

Schotter, Andrew; Sopher, Barry

Working Paper

Social learning and coordination conventions in inter-generational games: An experimental study

Working Paper, No. 2000-21

Provided in Cooperation with:

Department of Economics, Rutgers University

Suggested Citation: Schotter, Andrew; Sopher, Barry (2001) : Social learning and coordination conventions in inter-generational games: An experimental study, Working Paper, No. 2000-21, Rutgers University, Department of Economics, New Brunswick, NJ

This Version is available at:

<https://hdl.handle.net/10419/94234>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Social Learning and Coordination Conventions in Inter-Generational Games: An Experimental Study

Andrew Schotter and Barry Sopher*

Revised: November 22, 2001

Abstract

This is a paper on the creation and evolution of conventions of behavior in "inter-generational games". In these games 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 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.

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, far stronger than the type of learning subjects seem capable of doing simply by learning the lessons of history without the guidance offered by such advice.

JEL Classification: C91, C72

*This work was completed under N.S.F. 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 benefitted greatly from presentation at the McArthur Foundation, The Russell Sage Foundation, and seminars at The Wharton School, Harvard University, Washington University at St. Louis, The University of Delaware, C.U.N.Y. Graduate Center, the Economic Science Association, and the University of Pittsburgh. We would like to thank Steven Levitt 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 upon which the experiments were run.

1. Introduction

This is a paper on the creation and evolution of conventions of behavior in “inter-generational games”. In these games 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 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 each generation’s payoff 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.¹²

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 inter-act 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 of the information about the norms and conventions that have been established to their successors.

¹We use a non-overlapping generation structure and not an overlapping generations one because in most overlapping generation games of this type (see Salant (1991), Kandori (1989), Cremer (1986)) cooperation is achieved by each generation realizing 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 them and acts properly now knowing that they are being observed and will inter-act directly with that generation. In this literature, folk-like theorems are proven if the length of the overlap between generations is long enough. In our work, however, generations never overlap. What they 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 where the overlap is either minimal or non-existent..

²Except for the use of advice and the inter-dependence of our generational payoffs, our game has many of the features of Kalai and Jackson’s (1996) Recurring Games.

As we will 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 Edna Ullman-Margalit (1977) calls a "norm of 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. 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 etc. All of these confer unequal benefits to agents in the economy and perpetuate a status quo. Other examples include occupational segregation by sex where, certain occupations come to be predominantly female, seniority rules where 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 where 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 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 where generations of social agents pass on conventions of behavior they create during their lifetime to their successors.³⁴Such conventions may reinforce social inequality

³Our emphasis on this Lamarckian evolutionary process is in contrast to practically all work in evolutionary game theory which is predominantly Darwinian (see, for example, Kandori, Malait and Rob (1993) , Samuelson (1997), Vega-Redondo (1996) and Weibull (1995) 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 non-thinking programs (genes) hard-wired to behave in a particular manner. These agents either inter-act randomly or "play the field". The dynamics of the growth and or decay of these strategies is governed by some type of replicator-like dynamic (see Weibull (1995)) in which those strategies which receive relatively high payoffs increase in the population faster than those which 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 (ESS)? Are there cycles in which different strategies over run 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 inter-generational games except for the inter-generational communication and caring.

⁴Of course this point has already been made by Boyd and Richerson , (1985) , Cavalli Sforza

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 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, while 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 while 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 their ability to coordinate. Hence, in our inter-generational setting, it appears as if 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, i.e., 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 will proceed as follows: Section 2 presents our experimental design. In Section 3 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 4 is about social learning. It describes what happens in our experiments when we eliminate our subject's ability to pass on advice or see the history of their predecessors. Finally, in Section 5 we offer some conclusions and speculations for future work.

and Feldman (1981) and more recently Bisin (1998), all of whom have presented a number of interesting models where imitation and socialization, rather than pure absolute biological fitness, is the criterion upon which strategies evolve. We would include Young's (1996, 1998) work in this category as well.

