Games on Evolving Networks*

Complex Systems Summer School at Santa Fe Institute,
Santa Fe, NM, June, 2006.

Juan Calderón[1], Adrian de Froment[2], Matina Donaldson[3,4], Greg McInerny[5,6], and Nikhil Kaza[7]**

[1] Universidad de Los Andes, Bogotá
[2] Princeton University
[3] University of Washington, Seattle
[4] Max Planck Institute for Evolutionary Anthropology, Leipzig
[5] Centre for Ecology & Hydrology, Banchory
[6] Institute of Comparative Biology, University of Leeds
[7] University of Illinois at Urbana Champaign

Abstract. This paper explores the coevolution of network formation and cooperative behaviour. It has previously been shown that the structure of a community (for example, as formed by social relationships) can have a substantial impact on the evolution of cooperative behaviour. We hypothesize that the nature of the game played on such a network can alter the way that individuals make and break social connections, therefore leading to a feedback loop between the network structure and the behaviour in the game. We explore the dynamics of this process in a model where individuals form connections with others based on simple genetically encoded rules, and play one of a variety of two-player games with their connections. In order to explore the effects of different payoff structures in these games, simulations over a wide range of parameter space have been pursued, producing games including the Prisoner's Dilemma, the Stag Hunt, and Hawk-Dove. By comparing to traditional game-theoretic models, we show that allowing the network structure to evolve prevents defectors from completely taking over the population, and also broadens the circumstances under which cooperation can fix. A first analysis of the properties of the networks created in these evolving populations suggests a mechanism for these results: when defectors are common, networks are created which tend to retard the effects of selection and encourage drift, while when cooperators are common, the networks formed tend to enhance selection and drive the population more quickly to the fully cooperative outcome.

Introduction

Network qualities can have a significant influence on the outcome of locally played games and dynamics (Watts 1998, Ohtsuki et al. 2006, Santos et al. 2006). The sharing of information through social networks may have considerable group level benefits, whilst posing costs at the individual level (Dall et al. 2005, Wakano et al. 2004). For instance, acquisition of information from network members may be beneficial in a changing environment (Rogers 1988). However, the nature of interactions in such networks may in turn affect the way that the network is constructed. For example, the opportunity to defect and not pay the costs of a group task, such as foraging, may reduce the social connections within a group.

The amount of literature concerning the dynamics of processes occurring on networks far outweighs that investigating the processes that govern the network formation (e.g. see Newman et al. 2006). As such it is pertinent to ask how the rules that individuals within a social network use to create and dissolve connections give rise to the network structure. Furthermore, is there reciprocal

---

* We gratefully acknowledge the support of the Santa Fe Institute, which made this project possible. Special thanks go out to Dan Rockmore and Tom Carter for their efforts. We also gratefully acknowledge the inputs of various participants of the school at different stages of the project.

** Respective emails: ju-cald1@uniandes.edu.co, adriande@princeton.edu, donaldso@eva.mpg.de, gmci@ceh.ac.uk, nkaza@uiuc.edu

feedback between the dynamics of network formation, and the evolution of strategies used within the network?

The integration of game theoretic investigations onto network topologies reveals that cooperation does poorly as a strategy where neighbours are randomly determined (Watts 1998). If the network of links between players has small world properties, this increases the likelihood that co-operative individuals may succeed through "*banding together against the evils of an uncooperative world and scoring points by co-operating with each other*" (Watts 1998). Small world networks that increase the stability of cooperation can be generated and maintained by a mechanism such as the preferential attachment of new nodes to existing high degree nodes (e.g. Dorogovstev et al. 2000). However, it is not clear how behaviour which can maintain small world properties (such as preferential attachment) might emerge from the strategic interactions of game-playing agents.

Several models formulated to develop networks that reproduce features of the topology of food webs (Erdos and Reyni 1960, Cohen et al. 1990, Williams and Martinez 2000) do not directly use the laws that determine the units' (in this case, nodes represent species) behaviour, presenting us with a considerable inverse problem. We know that the mechanisms in such models can generate realistic topology, but not whether those mechanisms are the ones actually operating in nature. It will be of great utility to understand how selection arising from the processes occurring on networks can affect the properties of the networks themselves, because this may give some insight into how real evolutionary processes can shape network structure. It has been demonstrated that network topology in turn may have considerable effects on evolution of behaviour within a network (Lieberman et al. 2005). Taking a coevolutionary approach may provide greater realism as we treat the network as a dynamical entity that supports dynamic processes, and move towards the description of simple laws that underlie complexity.

The social links between individuals of group-living species can be represented as networks, and these networks change dynamically as individuals are born and die, migrate in or out of groups, and change the allocation of their time. The way that edges are formed in this context is crucial to the functioning of the social group as a whole. The investigation we present here is partially motivated by such patterns seen in primate groups. At a very general level of description, some primates such as yellow baboons (*Papio cynocephalus*) appear to form links preferentially with individuals a short social distance (here modelled as network path length) from their mothers, with the probability of interaction decreasing concomitantly with relatedness (Silk et al. ming). We explored the coevolution of cooperation and of a path-length based edge generating rule in a dynamic network. Using a strategic approach we investigated the effect of varying the type of game individuals were playing on the coevolutionary dynamics of network topologies and individual strategies in the game.

