

THE ROLE OF SOCIAL STRUCTURE IN THE MAINTENANCE OF COOPERATIVE REGIMES

Michael D. Cohen, Rick L. Riolo, and Robert Axelrod

ABSTRACT

We analyze the role of social structure in maintaining cooperation within a population of adaptive agents for whom cooperative behavior may be costly in the short run. We use the example of a collection of agents playing pairwise Prisoner's Dilemma. We call sustained cooperative behavior in such circumstances a 'cooperative regime'. We show that social structure, by channeling which agents interact with which others, can sustain cooperative regimes against forces that frequently dissolve them. We show in detail the process through which structured interaction in a population creates a 'shadow of the *adaptive* future', allowing even a small set of cooperative strategies to grow into a cooperative regime, a coherent, self-sustaining entity that is something more than the sum of the pairwise interactions among its members.

KEY WORDS • adaptation • cooperation • emergence • Giddens
• Prisoner's Dilemma • social networks

1. Introduction

In situations with Prisoner's Dilemma logic, where collectively beneficial actions are costly to individuals in the short run, theories based on assumptions of rational agents predict that cooperation should be difficult to sustain. At the same time, sustained voluntary cooperation is a recurring feature of real social systems. Cooperative action may have results that we admire, as when team members make sacrifices on behalf of their colleagues, or it may have results that we dislike, as when firms collude to fix prices in spite of competitive forces. But in either case, the existence of voluntary cooperation is not in doubt.

Cooperation is not universal, of course, and a major strength of rational agent theories is the explanation they provide for the difficulty of sustaining cooperation in a Prisoner's Dilemma situation. Nonetheless, there is a disjuncture between the outcomes expected under theories of rational individual action and the cooperative actions that are frequently observed in the world, and the disjuncture poses a serious challenge. How should we account for the cases where a durable 'cooperative regime' does form?

A principal response to this challenge has been to invoke the 'folk theorem' (e.g. Fudenberg and Maskin 1990), of which Axelrod's 'shadow of the future' (Axelrod 1984) can be seen as a special case. If one rational agent expects its present actions to affect the future behavior of the other, then a less valued outcome in the short run might be justified in view of a long run with increased expected value.

However, cooperation in the real world occurs in many situations that do not meet the conditions for an explanation based upon the shadow of the future. This approach assumes that the two interacting agents have strategies that remain constant over the interval of iterated play. But we know, for example, that trustworthy practices are maintained in securities trading pits even though strategies are changing rapidly during 'play' (Baker et al. 1984). The shadow of the future approach also assumes that the interaction is between agents who remain the same. But we know that norms are sustained over long periods of time in larger systems such as universities and legislatures, even though individual members may turn over rapidly and many shirk the costs of enforcing norms against violators (Axelrod 1986).

For worlds where interaction sequences may be short, where agents may be rapidly changing their strategies, or where agents may be turning over, the conventional shadow of the future approach does not suffice. How then do we account for the emergence in such settings of 'cooperative regimes', durable high levels of cooperative action within large populations of agents?

Our answer focuses on social structure, which determines the patterned character of interaction among agents or, 'who tends to interact with whom?' Though our approach is via formalized models, the work is much in the spirit of Giddens's structuration theory (1984). He defines social structure as 'the *patterning of interaction*, as implying relations between actors or groups, and the *continuity of interaction* in time' (Giddens 1979). Our simulations show

how the micro-level actions taken by local, changing, and self-interested actors can build and maintain a cooperative regime, which in turn stabilizes their actions and provides the conditions – Giddens might say ‘resources’ – that their strategies require (Kluever and Schmidt 1999). The continuity of interaction patterns, which we label as ‘context preservation’, plays a crucial role in fostering and sustaining cooperative regimes by altering the adaptive dynamics of agent populations.

Our argument proceeds as follows. We describe our framework for controlled experimentation with simulated populations of agents playing a pairwise iterated Prisoner’s Dilemma. We compare several sets of studies completed within this simulation environment to demonstrate the contribution of context-preserving social structure to the emergence of cooperative regimes. We present a micro-level analysis of the process by which cooperation does, or does not, emerge. It shows the central role played by social structure in creating a ‘shadow of the adaptive future’. We use this label to distinguish the effects of future interactions on the survival chances of variant strategies being generated by the adaptive processes of the present. A favorable shadow of the adaptive future implies conditions that amplify the spread of cooperative strategies. The paper concludes with a discussion of the potential implications of our results for concrete issues, such as the maintenance of trust in the Internet.

2. Experimental Framework

In our simulations a period consists of each of 256 agents¹ selecting four others to play iterated Prisoner’s Dilemma. We use the standard payoff matrix shown in Table 1. Pairs of agents play a game of length four, short enough to make cooperation difficult to achieve, but still possible (Riolo 1997a). An agent’s strategy is updated at the close of each period based on comparison with the other agents it encountered, according to rules which are described below.

Each agent’s strategy is represented as in Nowak and Sigmund (1989), by a triplet of real numbers $0 \leq y, p, q \leq 1$. The first of these, y , represents the probability that the agent will cooperate on the first move of a game. The second, p , is the probability that an agent will cooperate following a cooperative move by the agent with which it is playing. The last, q , is the probability that the agent will cooperate

Table 1. Prisoner's Dilemma Payoffs

	<i>Cooperate</i>	<i>Defect</i>
<i>Cooperate</i>	3, 3	0, 5
<i>Defect</i>	5, 0	1, 1

The first value is the row player's payoff, the second is for the column player.

after the other agent has defected. We initialize y as equal to p , but they may vary independently thereafter.² This large space of possible strategies contains Tit-for-Tat (TFT: 1,1,0), Always Cooperate (ALLC: 1,1,1), and Always Defect (ALLD: 0,0,0).

Our interest is in adaptive settings, where an agent's strategy may change over time. We have studied several processes of adaptation but, for simplicity, report only one here. At the close of each period, an agent takes as its own strategy for the next period the strategy of that agent among those it met whose average score per move was best – provided that score was strictly greater than the agent's own.

Processes of emulation are subject to errors, of course. An agent may inaccurately perceive the strategy of another it wants to copy, or may not judge correctly how well another's strategy has performed. We reflect these errors of emulation in our simulation by introducing two kinds of 'noise'. Agents reach the wrong conclusion 10 percent of the time in comparing their own recent per move score with that of the best performing agent they've encountered. This results in their occasionally copying an inferior strategy or failing to copy a superior one. Agents also make errors in the actual copying process. For each of the three variables, there is a 10 percent chance that during copying the variable value will be disturbed by Gaussian noise of mean 0 and standard deviation 0.4, with truncation of values falling above 1 or below 0.

