

Zones of Cooperation
in
Demographic Prisoner's Dilemma

Joshua M. Epstein

ABSTRACT

The emergence of cooperation in Prisoner's Dilemma (PD) games is generally assumed to require repeated play (and strategies such as Tit-For-Tat, involving memory of previous interactions) or features ("tags") permitting cooperators and defectors to distinguish one another. In the Demographic Prisoner's Dilemma, neither assumption is made: agents with finite vision move to random sites on a lattice and play a fixed culturally-inherited zero-memory strategy of cooperate (C) or defect (D) against neighbors. Agents are indistinguishable to one another--they are "tagless". Positive payoffs accrue to agents playing C against C, or D against C. Negative payoffs accrue to agents playing C against D, or D against D. Payoffs accumulate. If accumulated payoffs exceed some threshold, agents clone offspring of the same strategy onto neighboring sites and continue play. If accumulated payoffs are negative, agents die and are removed. *Spatial zones of cooperation emerge* .

The Prisoner's Dilemma

The Prisoner's Dilemma game raises the problem of cooperation in a stark form. Two strategies are available to each player: cooperate or defect. The payoff to mutual cooperation (R) exceeds the payoff to mutual defection (P). But the highest payoff (T) goes to one who defects against a cooperator, while the lowest payoff (S) goes to one who cooperates against a defector. The letters T, R, P and S are used to denote the Temptation to defect, the Reward for mutual cooperation, the Punishment for mutual defection, and the Sucker's payoff accruing to a sole cooperator. With $T > R > P > S$ the Prisoner's Dilemma payoff matrix is

	C	D
C	(R,R)	(S,T)
D	(T,S)	(P,P)

[1]

In a one-shot game, the dominant strategy is D for both players. Rationality yields an outcome that, from the individual's perspective, is suboptimal. Indeed, the Prisoner's Dilemma is seen as posing a fundamental problem in social science (and biology): *How can cooperation evolve in populations whose bilateral interactions are governed by the PD?*

Repeated Games

One important line of attack has involved repeated play. This is a bit of a misnomer--it really means repeated play *with memory*. The distinction is crucial for this research. If A plays B repeatedly with no recollection of previous engagements (as in the Demographic Game below) that is *not* a repeated game. A "strategy" in a repeated game is a rule that specifies a player's behavior given some history of interactions against the opponent in question. For example, Anatol Rapoport's Tit-For-Tat (hereafter TFT) strategy is : cooperate on the first move; thereafter, adopt whatever strategy your opponent played in your last interaction. TFT is a *memory one* strategy, and will sustain cooperation[1]. The iterated PD must be of indeterminate length or there is regress to pure defection. An extensive literature has developed cooperation-supporting strategies for the repeated PD.[2] A strong trend in this research has been toward long memory and what, for humans, would be high computational loads[3].

Evolutionary Games

At the other end of the cognitive requirements spectrum is classical Evolutionary Game Theory[4]. Here, agents do not consciously optimize over strategic alternatives. Rather, they inherit a fixed strategy (a phenotype) and then replicate depending on that strategy's payoff (fitness). One elegant feature of these *replicator dynamics* is that their evolutionarily stable strategies (ESSs)[5] can correspond to the strategies that would be adopted by fully-informed rational players of the game. For the Prisoner's Dilemma specifically, the classical replicator dynamics (see below) lead straight to a world of pure defection--precisely the dominant strategy in one-shot rational play. In a social science context, the basic evolutionary reasoning is nicely set forth by Frank (1993; 1994).

His exposition begins with the assumption "that everyone in the population is one of two types--cooperator or defector. A cooperator is someone who, possibly through intensive cultural conditioning, has developed a heritable capacity to experience a moral sentiment that predisposes him to cooperate. A defector is someone who either lacks this capacity or has failed to develop it." [6] From a modeling standpoint, then, individuals are "hard-wired" to execute a fixed strategy. Now, Frank continues, "suppose, for argument's sake, that cooperators and defectors look exactly alike, thus making it impossible to distinguish the two types. In this hypothetical ecology...individuals will pair at random...The expected payoffs to both defectors and cooperators therefore depend on the likelihood of pairing with a cooperator, which in turn depends on the proportion of cooperators in the population" [7]. If cooperators comprise a fraction c of the population and individuals are paired randomly, then the probability of a given cooperator being paired with another cooperator is c ; the probability of her being randomly paired with a defector is $(1-c)$. Expected payoffs in the two cases are $E[C]=cR+(1-c)S$ and $E[D]=cT+(1-c)P$. For expository purposes, we will assume the payoffs below [8]

