The Evolution of Beliefs in a Finitely Repeated Game*

by

Sheryl Ball **
Department of Finance and Economics
Boston University
704 Commonwealth Avenue
Boston, MA 02215

and

Roy Gardner
Department of Economics
Ballantine Hall
Indiana University
Bloomington, IN 47405

*Financial support for this research was provided by National Science Foundation Grant #8643031 and the Vice-President for Research, Indiana University. We are deeply indebted to Robert E. Forsythe for access to his data and software.

**Current address: Department of Business Economics and Public Policy, School of Business, Indiana University, Bloomington, IN 47405.

Abstract

Consider an agent playing a series of games against opponents drawn from a given population composed of various player types. The choice of a strategy then depends critically on the agent's beliefs about the distribution of types present in the population. Beliefs evolve with playing experience, and if play converges to an equilibrium, so do beliefs. In this paper we model evolution according to the replicator dynamics drawn from evolutionary stability theory. We examine a finitely repeated coordination game played against two types of randomly drawn opponents, human and robot. By controlling the distribution of types and initial beliefs, we achieve experimental control over the evolution of beliefs. Our major result is that in a wide variety of designs and treatments human subject play does converge to an evolutionarily stable strategy.

1. Introduction

Classical game theory is predicated on the axiom of completely rational play at Nash Equilibrium. In the last 30 years, experimental investigations of games have shown substantial and persistent deviations from fully rational play (see for example Guth, Schmittberger and Schwarze 1982; Roth and Schoumaker 1983; Dyer, Kagel and Levin 1986; Ball, Bazerman and Carroll 1990). In other cases repeated play approaches and sometimes even converges to rational outcomes (Alger 1985; Cooper, DeJong, Forsythe and Ross 1990). The standard interpretation of the latter result is that the subjects have learned how to play rationally. Learning, however, can take a variety of forms.

The most explicit form of learning takes place when subjects, being goal-directed by payoffs, strive for and achieve a Nash Equilibrium strategy relative to the subject population, ie, a behavioral equilibrium. A variety of learning models have been developed for the study of cultural evolution, of which game playing is an example (Boyd and Richerson, 1985).

An implicit form of learning is characterized by the evolutionary hypothesis, that behavior dynamics driven by payoff improvement will select for evolutionarily stable strategies (ESS) which themselves are Nash Equilibria (Hofbauer and Sigmund 1988). This form of learning on an evolutionary time scale has has been remarkably successful in biological applications (Gardner, Morris and Nelson 1987; Gardner and Morris 1989). Evolutionary time is typically on a scale of 10^4 years, whereas laboratory experimental games take place on a scale of 10^{-3} years. Despite the fact that humans do not actually evolve in the laboratory setting, the dynamic equations (the replicator equations) which govern the evolution of animal behavior are equally appropriate for describing human behavior at the population level (Friedman and Rosenthal 1986; Miller and Andreoni 1990).

Our aim in this work is to study the adaptive process by which human beings approach a behavioral equilibrium in a laboratory setting. Our instrument for this study is a simple coordination game with two evolutionarily stable equilibria. Equilibrium selection, according to the evolutionary dynamics, depends solely on initial conditions. We achieve control of initial conditions by designing a decision task in which all human subjects are of the same behavior type. Holt (1990) and El-Gamal, McKelvey and Palfrey (1991) both present interesting repeated games which estimate types, including belief types. The novel feature of our approach is control of the distribution of initial beliefs.

The paper is organized as follows: The next section describes the dynamic system for the evolution of behavior in games, in particular, coordination games. In section 3 we present our experimental design and experimental methods. Section 4 presents results. In a wide variety of experimental designs and treatments, human subject play converges to an ESS.

2. Evolutionary Models

In this section we study the evolutionary dynamics applied to the coordination games depicted in table 1. There are two players in each game, both of which are the same type. Each player has a choice of whether to pick strategy A or B. There are three Nash Equilibria: the Pareto superior equilibrium occurs when both players choose strategy B. This equilibrium set containing a unique Pareto superior equilibrium, is reminiscent of the bargaining under bilateral monopoly example in Hansen and Samuelson (1988).

