J Evol Econ (2009) 19:379–396
DOI 10.1007/s00191-008-0117-5

REGULAR ARTICLE

Reciprocity in evolving social networks

Tackseung Jun · Rajiv Sethi

Published online: 30 August 2008

Abstract We study the viability of conditional cooperation in a dynamically evolving social network. The network possesses the *small world* property, with high clustering coefficient but low characteristic path length. The interaction among linked individuals takes the form of a multiperson prisoners' dilemma, and actions can be conditioned on the past behavior of one's neighbors. Individuals adjust their strategies based on performance within their neighborhood, and both strategies and the network itself are subject to random perturbation. We find that the long-run frequency of cooperation is higher under the following conditions: (i) the interaction radius is neither too small nor too large, (ii) clustering is high and characteristic path length low, (iii) the mutation rate of strategies is small, and (iv) the rate of adjustment in strategies is neither too fast nor too slow.

Keywords Evolution of cooperation · Reciprocity · Small world networks

JEL Classification C72 · D64

T. Jun (✉)
Department of Economics, Kyung Hee University,
Seoul 130-701, South Korea
e-mail: tj32k@daum.net

R. Sethi
Department of Economics, Barnard College, Columbia University,
3009 Broadway, New York, NY 10027, USA
e-mail: rs328@columbia.edu

1 Introduction

The willingness of individuals to set aside their own self-interest in favor of conditional cooperation or reciprocity has been widely documented in experimental work (See Fehr and Gächter 2000). Accounting for such propensities in evolutionary models has therefore been an active area of research in economics. Early evolutionary approaches to this question were based on a very simple interaction structure, such as pairwise random matching in large populations (Trivers 1971; Axelrod and Hamilton 1981). More recently, attention has been paid to evolution in structured populations with local interaction (Nowak and May 1992; Eshel et al. 1998; Masuda and Aihara 2003; Szabo and Fath 2006). We build on this literature by examining the dynamics of reciprocity in small world networks, when both strategies and the network itself are evolving.

Evolutionary models of cooperation with local interaction have commonly been based on regular networks. For instance, Eshel et al. (1998) consider individuals arrayed in a circle and Nowak and May (1992, 1993) consider a regular two-dimensional lattice (see also Lindgren and Nordahl 1994; Kirchkamp 1995; Cohen et al. 1999).[1] Despite their analytical convenience and popularity, such networks do not correspond closely to real world social networks. Neither is it the case that social linkages are completely unstructured, as in the case of random networks. Social networks lie somewhere between the two extremes of complete regularity and complete randomness: they are highly clustered like regular networks, but any two individuals are connected by a relative small chain of links, as in random networks. These properties are possessed by *small world* networks (Watts 1999; Watts and Strogatz 1998), and the evolution of cooperation on such networks has begun to be explored (Kim et al. 2002; Abramson and Kuperman 2001; Szabo and Vukov 2004).

In this paper we allow the social network as well as the propensity of individuals to cooperate to evolve over time. Changes in the network arise through a process of random rewiring, which allows for existing links to be severed as new ones are formed. This reflects the fact in modern human populations social networks are altered as individuals switch locations, occupations, and memberships of clubs and other organizations. Changes over time also occur in the strategies used, or the extent of conditional cooperation, based on the payoffs to the various strategies adopted in one's neighborhood.

Specifically, our model has the following features. Individuals in a finite population are represented as nodes in a small world network. Each individual plays a multi-player Prisoners' Dilemma game involving all individuals with whom she is linked. Two actions are available: cooperation entails a private cost but yields a benefit to one's neighbors and defection costs nothing but yields no benefit either. Individual strategies are conditioned on the past

[1] Recently, Wu et al. (2006) consider a model of cooperation on a disordered square lattice, in which some links on the lattice are randomly rewired.

behavior of their neighbors as in Watts (1999): each individual cooperates if and only if the proportion cooperating in her neighborhood in the previous period exceeds some threshold. This threshold or "hardness" itself varies across individuals and evolves over time in response to observed payoffs in the neighborhood. When cooperation is more highly rewarded than defection in one's neighborhood, hardness is adjusted downwards, making cooperation more likely in future periods. Similarly, hardness is adjusted upwards when defection is more highly rewarded than cooperation in one's neighborhood. We also allow for occasional mutations in individual hardness, and for the periodic deletion and addition of links between randomly-selected individuals (while maintaining the small world property of the network). The key variable of interest is the long run frequency of cooperation, and the manner in which this is affected by changes in the underlying parameters of the model.

We find that the long run frequency of cooperation is highest when the degree of the network is neither too great nor too small. The intuition for this is as follows. The increase in degree has two effects. First, it increases the level of clustering: neighbors of an individual are more likely to be neighbors to each other. This makes it easier for a large cooperative cluster to be formed, which has a positive effect on survival of cooperation. On the other hand, an increase in degree leads to more heterogeneous neighborhoods. This is damaging for the evolution of cooperation since defectors in largely cooperative neighborhoods obtain very high payoffs, which cooperators in neighborhoods occupied largely by defectors perform badly. Hence an increase in behavioral heterogeneity within neighborhoods has a negative impact on the survival of cooperation. Overall, as long as degree of network is not too great, the former effect dominates the latter. However, if degree is too great, the latter effect starts to dominate and further increases in degree result in lower levels of overall cooperation.

We also find that more clustered networks and a shorter average distance between individuals induces a higher long-run frequency of cooperators. To see why, consider the effect of introducing random links. Such links bridge individuals even if they do not share neighbors, and hence reduce both overall clustering and the average distance between individuals. Since cooperators are less likely to survive when they are less clustered, more random links decrease the long run frequency of cooperators. This implies that small world networks are conducive to the survival of cooperation.