2. The Experiment: Design and Procedures

2.1. General Features

The general features of our inter-generational 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 are replaced by a next generation, $t+1$, who were able to view some or all of the history of what has 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 is being given. This feature obviously permits socialization. The payoffs to any subject in the experiment were equal to the payoffs earned by that generation during their lifetime plus a discounted payoff which depends 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 which makes telling the truth a dominant strategy). This is 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 or at 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 $1\frac{1}{2}$ hours, they engaged in three separate inter-generational games, a Battle of the Sexes Game (BOSG), an Ultimatum Game (UG) in which they were asked to divide 100 francs, and a Trust Game (TG) 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 only present the results of the following Battle of the Sexes Game.⁵

⁵The 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 while Players 2 and 5 play the Ultimatum Game and Players 3 and 4 play the Trust game. When they have finished their respective games, we would

Battle of the Sexes Game

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

As is true in all BOSG's, this game has two pure strategy equilibria. In one, (1,1), player 1 does relatively well and receives a payoff of 150 while 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 which is exacerbated by our generational structure since new generations may not want to adhere to a convention established in the past which is unfavorable to them. (There is also a mixed strategy equilibrium which we will 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 $1\text{fr} = \$.04$.

When subjects started to play the Battle of the Sexes Game, 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 was in 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 they suggested what they 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 each generational player 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, subjects were paid 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

rotate them in the next period so that in period 2 Players 2 and 4 play the Battle of the Sexes Game while Players 3 and 6 play the Ultimatum Game and Players 1 and 5 play the Trust game. The same type of rotation is carried out in period 3 so that at the end of the experiment each subjects has played each game against a different opponent who has not played with any subject he has played with before.

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 or her pure strategies.⁶ We rewarded subjects for their beliefs in experimental points which are converted into dollars at the end of the experiment as follows:

First subjects report their beliefs by entering a vector $\mathbf{r} = (r_1, r_2)$ indicating their belief about the probability that the other subject will use strategy 1 or 2.⁷ Since only one such strategy will actually be used, the payoff to player i when strategy 1 is chosen by a subject's opponent and \mathbf{r} is the reported belief vector of subject i will be:

$$\pi_1 = 20,000 - \left\{ ((100 - r_1)^2 + (r_2)^2) \right\}. \quad (2.1)$$

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

$$\pi_2 = 20,000 - \left\{ (100 - r_2)^2 + (r_1)^2 \right\}. \quad (2.2)$$

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 their opponent chooses 1, then the subject would have been best off if he or she had put all of their 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, i.e. predicting a particular pure strategy only to have your opponent choose another, yields a payoff of 0. 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.⁸ 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

⁶See Appendix 1 for the instructions concerning this part of the experiment.

⁷In the instructions r_j is expressed as numbers in $[0,100]$, so are divided by 100 to get probabilities.

⁸An identical elicitation procedure was used successfully by Nyarko and Schotter (1999).

played. (In fact, over the entire experiment subjects earned, on average, \$26 while 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.

2.2. 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 allowed to see. The second is inter-generational 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 they repeat the game) while the fourth indicates whether advice is allowable between generations. In all of our experiments each generation lives for one period or repeats the trust game only once and has a discount rate of $1/2$. Hence, they only differ 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 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 inter-generational game. The History-only Treatment, Treatment II, was identical to the Baseline except for the fact that no generation was able to pass advice onto their successors. They could see the entire history, however, so that this treatment isolated the impact of history. Advice-only Treatment was run for an additional 80 generations while the History-only Treatment was run for an additional 66 generations, each starting after generation 52 was completed in the Baseline. Hence, our Baseline was of length 81, the Advice-only Treatment was of length 78⁹ and the History-only

⁹One 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

Treatment was of length 66. Our experimental design can be represented by Figure 1:

[Figure 1 here]

3. Results

We will 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.

3.1. Stylized Facts of Social Evolution

The stylized facts of social evolution which 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 Peyton Young (1996) makes clear, over long periods of time one observes periods of punctuated equilibria where 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.

¹⁰

In our experiments departures from equilibria are sometimes caused by the advice handed down from one generation to the next. As we will see, there are times during the experiment where a convention appears to be relatively firmly established and yet there will be generational advice advocating a departure. In addition, there will be periods where 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

¹⁰There are a number of reasons for the disruption of these conventions. In Darwinian models of evolution random mutations can arise which, if persistent enough, can cause a disruption of the current equilibrium and drift towards a new one (see Kandori, Mailath, and Rob (1996), Young (1993), Fudenberg and Maskin (1990), Samuelson and Zhang (1992), and Samuelson (1991). 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 last temporarily, of existing conventions of behavior.

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 socialization in which present generations teach and pass on current conventions of behavior to the next generation. ¹¹

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 as they are often followed unthinkingly while 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.

3.2. 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 will then go on to investigate behavior in Treatments 1 and 2.

3.2.1. Actions in the Baseline Experiment: Punctuated Equilibria

Figure 2 presents the time series of actions generated by our 81 generation Baseline experiment.

