Social Learning and Coordination Conventions in Intergenerational Games: An Experimental Study

Andrew Schotter

New York University

Barry Sopher

Rutgers University

We investigate the creation and evolution of conventions of behavior in "intergenerational games" or games in which a sequence of non-overlapping "generations" of players play a stage game for a finite number of periods and are then replaced by other agents who continue the game in their role for an identical length of time. Players in generation t can offer advice to their successors in generation $t+1$. What we find is that word-of-mouth social learning (in the form of advice from laboratory "parents" to laboratory "children") can be a strong force in the creation of social conventions.

This work was completed under National Science Foundation grants SBR-9709962 and SBR-9709079. The financial support of the Russell Sage Foundation and the C. V. Starr Center for Applied Economics at New York University is also gratefully acknowledged. The paper has benefited greatly from presentation at the MacArthur Foundation, the Russell Sage Foundation, and seminars at the Wharton School, Harvard University, Washington University in St. Louis, the University of Delaware, City University of New York Graduate Center, Grinnell College, the Economic Science Association, and the University of Pittsburgh. We would like to thank Steven Levitt, Hiroki Tsurumi, James Ramsey, and an anonymous referee for their valuable suggestions. Finally, we would like to thank Sangeeta Pratap, Mikhael Shor, and Judy Goldberg for their valuable research assistance and Yevgeniy Tovshteyn for writing the program on which the experiments were run.

[*Journal of Political Economy*, 2003, vol. 111, no. 3]

I. Introduction

This is a paper on the creation and evolution of conventions of behavior in "intergenerational games." In these games a sequence of nonoverlapping "generations" of players play a stage game for a finite number of periods and are then replaced by other agents, who continue the game in their role for an identical length of time. Players in generation t are allowed to see the history of the game played by all (or some subset) of the generations who played it before them and can communicate with their successors in generation $t+1$ and advise them on how they should behave. Hence, when a generation t player goes to move, she has both history and advice at her disposal. In addition, players care about the succeeding generation in the sense that the payoff of each generation's players is a function not only of the payoffs achieved during their generation but also of the payoffs achieved by their children in the game that is played after they retire.[1]

Our motivation for studying such games comes from the idea that while much of game-theoretical research on convention creation has focused on the problem of how infinitely lived agents interact when they repeatedly play the same game with each other over time, this problem is not the empirically relevant one. Rather, as we look at the world around us, we notice that while many of the games we see may have infinite lives, the agents who play these games are finitely lived and play these games for a relatively short period of time. When they retire or die, they are replaced by others, who then carry on. When these transitions take place, each agent transmits all the information about the norms and conventions that have been established to successors.

As we shall see in the Battle of the Sexes Game studied here, the result of this cultural transmission may be a perpetuation of social and economic inequality or what Ullmann-Margalit (1977) calls a "norm of

[1] We use a nonoverlapping generations structure and not an overlapping generations one because in most overlapping generations games of this type (see Crémer 1986; Salant 1991; Kandori 1992) cooperation is achieved when the players of each generation realize that they must be nice to their elders since they will be old one day, and if the current young see them acting improperly toward their elders, they will not provide for them in their old age. The analysis is backward looking in that each generation cares about the generation coming up behind it, and players act properly now knowing that they are being observed and will interact directly with that generation. In this literature, folk-like theorems are proved if the length of the overlap between generations is long enough. In our work, however, generations never overlap. What players do is hope to behave correctly so that their children will see them as an example and act appropriately toward each other. Since they care about their children, adjacent generations are linked via their utility functions but not directly through strategic interaction. Hence, our model is a limiting type of overlapping generations model in which the overlap is either minimal or nonexistent.

Except for the use of advice and the interdependence of our generational payoffs, our game has many of the features of Jackson and Kalai's (1997) recurring games.

partiality" in which an equilibrium with uneven asymmetric payoffs is established as the norm of behavior for a group of people and passed on as the status quo from generation to generation through a process of socialization (see also Schotter 1981). For example, any situation involving a network externality is a candidate for such a norm. Examples here include the common adoption of the Windows operating system as opposed to Unix, the QWERTY keyboard, VHS video format, and so forth. All of these confer unequal benefits to agents in the economy and perpetuate a status quo. Other examples include occupational segregation by sex in which certain occupations come to be predominantly female, seniority rules in which privileges are conferred on certain people because of their years of service (a characteristic that may or may not be correlated with current merit), and any other situation in which privileges or property rights are awarded arbitrarily to some subset of the population and these privileges are perpetuated over time.

The evolutionary model we have in mind is more Lamarckian than Darwinian in that while Jean-Baptiste de Monet de Lamarck had the wrong model of biological evolution, believing that animals could pass on acquired traits to their successors, such a model may be a correct model of social evolution in which generations of social agents pass on conventions of behavior they create during their lifetime to their successors.[2] Such conventions may reinforce social inequality.

What we find is that word-of-mouth social learning (in the form of advice from parents to children) can be a strong force in the creation of social conventions, far stronger than the type of learning subjects seem capable of doing simply by learning the lessons of history without

[2] Our emphasis on this Lamarckian evolutionary process is in contrast to practically all work in evolutionary game theory that is predominantly Darwinian (see, e.g., Kandori, Malaith, and Rob [1993], Weibull [1995], Vega-Redondo [1996], and Samuelson [1997], just to name a few). In this literature, conventions are depicted as the equilibrium solution to some recurrent problem or game that social agents face. More precisely, in these models agents are depicted as nonthinking programs (genes) hard-wired to behave in a particular manner. These agents either interact randomly or "play the field." The dynamic of the growth or decay of these strategies is governed by some type of replicator-like dynamic (see Weibull 1995) in which those strategies that receive relatively high payoffs increase in the population faster than those that receive relatively low payoffs. The focus of attention in this literature is on the long-run equilibria attained by the dynamic. Does it contain a mixture of strategies or types? Is any particular strategy by itself an evolutionarily stable strategy? Are there cycles in which different strategies overrun the population for a while and then die out only to be replaced by others later on? An exception to this strand of work is the work of Jackson and Kalai (1997) on recurring games, which have a structure very close to our intergenerational games except for the intergenerational communication and caring.

Of course this point has already been made by Cavalli-Sforza and Feldman (1981), Boyd and Richerson (1985), and more recently Bisin and Verdier (2000), all of whom have presented a number of interesting models in which imitation and socialization, rather than pure absolute biological fitness, are the criteria on which strategies evolve. We would include Young's (1996, 1998) work in this category as well.

the guidance offered by such advice. Put differently, we find that in terms of coordinating subject behavior, having access to both parental advice and the complete history of the game being played is quite efficient; having access only to history is inadequate (i.e., subjects coordinate their behavior over half the time when they both get advice and see history whereas they coordinate less than one-third of the time when they are deprived of advice). Eliminating a subject's access to history while preserving his or her ability to get advice seems to have little impact on the ability to coordinate. Hence, in our intergenerational setting, it appears as though advice is a crucial element in the creation and evolution of social conventions, an element that has been given little attention in the past literature.

In addition to highlighting the role played by social learning in social evolution, the data generated by our experiments exhibit many of the stylized facts of social evolution, that is, punctuated equilibria, socialization, and social inertia. What this means is that during the experiment, social conventions appear to emerge over time, are passed on from generation to generation through the socializing influence of advice, and then spontaneously seem to disappear only to emerge in another form later in the experiment. (Such punctuated equilibria are also seen in the theoretical work of Young [1996, 1998], where people learn by sampling the population of agents who have played before and then make errors in best responding to what they have learned.) Some behavior is quite persistent, taking a long time to disappear despite its dysfunctional character.

In this paper we shall proceed as follows: Section II presents our experimental design. In Section III, we present the results of our experiments by first describing how our results illustrate the three properties of social evolution we are interested in: punctuated equilibrium, socialization, and inertia. Section IV is about social learning. It describes what happens in our experiments when we eliminate our subjects' ability to pass on advice or see the history of their predecessors. Finally, in Section V, we offer some conclusions and speculations for future work.

II. The Experiment: Design and Procedures

A. *General Features*

The general features of our intergenerational Battle of the Sexes Game were as follows: Subjects once recruited were ordered into generations. Each generation played the game once and only once with an opponent. After their participation in the game, subjects in any generation t were replaced by a next generation, $t+1$, who were able to view some or all of the history of what had transpired before them. Subjects in generation

t were able to give advice to their successors by suggesting a strategy and by explaining why such advice was being given. This feature obviously permits socialization. The payoffs to any subject in the experiment were equal to the payoffs earned by that generation's players during their lifetime plus a discounted payoff that depended on the payoffs achieved by their immediate successors. Finally, during their participation in the game, subjects were asked to predict the actions taken by their opponent (using a mechanism that makes telling the truth a dominant strategy). This was done in an effort to gain insight into the beliefs existing at any time during the evolution of our experimental society.