Since agents prefer to coordinate their strategy choices, the choice of A or B will depend on what they expect their opponent to do. Since an agent would prefer the outcome "both choose B" and she believes that her opponent also prefers this outcome, we should expect to see this equilibrium. If, on the other hand, an agent has reason to believe that her opponent prefers the outcome "both choose A", she should choose strategy A. If she is uncertain about what her opponent is likely to do, her assessment of the likely distribution of strategies A and B will determine her choice.

The evolutionary process which we study here is how beliefs about the choice of strategies being made in the population change with experience. The intuition behind these evolutionary dynamics is that if a strategy played by a certain player or type of player has an above-average payoff, then that strategy will be represented in the population to a greater extent in the next period. Following this intuition, one arrives at the following dynamic system:

$$\Delta x = x_i \, [F_i(x) - \mathbf{F}(x)] \quad i = 1, 2 \tag{1}$$

where

$$F_1(x) = a\,x_1 + b\,x_2 \tag{2}$$

$$F_2(x) = c\,x_1 + d\,x_2 \tag{3}$$

and

$$\mathbf{F}(x) = x_1 F_1(x) + x_2 F_2(x) \tag{4}$$

where x_1 denotes the percentage of players who choose strategy A and x_2 denotes the percentage of players who choose strategy B. The parameters satisfy the conditions $a>c$ and $d>b$. The reasoning behind equation 1 is as follows. Suppose a player chooses strategy A. Then his expected payoff is $F_1(x)$. If this is greater than the population average payoff $\mathbf{F}(x)$ then we should expect the proportion of individuals choosing strategy A to increase. Similarly, if $F_1(x)$ is less than $\mathbf{F}(x)$ we should expect the proportion of individuals choosing A to decrease. If $F_1(x)$ equals $\mathbf{F}(x)$ we should see no change.

An evolutionarily stable strategy is a behavior in a population which cannot be invaded by a mutant behavior. A Nash equilibrium is a dynamic equilibrium of (1). If $x^* = (x_1^*, x_2^*)$ is a Nash Equilibrium and x^* is a stable dynamic equilibrium of the system (1) then x^* is an ESS (Van Damme 1988). We now show that the only ESS of our design occur at $x^* = (1,0)$ or $x^* = (0,1)$.

The argument uses the Jacobian of the replicator dynamics. The Jacobian J(x) of (1) at population x is given by:

$$J(x) = \begin{bmatrix} x_1(1-x_1)(a-c) + (1-2x_1)[F_1(x) - F_2(x)] & x_1(1-x_1)(b-d) \\ x_2(1-x_2)(a-c) & x_2(1-x_2)(d-b) + (1-2x_2)[F_2(x) - F_1(x)] \end{bmatrix} \quad (5)$$

Evaluating the Jacobian (5) at $x^* = (1,0)$ one has the diagonal matrix $J(x^*)$.

$$J(x^*) = \begin{bmatrix} (c-a) & 0 \\ 0 & (c-a) \end{bmatrix}$$

The eigenvalues of $J(x^*)$ are the diagonal entries. For hyperbolic dynamic stability, all the eigenvalues must be negative. Since, by assumption, $a>c$ these eigenvalues are negative. This establishes (1,0) as an ESS. It is easy to show that (0,1) is also an ESS using the same technique. In that case the eigenvalues are (b-d), again negative.

The other candidate ESS is $[(b-d)/(b-a+c-d), (c-a)/(b-a+c-d)]$. At this equilibrium the Jacobian has zeros along the diagonal and positive entries off the diagonal; hence there is a positive eigenvalue and this Nash equilibrium is not evolutionarily stable. This discussion is summarized in the phase diagram of Figure 1, where the two boundary roots of the Δx_1 equation are locally stable while the intermediate root is not.