Model

The model consists of an agent-based simulation of an evolving social network, with agents represented as game-playing nodes on a graph (see Fig. 1). Each graph is composed of a constant number of nodes, K, each carrying three individual "genetic" traits. Individual nodes form edges with other nodes stochastically, according to probabilities encoded in these three genes. Each individual carries two genes that control the probability of connecting to nodes (1) having edges with its parent and (2) all other nodes. Each individual also carries a gene determining its game strategy, in the form of a probability of cooperating which remains constant across its lifetime. Individuals are constrained within each game to a binary decision of whether to cooperate or not, and do so according to their genetically encoded probability. Initial levels for the genetic parameters are drawn from a random

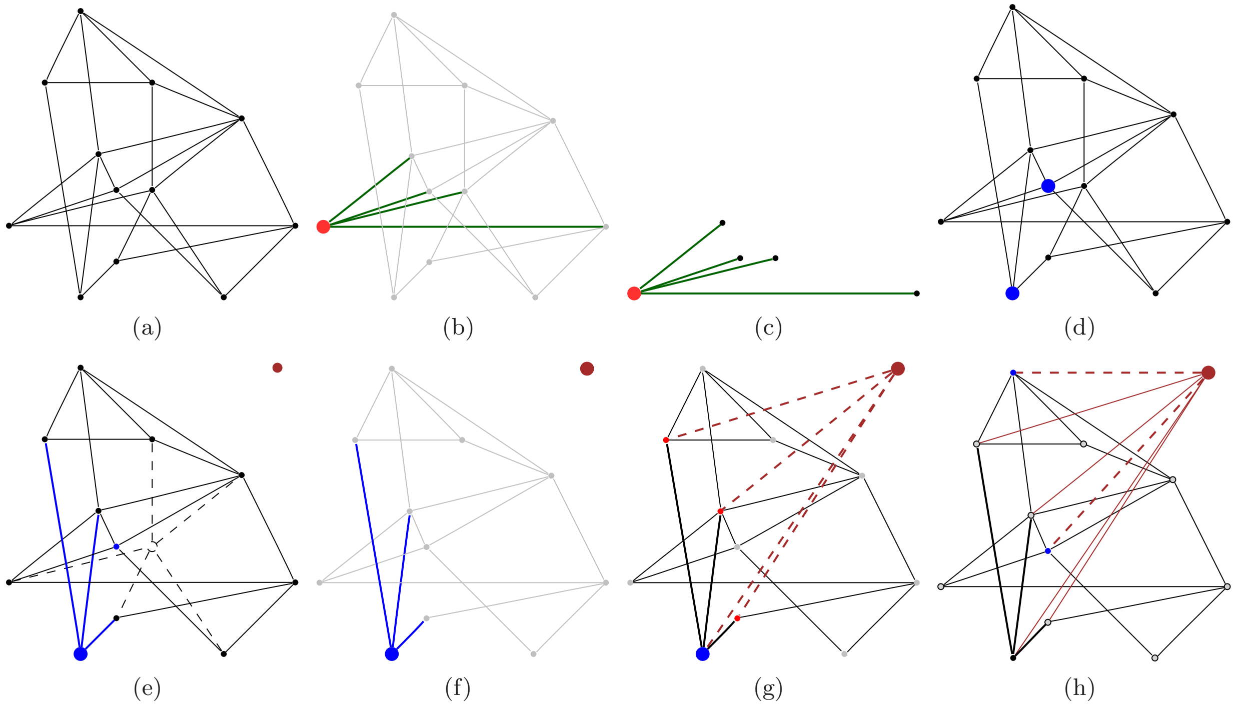

Fig. 1. Schematic of the model. (a) The graph is initialised with randomly generated edges between K nodes. (b) Each node plays N games with each neighboring node, where (c) each individual randomly chooses whether to cooperate or defect according to a gene determining the probability of cooperation. Fitness is recorded as the sum of all games played according to the payoff matrix determined for each simulation run. (d) A tournament between two random individuals determines which node gets to reproduce, and (e) the resulting offspring replaces a randomly chosen node in the network. (f) Offspring inherit all the parental node's genes with mutation, and then (g) form edges with nodes that neighbour the parent according to a genetically determined probability of close links and (h) all other nodes according to a gene which specifies the probability of far links. The new network is used to determine the game partners in the next generation.

uniform distribution between zero and one. Games are played between those nodes which share an edge. The number of games each node plays with each neighbour is given by N.

The fitness of an individual is taken to be the sum of its payoffs from the games it plays, and these fitness are recorded and used to resolve a tournament between two randomly chosen individuals. The node with the greatest fitness reproduces, and the offspring inherit parental genes with a 0.05 mutation (the resulting probabilities being bounded by 0 and 1). The offspring node is added to the group and a randomly selected node is removed. The new node makes connections to individuals near or far in social distance according to the genetically ascribed traits described above. For each experiment this cycle is repeated for g generations. This model was implemented independently in two different programming languages (JAVA and C++) to reduce the chance of errors in model building.

We investigated a family of two-player games which are familiar from many studies of cooperative behaviour: the Prisoner's Dilemma, the Stag Hunt and Hawk-Dove (also known as Chicken, or the Snowdrift game.) We explored a subset of the parameter space of game payoffs which is sufficient to fully characterize the behaviour in these three types of games (Santos et al. 2006): we fixed the payoffs when two individuals playing the same strategy meet, and varied the payoffs for individuals playing different strategies (see Fig. 2). The most well-known of the three games