Figure 2 Here

¹¹Replicator dynamics attempt this inter-generational transmission in a very specific and non-human manner but as a descriptive theory of social reality such a theory is quite poor. Other theories of social evolution, [see Boyd and Richerson (1985), Cavalli Sforza and Feldman (1981), and Bisin and Verdier (1998)] 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 will only model vertical as opposed to horizontal socialization. Still, what we see in front of us in the real world are 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.

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 $o_{11} = (\text{row}_1, \text{column}_1)$, $o_{12} = (\text{row}_1, \text{column}_2)$, $o_{21} = (\text{row}_2, \text{column}_1)$, $o_{22} = (\text{row}_2, \text{column}_2)$, where o_{ij} indicates an outcome where the row player chose action i and the column player action j . (We will denote these states as states 1, 2, 3, and 4, respectively).

To give a greater insight into the data we have divided the 81 generations into *regimes* based on a simple inspection of the outcomes observed. Regime 1 consists of period 1-25, regime 2, generations 26-45, regime 3, generations 46-66 and regime 4, generations 67-81. To support this separation we estimated a multinomial logit model describing the probability of being in any state. This model demonstrates that when we cross the boundaries of our proposed regimes, the probabilities of being in the various states change dramatically and in a manner consistent with our proposed regimes. This analysis will be presented below but first let us look more carefully at Figure 2.

Regime I (generations 1-25) 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 where we observed nine consecutive periods of (2,2), the longest run for any stage-game equilibrium in all 81 generations of the Baseline. Regime II (generations 25-45) we call the (1,1) Convention Regime because while in the first 25 generations we only saw the (1,1) equilibrium chosen twice, in Regime II it is chosen in 11 of the 21 generations. In addition, during this time the (2,2) equilibrium, which was so prevalent in Regime I, appears only once. If we look at the row players in this Regime II, they choose strategy 1, in 17 of the 21 generations indicating that at least in their minds they are adhering to the (1,1) convention in playing this game. Regime III (generations 46-66) we call a transition regime 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 (two and three respectively). It is interesting to note that during this time the row player is starting to play strategy 2 more frequently (choosing it 6 out of 21 times as opposed to 4 out of 21 times in Regime II). Finally, Regime IV (generations 67-81), appears to present evidence that the (2,2) equilibrium is reestablishing itself as a convention after a virtual absence over 42 generations. We say this because during these last 15 rounds we see the (2,2) equilibrium appearing in 10 out of 15 generations while it only appeared four times in the previous 42 rounds. Even more surprising, the row players, after a great resistance to playing row 2, (e.g., they only played it 10 times in 42 generations between generation 25

and 66), chose it 11 times in the last 15 rounds. 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 that the (2,1) state arises (7 times) as opposed to the (1,2) state (27 times). These results are tabulated in Table 1.

[Table 1 Here]

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 re-establish itself with any regularity until generations 67-81 (Regime IV). 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 10 (1,1) choices in 21 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 IV, where it appears that the (2,2) convention has reestablished itself also presents interesting evidence of the punctuated equilibrium phenomenon.

A formal test of the proposition that there are different conventions in place in the regimes that we propose above can be constructed on the basis of an estimated econometric model of the process determining the state. To do this, consider a multinomial logit response model in which the probability of state h occurring 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 that there is a structural change that occurs in the model. That is, a sudden change (punctuation) 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 across our proposed regimes, the resulting “unrestricted” 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. For econometric purposes, we define only three distinct structural regimes: one for Regimes I and IV where the (2,2) outcome is dominant (indexed by $k = 1$), one for Regime II where the (1,1) outcome is dominant

(indexed by $k = 2$), and one for Regime III, where neither equilibrium outcome dominates (indexed by $k = 3$). A standard likelihood ratio test between the restricted and unrestricted models serves as a test of structural change.

To substantiate our choice of regimes in a more formal manner, 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 period. Using a multinomial logit form for this probability yields,

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

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

where $k \in K = \{1, 2, 3\}$ indexes the 3 structural regimes defined above, J is the set of states indexed by j , 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 1 if the observed state in period t is in structural regime k , and equal zero otherwise. Then $r_{t,k} = f_{t,r}^m d_k$. The column history variables, $c_{t,k}$, are defined similarly. We call the model “restricted” when we impose the following: $b_{1r}^j = b_{2r}^j = b_{3r}^j$ and that $b_{1c}^j = b_{2c}^j = b_{3c}^j$, implying that there are no structural regimes. The “unrestricted” model allows these coefficients to vary across structural regimes.

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-squared measure.¹² We then proceeded to estimate the unrestricted model, using the five-period moving average, and computed the likelihood ratio test for the restrictions.