	C	D
C	(5, 5)	(-6, 6)
D	(6, -6)	(-5, -5)

[2]

Then, we have average returns of $E[D]=-5+11c$ for defectors and $E[C]=-6+11c$ for cooperators. For all c in $[0,1]$, $E[D]>E[C]$ and Frank's summary claims apply: "Since defectors always receive a higher payoff here, their share of the population will grow over time. Cooperators, even if they make up almost all of the population to begin with, are thus destined for extinction. *When cooperators and defectors look alike, genuine cooperation cannot emerge.*" [9] (emphasis added)

The same conclusion follows in the closely related classical Replicator Dynamics (Maynard-Smith 1982; Hofbauer and Sigmund 1988; Weibull 1995; Samuelson 1997).[10] Here, the relative frequency z_i of a strategy i grows in accordance with

$$\frac{dz_i}{dt} = z_i [(Return\ to\ Pure\ i) - (Average\ Return)].$$

[3]

For a symmetric payoff matrix A, this becomes

$$\frac{dz_i}{dt} = z_i[(Az)_i - z^T Az], \text{ where, for } n \text{ strategies, } z = (z_1, \dots, z_n).$$

[4]

For our two-strategy Prisoner's Dilemma, this z vector is simply (x, 1-x), where $x \in [0, 1]$ is the relative frequency of defection. For the symmetric matrix [1] above, A takes the form

$$A = \begin{pmatrix} P & T \\ S & R \end{pmatrix}$$

and, expanding [4], the dynamics are given by

$$\frac{dx}{dt} = (T - R)x + (P + 2R - S - 2T)x^2 + (T + S - P - R)x^3,$$


a cubic whose equilibria are $x=0$, $x=1$, and $x=(R-T)/(R-T+P-S)$. For our payoff orderings, $x=0$ (defection frequency zero) has eigenvalue $T-R > 0$, making it unstable, while $x=1$ (universal defection) has eigenvalue $S-P < 0$ making it asymptotically stable. With the numerical assumptions from [2], A becomes

$$A = \begin{pmatrix} 5 & -6 \\ 6 & -5 \end{pmatrix}$$

[5]

With this A matrix, the phase diagram for the replicator dynamics [4] on [0,1] are given in figure 1.

Figure 1. Phase Diagram for PD Replicator Dynamics



Just as in Frank's example, even the slightest perturbation from pure cooperation ($x=0$) results ultimately in pure

defection ($x=1$).

[11]

Demographic Games

I wish to extend this evolutionary literature with a class of agent-based models that I call Demographic Games. In the Demographic Prisoner's Dilemma, (1) agents inherit a fixed strategy of cooperate or defect [12] and (2) agents are indistinguishable to one another, just as in the models above. Yet, *cooperation can emerge and endure*.

"Demographic Games" seems an appropriate name for this class of models because they involve spatial, evolutionary, and population dynamics. The space on which agents interact is a 30 by 30 lattice of sites. Periodic boundary conditions obtain. Visually, agents who exit the lattice on the right/bottom reenter from the left/top. [13] Agents move around this space [14], interact with Von Neumann neighbors, and have offspring. Each agent is an object whose main attributes are vision, wealth, age, and strategy. Vision is the distance an agent can see, looking north, south, east, or west. In the evolutions presented below, vision is one. Agents are born with a strategy of cooperate (C) or defect (D). The agents' sole rule of behavior is as follows: *Choose a random site within your vision; go there and play your strategy against a random neighbor*. The agent and the neighbor then receive payoffs from some game matrix, here the Prisoner's Dilemma bimatrix [1], where $T > R > P > S$. For the Demographic PD, we introduce negative payoffs [15]. Specifically, we require that $T > R > 0 > P > S$, as in [2], which we shall use here.