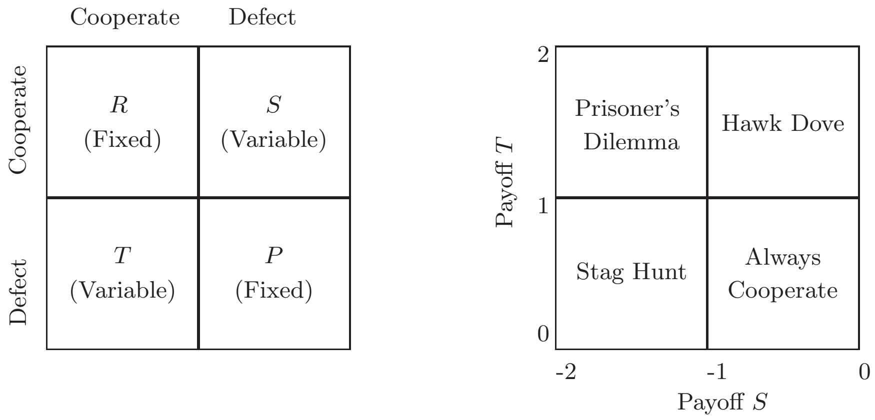

(a) (b)

Fig. 2. Parameterization of game payoffs. (a) A generalized payoff matrix for cooperative games, with the payoffs for the row player indicated as follows: R (reward), S (sucker's payoff), T (temptation to defect) and P (punishment.) (b) In the simulation runs we fixed $R = 1$ and $P = -1$ while varying S and T continuously, yielding a parameter space which can be divided into four regions corresponding to four different types of games.

is the Prisoner's Dilemma, which has a single dominant strategy: always defect. Hawk-Dove has a single mixed-strategy equilibrium, the precise mixture depending upon the payoffs; this mixture may be interpreted in two ways. In a population-level mixed strategy equilibrium, different individuals within the population play different strategies. On the other hand, in an individual-level mixed strategy, each individual in the population plays the same mixed strategy, cooperating and defecting randomly according to some fixed proportion. Since our agents' strategies consist of probabilities of cooperation, both of these interpretations are possible. Finally, the Stag Hunt game has two equilibria, one where all individuals cooperate and another where all defect; which one is reached depends not only on the payoffs, but also on the initial makeup of strategies in the population. In evolutionary game theory, these equilibria are taken as predictions for the long-term evolutionary outcome of a population playing such games. Our model differs from traditional game-theoretic models, however, in that the games are played on a social network, which is created by rules at the individual level—and these rules are permitted to evolve. How does this change the long-term cooperative behaviour of an evolving population in each of these types of games?

Methods

In order to explore the effects of different payoff structures in these games, we ran simulations over a wide range of parameters, encompassing Prisoner's Dilemma, Stag Hunt, and Hawk-Dove as well as a region in which the dominant strategy is to cooperate (see Fig. 2.) We kept track of several measures designed to distinguish between the three different types of equilibrium structure exhibited by these games. For each parameter set we ran 10 independent simulations of 20,000 generations each, and recorded the following statistics:

1. Average strategy in the population. We measured the average strategy over all individuals in the population at the end of each simulation, and then averaged over the different simulations. This indicates which equilibrium was reached, if all individuals are fixed for one strategy, but does not distinguish between individual-level mixed equilibria, population-level mixed equilibria, and multiple pure equilibria which may be reached in different simulations.

2. Variance within the population. We measured the variance in strategy within the population at the end of each simulation, and then averaged over the different simulations. This will be high when a population-level mixed equilibrium is reached, as might be expected for Hawk-Dove games, or when the simulation has not really converged.
3. Variance over time. We measured the average strategy in the population at 100-generation time points over the course of the simulation, and looked for the variance over time in each simulation. This variance was then averaged over the different simulations. This measure will be high when the simulation does not converge quickly because selection is very weak, or because it keeps switching between multiple equilibria.
4. Variance between simulations. We measured the average strategy in the population at the end of each simulation, and then looked at the variance over the different simulations. This measure will be high when the simulation has not really converged, or when it does not reliably converge to the same equilibrium every time. This latter may occur if there are multiple equilibria, and which one is reached depends upon the initial conditions in the population, such as in Stag Hunt games.

We used this set of four statistics to describe the evolution of each of the three genetic variables: one dictating the probability of cooperation, and two describing the rules for creating links. We also compared the cooperative behaviour in our model, where rules for creating links are allowed to evolve, with the cooperative behaviour seen in a simplified version where each node always connects to every other node. This simplified version corresponds to traditional game-theoretical models in which all individuals in a population are equally likely to meet.

Finally, in order to explore how the processes interact to give rise to the evolutionary patterns we observe, we looked at several representative parameter sets in more detail. We looked at the change over generations in the way that properties of individuals are correlated, in one simulation for each of the four types of games. Specifically, we looked for a relationship between the probability of cooperation and the probability of close and far links, between the probability of cooperation and the degree of a node, and between the degree of a node and its fitness. In order to better understand how the structure of the graph changes over time, we also looked at the dynamics of the assortative mixing coefficient, which is a measure of how often nodes with similar degree tend to be connected with one another(Newman 2002).

In all simulations, the number of nodes, K, was fixed at 25, while the number of games played per generation, N, was fixed at 10.

