

Social Learning and Coordination Conventions in Inter-Generational Games: An Experiment in Lamarckian Evolutionary Dynamics

Andrew Schotter and Barry Sopher[Ⓜ]

January 2, 2000

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 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.

[Ⓜ]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, Washington University at St. Louis, The University of Delaware, C.U.N.Y. Graduate Center, The Economic Science Association, and the University of Pittsburgh. In addition, the authors 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. (This might be like a CEO who has a long-term compensation package which extends beyond the date of his or her retirement and is based on the performance of the firm (and the succeeding CEO) after that date).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 (i.e. there may always be a G.M. and a Ford playing a duopoly game against each other or super powers playing a geo-political game

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.

with each other) the agents who play these games are ...nitely 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. For example, in a duopoly, at any time each firm is run by a C.E.O. who is in charge of the strategy choices for the firm. When she retires the C.E.O. instructs her replacement as to what to expect from the other firm etc. When these transitions take place, each C.E.O. transmits all of the information about the norms and conventions that have been established by the firms in their previous inter-action. The "culture" of the market is passed on in a Lamarckian manner in the sense that conventions created during one generation can be passed on to the next through a process of socialization just as Lamarck (incorrectly) thought that physical characteristics could be acquired and then passed on in a non-genetic manner.³ We are interested in these transitions and the evolutionary dynamics they imply.⁴

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

³Of course this point has already been made by Boyd and Richerson , (1985) , Cavalli Sforza 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.

⁴Our 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.

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. We also present a model called the "The Bounded-Memory Advice Giving and Following Model" which captures what we feel are the salient features of the advice giving and receiving behavior we observed in our Baseline experiment. Section 4 is about social learning. It starts out describing what happens in our experiments when we eliminate our subject's ability to pass on advice or see the history of their predecessors. It then presents a set of simple models all of which attempt to capture the behavior generated by our experiments and characterize it as a Markov chain. We then present a set of simple tests designed to select across these models. Finally, in Section 5 we offer some conclusions and speculations for future work.

2. The Experiment: Design and Procedures

2.1. General Features

Given our discussion above, it should be clear that any experiment on inter-generational games would have to contain certain salient features. For example, subjects once recruited should be ordered into generations in which each generation will play a pre-specified game repeatedly with the same opponent for a pre-specified length of time, T . After their participation in the game, subjects in any generation t should be replaced by a next generation, $t+1$, who will be able to view some or all of the history of what has transpired before them. Subjects in generation t will be able to give advice to their successors either in the form of suggesting a strategy, if the strategy space is small enough, or writing down a suggestion as to what to do and explaining why such advice is being given. This feature obviously permits socialization. The payoffs to any subject should be equal to the payoffs earned by that generation during their lifetime plus a discounted payoff which depends on the payoffs achieved by their successors (either immediate or more distant future). Finally, during their participation in the game, subjects should be 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 since the objects of societal evolution are both beliefs (social norms) and actions (social conventions based on norms).

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.

The experiment had three periods. In each period a subject would play one

of the three games with a different opponent. For example, consider the following table:

		Rotation Scheme For Subjects		
		Game		
		Battle of Sexes	Ultimatum	Trust
Period	Matches			
Period 1	Row	1	2	3
	Column	6	5	4
Period 2	Row	2	3	1
	Column	4	6	5
Period 3	Row	3	1	2
	Column	5	4	6

In this table we see six players performing our experiment in three periods. In period 1, Players 1 and 6 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 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. Each generation played the game once and only once and their payoff was equal to the payoff they received during their generation plus an amount equal to 1/2 of the payoff of their successor in the generation $t+1$ that followed them. (Payoffs were denominated in terms of experimental francs which were converted into U.S. dollars rates which varied according to the game played. The design was common knowledge among the subjects except for the fact that the subjects did not know the precise rotation formula used. They did know they would face a different opponent in each period, however.

