Beyond Geography: Cooperation with Persistent Links in the Absence of Clustered Neighborhoods

Robert Axelrod
Gerald R. Ford School of Public Policy
University of Michigan

Rick L. Riolo
Center for the Study of Complex Systems
University of Michigan

Michael D. Cohen
School of Information
University of Michigan

Electronic communication allows interactions to take place over great distances. We build an agent-based model to explore whether networks that do not rely on geographic proximity can support cooperation as well as local interactions can. Adaptive agents play a four-move Prisoner's Dilemma game, where an agent's strategy specifies the probability of cooperating on the first move, and the probability of cooperating contingent on the partner's previous choice. After playing with four others, an agent adjusts its strategy so that more successful strategies are better represented in the succeeding round. The surprising result is that if the pattern of interactions is selected at random, but is persistent over time, cooperation emerges just as strongly as it does when interactions are geographically local. This has implications for both research on social dynamics, and for the prospects for building social capital in the modern age.

For over a century, the broadening patterns of interaction among people have raised concern that modernity would destroy the basis of community (Tonnies, 1887/1957). Recently, the ability of electronic communication to support distant interactions that fragment solidarities and de-emphasize local patterns of interaction (Wellman et al.,1996) has heightened this concern. One basis of concern is that the patterns of interaction in physical space are densely clustered, as in a neighborhood, where each individual tends to interact with a large proportion of the other individuals in the neighborhood; this clustering may be lost in other patterns of interaction. It is well known, however, that having a few distant connections can foster the diffusion of valuable information in a social network (Milgram, 1967; Watts, 1999; Watts & Stoats, 1998). What is not known is whether social networks that are not densely clustered (i.e., in which there is little overlap among the interacting individuals), can still foster cooperation among ego-

ists (individuals motivated by rational self-interest) based upon reciprocity. In this article, we use a computer simulation model to examine the question of whether social networks with distant links can perform as well as networks based on local geography. This model portrays individuals (agents), represented by nodes in a spatially distributed network, and the interactions among agents, represented by reciprocal links among nodes.

Our work builds on the work of Holland (1992, 1995), who has shown that in a complex adaptive system[1] the pattern of interactions among the agents can have a strong effect not only on the success of individual agents, but also on the performance of the system as a whole. Holland has addressed this theme in many contexts over the years, most often by exploring the role of labels ("tags") that may be used to identify agents as desirable or undesirable interaction partners. Tags can thus establish patterns of interactions among

---

[1]The notion of complex adaptive system is somewhat different from the notion of nonlinear dynamical system. A nonlinear dynamical system refers to a system that evolves in time, and in which interactions among elements are nonlinear. A complex adaptive system is a system composed of many elements, in which the elements adapt to each other and the environment either through learning or evolution. Although most complex adaptive systems are nonlinear dynamical systems, nonlinear dynamical systems do not necessarily learn or evolve and are not necessarily composed of many elements.

agents by allowing for the selection of interaction partners (Holland, 1995, 1998; Holland, Holyoak, Nisbett, & Thagard, 1986). Spatial rules for the selection of interaction partners have also been investigated in this line of work. It is usually assumed that agents are biased toward interaction with other nearby agents (Holland, 1992, 1995). Local interaction rules, in turn, have been shown to promote cooperation (Riolo, Cohen, & Axelrod, 2001; Sigmund & Nowak, 2001). This is because local interactions tend to produce clusters of cooperating individuals, who are less exposed to exploitative agents who might otherwise take advantage of the tendency to cooperate.

In this article, we show that even random networks can establish and sustain cooperation, provided that the network of interactions is stable over time. To appreciate the difference between a random network and one based on geography, consider the contrast between two pure types. In both networks, each person has exactly four "neighbors." First, let us consider a network based on geography, such as the people in a small town. In this two-dimensional lattice, the neighborhoods are correlated (i.e., systematically overlapping). For example, a person has two neighbors in common with whomever is one step north and one step east. The correlational structure of such a geographic network means the number of people at a certain distance from any given person does not grow very quickly and that there is a large overlap in the people who interact with each other. In fact, the number of others who are exactly d steps away form a diamond in this lattice network, and their number increases linearly: $N(d) = 4d$.