3. Experimental Design

We employed a total of 4 experimental games based on different values of the payoff parameters (see table 2). In each of the experimental games there is a dominant and a dominated pure strategy Nash equilibrium. By ESS considerations noted above, we hypothesize the following:

Hypothesis 1.1: The mixed strategy equilibrium will not be played in any game. In games 1 and 2, the payoff dominated equilibrium is also the maxi-min equilibrium as well as the risk-dominant equilibrium, hence should be selected according to Harsanyi-Selten (1988) equilibrium selection theory. The empirical results of Cooper, DeJong, Forsythe and Ross (1991), where subjects play a dominated maxi-min equilibrium when its payoffs are within 20% of those of the dominant equilibrium also point in the direction of:

Hypothesis 1.2: Play will evolve to the risk-dominant equilibrium in games 1 and 2.

In games 3 and 4 both equilibria are maxi-min; thus the previous hypothesis fails to select between them. However, by appealing to the principal of payoff dominance in Harsanyi-Selten, we expect:

Hypothesis 1.3: Play will evolve to the dominant equilibrium in games 3 and 4.

This concludes the discussion for all-human subject populations.

Experiments consisted of two stages. In the first stage, subjects participated in five or more periods in which they knew they were paired only with human subjects. This stage served 3 purposess. It provided a sequence of training periods in which subjects could familiarize themselves with the logistics of the experiment. It gave subjects evidence on the behavior of other human subjects. Finally, it allowed us to test hypotheses 1.1-1.3. Design of the second stage, in which we explore the evolution of beliefs, is discussed below.

Our principal aim in these experiments is to study how agents' beliefs about opponents' types predict which of two strategies he or she will employ in a population composed of two types of agents. Ideally, one would conduct experiments with two types of human subjects. Unfortuanately, this is hard to do because of the difficulty in identifying distinct human subject types. The two subject types in this experiment are humans and robots. Human subjects are free to choose either strategy, while robot subjects are programmed to play a fixed strategy which is chosen by the experimenter and made known to the subjects.

The choice of a strategy in these games depends on an agent's beliefs about what his opponent is likely to do. An agent who is playing a series of games against opponents drawn from a given population will have beliefs about the distribution of types present in the population. In order to study the evolution of beliefs, it is critical to induce initial beliefs about the proportion of robots in any experiment. Without control, comparison of behavior across individuals becomes very difficult. To insure that subjects began the experiment with uniform priors, they were told that, across experiments, the chance of hitting a human subject was fifty percent, although the chance in their particular experiment might be different from fifty percent. Although it is not possible to guarantee that subjects fully understood and believed this information, we are confident that we had a high degree of control because: 1) 50% is a focal point, and 2) 50% is the probability most easily understood. Of 60 subjects participating in the experiment, 55 reported 50% as their initial assessment of the probability of facing a robot opponent and no subject reported an initial belief below 40% or above 60%.

Although it is true that the prior probability of facing a robot in this entire set of experiments is 50%, in no individual experiment was this the case. The treatment "few robots" refers to 25% robots; "many robots," 75%. An equal number of experiments was conducted with each treatment. An additional treatment is self explanatory: "robots play A," robots play B." The entire stage two design is presented in table 3.

In stage 2, the stage 1 hypotheses must be amended to allow for the presence of robots in the population. The evolutionary hypothesis (1.1) continues to apply in this stage:

Hypothesis 2.1: Play will converge to an ESS in all games.

An immediate caution is necessary for the ESS hypothesis. To the extent that the "matching law" (Siegel, 1961) applies to our subjects, convergence to any Nash Equilibrium is jeopardized.

Given the multiplicity of ESS in these designs, a sharpening of the above hypothesis is necessary. Initial beliefs about the proportion of robot subjects and the behavior of human subjects are crucial to the subsequent evolution of play. Consider experiments 1 and 3, in both of which few robots play B. Human subjects start from a prior belief that the probability of hitting a B playing robot is 50% and that human subjects play A. The initial proportion of ail subjects playing A is sufficiently high (x_1 is sufficiently close to 1) that the ESS with $x_1 = 1$ is the local attractor (see figure 1). The sharpened ESS hypothesis says:

Hypothesis 2.2: Play will converge to the risk-dominant equilibrium in experiments 1 and 3.