The restricted and unrestricted estimates are contained in Table A1 in the appendix. Since the coefficients are not of particular interest here, we focus on the

¹²Specifically, the pseudo R-squared = .17 for the five period moving average, and it declined as one either lengthened or shortened the moving average window, where we allowed the window to range from $m=1$ to $m=10$.

goodness of fit measures and the likelihood ratio test for the structural change hypothesis. Table 2 summarizes these results ¹³ Note that each model is significant: the model chi square rejects the hypothesis that all coefficients but the constant terms are zero in each case. The likelihood ratio chi square for the structural change hypothesis is also significant: the restriction that coefficients are equal across regimes is rejected at the 4% significance level.

[Table 2 Here]

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.

Figure 3 Here

This figure presents strong evidence that behavior changed dramatically as generations moved across our regimes. For example, note that in our Regime I the probability of being in state (2,2) approaches .90 for a substantial portion of time and then drops precipitously as the experiment enters our Regime II. In fact, during almost the entire length of Regime II that probability is practically zero indicating, as we did above, that the (2,2) equilibrium abruptly disappears in Regime II. Note however, the re-emergence of the (2,2) state in our Regime IV.

The evidence for the (1,1) state is less strong but still 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 to almost .50 for regime II only to disappear once more for the remainder of the experiment. 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 our Regime II we claim as support for our division of generations into regimes using our "eyeball metric".

3.2.2. Inertia and norms of inequality

With respect to inertia, there are really two types of social inertia one can discuss. One, which we will call equilibrium inertia, is the inertia that leads people

¹³Since there are 3 equations to estimate (one less than the total number of states) there are, besides the constant terms, 2 (row and column history) x 3 (equations) x 3 (regimes) = 18 coefficients to estimate in the unrestricted model, and 2 (row and column history) x 3 (equations) = 6 coefficients to estimate in the restricted model, apart from the constant terms. Thus the goodness of fit chi-squared statistic has 18 degrees of freedom for the unrestricted model and 6 degrees of freedom for the restricted model, while the chi-squared statistic for the test of structural change has 18-6=12 degrees of freedom.

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 a 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 where each player thinks that his or her opponent is going to adhere with probability 1, then deviating can never be beneficial since if you continue to adhere you will get 50 today plus one half of fifty tomorrow, while deviating will yield 0 today and if 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 the player 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 inter-generational utility function. (We will be able to explain this disappearance later when we talk about advice).

Note that such conventions establish what Edna Ullman-Margalit (1977) calls “norms of partiality” in which seemingly symmetric agents select an equilibrium to a game which 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 are 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 is the result of pressures that arise as one type of agent realizes that the cause of the inequality they face is purely arbitrary and hence not fair. It appears as if subjects deviate to further the cause of their type at a sacrifice to themselves. (Another violation of Darwinian evolutionary theory.)

14

¹⁴This 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 3 or 4 periods to a generations’s lifetime so there is still

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 and III, 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 32 out of 42 generation between generation 25 and 66. They persisted in doing so despite the fact that this behavior led to a disequilibrium outcome in 25 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 you were in that state in period t .

Table 3 presents the probabilities:

[Table 3 Here]

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 a 0.81 probability of remaining in the (2,2) state if one reached it. In Regime II, while the (1,1) state was observed 11 out of 21 times, many of these instances were isolated instances that were not repeated. Still, the continuation probability was 0.30. More remarkable is the fact that none of the other states ever repeated themselves during the entire regime. Regime III demonstrated a dramatic ability to remain in the disequilibrium state (16 out of 21 times) with no persistence to the (2,1) state but a continuation probability for the (1,2) state of 0.50. Finally, Regime IV showed the return of the (2,2) state and its persistence (5 out of 9 times) while no other state appeared to have any durability whatsoever.

3.2.3. 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.

work to be done here.

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

[Tables 4 and 5 Here]

What Advice Was Given¹⁵ 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% 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, while it occurs 90% 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), that state which is best for the column player, we see 100% of the column players (30 out of 30) suggesting a choice of 2, while when the state is (1,1) 87.5% 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.

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 that equilibrium which was best for his or her opponent. We call this the **submissive disequilibrium** state since both subjects yielded to the other and chose that state which 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 while the column player suggested switching and choosing 1 in five of the seven cases. When the

¹⁵In Schotter and Sopher (2000a) 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 a subject, i.e., they received a 150 payoff, they tended to leave low level messages which were not supported by strategic reasoning. However, the subject receiving the low payoff tended to leave more highly reasoned and strategic advice.