Payoffs (wealths) accumulate. Since our payoff matrix has negative entries, [16] an agent's accumulated wealth may go negative. In that event the agent "dies"--it is removed from play and its wealth disappears. By the same token, if an agent's accumulated wealth exceeds some positive threshold and there is an unoccupied site within the agent's vision, the agent has an offspring, who begins life on one of these sites with a nominal initial endowment subtracted from the parent's wealth. [17] Here (with vision one) these are neighboring sites. Progeny are born with a fixed strategy; in the zero-mutation case, it is the parent's strategy. [18]

Perhaps it is worth emphasizing that, in adopting this assumption of a fixed agent strategy, we are not claiming that human strategies are literally hard-wired genetically. Rather, for modelling purposes, we are assuming that they are culturally transmitted from parents to children--vertically transmitted--with high fidelity, like certain religious or ethnic affiliations, tastes, and native tongues. ¹⁹ Below we consider the effect of degradation (mutation) in this vertical transmission fidelity. In a more elaborate model, horizontal (intragenerational) transmission would also be included. [20] However, here we follow Martin Shubik's injunction: "start with radical simplification. . .do not reject the simplest models because they are *a priori* too simple. Reject them when a quick investigation shows that the phenomenon of interest to you cannot appear at this level of simplicity." [21] The phenomenon of interest here is persistent cooperation, and we demonstrate that an extremely simple spatial model with vertical transmission suffices to generate it.

Run 1. No Maximum Age, Zero Mutation

For our first run of the model, we impose no upper bound on agent lifespans. Initially, one hundred agents are assigned random fixed strategies (C or D) and random initial positions. The five panels of figure 2 show the spatial situation at selected times illustrative of the main points. Cooperators are black, defectors white. A time step represents one cycle through the agent list. The agents are processed serially (updated asynchronously) and the agent list is randomized after every cycle. [22]


Figure 2. An Evolution To Cooperative Dominance



Panel 1 gives the early situation, with a few random agents scattered about the lattice. As cooperators randomly encounter one another, the positive payoffs associated with their interactions accumulate and they "clone" cooperator offspring onto neighboring sites. Neighborhoods of cooperation are thus formed. The first of these have clearly taken shape by $t=15$, as seen in Panel 2. By $t=50$, a stable ratio of cooperators to defectors (approximately 5 to 1), has set in, as shown in Panel 3. And this slightly noisy equilibrium persists. Panels 4 and 5 give the spatial configuration at $t=100$ and $t=1000$. In stark contrast to the replicator dynamics picture, cooperators are not annihilated. Indeed, they endure and predominate.

The time series showing aggregate cooperator and defector population dynamics is given in figure 3.


Figure 3. Aggregate Population Time Series



Note: Cooperators dark, defectors light

A phase diagram of these same data is given in figure 4. The x and y coordinates of each point correspond respectively to the cooperator and defector populations at that time. We see direct attraction to a (dark) invariant region of phase space corresponding to the noisy equilibrium noted earlier.

Figure 4. Phase Diagram



The basic point, then, is that *cooperation can emerge and flourish in a population of tagless agents playing zero-memory fixed strategies of cooperate or defect in this demographic setting.*

Statistical Analysis

Moreover, this result is robust. As Axelrod (1997b) writes, "...the analysis of a single run can be misleading. In order to determine whether the conclusions from a given run are typical it is necessary to do several dozen simulation runs using identical parameters (using different random number seeds)...While it may be sufficient to describe detailed history from a single run, it is also necessary to do statistical analysis of a whole set of runs to determine whether the inferences drawn from the illustrative history are really well founded. The ability to do this is one major advantage of simulation: the researcher can rerun history to see whether the particular patterns observed in a single run are idiosyncratic or typical."

Table 1 gives the result of such a statistical analysis.^[23] The model was rerun n=30 times with the same initial conditions and parameter values, but with a different random seed each time. In each run, data was sampled at t= 500 (after the essentially steady state had set in).