We also show that the long run frequency of cooperation is highest when the adjustment rate of hardness is neither too great nor too small. The intuition for this is as follows. If the adjustment rate of hardness is too fast, then cooperators may switch to defectors, before their neighboring defectors are converted to cooperators, which makes it hard for a sufficiently-large cluster of cooperators to be formed. On the other hand, if the adjustment rate of hardness is too slow, then individuals are unwilling to change their actions, which has more detrimental effect on cooperation than defection, since in this case, sufficiently-large cluster of cooperators must exist initially for them to survive and spread. Finally, we find that, with more mutations, which randomly change the hardness of selected individuals, a neighborhood is more likely to become heterogeneous, and therefore it is less likely that cooperative clusters will be formed.

2 Social networks

We represent the members of a population as vertices of an undirected graph Γ. Let N denote the set of individuals $\{1, 2, ..., n\}$ where $n = |N|$, and let ij denote a link connecting vertices i and j. Let E denote the set of links $\{ij \mid i, j \in N\}$ of graph Γ. Then the network is defined by $\Gamma = \langle N, E \rangle$.

The graph is *simple* if there is at most one link between the same pair of vertices. It is *connected* if, for every two individuals i and j with $j \neq i$, there is a sequence of vertices $i, i_1, i_2, ..., i_m, j$ such that $ii_1, i_1 i_2, ..., i_m j \in E$, where $m \geq 0$. Such sequence is called a *path* between i and j. We restrict attention to simple and connected graphs. We say that i and j are *neighbors*, or *linked*, if $ij \in E$. Let $N(i) = \{j \mid ij \in E\}$ denote the set of i's neighbors, called *open neighborhood* or simply *neighborhood* and $k_i = |N(i)|$ denote the total number of neighbors. We also call $N(i) \cup \{i\}$ as *closed neighborhood* of individual i. The *degree* of the graph $k(\Gamma) = \sum_i k_i/n$ is the average number of neighbors per individual. The number of links in the shortest path between i and j is the *distance* $d(i, j)$ between these two vertices.

Important structural properties of the network can be captured by the *characteristic path length* and the *clustering coefficient* of the graph Γ. The *characteristic path length*, denoted $L(\Gamma)$, is the average distance between any two vertices, that is,

$$L(\Gamma) = \frac{1}{n(n-1)/2} \sum_{i=1}^{n} \sum_{j>i}^{n} d(i, j).$$

The clustering coefficient C_i of the open neighborhood $N(i)$ of vertex i is

$$C_i = \frac{1}{k_i(k_i - 1)/2} \left[\frac{1}{2} \sum_{j \in N(i)} \sum_{l \in N(j)} 1_{\{l \in N(i)\}} \right], \tag{1}$$

where 1_{φ} is equal to 1 if φ is true and zero otherwise. The denominator of Eq. 1 is the total number of all possible links in $N(i)$. The numerator is the number of links in $N(i)$, that is, the total number of connections among members of $N(i)$. Hence C_i measures the extent to which neighbors of i are also neighbors of each other, that is, C_i measures the extent of *cliquishness* among the individuals in $N(i)$. The *clustering coefficient* $C(\Gamma)$ of the graph Γ is simply the average of C_i over all vertices, that is.,

$$C(\Gamma) = \frac{1}{n} \sum_{i \in N} C_i.$$

For instance, if everyone knows everyone else, then $C = 1$. To summarize, $L(\Gamma)$ measures the typical distance between two vertices, and $C(\Gamma)$ measures the cliquishness among the members of a typical neighborhood.

Two types of networks have been extensively studied in the literature. At one extreme is the class of *regular* networks, in which each individual i is linked with individuals $i \pm h$ for $1 \leq h \leq k/2$ (mod n). The simplest example of this is a ring network ($k = 2$), in which the vertices are arrayed in a circle and each individual has exactly two neighbors. At the other extreme is the class of *random* networks, in which each individual is linked with k randomly-selected individuals on average.

Neither regular nor random networks are adequate representations of social networks in the real world. In practice, the likelihood that two individuals are linked will be significantly higher if they have one or more neighbors in common. This implies a high clustering coefficient, a property that regular networks (but not random networks satisfy). However, it is also the case that in most real world networks, the smallest path between two randomly selected vertices is very small: 3.5 steps in the a network of actors, 9.5 in a network of scientific collaboration, and just 19 in the Internet, which contains more than 800 million webpages (Barabási 2003). But for a regular network of degree k, $L(\Gamma) = n(n + k - 2)/2k(n - 1)$, which scales linearly with respect to n. Therefore for large n, the average distance between two vertices is very large in regular networks. In contrast, the average distance between two vertices in random networks is small and increases only in the log of population size. Hence social networks lie somewhere between the two extremes networks, satisfying the following properties: (i) the characteristic path length L is short, as in random networks, and (ii) the clustering coefficient C is high, as in regular networks.

3 The model

3.1 Network construction

We construct a social network of (average) degree k using an algorithm based on the α-model of Watts (1999). Starting with a ring lattice, additional links are created in the following manner. Fix a vertex i and for each vertex $j \neq i$, set $R_{ij} = 0$ if $ij \in E$ and, if $ij \notin E$, set

$$R_{ij} = \begin{cases} 1 & \text{if } m_{ij} \geq k \\ \left(\frac{m_{ij}}{k}\right)^{\alpha}(1 - p) + p & \text{if } m_{ij} \in \{1, ..., k - 1\} \\ p & \text{if } m_{ij} = 0, \end{cases} \tag{2}$$

where $m_{ij} = |\{l \in N \mid il \in \Gamma \wedge jl \in \Gamma\}|$ is the number of common neighbors shared by individuals i and j, $p \ll \binom{n}{2}^{-1}$ and α is tunable parameter of the