In contrast, consider a network in which every person is connected to four others chosen at random from a large population, as might be the case in a large corporation that relies heavily on electronic communications. In this case, the network fans out in a tree structure. A person has four immediate neighbors (in terms of communication links, but not necessarily spatially), each of whom has three other neighbors, and so on. For a random network in an infinite population, the number of others at a given network distance expands exponentially as a function of distance: $N(d) = 4*3^{d-1}$. Thus, for a random network in a large finite population, two people may share more than one neighbor, but this will be rare. Rapid fan out of links in a random network can help the diffusion of information. On the other hand, rapid fan out as well as the lack of neighborhood clustering raises questions about the prospects for prosocial behavior in an uncorrelated structure.

The standard formal paradigm for analyzing the problem of prosocial behavior is the Prisoner's Dilemma Game (PDG; Axelrod, 1984; Axelrod & Hamilton, 1981). In the PDG, each player has two choices: cooperate or defect. In any given move, the players each receive R points if both cooperate, and

only P points if both defect. A defector exploiting a cooperator gets T points, while the cooperator gets S (the relationships among the various payoffs in this game are: T > R > P > S and 2R > T + S). This pattern of payoffs insures that in a single move, it is always better for a player motivated by self-interest to defect. Since this save is true for both players, the unfortunate result is mutual defection. Agents can maximize their personal gains, therefore, by suspending their immediate self-interest in favor of smaller but systematic gains achieved through sustained mutual cooperation. The agents' ability to discover this rule has been demonstrated in numerous computer simulations, employing both evolutionary mechanisms and learning mechanisms such as imitation (e.g., Axelrod, 1984; Axelrod & Dion, 1988; Axelrod & Hamilton, 1981; Messick & Liebrand, 1995).

One of the most rebust strategies in the PDG is "tit for tat," in which the individual starts with cooperative behavior, and on each succeeding interaction simply imitates the previous behavior of his or her interaction partner. Although this strategy does not allow the individual to win in direct play with any particular partner, in the long run it produces good outcomes because it elicits and sustains cooperation with cooperative agents, while punishing exploitative agents with defection. More generally, although it never pays to cooperate in a single interaction, the shadow of the future allows cooperation based upon reciprocity to emerge and be sustained in a population of egoists.

In psychology, the PDG has inspired several lines of research into the factors that foster versus inhibit cooperative behavior among individuals or groups (cf. Pruitt, 1998). Thibaut and Kelley (1959), for example, investigated how variations in the payoff matrix for interacting agents affected their cooperative versus competitive tendencies. This approach was subsequently generalized to an *N*-person dilemma (Messick & Brewer, 1983), such as resource depletion in the "tragedy of the commons" (e.g., Hardin, 1968). Individual behavior and preferences in the PDG, meanwhile, have been utilized as diagnostic tools for social orientations, such as individualism, altruism, cooperation, and competition (e.g., McClintock & Liebrand, 1988). Yet other lines of research have examined a wide panoply of psychological and social factors, including ingroup-outgroup bias, the prospect of future interaction, individual versus group interaction, group size, communication structure, and norms (cf. Pruitt, 1998; Schopler & Insko, 1992).

These research programs have generated empirical generalizations that are consistent with the possibility that cooperation can emerge among spatially distributed individuals who have sustained interactions. Direct evidence on this point, however, has yet to be obtained. The simulations we report in this article are intended to provide such evidence.

Table 1. *Cooperation in Three Kinds of Networks*

Neighborhood structure	Performance (± s.d.)	Instability (± s.d.)	No. Successful Runs (of 30)	Time to Threshold	Retention
Temporary Random	1.098 ± .007	0.031 ± .026	1	1923	.000
Two Dimensional	2.557 ± .009	0.074 ± .006	30	24	.997
Persistent Random	2.575 ± .010	0.078 ± .006	30	26	.995

Note: Performance is the average score of the 30 runs of each type. It is calculated only over the final 1000 periods to avoid the effects of the initial conditions. *Instability* is the variation over time within a single run, measured by the standard deviation of population average scores within a run over the final 1000 periods. *No. Successful Runs* is the number of runs out of 30 that ever attained an average population score of 2.3. Time to threshold is the average number of periods until the threshold is reached. Data for this and the next column are calculated only for the runs, which succeeded in reaching the threshold. *Retention* is the proportion of periods after the threshold is first reached in which the population average is above the threshold.