Results

First of all, we confirm that the results from the simplified model where all individuals are connected conform to the predictions from traditional game theory (see first column, Fig. 3.) Under Prisoner's Dilemma conditions, shown in the upper left-hand quadrant of each graph, all individuals defect; in the opposite quadrant all individuals cooperate. Since all three variance statistics are zero in both these quadrants, we conclude that the populations quickly and reliably converge to a single pure-strategy equilibrium. In the Hawk-Dove region, shown in the upper right hand quadrant, the average population strategy ends up a mixed one, but does this occur at the individual level or at the population level? Since the variance within the population is relatively large here, we can infer that at least some of the mixing occurs at the population level; since the other two variances are zero we see that convergence to the mixed strategy is quick and reliable. In the Stag Hunt region, shown in the lower left-hand quadrant, the population usually converges to either all cooperators or all defectors, but in a small region of parameter space some variation is seen. All this variation

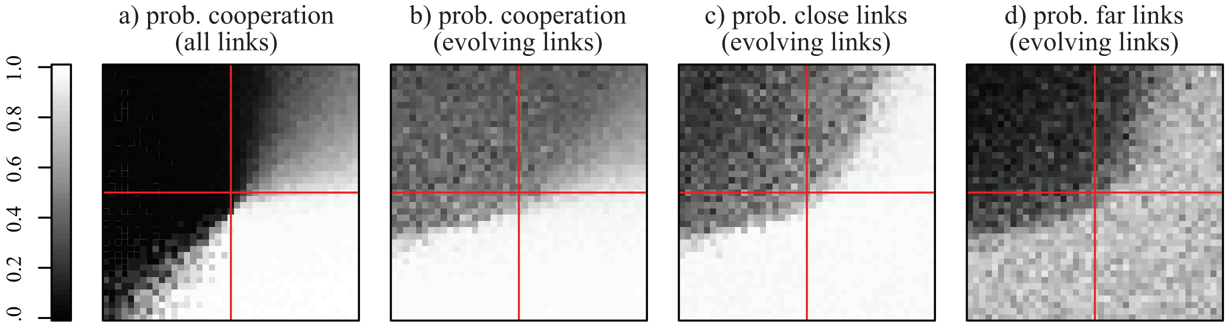

1) average strategy in the population

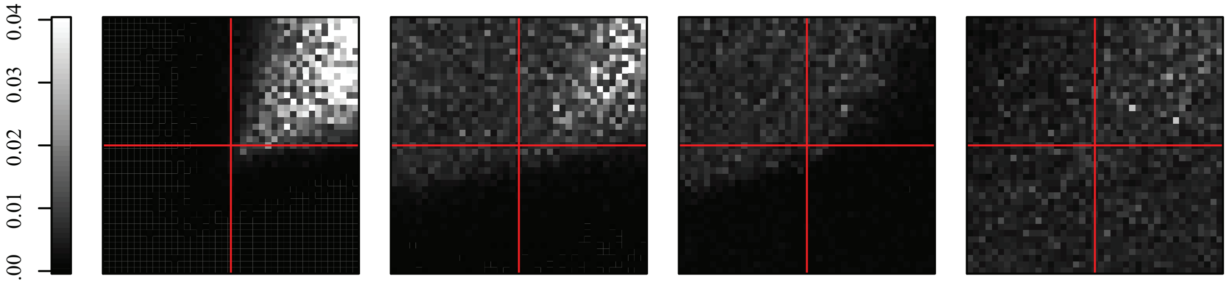

2) variance within the population

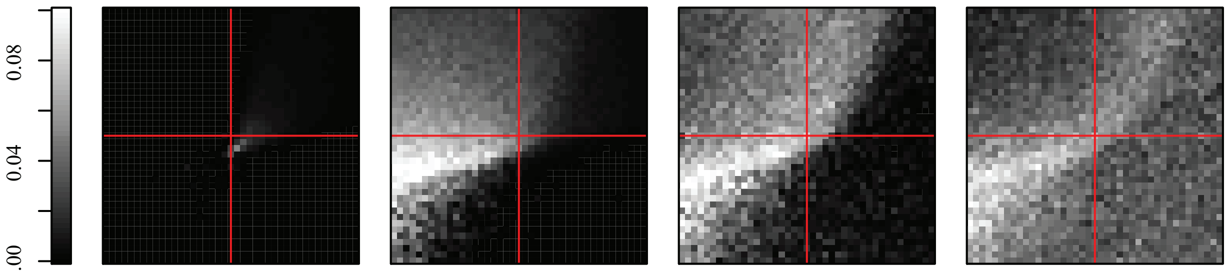

3) variance over time

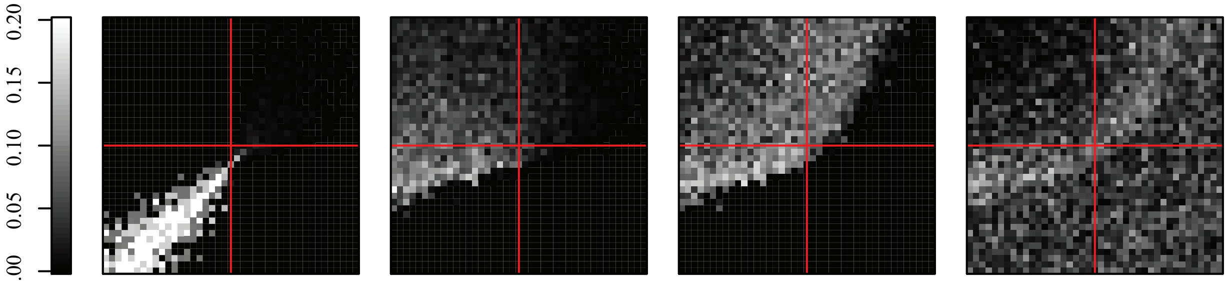

4) variance between simulations

Fig. 3. Evolution of strategies in different kinds of games. Each plot is made over parameters of the payoff matrix, as illustrated in Fig. 2b. The first column of plots (a) shows how the probability of cooperation evolves when all nodes always make connections to all others. This is used as a comparison for our model, in which three variables evolve: (b) the probability of cooperation, shown in the second column, (c) the probability of connecting with the parent's neighbours, shown in the third column, and (d) the probability of connecting with nodes not linked to the parent, shown in the fourth column. Each row displays a different statistic: (1) the average strategy in the population at the end of the simulation, (2) the variance of the strategy in the population at the end of the simulation, (3) the variance of the average population strategy, measured over generations within a simulation, and (4) the variance over different simulations of the average population strategy at the end of the simulation (see Methods.)

occurs between simulations, so we infer that the population quickly converges to one of the two pure-strategy equilibria, depending on the initial configuration of the population, and stays there.