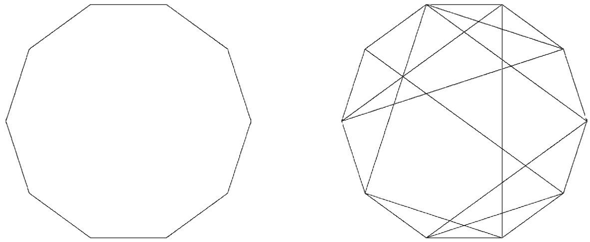

Fig. 1 Creation of a graph using α-model: an initial ring network (*left*) and α-model after one around (*right*)

model. Then the probability that vertex i will be linked to vertex j, denoted by P_{ij}, is computed as follows:

$$P_{ij} = R_{ij} / \sum\nolimits_{l \neq i} R_{il}. \tag{3}$$

According to P_{ij} in Eq. 3, a vertex is selected and it is linked to vertex i. Then this procedure is repeated until all vertices have exactly one chance to add a link. This completes one round of the algorithm. To generate a small world network with average degree $k \geq 4$ (where k is even), we repeat this procedure for $k/2 - 1$ rounds.

Figure 1 illustrates an example of a network created using the α-model with $n = 10$ and $k = 4$, which requires just a single round of connections. Note that some individuals are more 'linked' than others in the resulting network although the average number of connections per vertex is predetermined.

The parameter α affects the structural properties of the resulting network in a systematic way. When α is small but not zero, links will be preferentially attached between individuals who share at least one neighbor. The result is a network with a high degree of clustering and high characteristic path length, which resembles a regular network. At the other extreme, as $\alpha \to \infty$, the resulting network approximates a random graph with negligible clustering and low characteristic path length.[2] Intermediate values of α generate networks with the small world property. Figure 2 shows the effect of changing α on the structure of a network with $n = 80$ and $k = 6$.

The manner in which the characteristic path length and clustering coefficient of the network vary with α is shown in Fig. 3 (based on $n = 1,000$ and $k = 8$).[3] For $\alpha = 0$, no distinction is made in the network construction algorithm between pairs of individuals with just one neighbor in common, and pairs with more than one neighbor in common. The resulting graph has a small clustering coefficient as well as small characteristic path length. As α increases, but remains small, links begin to form preferentially between individuals who

[2]There are some structural differences between a random network and the limiting case of the α-model. However, when $k \gg 1$, the limit very closely approximates a random graph. See Watts (1999, pp. 51–52) for details.

[3]In Fig. 3, $L(\Gamma)$ and $C(\Gamma)$ are normalized by dividing them by $L(\Gamma)$ for $\alpha = 0$ and $C(\Gamma)$ for $\alpha = 0$, respectively.

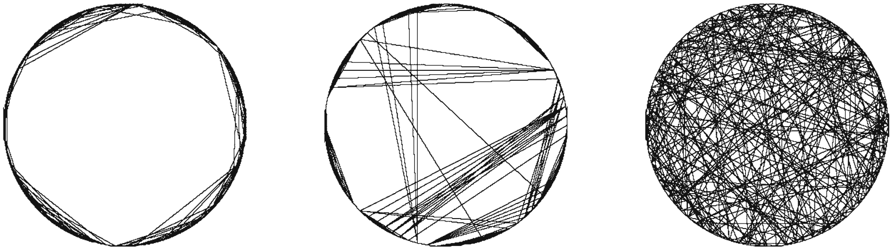

Fig. 2 Effect of changing α on the structure of network: $\alpha = 0, 5, 20$ from the left

have more than one neighbor in common. This results in a graph with high clustering and high characteristic path length. This effect may be seen in the upward sloping segment of the clustering coefficient and characteristic path length curves for $0 < \alpha < 5$ in Fig. 3. As α increases beyond this range, the probability of a link between individuals with less than k but greater than one common neighbors decreases. Hence for fixed k, more links will be created between individuals with no common neighbors. The defining properties of small world networks – short characteristic path length and high clustering coefficient – emerge when α is about 7.[4]

3.2 Payoffs, actions, and dynamics

In every period, each person chooses one of two actions: cooperate or defect. Any individual i choosing to cooperate incurs a private cost γ and provides an aggregate benefit $\beta > \gamma$, which is shared equally among her neighbors $N(i)$. Defection results in no costs or benefits. Without loss of generality, we normalize $\gamma = 1$ and interpret β as the benefit-cost ratio. Since $\beta > 1$, efficiency requires the cooperative action to be taken by all players. From the perspective of any individual, however, defection yields a higher payoff regardless of the actions taken by her neighbors. This is a multi-player Prisoner's Dilemma with local interaction, of the kind studied by Bergstrom and Stark (1993), Eshel et al. (1998), Nowak and May (1993), and Albin and Foley (2001).

Let $s_i(t) = 1$ if individual i cooperates at time t and $s_i(t) = 0$ otherwise. The vector $s(t) = (s_1(t), ..., s_n(t))$ is the state of the actions at time t. Let $\mathcal{S} \equiv \{0, 1\}^n$ denote the set of all states of the actions. The payoff to player i at time t is

$$\pi_i(t) = -s_i(t) + \beta \sum_{j \in N(i)} \frac{s_j(t)}{k_j}, \tag{4}$$

[4]The fact that the characteristic path length and clustering coefficient cannot be tuned independently is a disadvantage of the α-model. However, varying one of these while holding the other constant would cause the average degree k to change, and this would prevent us from identifying the independent effects of changes in the average degree on the incidence of cooperation.

