

Agent Mobility and the Evolution of Cooperative Communities

An artificial world is constructed that is based upon a spatial iterated prisoner's dilemma game. Several additional features are introduced into this model, the key feature being the ability of agents to move around in their world. Movement is a mechanism for exit or noncompulsory play. When agents can move, high levels of cooperation are achieved more frequently and are considerably more stable than when they cannot move. Also, when cooperative worlds occur, they are generated and sustained by the formation of networks of densely connected "cooperative" agents that can withstand invasion and parasitism by noncooperative agents.

STEPHEN J. MAJESKI, GREG LINDEN, CORINA LINDEN, AND AARON SPITZER

Stephen Majeski is Professor of Political Science at the University of Washington. Greg Linden is affiliated with the Department of Computer Science and Engineering at the University of Washington and Amazon.com. Corina Linden and Aaron Spitzer are graduate students in Political Science at the University of Washington.

INTRODUCTION

We are concerned with whether and how cooperation evolves in social worlds characterized by the presence of selfish agents engaged in repeated relations without central authority. This fundamental problem has been the focus of many studies across all the social sciences as well as philosophy, biology, and computer science. Based on the seminal works of John Maynard Smith [1], and Robert Axelrod [2], the iterated prisoner's dilemma (IPD) has become the central metaphor for the evolution of cooperation in populations of selfish agents without central authority.

To investigate whether and how agents in an IPD can overcome the individual rational choice to defect to achieve the socially optimal outcome of mutual cooperation, we construct an artificial world and then analyze characteristics of that world via simulation. The artificial world that we construct is an anarchic, competitive place populated by agents with limited cognitive and social capabilities whose resources are being constantly drained. Into this harsh environment we introduce agent movement and constrained interactions in the form of spatial neighborhoods, features that we anticipate will foster cooperation.¹ We are interested in several properties of this artificial world. Does it become highly cooperative? If cooperation emerges, is it stable? Do agent movement and the type of movement facilitate the generation of a stable cooperative world? Are there patterns in the kinds of strategies that help generate a cooperative world or in the kinds of strategies that evolve in a cooperative world?

STRUCTURE OF THE ARTIFICIAL WORLD

The basic architecture of the artificial world is an IPD.² Agents have two choices: cooperate (C) and defect (D). For each interaction, both agents receive a payoff of

CC if both cooperate and DD if both defect. If one agent defects while the other cooperates, then the exploiter receives DC and the sucker receives CD. To conform to the prisoner's dilemma, agent preferences across payoffs adhere to the following standard inequalities: $DC > CC > DD > CD$ and $2CC > DC + CD$. Each agent is represented by a strategy specifying how the agent behaves as it interacts with other agents. Agent strategies are restricted to those employing just the previous interaction with the other agent(s) to determine current choices.³ Strategies are probabilistic, defined by the conditional probabilities to cooperate (p_1, p_2, p_3, p_4), given that the outcome of the previous interaction was CC, CD, DC, or DD, respectively, with the addition of a cooperate or defect choice stipulation when an agent interacts with another agent for the first time. The payoffs used in the simulation are $CC = 1, CD = -3, DC = 3,$ and $DD = -1$, although additional simulations with different payoffs also were implemented to assess the sensitivity of the results to changes in the payoffs. The initial mix of agent strategies is a set of identical strategies, all with the following set of conditional probabilities (.5, .5, .5, .5). Also, half the agents cooperate the first time they interact with another agent, and the other half defect the first time they interact with another agent. The purpose of choosing this "random strategy mix" is to avoid an initial selection of strategy mixes that is biased toward generating either a cooperative or a noncooperative world.⁴

Several features are added to the basic IPD to produce the artificial world. First, in most social contexts, agents are located at or occupy some place or position in their world at any given moment in time and, as a result, most in-

teractions among social units remain dictated by spatial proximity.⁵ The addition of a spatial component to the artificial world introduces the possibility of spatial variability.⁶ Therefore, an explicit spatial dimension is introduced by constructing a set of toroidal worlds (a 20×20 grid of cells) consisting initially of 60 agents that are randomly as-

Specifically, agents move only when they receive a nonpositive total payoff for the current round. Agents receiving a positive payoff in the current round are "satisfied" with their location and remain there. Because there is a nontrivial risk associated with movement, agent movement is probabilistic. In practical terms, if agents were to move with certainty every time a nonpositive payoff were obtained, then far too much movement would occur and chaos would reign.

signed locations on the grid.⁷ Each cell can contain, at most, one agent. For each round of the simulation, agents interact with all agents who occupy the four nondiagonal cells that immediately surround the agent: a von Neumann neighborhood.