Now we can compare the results of our model with the predictions of evolutionary game theory by looking for differences between the first two columns of Fig. 3. The first thing to notice is that complete defection never seems to be a stable outcome in the long term. Over all parameter values, the average probability of cooperation is always either the same or greater under our model than it is under the classical model, and the greatest changes occur in the two left-hand quadrants, the Prisoner's Dilemma and the Stag Hunt. In the Prisoner's Dilemma quadrant, the average strategy at the endpoint is now quite moderate—and quite variable, by all measures. This suggests that either selection on cooperative behaviour is rather weak, or that it is constantly changing directions. In the Stag Hunt region, the conditions under which full cooperation seems to be a stable long-term outcome have greatly expanded, and where it is not stable, the population seems to swing back and forth between all cooperators and all defectors (because the variance over time and between simulations are relatively large, while the variation within the population is small.) On the other hand, in the lower right-hand quadrant, where cooperation is a dominant strategy, the population quickly and reliably evolves to a fully cooperative state, just as in the classical model. In the Hawk-Dove region, again the population-average results are quite similar to the simpler model, but the mixed strategy seems to be created less often at the population level.

How do the rules for making connections evolve to cause these changes in the evolution of cooperative behaviour? In the third column of Fig. 3 we see the probability of connecting with nodes which were connected to the parent node; when this is high, such nodes are more likely to be related to the focal node, and thus have a similar probability of cooperation. Unsurprisingly, in the regions where full cooperation is the rule, the population always ends up making connections to all the parent's neighbours; in the other regions the probability of close links is much lower, and quite variable at all levels. The probability of connecting with all other nodes is generally lower and more variable, though it follows the same clear pattern: low where defectors are common, and higher where cooperators take over. This is because adding connections to cooperators has a positive effect, while adding connections to defectors has a negative effect; therefore when cooperators are common having more connections of any kind is good, while when defectors are common, having fewer connections is better. Considered together, the patterns seen in these two genes clearly divide the parameter space into two distinct regions that—we will see—display very different network structures.

In the usual Stag Hunt game, the population quickly converges to either all cooperators or all defectors, with a narrow range of parameters where either one is possible. When the network construction rules can evolve, the conditions under which cooperation is a long-term stable outcome greatly expand, and where cooperation is not stable, instead of converging immediately to all defectors, the population instead seems to switch between the two extremes repeatedly. One explanation for these two facts is that adding the evolving network structure somehow makes the all-defector equilibrium a bit less stable, and easier to escape, while the all-cooperator equilibrium becomes even more stable, and more difficult to escape. When there are multiple equilibria in a game, a good long-term prediction of which equilibrium will be observed depends upon the stability of each equilibrium against invasion from the other (Kandori et al. 1993). This would make the cooperative equilibrium the more likely long-term result over a larger range of parameter space.

We explored the details of the network structure more carefully in four simulations, chosen to represent each of the four types of games: Prisoner's Dilemma ($T = 2, S = -2$), Stag Hunt ($T = .6, S = -2$), Hawk-Dove ($T = 2, S = 0$), and All Cooperate ($T = 0, S = 0$) (see Figs. 4-7). Assortativity, the extent to which nodes of similar degree have edges linking them, has important implications for the dynamics of interactions on networks and is measured as the Pearson correlation

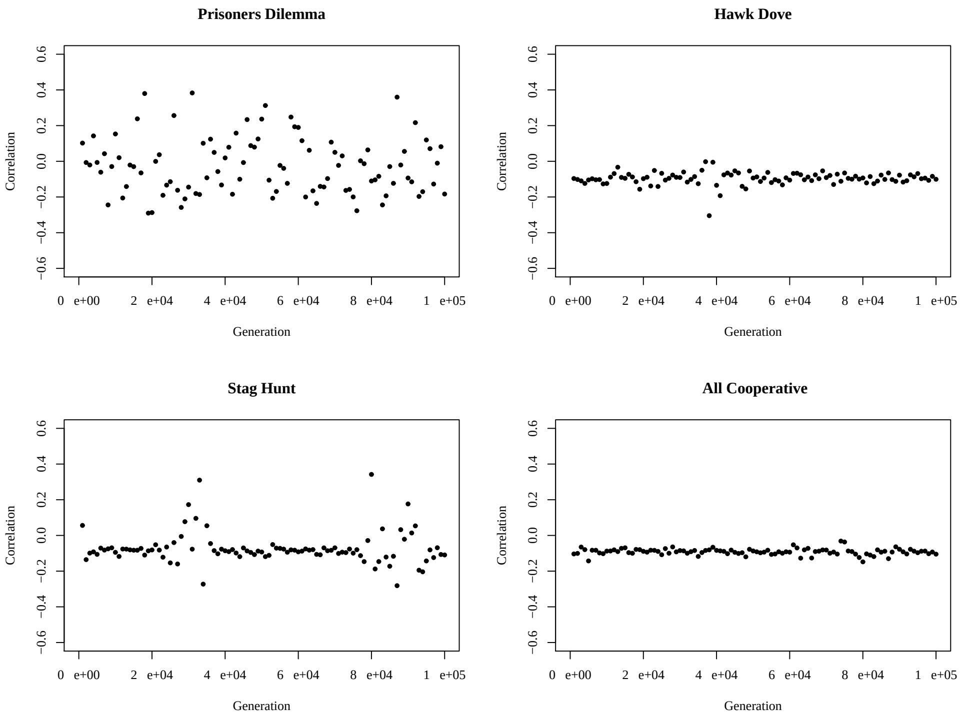

Fig. 4. Assortative mixing coefficients of the networks in various games.