Table 1 Statistics


<u>Measure</u>	<u>Cooperators</u>	<u>Defectors</u>
Range	(752, 806)	(93, 148)
Mean	779	121
Standard Deviation	15	15
95% CI for the mean	(773,784)	(115,126)

Cooperator levels ranged from 752 to 806, while Defector levels ranged from 93 to 148. Sample distributions for both populations are tightly clustered about their means, as reflected in the small standard deviations; 95% confidence intervals for the mean, in turn, are narrow.

Run 2. Maximum Age, Zero Mutation

In Run 1, agents could live forever. A cooperator living for a long time in the interior of a cooperative zone could amass vast *wealth*. An enormous number of encounters with defectors would then be needed to eliminate him. This is not entirely unrealistic: once a human becomes extremely wealthy, a great many small setbacks are required before he is ruined financially. But, just to explore whether our results depend sensitively on this assumption, let us impose a maximum lifetime of 100 on the agents^[24]. All other settings are as before. The time series for this case is shown in figure 5.

Figure 5. Population Time Series with Maximum Age



Note: Cooperators dark, defectors light

The time series for this realization is clearly more oscillatory than Run 1. Not surprisingly, the cross-sectional statistics for 30 runs (sampling again at $t=500$) also exhibit greater variability, and for the defectors a reduction in the mean, as displayed in Table 2.

Table 2 Statistics


<u>Measure</u>	<u>Cooperators</u>	<u>Defectors</u>
Range	(708,846)	(45, 160)
Mean	784	99
Standard Deviation	29	25
95% CI for the mean	(773,794)	(90,108)

Clearly, cooperation persists as before. Now, as noted, we see in figure 5 the emergence of some oscillatory dynamics. We can amplify these by reducing one of the entries in our payoff matrix.

Run 3. Mutual Cooperation Payoff Reduced from 5 to 2

Specifically, with mutation still at zero, and leaving all else fixed as before, let us begin to "dial down" the mutual cooperation payoff (the value of R), first from 5 to 2. A typical realization is shown in figure 6.

Figure 6. Population Time Series with R=2




Note: Cooperators dark, defectors light

Cooperators do worse than in figure 5, defectors do better, and the oscillations are now more evident.

Run 4. Mutual Cooperation Payoff Reduced to 1


This oscillatory dynamic is more pronounced if we reduce R further to 1. The panels of figure 7 summarize the spatial story. As before, cooperators are black, defectors white.

Figure 7. Evolution of Cooperative Zones



Panel 1 shows the initial agents in their random starting positions. By $t=40$, cooperative neighborhoods have begun to take shape, as shown in Panel 2. Cooperators dominate by $t=60$, as shown in Panel 3. However, the cooperative zone is bordered by defectors. These surrounding defectors gradually "eat away" at the cooperative region and by $t=160$ have nearly annihilated it, as shown in Panel 4. In time, the defectors will have few cooperators with whom to interact; most interactions are then with other defectors. But since these interactions carry negative payoff, the fractious defectors are no sooner the majority than they begin to kill one another off, making way for a resurgence of cooperation, evident by $t=240$, shown in Panel 5. This is a typical cycle. Another such cycle is shown in Panels 6, 7, and 8, corresponding to times 560, 625, and 700. A time series of this evolution is shown in figure 8.


Figure 8. Oscillatory Dynamics



Note: Cooperators dark, defectors light

Figure 9 offers a phase diagram of the same dynamics.


Figure 9. Phase Diagram



We see, essentially, a perturbed limit cycle. The dynamics are reminiscent of predator-prey cycles. And, indeed, defectors "feed" on cooperators in the sense that they require interactions with cooperators to accumulate positive payoff, quite as if they were predators and cooperators were prey. In panel 4 of figure 7, we clearly see a cooperative neighborhood ringed by defector parasites. In the simulation proper, it looks very much like a predatory white species "eating through" a black prey concentration. But defectors need cooperators for a less obvious reason, also illustrated in this run: *cooperators function to separate defectors from one another spatially*. Because defector-defector interactions carry negative payoff, defector neighborhoods seldom take shape, and are highly self-destructive--and hence ephemeral--when they do. Indeed, their self-destruction is what clears the way for the resurgence of cooperative zones.


I have not attempted to count cooperative zones over the course of this run. A rough index of cooperative blotchiness is a count of cooperators all of whose eight Moore neighbors are also cooperators; I call these *surrounded cooperators*. For example, the 14-cooperator configuration in figure 10 would count as 2 (circled) surrounded cooperators.