As a result of this design, when we were finished running one group of six subjects through the lab we generated three generations of data on each of our three games since, through rotation, each player played each game once and was therefore a member of some generation in each game. Thus for the set-up cost of one experiment we generated three generations worth of data on three different inter-generational games at once. Still, our experimental design is extremely time

and labor intensive requiring 152 hours in the lab to generate the data we report on here.⁵

In this paper we will report the results of only the Battle of the Sexes Game played. This game had the following form:

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 1fr = \$.04.

The procedures used in playing all games were basically the same. When subjects started to play any of the three games, 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. 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.

⁵As far as we know, this is the record for economics experiments.

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 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 $r = (r_1; r_2)$ 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 r is the reported belief vector of subject i will be:

$$\frac{1}{2}_1 = 20,000 - ((100 - r_1)^2 + r_2^2) \quad (2.1)$$

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

$$\frac{1}{2}_2 = 20,000 - (100 - r_2)^2 + r_1^2 \quad (2.2)$$

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 $r = (r_1; r_2)$. 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.

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).

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 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 co-ordinate 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 a set of parameters $P = \{j; L_{h_t}; \pm; l, a\}$, where j is the stage game to be played over time, L_{h_t} is the length of the history h_t that the generation t player is allowed to see, with $L_{h_t} = t-1$ being the full history up until generation t , and $L_{h_t} = 1$, being only the last generation's history, \pm is the degree of inter-generational caring or the discount rate, l is the number of periods generation t lives before retiring, i.e., how many times they repeat the stage game with each other, and finally a is a 0-1 variable which takes a value of 1 when advice is allowed to be offered between any generation t and $t+1$ and 0 when it is not. In our Baseline BOSG experiment we set $L_{h_t} = t-1$, $\pm = 1/2$, $l = 1$, and $a = 1$ so subjects could pass advice to their successor, see the full history of all generations before them, live for only one period before retiring. They received a payoff which was equal to what they received in their one play of the game plus $1/2$ of what their successors received. This Baseline experiment was run for 81 generations. However, at period 52 we took the history of play and started two separate new treatments at that point which generated a pair of new independent histories. In Treatment I we set $L_{h_t} = 1$ so that before any generation made its move it could see only the last generation's history and nothing else. (All other parameters we kept the same). This treatment isolated the effect of advice on the play of the inter-generational game. 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. Treatment I was run for an additional 80 generations while Treatment II was run for an additional 66 generations, each starting after generation 52 was completed in the Baseline. Hence, our Baseline

was of length 81, our Treatment I of length 81⁹ and our Treatment II 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. Here, we build, estimate and test an extremely simple Markovian social learning model called the Stochastic Advice Model that does a remarkably good job of organizing our data.

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. There 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.

In our experiments we have another source of noise and that is the advice offered by one generation to the next. As we will see, there are times during

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

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

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 or norms 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

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 four parts or Regimes.

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: 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

To verify that these are in fact different regimes, we performed a two types of statistical tests. First we tested the null hypothesis that there was no difference in the distribution of row (or column) choices across our four regimes. Using a Kruskal-Wallis test we found we can reject this hypothesis for both the row and column players. In both cases the test statistic, K-W, is distributed as a $\hat{A}^2(3)$: For the row player K-W = 14.1, p-value = .00 while for the column player K-W = 13.1, p-value = .00. In order to investigate the source of these differences, we then performed pair-wise Wilcoxon Rank-Sum tests between our four regimes. For the row player, these tests detect no differences between Regime I and IV or between Regimes II and III, but there are significant differences between Regimes I and II,

Regimes I and III, Regimes II and IV, and Regimes III and IV at at least the 5% level.¹⁰ For the column player there are significant differences between Regimes I and II, II and III, and II and IV only, again at least at the 5% level.

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.

3.2.2. Inertia

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 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

	Row	Column
I v II	2.97 (.00)	3.34 (.00)
I v III	2.42 (.02)	0.89 (.37)
¹⁰ I v IV	-.13 (.90)	.26 (.80)
II v III	.53 (.60)	2.38 (.02)
II v IV	-2.74 (.01)	-2.70 (.01)
III v IV	-2.26 (.02)	-.53 (.60)

Each cell contains a z-statistic and the associated p value for the test between the two regimes listed in the first column of each row.

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).

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.