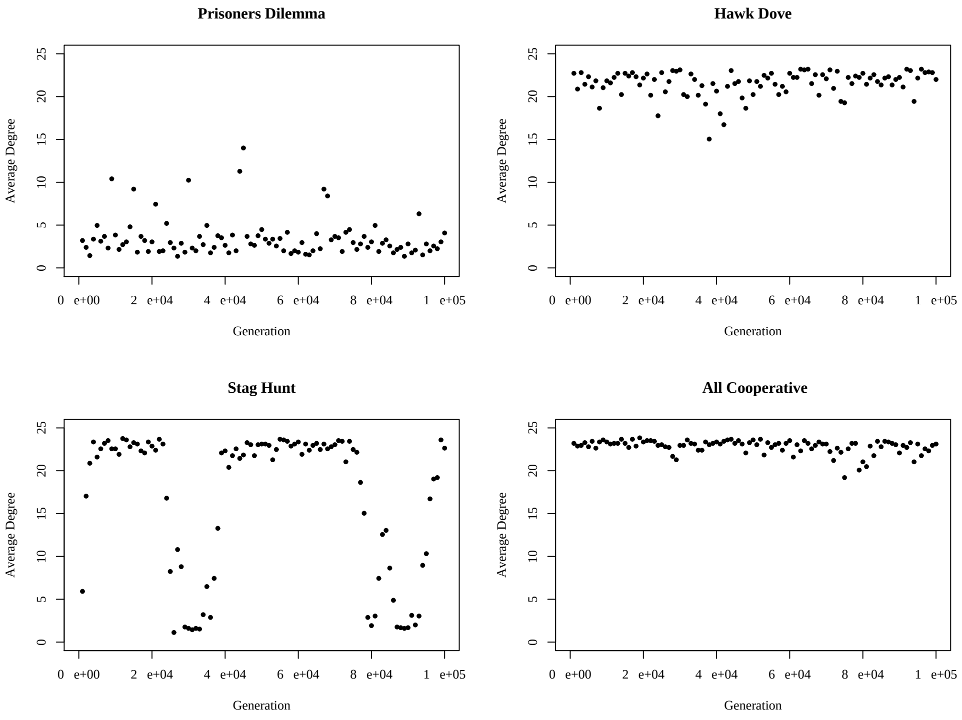

Fig. 5. Average degree in the population in various games.

between the degrees of the vertices at either end of all the edges in a graph (Newman 2002). Under the Prisoners' Dilemma regime we find no pattern in the correlation, but under the Hawk-Dove, Stag Hunt and all cooperative regimes there is a negative correlation between the degrees of the nodes at either end of an edge (Fig. 4 and table 1). This constitutes disassortative connection between the individuals in our simulations: there is a tendency for high degree vertices to be connected with low degree vertices. Although it is still unclear exactly why we see this negative correlation forming, it does seem to be strongly related to the average degree in the population (Fig. 5). In the regions of parameter space where the probability of close links is fixed (all cooperative, plus much of Hawk-Dove and Stag Hunt) the average degree is quite high and the assortativity is slightly negative. Everywhere else (mainly the Prisoner's Dilemma region) the average degree in the graph is low, and the assortative mixing coefficient seems to vary randomly.

Table 1. Assortative mixing coefficients

	Mean	**Variance**
Prisoners Dilemma	-0.030	0.0291
Hawk Dove	-0.092	0.0007
Stag Hunt	-0.084	0.0060
All Cooperative	-0.091	0.0004

The pattern in the within-population correlation between the fitness of a node and its degree (Fig. 6) can be predicted by the average probability of cooperation in the population (Fig. 7). There was a positive correlation between the fitness of a node and its degree in the Hawk-Dove and all cooperative regimes, and a negative one in the Prisoners' Dilemma game, while in the Stag Hunt game the correlation oscillated between negative and positive during the course of a simulation run. In the Prisoners' Dilemma quadrant individuals are more likely to be playing a defector than a cooperator, and so we expect higher degree nodes to lose more and have a lower total payoff and fitness. In the Stag Hunt, as the population switches between the defect and cooperate equilibria, the payoff benefit of being highly connected oscillates along with it. Under the Hawk-Dove game there is higher variance in the correlation because there is variance in the average probability of cooperation in the population, whereas in the all cooperate regime the variance is lower because the population settles at the equilibrium of universal cooperation. Simply put, where there are more cooperators you do well to play more games and pocket the winnings.

In the Hawk-Dove region, where a mixture of cooperating and defecting individuals is maintained, one might expect that the optimal strategy for forming links would differ between the types. Cooperative individuals should preferentially attach to their mother's connections, which are more likely to be related to them, and thus more likely to show similar levels of cooperation. On the other hand, individuals who defect more often might rather connect to random individuals than to their relatives. However, there were no within-population correlations between the probability of cooperation and either of the other two genetic traits (probability of forming links to close individuals and probability of forming links to far individuals), or between the probability of cooperation and the degree of a node (not shown.) This indicates that the stable mixed equilibrium in the Hawk-Dove region does not seem to consist of cooperators who make connections with their close kin and defectors who make farther-ranging connections. In fact, while there is clearly stabilizing selection on the probability of cooperation and of creating close links, there may be relatively little selection on the probability of creating far links once the equilibrium has been reached.