In the studies reported below we systematically vary the social structure that controls the meetings among the agents. Iterative play occurs on a quick time scale within each period. Adaptation occurs on a slower time scale, over a succession of periods as the inter-agent emulation process spreads strategies that are succeeding and extinguishes those that are not. A complete description of the experimental procedures used to produce the results presented in this paper is given in the Appendix.³

3. Results

3.1. Cooperation Levels in Alternative Social Structures

We ran 30 replications of each experimental condition, each starting with a different random number seed. Thus each replication generates a distinct ‘population history’ for the given social structure. The first row of Table 2 reports a study in which each agent played with four other agents chosen randomly in each period. (Thus, on average, an agent played eight others, exactly four as the chooser of partners, and, on average, four as the chosen other.) A population played for 2500 periods, starting from strategies that were initialized randomly.

We have devised many measures on these data, but report only three here:

1. the proportion of 30 populations that achieved a high level of cooperation at any time in their history (‘Attain High C’);
 2. the average score attained by all the populations over the last 1000 periods of their histories (after the effects of random initialization have long died away) (‘Mean Payoff’);
- and

Table 2. Cooperation for Different Social Interaction Conditions

<i>Social structure</i>	<i>Attain high C</i>	<i>Mean payoff</i>	<i>Remain high</i>
RWR	0.30	1.091	0.015
2DK	1.00	2.557	0.997
FRNE	1.00	2.575	0.995
FRN	1.00	2.480	0.942
FFR 0.1	1.00	2.385	0.844
FFR 0.3	1.00	2.100	0.402
FFR 0.5	0.93	1.257	0.061

RWR is ‘Random With Replacement’ each time step; 2DK = ‘2 Dimensions, Keeping’ neighbors for the entire run; FRNE is ‘Fixed Random Neighbors, Equal number and symmetrical; FRN is ‘Fixed Random Neighbors’, not equal numbers nor symmetrical; FFR 0.1 is ‘FRN with Fraction RWR = 0.1’, and similarly for the other FFR cases. ‘Attain High C’ is the proportion of histories (runs) that achieved high levels of cooperation. ‘Mean Payoff’ is the average score per move attained by all populations over the last 1000 steps of all runs. ‘Remain High’ is the proportion of time the populations remain at high levels of cooperation once they achieve it.

3. the proportion of time spent in conditions of high cooperation among those populations that achieved it ('Remain High').

As the first row in Table 2 shows, high levels of cooperation were attained in only a few of the populations where contact patterns were randomized each period (the case we label 'RWR' to indicate that agents play others chosen randomly with replacement); and even when a high cooperation level occurred, it was only momentary. Among those populations that experienced high cooperation, less than 2 percent of their subsequent time was spent in that state. As a result, the mean payoff was dismal. It hovered just above the theoretical minimum mostly because of accidental mutual cooperations stemming from the noisy character of both emulation and strategy execution. This is a Hobbesian war of all against all. Sustainable cooperative regimes did not emerge. Brief bursts of cooperation did occur among the adapting agents, but they were evanescent fads that died out immediately.

By contrast, consider the second row of Table 2. It reports the result of an equivalent simulation except for a change in the 'social structure' governing agent interaction. Agents were assigned permanent positions in a 16×16 two-dimensional toroidal lattice. Each agent chooses to play a game each period with each of its four neighbors who are 'kept' through the entire history. (The case is therefore labeled '2DK'.) Thus each agent played exactly eight times: four as chooser and four when it is chosen by each of its neighbors.

In this case all 30 populations attained high cooperation levels, and, once that happened, the cooperative regime was very successfully maintained. Over 99 percent of the subsequent experience was of widespread cooperative action. The resulting average score is quite high, near the maximum possible given the noise in strategy emulation and execution.

The result that two-dimensional spatial embedding favors cooperation in Prisoner's Dilemma is not new. Indeed, it might be called robust. It has been reported by a number of investigators across settings that vary substantially in their details (Axelrod 1984; Nowak and May 1992; Lomborg 1996). However, we can – and should – go beyond attributing cooperative regimes to spatial embedding of agents. It is necessary to ask, 'What are the fundamental properties of spatial embedding that so dramatically alter dynamic processes within these populations?' Unless we pursue

this deeper question, we will not even be able to judge whether the observed results would be likely to hold for real actors in real physical spaces (which, after all, may deviate substantially from toroidal lattices). Beyond that, we will be unable to infer the relevance of the results to actors located in other, more abstract spaces, such as social networks and organizational hierarchies.

Pursuing this question reveals that our randomly paired and our spatially embedded populations differ in two important properties, not one. First, the imposition of fixed spatial location means that agent pairings do not change – even though the adapting agents may change their strategy after any period. Second, it means that the contact networks of agents are correlated. Paired agents have neighbors who are themselves paired. This feature of the interactions, called ‘clustering’ in the graph theory literature (Watts 1999), will rarely occur by chance in random pairing.⁴ Which of these two properties is responsible for the difference we have observed? Does it occur because clustering fosters the diffusion of cooperative strategies, or because a fixed interaction pattern is favorable?

Our experimental framework allows us to answer this question very cleanly. In a third experimental condition, we paired agents randomly and left those initial pairings fixed for the full duration of the population history. We also forced the pairings to be symmetric in order to reproduce that further characteristic of two-dimensional lattice play. (We label this fixed random network with equal number of other players as ‘FRNE.’) Thus our simulated agents are in a social structure that has the same fixity of interaction but none of the clustering of the previous case. In the third row of Table 2 it can be seen that the results for fixed random pairing are nearly indistinguishable from the two-dimensional study.⁵

Before we can conclude that it is context preservation (the persistence of interaction pattern) rather than the clustering (correlated interaction profiles) that lies behind the observed effect of spatial embedding, there is a very subtle alternative explanation that we need to rule out. Our original cases of random mixing and spatial embedding actually differed in a third way. Spatially embedded agents played exactly eight games per period with exactly four others. Random pairing matches an agent into eight games on average, but not eight games exactly. There can be variance since an agent will be randomly chosen four times by others only on average. Also, the games are typically with eight distinct others, rather than two games each against exactly four others. To rule out

these subtleties as the sources of our differences, we used an initial random pairing that remained fixed and did not retain the symmetry of play from the two-dimensional case (FRN). The fourth row of Table 2 shows that this makes only a minor difference that does not alter the substantive conclusions. When initial pairings are determined randomly without symmetry but are left fixed, high levels of cooperation are again attained by all populations and are maintained consistently to produce high average scores. Context-preserving social structure suffices for the emergence and maintenance of cooperative regimes, without paired agents having correlated networks.