All social entities consume various resources to sustain themselves. All ecologies can support a finite number of agents and, as more agents compete, the economic and environmental costs of available resources increase. Therefore, as the artificial world becomes more populated, the cost of surviving increases.

Second, while compulsory play is a typical characteristic of most IPD analyses, there are probably few social situations in which agents are forced to interact with some set of other agents. Agents usually can choose to exit, withdraw, or refuse to play, suggesting that many kinds of social interactions appear to be better represented as forms of noncompulsory play.⁸ We focus on agent movement as form of noncom-

pulsory play. When agents find themselves in an unprofitable and undesirable situation, they often move to free themselves from the negative consequences of that situation.⁹

Since agents are assumed to have little information about their environment, they only move when they are doing poorly at their current location, and the worse they are doing, the more likely it is that they will move. Specifically, agents move only when they receive a nonpositive total payoff for the current round. Agents receiving a positive payoff in the current round are "satisfied" with their location and remain there. Because there is a nontrivial risk associated with movement, agent movement is probabilistic. In practical terms, if agents were to move with certainty every time a nonpositive payoff were obtained, then far too

much movement would occur and chaos would reign. The likelihood that an agent will move increases with the magnitude of the negative total payoff of the current iteration. An agent moves if the change in energy (ΔE , defined later) is greater than or equal to a randomly generated integer between 1 and 100. Two types of movement are introduced: local, in which the agent moves to an empty adjacent cell; and global, in which the agent moves to a randomly selected open cell anywhere on the grid.

Third, we incorporate the notion of an environmental carrying capacity into the artificial world by introducing a cost of survival for agents. All social entities consume various resources to sustain themselves. All ecologies can support a finite number of agents and, as more agents compete, the economic and environmental costs of available resources increase. Therefore, as the artificial world becomes more populated, the cost of surviving increases.¹⁰ The number of agents is restricted to a fixed

range by applying to every agent each iteration a cost of surviving that is dependent on population size. The formula for the cost of surviving, α , is

$$\alpha = k + 4 \cdot (DC + CC) \cdot N / (X \cdot Y)$$

in which k is a constant, DC and CC are the IPD payoffs, N is the number of agents in the world, and X is the width and Y is the height of the world grid. The cost of surviving indirectly allows the simulation to select the percentage of the population with the highest energy levels for reproduction and the lowest percentage for elimination. The change in energy (ΔE) for each iteration of an agent is the sum of all interactions minus the cost of surviving:

$$\Delta E = \sum_{i=1}^4 A_i - \alpha$$

where A_i is the payoff from the interaction in the i th direction and α is the cost of surviving. An agent is eliminated from the simulation when its energy falls below 0.

Fourth, because all individuals die and all social units, that we are aware of, eventually fall apart, disband, go bankrupt, are taken over, or are overrun, we assume that agents have a limited existence or life span. An agent has a probability Γ of elimination on each iteration,

$$\Gamma = (A - T) / M$$

where A is the age of the agent, T is a constant for the minimum life span, and M is a constant where $T + M$ is the maximum life span. Once an agent reaches the minimum life span, then it has a nonzero and increasing probability of elimination until it reaches the maximum life span, and then it is eliminated with a probability of 1.0. Of course agents can be eliminated at any time if their energy level falls below 0. In the basic simulation results that will be reported later, the minimum life span of an agent (T) is 2048 iterations and the maximum life span ($T + M$) is 6144 iterations.

TABLE 1

Simulation Results

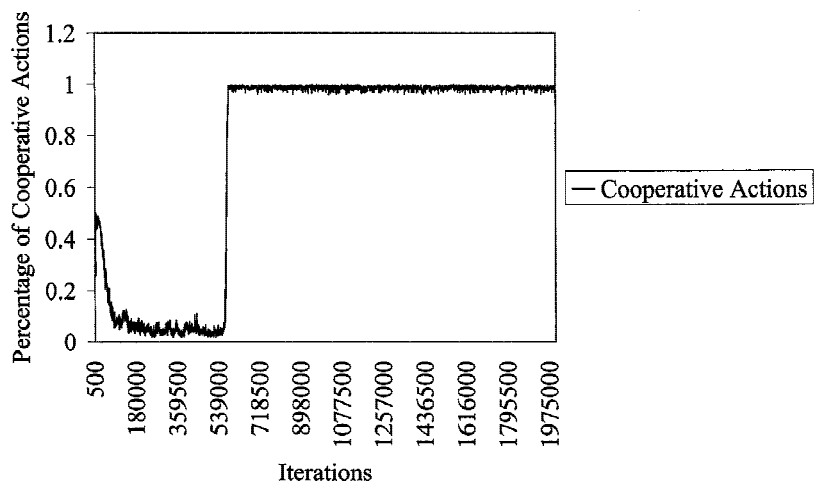
Outcomes	No Movement	Global Movement	Local Movement	Total
Stable Cooperation	11 (37%)	24 (80%)	24 (80%)	59 (65%)
No Cooperation	6 (20%)	6 (20%)	3 (10%)	15 (17%)
Punctuated Equilibrium	13 (43%)	0 (0%)	3 (10%)	16 (18%)