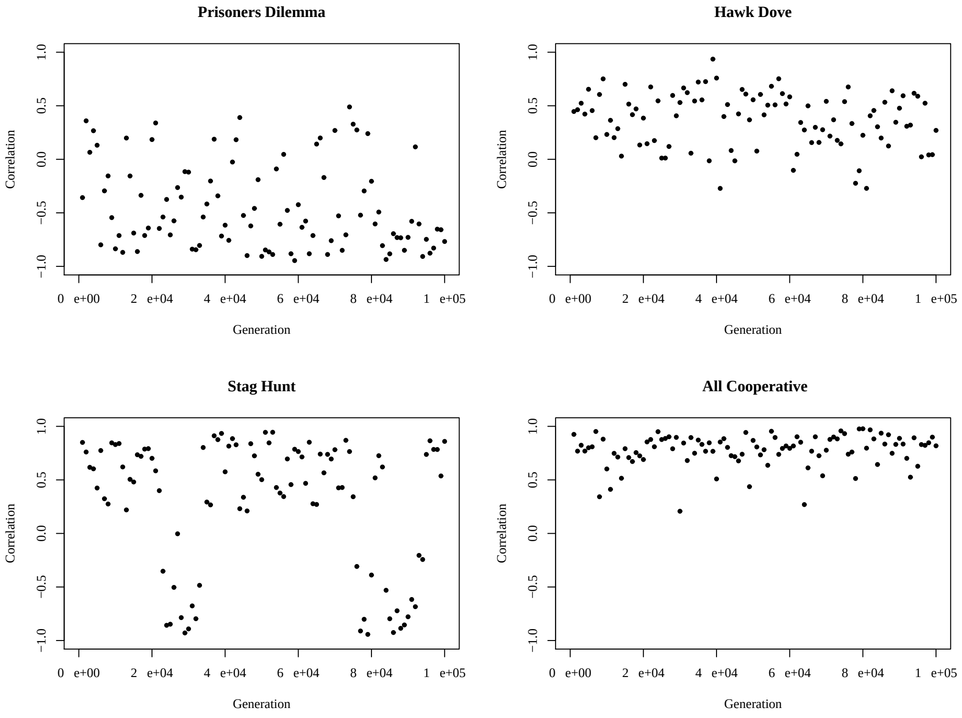

Fig. 6. Correlation between degree and payoff of a node within the population in various games.

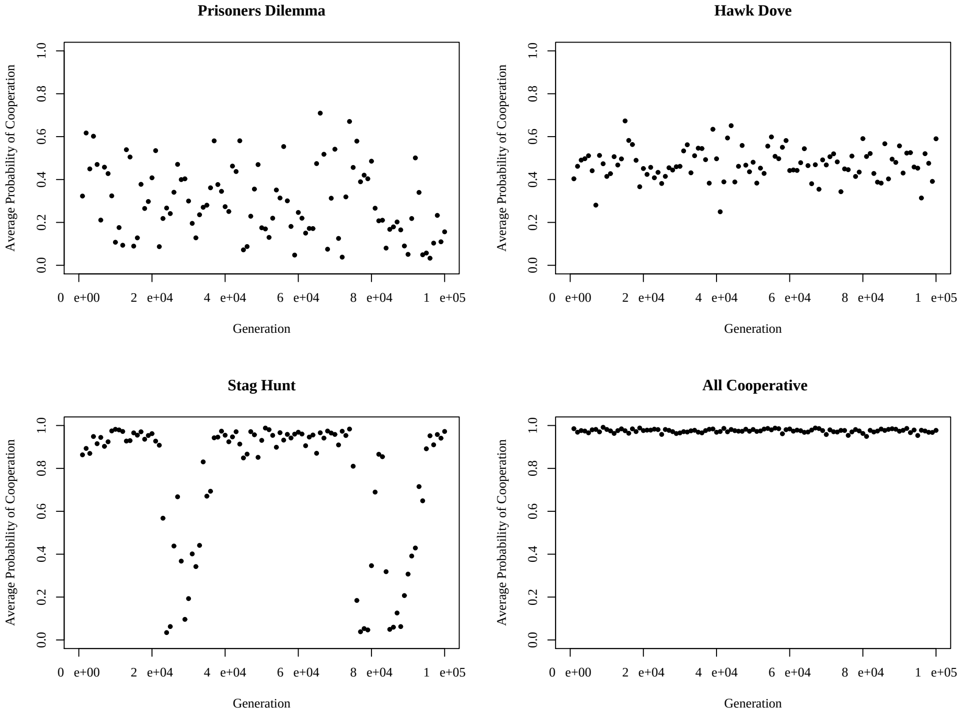

Fig. 7. Average probability of cooperation in various games.

Conclusions

We explored the coevolution of a path-length based edge generating rule and of the behaviour of individuals in various two-player games played on a dynamic network. Through simulations across a range of parameters encompassing four commonly studied game structures, we demonstrated a feedback loop between network topology and game dynamics which acts to encourage the evolution of cooperation under a wide variety of circumstances. The populations tended to spend most of their time in one of two states: a stable, highly-connected, all-cooperative equilibrium, or an unstable state with many defectors and few connections.