Now consider experiments 2 and 4, in both of which many robots play B. Human subjects start from a prior belief that the probability of hitting a B playing robot is 50%

and that human subjects play A. The initial proportion of all subjects playing A is sufficiently low (x_1 is sufficiently close to 0) that the ESS with $x_1 = 0$ is the local attractor (see figure 1). The sharpened ESS hypothesis now says:

Hypothesis 2.3: Play will converge to the payoff-dominant equilibrium in experiments 2 and 4.

In experiments 5-8 the maxi-min effect is no longer present to prevent convergence to the payoff-dominant equilbrium in the all-human stage of the experiments. In experiment 8, even though there are many A playing robots, payoff dominance is so strong that we expect the payoff-dominant equilibrium to emerge. Thus one has the following:

Hypothesis 2.4: Play will converge to the payoff-dominant equilibrium in experiments 5, 7 and 8.

Finally, in experiment 6, with many robots playing A and much weaker payoff dominance we expect:

Hypothesis 2.5: Play will converge to the payoff-dominated equilibrium in experiment 6.

These sharpenings of the ESS hypothesis leads to a unique equilibrium selection for every experiment.

Since there are 4 games and 4 treatments, there were a total of 16 possible experiments which we could have conducted (see table 3). Of these, we deemed only 8 interesting enough to run. The common feature among the 8 which we did not run was that robot players would also be playing the same equilibrium we hypothesized human subjects to play in the stage 1 designs. Thus decision making would be taking place under constant condition of reinforcement. We ran two pilot studies in which subjects converged to the expected equilibrium within 5 periods in stage 2, which confirmed this decision.

Experiments were run by hand employing Indiana University undergraduate students recruited from introductory economics and business classes. Subjects were guaranteed $3.00 for showing up and earned an average of approximately $15.00 for one hour's participation. At the beginning of the experiment, instructions regarding both stages of the experiment were read aloud as the subjects read along.[1] Subjects were free to ask questions at any time during the experiment. A predetermined random pairing of players was used to implement the random matching assumption of ESS theory.

4. Results and Discussion

We begin our discussion of results with the all-human portion of our design, stage 1. Our overall hypothesis in this stage (1.1), that play would converge to an ESS is confirmed for all 4 games. The speed of convergence is often impressive (see figures 2-9). Game 1 converged to the maxi-min equilibrium in both experiments 1 and 2, as predicted by hypothesis 1.2. Game 2 converged more slowly to the dominant equilibrium, contrary to the prediction of hypothesis 1.2. This is rather surprising in light of the results of Cooper et al. (1990).

Games 3 and 4 converged quickly to the dominant equilibrium as predicted by hypothesis 1.3. In fact, virtually all subjects played B in the first period of stage one with these games.

We ran all 8 designs we deemed interesting in stage 2 (recall table 3). There are two treatments at work in stage 2; few or many robots, robots play A, robots play B. In some but not all cases, play did converge to an ESS (hypothesis 2.1).

[1] Instructions are available from the authors upon request.

Convergence, when it did occur, was slower than in stage 1, reflecting the added complexity of the decision task (see figures 2-9).

In particular, play converged to an ESS in experiments 3 and 6. This result is not particularly surprising in experiment 3, where robots choose the same strategy as humans did in stage 1 of the experiment. We reject hypothesis 2.2 in this case. In experiment 6 play converged to the payoff dominant equilibrium in stage 1. The addition of a large number of A-playing robots caused stage 2 play to converge to the payoff dominated equilibrium (hypothesis 2.5).

In the cases where play failed to converge, we observed that subjects continued to play the non-hypothesized strategy with a frequency equal to that of the robots. This resembles the result of Siegel (1963) where subjects, under varying conditions of reinforcement, attempt to match probabilities. Here the probabilities were matched at the population level, demonstrating that the human versus robot results in the literature extend to the population level. For example, in experiment 7 (figure 8), convergence to the equilibrium where all subjects play B was predicted, but 25% A-playing robots were present, and subjects continued to play A 25% of the time throughout stage 2. Similar statements could also be made for experiments 1 and 8.