The point established should be stated with care. We do not contend that clustering could never contribute to the maintenance of cooperative regimes. (Indeed, we have run some other experiments not analyzed here in which small positive effects for clustering are visible.) Rather, our point is that simple context preservation is a powerful factor, sufficient to sustain a cooperative regime against short shadows of the future and rapid adaptive processes, factors that overwhelm cooperation in a world of random mixing (and, equivalently, in a tournament world where all agents meet all others with equal frequency).

The last three lines of Table 2 deepen this point. We report another experiment, which takes as its base the fixed random network (FRN) case just discussed. In the fifth row of Table 2 we show results based on 30 populations in which agents played with fixed partners chosen initially at random, as in the FRN case, but in each period a new randomly chosen partner was substituted (temporarily, as in the RWR case) for each fixed network partner with probability 0.1. (We use the label 'FFR-0.1' for this case of Fixed random network with Fraction Random with replacement of 0.1.) This addition of minor 'noise' to the basically fixed pairing of partners makes only a slight difference, so the fifth row of Table 2 strongly resembles the preceding two cases. However, this experimental structure provides us with a parameter that allows us to smoothly vary the context-preserving property of the social structure. At a substitution probability of zero, the model is identical to the FRN case. At a probability of 1.0, it is identical to pure RWR. Every pairing is chosen at random every period. The sixth and seventh rows of Table 2 show the results for the intermediate substitution probabilities of 0.3 and 0.5. The first of these variants causes a marked departure from the FRN dynamics. The second

increase in the substitution probability is sufficient to produce histories similar to the pure RWR case. As we ‘turn the dial’, random substitution of new partners into the fixed network increases, and therefore context preservation declines. Around a parameter value of 0.3 the dynamics shift. The ability to sustain cooperative regimes weakens, and then, at levels of 0.5 and above, it collapses. These results for sets of population histories (i.e. individual runs of the model) strongly suggest that context-preservation plays an essential role in the emergence and maintenance of cooperative regimes.

3.2. Detailed Investigation of Processes Leading to a Cooperative Regime

To understand this fundamental pattern, we trace through qualitatively similar dynamics that are found in the detailed events within our simulated histories. We show that context-preserving social structure reshapes this basic pattern into a cooperative regime by maintaining interaction patterns among agents as their strategies adapt. In effect, this creates a *shadow of the adaptive future*. This concept, distinct from the familiar shadow of the future, is defined on the time scale of adaptation rather than that of iterative play. Games here are only four moves long, too short for the effective operation of the traditional shadow notion. What drives the shadow of the *adaptive future* is the benefit derived from keeping together agents whose strategies developed through emulating each other. Since there is no best strategy in Prisoner’s Dilemma (independent of the others’ strategies), it is crucial for cooperation that a strategy encounter others with which it is compatible. A context-preserving social structure means that strategies resulting from today’s interactions will be near each other tomorrow. This fosters the emergence of mutually compatible strategies.

We use the label ‘*friendly*’ for strategies with high values of p , those tending strongly to respond to cooperation with cooperation. The label ‘*provocable*’ characterizes strategies with low q , that is, those with a strong tendency to follow a defection with a defection.⁶

In a context-preserving social structure, if an agent does well by employing a friendly strategy, it will be copied by its neighbors, and in the succeeding periods the original strategy will be interacting with its compatible ‘imitators’. Friendly strategies such as Tit-for-Tat that diffuse through a context-preserving social structure will find themselves interacting with other friendly strategies. If an

unfriendly (low p) strategy does well and is emulated by neighboring agents, the resulting cluster of unfriendly strategies will be interacting in the following periods, lowering each other's scores.

Friendly strategies are vulnerable to exploitation. In a randomly mixing world, adaptation drives them out in favor of unfriendly – and provocable – alternatives. But in a context-preserving world, the friendly strategies resulting from adaptation are more likely to encounter other friendly strategies. Moreover, unfriendly strategies that remain near their own kind will do *less* well as a result. Context-preservation allows the adaptive future to cast its shadow. The stage is thus set for the emergence of a cooperative regime. The agents do not require foresight for this to happen. Indeed, our simulated agents have no expectations about the future. The process is driven entirely by the differential survival of adaptations that occur in what happen to be favorable contexts.

The differences between average scores in our matched population histories demonstrate that persisting patterns of interaction are clearly a major determining factor for the emergence of cooperative regimes. But, once again, we should press deeper if we are to better understand when this harnessing of a shadow of the adaptive future might be expected in real settings. This time we exploit another advantage of simulation, the opportunities it presents for controlled examination of dynamic sequences of events *within* single histories. It allows us to explicate in detail the processes that create and sustain a cooperative regime. This (stochastic) transition is observed repeatedly in all our simulations that engender cooperative regimes, and is not found in those that do not.

First, we describe dynamics that occur in all our simulations: an initial phase of learned mutual defection, bursts of strategies that are friendly and provocable, and 'predation' of non-provocable strategies by unfriendly ones. In section 3.3 we then describe the dynamics that distinctively characterize the transition to sustained cooperation.

3.2.1. The Initial Phase. Since we always begin a simulation with a random population of strategies located in randomly chosen positions, the first periods are very similar in all cases. Figure 1 shows the first 50 periods of average payoff for each of our first four structural variations, RWR, 2DK, FRNE, and FRN. All show an initial collapse from an unsustainable level of substantial mutual cooperation that occurs in our randomized initial configuration. With

minor sampling fluctuation, the first period results are the same in all conditions: each of the four cells of the Prisoner's Dilemma matrix is realized one quarter of the time, making average payoff 2.25. However, since at random initialization there is no structure to the interaction of strategies, all four cases evidence a rapid decline, as agents switch to the unfriendly, provokable strategies that did best in the initial random encounters.