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 while 9 suggested standing pat and choosing row 1. For the column players 15 suggested switching to the submissive strategy (column 1), while 12 suggested standing pat and continuing to choose column 2.

When Was Advice 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 and 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 be followed quite often but the degree to which it is followed varies depending on the state last period. On average, for the row players it is followed 68.75% of the time while for the column player it was followed 70% of the time. When the last period state was (2,2), row players followed the advice given to them 73.3% of the time (strangely agreeing to follow advice to switch to the row 1 strategy three out of the three times), while column subjects followed 86.6% 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% of the time while row player adhered 68% of the time.

One question that arises here is how powerful is advice 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 what 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 Tables 6a and 6b:

[Tables 6a and 6b Here]

What we can conclude from these tables 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 where the best response prescription was different than the advice given and 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 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 possesses (the only informational difference between a generation t and 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 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 subjects' 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, 3 of the 30 rounds where the (2,2) equilibrium was in place, the row player chose not to give advice to his successor to adhere to it, while in 2 of 16 instances where 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 and, as we will see, leads us to model it as an irreducible finite state Markov chain.

3.3. Beliefs

As described above, before each generational subject makes his or her choice, they were asked to state their beliefs about what they felt the probability was that their opponent would use strategies 1 or 2. The time path of these belief vectors are presented in Figures 4a and 4b where we present the probability that

each generational subject felt his opponent would choose strategy 1.

[Figure 4a and 4b here]

Note that in Figures 4a and 4b we have placed a straight line which indicates the critical belief value which is such that if the beliefs that your opponent is going to choose strategy 1 with a higher probability than critical value, a best response for you is to choose strategy 1 as well. (We have also placed a curved line which we will explain shortly but which we will ignore at the moment). As we see, beliefs of both subjects seem to exhibit a type of over-confidence bias in the sense that overwhelmingly both subjects appear to believe that their opponent is going to choose that strategy which is consistent with that equilibrium which is best for them. More precisely, in only 26 on the 81 generations did row subjects believe that their opponent was so likely to choose row 2 so as to lead them to choose 2 as a best response. For column players the situation was even worse with beliefs only consistent with 15 row-1 best responses. Obviously, if these beliefs are based on the history of play of the game, each can not 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 (1998)). 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 while the column player 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 ($p = 0.00$), ($z = 4.93$). There appears to be some convergence, however, since in Regimes III and IV these same Signed-Rank tests fail to reject the hypothesis that the distributions are equal (Regime III: $z = 1.34$, p -value 0.18: Regime IV: $z = -.34$, p -value 0.73).

For column players, a Signed-Ranks 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 0.70) of the experiment or in any of the Regimes, (Regime I, $z = 0.70$, p-value 0.48, Regime II, $z = 1.55$, p-value 0.12, Regime III, $z = -1.16$, p-value 0.24, Regime IV, $z = -1.36$, p-value 0.17).

4. 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 only see the play of the generation before them. They could, however, receive advice just as did subjects 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.¹⁶

These treatments furnish a controlled experiment which 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 which involves agents taking actions which 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 their 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 do 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 which the receiver has not

¹⁶This was done by forbidding them to write any instructions on the screen despite the fact that they were prompted to.

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.

Figures 5a-5c plots the time series generated by these two treatments along with our original Baseline treatment (a repeat of Figure 2).

[Figures 5a-5c here]

As we can see from these graphs, removing history has a very different impact on the path of play than does removing advice. Consistent with what we have noticed above, players in inter-generational 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 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 $\frac{20}{39} = .512$). 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. While in the Baseline we observe equilibrium outcomes 47 out of 81 time, when we eliminate advice, as we do in Treatment II, we only observe coordination in 19 out of 66 periods. When we allow advice but remove history, Treatment I, coordination is restored and occurs in 39 out of 81 generations.¹⁷

¹⁷A 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 if they were generated by the same stochastic process generating the data observed in Treatments I and II. More precisely, treating the data as if it were generated by a one-state Markov chain, for each experiment we can estimate the probability of transiting from any of our four states $\{(1,1), (1,2), (2,1) \text{ and } (2,2)\}$ to the other. A simple counting procedure turns out to yield maximum-likelihood estimates of these transition probabilities. Doing so would generate a 4 x 4 transition matrix for each experimental treatment. These transition matrices are presented

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 parent's advice almost 50% of the time when the advice differs from the best response to their own beliefs.¹⁸

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, Dejong, Forsythe and Ross (1989), it appears to aid coordination in the sense that the amount of equilibrium occurrences in our Baseline (58%) and Treatment I (49%) where advice was present is far greater than that of Treatment II (29%) 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, Dejong, Forsythe and Ross (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 contained in the history of the game, but are instead following the advice given to them by their predecessor who is as just about much a neophyte as they are.

in the Appendix to Schotter and Sopher (2000a).