For instance under the Prisoner's Dilemma regime we found that allowing the edge-forming rule to evolve led to the disappearance of the all-defect equilibrium. Here the coevolutionary interaction undermines the stability of defection through the feedback between structure and game dynamics. We also found that the range of parameters leading to an all-cooperative outcome was greatly expanded in the Stag Hunt regime. Lieberman et al. (2005) clearly demonstrated how network structure may affect evolution on their graphs through amplification or suppression of selection in contrast to the processes of genetic drift. We found the variance in all three genetic variables in the non-cooperative regions to be higher, suggesting that selection is less effective there than in the regions where cooperation fixes. The kinds of graphs formed in the less cooperative regions may tend to dampen the effects of selection, thereby destabilizing the all-defect equilibrium. This could explain the puzzling observation that in the Prisoner's Dilemma game, the all-defect equilibrium is never stable. We also speculate that selection may be enhanced in the more cooperative regions, further stabilizing the all-cooperate equilibrium. This would explain the enhanced cooperative behaviour found in the Stag Hunt game.

This is supported by our finding of two basic network types. Across a large portion of the parameter space which encompasses the all cooperate region and most of the Hawk-Dove and Stag Hunt regions, the networks formed were highly connected and mildly disassortative (Fig. 4: Stag Hunt, Hawk-Dove, and All Cooperate). However, in the remaining region the networks were not well connected, with the assortativity varying randomly over time (Fig. 4: Prisoner's Dilemma). Where full cooperation is the rule, connections are made to all the parent's neighbours, with a high tendency of connecting to other nodes as well. This creates a well-connected graph which in turn reinforces the stability of the cooperative outcome. Connections are decreasingly made with parental and non parental neighbours as defection pervades the population, with a concomitant increase in the variation within the population and over time, so the population never reaches full defection (Fig. 3).

It is well known that the simple addition of spatial processes can produce heterogeneity in interactions that may stabilize greater diversity in model systems and produce novel insights (for behavioural systems and ecological systems, respectively, see e.g. Jansen and van Baalen 2006, Murrell et al. 2002). We have presented a similar process acting within the novel context of an evolving graph. In contrast to highly symmetric, static spatial models, within an evolving network the process of selection may be accentuated or hindered as the topology is engineered in concert with its dynamics.

Some further work to elucidate the finer scale processes is warranted, along with a thorough sensitivity analysis of the payoff structure. Simulations where the network structure or strategies are fixed after pseudo-equilibrium is reached could test how selection acts on both processes. Without this some of the previous explanation put forward as to how cooperation becomes an increasingly stable equilibrium will remain speculative. There is also considerable dynamism found in the Stag Hunt portion of the payoff structure that may remain unexplained without a more detailed analysis of the causal relationships between the individual statistics we have examined.

We have shown that "*scoring points by co-operating with each other*" (Watts 1998) may reinforce the network structure and produce a closely knit group. However it is likely that such processes may not be acting in isolation. For instance, links may also represent potential vectors for costly processes such as disease and parasite transmission, as well as other benefits such as information sharing. Further work investigating the interplay between processes simultaneously utilising a graph is a natural next area of study. How would the degree of antagonism between the directions of selection of multiple processes affect the properties of the network, and the evolution of strategies used on those networks? Could a simple mutually beneficial activity, like grooming, stabilize the structure of the interaction network enough that cooperation under more challenging circumstances, such as group hunting, could be maintained?

Bibliography

Cohen, J. E., F. Briand, and C. M. Newman (1990). *Community food webs: Data and theory.* New York, NY: Springer-Verlag.

Dall, S. R. X., L. A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology and Evolution 20*, 187–193.

Dorogovstev, S. N., J. F. F. Mendes, and A. N. Samukhin (2000). Structure of growing networks with preferential linking. *Physical Review Letters 85*(21), 4633–4632.

Erdos, P. and A. Reyni (1960). On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kut. Int. Kozl 5*, 17–61. Reprinted in the art of counting: selected writings (ed. J. Spencer), pp. 574-617. Cambridge: M.I.T. Press (1973).

Jansen, V. and M. van Baalen (2006). Altruism through beard chromodynamics. *Nature 440*, 663–666.

Kandori, M., G. J. Mailath, and R. Rob (1993). Learning, mutation and long run equilibria in games. *Econometrica 61*, 29–56.

Lieberman, E., C. Hauert, and M. A. Nowak (2005). Evolutionary dynamics on graphs. *Nature 433*, 312–316.

Murrell, D., D. Purves, and R. Law (2002). Intraspecific aggregation and species coexistence. *Trends in Ecology and Evolution 17*(5), 211–211.

Newman, M. E. J. (2002). Assortative mixing in networks. *Physical Review Letters 89*(20), 208701.

Newman, M. E. J., A. L. Barabasi, and D. J. Watts (2006). *The structure and dynamics of networks.* Princeton Studies in Complexity. Princeton, NJ: Princeton University Press.

Ohtsuki, H., C. Hauert, E. Lieberman, and M. Nowak (2006). A simple rule for the evolution of cooperation on graphs and social networks. *Nature 441*, 502–505.

Rogers, A. (1988). Does biology constrain culture? *American Anthropologist 90*(4), 189–831.

Santos, F. C., J. M. Pacheco, and T. Lenaerts (2006). Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. of National Academy of Sciences 103*, 3490–3494.

Silk, J., J. Altmann, and S. Alberts (forthcoming). Social relationships among adult female baboons (*Papio cynocephalus*): I. variation in the strength of social bonds. *Behavioral Ecology & Sociobiology*.

Wakano, J., K. Aoki, and M. Feldman (2004). A mathematical analysis of social learning. *Theoretical Population Biology 66*, 249–258.

Watts, D. J. (1998). *Small worlds: The dynamics of networks between order and randomness.* Princeton Studies in Complexity. Princeton, NJ: Princeton University Press.

Williams, R. J. and N. D. Martinez (2000). Simple rules yield complex food webs. *Nature 404*, 180–183.