As time passes, the three context-preserving conditions separate from the random mixing case. Their performance falls, but not as deeply, and it begins to recover, while the random pairing system never does. (Brief bursts of cooperation in some individual histories appear here as minor fluctuations of the dismal average of the 30 populations.)

The left column of Figure 2 shows the full 2500-period history for a representative individual population drawn from each of the five conditions RWR, FFR-0.5, FFR-0.3, FFR-0.1, and FRN. The typical history shown for random pairing in every period (RWR) is at the top left of Figure 2. The four subsequent histories in the left column are typical cases for progressively increasing levels of




Figure 1. Average score per move over the initial 50 moves for 4 experimental conditions: RWR, 2DK, FRNE, and FRN (see text for descriptions). The value for each period in each figure is the average score per move for all individuals in that period over all 30 histories (replications) for that condition.

context-preservation. The striking difference in these histories is the result of a transition to cooperation that is enacted reliably in the population with fixed interactions and is absent in the population where pairings are repeatedly randomized. Corresponding to each typical average-payoff history shown in Figure 2 is a 'p-q plot', constructed from data on changes in p and q during every period of each of the 30 histories in an experimental case. Using every occasion when the population average of p and of q lay within a small square region of the surface, the arrow points to the average values of p and q that resulted after adaptation processes were complete. Thus an arrow pointing up and to the left indicates that whenever p and q were near the base of the arrow, population average p fell and population average q rose. The circles on the plots indicate the amount of time that the populations spent in the various cells. The largest circle is the most frequently visited cell. The other cells have circles whose areas indicate the proportion of time spent there relative to the most frequently visited cell.

Note that Figures 3 and 4 are enlarged versions of the p-q plots for the RWR and FRN cases, respectively (i.e. the plots at the top and bottom of the right column of Figure 2). Figures 3 and 4 clearly show the direction and size of the arrows (and circles) in these two cases. The general pattern of counter-clockwise flow seen in those two extreme cases is also found in all the p-q plots for the cases with intermediate, and increasing, context preservation (FFR-0.5, FFR-0.3, FFR-0.1), as shown in Figure 2. In the following sections we describe the differences in these plots, which reflect the effects of the different degree of context preservation across the cases shown in Figure 2.

3.2.2. Bursts of Friendly Strategies. We begin at the top of Figure 2 by noting that even our repeatedly mixing RWR populations do have bursts of friendly strategies that correspond to increases in the population's average value of p . These bursts provide momentary opportunities for a cooperative regime. The single history displayed at the top left of Figure 2 shows a brief spike in average score per move. The p-q plot, based on all 30 histories, displays one large circle and two smaller ones in the lower left hand corner of the space. This shows that the populations spent most of their time with low values of p and of q . The thin trail of short arrows running to the right along the lower edge of the p-q diagram corresponds to the occasional bursts of high- p low- q (friendly, but provokable) strategies associated with spikes like the one in the

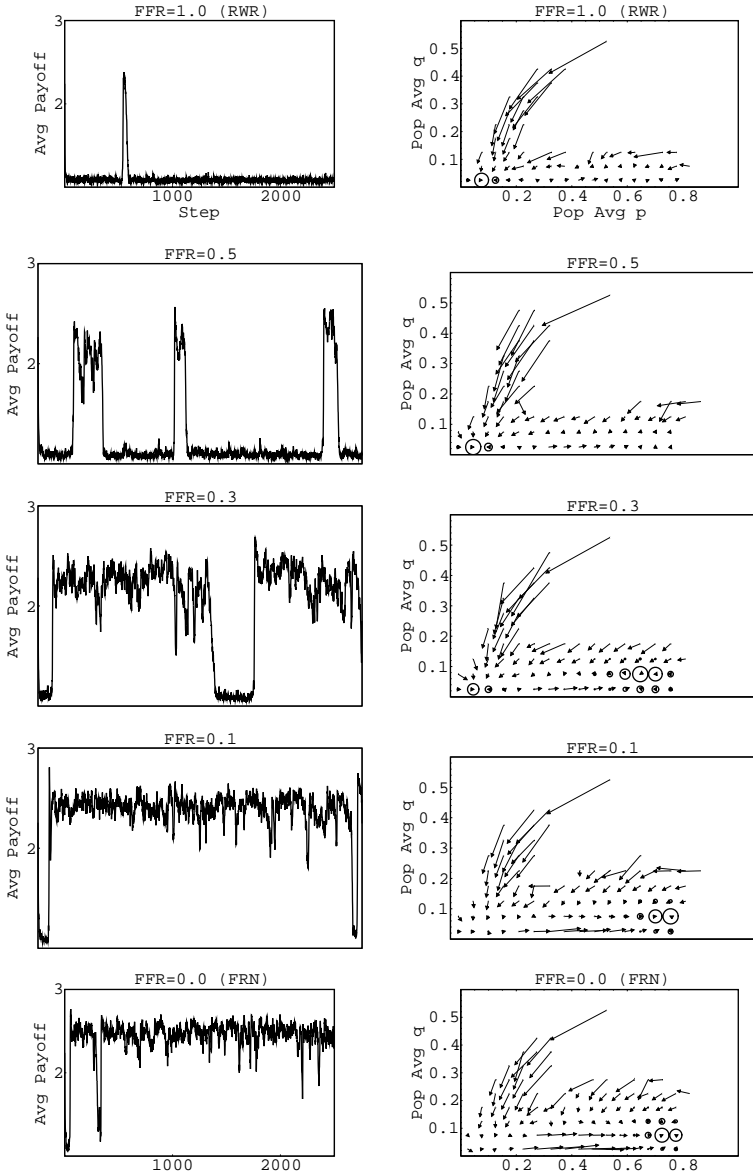


Figure 2. Average score per move (left column) for typical single runs, and p-q plots (right column) for a series of cases with decreasing (down the page) fraction of temporary random connections (RWR) added to fixed random connections (FRN). The p-q data are averaged over all 30 runs for each case.




Figure 3. Larger version of the p - q plot for the temporary Random With Replacement (RWR) case in Figure 2 (the top row, right plot).




Figure 4. Larger version of the p - q plot for the Fixed Random Network (FRN) case in Figure 2 (the bottom row, right plot).

history. The line of leftward arrows higher in the diagram corresponds to the immediate fall in population average p as the burst of friendly strategies is wiped out, a process we will describe in the next section.