The experiment was run at both the Experimental Economics Laboratory of the C. V. Starr Center for Applied Economics at New York University and the Experimental Lab in the Department of Economics at Rutgers University. Subjects were recruited, typically in groups of 12, from undergraduate economics courses and divided into two groups of six, with which they stayed for the entire experiment. During their time in the lab, for which they earned approximately an average of $26.10 for about one and a half hours, they engaged in three separate intergenerational games, a Battle of the Sexes Game (BOSG), an Ultimatum Game in which they were asked to divide 100 francs, and a Trust Game as defined by Berg, Dickhaut, and McCabe (1995). All instructions were presented on the computer screens, and questions were answered as they arose. (There were relatively few questions, so it appeared that the subjects had no problems understanding the games being played, which purposefully were quite simple.) All subjects were inexperienced in this experiment.

In this paper we present only the results of the following Battle of the Sexes Game.[3]

	Column Player	
Row Player	1	2
1	150, 50	0, 0
2	0, 0	50, 150

As is true in all BOSGs, this game has two pure-strategy equilibria. In one, (1,1), player 1 does relatively well and receives a payoff of 150

[3] The actual experiment performed had three periods. In each period a subject would play one of the three games with a different opponent. For example, in period 1, players 1 and 6 might play the Battle of the Sexes Game whereas players 2 and 5 play the Ultimatum Game and players 3 and 4 play the Trust Game. When they had finished their respective games, we would rotate them in the next period so that in period 2 players 2 and 4 play the Battle of the Sexes Game, players 3 and 6 play the Ultimatum Game, and players 1 and 5 play the Trust Game. The same type of rotation was carried out in period 3 so that at the end of the experiment each subject had played each game against a different opponent who had not played with any subject he had played with before.

whereas player 2 does less well and receives a payoff of 50. In the other equilibrium, (2,2), just the opposite is true. In disequilibrium all payoffs are zero. The convention creation problem here is which equilibrium will be adhered to, and the problem is that because each type of player favors a different equilibrium, there is an equity issue that is exacerbated by our generational structure since new generations may not want to adhere to a convention established in the past that is unfavorable to them. (There is also a mixed-strategy equilibrium, which we shall ignore for the present, and a coordinated alternating equilibrium, which we see no evidence of in our data.) The conversion rate of francs into dollars here is 1 franc = \$0.04.

When subjects started to play the BOSG, after reading the specific instructions for that game, they would see on the screen the advice given to them from the previous generation. In the BOSG, this advice took the form of a suggested strategy (either 1 or 2) as well as a free-form message written by the previous generational player offering an explanation of why he suggested what he did. No subjects could see the advice given to their opponent, but it was known that each side was given advice. In the baseline experiment, it was also known that each generational player could scroll through the previous history of the generations before it and see what generational players of each type chose and what payoff they received. They could not see, however, any of the previous advice given to their predecessors. Finally, before they made their strategy choice, they were asked to state their beliefs about what they thought was the probability that their opponent would choose any one of his or her two strategies.

To get the subjects to report truthfully, we paid subjects for their predictions according to a proper scoring rule, which gave them an incentive to report their true beliefs. More specifically, before subjects chose strategies in any round, they were asked to enter into the computer the probability vector that they felt represented their beliefs or predictions about the likelihood that their opponent would use each of his of her pure strategies.[4] We rewarded subjects for their beliefs in experimental points, which were converted into dollars at the end of the experiment as follows.

First, subjects reported their beliefs by entering a vector $\mathbf{r} = (r_1, r_2)$ indicating their belief about the probability that the other subject would use strategy 1 or 2.[5] Since only one such strategy would actually be used,

[4] See App. B for the instructions concerning this part of the experiment.

[5] In the instructions, r_i is expressed as numbers in [0, 100], so they are divided by 100 to get probabilities.

the payoff to player i when strategy 1 was chosen by a subject's opponent and $\mathbf{r}$ was the reported belief vector of subject i would be

$$\pi_1 = 20{,}000 - [(100 - r_1)^2 + (r_2)^2]. \tag{1}$$

The payoff to subject i when strategy 2 was chosen was, analogously,

$$\pi_2 = 20{,}000 - [(100 - r_2)^2 + (r_1)^2]. \tag{2}$$

The payoffs from the prediction task were all received at the end of the experiment.

Note what this function says. A subject starts out with 20,000 points and states a belief vector $\mathbf{r} = (r_1, r_2)$. If the opponent chooses 1, then the subject would have been best off if he or she had put all of his or her probability weight on 1. The fact that he or she assigned it only r_1 means that he or she has, ex post, made a mistake. To penalize this mistake, we subtract $(100 - r_1)^2$ from the subject's 20,000-point endowment. Further, the subject is also penalized for the amount he or she allocated to the other strategy, r_2, by subtracting $(r_2)^2$ from his or her 20,000-point endowment as well. (The same function applies symmetrically if 2 is chosen.) The worst possible guess, that is, predicting a particular pure strategy only to have one's opponent choose another, yields a payoff of zero. It can easily be demonstrated that this reward function provides an incentive for subjects to reveal their true beliefs about the actions of their opponents.[6] Telling the truth is optimal; however, this is true only if the subjects are risk neutral. Risk aversion can lead subjects to make a "secure" prediction and place a .50 probability of each strategy. We see no evidence of this type of behavior.

We made sure that the amount of money that could potentially be earned in the prediction part of the experiment was not large in comparison to the game being played. (In fact, over the entire experiment, subjects earned, on average, \$26, whereas the most they could earn on all of their predictions was \$6.) The fear here was that if more money could be earned by predicting well rather than playing well, the experiment could be turned into a coordination game in which subjects would have an incentive to coordinate their strategy choices and play any particular pure strategy repeatedly so as to maximize their prediction payoffs at the expense of their game payoffs. Again, absolutely no evidence of such coordination exists in the data of the BOSG.

B. *Parameter Specification*

The experiments performed can be characterized by four parameters. The first is the length of the history that each generation t player is

[6] An identical elicitation procedure was used successfully by Nyarko and Schotter (2002).

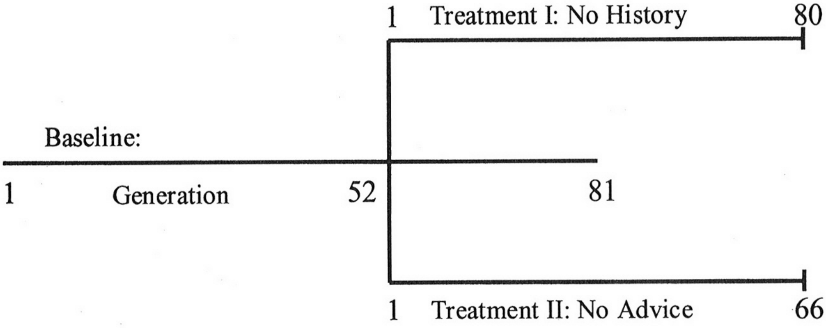

FIG. 1.—Experimental design

allowed to see. The second is the intergenerational discount rate, indicating the fraction of the next generation's payoff to be added to any given generational player's payoff. The third is the number of periods each generation lives for (i.e., the number of times the players repeat the game), and the fourth indicates whether advice is allowable between generations. In all our experiments, each generation lives for one period or repeats the Battle of the Sexes Game only once and has a discount rate of one-half. Hence, they differ only on the basis of the length of history the subjects are allowed to view before playing and whether they are able to get advice from their predecessor or not. In the baseline experiment, subjects could pass advice to their successor and see the full history of all generations before them. This baseline experiment was run for 81 generations. However, at period 52 we took the history of play and started two separate and independent new treatments at that point, which generated a pair of new histories. In the Advice-Only Treatment (Treatment I), before any generation made its move, it could see only the last generation's history and nothing else. This treatment isolated the effect of advice on the play of the intergenerational game. The History-Only Treatment, Treatment II, was identical to the baseline except for the fact that no generation was able to pass advice on to its successors. The players could see the entire history, however, so that this treatment isolated the impact of history. The Advice-Only Treatment was run for an additional 80 generations, and the History-Only Treatment was run for an additional 66 generations, each starting after generation 52 was completed in the baseline. Hence, our baseline had a length of 81, the Advice-Only Treatment a length of 80,[7] and the History-Only Treatment a length of 66. Our experimental design can be represented by figure 1.

[7] One generation was lost because of a computer crash. The lost generation was the third (last) period of a session. We were able to reconstruct the relevant data files.

III. Results

We shall analyze our results by first seeing how they illustrate what we consider to be the three basic stylized facts of social evolution: punctuated equilibria, inertia, and socialization. After this we investigate the role of social learning in our experiment by taking a close look at the role played by advice.

A. *Stylized Facts of Social Evolution*

The stylized facts of social evolution that we wish to study in our experiment are as follows.

1. Punctuated Equilibria

If one looks at the history of various societies, one sees certain regularities in their development. First, as Young (1996) makes clear, over long periods of time one observes periods of punctuated equilibria in which certain conventions of behavior are established, remain perhaps for long periods of time, but eventually give way to temporary periods of chaos, which then settle down into new equilibria.[8]

In our experiments, departures from equilibria are sometimes caused by the advice handed down from one generation to the next. As we shall see, there are times during the experiment in which a convention appears to be relatively firmly established, and yet there will be generational advice advocating a departure. In addition, there will be periods in which a convention also seems firmly established and advice will be given to adhere to it, only to be ignored. Each of these phenomena causes a disruption in the chain of social learning that is passed on from generation to generation and can cause spontaneous breakdowns of what appear to be stable social conventions.