Experiment 2 failed to converge or exhibit the probability matching phenomenon discussed above. The data suggest that given a sufficient amount of time behavior would have converged to the predicted equilibrium (hypothesis 2.5). Logistical considerations caused the experiment to be concluded before convergence could take place.

5. Discussion

The results of this paper are reminiscent of a large body of literature on critical mass and frequency dependence (Schelling, 1978). In our model the out of equilibrium dynamics are explicit rather than implicit. But the qualitative properties of the two-strategy model are the same; 1) Multiple Equilibria, 2) Possible Inefficiency, 3) Dynamic Instability of an Inefficient Equilibrium, and 4) Path Dependence. A recent survey of a general class of such adaptive processes is found in Arthur (1988).

Our results can be used to help understand how the proportions of competing new technologies change over time. Consider the market for IBM compatible or Apple Macintosh personal computers. One consideration for an individual who is about to purchase such a computer is the features. A second important consideration, since files can not be readily transferred between the two computer types, is the distribution of these two "computer use strategies" in the population of individuals with whom he interacts. Suppose that the payoff to coordinating on the Macintosh exceeds the payoff to coordinating on an IBM compatible. Since IBM compatible machines were available first this strategy was widely represented in the population. This means that even if similar software had been immediately available upon the introduction of the Macintosh, it was still have been unlikely that it would evolve as the predominant computer use strategy. Similar examples can be seen in the markets for video cassette recorders (VHS vs. Beta format) and in the development of mini-compact-disc players. Some of these cases may involve potentially large efficiency losses to the economy.

Table 1

Coordination Game

	A	B
A	a,a	b,c
B	c,b	d,d

Table 2

Stage 1 Experimental Design

Game 1		Game 2	
50,50	49,0	50,50	19,0
0,49	60,60	0,19	60,60

Game 3		Game 4	
40,40	0,0	10,10	0,0
0,0	60,60	0,0	60,60

Table 3

Stage 2 Experimental Design

	Game 1	Game 2	Game 3	Game 4
Few Robots Robots Play A			5	7
Few Robots Robots Play B	1	3		
Many Robots Robots Play A			6	8
Many Robots Robots Play B	2	4		