Fifth, agents reproduce themselves, in the sense that they create a replication of themselves. Since we are attempting to model social units and not just individuals, we have opted for asexual reproduction. The “genetic material” (here, the strategy), therefore, comes from one “parent,” and there is no “crossover” of genetic material. Replication requires a certain level of energy (maturity, size, power, and wealth), and it costs the agent a significant amount of energy. Therefore, agents must reach a fixed level of energy (ρ) before they can replicate. Once an agent replicates itself, the energy of the agent and the replicated agent are both set to ($\rho/2$). ρ is set to 1000 for the basic simulation runs and is varied for additional simulation runs. For all simulations in this analysis, the replicated agents are placed in a randomly selected open cell

on the grid. Our approach to setting reproduction thresholds and to determining the relative fitness of the members of the population has the advantage of performing the reproduction/elimination calculation at every iteration instead of making periodic sweeps through the population, more gracefully modifying the population of the simulation.

Sixth, to give the artificial world some dynamic, we introduce strategy mutation. There are a number of ways to vary an agent’s strategy or to introduce new strategies. Agent strategies can change via imitation, learning, or innovation, or they can emerge generationally via mutation. In a world where agents face a noisy environment, have difficulty observing the behavior of other agents for a significant period of time, where it is often **not** in the interest

FIGURE 1



Cooperative actions with local movement.

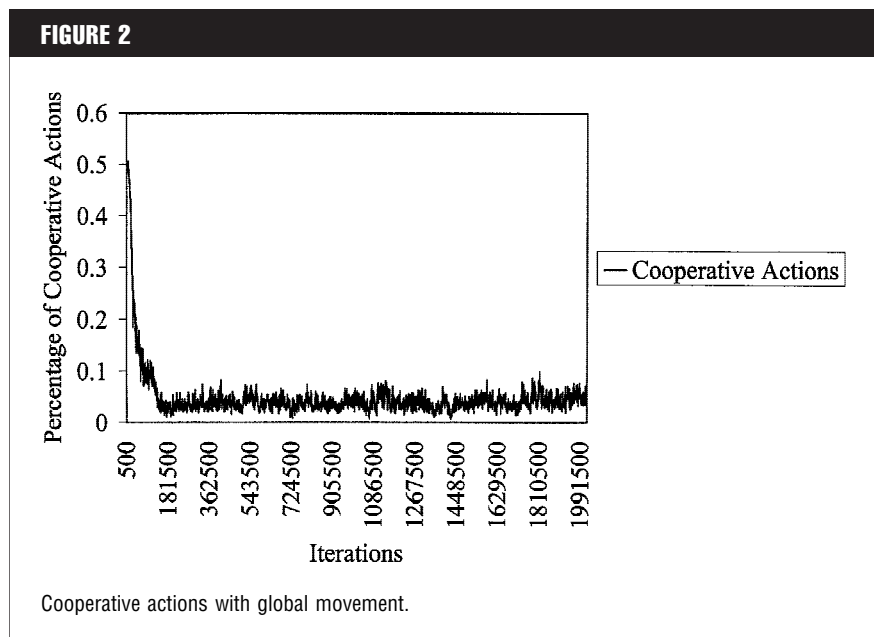
of agents to reveal their strategies, and where there are no reliable mechanisms to transmit that information, change via imitation or learning seems less appropriate than change by either innovation or mutation.¹¹ The introduction of new strategies occurs when replication occurs. There is a fixed (20%) chance that a strategy mutates during the replication process. When a mutation occurs, the agent's strategy is modified by changing each of the p_i in the strategy of the parent by $[-\delta, \delta]$, where $0 < \delta < 1$. Specifically, if δ is set to (0.1), as is the case for the basic simulation results, then the actual value to change the p_i is randomly selected from a uniform distribution over the interval (-0.1) to (0.1) .