2. Socialization

Another stylized fact of social evolution that we wish to capture in our design is the fact that such evolution is maintained by a process of

[8] There are a number of reasons for the disruption of these conventions. In Darwinian models of evolution, random mutations can arise that, if persistent enough, can cause a disruption of the current equilibrium and drift toward a new one (see Fudenberg and Maskin 1990; Samuelson and Zhang 1992; Kandori et al. 1993; Young 1993; Samuelson 1997). In Young's (1996) model, the cause of disruption is not mutation but rather noise. While various equilibria are more or less resistant to such shocks, noise or mutation can lead to the disappearance, at least temporarily, of existing conventions of behavior.

socialization in which present generations teach and pass on current conventions of behavior to the next generation.[9]

3. Inertia

Because so much behavior is tradition or convention based, there is a lot of inertia built into human action. The world is as stable as it is because people are, to some extent, blindly following the rules and conventions taught to them by their parents or mentors. Social conventions are hard to disrupt since they are often followed unthinkingly, whereas they are sometimes hard to establish because people seem overly committed to past patterns of behavior. Finally, if beliefs are sticky or move sluggishly, inertia will be even harder to overcome since people will find it hard to learn from their mistakes in the past.

B. *Results in the Baseline Experiment*

Since we designed our experiments to allow us to observe not only the actions of subjects but their beliefs and the advice they give each other, let us present these one at a time for the baseline experiment. We shall then go on to investigate behavior in Treatments I and II.

1. Actions in the Baseline Experiment: Punctuated Equilibria

Figure 2 presents the time series of actions generated by our 81-generation baseline experiment. Note that in this figure we have time on the horizontal axis and the actions chosen by our generation pair on the vertical axis. Hence there are four possible action pairs that we can observe: o_{11} = (row$_1$, column$_1$), o_{12} = (row$_1$, column$_2$), o_{21} = (row$_2$, column$_1$), and o_{22} = (row$_2$, column$_2$), where o_{ij} indicates an outcome in which the row player chose action i and the column player action j. (We shall denote these states as states 1, 2, 3, and 4, respectively.)

To give a greater insight into the data, we divide the 81 generations into *regimes* on the basis of estimates of *structural breaks* in the context of a multinomial response model of the probabilities of the different states. The estimation procedure will be described in detail below, but

[9] Replicator dynamics attempt this intergenerational transmission in a very specific and nonhuman manner, but as a descriptive theory of social reality, such a theory is quite poor. Other theories of social evolution (see Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985; Bisin and Verdier 2000) use imitation as the socialization mechanism, and in that sense they are closer to the model we employ here, except for the fact that we shall model only vertical as opposed to horizontal socialization. Still, we see in front of us in the real world such things as tradition and convention-based behavior, which are taught and passed on explicitly by one generation to another. It is this process we wish to capture in our experiments.

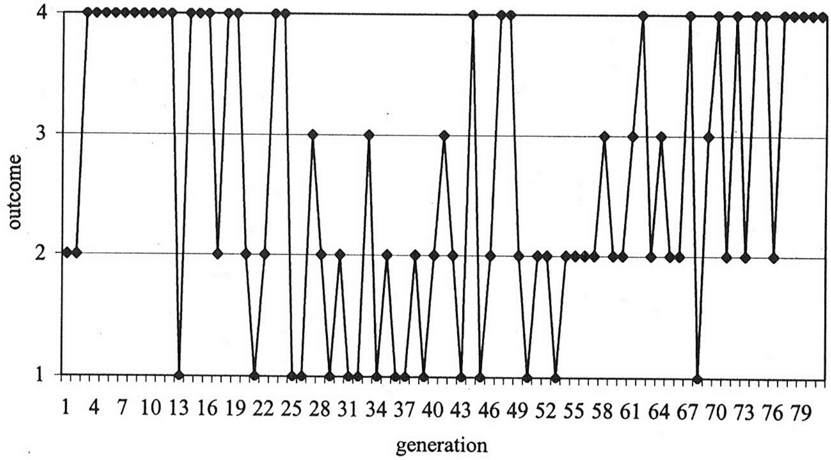

FIG. 2.—Baseline outcomes

we briefly summarize the results here. Our estimates indicate the existence of five distinct regimes. Regime I consists of generations 1–24; regime II, generations 25–37; regime III, generations 38–46; regime IV, generations 47–60; and regime V, generations 61–81. The model used to identify these regimes demonstrates that when we cross the boundaries of our estimated regimes, the probabilities of being in the various states change quite dramatically.

Regime I (generations 1–24) we call the (2,2) convention regime since during this time period we observed 17 periods in which the (2,2) equilibrium was chosen along with one stretch of time in which we observed nine consecutive periods of (2,2), the longest run for any stage-game equilibrium in all 81 generations of the baseline game. Regime II (generations 25–37) we call the (1,1) convention regime because while in the first 24 generations we saw only the (1,1) equilibrium chosen twice, in regime II it was chosen in eight of the 13 generations. In addition, during this time the (2,2) equilibrium, which was so prevalent in regime I, disappears completely. If we look at the row players in regime II, they choose strategy 1 in 11 of the 13 generations, indicating that at least in their minds they are adhering to the (1,1) convention in playing this game. Regimes III and IV (generations 38–60) we call transition regimes since the generational players spend most of their time in a disequilibrium state with infrequent occurrences of the (1,1) equilibrium and the (2,2) equilibrium (four and three, respectively). Finally, regime V (generations 61–81) appears to present evidence that the (2,2) equilibrium is reestablishing itself as a convention after a virtual absence over 35 generations. We say this because during these last 20

TABLE 1
CHOICES OF ROW AND COLUMN PLAYER BY REGIME
A. CHOICES BY STATES AND REGIME

Regime	(1,1)	(1,2)	(2,1)	(2,2)	Total
I	2	5	0	17	24
II	8	3	2	0	13
III	3	4	1	1	9
IV	2	9	1	2	14
V	1	6	3	11	21
Total	16	27	7	31	81

B. CHOICES BY REGIME

Regime	Row 1	Row 2	Column 1	Column 2
I	7	17	2	22
II	11	2	10	3
III	7	2	4	5
IV	11	3	3	11
V	14	7	4	17
Total	50	31	23	58

generations we see the (2,2) equilibrium appearing in 10 out of 20 generations, whereas it appeared only three times in the previous 35 generations. Even more surprising, the row players, after a great resistance to playing row 2 (e.g., they played it only seven times in 35 generations between generations 25 and 60), chose it 14 times in the last 20 generations. In total there were 47 periods of stage-game equilibrium played and 34 periods of stage-game disequilibrium. Note finally that there is a great asymmetry in the number of times in which the (2,1) state arises (seven times) as opposed to the (1,2) state (27 times). These results are tabulated in table 1.

The time series presented in figure 2 offers strong evidence for the existence of the punctuated equilibrium phenomenon. Regime I is clearly a period of time over which the (2,2) equilibrium is firmly established. In fact, round 13, where both row and column deviate simultaneously, does not seem to disrupt the convention, which continues for three more periods after this deviation occurs. What is then surprising in regime II is how completely this convention disappears, never to reestablish itself with any regularity until generations 61–81 (regime V). While regime II does not present as clear a picture of the existence of a convention (the (1,1) outcome, while frequent, is not persistent), the absence of any (2,2) choices, along with the appearance of eight (1,1) choices in 13 generations and the persistent choice of the row player for row 1, creates a strong case for dubbing it the (1,1) convention regime. Regime V, where it appears that the (2,2) convention has reestablished itself, also presents interesting evidence of the punctuated equilibrium phenomenon.

A formal method for determining the number and location of structural break points in the data can be implemented in the context of an estimated econometric model of the process determining the state. To do this, consider a multinomial logit response model in which the probability that state h occurs in generation t is estimated as a function of the recent history of play by the row player and the column player. In the context of such a fitted model, the idea of punctuated conventions suggests sudden changes (punctuations) in the estimated coefficients on the row and column player choice histories. Therefore, if we allow the estimated coefficients on the row and column player history variables to vary over discrete intervals of time, the resulting model can be compared to a "restricted" model in which the coefficients on the row history variables and column history variables (defined for each regime) are restricted to be equal to one another across regimes. We employ the likelihood-based estimation procedure proposed by Quandt (1958) to select the best-fitting model with K structural breaks.[10] We estimate break points for different values of K and then select the number of breaks according to the Akaike information criterion. We then test this "best estimate" of the number of breaks and their locations against the alternative hypothesis of no breaks with a likelihood ratio test.