Figure 1

Phase Diagram

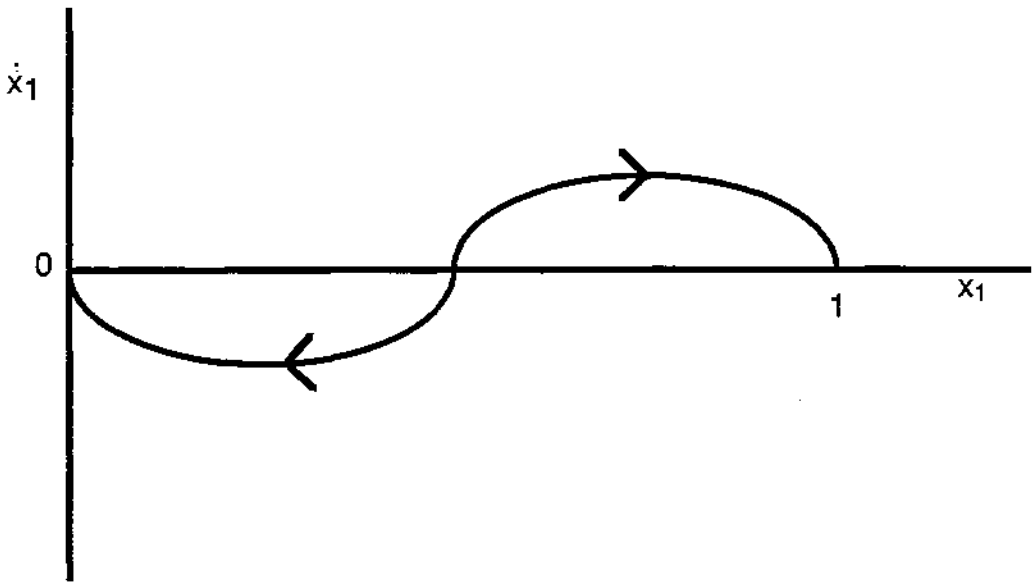

Figure 2

Experiment 1 Data

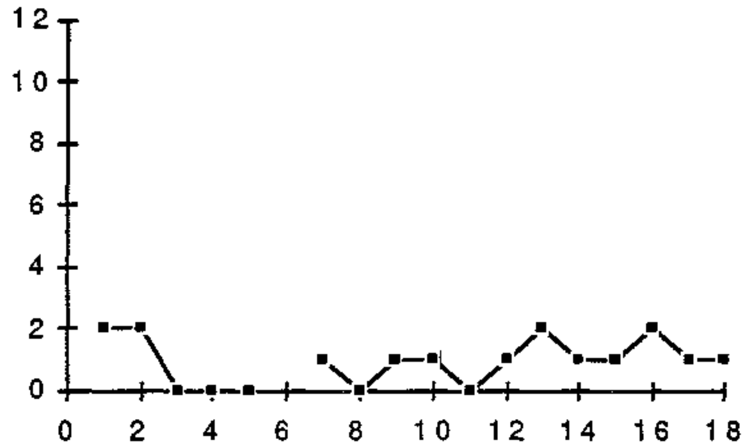

Figure 3

Experiment 2 Data

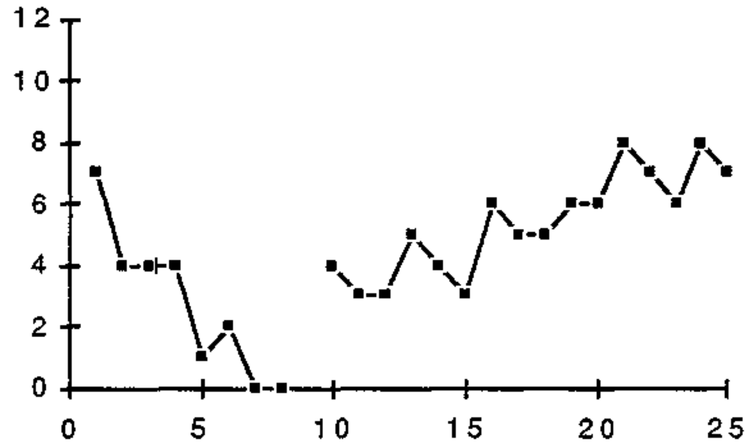

Figure 4

Experiment 3 Data

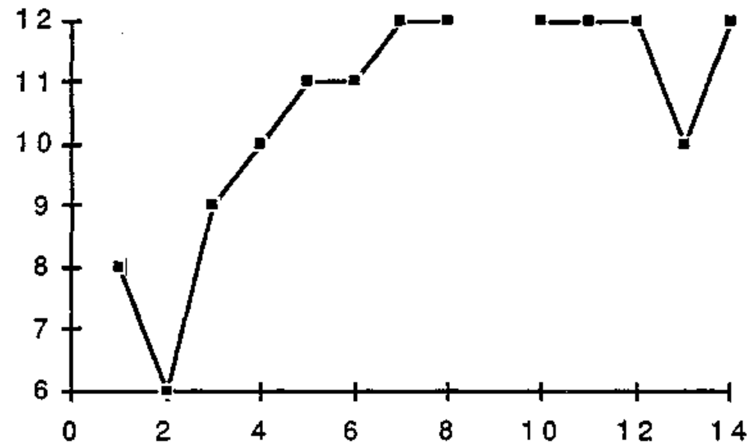

Figure 7

Experiment 6 Data

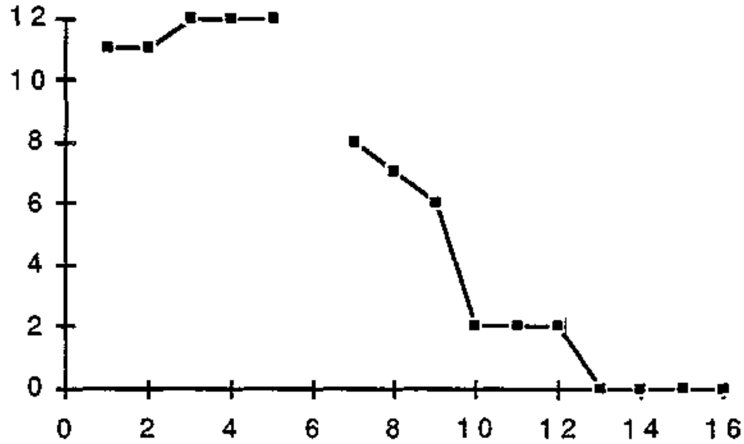

Figure 8

Experiment 7 Data

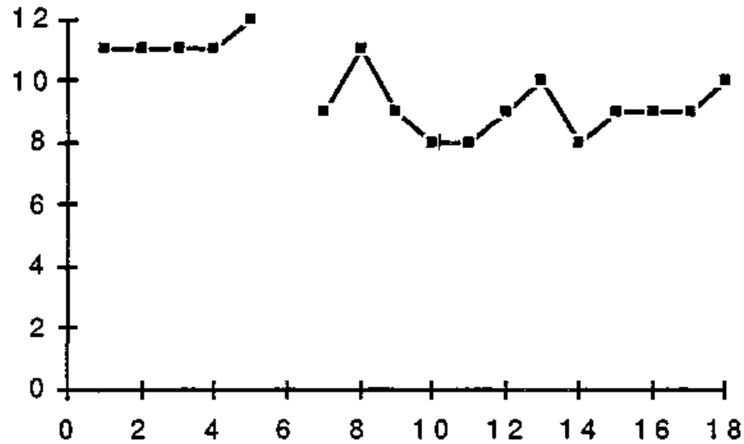

Figure 9

Experiment 8 Data

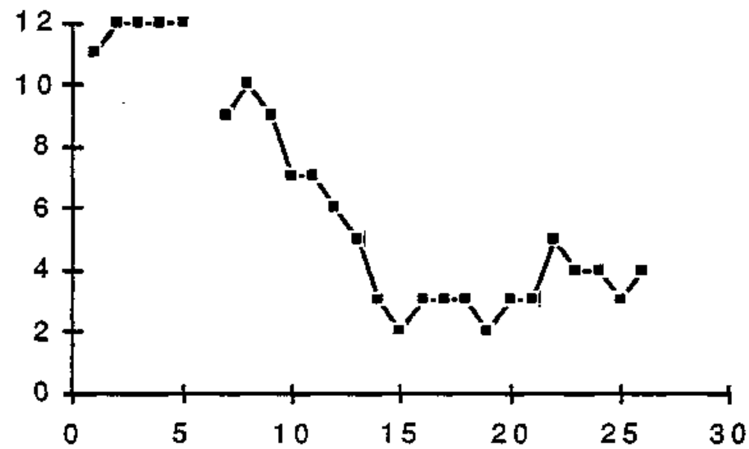

Bibliography

Alger, Dan. Investigating Oligopolies Within the Laboratory. Washington, D. C: Federal Trade Commission, 1985.

Arthur, Brian. "Self Reinforcing Mechanisms in Economics." in Phillip W. Anderson, Kenneth J. Arrow and David Pines, eds. The Economy as an Evolving Complex System: The Proceedings of the Evolutionary Paths of the Global Economy Workshop, Held September 1987 in Santa Fe, New Mexico, Redwood City: Addison Wesley, 1988.

Ball, Sheryl B., Max H. Bazerman and John S. Carroll. "An Evaluation of Learning in the Bilateral Winner's Curse." Organizational Behavior and Human Decision Processes, Vol. 48 (1990), 1-22.

Boyd, Robert and Peter J. Richerson. Culture and the Evolutionary Process. Chicago: University of Chicago Press, 1985.

Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross, "Selection Criteria in Coordination Games: Some Experimental Results," American Economic Review, Vol. 80 (1990), 218-233.

Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross, "Communication in Coordination Games." Quarterly Journal of Economics, forthcoming 1991.

Crawford, Vincent P., "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination." Working paper, September 1990.

Dyer, Douglas, John H. Kagel and Dan Levin. "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis." The Economic Journal, Vol. 99 (1989), 108-15.