We might describe the upward course of a burst this way: when nearly all strategies are unfriendly and provocable (in essence, 'always defect'), noisy emulation processes will produce in every period a few new strategies that are either friendly and/or non-provocable. Either variant strategy will be expected to do badly in a predominantly defecting world, and hence will be unlikely to be retained or emulated in subsequent periods. However, chance also guarantees that there will occasionally be random meetings of two variant friendly strategies. (Also, there will be meetings of non-provocable variants, but unless they are also friendly, they fare poorly and immediately vanish.) A random pairing of two friendly strategies does well enough to be emulated. Playing each other, the agents with high- p strategies may average as much as three points per move. They may average as little as 0.75 per move with always-defecting others, and there will typically be seven of those. But the average performance will work out to $8.25/8$, which will be superior to the average of $8/8$ that always-defects obtain with each other.

In this situation, where two friendly variant strategies have met by chance, they probably will be retained by their agents into the next time period. Moreover, the friendly strategy will be emulated by the 14 other agents whom those two encountered (or, more precisely, by 90 percent of them, on average, since performance comparisons involve noise). Thus, the succeeding period will contain the usual sprinkling of new friendly strategies created by chance emulation errors, and about 12–13 systematically transmitted strategies with high p .

At this substantially increased density, subsequent pairings of friendly strategies are considerably more likely. At the increased density expected in the period following a single chance friendly pairing, there is a non-trivial chance of two such pairings. The logic of increased density holds across several higher levels of friendly pair occurrence, so that any period with a small number of these events has a modest probability of being succeeded by a period with a higher count. The most common outcome for any given burst of dense friendliness is still for the nascent cooperative regime to be extinguished by the predation process described in the next section.

But the probabilities are such that 9 of our 30 histories reached an average score over 2.3, although their time in such cooperative circumstances was extremely brief.

3.2.3. Predation of the Friendly and Non-Provocable. The failure to sustain cooperation in these cases of repeated randomization can be understood from the p-q plot at the top right of Figure 2, for the RWR case. What is evident in the diagram is that with bursts of friendly strategies, average p can rise quickly to substantial levels. This shows as the rightward arrows along the bottom edge. But in these histories, where context is not preserved, the increase of friendly strategies is only momentary. Figure 2 shows their decline in the line of leftward pointing arrows slightly higher in the p-q space.

The RWR p-q plot (the top right of Figure 2, or Figure 3) also shows that at any increased q level, p levels will begin to fall. This is because any of the remaining strategies with low p and q will do extremely well against another with higher q. A strategy that is completely unfriendly and provocative always defects and scores 5 points per move against a maximally friendly and non-provocable strategy, which always cooperates. If the density is sufficient for it to meet two non-provocable others, and its six other encounters are with friendly and provocative strategies (which correspond to Tit-for-Tat), it will average⁷ – using pure types for approximation – $(2*5 + 6*2)/8 = 2.75$, while each of the others will average no more than $(0.75 + 7*3)/8 < 2.72$. The always-defect strategy will be emulated by all eight others, reducing the population average for p.

In our p-q plots this gives rise to the leftward arrows seen at higher levels of q in all histories, no matter what the social structure. Once there is a sufficient density of friendly, unprovocable strategies, successful predatory exploitation by unfriendly strategies is inevitable. However, in context-preserving structures, falling levels of friendliness are subsequently restored, while in randomly mixing systems they are not. This crucial difference is the focus of the following section.

3.3. The Emergence of Cooperative Regimes

What distinguishes the detailed sequences of events in the cases that form and maintain cooperative regimes from the cases where this does not occur? The crucial difference is adaptive dynamics at

low levels of q , just above the minimum. In all histories there are bursts of friendly strategies when q is extremely low. But what happens as a few friendly and unprovocable individuals enter the population, leaving retaliation to other agents? This always happens when there are errors in emulation, but what consequences it has will depend on the degree of context-preservation.

The succeeding rows of Figure 2 show the dynamics as context preservation steadily increases up to the fifth row, corresponding to the pure Fixed Random Network (FRN), where pairings are not shuffled at all at the close of each period. Here, a pair that cooperated will be paired again (except when noise disrupts one of the strategies). In such a context-preserving structure, the strategies a successful agent faces in the future are emulations of the strategy that agent employed in the past.

The second, third and fourth rows of Figure 2 show typical histories and p - q plots as we 'turn the dial' to increase context-preservation. Transitions to high scores per move become more common in the typical histories, and more persistent. Time spent in the high- p , low- q , region increases, as shown by the circles in the right-hand region of the p - q plots. Indeed, the third row of Figure 2, for the case with substitution probability 0.3, shows a bi-stable condition in which long stretches at high- p alternate with long stretches at low- p .

With this array we can see a crucial difference between context-preserving cases, in rows four and five of Figure 2, and the non-preserving cases in rows one and two. In the non-preserving cases of the first two rows, the arrows for population average q slightly above minimum (the second line of arrows from the bottom) show a leftward flow. The slightest increase of non-provocability in the population is instantly exploited and unfriendly strategies are copied as a result. However, as context-preservation increases, moving toward the last rows in Figure 2, this flow reverses. In these cases, when the population has the same slightly elevated level of q , it is carried toward higher p values. The population spends its time in the right-hand corner, rather than the left one. We have a cooperative regime.

We can examine precisely what happens to any population with a history that passes through a crucial region of the p - q space, where p is in the interval $[0.3, 0.35]$ and q in $[0.05, 0.10]$. Where there is random mixing in every period (RWR - top row), the average value of p for a population in this region moves down by -0.016 . In the fifth row, where there is random initial pairing and no further

shuffling (FRN), the populations visiting the same crucial p - q region saw p values increase by 0.052. In the same region of p - q space, randomly mixing and context-preserving populations are headed in opposite directions.

We can get an immediate insight into the cause of this key difference by examining further data on this p - q region of the two conditions. Figures 5 and 6 show the relationship between the p value of an agent's strategy (its friendliness) and the average p values of the strategies it encountered. The data are for all individual agents during periods of any population history (for a given experimental condition) where the average values of p and q were in the p - q region defined above. Figure 5, for the repeated random mixing case (RWR), shows that the p of encountered agents varied widely from the agent's own p value. Figure 6, for the Fixed Random Network case (FRN), shows a much tighter correlation of the p values of self and other. While the regression of own p on average p -value of others is not significant for RWR, it is highly significant for FRN (slope = 0.1580; $F = 717.20$; $N = 5376$).