We first estimated a multinomial logit model for the state on a moving average of the row and column player choices. That is, the probability that any state is observed in period t is a function of the relative frequencies with which the row and column players have used their various strategies over the last m periods. Using a multinomial logit form for this probability yields

$$P_h(t) = \frac{\exp\left(b_o^h + \sum_{k\in S_K} b_{kr}^h r_{t,k} + \sum_{k\in S_K} b_{kc}^h c_{t,k}\right)}{1 + \sum_{j\in J}\exp\left(b_o^j + \sum_{k\in S_K} b_{kr}^j r_{t,k} + \sum_{k\in S_K} b_{kc}^j c_{t,k}\right)},$$

$$h \in J = \{(1,1),\ (1,2),\ (2,1)\},\ P_{(2,2)} = 1 - \sum_{j\in J} P_j,$$

where $k \in \{1, 2, 3, \ldots, K\} = S_K$ indexes the different possible structural regimes, J is the set of states indexed by j, and h is any particular state ((2,2) is the base state). The row history variables, $r_{t,k}$, are defined as follows. Let $f_{t,r}^m$ be the relative frequency with which the row player has chosen action 1 in the previous m periods before t (periods $t - m - 1$ to $t - 1$), and let d_k be a dummy variable equal to one if the observed state in period t is in structural regime k, and equal to zero otherwise. Then $r_{t,k} = f_{t,r}^m d_k$. The column history variables, $c_{t,k}$, are defined similarly.

[10] Though Quandt's article considered a linear model with a single break point, the generalization to multiple break points and a general maximum likelihood model is immediate. Quandt mentions the multiple break point case in a footnote. The method is based on analysis of the likelihood function only and is not restricted to linear models.

TABLE 2
STRUCTURAL BREAK ESTIMATES

Breaks Assumed	Log Likelihood	Pseudo R^2	Breaks	Akaike Information Criterion
0	−81.25	.17		2.14
1	−73.35	.25	61	2.25
2	−68.16	.30	47, 61	2.11
3	−62.88	.35	43, 47, 61	2.13
4	**−55.14**	**.43**	**25, 38, 47, 61**	**2.08**
5	−53.55	.45	25, 34, 44, 48, 61	2.20

NOTE.—The likelihood ratio test for structural change is $\chi^2 = 52.22$ (18 degrees of freedom). Probability greater than χ^2 is .00.

We call the model "restricted" when we impose the following: $b^j_{kr} = b^j_{k'r}$ and $b^j_{kc} = b^j_{k'c}$ for all k and k', implying that there are no structural regimes. The "unrestricted" model allows these coefficients to vary across structural regimes.[11]

In the estimation we use $m = 5$ to construct the row and column player history variables. We decided on this by comparing the fit for the restricted model for different values of m (from $m = 1$ to $m = 10$) on the basis of the pseudo R^2 measure.[12] We then proceeded to estimate the break points in the unrestricted model for different values of K, using the five-period moving average, selecting for each K the best-fitting (highest likelihood) set of break points. We do this exercise for $K = 1, 2, 3, 4$, and 5.[13] The results are summarized in table 2. The table shows, for each number of break points, K, the value of the log likelihood function, the pseudo R^2 measure of goodness of fit, the estimated break points, and the value of the Akaike information criterion (used to select the best-fitting number of breaks). The break points are the points at which a new set of coefficients on the regressors take effect. Thus there are $K + 1$ regimes implied by a set of break points t_1, t_2, …, t_K, namely, (1 to $t_1 - 1$), (t_1 to $t_2 - 1$), …, and (t_K to 81). The best-fitting model according to the Akaike information criterion is the one with four breaks. It was only at the level of four break points that the

[11] For $K > 1$, we do impose one restriction in the "unrestricted" model, that the coefficients in the first and last regimes are the same. This seemed a natural thing to do since the data so clearly begin and end with the (2,2) equilibrium being chosen most frequently. This restriction actually improves the fit according to the Akaike information criterion.

[12] Specifically, the pseudo R^2 is .17 for the five-period moving average, and it declined as one either lengthened or shortened the moving average window (we allowed the window to range from $m = 1$ to $m = 10$).

[13] Estimating the break points involves estimating the multinomial logit model discussed above for every possible configuration of the K break points over the sample period and then selecting the set of break points yielding the highest log likelihood for the multinomial model. This is straightforward enough, though obviously cumbersome. We imposed the restriction in conducting this "grid search" that each regime must be at least four periods long.

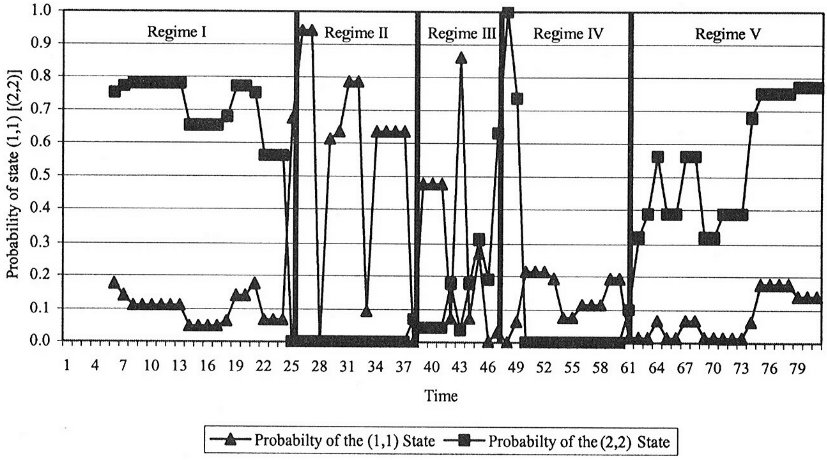

FIG. 3.—Estimated state probabilities

rather obvious break in the data at around generation 25 entered the picture, and this leads to a substantial improvement in the fit.

The restricted and unrestricted estimates for the multinomial logit model with $K = 4$ are contained in Appendix table A1. Since the coefficients are not of particular interest here, we report only the likelihood ratio test for the structural change hypothesis in the note to table 2.[14] The likelihood ratio χ^2 for the structural change hypothesis is significant: the restriction that coefficients are equal across regimes is rejected at the 0 percent significance level.

The estimated probabilities for the equilibrium states (1,1) and (2,2) derived from the unrestricted estimated multinomial logit model are presented in figure 3. This figure presents strong evidence that behavior changed dramatically as generations moved across our regimes. For example, note that in regime I the probability of being in state (2,2) is consistently high and then drops precipitously as the experiment enters our regime II. In fact, during almost the entire length of our regimes II and III, that probability is practically zero. Note, however, the reemergence of the (2,2) state in regime V.

The evidence for the (1,1) state is not so uniformly strong but is still

[14] Since there are three equations to estimate (one less than the total number of states), there are, besides the constant terms, 2 (row and column history) × 3 (equations) × 5 (regimes) = 30 coefficients to estimate in the unrestricted model, and 2 (row and column history) × 3 (equations) = 6 coefficients to estimate in the restricted model, apart from the constant terms. We imposed the restriction in the models with two or more breaks that the first and last regimes have the same coefficients, so in fact there are only 24 estimated parameters in the so-called unrestricted model. Thus the χ^2 statistic for the test of structural change has $24 - 6 = 18$ degrees of freedom.

convincing. Note here that the initial probability of being in state (1,1) in regime I is practically zero for the first 25 generations and then rises dramatically in regimes II and III, only to dwindle away in regimes IV and V. The abrupt changes in the probability of the (2,2) state as we cross our regime boundaries along with the low probabilities for the (1,1) state in all but regimes II and III support the idea that there are distinct equilibrium regimes.

2. Inertia and Norms of Inequality

With respect to inertia, there are really two types of social inertia one can discuss. One, which we shall call equilibrium inertia, is the inertia that leads people to adhere to a convention simply because it has existed for a long time in the past despite the fact that it may not be the best equilibrium for their particular group. For example, in our experiment the (2,2) convention is obviously the best convention for the column chooser. Hence, when a row player enters the game and observes (as in regime I) that this convention has been in place for a very long time, and hence is likely to be chosen by the other side, there are a great many forces leading such a player to continue adhering to the convention. Given these forces, it is actually surprising that the (2,2) convention ever disappeared after round 24. In fact, if the (2,2) convention is a strong convention in which each player thinks that his or her opponent is going to adhere with probability one, then deviating can never be beneficial since if one continues to adhere, one will get 50 today plus one-half of 50 tomorrow. Deviating will yield zero today, and if the player is successful in breaking the (2,2) convention and shifting it to the (1,1) convention in period $t = 1$ (an event that is rather unlikely given that we are talking about a strong convention), then he or she will get one-half of 150 tomorrow. In either case, the payoff will be 75, so that there is no positive incentive to deviate unless one cares about generations beyond next period, a consideration that was ruled out by our intergenerational utility function. (We shall be able to explain this disappearance later when we talk about advice.)