El-Gamal, Mahmoud A., Richard D. McKelvey, and Thomas R. Palfrey. "A Bayesian Sequential Experimental Study of Learning in Games." California Institute of Technology, Social Science Working Paper 757, February 1991.

Friedman, James W. and Robert W. Rosenthal, "A Positive Approach to Non-Cooperative Games." Journal of Economic Behavior and Organization, Vol 7, (1986), pp. 235-251.

Gardner, Roy and Molly R. Morris, "The Evolutionary Stability of Bluffing in a Class of Extensive Form Games." in Selton Ed. Game Equilibrium Models, forthcoming.

Gardner, Roy, and Molly R. Morris, "The Evolution of Bluffing in Animal Contest: an ESS Approach." Journal of Theoretical Biology, Vol 137, (1989), pp. 235-243.

Gardner, Roy, Molly R. Morris, and Craig E. Nelson "Conditional Evolutionary Stable Strategies." Animal Behavior. Vol 35 (1987) pp. 507-517.

Guth, Werner, Rolf Schmittberger and Bernd Schwarze. "An Experimental Analysis of Ultimatum Bargaining." Journal of Economic Behavior and Organization, Vol. 3 (1982), 367-388.

Hansen, Robert H., and William F. Samuelson, "Evolution in Economic Games." Journal of Economic Behavior and Organisation, Vol 10 (1988).

Harsanyi, John, and Reinhard Selten, A General Theory of Equilibrium Selection in Games. Cambridge: MIT Press, 1988.

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

Holt, Debra J., "An Empirical Model of Strategy Choices in Coordination Games." Carnegie Mellon University GSIA Working Paper 1990-34, November 1990.

Miller, John H., and James Andreoni, "A Coevolutionary Model of Free Riding Behavior: Replicator Dynamics as an Explanation of the Experimental Results." Working paper, 1990.

Roth, Alvin E., and Francoise Schoumaker. "Expectations and Reputations in Bargaining: An Experimental Study." American Economic Review, Vol. 73 (1983), 362-72.

Schelling, Thomas C. Micro motives and Macrobehavior, New York: Norton, 1978.

Selten, Reinhard, "Bounded Rationality." Journal of Institutional and Theoretical Economics, Vol. 146, (1990), pp. 649-58.

Siegel, Sidney. "Decision Making and Learning under Varying Conditions of Reinforcement." Annals of the New York Academy of Sciences, Vol. 89 (1961), 766-783.

Van Damme, Eric. Stability and Perfection of Nash Equilibrium. New York: Springer-Verlag, 1988.

Workshop in Political Theory and Policy Analysis

513 N. Park | Indiana University | 812-855-0441

Colloquium Presentation
September 23, 1991

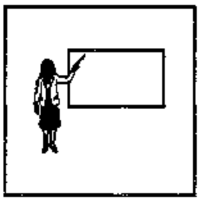

Professor Roy Gardner, Department of Economics, Indiana University, and Sheryl Ball, Department of Business Economics and Public Policy, School of Business, Indiana University, will he the speakers for the Workshop Colloquium on Monday, September 23, 1991. Their presentation is entitled "The Evolution of Beliefs in a Finitely Repeated Game." An abstract of their paper is provided below.

Consider an agent playing a series of games against opponents drawn from a given population composed of various player types. The choice of a strategy then depends critically on the agent's beliefs about the distribution of types present in the population. Beliefs evolve with playing experience, and if play converges to an equilibrium, so do beliefs. In this paper we model evolution according to the replicator dynamics drawn from evolutionary stability theory. We examine a finitely repeated coordination game played against two types of randomly drawn opponents, human and robot. By controlling the distribution of types and initial beliefs, we achieve experimental control over the evolution of beliefs. Our major result is that in a wide variety of designs and treatments human subject play does converge to an evolutionarily stable strategy.

A copy of their paper is available by calling the above telephone number. Colloquium sessions begin at 12 noon and adjourn promptly at 1:30 p.m. You are welcome to bring your lunch. Coffee is provided free of charge, and soft drinks are available. We hope you will be able to join us!