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 . Later we will see that these continuation probabilities are the diagonal elements of the Markov transition matrix we will estimate from our data.

Table 2 presents the probabilities:

Table 2: Continuation Probabilities by Regime

	1,1	1,2	2,1	2,2
Regime I:	0	0:166	NA ^a	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

^aNo (2,1) state occurred in Regime I

Since conventions are persistent states our intention in presenting Table 2 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.

To discuss advice we will present a summary of how advice was given in Table 3, and under what circumstances it was followed Table 4.

Table 3: 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 4: Advice Adherence Conditional on Last Period's State

State Last Period: (1,1)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
Advice 1	11	5	5	9
Advice 2	0	0	1	1
Total	11	5	6	10

State Last Period: (1,2)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
Advice 1	7	2	10	5
Advice 2	10	8	10	2
Total	17	10	20	7

State Last Period: (2,1)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
Advice 1	5	2	0	2
Advice 2	0	0	4	1
Total	5	2	4	3

State Last Period: (2,2)

	Row Player		Column Player	
	Followed	Rejected	Followed	Rejected
Advice 1	3	0	0	0
Advice 2	19	8	26	4
Total	22	8	26	4

What Advice Was Given Table 3 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 greedy disequilibrium state occurred, advice was more diverse. In 18 of the 27 occurrences of this disequilibrium state, the row player suggested switching to the submissive strategy of choosing two 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 strategy 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 4 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 5a and 5b:

Table 5a: 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 5b: 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

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.

3.2.4. The Bounded-Memory Advice Giving and Following Model

A full model of our subjects' advice following and giving behavior in the inter-generational Battle of the Sexes game would have to contain three elements. First

it would have to describe how beliefs are formed given the history of the game. Second, it would have to explain how subjects choose to give advice. Finally, it would have to describe how subject's choose to follow or disobey their advice given their beliefs. While each of these relationships can be complex, we chose here to present as simple a model as we can which will capture what we feel are the salient features of the phenomenon under consideration.

To do this we start out with the assumption that subjects are bounded in either the memory or their ability to process long histories of actions. Hence, we assume that our subjects are capable of only using data which is at most four generations old. We define the belief held by agent i , $i=\{\text{row } (r), \text{ Column } (c)\}$ at time t that his or her opponent, j , will choose strategy k , $k=\{1,2\}$ as

$$b_{it(s_j(k))} = \frac{n_k^4}{4}, \quad (3.1)$$

where n_k^4 is the number of k 's that subject j has chosen in the last four rounds. $b_{it(s_j(k))}$ is therefore simply the fraction of times in the last four generations that one's opponent has chosen the strategy k .

When a subject goes to choose, however, he or she has both beliefs and advice at his or her disposal and, in the process of decision making, these must be traded off. To do this we make the following assumption. When a subject holds beliefs that make him or her indifferent between his two different strategies (i.e., when he holds "equilibrium beliefs"), then he or she follows advice with some base propensity which depends on the subject's predisposition to follow advice in general. To the extent, however that history indicates that his or her received advice is not consistent with the subject's best response to recent history, he or she will discount such advice and visa versa. Still, given our data we want to retain the fact the no matter what a subject's beliefs and no matter what advice they are given, choice is stochastic.

To capture this feature we posit a choice function of the logistic form. Before we describe this choice function, however, remember that in the Battle of the Sexes Game there are two types of actions; those that are consistent with the equilibrium that is best for the row chooser (Row 1, Column 1) and those that are consistent with the equilibrium that is best for the Column chooser (Row 2, Column 2). Calling these strategies the subject's "better" (b) and "worse" (w) strategies, respectively, we formulate our logistic equation as defining the probability that subject i at time t will choose his better strategy as:

$$\text{Pr}_{it}(s(b)) = \frac{e^{(\beta + \gamma)(b_{it}(s_j(b)) - b^*) + \gamma D}}{1 + e^{(\beta + \gamma)(b_{it}(s_j(b)) - b^*) + \gamma D}} \quad (3.2)$$