Note that such conventions establish what Ullmann-Margalit (1977) calls "norms of partiality" in which seemingly symmetric agents select an equilibrium to a game that favors one type of agent and then pass this norm, or what we call convention, on to succeeding generations. The point is that the agents being favored today are no more worthy of preferential treatment than their cohorts, yet a quirk of history (a path dependency) has given one mode of behavior saliency. We suspect that the pressures to deviate from established equilibria and the resulting punctuated equilibria discussed above are the result of pressures that arise as one type of agent realizes that the cause of the inequality

TABLE 3
CONTINUATION PROBABILITIES BY REGIME

	1,1	1,2	2,1	2,2
Regime I	0	.17	...*	.81
Regime II	.38	0	0	0
Regime III	0	0	0	0
Regime IV	0	.56	0	.50
Regime V	0	.17	0	.45
Total	.19	.26	0	.63

* No (2,1) state occurred in regime I.

they face is purely arbitrary and hence not fair. It appears as though subjects deviate to further the cause of their type at a sacrifice to themselves (another violation of Darwinian evolutionary theory).[15]

Another type of social inertia exists when people are recalcitrant and persist in behavior that is clearly detrimental to them. For example, in regimes II–IV, the row players, apparently in an effort to move the convention from (2,2) to (1,1), which is better for them, persisted in choosing row 1 30 out of 35 generations between generations 25 and 60. They persisted in doing so despite the fact that this behavior led to a disequilibrium outcome in 18 of those generations. Obviously, they felt that their efforts might establish the (1,1) equilibrium favorable to them as a convention even if they would not benefit directly from it.

To give a different picture of the persistence of both equilibrium and disequilibrium states, we calculated a continuation probability for each of our four states in each of the regimes listed above. More precisely, a continuation probability defines a conditional probability of being in any given state in period $t + 1$ given that one was in that state in period t.

Table 3 presents the probabilities. Since conventions are persistent states, our intention in presenting table 3 is to give some indication as to what states seem to form conventions in each of these regimes. For example, in regime 1 the (2,2) state is remarkably persistent, indicating an .81 probability of remaining in the (2,2) state if one reached it. In regime II, while the (1,1) state was observed eight out of 13 times, many of these instances were isolated instances that were not repeated. Still, the continuation probability was .38. More remarkable is the fact that none of the other states ever repeated themselves during the entire

[15] This conjecture is supported by the results of some of our pilot experiments run on generational subjects who played the game 10 times before being replaced. Such subjects easily established an alternating convention in which they successively alternate between choosing (1,1) and (2,2). This has the effect of equalizing the payoffs to subject types (row or column) and makes adherence easier in the long run. The same pilots indicated that such conventions do not get established when there are only three or four periods to a generation's lifetime, so there is still work to be done here.

TABLE 4
ADVICE OFFERED CONDITIONAL ON THE STATE

State	Row 1	Row 2	Column 1	Column 2
1,1	16	0	14	2
1,2	9	18	15	12
2,1	7	0	2	5
2,2	3	28	0	31

regime. Regimes III and IV demonstrated a dramatic ability to remain in the disequilibrium state (15 out of 23 times), with no persistence in the (2,1) state but a continuation probability for the (1,2) state of .55. Finally, regime V showed the return of the (2,2) state and its persistence of .45.

3. Socialization in the Baseline

The type of Lamarckian evolution we are interested in here relies heavily on a process of social learning for its proper functioning. The transmission of conventions and "culture" through advice is permitted in our experiments and turns out to be extremely important to the functioning of our experimental societies.

To discuss advice, we present a summary of how advice was given in table 4 and under what circumstances it was followed in table 5.

What advice was given.[16]—Table 4 presents the type of advice that was offered subjects by their predecessors conditional on the state. Note the conservatism of this advice. When a stage-game equilibrium state has been reached, no matter which one, subjects overwhelmingly tell their successors to adhere to it. For the row player this occurs 100 percent of the time (16 out of 16 times) when the stage-game equilibrium is the (1,1) equilibrium, the equilibrium that is best for the row player; it occurs 90 percent of the time (27 out of 30 times) when the state is (2,2). For the column player a similar pattern exists. When the state is (2,2), the state that is best for the column player, we see 100 percent of the column players (31 out of 31) suggesting a choice of 2; when the state is (1,1), 87.5 percent of the subjects suggest that their successors adhere to the (1,1) equilibrium despite the fact that it gives the opponent the lion's share of the earnings.

[16] In Schotter and Sopher (2000), we investigate the content of the advice given by coding it and investigating how it changes depending on the state of the game. What we find is that the detail with which messages are written depends on the state of the game. When an equilibrium state existed last period that determined a good outcome for subjects, i.e., they received a 150 payoff, they tended to leave low-level messages that were not supported by strategic reasoning. However, the subjects receiving the low payoff tended to leave more highly reasoned and strategic advice.

TABLE 5
ADVICE ADHERENCE CONDITIONAL ON LAST PERIOD'S STATE

Advice	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
	A. State Last Period (1,1)			
1	11	5	5	9
2	0	0	1	1
Total	11	5	6	10
	B. State Last Period (1,2)			
1	7	2	10	5
2	10	8	10	2
Total	17	10	20	7
	C. State Last Period (2,1)			
1	5	2	0	2
2	0	0	4	1
Total	5	2	4	3
	D. State Last Period (2,2)			
1	3	0	0	0
2	19	8	26	4
Total	22	8	26	4

When the last-period state was a disequilibrium state, behavior was more erratic and differed across row and column players. Note that there are two types of disequilibrium states. In one, the (2,1) state, each subject chose in a manner consistent with the equilibrium that was best for his or her opponent. We call this the *submissive disequilibrium* state since both subjects yielded to the other and chose the state that was best for his or her opponent. The (1,2) state is the *greedy disequilibrium* state since here we get disequilibrium behavior in which each subject chooses in a manner consistent with his or her own best equilibrium. In the submissive disequilibrium state, (2,1), both the row and column subjects overwhelmingly suggest a change of strategy for their successors in which they suggest a greedy action next period. More precisely, in the seven such instances of the submissive disequilibrium state, the row player gave advice to switch and choose row 1 in all seven instances, whereas the column player suggested switching and choosing 2 in five of the seven cases. When the greedy disequilibrium state occurred, advice was more diffuse. In 18 of the 27 occurrences of this disequilibrium state, the row player suggested switching to the submissive strategy of choosing row 2, and nine suggested standing pat and choosing row 1. For the column players, 15 suggested switching to the submissive strategy (column 1), and 12 suggested standing pat and continuing to choose column 2.

When advice was followed.—In order for an equilibrium convention to persist, it must be the case that either all generations advise their successors to follow the convention and their advice is adhered to, or their advice deviates from the dictates of the equilibrium and it is ignored. What we find when we look at the behavior of subjects is that they overwhelmingly tended to follow the advice they were given but not sufficiently strongly to prevent periodic deviations; hence the punctuated equilibrium behavior we discussed above. More precisely, table 5 presents the frequency with which advice was followed conditional on the state in which it was given.

These tables present some interesting facts. First of all, advice appears to have been followed quite often, but the degree to which it was followed varies depending on the state last period. On average, for the row players, it was followed 68.75 percent of the time, whereas for the column players it was followed 70 percent of the time. When the last-period state was (2,2), row players followed the advice given to them 73.3 percent of the time (strangely agreeing to follow advice to switch to the row 1 strategy three out of the three times), whereas column subjects followed 86.6 percent of the time (here all advice was to choose column 2). When the last-period state was the (1,1) equilibrium, column subjects chose to follow it only 37.5 percent of the time, whereas row players adhered 68 percent of the time.

One question that arises here is how powerful advice is when compared to the prescriptions of best-response behavior. For example, it may be that subjects follow advice so often because the advice they get is consistent with their best responses to their beliefs, so following advice is simply equivalent to best responding. In our design we are fortunate in being able to test this hypothesis directly since for each generation we have elicited their beliefs about their opponent and hence know their best response and also the advice they have received. Hence it is quite easy for us to compare them, and this is what we do in table 6.

What we can conclude from this table is quite striking. When advice and best responses differ, subjects are about as likely to follow the dictates of their best responses as they are those of the advice they are given. For example, for the row players, there were 28 instances in which the best-response prescription was different from the advice given; of those 28 instances, the advice was followed 15 times. For the column players, there were 34 such instances, and in 17 of them the column player chose to follow advice and not to best respond. These results are striking since the beliefs we measured were the players' posterior beliefs after they had both seen the advice given to them and the history of play before them. Hence, our beliefs should have included any informational content contained in the advice subjects were given, yet half of the time they still persisted in making a choice that was inconsistent

TABLE 6
FOLLOWING ADVICE WHEN ADVICE AND BEST RESPONSES DIFFER AND ARE EQUAL

	Row		Column	
	Follow	Reject	Follow	Reject
	A. When Advice and Best Responses Differ			
State last period (1,1)	0	3	3	8
State last period (1,2)	4	5	11	6
State last period (2,1)	0	0	0	2
State last period (2,2)	11	5	3	1
Total	15	13	17	17
	B. When Advice Equals Best Responses			
State last period (1,1)	11	2	3	2
State last period (1,2)	13	5	9	1
State last period (2,1)	5	2	4	1
State last period (2,2)	11	3	23	3
Total	40	12	39	7

with their best response. Since advice in this experiment was a type of private cheap talk based on little more information than the next generation already possessed (the only informational difference between a generation t and a generation $t + 1$ player is the fact that the generation t player happened to have played the game once and received advice from his predecessor, which our generation $t + 1$ player did not see directly), it is surprising that it was listened to at all.