Fig. 3 Scaled characteristic path length and clustering coefficient: $n = 1,000$ and $k = 8$

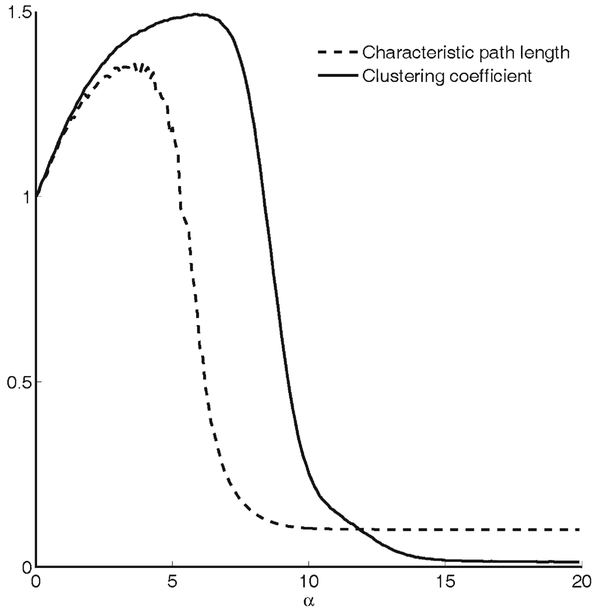

where $s_j(t)/k_j$ is the benefit conferred by neighbor j onto i if j cooperates. Consider any individual i with at least one neighbor taking a different action than i herself does. The average payoff of defectors in the closed neighborhood of individual i is

$$\bar{\pi}_i^d(t) = \frac{\sum_{j\in N(i)\cup\{i\}} \left(1 - s_j(t)\right) \pi_j(t)}{\sum_{j\in N(i)\cup\{i\}} \left(1 - s_j(t)\right)} = \frac{\beta \sum_{j\in N(i)\cup\{i\}} \sum_{l\in N(j)} \left(1 - s_j(t)\right) s_l(t)/k_l}{\sum_{j\in N(i)\cup\{i\}} \left(1 - s_j(t)\right)}. \tag{5}$$

By the definition of $s_j(t)$, the denominator of Eq. 5 is simply the number of defecting neighbors of individuals $i/$, including i herself, at time t. The numerator of Eq. 5 is the sum of payoffs of those defectors in the denominator. Similarly, the average payoff of cooperators in the closed neighborhood of individual i is:

$$\bar{\pi}_i^c(t) = \frac{\sum_{j\in N(i)\cup\{i\}} s_j(t)\pi_j(t)}{\sum_{j\in N(i)\cup\{i\}} s_j(t)} = \frac{\beta \sum_{j\in N(i)\cup\{i\}} \sum_{l\in N(j)} s_j(t) s_l(t)/k_l}{\sum_{j\in N(i)\cup\{i\}} s_j(t)} - 1. \tag{6}$$

The share of cooperators in the neighborhood of i is given by

$$\rho_i(t) = \frac{1}{k_i} \sum_{j\in N(i)} s_j(t).$$

We assume that each individual $\acute{\imath}$'s action in any given period can be conditioned on the past behavior of her neighbors in a manner that reflects some degree of reciprocity. Specifically, individuals cooperate if the level of prior cooperation in their neighborhood exceeds some threshold, where the threshold itself may vary across individuals. Following Watts (1999), we refer to the threshold as an individual's *hardness*. Those with higher levels of hardness are more reluctant to cooperate in the sense that a greater proportion of neighborhood cooperation is necessary in order to induce them to cooperate. This intrinsic preference is not observable to other individuals. Let $h_i \in [0, 1]$ denote the hardness of individual i. The vector $h(t) = (h_1(t), ..., h_n(t))$ is the propensity distribution at time t. Putting this together with the profile of actions, the state of the system at time t is $(h(t), s(t))$.[5]

Note that the model permits unconditional cooperation and defection when hardness is sufficiently low and high, respectively.[6] The choice of action by individual i in period $t+1$ is given by:

$$s_i(t+1) = \begin{cases} 1 & \text{if } h_i(t) < \frac{\rho_i(t)k_i+1}{k_i+2} \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Equation 7 implies that even if there are no cooperators in i's neighborhood (that is, if $\rho_i(t) = 0$), she will still cooperate if $h_i(t) < 1/(k+2)$. Similarly, $h_i(t) \geq (k+1)/(k+2)$, she will defect even if all neighbors cooperated in the previous period. Intermediate values of $h_i(t)$ correspond to varying degrees of conditional cooperation.

We assume that individual hardness changes in response to cooperativeness of neighbors. (In contrast, Watts (1999) assumes that individuals are endowed with a fixed level of hardness.) It is reasonable to assume that individual hardness responds to the average past performance of actions in one's neighborhood. In other words, when the average payoff of cooperators in the neighborhood of individual i is higher than that of defectors, the hardness of individual i is lowered, so that she is more likely to take the cooperative action in the next period. The opposite is the case if defectors outperform cooperators. Specifically:

$$h_i(t+1) = \begin{cases} \delta h_i(t) & \text{if } \bar{\pi}_i^a(t) > \bar{\pi}_i^e(t) \\ h_i(t) + (1-h_i(t))(1-\delta) & \text{if } \bar{\pi}_i^a(t) \leq \bar{\pi}_i^e(t), \end{cases} \quad (8)$$

where $\delta \in [0, 1]$ is the rate of adjustment of hardness.[7]

The evolutionary dynamics may be summarized as follows. Starting with initial set of actions and hardness $(h(0), s(0))$, individuals take actions and receive payoffs in accordance with Eq. 4. Then actions and hardness are updated according to Eqs. 7 and 8 respectively. This process is repeated over

[5]The notion of hardness allows for a multiperson generalization of the Tit-for-Tat strategy. In a two-person repeated interaction, Tit-for-Tat cooperates in the first period and then simply mimics the last action of the opponent. In a multiperson setting, opponent actions are multidimensional. Hardness allows us to aggregate these actions in a natural way, by identifying the threshold proportion of cooperation in one's neighborhood that is sufficient to induce a player to cooperate.