To test if 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 T the transition matrix estimated from our Baseline data and P^k the transition matrix defined by our k^{th} treatment, i.e., $k = \{I, II\}$. Denote $p_{ij}^{P^k}, j = i = \{1, 2, 3, 4\}$, as the transition probability from state i to state j in matrix 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 chi-square test (See Schotter and Sopher (2000a) 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 ($\chi^2(12df) = 27.6521, (p= 0.000)$) or Treatment II ($\chi^2(9df) = 59.4262, (p= 0.000)$). Hence, if the process generating our data can be considered Markovian, it would appear as if imposing different informational conditions on the subject significantly changed their behavior.

¹⁸There is no sense, then, that parents in our experiment are in any way "experts" as in the model of Ottaviani and Sorensen (1999).

5. Conclusions

This paper, utilized an experimental approach to investigate the process of convention creation and transmission in inter-generational games. It has modeled the process as a Lamarckian one in which non-overlapping 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, 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 while 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.

6. Bibliography

References

- [1] Berg, J., Dickhaut, J. and McCabe, K., (1995), "Trust, Reciprocity, and Social History", *Games and Economic Behavior*, Vol. 10, pp. 122-142.
- [2] Boyd, R. and Richerson, P.J., 1985, *Culture and Evolutionary Process*, University of Chicago Press, Chicago, Illinois.
- [3] Bisin, A. and Verdier, T., "Cultural Transmission, Marriage and the Evolution of Ethnic and Religious Traits", Economic research report, C. V. Starr Center for applied Economics, New York University, RR# 98-39, November 1998.
- [4] Cavalli Sforza, L and Feldman, M., (1981), *Cultural Transmission and Evolution: A Quantitative Approach*, Princeton University Press, Princeton, N.J.

- [5] Cooper, R., Dejong, D., Forsythe, R., and Ross, T., (1989), "Communication in the Battle of the Sexes Game: Some Experimental Results", *Rand Journal of Economics*, vol. 20, pp. 568-587.
- [6] Cooper, R., Dejong, D., Forsythe, R., and Ross, T., (1992), "Communication in Coordination Games", *Quarterly Journal of Economics*, vol. 107, pp. 738-771.
- [7] Crawford, V., (1991), "An Evolutionary Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination", *Games and Economic Behavior*, vol. 3, pp. 25-59.
- [8] Cremer, J., (1986), "Cooperation in Ongoing Organizations", *Quarterly Journal of Economics*, vol. 101, pp., 33-49.
- [9] Fudenberg, D. and Harris, (1992) , "Evolutionary Dynamics in Games With Aggregate Shocks", *Journal of Economic Theory*, vol. 57, pp. 420-441.
- [10] Fudenberg, D., and Maskin, E., (1990), "Evolution and Cooperation in Noisy Repeated A Games", *American Economic Review*, vol. 80, pp. 274-279.
- [11] Jackson, M. and Kalai, E., (1997), "Social Learning in Recurring Games", *Games and Economic Behavior*, Vol. 21, pp. 102-134.
- [12] Kandori, M., (1992) "Repeated Games Played by Overlapping Generations of Players", *Review of Economic Studies*, Vol. 59, pp., 81-92.
- [13] Kandori, M., Mailath, G., and Rob, R., (1993), "Learning, Mutation, and Long Run Equilibria in Games", *Econometrica*, vol. 61, No. 11, pp. 29-56.
- [14] Lewis, D., (1969) *Convention: A Philosophical Study*, Harvard University Press, Cambridge Massachusetts..
- [15] Ottaviani, M. and Sorensen, P., (1999), "Professional Advice", Mimeo, Department of Economics, University College London.
- [16] Nyarko, Y. and Schotter, A., "An Experimental Study of Belief Learning Using Real Beliefs" Economic Research Report 98-39, C. V. Starr Center for Applied Economics, New York University, December 1998.
- [17] Okuno-Fujiwara, M. and Postlewaite, (1995) A., "Social Norms and Random Matching", *Games and Economic Behavior*, Vol. 9, pp. 79-109.