One of the striking aspects to this advice-giving and advice-receiving behavior is how it introduces a stochastic aspect into what would otherwise be a deterministic best-response process. If advice was always followed, or at least followed when it agreed with a subject's best response and if beliefs were such that both subjects would want to choose actions consistent with the (1,1) or (2,2) state, then these states, once reached, would be absorbing. However, we see that neither of these assumptions is supported by our data. Despite the fact that the (2,2) state was observed nine times in a row in regime 1 and despite the fact that choosing 2 was a best response to subjects' stated beliefs, we observed in generation 13 a completely unexplained deviation. In addition, in three of the 30 rounds in which the (2,2) equilibrium was in place, the row player chose not to give advice to his successor to adhere to it, whereas in two of 16 instances in which the (1,1) equilibrium was in place, the column subject chose to offer advice to choose 2. Such behavior makes the process we are investigating more complex.

C. *Beliefs*

As described above, before generational subjects each made their choice, they were asked to state their beliefs about what they felt the

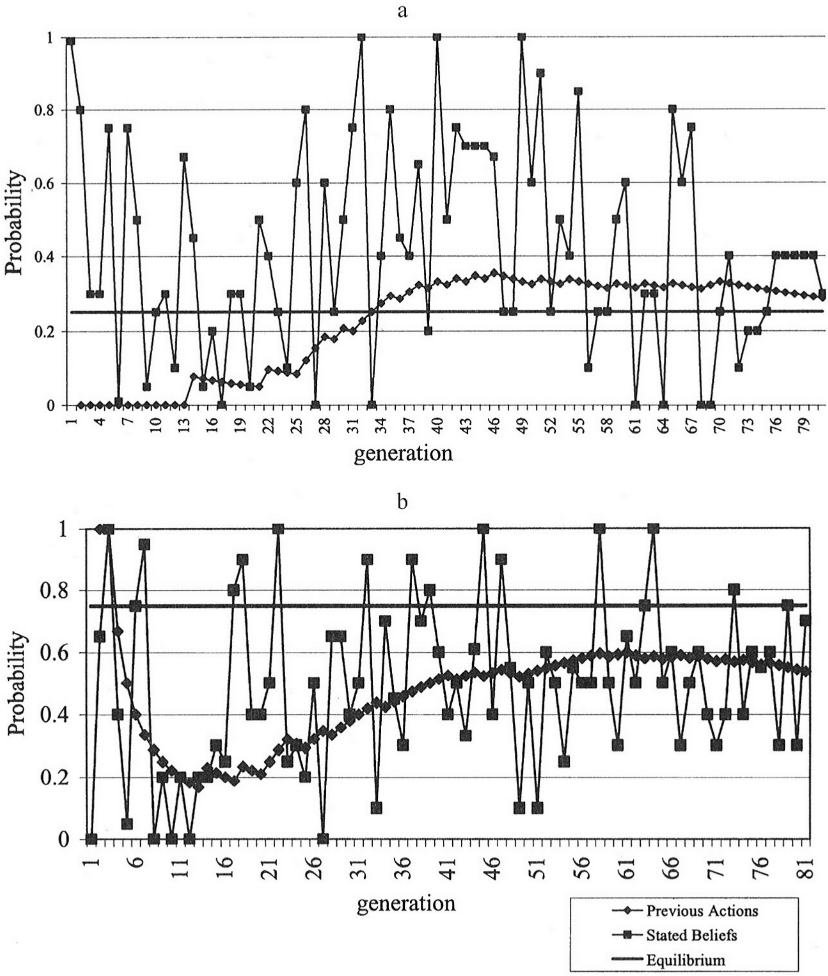

FIG. 4.—*a,* Row's beliefs about column—probability column chooses 1. *b,* Column's beliefs about row, baseline.

probability was that their opponent would use strategy 1 or 2. The time paths of these belief vectors are presented in figure 4, where we present the probability that each generational subject felt his opponent would choose strategy 1.

Note that in figures 4*a* and 4*b* we have placed a straight line that indicates the critical belief value, which is such that if one believes that one's opponent is going to choose strategy 1 with a higher probability than that critical value, a best response is to choose strategy 1 as well. (We have also placed a curved line, which we shall explain shortly but shall ignore at the moment.) As we see, beliefs of both subjects seem to exhibit a type of overconfidence bias in the sense that overwhelmingly both subjects appear to believe that their opponent is going to choose the strategy that is consistent with the equilibrium that is best for them. More precisely, in only 26 of the 81 generations, row subjects believed that their opponent was so likely to choose row 2 as to lead them to choose 2 as a best response. For column players the situation was even worse, with beliefs consistent only with 15 row 1 best responses. Obviously, if these beliefs are based on the history of play of the game, each cannot be correct.

To demonstrate how historical beliefs would differ, we have calculated the empirical beliefs of subjects in this game (i.e., beliefs that the probability that a player will play a strategy is equal to the fraction of time that player has played that strategy in the past) and superimposed them on the graphs as well. While empirical beliefs are a very drastic form of historical belief, giving equal weight to each past observation, they still may be useful as a point of contrast to the stated beliefs we received from our subjects. As we can see, there is little connection between these historical (empirical) beliefs and the stated beliefs of our subjects. (These results replicate the same finding for repeated zero-sum games presented previously in Nyarko and Schotter [2002].) As we see, for the row players the empirical beliefs seem to do a good job at converging to the theoretical equilibrium beliefs as time proceeds, whereas the column players' empirical beliefs appear to be converging to a value considerably less than the theoretical equilibrium value. In either case, however, subject beliefs appear to be more optimistic about the chances of achieving one's preferred equilibrium than is warranted by the data.

In fact for the row player we can reject the hypothesis of the equality of the distributions of stated and empirical beliefs for the 81 generations of the experiment ($p = .00$, $z = 4.93$).

For column players, a signed-rank test fails to reject the hypothesis that the distributions of stated and empirical beliefs are equal either over the entire 81-generation horizon ($z = 0.39$, p-value .70) of the experiment or in any of the regimes (regime I: $z = 0.70$, p-value .48; regime II: $z = 1.55$, p-value .12; regime III: $z = -1.16$, p-value .24; regime IV: $z = -1.36$, p-value .17).

IV. The Advice Puzzle: Social and Belief Learning in Treatments I and II

Starting in generation 52 we introduced two new treatments into our experiment. In Treatment I we "took away history" by having successive generations of players play without the benefit of being able to see any history beyond that of their parent generation. What this means is that subjects performing this experiment knew only that the game they were playing had been played before, possibly many times, but that they could see the play only of the generation before them. They could, however, receive advice just as subjects did in our baseline. This treatment was run independently of the baseline and Treatment II, except for the common starting point in period 52. In Treatment II we "took away advice" by allowing subjects to view the entire history of play before them, if they wished, but not allowing them to advise the next generation.[17]

These treatments furnish a controlled experiment that allows us to investigate the impact of social learning, in the form of advice giving and following, on subjects' ability to attain and maintain an equilibrium convention of behavior in this game. Such learning is in contrast to the more frequently studied belief learning, in which agents take actions that at any time are best responses to the beliefs they have about the actions of their opponents. In our experiment we can easily test these two types of learning since we have elicited the beliefs of agents at each point during the game. Hence, if each generation forms its beliefs in light of history and then best responds to them, the addition of advice should have no impact on the frequency and persistence of equilibrium behavior among the subjects. This is especially true since in our experiments the people giving advice barely have more information at their disposal than the ones receiving it. (The only difference in their information sets is that the advice giver has received advice from his or her parental generation that the receiver has not seen.)

More precisely, if advice giving were not essential to convention building, then we should not observe any difference in the number of times our subjects achieved an equilibrium when we compare Treatment II (the full history/no-advice experiment) to our baseline experiment, where subjects had access to both. Furthermore, if history was not essential for coordination but advice was, then eliminating history and allowing advice, as we did in Treatment I, should lead to identical amounts of cooperation as observed in the baseline.

Figure 5 plots the time series generated by these two treatments along with our original baseline treatment (a repeat of fig. 2). As we can see from these graphs, removing history has a very different impact on the path of play than removing advice. Consistent with what we have noticed above, players in intergenerational games appear much more successful in achieving equilibrium behavior (or establishing a convention) when advice is present even if they have no access to the history of play before them. History, with no accompanying advice, appears to furnish less of

[17] This was done by forbidding them to write any instructions on the screen despite the fact that they were prompted to.