This substantial difference in who meets whom has transforming consequences. In this key region of the p - q parameter space, the random mixing regime frequently brings friendly strategies together with unfriendly ones. Unfriendly strategies prosper and spread, and the population moves to lower p . In the context-preserving cases, friendly strategies that succeed are emulated by their continuing neighbors, and subsequently interact much more with their own kind. They prosper further. An unfriendly strategy that has a success interacts in the next period not with more friendly strategies on which it can prey, but with many of its own kind, since it has been emulated by continuing neighbors. The unfriendly strategy no longer prospers and is not emulated. The two effects move the population to higher average p . With context-preservation a channel is opened that allows the continuous restoration of a cooperative regime whenever predation of non-provocables threatens to carry it to the Hobbesian regime of unfriendly provocability.

By exploiting simulation tools we have been able to demonstrate cooperative regimes in populations, and then to probe the detailed dynamics that distinguish those cases from the histories of other populations that languish in mutual defection. Our results provide a precise demonstration of an effect well appreciated by Giddens, that continuity of interaction is a crucial aspect of social structure.




Figure 5. For the temporary Random With Replacement (RWR) case there is no relation between an agent's own friendliness (p) and the average friendliness of the other agents it played in a time step in which the average population p and q values were in the ranges $[0.30,0.35]$ and $[0.05,0.10]$, respectively. (The line is the regression line through the 7168 points for this case.)

Such continuity preserves the contexts in which strategies have worked, allowing later variants and emulations of those strategies to find a supportively similar environment.

4. Implications

Our results are consequential – both in their implications for theory and for the illumination of important contemporary issues.

A deep implication of our results is the demonstration that the sociology of a system contributes massively to its adaptive dynamics. This is a point that goes far beyond Prisoner's Dilemma,




Figure 6. For the Fixed Random Network (FRN) case, there is a positive relation between an agent's own friendliness (p) and the average friendliness of the other agents it played in a time step in which the average population p and q values were in the ranges $[0.30, 0.35]$ and $[0.05, 0.10]$, respectively. (The line is the regression line through the 5376 points for this case.)

applying just as well to evolutionary accounts of social roles and routines. Game theory, even though it has made great progress in admitting processes of learning and adaptation (Fudenberg and Levine 1998), has paid far too little attention to the details of who interacts with whom. But these simulations demonstrate that those details can be decisive, by themselves completely reversing the course of system evolution. It appears that the continuing search in game theory for improved equilibrium concepts can usefully be augmented with work on the underlying interaction structures that may be equally determinative of which outcomes are realized.

A related theoretical implication is the major opportunity to advance our understanding of social processes by conceiving them

as occurring on – or in – networks with varying structures. There are considerable challenges here, since closed-form mathematical techniques are not always available. But some progress can be made with simulation, as we have shown. And it is encouraging that we have been able to approximate the dynamics of quite different networks by varying the rate of substitution in a random network. Watts (1999) shows in a similar vein that many qualitative properties of dynamics on networks can be studied by varying the amount of random perturbation of highly clustered networks. The two lines of work are complementary. His studies are comparisons of dynamics on fixed networks that differ in the proportion of random connections; our studies are of dynamics on changing networks. However, together they suggest that workable approximations for large classes of network structures may be found. With such approximations many important insights into dynamics within structure may become reachable.

With respect to contemporary social concerns, work such as ours and other work on dynamics on networks, such as that of Watts (1999), opens the door to a better-grounded understanding of the impacts of the massive communications upheaval now occurring. Falling costs of communication, over the Internet and by telephone, are dramatically raising the rates of interaction among physically and socially dispersed agents, and therefore could be decreasing the clustering of their social networks. The issues this raises are not new, stretching back to the concern of (Tönnies 1887) for the waning of *Gemeinschaft*. But the contemporary growth of connection to distant agents is dramatic, and a reason for renewed attention. Our results suggest that the simplest projection may be inadequate. Decreased clustering may not entail the loss of cooperative regimes. In our simulations, context preservation and the shadow of the adaptive future suffice to overcome the effects of random network connectivity and constantly changing strategies.⁸ The results provide only an 'existence proof', of course. But even if we cannot guarantee cooperative regimes, it is a gain to understand how they are possible in the novel conditions of the distant social interactions now unfolding all around us.

NOTES

For financial support, we thank the Intel Corporation and the University of Michigan College of LS&A College Enrichment Fund. For computing facilities, we thank

the University of Michigan Center for the Study of Complex Systems. Correspondence should be addressed to Michael D. Cohen (e-mail: mdc@umich.edu)

1. Although we have not done the full analysis reported here for population sizes other than 256, our simulation runs for populations up to 4096 agents display very similar aggregate statistics to those reported here and appear to us to sustain all the relationships we report for the 256 agent case.
2. While the y values do in fact vary from the p values in individual agents, there is a strong correlation between agents' y and p values, and the average y closely tracks the average p . Thus to simplify the analysis, we only report analysis of the p values.
3. Complete details for a wider set of experiments, including other strategy representations, other adaptive methods, and other social structures, are available in Cohen, Riolo, and Axelrod (1999). Here we discuss only a subset of the experimental conditions and results that are especially pertinent to the maintenance of cooperative regimes when agents are embedded in different social structures. A very careful reader will notice very minor discrepancies that are due to improved treatment of infrequent tie events programmed since completion of the working paper. Those changes did not produce any differences that altered the basic results reported in either paper.
4. To appreciate the difference between a network based on geography and one based on arbitrary links, consider the contrast between two pure types. In both networks, each person has exactly four 'neighbors'. In a two-dimensional lattice, the neighborhoods are correlated. For example, a person has two neighbors in common with whoever 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. 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. In this case, the network fans out in a tree structure. A person has four immediate neighbors, 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: $N(d) = 4 \times 3^{d-1}$. For a random network in a large finite population, two people may share more than one neighbor, but this will be rare. Rapid fanout of links in a random network can help the diffusion of information. On the other hand, rapid fanout as well as the lack of neighborhood clustering raises questions about the prospects for prosocial behavior in an uncorrelated structure. The table in the Appendix reports on the fanout measured in some typical graphs used in the studies reported here; in short, the fanout for random graphs of size 256 is very close to that predicted above for d up to about 4, after which the finite size effects begin to reduce the fanout at an increasing amount. And as expected, for larger graphs the fanout approximates the theoretical limit up to higher d values.
5. Statistically, the FRNE populations actually have a better average score than the 2DK populations, but the magnitude of the difference is not important for our argument, and the direction of the difference actually makes the clustered spatial embedding slightly inferior to its fixed random equivalent.