Simulations

In our simulations, a period consists of each of 256 agents selecting each of its four neighbors to play a game of exactly four moves each. The neighbors also select the agent. This symmetry implies that everyone plays exactly eight games in a period. To make cooperation relatively difficult to achieve, we use games of only four moves each. An agent's strategy is updated after each period based on its experience in comparison to its four neighbors' experience (see the following). A run consists of 2500 periods. We are interested in the effects of how neighborhoods are structured. Our three neighborhood structures are:

Temporary random network: The four neighbors are chosen at random at the start of each period.[2]

Two-dimensional network: At the start of each run, one agent is assigned to each location in a 16 × 16 lattice. Neighborhoods are the four adjacent sites of each location. In this case the neighborhoods are correlated.

Persistent random network: This is the same as the temporary random network, except that the network is retained for an entire run of 2500 periods rather than being shuffled after each period. Like the two-dimensional case, the neighborhoods are persistent. But unlike the two-dimensional case, there is no correlational structure to the neighborhoods.

We have also run experiments with a number of other neighborhood structures, including neighborhoods induced when agents select who to play based on arbitrary tags, using tagging-mechanisms similar to those described in Holland (1995); for details, see Cohen, Riolo, & Axelrod (1999, 2001).

The stochastic strategies used by agents are given by three parameters (y, p, q), where y is the probability of cooperating on first move, and p and q are the conditional probabilities to co-operate, given that the other player's previous move was a C (cooperation) or D (defection) (Nowak & Sigmund, 1989). This class contains, for example, Tit for Tat (1,1,0), ALLC (1,1,1)

and ALLD (0,0,0) as extremal representatives. For our runs, we use standard payoff values of R = 3, P = 1, T = 5 and S = 0.

The initial population of agents is spread evenly across the (p, q) space, one agent at each of the 256 combinations of equally spaced levels of p (p = 1/32, 3/32 … 31/32) and q (q = 1/32, 3/32 … 31/32). The y value is initially set to the p value.

The basic learning mechanism is imitation. Specifically, at the end of each period, each agent identifies the most successful of its four neighbors in the current period. It then adopts that neighbor's strategy if that neighbor did strictly better than itself in the current period. To be realistic, the update rule takes account of two sources of possible error. First, there is a ten-percent chance the agent makes a mistake when deciding whether or not its best neighbor did better than itself. Second, regardless of which of the two strategies is adopted, for each of the three parameters there is a ten-percent chance that Gaussian noise is added (M = 0, SD = 0.4, bounded by 0 and 1). Adding noise in the update rule allows new strategies to be introduced into the population, thereby increasing the potential for adaptation.

Results

The results shown in Table 1 are based on 30 replications for each condition. The average score for a population can vary between P = 1 and R = 3. Since the agents use stochastic strategies and make errors of comparison and copying when updating their strategies, a convenient threshold for regarding a population as successful is a mean score of 2.3 points.[3] As expected, populations with temporary random neighborhoods do hardly better than mutual defection, have little instability within runs, only rarely reach the threshold of success, and are not able to sustain cooperation even then. In contrast, the populations in a two-dimensional setting achieve and sustain high levels of cooperation, do so with little instability, always attain the threshold of success, attain the threshold of success very quickly, and consistently stay above the threshold once attained. Most important, the populations with persistent random networks perform

[2]Our algorithm to maintain symmetry in the random neighborhoods is the following. After all agents have four symmetric neighbors, each agent randomly swaps a neighbor with 256 randomly chosen other agents, with each swap subject to the constraint that each agent have four unique neighbors (none of whom is itself) and each neighbor relation is symmetric.

[3]This sets the threshold of success higher than 2.25, the expectation of a random initial population.

343

just as well as populations in the geographic setting of a two-dimensional network.