This function describes the probability that a subject will choose his better strategy as a function of his belief, $b_{it}(s_j(b))$, and a dummy variable which takes a value of 1 if the subject was told to choose his better strategy and 0 if he or she was not told to do so. (The dependent variable takes on a value of 1 if the better strategy was chosen and zero otherwise). Here the b^* term, which is subtracted from the belief value in the exponent, is the "indifference belief" or that belief that if held by a subject would make him or her indifferent between either strategy 1 or strategy 2. In our experiment $b^* = .25$, so that if either the row or column player believes that his or her opponent will choose his or her favored strategy with probability of .25, he or she will be indifferent between choosing either row (or column) 1 or 2. With the dummy, this one equation specifies the probability of making any choice given any piece of advice, i.e., the probability of choosing the better strategy if you are told to do so as well as the probability of choosing the worse strategy if you are told to do the opposite etc..

Finally, we must specify an advice function. Here again we make as simple an assumption as we can. What we want is an advice function that suggests choosing a strategy in a manner that is monotonically increasing in that strategy's expected payoff given the recent history of choices by one's opponent. What we mean here is that say you are a row player who has just finished playing the Battle of the Sexes Game and assume that the column choosers have chosen strategy 1 for three out of the last four generations including your own (which you will observe before you give advice). Since the expected payoff from choosing strategy 1 is monotonically increasing in this fraction, we would expect that you would give advice to choose strategy 1 more often than if the fraction of 1's chosen recently were smaller. This is captured in the following simple logistic advice function which specifies the probability of offering advice to choose the better strategy b as a function of the belief that the next generation's opponent will choose it as well.

$$\text{Adv}_{it}(s(b)) = \frac{e^{\beta + \gamma(b_{it+1}(s_j(b)))}}{1 + e^{\beta + \gamma(b_{it+1}(s_j(b)))}} \quad (3.3)$$

This is a single belief function for all subjects, whether they are row or column. Disaggregating on this basis had no effect. Finally, note that $b_{it+1}(s_j(b))$ is the belief held by subject t at the end of his lifetime and including the actions of his

opponent during their inter-action. (It differs from $b_{it(s_j(b))}$ which was used in the choice function (3.2) above.)

We estimated this model using the data generated by our Baseline experiment. This is the only experiment for which we could make such an estimate since it is the only one for which we have both advice and four-period histories. Further, in previous formulations of the model we allowed coded advice as an explanatory variable (see below) in the choice equation 3.2 as well as including a dummy variable to capture the potential difference between row and column behavior. Neither of these treatments affected the results in statistically significant ways, and, by Occum's Razor, we present only our simplest model here.

The results of this estimation are presented in the tables 6-7 below:

Table 6: Logit Choice Model for Good Equilibrium Choices

	Coef	Std. Err	t	P> t	[95% Conf. Interval]
belief	2.473	.8031	3.080	.0002	.8995 - 4.047
advice	.9495	.4083	2.325	.0200	.1491 - 1.749
cons	-.1992	.2433	-.0819	0.413	-.6715 - -.2771
obs =	154				
LL =	-86.81				

Table 7: Advice Behavior - - Probability of Giving Advice To Choose Better Strategy

	Coef	Std. Err.	t	P> t	[95% Conf. Interval]
belief (t + 1)	4.8457	.8354	5.800	.000	3.208-6.483
cons	-1.668	.3439	-4.850	.000	-2.342 -.9939
obs =	154				
LL =	-81.58				

In order to discuss the results of these regressions in a more meaningful manner we transform them into probabilities. In Table 8 we present the probability that any particular piece of advice will be given. In Table 9 we present the advice following and disobeying relationships. Here we calculate the implied probabilities of following (disobeying) the advice to choose the better (or worse) strategy given that a subject was told to do so conditional on the subject's beliefs.

Table 8: Probability of Offering Advice to Choose Better Strategy

Belief*	Probability of Offering Advice to Choose The Better Strategy
1	.96
.75	.88
.50	.68
.25	.39
0	.16

* This is the updated belief that your opponent will choose the strategy which is better for you.