[6]This unconditional response is not only realistic but also technically important when we introduce mutation in hardness. For example, without the unconditional response, mutations in hardness in a homogenous neighborhood cannot lead to changes in actions.

[7]The manner in which hardness evolves ensures that $h_i(t) \in [0, 1]$ for all t.

several periods. The proportion of individuals cooperating in period t is given by $\rho(t) = \sum_{i \in N} s_i(t)/n$, and the average value of hardness in the population at time t is $\sigma(t) = \sum_{i \in N} h_i(t)/n$. We are interested in the long run values of these variables.

In addition to the deterministic dynamics, we allow for mutations in network structure. Almost all social networks change over time as people sever existing relationships and make establish new ones. New contacts may be made systematically through one's existing network of contacts, or may arise for more idiosyncratic reasons. When individuals join a club or move to a different city for employment reasons, two things happen naturally: new connections are formed that may or may not be ordered by the existing network, and some existing links are broken. Hence the evolution of a network can be driven by a mixture of (pure) random connections and connections through the existing network, where the former effect is typically small relative to the latter.

In order to capture this phenomenon, we allow for random *rewiring* of the network, using an algorithm based on Watts' (1999, p.67) β-model. To rewire connections periodically while maintaining the small world property, we proceed as follows. A vertex i is randomly chosen and we disconnect a randomly-chosen connection between vertex i and vertex $j \in N(i)$. (We exclude two geographically immediate neighbors from this disconnection to prevent the vertex i from being isolated from the rest of the network.) We then add a link between vertex i and a vertex $l \in N \backslash N(i)$ according to the probability defined in Eq. 3. The vertex i can be re-linked j, although this is unlikely to occur. This process preserves a high clustering coefficient, and extensive simulations confirm that even after repeated application, the network retains the small world property.[8]

Finally, we also allow mutations in individual propensities to cooperate. With small probability $\varepsilon > 0$, each individual's hardness is replaced by a number drawn from a uniform distribution over the unit-interval. This possibility arises at the end of a round (after actions have been taken). Without loss of generality, we rewire a random vertex and mutate hardness at the same frequency $\omega > 0$. Periodic rewiring and mutation implies that no subset of states will be absorbing in the resulting Markov process, which allows us to estimate the unique level of cooperation in the steady state distribution, i.e., a distribution obtained with the transient behavior in early stage of evolution disregarded.[9]

[8] An important determinant of real world social networks is the *endogenous* formation and breakage of links (Jackson and Wolinsky 1996). One might expect that individuals would seek to form links with those with a higher propensity to cooperate, and break links with defectors. These effects, which we neglect here, would make it easier for cooperative clusters to form and spread. We show that despite the unbiased formation and breakage of links, the incidence of cooperation can be significant.

[9] Note that mutations in the propensity to cooperate are sufficient to ensure that there are no absorbing states, so random rewiring of the network is not necessary for this purpose. We allow for changes in network structure because it occurs in practice, and is likely to affect the extent to which cooperation can persist in the long run.

4 Results

4.1 A baseline

We begin by replicating the results of Eshel et al. (1998). Suppose $n = 1,000$ and $k = 2$ (a ring network), with $\varepsilon = \omega = 0$ (so there is no rewiring and no mutation in hardness). The hardness of each individual is constrained to lie in $\{0, 1\}$, which is achieved by setting $\delta = 0$. Hence

$$h_i(t+1) = \begin{cases} 0 & \text{if } \bar{\pi}_i^a(t) > \bar{\pi}_i^e(t) \\ 1 & \text{if } \bar{\pi}_i^a(t) \leq \bar{\pi}_i^e(t). \end{cases}$$

Since hardness is constrained to lie in $\{0, 1\}$, average hardness is simply the proportion of defectors in the population. In each period, each individual switches to whichever action earned the highest average payoff in her neighborhood in the previous period. This is exactly the model of Eshel et al. (1998). There are two kinds of absorbing states: one in which all individuals defect and another set of states in which at least three-fifths of individuals cooperate.

Adding mutations to hardness in this model (while holding constant network structure) results in a unique steady state distribution. Eshel et al. (1998) show that for sufficiently small mutation rates, the system spends most of its time in predominantly cooperative states. This can be illustrated by examining time series plots for the proportion of cooperation and average hardness. For a mutation rate $\varepsilon = .005$, Fig. 4 shows that hardness remains low and cooperation is widespread and persistent.[10] The long run frequency of cooperation is approximately 80%, consistent with the analytical results of Eshel et al. (1998).[11]

4.2 Steady state distribution

We are interested in the incidence of cooperation in the long run, and the manner in which this is affected by the underlying parameters. The complexity of the model precludes an analytical treatment and we therefore rely on simulations. We set $n = 1,000$, $k = 8$, and $\alpha = 7$ as benchmark parameters. The remaining parameters of the model are as follows: the mutation probability $\varepsilon = 0.005$, the rate of adjustment in hardness $\delta = 0.9$, the baseline probability $p = 10^{-10}$, and the benefit-cost ratio $\beta = 4.91$.[12] Finally, $\omega = 2$, so mutations in individual hardness and network rewiring occur once every two periods.

[10] The simulation was run for 100, 000 periods and first 30, 000 observations deleted. There are 300 periods between adjacent data points in the figure, to smooth out extreme short run volatility.

[11] Eshel et al. prove that the level of cooperation must lie between 60% and 100%. It is interesting that we find cooperation to lie roughly at the midpoint of this range.

[12] In a regular network with degree k, a large enough cluster of cooperators survives in the steady state if the benefit-cost ratio β exceeds than $4k/(3(k/2) - 1)$; see Jun and Sethi (2007). Accordingly, we choose a value higher than this threshold.