Another way to compare the three conditions is to look at the average performance over time in typical runs. This comparison is shown in the three panels of Figure 1. In all three conditions, the population average begins at 2.25 points because the initial conditions provide a random set of strategies, resulting in all four outcomes being attained with equal likelihood. In the temporary random neighbor case (Fig. 1a), the population average quickly drops, reflecting an evolution toward mutual punishment. While the population average occasionally increases for short stretches as clusters of cooperation occur, these clusters cannot be sustained when the neighborhoods are constantly reshuffled. In contrast, both the two-dimensional case (Fig. 1b) and the persistent random neighborhood case (Fig. 1c) quickly attain high scores as cooperation based on reciprocity becomes established. Although there is some fluctuation due to mutation in these two cases, the high levels of cooperation are never lost.

Figure 2 provides a way to visualize the forces producing population histories such as those shown in Fig-
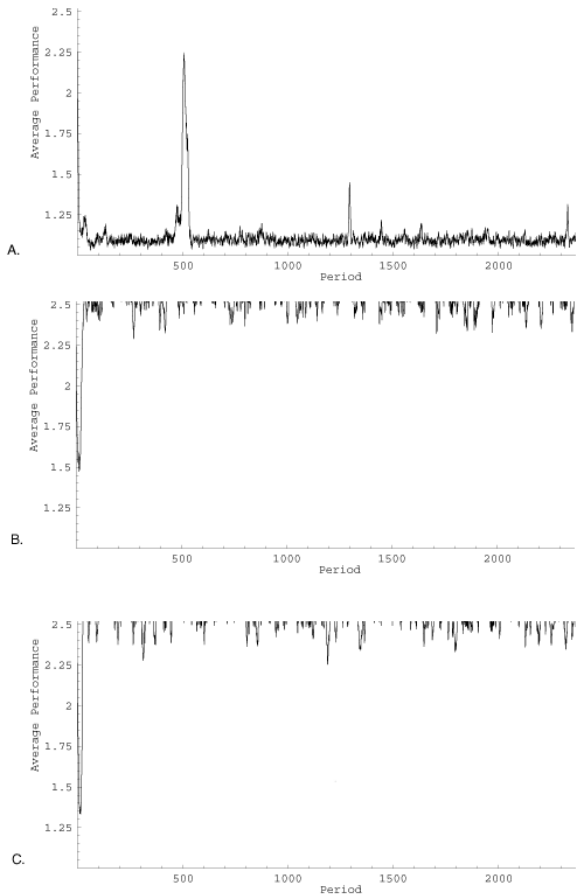

Figure 1. Average Performance Over Time of a Typical Run for Each of Three Neighborhood Structures
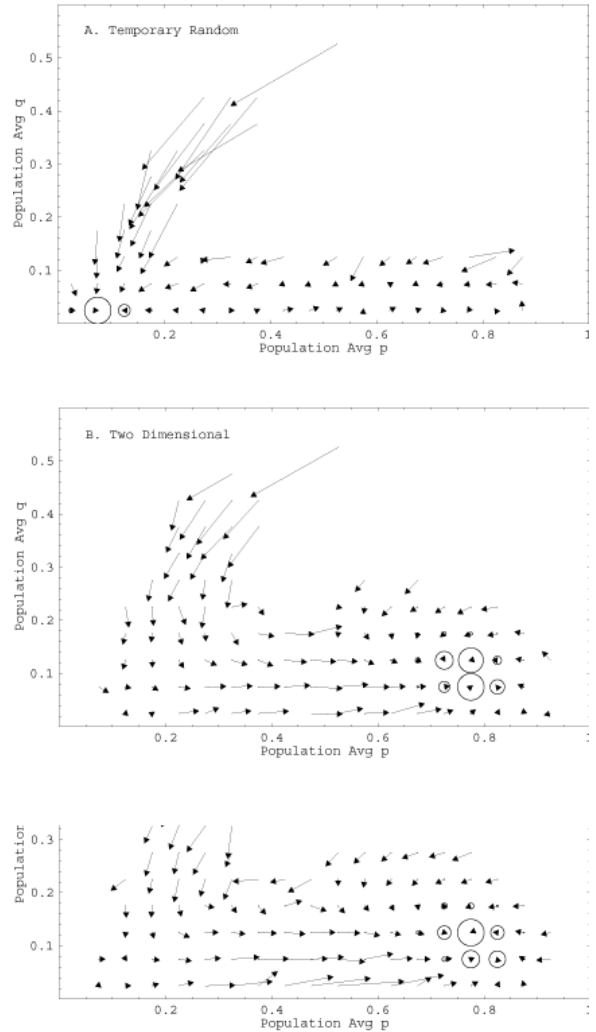