Figure 10. Two Surrounded Cooperators



The time series for this index is plotted in figure 11, and reflects the rise and fall of cooperative zones in Run 4.

Figure 11. Time Series of Surrounded Cooperators




In contrast to Runs 1 and 2, different random seeds *do* produce different outcomes for these *same* payoffs. While the oscillatory dynamic is robust to such variation, the oscillations are extreme enough that qualitatively different long-run outcomes are possible. In some cases, cooperators and defectors coexist. In others, cooperators inherit the earth, while in others they are the first to go extinct (followed by the defectors, since the latter's mutual interactions carry strictly negative payoffs). In those realizations where cooperators inherit the earth, they do so for an interesting and counterintuitive reason. With the higher R value (the payoff to mutual cooperation) of 5, the cooperators never "thin out" enough spatially to let the defectors annihilate each other, while with $R=1$ they may, paving the way for their own monopoly. In such cases, cooperators ultimately do better with a low payoff ($R=1$) than with a high one ($R=5$)! This counterintuitive result is reminiscent of altruistic behavior in which one generation (here the generation of cooperators that thins out) sacrifices itself for the long-run benefit of the species. But it is not robust.

Run 5. Maximum Age, Mutation

Let us now restore our initial value of $R=5$, the payoff to mutual cooperation, and explore the sensitivity of our basic result--cooperator persistence--to mutation, defined as the probability that an agent will have a strategy different from it's parent's (e.g., that a cooperator's offspring will be a defector).

In this spatial setting, cooperator mutation can introduce defector "invaders" into the very heart of cooperative zones. One might well imagine that these will then spread until cooperators are annihilated. In fact, cooperator persistence will withstand very high mutation rates. In figure 12, all else is as in Run 2 (finite lives with maximum age of 100), except that a 50 percent mutation rate is assumed.

Figure 12. Cooperation Withstands Mutation




Note: Cooperators dark, defectors light

While the dynamics are much more oscillatory than in the first case, and the cooperator-defector ratio closer to 1, the long-term persistence of cooperation is intact through 10 thousand cycles.

Simulation and Sufficiency

Simulation is a particularly direct tool when the aim is to establish that some set of micro assumptions is *sufficient to generate* a macro phenomenon of interest. [25] In this case, the macro phenomenon is the persistence of cooperation. A widely accepted view is that "when cooperators and defectors look alike, genuine cooperation cannot emerge." [26] In the simulations above, cooperators and defectors have no "tags"; they do "look alike" in the relevant sense that they are indistinguishable *to one another*. Yet, on certain micro assumptions, cooperation emerges and endures. This basic sufficiency result is robust and withstands the introduction of a maximum age and high mutation rates. Of course, there are variations it will not withstand. Specifically, if we replace space and local interactions with soup (equiprobable random agent pairings), then the system runs to pure defection despite negative payoffs, quite consistent with the related replicator dynamics reviewed earlier. Moreover, space alone is not sufficient to ensure cooperation; with all payoffs shifted up by 6, so that $T=12$, $R=11$, $P=1$, and $S=0$, the spatial system again converges to pure defection, as shown in figure 13.

Figure 13. Space and Non-Negative Payoffs



Note: Cooperators dark, defectors light

Negative Payoffs Unnecessary But Appealing

Negative payoffs (for P and S) are not necessary for cooperation to persist, or even monopolize [27]. However, with all payoffs non-negative, the only way for agents to die is through an exogenous "old age" (i.e., a random maximum age assigned when the agent is initialized). With negative payoffs, agents can die through an endogenous series of infelicitous social interactions. To "bottom-up" modellers, this endogeneity is an appealing feature of negative payoffs. However, it should be noted that, since defector-defector interactions then carry negative payoffs, a society composed purely of defectors is not viable in the long run. Of course, with mutation, defectors may persist--by spawning cooperators to "save them from themselves", as it were. Negative payoffs admit an alternative interpretation, which will lead to an interesting conjecture.