6. The friendliness property of a strategy is not identical to the property labeled 'niceness' in Axelrod (1984). Strategies that are *nice* are not the first to defect. Provocability as used here is similar to the earlier usage.
7. The unfriendly and provocable always-defect receives 5 points per move from each of its meetings with non-provocable others, and the always-defect receives 2 points per move from each of its six games with friendly and provocable others, i.e. 5 points in the first move and 1 point on each of the three subsequent moves.
8. We elaborate this point in Axelrod, Riolo, and Cohen (2000).

REFERENCES

- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, Robert. 1986. 'An Evolutionary Approach to Norms.' *The American Political Science Review* 80: 1095-111.
- Axelrod, Robert, Rick L. Riolo, and Michael D. Cohen. 2001. 'Beyond Geography: Cooperation with Persistent Links in the Absence of Clustered Neighborhoods.' In *Adaptation and Evolution: Essays in Honor of John H. Holland*, eds L. Booker, S. Forrest, M. Mitchell, and R. Riolo, pp. ???. Location: Publisher (forthcoming).
- Baker, Wayne E., Patricia Adler, and Peter Adler. 1984. 'Floor Trading and Crowd Dynamic.' In *The Social Dynamics of Financial Markets*, eds Patricia A. Adler and Peter Adler. Greenwich, CN: JAI Press.
- Cohen, Michael D., Rick L. Riolo, and Robert Axelrod. 1999. 'The Emergence of Social Organization in the Prisoner's Dilemma: How Context Preservation and other Factors Promote Cooperation.' The Santa Fe Institute Technical Report 99-01-002. Santa Fe, NM.
- Fudenberg, Drew and Eric Maskin. 1990. 'Evolution and Cooperation in Noisy Repeated Games.' *American Economic Review* 80: 274-9.
- Fudenberg, Drew and David Levine. 1998. *Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Giddens, Anthony. 1979. *Central Problems in Social Theory: Action, Structure, and Contradiction in Social Analysis*. Berkeley, CA: University of California Press.
- Giddens, Anthony. 1984. *The Constitution of Society: Outline of a Theory of Social Structuration*. Cambridge, MA: Polity Press.
- Cluever, Juergen and Joern Schmidt. 1999. 'Topology, Metric, and Dynamics of Social System.' *Journal of Artificial Societies and Social Simulation* 2(3). <http://jasss.soc.surrey.ac.uk/2/3/contents.html>
- Lomborg, Bjorn. 1996. 'Nucleus and Shield: The Evolution of Social Structure in the Iterated Prisoner's Dilemma.' *American Sociological Review* 61: 278-307.
- Nowak, Martin and Robert M. May. 1992. 'Evolutionary Games and Spatial Chaos.' *Nature* 359: 826-9.
- Nowak, Martin and Karl Sigmund. 1989. 'Oscillations in the Evolution of Reciprocity.' *Journal of Theoretical Biology* 137: 21-6.
- Riolo, Rick L. 1997a. 'The Effects and Evolution of Tag-Mediated Selection of Partners in Populations Playing the Iterated Prisoner's Dilemma.' In *Proceedings of*

- the 7th International Conference on Genetic Algorithms (ICGA97)*, ed. Thomas Bäck, pp. 378–85. San Francisco CA: Morgan Kaufmann.
- Riolo, Rick L. 1997b. 'The Effects of Tag-Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoner's Dilemma.' The Santa Fe Institute Technical Report 97-02-016. Santa Fe, NM.
- Watts, Duncan. 1999. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton, NJ: Princeton University Press.
- Tönnies, F. 1887. *Gemeinschaft und Gesellschaft (Community and Society)*, trans. C. P. Lewis, 1957. East Lansing, MI: Michigan State University Press.

MICHAEL D. COHEN is Professor of Information and Public Policy at the University of Michigan. He has also served as an external faculty member of the Santa Fe Institute. His most recent book (with Robert Axelrod) is *Harnessing Complexity: Organizational Implications of a Scientific Frontier*.

ADDRESS: Michael D. Cohen, School of Information, University of Michigan, Ann Arbor, MI 48109 USA [email: [mdc,rlriolo,axe]@umich.edu]

ROBERT AXELROD is the Arthur Bromage Distinguished University Professor of Political Science and Public Policy at the University of Michigan. He is a member of the National Academy of Sciences and a MacArthur Prize Fellow. Among his books are *The Evolution of Cooperation* (1984) and *Harnessing Complexity* (2000) with Michael D. Cohen.

RICK RIOLO is Associate Research Scientist and Director of the Computer Lab of the Centre for the Study of Complex Systems at the University of Michigan.

Appendix: Experimental Methods

A single *run* of the model lasts 2500 periods (after the initial one). The result of a run is a *history* of agent activity and change, as well as a history of resultant population-level aggregate measures. Because these models include many stochastic processes, we generate 30 histories for each case (i.e. for each combination of experimental parameter values), each run beginning with a different seed for the random number generator.

Abstractly, a run of the model, to generate a single history, proceeds as follows:


```

Population(0) = GenerateRandomPopulation
for each Period T
  for each individual A from Population(T)
    for 4 IPD games
      X = GetOtherToPlay
      PlayPrisonersDilemma( A, X, 4 Moves )
      UpdateCummulativePayoffs( A, X )
    endfor
  endfor
  NormalizePayoffs
  Population(T+1) = ApplyAdaptiveProcedure( Population(T) )
endfor

```

Note that we picked four moves per game because under the general conditions we are using, the shadow of the future imposed by that game length makes the attainment of cooperation difficult but not impossible (Riolo 1997a, b). Since the strategies do not take the move number into account, four moves per game corresponds to $w = 0.75$ in Axelrod (1984).

In our simulations a *period* consists of each of 256 agents in the population selecting four others to play iterated Prisoner's Dilemma (IPD) games. An IPD *game* consists of four moves played between agent A and one other agent, X. We use the standard payoff matrix shown in Table 1. Each agent's strategy is represented as in Nowak and Sigmund (1989), by a triplet of real numbers $0 \leq y, p, q \leq 1$. The first of these, y , represents the probability that the agent will cooperate on the first move of a game. For moves after the first move, the second, p , is the probability that an agent will cooperate following a cooperative move by the agent with which it is playing. The last, q , is the probability that the agent will cooperate after the other agent has defected. We initialize y as equal to p , but they may vary independently thereafter.