Figure 2. A (p, q) Phase Plot for Three Neighborhood Structures. The p value is the probability of cooperating after the other player cooperated, and q is the probability of cooperating after the other player defected. The arrows indicate the direction of the change in population average values of p and q one period after the population is in a given pq bin. The size of a circle indicates the number of periods in which the population remains in that bin. The largest circle in each panel corresponds to the cell in which the population spent the most time. The areas of other circles are scaled relative to that cell.

ure 1. For each neighborhood structure, it shows a diagram of movement in the (p, q) space as an indication of how the strategies change over time. For each small region (there are 20×20 "bins") of the space in the diagram, we have determined all the periods in which the population average values of p and q are in that region.[4] For each such bin, we have found the average of the p and q values in the immediately succeeding period. The arrow originating from the center of each region

---

[4]We are ignoring y in this analysis because for the most part y is correlated with p.

shows the average one period movement of p and q for populations that were ever in that region. Each panel uses data from all 30 replications of the respective condition. In all conditions, both p and q initially fall. This is a trace of the collapse from the initial state of the population, in which the p and q values of the agents are evenly distributed across the entire (p, q) space. In that initial environment, agents with strategies closer to ALLD (p = 0 and q = 0) do better than more cooperative strategies.

With temporary random neighborhoods (Fig. 2a), the population spends most of its time near the Always Defect corner, as indicated by the large circles in the corresponding cells. Occasional increases in p (the probability of cooperating after the other player cooperated) are quickly reversed. In both two dimensional and persistent random cases, however, the population quickly moves out of the Always Defect corner and sustains cooperation with strategies akin to Tit for Tat (high p and low q). Moreover, the arrows indicate that both of these conditions develop counterclockwise movement when they depart from the region near the Tit for Tat corner: an increase in q (toward unconditional cooperation) results in a decrease of p (as the more gullible strategies are exploited), a decline in q (as the gullibles are eliminated), and a return to the more stable region of high p and low q (near the Tit for Tat corner). Thus the persistent random network setting is very much like the two-dimensional setting and very different from the temporary random setting not only in terms of its overall performance (Table 1), but also in the dynamics by which that performance is achieved (Figures 1 and 2).

Discussion

The key finding is that a persistent random network displays virtually the same pattern of success as a two-dimensional setting. This result establishes the importance of distinguishing the multiple characteristics of non-geographic interaction patterns. Clustering (i.e., the correlation of linkage patterns) is not needed to establish and sustain cooperation over networks, provided that the linkages are stable. The ability of stable neighborhood structures to support cooperation need not be lost in the geographically dispersed social networks that are now arising all around us.

Our results have both methodological and substantive implications. A methodological implication is that the dynamics of the pattern of interactions can have a powerful effect on the behavior of a system. This is especially important for game theory, which is a leading theoretical tool in the social sciences. Game theory has made great progress in the analysis of learning and adaptation (Fudenberg & Levine, 1998), but has paid little attention to who interacts with whom. Our simula-

tions show that these details can be decisive. In particular, whether a random network can sustain cooperation depends on whether the pattern of interactions is persistent or changing. This suggests that the continuing search in game theory for improved equilibrium concepts can be usefully augmented by a study of the dynamics of the underlying interaction structures (Cohen et al., 2001).

Our finding that a persistent random network can support cooperation as well as a geographic network can has implications for the development of social capital, which Putnam (2000) has defined in terms of the "connections among individuals – social networks and the norms of reciprocity and trustworthiness that arise from them" (p. 19). High levels of social capital in a society depend on a dense network of reciprocal social relationships. Social capital embodied in norms and networks of civic engagement supports such things as successful education, effective government, economic development, and a generally healthy civil society (Putnam, 1993a, 1993b).