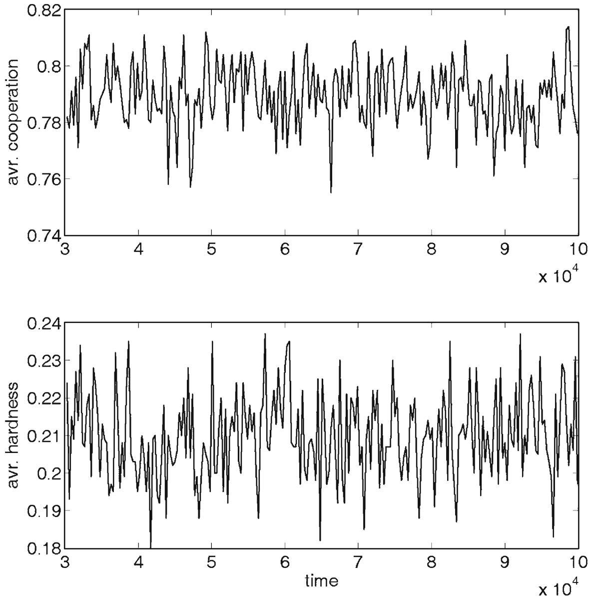

Fig. 4 Evolution of cooperation and hardness for a regular ring network with mutation and rewiring disallowed, corresponding to Eshel et al. (1998)

We keep track of a sequence of average cooperation $\rho(t)$ over time. The average proportion of cooperating individuals in the $T - t$ periods is given by

$$\hat{\rho} = \frac{1}{T-t} \sum_{\tau=t+1}^{T} \rho(\tau)$$

where $T = 100,000$ is the total length of the simulation and $t = 30,000$. That is, the first $30,000$ observations are excluded from the computation of average cooperation in order to eliminate the effects of initial conditions. Figure 5 shows that the long-run average frequency of cooperation $\hat{\rho}$ is about 0.69 and the frequency distribution is mostly concentrated around the mean.

4.3 Comparative statics

In this section, we explore the manner in which the long-run frequency of cooperation depends on the underlying parameters of the model.[13]

4.3.1 Changes in k

As shown in Fig. 6, the long run frequency of cooperation increases as degree of the network k increases, provided that the degree is not to large. As k

[13]For each of the following four cases, we change only one parameter of the model, while fixing all other parameters at the values of Section 4.2.

Fig. 5 Frequency distribution of cooperation: $n = 1,000$ and $k = 8$

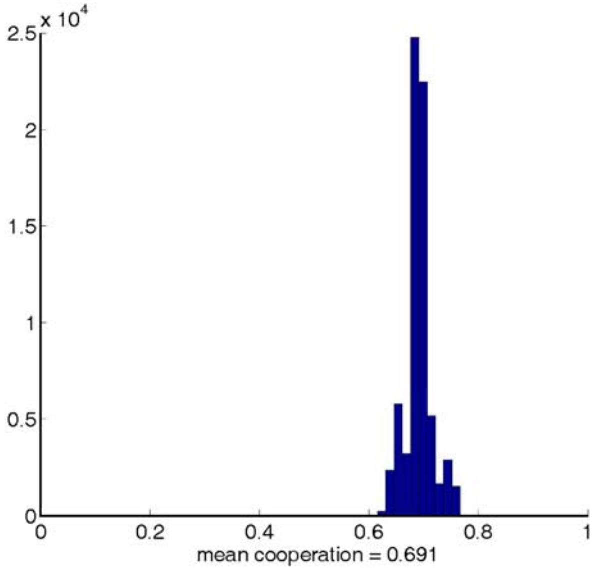

rises from 4 to 34, the frequency of cooperation rises from 0.075 to 0.999, but for $k \geq 34$, the frequency is almost zero.[14] This result can be understood as follows. Increases in the degree has two effects. First, it changes the structure of the network such that individuals have more neighbors on average, and so the number of connected individuals who do not share any other common neighbors declines. If α is relatively small, but not zero, then Eq. 2 implies that an increase in k reduces the probability of a link between individuals with less than k but greater than one common neighbors. This in turn increases the probability that an individual with whom more than k neighbors are shared is selected as a neighbor. Overall, this will make the network more clustered, and more clustered networks induce more cooperation.

Second, an increase in degree affects the distribution of actions in each neighborhood. As k increases, it is more likely that any given individual has a defecting neighbor, earning higher payoffs than surrounding cooperators. This phenomenon, the "puncturing" of cooperative clusters by single defector, has been noted in the literature (Watts 1999; Kim et al. 2002; Abramson and Kuperman 2001) and can lead to the unraveling of cooperation. In contrast, single cooperators who find themselves in clusters of defectors will earn very low payoffs. Therefore the increase in heterogeneity within neighborhoods will have negative impact on the survival of cooperation.

[14] Although not shown, the frequency is almost zero for values of $k \geq 34$.

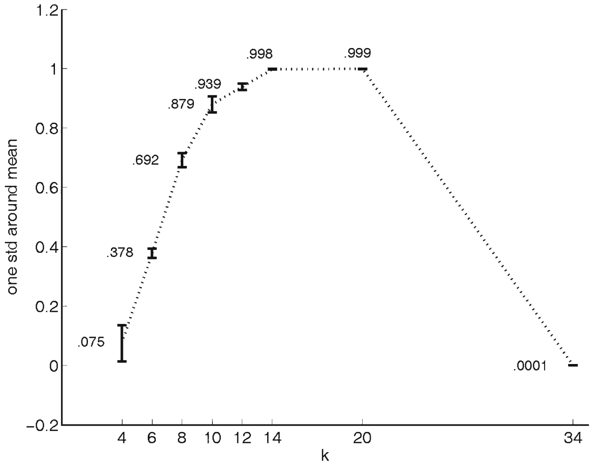

Fig. 6 Effect of k on mean frequency of cooperators

As long as k is not too large, the effect of greater clustering dominates the effect of greater neighborhood heterogeneity. The reverse is true when k gets large, since further increases in degree have little impact on the extent of clustering.