RESULTS

Original Configuration

To address the set of questions posed in the Introduction, 90 simulations of this artificial world, as initially configured in section 2, were run: 30 with no movement, 30 with local movement, and 30 with global movement.¹² Several general patterns emerged. First, as presented in Table 1, most simulation runs (83%) achieve a high level of cooperation (the average cooperation rate among agents is over 95%) at some point during the course of the simulation run. Fifty-nine simulation runs (65%) eventually ended up in what appears to be a stable cooperative equilibrium (a high level of cooperation is achieved at some point in the simulation and is maintained until the end of the simulation), and 16 simulation runs (18%) ended up in a punctuated equilibrium (long periods of stable, high levels of cooperation punctuated by periodic massive dips to near universal defection), a pattern that is similar to that found by Nowak and Sigmund [10]. Only 17% (15 runs) ended up in a stable, noncooperative equilibrium (the average cooperation rate quickly declined and stayed at less than 5% for the entire simulation run).¹³

Second, somewhat surprisingly, the introduction of movement does not seem to have a significant impact on



whether or not high levels of cooperation are achieved. High levels of cooperation are not achieved (“no cooperation” in Table 1) about the same percentage of the time for all three types of movement. However, the introduction of movement does have an important impact on the stability of cooperative equilibria. Simulations with local or global movement achieve stable cooperation 80% of the time, while those simulations with no movement achieve stable cooperation only 37% of the time.

Simulations with no movement experience punctuated equilibria 43% of the time, while those simulations with some form of movement experience punctuated equilibria only 5% of the time. Examples of typical patterns of stable cooperation, no cooperation, and punctuated equilibrium are depicted in Figures 1–3, respectively.

Third, the transition from near universal defection to near universal cooperation is always characterized by two features: networks of cooperative




TABLE 2

Summary of All Sensitivity Results

Outcomes	No Movement	Global and Local Movement	Total
Stable Cooperation	25 (39%)	79 (83%)	104 (65%)
No Cooperation	10 (15%)	6 (6%)	16 (10%)
Punctuated Equilibrium	30 (46%)	10 (11%)	40 (25%)
Total	n = 65	n = 95	n = 160

agents form, and these networks are composed of agents that employ versions of the Grim strategy (a pure form of the Grim strategy is [1.0, 0.0, 0.0, 0.0]).


Sensitivity Analysis

A total of 160 additional simulations were run to assess the robustness of the simulation results to variations in payoffs, mutation rates, mutation magnitudes, reproduction thresholds, and life

span lengths. First, 30 simulations were run with two alternative payoff matrices. The payoffs of the first alternative matrix are CC = 1, CD = -5, DC = 5, and DD = -1. This reconfiguration serves to increase both the potential payoff from exploiting one's partner and the potential loss from being exploited. The second alternative matrix is CC = 2, CD = -3, DC = 3, and DD = -2. This matrix broadens the difference in payoffs between the mutual cooperation payoff

CC and the mutual defection payoff DD to four points, as opposed to two in the original matrix. Second, 30 simulations were run varying both the length of the agent's life span (shorter, 2000 iterations, and longer, 6000 iterations) as well as whether the life span occurred earlier (from 1000 to 5000 iterations) or later (from 3000 to 7000 iterations). Third, 30 simulations were run with lower (10%) and higher (40% and 50%) mutation rates. Fourth, 30 simulations

The simulation results indicate that the introduction of movement (either global or local) increases substantially the likelihood that stable, highly cooperative worlds are generated; 80% of the time for simulations with movement versus 37% without movement.

FIGURE 4

Nodes of cooperation emerge. Grim (red) agents begin to dominate over more noncooperative (darker shaded) agents. Turn number 306,202 of a simulation with local movement.

were run varying the magnitude of change in strategy probabilities due to mutation; specifically, $\delta = (0.3)$, and $\delta = (0.5)$. Fifth, 40 simulations were run with two alternative reproduction thresholds: one higher, requiring 2000 units of energy to reproduce and one lower, requiring 800 units of energy to reproduce.

The results of these sensitivity runs add considerable additional support to the general results regarding the effects of movement. As reported in Table 2, for the entire set of sensitive runs, agent movement is more likely to generate stable cooperation (83% versus 39%), is less likely to produce punctuated equilibria (11% versus 46%), and is less likely to produce no cooperation (6% versus 15%) than when agents cannot move. In addition, the pattern between the type of movement and the three types of outcomes holds within each of the five types of sensitivity variations.¹⁴

DISCUSSION

Movement


The simulation results indicate that the introduction of movement (either global or local) increases substantially the

likelihood that stable, highly cooperative worlds are generated: 80% of the time for simulations with movement versus 37% without movement. When agents cannot move, there is a high probability that they are either isolated and experience no interaction or become trapped in unprofitable relationships and are unable to find more cooperative partners. In either case, the agent's energy is drained to 0, and it is eliminated. Movement makes it more likely that agents will sever unproductive relationships and will more readily reestablish productive relationships.¹⁵

Agent movement appears to be an effective mechanism to help build and maintain large networks of cooperative agents because it appears to increase the stability of the cooperative networks. Cooperative networks can be fragile. Networks of cooperative agents fall apart when they become fragmented and are separated into smaller parts. Fragmentation typically occurs when one or more members of the network are "invaded" by noncooperative agents. Interactions with the noncooperative agents can produce energy losses that can lead the member of the cooperative network to move and in some cases to be eliminated. This can allow noncooperative agents into the network and can lead to its fragmentation.