An Alternative Formulation of Negative payoffs [28]

At the outset, we assumed a payoff ordering:

$T > R > 0 > P > S$. [6]

Negative per period payoffs can accumulate until an agent's aggregate holding goes negative, at which point it dies. Equivalent mathematically, but quite different conceptually, is a set-up with all payoffs non-negative (as is perhaps more typical in the literature), but with some fixed "global" metabolic rate imposed on all agents after every interaction.[29] For example, if we add $|S|$ to all payoffs in [7] and denote the new values with asterisks, we have

$T^* > R^* > P^* > S^* = 0$. [7]

Now, taking as the unstarred values the entries in matrix A (6,5,-5,-6), we obtain (adding $|S|=6$ to each), the nonnegative payoffs 12,11,1,0. With metabolism set at zero, these are the *ex post* payoffs as well and we generate the immediately preceding run (figure 12), in which cooperators are annihilated. However, with the global metabolic rate hiked to 6, we recover, in effect, our initial payoffs and generate the run in which cooperators are *not* annihilated.

Necessity the Mother of Cooperation?

Now, one can think of this global metabolic rate as a kind of "environmental" or "selective" pressure. When the environment is hospitable--when selection pressure is weak--there is no particular advantage to cooperation, and defectors rule. However, as we turn up the selection pressure, we find that cooperators are more persistent. This makes some sense. After all, if there's an abundance of food lying around, we don't need to cooperate agriculturally. A conjecture, then, is that, in the demographic setting we have been exploring, *cooperation acquires selective advantage as selection pressures increase*. Or, more prosaically, necessity is the mother of cooperation.

Concluding Thoughts

Perhaps, in conclusion, it deserves emphasis that all the results reported above involve a particular space (a 30 x 30 torus), a small set of payoffs (and other numerical assumptions), and so on. The only claims that can be advanced definitively are that this specific complex of assumptions is *sufficient to generate cooperative persistence* on the time scales explored in the research.[30] As it happens, that is fairly notable given the existing literature. Obviously, it would be worthwhile to undertake a large-scale computational study sweeping the parameter space of the model, exploring the sensitivity of the main results and, if possible, to assess their generality mathematically. Meanwhile, the sufficiency results and basic dynamics were deemed notable enough, and the conjecture interesting enough, to warrant wider dissemination and discussion.

Java Implementation

The source code was written in C++ in the Metroworks Code Warrior development environment for the Macintosh. Readers interested in running the model under their own assumptions--payoffs, mutation rate, maximum lifetimes, and so forth--may do so in JAVA at <www.brook.edu/dynamics/wip/pd.htm>.

References

Ashlock, Dan, Mark D. Smucker, E. Ann Stanley and Leigh Tesfatsion. 1996. "Preferential partner selection in an evolutionary study of Prisoner's Dilemma." *BioSystems* 37: 99--125.

Axelrod, Robert. 1980a. "Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution* 24: 3--25.

Axelrod, Robert. 1980b. "More Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution* 24: 379--403.

Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, Robert and William D. Hamilton. 1981. "The Evolution of Cooperation." *Science* 211: 1390--96.

Axelrod, Robert. 1987. "The Evolution of Strategies in the Iterated Prisoner's Dilemma." In *Genetic Algorithms and Simulated Annealing*, edited by Lawrence Davis. London: Pitman; Los Altos, CA: Morgan Kaufman.

Axelrod, Robert. 1997a. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization." *Journal of Conflict Resolution* 41: 203--26.

Axelrod, Robert. 1997b. "Advancing the Art of Simulation in the Social Sciences." Santa Fe Institute Working Paper 97-05-048.

Berninghaus, Siegfried K., and Ulrich Schwalbe. 1996. "Evolution, interaction, and Nash equilibria." *Journal of Economic Behavior and Organization* 29 (1): 57--85.

Binmore, Ken. 1992. *Fun and Games: A Text on Game Theory*. Lexington, MA: D. C. Heath and Company.

Binmore, Ken and Larry Samuelson. 1990. "Evolutionary Stability in Repeated Games Played by Finite Automata." Center for Research on Economic and Social Theory, Working Paper 90-17, Department of Economics, University of Michigan, Ann Arbor, MI.

Boyd, R., and P. J. Richerson. 1985. *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Brams, Steven J. 1985. *Superpower Games: Applying Game Theory to Superpower Conflict*. New Haven: Yale University Press.