4.3.2 Changes in α

As shown in Fig. 7, as α rises from 4 to 22, the long run frequency of cooperation initially rises but then declines, reaching its peak at around $\alpha = 9$. The reason for this is evident from a closer examination of Fig. 3: the nomonotonicity of the clustering coefficient in α is reflected in the nonmonotonicity of cooperation in α.

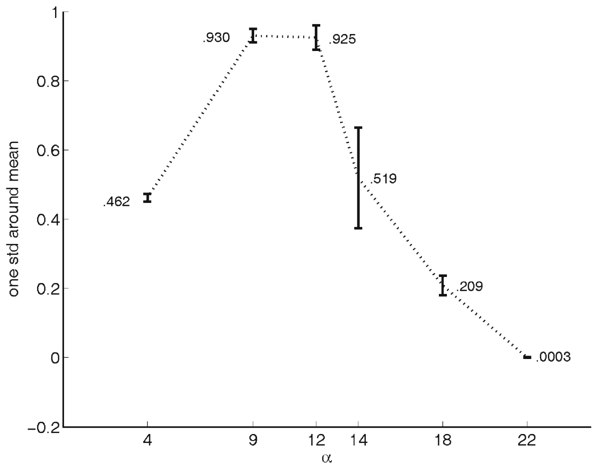

Fig. 7 Effect of α on mean frequency of cooperators

As noted above, we cannot explore the *independent* effects of changes in clustering and characteristic path length on cooperative behavior while holding constant the average degree k of the network. The effects of clustering are fairly intuitive, but the effects of characteristic path length less so. High levels of clustering allow cooperative neighborhoods to form and survive, and hence allow for greater long run levels of cooperation. We expect that a low characteristic path length is also conducive to cooperation, since it allows such behavior to spread rapidly through the network once a cooperative cluster has been formed.

4.3.3 Changes in ε

Figure 8 reveals that an increase in the mutation probability decreases the long-run frequency of cooperation. As ε rises from 0.006 to 0.1, the long run frequency of cooperation falls from 0.585 to 0.071. Mutations randomly select a level of hardness for some individuals, causing neighborhood to become more heterogeneous, and cooperative clusters of sufficiently large size are therefore less likely to form and persist.

4.3.4 Changes in δ

Figure 9 shows that the long run frequency of cooperation initially rises but then declines, reaching its peak at $\delta = .8$. As we argued, if the adjustment rate of hardness is too fast, then the likelihood that large cooperative clusters are formed is small, making it harder for cooperation to persist. However, if

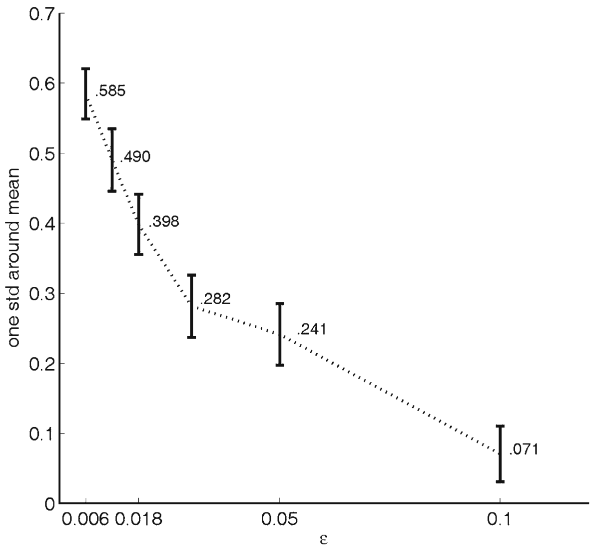

Fig. 8 Effect of ε on mean frequency of cooperators

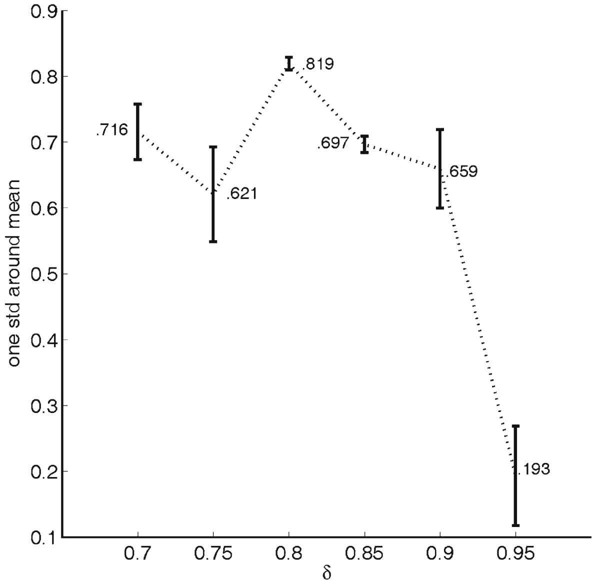

Fig. 9 Effect of δ on mean frequency of cooperators

the adjustment rate of hardness is too slow, then individuals are unwilling to change their actions, which has more detrimental effect on cooperators than defectors, since in this case, sufficiently-large cluster of cooperators must exist initially for them to survive and spread.

4.4 Clustering of strategies

Although significant levels of cooperation can persist with relatively low variance at the aggregate level, there is much greater instability at the local level. Any given node in the network switches back and forth between cooperation and defection, as clusters of cooperative behavior emerge, expand, are punctured and collapse. This dynamism at the local level implies that connected nodes have *correlated* behaviors: if two individuals are neighbors, the likelihood of their choosing different actions is small.