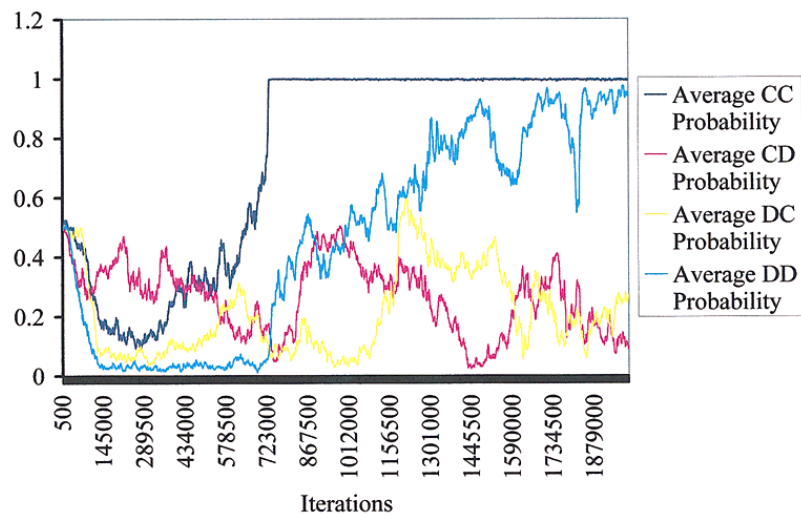
While movement helps generate high levels of cooperation and the formation and maintenance of cooperative clusters of agents, too much movement is very detrimental to the development of high levels of cooperation.¹⁶ As noted earlier, if agents were to move with certainty every time a nonpositive payoff were obtained, then far too much movement would occur, making the development of both cooperative clusters and a cooperative world impossible. Indeed, if the probability of the movement of agents with nonpositive payoffs is adjusted so that, on average, agents have a 10% chance of moving rather than the original 1% chance, then cooperation is never achieved, regardless of the type of movement and the mutation magnitudes. Increasing the probability of movement, even when agents have

FIGURE 5



Much larger nodes of cooperation populated by Grim agents evolve. Turn number 308, 544 of the same simulation run reported in Figure 4.

FIGURE 6



Strategy mix with local movement.

nonpositive payoffs, makes the generation of cooperative clusters impossible. Cooperative clusters can only develop and be maintained if agents have a relatively low probability of moving even if they receive a nonpositive payoff. Movement is only beneficial if an agent really is located in a position that generates a significant string of nonpositive payoffs.

Cautious Cooperation: The "Grim Transition"

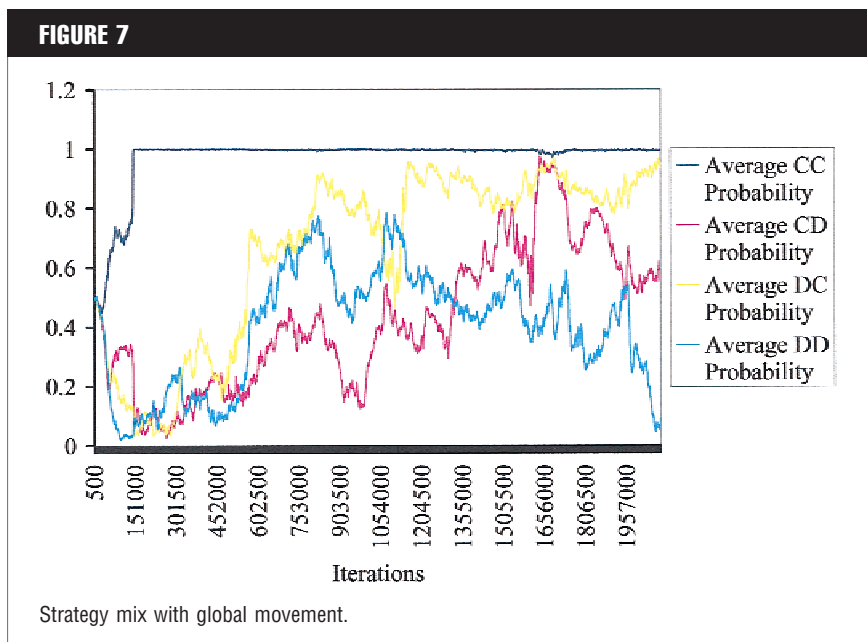
When there is a transition from a non-cooperative to a highly cooperative world, a version of the Grim strategy dominates numerically and appears to be essential for this transition. Indeed, the Grim strategy dominates all (75) transitions to cooperation for the originally configured results and all transitions to cooperation for the sensitivity analysis simulation runs. In the transi-