Putnam (2000), among others, has recently argued that there has been a serious decline in the amount of social capital in America, with an attendant decline in the health of important social institutions. Now that social, political, and economic networks extend far beyond local geographic constraints, the question arises whether social capital can be built and sustained under these conditions. For example, numerous commentators have questioned whether the growth of the Internet and electronic communication is weakening networks of social relationships. Our research shows that the answer to these questions lies, at least partially, in whether the new networks are persistent or transient. To the extent that new networks of distant interactions tend to be stable, social capital can flourish.

References

Axelrod, R. (1984). *Evolution of cooperation*. New York: Basic Books.

Axelrod, R., & Dion, D. (1988). The further evolution of cooperation. *Science, 242*, 1385–1390.

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science, 211*, 1390–1396.

Cohen, M. D., Riolo, R. L., & Axelrod, R. (1999). *The emergence of social organization in the prisoner's dilemma: How context preservation and other factors promote cooperation*. (Santa Fe Institute Technical Report No. 99–01–002).

Cohen, M. D., Riolo, R. L., & Axelrod, R. (2001). The role of social structure in the maintenance of cooperative regimes. *Rationality and Society, 13*, 5–32.

Fudenberg, D., & Levine, D. (1998). *Theory of learning in games*. Cambridge, MA: MIT Press.

Hardin, G. (1968). The tragedy of the commons. *Science, 162*, 1243–1248.

Holland, J. H. (1992). *Adaptation in natural and artificial systems*. Cambridge, MA: MIT Press.

Holland, J. H. (1995). *Hidden order: How adaptation builds complexity.* Reading, MA: Addison Wesley.

Holland, J. H. (1998). *Emergence: From chaos to order.* Reading, MA: Addison Wesley.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. (1986). *Induction: Processes of inference, learning and discovery.* Cambridge, MA: MIT Press.

Messick, D. M., & Brewer, M. B. (1983). Solving social dilemmas: A review. In L. Wheeler & P. Shaver (Eds.), *Review of personality and social psychology* (Vol. 4, pp. 11–44). Beverly Hills, CA: Sage.

Messick, D. M., & Liebrand, V. B. G. (1995). Individual heuristics and the dynamics of cooperation in large groups. *Psychological Review, 102*, 131–145.

McClintock, C. G., & Liebrand, W. B. G. (1988). Role of interdependence structure, individual value orientation, and another's strategy in social decision making: A transformational analysis. *Journal of Personality and Social Psychology, 55*, 396–409.

Milgram, S. (1967, May). The small world problem. *Psychology Today, 1, 62–67.*

Nowak, M. A., & Sigmund, K. (1989). Oscillations in the evolution of reciprocity, *Journal of Theoretical Biology, 137*, 21–26.

Pruitt, D. G. (1998). Social conflict. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology, 4th edition* (Vol. 2, pp. 470–573). Boston: McGraw-Hill.

Putnam, R. D. (1993a.) *Making democracy work: Civic traditions in modern Italy.* Princeton, NJ: Princeton University Press.

Putnam, R. D. (1993b). Social capital and public affairs. *The American Prospect, 13,* 1–8.

Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community.* New York: Simon & Schuster.

Riolo, R., Cohen, M. D., & Axelrod, R. (2001). Evolution of cooperation without reciprocity. *Nature, 414*, 441–443.

Schopler, J., & Insko, C. A. (1992). The discontinuity effect in interpersonal and intergroup relations: Generality and mediation. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 3, pp. 121–151). Chichester, England: Wiley.

Sigmund, K., & Nowak, M. A. (2001). Tides of tolerance. *Nature, 414*, 403–404.

Thibaut, J. W., & Kelley, H. H. (1959). *The social psychology of groups.* New York: Wiley.

Tonnies, F. (1957). *Gemeinnshaft und Gesellshaft* [Community and Society]. (C. P. Lewis, Ed. & Trans). East Lansing: Michigan State University Press. (Original work published 1887)

Watts, D. J. (1999). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology, 105*, 493–527.

Watts, D. J., & Stoats, S. H. (1998). Collective dynamics of "small-world" networks. *Nature, 393*, 440–442.

Wellman, B., et al. (1996). Computer networks as social networks. *Annual Review of Sociology, 22*, 213–238.