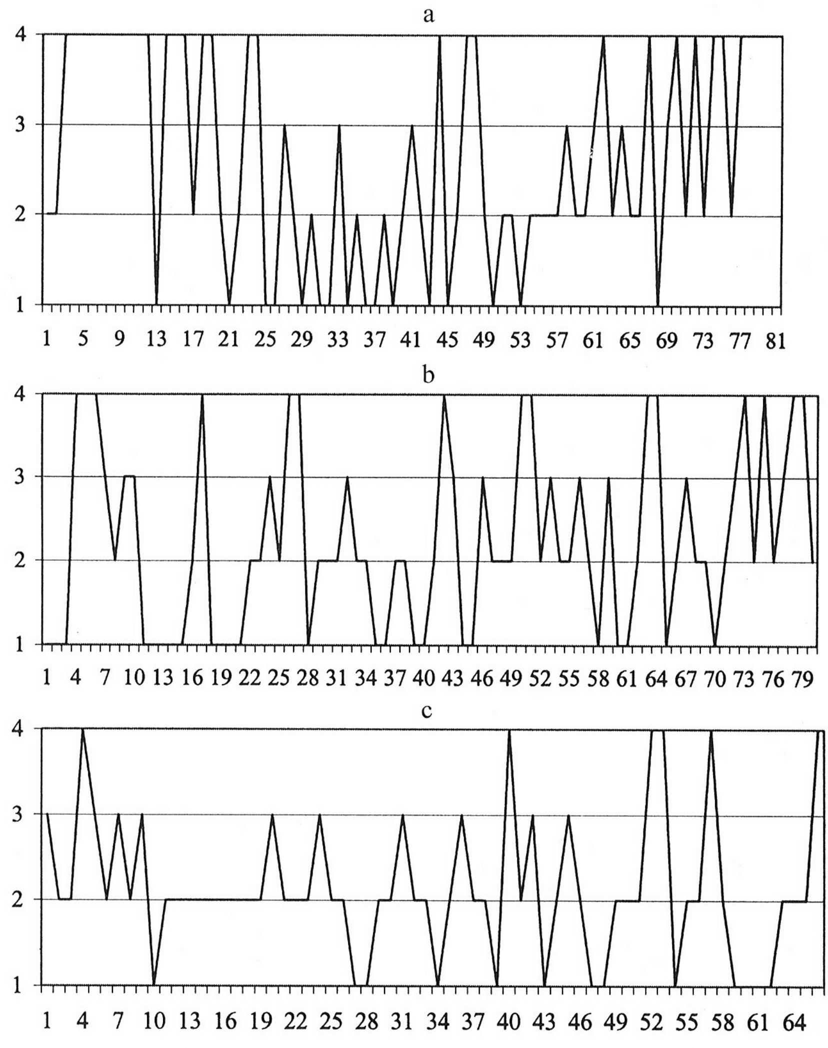

FIG. 5.—*a,* Baseline outcomes. *b,* Treatment I outcomes. *c,* Treatment II outcomes.

a guide to coordinated behavior. More precisely, as we see, Treatment I was successful in reaching a stage-game equilibrium in 39 out of 80 generations, and when equilibrium was reached, subjects maintained it, on average, for 1.95 generations in a row. (The continuation probability was $20/39 = .512$.) In Treatment II, equilibria to the stage game appeared rather infrequently, in just 19 out of 66 generations, with a continuation probability of .315 and a mean persistence of 1.58. Hence, there is a dramatic drop in the frequency of coordination when advice is removed. In the baseline we observe equilibrium outcomes 47 out of 81 times.[18]

These results raise what we call the "Advice Puzzle," which is composed of two parts. Part 1 is the question of why subjects would follow the advice of someone whose information set contains virtually the same information as theirs. In fact, the only difference between the information sets of parents and children in our baseline experiment is the advice that parents received from their parents. Other than that, all information is identical, yet our subjects defer to their parents' advice almost 50 percent of the time when the advice differs from the best response to their own beliefs.[19]

Part 2 of our paradox is the puzzle that despite the fact that advice is private and not common knowledge cheap talk, as in Cooper et al. (1989), it appears to aid coordination in the sense that the amount of equilibrium occurrences in our baseline (58 percent) and Treatment I (49 percent), where advice was present, is far greater than that of Treatment II (29 percent), where no advice was present. While it is known that one-way communication in the form of cheap talk can increase coordination in Battle of the Sexes Games (see Cooper et al. 1989) and that two-way cheap talk can help in other games (see Cooper et al. (1992), how private communication of the type seen in our experiment works is an unsolved puzzle for us.

Finally, note that the desire of subjects to follow advice has some of the characteristics of an information cascade since in many cases subjects are not relying on their own beliefs, which are based on the information

[18] A more formal way to compare the impact of these treatments on the behavior of our subjects is to compare the state-to-state transition matrices generated by our baseline data and test to see whether they were generated by the same stochastic process generating the data observed in Treatments I and II. More precisely, treating the data as though they were generated by a one-state Markov chain, for each experiment we can estimate the probability of transiting from any of our four states {(1,1), (1,2), (2,1), and (2,2)} to the other. A simple counting procedure turns out to yield maximum likelihood estimates of these transition probabilities. Doing so would generate a 4×4 transition matrix for each experimental treatment. These transition matrices are presented in the appendix to Schotter and Sopher (2000). To test whether the transition probabilities defined by our baseline data are generated by a process equivalent to the one that generated the data in Treatments I and II, we use a χ^2 goodness-of-fit test. More precisely, call $\mathbf{T}$ the transition matrix estimated from our baseline data and $\mathbf{P}^k$ the transition matrix defined by our kth treatment, i.e., $k = \{I,II\}$. Denote $p_{ij}^{\mathbf{P}^k}$, $j = i = \{1, 2, 3, 4\}$, as the transition probability from state i to state j in matrix $\mathbf{P}^k$. To test whether the transition probabilities estimated for any one of our treatments has been generated by a process with transition probabilities equal to those of our baseline experiment, we employ a χ^2 test (see Schotter and Sopher [2000] for details). We find that we can reject the hypothesis that the same process that generated the baseline data also generated the data observed in either Treatment I (χ^2(12 df) = 27.6521 [p= .000]) or Treatment II (χ^2(9 df) = 59.4262 [p= .000]). Hence, if the process generating our data can be considered Markovian, it would appear as though imposing different informational conditions on the subjects significantly changed their behavior.

[19] There is no sense, then, in which parents in our experiment are in any way "experts" as in the model of Ottaviani and Sorensen (1999).

contained in the history of the game, but are instead following the advice given to them by their predecessor, who is as just as much a neophyte as they are.

V. Conclusions

This paper utilized an experimental approach to investigate the process of convention creation and transmission in intergenerational games. It has modeled the process as a Lamarckian one in which nonoverlapping generations of players create and pass on conventions of behavior from generation to generation. These conventions tend to perpetuate social inequality. Since the process is stochastic, however, it exhibits punctuated equilibria in which conventions are created, are passed on from one generation to the next, but then spontaneously disappear. In this process, several stylized facts appear.

Probably the most notable feature of our results is the central role that the advice, passed on from one generation to the next, plays in facilitating coordination across and between generations. It appears that relying on history and the process of belief learning is not sufficient to allow proper coordination in the Battle of the Sexes Game played by our subjects. For a reason yet left unexplained, advice, even in the absence of history, appears to be sufficient for the creation of conventions whereas history, in the absence of advice, does not. This implies that social learning may be a stronger, and belief learning a weaker, form of learning than previously thought.

Appendix A

TABLE A1
COEFICIENT ESTIMATES FOR THE TEST FOR STRUCTURAL CHANGE

	Unrestricted Model		Restricted Model	
	Estimate	Probability >z	Estimate	Probability >z
	A. State = Row Equilibrium (1,1)			
Row history, regimes I and V	1.24	.72	3.27	.08
Row history, regime II	56.57	.00		
Row history, regime III	16.54	.20		
Row history, regime IV	105.96	1.00		
Column history, regimes I and V	−4.62	.45	3.56	.09
Column history, regime II	−17.08	1.00		
Column history, regime III	−20.33	.32		
Column history, regime IV	−214.57	1.00		
	B. State = Selfish (1,2)			
Row history, regimes I and V	−1.12	.68	3.14	.05
Row history, regime II	−47.98	1.00		
Row history, regime III	−.17	.98		
Row history, regime IV	106.61	1.00		
Column history, regimes I and V	7.08	.03	2.73	.00
Column history, regime II	86.75	.00		
Column history, regime III	5.88	.56		
Column history, regime IV	−209.85	1.00		
	C. State = Altruist (2,1)			
Row history, regimes I and V	15.73	.03	7.21	.00
Row history, regime II	74.40	1.00		
Row history, regime III	21.64	.33		
Row history, regime IV	197.83	1.00		
Column history, regimes I and V	4.38	.53	.95	.72
Column history, regime II	−27.42	.00		
Column history, regime III	−13.93	.70		
Column history, regime IV	−621.79	1.00		

NOTE.—Column equilibrium (2,2) is the baseline state.

Appendix B

Instructions

The following are the instructions to the Battle of the Sexes Game as they appeared on the computer screen for subjects. They are preceded by a set of general instructions, which explain the overall procedures for the three games each subject was to play. After a subject finished playing this game, he would proceed to another game (unless this was the last game he played).

Since these are generic instructions, things such as conversion rate of experimental currency to dollars have been left blank.

Specific Instructions

Introduction

In this decision problem you will be paired with another person. When your participation in this decision problem is over, you will be replaced by another participant who will take your place in this decision problem. Your final payoff in the entire decision problem will be determined both by your payoff in the decision problem you participate in and by the payoff of your successor in the decision problem he/she participates in.

The currency in this decision problem is called francs. All payoffs are denominated in this currency. At the end of the decision problem your earnings in francs will be converted into real U.S. dollars at a rate of 1 franc = \$*x.xx*.

Your Decision Problem

In the decision problem you participate in there will be — round(s).