When there is a transition from a noncooperative to a highly cooperative world, a version of the Grim strategy dominates numerically and appears to be essential for this transition.

tion period, cooperative nodes of a "stingy" form of Grim (e.g., [0.85, 0.05, 0.05, 0.05]) begin to emerge. This is illustrated in Figure 4, a screen capture of the spatial grid at turn 306,202 of a simulation with local movement. Cooperative nodes of Grim strategies (red shaded) begin to emerge and are attacked by highly noncooperative strategies (darker shaded).¹⁷ As Grim strategies come to dominate the world, the strategies themselves transform and become somewhat more generous. The probability of cooperation after a CC interaction approaches 1.0, and the other probabilities increase to between 0 and .3 overall (see the screen capture of the spatial grid at turn 308,544 in Figure 5).

"Communities" of Cooperation

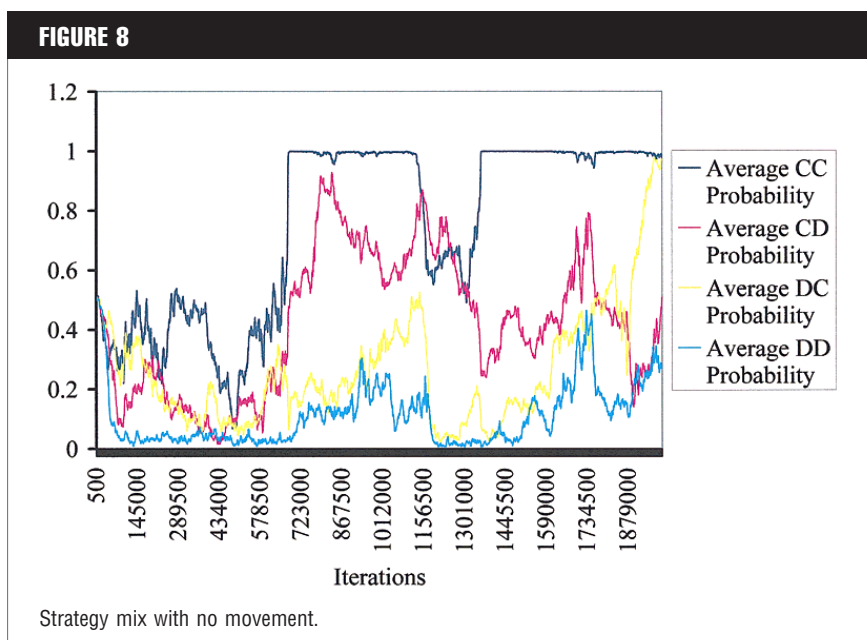
The emergence of small, relatively stable nodes or networks of agents characterize every transition from uncooperative to cooperative worlds. The



importance of the formation of networks is strengthened by the fact that when the probability of agent movement is made higher, cooperation does not occur and no networks or clusters form even when Grim agents dominate. The networks of cooperative agents form what could be considered to be primitive social networks or the basic foundations of "communities."¹⁸

Once cooperative communities become large and dense, agents employing all-D strategies may cause some

movement around the edges of the community, but these agents are unable to break up the cooperative nodes. Agents on the periphery of the community may be driven to move by the negative interactions with all-D strategies, but all-D agents find such interactions unprofitable as well and will also soon move away. The agents in the heart of the community remain untouched, continue to profit from the CC interactions, and are in a situation to reproduce at a faster and more consistent



The importance of the formation of networks is strengthened by the fact that when the probability of agent movement is made higher, cooperation does not occur and no networks or clusters form even when Grim agents dominate. The networks of cooperative agents form what could be considered to be primitive social networks or the basic foundations of "communities."

rate than other agents. This structure of cooperation in which an external group shields an internal core is similar to that reported by Lomborg [17].

Changes in Dominant Cooperative Strategy

Once cooperative networks are established and interactions are almost entirely characterized by mutual cooperation, Grim-like strategies often lose their numerical dominance. This occurs because highly noncooperative strategies such as all-D usually have been eliminated, thereby diminishing the danger of exploitation by these strategies and allowing room for strategies that are a little more forgiving than Grim.¹⁹ These strategies are able to infiltrate the Grim communities by having, like Grim, a high CC probability. No one strategy clearly emerges. Indeed, the emergence of strategies that are Pavlov-like (see Figure 6) and TFT-like (see Figure 7) and of a strategy that is the combination of Grim coupled with high CD probabilities (i.e., [0.99, 0.87, 0.03, 0.08]) are all about equally likely to emerge, probably because it is the CC probability that is most important. The profile of strategies that come to dominate or emerge appears to be quite similar across the three types of movement.