Brams, Steven J., and D. Marc Kilgour. 1988. *Game Theory and National Security*. New York: Basil Blackwell.

Casti, John L. 1992. *Reality Rules: II. Picturing the World in Mathematics*. New York: John Wiley & Sons, Inc.

Cavalli-Sforza, L. L., and M. W. Feldman. 1981. *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton, NJ: Princeton University Press.

Downs, George W., and David M. Rocke. 1990. *Tacit Bargaining, Arms Races, and Arms Control*. Ann Arbor, MI: University of Michigan Press.

Epstein, Joshua M., and Robert Axtell. 1996. *Growing Artificial Societies. Social Science from the Bottom Up*. Washington: Brookings/MIT Press.

Epstein, Joshua M., and Robert Axtell. 1997. "Artificial societies and generative social science." *Artificial Life and Robotics* 1 (1): 33--34.

Feldman, Barry, and Kai Nagel. 1993. "Lattice Games with Strategic Takeover." In *Lectures in Complex Systems*, edited by L. Nadel and D. Stein, SFI Studies in the Sciences of Complexity, Lect. Vol. V. Reading, MA: Addison Wesley.

Feldman, Richard M., and Ciriaco Valdez-Flores. 1996. *Applied Probability and Stochastic Processes*. Boston, MA: PWS Publishing Company.

Frank, Robert H. 1988. *Passions Within Reason: The Strategic Role of the Emotions*. New York: W. W. Norton and Company.

_____ 1994. *Microeconomics and Behavior*. New York: McGraw-Hill, Inc.

Guttman, Joel M. 1996. "Rational Actors, tit-for-tat types, and the evolution of cooperation." *Journal of Economic Behavior and Organization* 29 (1): 27--56.

Hofbauer, Josef and Karl Sigmund. 1988. *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press.

Holland, John H. 1995. *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Helix Books.

Huberman, B. and N. Glance. 1993. "Evolutionary games and computer simulations." *Proceedings of the National Academy of Sciences* 90 (16): 7715--18.

Law, Averill M., and W. David Kelton. 1991. *Simulation Modeling and Analysis* (2nd ed.). New York: McGraw-Hill.

Lindgren, Kristian, and Mats G. Nordahl. 1994. "Cooperation and Community Structure in Artificial Ecosystems." *Artificial Life* 1: 15--37.

Lomberg, Björn. 1996. "Nucleus and Shield: The Evolution of Social Structure in the Iterated Prisoner's Dilemma." *American Sociological Review* 61 (2): 278--307.

Macy, Michael. 1996. "Natural selection and social learning in prisoner's dilemma." *Sociological Methods and Research* 25 (1): 103--138.

Maynard Smith, John. 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

Miller, John H. 1996. "The coevolution of automata in the repeated prisoner's dilemma." *Journal of Economic Behavior and Organization* 29 (1): 87--112.

Nowak, Martin A., Robert M. May, and Karl Sigmund. 1995. "The Arithmetics of Mutual Help." *Scientific American* 272 (6): 76--81.

Oliphant, Michael. 1994. "Evolving Cooperation in the Non-Iterated Prisoner's Dilemma: The Importance of Spatial Organization." In *Artificial Life IV*, edited by Rodney A. Brook and Pattie Maes. Cambridge, MA: MIT Press.

Rapoport, Anatol. 1996. "Prisoner's Dilemma: Reflections and recollections." *Simulation and Gaming* 26 (4): 489--503.

Samuelson, Larry. 1997. *Evolutionary Games and Equilibrium Selection*. Cambridge, MA: MIT Press.

Schelling, T. C. 1978. *Micromotives and Macrobehavior*. New York: W. W. Norton and Company.

Shubik, Martin. 1996. "Simulations, Models and Simplicity." *Complexity* 2 (1): 60.

Tesfatsion, L. 1995. "A Trade Network Game with Endogenous Partner Selection." Economic Report 36. Iowa State University, Department of Economics.

Weibull, Jörgen W. 1995. *Evolutionary Game Theory*. Cambridge, MA: MIT Press.

