which would appear as a straight line in the log-log plot of Fig 5. Although a reasonable fit with a single law appears to be difficult, these empirical distributions are closer to exponentials, in particular the curve for $T = 1.7$. It can be observed that the distribution is broader the higher T is. In fact, although cooperation is attained nearly everywhere in the game's configuration space, higher values of the temptation T mean that agents have to rewire their links more extensively, which results in a higher number of neighbors for some players, and thus it leads to a longer tail in the CDDF.




Figure 6: Empirical cumulative degree distribution functions for three different values of the temptation q . Lin-log scales.

The influence of the q parameter on the shape of the degree distribution functions is shown in Fig. 6 where average curves for three values of q , $T = 1.7$, and $S = 0.2$, are reported. For high q , the cooperating steady-state is reached faster, which gives the network less time to rearrange its links. For lower values of q the distributions become broader, despite the fact that rewiring occurs less often, because cooperation in this region is harder to attain and more simulation time is needed.

Cooperator Clusters

From the results of the previous section, it appears that a larger amount of cooperation than what is predicted by the standard theory for mixing populations can be reached when ties can be broken and rewired. We have seen that this dynamics causes the graph to acquire local structure, and thus to lose its initial randomness in terms of links. In other words, the network self-organizes in order to allow players to cooperate as much as possible. At the microscopic, i.e. agent level, this happens through the formation of clusters of players using the same strategy. Fig. 7 shows one typical cooperator cluster.




Figure 7: A typical cooperator cluster. Links to the rest of the network have been suppressed for clarity. The size of a node is proportional to its connectivity in the whole graph. The most connected central cooperator is shown as a square.

In the figure one can clearly see that the central cooperator is a highly connected node and there are many links also between the other neighbors. Such a tightly packed structure has emerged to protect cooperators from defectors that, at earlier times, were trying to link to cooperators to exploit them. These observations help understand why the degree distributions are long-tailed (see previous section), and also the higher values of the clustering coefficient in this case.

Conclusions

In this paper we have introduced a new dynamical population structure for agents playing a series of two-person Hawks and Doves game. The most novel feature of the model is the introduction of a variable strength of the bi-directional social ties between pairs of players. These strengths change dynamically and independently as a function of the relative satisfaction of the two end points when playing with their immediate neighbors in the network. A player may wish to break a tie to a neighbor and the probability of cutting the link is higher the weaker the directed link strength is. The ensemble of weighted links implicitly represent a kind of memory of past encounters although, tech-

nically speaking, the game is not iterated. The model takes into account recent knowledge coming from the analysis of the structure and of the evolution of social networks and, as such, should be a better approximation of real social conflicting situations than static graphs such as regular grids. In particular, new links are not created at random but rather taking into account the “trust” a player may have on her relationally close social environment as reflected by the current strengths of its links. This, of course, is at the origin of the de-randomization and self-organization of the network, with the formation of stable clusters of cooperators. The main result concerning the nature of the pseudo-equilibrium states of the dynamics is that cooperation is greatly enhanced in such a dynamical artificial society. This is encouraging, as the Hawks-Doves game is a paradigm for a number of social and political situations in which aggressivity play an important role. The standard result is that bold behavior does not disappear at evolutionary equilibrium. However, we have seen here that a certain amount of plasticity of the networked society allows for cooperation to be attained. Although the model is an extremely abstract one, it shows that there is place for peaceful resolution of conflict. Ongoing and future work for which there is no space here will deal with the stability of the system against massive and targeted defector invasions in a society of cooperators. Other strategy evolution models based on more refined forms of learning than simple imitation should also be investigated.

Acknowledgements. This work is funded by the Swiss National Science Foundation under grant number 200021-111816/1. We gratefully acknowledge this financial support.

References

- Amaral, L. A. N., Scala, A., Barthélemy, M., and Stanley, H. E. (2000). Classes of small-world networks. *Proceedings of the National Academy of Sciences USA*, 97(21):11149–11152.
- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, Inc., New-York.
- Hauert, C. and Doebeli, M. (2004). Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature*, 428:643–646.
- Hebb, D. O. (1949). *The Organization of Behavior*. Wiley, New York.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.
- Huberman, B. A. and Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences USA*, 90:7716–7718.
- Killingback, T. and Doebeli, M. (1996). Spatial evolutionary game theory: hawks and doves revisited. *Proceedings of the Royal Society of London B*, 263:1135–1144.
- Lindgren, K. and Nordahl, M. G. (1994). Evolutionary dynamics of spatial games. *Physica D*, 75:292–309.
- Luthi, L., Giacobini, M., and Tomassini, M. (2006). A minimal information prisoner’s dilemma on evolving networks. In Rocha, L. M., editor, *Artificial Life X*, pages 438–444, Cambridge, Massachusetts. The MIT Press.
- Luthi, L., Pestelacci, E., and Tomassini, M. (2008). Cooperation and community structure in social networks. *Physica A*, 387:955–966.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- Newman, M. E. J. (2001). Scientific collaboration networks. I. network construction and fundamental results. *Phys. Rev E*, 64:016131.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45:167–256.
- Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359:826–829.
- Poundstone, W. (1992). *The Prisoner’s Dilemma*. Doubleday, New York.
- Santos, F. C. and Pacheco, J. M. (2005). Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys. Rev. Lett.*, 95:098104.
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. (2006a). Cooperation prevails when individuals adjust their social ties. *PLOS Comp. Biol.*, 2:1284–1291.
- Santos, F. C., Pacheco, J. M., and Lenaerts, T. (2006b). Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl. Acad. Sci. USA*, 103:3490–3494.
- Tomassini, M., Luthi, L., and Giacobini, M. (2006). Hawks and doves on small-world networks. *Phys. Rev. E*, 73:016132.
- 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.
- Zimmermann, M. G. and Eguíluz, V. M. (2005). Cooperation, social networks, and the emergence of leadership in a prisoner’s dilemma with adaptive local interactions. *Phys. Rev. E*, 72:056118.