Collapse of Cooperation and Cooperative Communities

Large clusters of cooperative agents sometimes do collapse, and this usually occurs because the networks of Grim-like agents evolve into networks made up of agents with strategies that are "too nice."²⁰ The typical pattern that leads to collapse is the upward drift in the probability of cooperating following a CD outcome.²¹ In these cases, average CD probabilities drift up to over 0.8 (Figure

8). These agents, having a strategy mix such as [0.99, 0.85, 0.10, 0.07], can easily be taken advantage of, allowing nasty strategies such as all-D or mean TFT to invade. This leads to a long and slow (up to 100,000 iterations) decline to near total defection. Lomborg [17] and Nowak and Sigmund [10] discovered a similar pattern: The generation of too many nice strategies makes successful exploitation too easy, leading to the col-

lapse of cooperation. Lomborg [17] likened this situation to the "classic predicament of societies going 'soft.' "

Summary

The artificial world that we have constructed is based on a spatial IPD that places agents with limited capabilities (limited memory and cognition, no means of verbal communication, and no shared history or beliefs) into a harsh world (an anarchic, competitive world where the environment consistently drains their limited resources). Yet, even in such a world, cooperative networks usually eventually emerge, and some, typically associated with the ability of agents to move, can be quite stable, leading to long-term high levels of cooperation.

NOTES

1. Agent movement and interaction in spatial neighborhoods affect partner selection, a feature shown to be relevant to the emergence of cooperation [3–7].
2. As with most agent-based modeling approaches [8], we begin with an explicit set of assumptions about some phenomena and use them to generate simulated data. We search for patterns in the simulated data, particularly the large-scale effects from the interactions of locally interacting agents or what are often referred to as "emergent properties" of the system (see [9]).
3. We make this restriction because it makes our analysis comparable to a considerable amount of previous work making the same restriction.
4. Following Nowak and Sigmund [10], the combination of seeding the initial strategy pool with one strategy and the application of a genetic algorithm allow for complex emergent behavior.
5. See Nowak et al. [11] for an argument about the importance of considering social dilemmas as occurring in social space and for a discussion of various geometries to represent social space.
6. Minnis [12] suggests that human populations employ a number of coping strategies to exploit spatial variability.
7. Dugatkin and Wilson [4], Nowak and May [13], Oliphant [14], Lindgren [15], Lindgren and Nordahl [16], and Lomborg [17] have introduced a spatial component into the IPD or PD environment.
8. A number of scholars have incorporated noncompulsory play into the IPD (see [18–21]).
9. Mobility is one of a number of important coping strategies or buffer mechanisms that human populations use to exploit the favorable aspects of temporal and spatial variability to mitigate the risk of scarcity [12, 22].
10. Our perspective on the relevance of an environmental carrying capacity is based on the work of Hardin [23,24] and Clayton and Radcliffe [25], who argue that those human populations do have cultural carrying capacities. Note as well that our cost of survival mechanism is similar to the global metabolic rate of Epstein [26].
11. When innovation is characterized as a process whereby a very small percentage of agents randomly vary their current strategy (e.g., [17]), then it has properties similar to generational change via mutation.
12. Each simulation was run for 2 million iterations, a period of time sufficiently long for the simulations to stabilize and either to reach a highly cooperative state or to remain in a conflictual state, and also to observe the collapse of high levels of cooperation.
13. All 90 experienced an initial and rapid decline to near universal defection, and it takes a considerable length of time before a transition to high levels of cooperation occurs. This is primarily due to the fact that simulations begin in a relatively noncooperative world populated by a set of random strategy agents. When the magnitude of change in strategies created by mutations is small, it is very difficult to generate moderately or highly cooperative agents, and the simulation is quite likely to remain mired in defection for a long period of time.

14. The results for the lower mutation rate are somewhat different from the main results reported earlier. As anticipated, cooperation is more difficult to achieve with a lower mutation rate. But this is the case only when agents cannot move. Stable cooperation is still achieved 80% of the time with movement and only 60% without movement. Also, as anticipated, when mutation magnitudes are increased, cooperative worlds are much more likely to be generated across all three types of movement and high levels of cooperation are reached consistently more quickly.
15. We anticipated that local movement would be more effective at generating and maintaining cooperation than global movement because movement to an adjacent location allows agents to stay in the same geographic locale, making it more likely that an agent can retain some of the same beneficial relationships that it was involved in previously; "membership" in a cooperative cluster can be preserved. However, local movement provides a slightly greater likelihood of achieving a highly cooperative world than global movement, but such worlds are slightly more likely to collapse.
16. This result is consistent with the work of Hutson and Vickers [6] and Ferriere and Michod [7]. They also discovered, by numerical and analytical means, respectively, that while movement helps cooperative or altruistic strategies invade a population of defectors, too much movement is detrimental.
17. It is Grim, rather than TFT, as Nowak and Sigmund [10] observed, that appears to succeed in invading the world of agents that are essentially almost pure all-D, because, in addition to being provokable, retaliatory, and nice, they are only slightly forgiving (DC probabilities range from 0.05 to 0.3).
18. These networks are similar to what Taylor [27] refers to as egalitarian anarchic communities because they exhibit a limited form of reciprocity and direct and many-sided relationships. However, they lack Taylor's third characteristic of community, beliefs and values in common.
19. See Taylor [27] for a discussion of the importance of the norm of reciprocity in communities.
20. This should not be surprising because Lorberbaum [28] demonstrated that no strategy is an evolutionarily stable strategy in the infinitely repeated prisoner's dilemma.
21. The evolution of strategies with high CD probabilities in both highly cooperative and highly defection-dominated worlds appears to be a function of evolutionary drift.