In this problem the row chooser must choose a row and the column chooser must choose a column. There are two rows (1 and 2) and two columns (1 and 2) available to choose from, and depending on the choices of the row and column choosers, a payoff is determined. For example, if the row chooser chooses 1 and the column chooser also chooses 1, then the payoffs will be the ones written in the upper left hand corner of the matrix. (Note that the first number is the payoff for the row chooser while the second number is the payoff for the column chooser.) Here the row chooser will earn a payoff of 150 while the column chooser will earn 50. If the row chooser chooses 2 and the column chooser also chooses 2, then the payoffs will be the ones written in the lower right hand corner of the matrix. Here the row chooser will earn a payoff of 50 while the column chooser will earn 150. If 1 is chosen by a row chooser and 2 by a column chooser (or vice versa), each chooser will get a payoff of zero.

To make your decisions you will use a computer. If you are the row (column) chooser and want to choose any specific row (column), all you need to do is use the mouse to click on any portion of the row (column) you wish to choose. This will highlight the row (column) you have chosen. You will then be asked to confirm your choice by being asked:

Are you sure you want to select row (column) 1 (2, 3, etc.)?

When the row and column choosers have both confirmed their choices, the results of your choices will be reported to both choosers. At this point the computer will display your choice, your pair member's choice, and your payoff for that round by highlighting the row and column choices made and having the payoffs in the selected cell of the matrix blink.

Your Payoff and Your Successor

After you have finished your participation in this decision problem, you will be replaced by another participant who will take your place in an identical decision problem with another newly recruited participant. Your final payoff for this decision problem will be determined both by your payoff in the decision problem you participate in and by the payoff of your successor in the decision problem that he/she participates in. More specifically, you will earn the sum of your payoffs in the decision problem you participate in plus an amount equal to one-half of the payoff of your successor in his/her decision problem.

Advice to Your Successor

You will also receive one-half of the payment earned by your successor. Since your payoff depends on how your successor behaves, we will allow you to give advice to your successor in private. The form of this advice is simple. You simply suggest an action, 1, or 2, or 3, etc. for your successor by writing in the advice form below what you think he/she should choose. You are also provided with a space where you can write any comments you have for them about the choice they should make. In addition, you can, if you wish, tell your successor the advice given to you by your predecessor as well as any history of your predecessors which you saw but your successor might not see.

To give advice, click on the "Leave the Advice!" button. You will then see on the screen the following advice form which provides you an opportunity to give advice to your successor.

Note that except if you are the first person ever to do this decision problem, when you sit down at your computer you will see the advice your predecessor gives you.

History

When you sit down at your computer you will also see the history of all previous pairs who have participated in this decision problem before.

To see this history information click on the "History" button located at the bottom of the Advice Box.

Note, finally, that all other successors will also see the advice of their predecessors, and the history of the decision problem that their predecessors participated in. You will not, however, see the advice given to the person you are paired with by his/her predecessor.

Predicting Other People's Choices

At the beginning of the decision problem, before you choose your row or column, you will be given an opportunity to earn additional money by predicting the choices of your pair member in the decision problem. A prediction form will appear when you need to make a prediction as follows: This form allow you to make a prediction of the choice of your pair member by indicating what the chances are that your pair member (the column or row chooser) will choose 1, or 2, or 3, etc. For example, suppose you are a row chooser and you think there is a 40% chance that your pair member will choose 1, and hence a 60% chance that 2 will be chosen. This indicates that you believe that 1 is less likely to be chosen than 2, but that there is still a pretty good chance of 1 being chosen. If this is your belief about the likely choice of your pair member, then click in the space next to the entry 1 and type the number (40). Then click in the space provided next to the entry 2 and type (60). Note that the numbers you write must sum up to 100. For example, if you think there is a 67% chance that your pair member will choose 1 and a 33% chance he/she will choose 2, type 67 in the space next to the entry 1 and 33 in the space next to the entry 2.

At the end of the decision problem, we will look at the choice actually made by your pair member and compare his/her choice to your predictions. We will then pay you for your prediction as follows:

Suppose you predict that your pair member will choose 1 with a 60% chance and 2 with a 40% chance. In that case you will place 60 next to the entry 1 and 40 next to the entry 2. Suppose now that your pair member actually chooses 2. In that case your payoff will be

$$\text{Prediction Payoff} = [20{,}000 - (100 - 40)^2 - (60)^2].$$

In other words, we will give you a fixed amount of 20,000 points from which we will subtract an amount which depends on how inaccurate your prediction was. To do this when we find out what choice your pair member has made (i.e. either 1 or 2), take the number you assigned to that choice, in this case 40 on 2, subtract it from 100 and square it. We will then take the number you assigned to the choice not made by your pair member, in this case the 60 you assigned to 1, and square it also. These two squared numbers will then be subtracted from the 20,000 francs we initially gave you to determine your final point payoff. Your point payoff will then be converted into francs at the rate of 1 point = — francs.

Note that the worst you can do under this payoff scheme is to state that you believe that there is a 100% chance that a certain action is going to be taken and assign 100 to that choice when in fact the other choice is made. Here your payoff from the prediction would be 0. Similarly, the best you can do is to guess correctly and assign 100 to that choice which turns out to be the actual choice chosen. Here your payoff will be 20,000.

However, since your prediction is made before you know what your pair member actually will choose, the best thing you can do to maximize the expected size of your prediction payoff is to simply state your true beliefs about what you think your pair member will do. Any other prediction will decrease the amount you can expect to earn as a prediction payoff.

Summary

In summation, this decision problem will proceed as follows. When you sit down at the terminal you will be able to see the decisions that have been made by the previous pairs who have participated in this decision problem, and you will be able to see the advice that your immediate predecessor has given you. You will then be asked to predict what your pair member will do by filling out the prediction form. After you do that, the decision box will appear on the screen and you will be prompted to make your decision. You will then be shown the decision made by the person you are paired with, and you will be informed of your payoff. Finally, you will fill out the advice form for your successor.

References

Berg, Joyce; Dickhaut, John W.; and McCabe, Kevin A. "Trust, Reciprocity, and Social History." *Games and Econ. Behavior* 10 (July 1995): 122–42.

Bisin, Alberto, and Verdier, Thierry. "'Beyond the Melting Pot': Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits." *Q.J.E.* 115 (August 2000): 955–88.

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

Cavalli-Sforza, Luigi L., and Feldman, Marcus W. *Cultural Transmission and Evolution: A Quantitative Approach.* Princeton, N.J.: Princeton Univ. Press, 1981.

Cooper, Russell; DeJong, Douglas V.; Forsythe, Robert; and Ross, Thomas W. "Communication in the Battle of the Sexes Game: Some Experimental Results." *Rand J. Econ.* 20 (Winter 1989): 568–87.

———. "Communication in Coordination Games." *Q.J.E.* 107 (May 1992): 739–71.

Crémer, Jacques. "Cooperation in Ongoing Organizations." *Q.J.E.* 101 (February 1986): 33–49.

Fudenberg, Drew, and Maskin, Eric. "Evolution and Cooperation in Noisy Repeated Games." *A.E.R. Papers and Proc.* 80 (May 1990): 274–79.

Jackson, Matthew O., and Kalai, Ehud. "Social Learning in Recurring Games." *Games and Econ. Behavior* 21 (October–November 1997): 102–34.

Kandori, Michihiro. "Repeated Games Played by Overlapping Generations of Players." *Rev. Econ. Studies* 59 (January 1992): 81–92.

Kandori, Michihiro; Mailath, George J.; and Rob, Rafael. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61 (January 1993): 29–56.

Nyarko, Yaw, and Schotter, Andrew. "An Experimental Study of Belief Learning Using Elicited Beliefs." *Econometrica* 70 (May 2002): 971–1005.

Ottaviani, Marco, and Sorensen, Peter. "Professional Advice." Manuscript. London: Univ. Coll. London, Dept. Econ., 1999.

Quandt, Richard E. "The Estimation of the Parameters of a Linear Regression System Obeying Two Separate Regimes." *J. American Statis. Assoc.* 53 (December 1958): 873–80.

Salant, David J. "A Repeated Game with Finitely Lived Overlapping Generations of Players." *Games and Econ. Behavior* 3 (May 1991): 244–59.

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

Samuelson, Larry, and Zhang, Jianbo. "Evolutionary Stability in Asymmetric Games." *J. Econ. Theory* 57 (August 1992): 363–91.

Schotter, Andrew. *The Economic Theory of Social Institutions.* Cambridge: Cambridge Univ. Press, 1981.

Schotter, Andrew, and Sopher, Barry. "Social Learning and Coordination Conventions in Inter-generational Games: An Experiment in Lamarckian Evolutionary Dynamics." Economic Research Report no. 2000-01. New York: New York Univ., Starr Center Appl. Econ., January 2000.

Ullmann-Margalit, Edna. *The Emergence of Norms.* Oxford: Clarendon, 1977.

Vega-Redondo, Fernando. *Evolution, Games, and Economic Behaviour.* Oxford: Oxford Univ. Press, 1996.

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

Young, H. Peyton. "The Evolution of Conventions." *Econometrica* 61 (January 1993): 57–84.

———. "The Economics of Convention." *J. Econ. Perspectives* 10 (Spring 1996): 105–22.

———. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions.* Princeton, N.J.: Princeton Univ. Press, 1998.