- [18] Salant, D.,(1988), "A Repeated Game with Finitely Overlapping Generations of Players", *Games and Economic Behavior*,
- [19] Samuelson, L., (1997), *Evolutionary Games and Equilibrium Selection*, MIT Press, Cambridge, MA.
- [20] Samuelson, L. and Zhang, J, "Evolutionary Stability in Asymmetric Games", *Journal of Economic Theory*, Vol. 57, pp 363-391.
- [21] Schotter, A. (1981), *The Economic Theory of Social Institutions*, Cambridge University Press, Cambridge England.
- [22] Schotter, A. and Sopher B. "Social Learning and Coordination Conventions in Inter-generational Games: An Experiment in Lamarckian Evolutionary Dynamics", *Economic Research Report #2000-01*, C.V. Starr Center for Applied Economics, New York University, January 2000.
- [23] Ullman-Margalit, E.,(1977), *The Emergence of Norms*, Oxford University Press, Oxford, England.
- [24] VanHuyck, J., Battalio, R., and Beil, R., (1990), "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure", *American Economic Review*, Vol. 80, pp. 234-248.
- [25] Vega-Redondo, F. (1996), *Evolution, Games and Economic Behavior*, Oxford University Press, Oxford, England.
- [26] Weibull, J., (1995), *Evolutionary Game Theory*, MIT Press, Cambridge, Massachusetts..
- [27] Young, H.P., (1993), "The Evolution of Conventions", *Econometrica*, vol. 61, No. 1, pp. 57-84.
- [28] Young, H.P., (1996) , "The Economics of Conventions", *Journal of Economic Perspectives*, vol. 10, No. 2, pp. 105-122
- [29] Young, H.P., (1998), *Individual Strategy and Social Structure*, Princeton University Press, Princeton New Jersey.

**Table 1: Choices of Row and Column
Player by Regime**

Choices by States and Regime

Regimes	(1,1)	(1,2)	(2,1)	(2,2)	Total
I	2	5	0	17	24
II	11	6	3	1	21
III	2	13	3	3	21
IV	1	3	1	10	15
Total	16	27	7	31	81

Choices by Regime

Regime	Row 1	Row 2	Column 1	Column 2
I	7	17	2	22
II	17	4	14	7
III	15	6	5	16
IV	4	11	2	13
Total	43	38	23	58

Table 2: Test for Structural Change

	Unrestricted Model	Restricted Model
N	76	76
Model X^2	54.15 (18 d.o.f.)	32.26 (6 d.o.f.)
Prob> X^2	.00	.00
Log Likelihood	-70.30	-81.25
Pseudo R^2	.28	.17
LR Test for Structural Change: $X^2 = 21.89$ (12 d.o.f.)		
Prob> X^2 .04		

Table 3: Continuation Probabilities by Regime

	1,1	1,2	2,1	2,2
Regime I:	0	0.166	NA*	0.812
Regime II:	0.3	0	0	0
Regime III:	0	0.50	0	0.33
Regime IV:	0	0	0	0.555
Total	0.187	0.259	0	0.633

*No (2,1) state occurred in Regime I

Table 4: Advice Offered Conditional of 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

Table 5: Advice Adherence Conditional on Last Period's State

State Last Period: (1,1)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
	1	11	5	9
Advice	2	0	0	1
Total		11	5	10

State Last Period: (1,2)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
	1	7	2	10
Advice	2	10	8	10
Total		17	10	20

State Last Period: (2,1)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
	1	5	2	0
Advice	2	0	0	4
Total		5	2	4

State Last Period: (2,2)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
	1	3	0	0
Advice	2	19	8	26
Total		22	8	26

Table 6a: Following Advice When Advice and Best Responses Differ

	Row		Column	
	Follow	Reject	Follow	Reject
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
	15	13	17	17

Table 6b: Following Advice When Advice Equals Best Responses

	Row		Column	
	Follow	Reject	Follow	Reject
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
	40	12	39	7

Table A1: Coefficient Estimates for Test for Structural Change*

	Unrestricted Model	Restricted Model
State=Row Equilibrium (1,1)	Estimate (Prob.> z)	Estimate (Prob.> z)
Constant term	-2.18 (.03)	-3.21 (.00)
Row History, regime 1	2.68 (.39)	3.27 (.08)
Row History, regime 2	2.70 (.65)	
Row History, regime 3	6.09 (.09)	
Column History, regime 1	-3.40 (.52)	3.56 (.09)
Column History, regime 2	4.03 (.60)	
Column History, regime 3	-9.69 (.35)	
State= Selfish (1,2)		
Constant term	-2.01 (.02)	-2.35 (.00)
Row History, regime 1	-0.12 (.97)	3.14 (.05)
Row History, regime 2	-1.48 (.81)	
Row History, regime 3	4.87 (.04)	
Column History, regime 1	4.63 (.23)	2.73 (.00)
Column History, regime 2	8.02 (.31)	
Column History, regime 3	0.76 (.85)	
State=Altruist (2,1)		
Constant term	-4.80 (.03)	-5.87 (.00)
Row History, regime 1	0.58 (.92)	7.21 (.00)
Row History, regime 2	11.51 (.15)	
Row History, regime 3	9.90 (.01)	
Column History, regime 1	8.69 (.25)	0.95 (.72)
Column History, regime 2	-6.44 (.58)	
Column History, regime 3	-10.22 (.34)	