REFERENCES

1. Smith, J.M. *Evolution and the Theory of Games*; Cambridge University Press: Cambridge, 1982.
2. Axelrod, R. *The Evolution of Cooperation*; Basic Books: New York, 1984.
3. Peck, J.R. *J Theor Biol* 1993, 162, 195–228.
4. Dugatkin, L.A.; Wilson, D.S. *The American Naturalist* 1991, 138, 687–700.
5. Wilson, D.S.; Pollock, G.B.; Dugatkin, L.A. *Evolutionary Ecology* 1992, 6, 331–341.
6. Hutson, V.C.L.; Vickers, G.T. *Philos Trans R Soc London, Ser B* 1995, 348, 393–404.
7. Ferriere, R.; Michod, R.E. *The American Naturalist* 1996, 147, 692–717.
8. Axelrod, R. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*; Princeton University Press: Princeton, NJ, 1997.
9. Epstein, J.M.; Axtell, R. *Growing Artificial Societies: Social Science from the Bottom Up*; The Brookings Institution: Washington, D.C., 1996.
10. Nowak, M.; Sigmund, K. *Nature* 1993, 364, 56–58.
11. Nowak, A.; Latane, B.; Lewenstein, M. In *Social Dilemmas and Cooperation*; Schulz, U.; Albers, W.; Mueller, U., Eds.; Springer-Verlag: Berlin, 1994; pp 269–289.
12. Minnis, P.E. In *Evolving Complexity and Environmental Risk in the Prehistoric Southwest*; Tainter, J.A.; Tainter, B.B., Eds.; Addison Wesley: Redding, MA, 1996; pp 57–78.
13. Nowak, M.; May, R. *Nature* 1992, 354, 826–829.
14. Oliphant, M. In *Artificial Life IV*; Brook, R.; Maes, P., Eds.; MIT Press: Cambridge, MA, 1994.
15. Lindgren, K. In *Artificial Life II*; Langton, C.G.; Farmer, J.D.; Rasmussen, S.; Taylor, C., Eds.; Addison-Wesley: Redding, MA, 1991; pp 295–312.
16. Lindgren, K.; Nordahl, M. In *Artificial Life: An Overview*; Langton, C.G., Ed.; MIT Press: Cambridge, MA, 1995; pp 15–37.
17. Lomborg, B. *Am Soc Rev* 1996, 61, 278–307.
18. Kitcher, P. *J Philos* 1993, 497–516.
19. Batali, J.; Kitcher, P. In *Proceedings of the Alife IV Workshop*, 1994.
20. Stanley, E.A.; Ashlock, D.; Tesfatsion, L. In *Artificial Life III; Proceedings Volume 17, SFI Studies in the Sciences of Complexity*; Addison-Wesley: Redding MA, 1994; 131–175.
21. Orbell, J.M.; Dawes, R.M. *Am Soc Rev* 1993, 58, 787–800.
22. Halstead, P.; O'Shea, J. In *Bad Year Economics: Cultural Responses to Risk and Uncertainty*; Halstead, P.; O'Shea, J., Eds.; Cambridge University Press: Cambridge, UK, 1989; pp 1–5.
23. Hardin, G. In *Managing the Commons*; Hardin, G.; Baden, J., Eds.; W.H. Freeman and Company: San Francisco, 1977; pp 112–125.
24. Hardin, G. In *Ecological Economics: The Science and Management of Sustainability*; Costanza, R., Ed.; Columbia University Press: New York, 1991; pp 47–57.
25. Clayton, A.; Radcliffe, N. *Sustainability: A Systems Approach*; Earthscan: London, 1996.
26. Epstein, J.M. *Complexity* 1998, 4, 36–48.
27. Taylor, M. *Community, Anarchy and Liberty*; Cambridge University Press: Cambridge, 1982.
28. Lorberbaum, J. *J Theor Biol* 1994, 168, 117–130.