For period 0, the initial population is generated by evenly distributing the agents throughout the strategy space. In each subsequent period each individual plays IPD games with a small sample of other agents from the population; which others (and the exact number) each agent ends up playing in a period depends on the selection procedure (GetOtherToPlay), which is determined by the *social structure* being studied, as described below. As a result of these games, each individual accumulates an average payoff per move for the period. After all agents have played, agents are changed by the adaptive procedure being used for the run, as described below. Then the next period commences.

In this report we systematically vary the *social structure* that controls which others an agent plays each period (i.e. it controls what X is returned by the GetOtherToPlay procedure). We sometimes refer to the others an agent plays as its *neighbors*. In this paper we report on four social structures (or mixtures of these four):

RWR–Random With Replacement: This interaction process is one kind of complete-mixing model often utilized in analytic population-level models. Each agent *A* plays four others chosen with equiprobability from the entire population (not including *A* itself). The others are chosen with replacement, so it is possible (but rare) that an agent would play the same other twice in one period. Note, however, that while on average each agent *A* will play 8 others in a period (exactly 4 others *A* chooses, with *A* being chosen by 4 others, on average), there is some variance in the number of others the agents play, with some agents playing more or less than 8 games per period.

2DK–Two-Dimensional torus, Keep Neighbors: In contrast to *RWR*, *2DK* provides complete continuity and clustering of neighborhoods. At the start of each run, the agents are placed at random on a 16 by 16 torus. Each period each agent plays its 4 NEWS (North-East-West-South) neighbors. Since those agents also have 4 NEWS neighbors (i.e. the neighbor relation is symmetrical), each agent plays exactly 4 other agents, playing each twice during a period. The agents never move, so they keep the same neighbors over all periods in a run.

FRN–Fixed Random Neighbors: This social structure is like *RWR*, in that neighbors are picked at random, non-symmetrically, but it is also like *2DK*, in that the neighbors remain the same throughout each run. At the start of each run, each agent *A* picks 4 other agents to be its neighbors, chosen with equiprobability and with replacement from the entire population (not including *A* itself). Thus as with *RWR*, on average each agent plays 8 different agents each period, but there is considerable variation with some playing more and some playing fewer others. However, unlike *RWR*, with *FRN* these neighbors are retained for the entire run, so the agents play the same set of others every period.

FRNE–Fixed Random Neighbors, Equal numbers: This social structure is designed to be like *2DK* in that each agent has exactly 4 other agents as (symmetric) neighbors, but each agent's neighbors are chosen at random, so that there are none of the correlations between neighbors-of-neighbors that are induced by a *2DK*-type topology. As in *2DK*, the agents keep their neighbors for the entire run. Thus at the start of each run, each agent *A* randomly selects 4 other agents to be its neighbors, and they in turn choose *A* to be one of their neighbors. After all agents have 4 symmetric neighbors, the neighbors are randomly mixed to ensure there are no correlations induced by the order in which agents choose others. The randomization is done as follows: After all agents have 4 symmetric neighbors, each agent randomly swaps a neighbor with 256 randomly chosen other agents, with each swap subject to the *FRNE* constraints that each agent have 4 unique neighbors and each neighbor relation is symmetric. Note that this method of constructing the neighborhood relationship is similar to the algorithm Watts describes for his β -model of relational graphs (Watts 1999), for $\beta = 1.0$.

In order to verify that the *FRNE* neighborhoods generated as described

above indeed have less clustering (correlation) than 2DK neighborhoods, we measured the distribution sequence (Watts 1999: 32), i.e. the number of agents exactly d steps away (and no less) from each agent. As mentioned in note 4, for infinite 2DK graph with 4 NEWS neighbors for each agent, the number of different agents exactly d away (and no less) increases linearly, $N(d) = 4d$. But for a graph in which each agent's 4 neighbors are chosen at random from a large (infinite) population, the number of others at a given network distance expands exponentially: $N(d) = 4 \times 3^{d-1}$. Table A1 gives the number of agents d steps away, averaged over all agents in a typical FRNE population of 256 agents from our experiments. (Results from other graphs are almost identical.) The table also gives the expected number d away for 2DK and for random graphs constructed with an infinite population.

Table A1.

d	Avg. $N(d)$	Predicted random	2DK
1	4.00	4	4
2	11.74	12	8
3	32.38	36	12
4	74.59	108	16
5	98.42	324	20
6	33.23	972	24

The number of agents found to be d steps away from a given agent in an FRNE social structure constructed as described in the Appendix, averaged starting from all 256 agents in a population. Also shown are the number of agents predicted to be d away from a given agent in 2DK and random graphs of infinite size.

As can be seen the number of agents d away in the FRNE graph grows approximately exponentially, at least until $d = 3$ or 4, after which the finite size effects of the small population begin to grow. The fanout is clearly much broader than that predicted for a 2DK neighborhood structure. Similar results are found for other FRNE graphs we have tested; and, as expected, for larger populations (e.g. 4096), the measured $N(d)$ remains close to the infinite population value for even larger values of d .

At the end of each period, all agents are subjected to an *adaptive procedure* which may change their strategies. Each agent A identifies the most successful of its neighbors (i.e. the agents A played in the current period). The most successful is the neighbor agent X with the highest average payoff per move resulting from *all* games X played during the current period. Ties are broken randomly if more than one neighbor has the same highest average payoff per move among A's neighbors. The agent A then updates its strategy by adopting X's strategy if X did strictly better than A in the current period, i.e. if X's average payoff per move is higher than A's average payoff per move (again, over *all* games each played in the current

period). The update rule includes two sources of possible error. First, there is a 10 percent chance the agent makes a mistake when deciding whether or not its best neighbor did better than itself. This results in the agent's occasionally copying an inferior strategy or failing to copy a superior one. Second, regardless of which of the two strategies (its own or its best neighbors) is adopted, for each of the three parameters (y,p,q) there is a 10 percent chance that Gaussian noise is added to the parameter (mean 0, s.d. 0.4, bounded by 0 and 1), i.e. sometimes resulting in a strategy that is different from the one being adopted.