Wu, Jianzhou and Robert Axelrod. 1995. "How to Cope with Noise in the Iterated Prisoner's Dilemma." *Journal of Conflict Resolution* 39 (1): 183--89.

Young, H. Peyton. 1993. "The Evolution of Conventions." *Econometrica* 61: 57--84.

Young, H. Peyton. and Dean Foster. 1991. "Cooperation in the Short and in the Long Run." *Games and Economic Behavior* 3: 145--156.

The author is a Senior Fellow in Economic Studies at the Brookings Institution and a Member of the Santa Fe Institute External Faculty. He thanks Steven McCarroll for research assistance and Henry Aaron, Robert Axelrod, Robert Axtell, Peyton Young, Daniel Shapero, and John Steinbruner for valuable discussions. For production assistance, he thanks Trisha Brandon.

[1] See Axelrod (1984).

[2] See Axelrod (1984), Binmore (1992), and Lomberg (1996).

[3] For example, see Lindgren and Nordahl (1994).

[4] See Maynard Smith (1982), and Weibull (1995)

[5] For a thorough development of this equilibrium concept, see Binmore (1992).

[6] Frank (1994: 262).

[7] Frank (1994: 263). In the model below, agents look indistinguishable to one another, but do not pair at random.

[8] Frank uses different numerical values. These do not affect Frank's argument and will be used for illustrative purposes throughout.

[9] Frank (1994: 265). This is one motivation for introducing features, "tags," allowing cooperators and defectors to distinguish one another. On tags, see Holland (1995). See also Tesfatsion (1995).

[10] For powerful extensions see Young (1993) and Young and Foster (1991).

[11] In a more deliberate development, one would distinguish the evolutionarily stable strategies from the asymptotically stable equilibria. However, for $n=2$, they are the same. See Binmore (1992: 428--9).

[12] Mutation is introduced below.

[13] Topologically, the space is a torus.

[14] They are not fixed Cellular Automaton sites, as in Nowak and May and Sigmund (1995), Feldman and Nagel (1993), Lindgren and Nordahl(1994).

[15] A discussion of this assumption and an equivalent formulation involving non-negative payoffs are presented below.

[16] See Rapoport (1966).

[17] By contrast with this use of a simple threshold, the standard replicator dynamic assumes that a phenotype's (i.e. a strategy's) frequency grows as the difference between the strategy's fitness and the population's average payoff. For the runs discussed here, the threshold is 10 units of accumulated payoff and offspring inherit an initial endowment of 6 units. An agent's initial age is a random integer between 1 and the maximum age.

[18] Oliphant (1995) also studies the evolution of cooperation in a spatial model of the non-iterated prisoners' dilemma. That interesting model differs from the demographic game in fundamental respects, including the use of a genetic algorithm, a fitness-biased probability of individual play, fixed population size (no population dynamics), and gaussian interaction probabilities on a different (one dimensional) space.

¹⁹ On cultural transmission, see Cavalli-Sforza and Feldman (1981) and Boyd and Richerson (1985).

[20] For models including horizontal transmission, see Epstein and Axtell (1996) and Axelrod (1997a).

[21] Shubik (1996)

[22] On the importance of asynchronous updating in avoiding artifacts, see Huberman and Glance (1993).

[23] On the statistical analysis of simulation output, see Law and Kelton (1991) and Feldman and Valdez-Flores (1996).

[24] This actually ends up being slightly favorable to the defectors since the richest possible cooperator would then have 500 while the richest possible defector would have 600.

[25] On generative sufficiency, see Epstein and Axtell (1996: 20, 177) and Epstein and Axtell (1997). See also the pioneering work of Schelling (1978).

[26] Frank (1994: 265).

[27] For instance, with all payoffs hiked by ten (so that $T=16$, $R=11$, $P=5$, $S=4$) and the maximum lifetime reduced from 100 to 10 cycles, and using the same random seed, we again see an evolution to cooperative monopoly. I thank Henry Aaron for these settings.

[28] I thank Robert Axtell for this line of thought.

[29] By metabolic rate, I mean a fixed decrement to accumulated payoff per cycle.

[30] Moreover, computer runs, even out to $t=$ Avagadro's Number, may not prove anything about limiting behavior as .