* Column Equilibrium (2,2) is the baseline state

Figure 6.1:

Figure 1: Experimental Design





Figure 2: Baseline Outcomes



Outcomes: 1 = (Row chooses 1, Column chooses 1)
2 = (Row chooses 1, Column chooses 2)
3 = (Row chooses 2, Column chooses 1)
4 = (Row chooses 2, Column chooses 2)

Figure 3: Estimated State Probabilities




Figure 4a: Row's Beliefs About Column - Probability Column Chooses 1




Figure 4b: Column's Beliefs About Row, Baseline

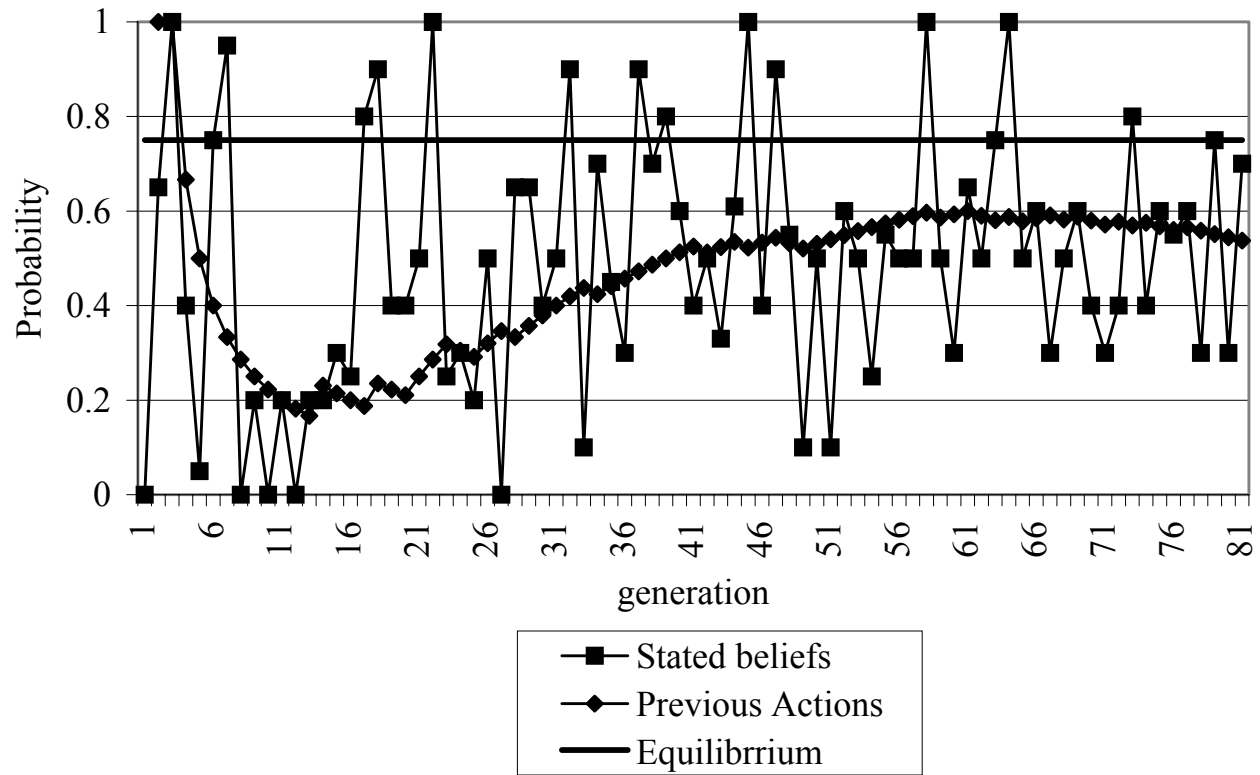


Figure 5a: Baseline Outcomes




Figure 5b: Treatment I Outcomes




Figure 5c: Treatment II Outcomes