This can be illustrated by examining pairs of individuals who remain connected throughout the process. Since the network is constructed by started with a ring lattice, and the links forming the ring lattice are never severed, each individual has at least two connections that always remain intact. Any pair of always connected nodes must be in one of four possible states at any time: CC, CD, DC, DD. By looking at all pairs of perpetually connected nodes, we can construct a transition matrix that identifies the likelihood of moving from each of these states to each of the others. The following matrix is obtained from a representative simulation:[15]

	CC	CD	DC	DD
CC	0.988	0.006	0.005	0.001
CD	0.063	0.880	0.002	0.055
DC	0.062	0.002	0.882	0.054
DD	0.005	0.031	0.028	0.936

Note that there is a lot of inertia even at the local level, with a high likelihood that a pair will remain in the state inherited from the previous period. This is because behavior is mediated by changes in hardness, which evolves slowly, and only alters behavior once certain thresholds are crossed. Nevertheless, it is clear that pairs which are behaviorally identical (both cooperating or both defecting) are more likely to remain in their inherited states relative to pairs that are behaviorally heterogeneous.

Given the transition matrix, one can compute the limiting proportions of each of the four states in the invariant distribution. For this particular example, we get

$$(CC, CD, DC, DD) = (0.743,\ 0.070,\ 0.062,\ 0.125).$$

This reveals both a high level of cooperation in the aggregate, as well as significant clustering and local homogeneity. An individual who is paired with someone taking a different action is much more likely to change behavior relative to one who is not. This gives us a glimpse into the kind of movements at the local level that sustain high levels of cooperation in the aggregate.

5 Conclusions

We have considered the evolution of strategies in social networks with the small world property, when interaction takes the form of a multi-person prisoners' dilemma. Strategy adjustment is made in response to differential payoffs in one's neighborhood, as is standard in evolutionary games, and is subject to occasional random mutation. The network itself changes over time as links are randomly created and broken. We find that the long-run frequency of cooperation is higher under the following conditions: (i) the interaction radius is neither too small nor too large, (ii) the network is more clustered and the average distance between individuals is shorter, (iii) the mutation rate of hardness is smaller, and (iv) the adjustment rate of hardness is neither too slow nor too fast.

[15]The matrix is based on parameter values $n = 1,000$, $k = 16$, $\alpha = 7$, $\delta = 0.85$, $\varepsilon = 0.005$, $p = 10^{-10}$, $\beta = 4.783$, $\omega = 2$, $T = 100,000$, and $t = 30,000$.

The model could be extended in several directions. First, we assume in this paper that links are formed or broken *exogenously*. In fact, individuals have some degree of autonomy over those with whom they are linked, and may make or break connections based on past behavior. Second, the model could be extended to incorporate the possibility that interaction neighborhoods and learning neighborhoods may be different. Under some circumstances it may be possible to observe behaviors and payoffs outside one's interaction neighborhood, and adjust strategies based on this information. Both these extensions seem promising as directions for future research.

Acknowledgements The first author acknowledges financial support from Kyung Hee University (KHU-20050406). We thank Beom Jun Kim for help with simulations and for comments on an earlier draft, and an anonymous referee for a number of helpful suggestions.

References

Abramson G, Kuperman M (2001) Social games in a social network. Phys Rev E 63:030901-1-4

Albin PS, Foley DK (2001) The co-evolution of cooperation and complexity in a multi-player, local-interaction prisoners' dilemma. Complexity 6:54–63

Axelrod R, Hamilton WD (1981) The evolution of cooperation. Science 211:1390–1996

Barabási AL (2003) Linked: how everything is connected to everything else and what it means. Plume, New York

Bergstrom T, Stark O (1993) How altruism can prevail in an evolutionary environment. Am Econ Rev Pap Proc 83:149–155

Cohen M, Riolo R, Axelrod R (1999) The emergence of social organization in the prisoner's dilemma: how context-preservation and other factors promote cooperation. Working paper 99-01-002. Santa Fe Institute

Eshel IL, Samuelson L, Shaked A (1998) Altruists, egoists, and hooligans in a local interaction model. Am Econ Rev 88:157–179

Fehr E, Gächter S (2000) Fairness and retaliation: the economics of reciprocity. J Econ Perspect 14:159–181

Jackson MO, Wolinsky A (1996) A strategic model of social and economic networks. J Econ Theory 71:44–74

Jun T, Sethi R (2007) Neighborhood structure and the evolution of cooperation. J Evol Econ 17:623–646

Kim BJ, Trusina A, Holme P, Minnhagen P, Chung JS, Choi MY (2002) Dynamic instabilities induced by asymmetric influence: prisoner's dilemma game in small world networks. Phys Rev E 66:021907-1-4

Kirchkamp O (1995) Spatial evolution of automata in the prisoner's dilemma. Discussion paper B-330. Rheinische Freidrich Wilhelms Universität Bonn

Lindgren K, Nordahl MG (1994) Evolutionary dynamics of spatial games. Physica D 75:292–309

Masuda N, Aihara K (2003) Spatial prisoner's dilemma optimally played in small-world networks. Phys Lett A 313:55–61

Nowak MA, May RM (1992) Evolutionary games and spatial chaos. Nature 359:826–829

Nowak MA, May RM (1993) The spatial dilemmas of evolution. Int J Bifurc Chaos 3:35–78

Szabo G, Fath G (2006) Evolutionary games on graphs. arXiv:cond-mat/0607344v2

Szabo G, Vukov J (2004) Cooperation for volunteering and partially random partnerships. Phys Rev E 69:036107-1-7

Trivers RL (1971) The evolution of reciprocal altruism. Q Rev Biol 46:35–57

Watts D (1999) Small worlds. Princeton University Press, Princeton

Watts D, Strogatz SH (1998) Collective dynamics of 'small-world' networks. Nature 393:440–442

Wu Z, Xu X, Huang Z, Wang S, Wang Y (2006) Evolutionary prisoner's dilemma game with dynamic preferential selection. Phys Rev E 74:021107