Table 9a: Advice Following

Belief*	Probability of Choosing Better Strategy if told to	Probability of Choosing Worse Strategy if told to
1	.93	.16
.75	.88	.26
.50	.80	.40
.25	.68	.55
0	.53	.69

* This is the belief that your opponent will choose the strategy which is better for you.

Table 9b: Advice Disobeying

Belief*	Probability of Choosing Better if told Worse	Probability of Choosing Worse if told Better
1	.84	.07
.75	.74	.12
.50	.60	.20
.25	.45	.32
0	.37	.47

* This is the belief that your opponent will choose the strategy which is better for you.

To discuss these results let us first look at the advice probabilities in Table 8. Note that these probabilities are similar to the ones we observed before in

our descriptive tables. The probability of offering advice to choose the better strategy given that you have seen your opponent over the last four periods always choosing in a manner complementary to it is .96. Likewise, the probability of offering advice to your successor to choose the better strategy given that you have never seen the complementary strategy chosen by your opponent is only .16. Note finally, that economic theory is violated in some sense by these results since, in looking at the probability of offering advice to choose the better strategy when holding indifference beliefs (.25), subjects only offer that advice with probability .39. If subjects were truly indifferent we might expect them to choose it with probability .50. (This may be a result of the fact that these four period historical beliefs are not a good proxy for the subjects' real, or elicited, beliefs or of the fact that subjects do not realize the .25 is their indifference point.)

The most interesting results are found in the advice following behavior of subjects. Here we observe a true bias which is consistent with the over-confidence bias we observed when we looked at the beliefs of subjects. Here the bias is a "better-strategy bias" or one in which subjects are more eager to choose their better strategy conditional on equivalent historical evidence than they are to choose their worse strategy. For example, looking at Table 9a we notice that if a subject is told to choose his or her better strategy and has only seen his opponent choose the complementary strategy to it for the last four periods, the probability that he or she will follow that advice is .93. In a symmetric situation for the worse strategy, i.e. being told to choose the worse strategy in a situation where the subject has always seen the complementary worse strategy chosen by his or her opponent, the probability of following that advice is only .69. Note also that the probability of choosing the better strategy having been told to do so and given a belief of 0.0 is .53, while a subject would only choose his worse strategy under similar circumstances with probability .16. Similar asymmetries exist for advice disobeying. Here when you have a zero belief that your opponent will choose the better strategy and you are told to choose the worse action, you disobey with probability .37 while you would choose the worse action with only probability .07 if you were told to do the better thing, given a zero belief that your opponent chose the worst thing. .

These results, combined with the overconfidence that subjects have in their elicited beliefs that their opponent will choose an action that is best for them, which we will comment on later in the paper (see Section 3.3 below) explain why coordination in these games, although frequent, is not stable or persistent.

In principal, these calculations can be used to construct a Markov chain de-

scribing the transitions between pairs of four-period row-column histories into themselves. We could then use this matrix to estimate the long-run probabilities associated with each of these states. ¹¹ However, there are 25 such four-period history-pair states ¹² and 625 transitions amongst them. Given our 81 generations of data, we can not hope to estimate these transitions as we have above and so we will leave this exercise for future experiments and concentrate in Section 5 on presenting set of simple one-state transition matrices which capture some of the behavior observed in our experiments. Note, however, that if such a matrix were constructed, our action probabilities imply that the stochastic process governing choice in this model would have no absorbing states since no matter what the history all strategies have a strictly positive chance of being chosen. This would imply that in such a model we are likely to get the punctuated equilibria that we observed in our data.

3.2.5. The Content of Advice

Recall that in addition to allowing our subjects to suggest strategies for their successors, we also allowed them to write free-form messages to them explaining why they are suggesting the strategies they are. In addition, we know that advice was taken quite seriously and followed many times to the exclusion of beliefs. What is not known, however, is how important these messages were in determining behavior. To help us understand the process of advice taking better we decided to code the advice data in the following manner.

First three people each independently read all of the advice messages and placed them in one of five different categories depending on the degree of higher order rationality they included. More precisely, advice which was null advice, i.e., a blank page, or mere gibberish, i.e. advice which had no content relevant to the game being played, was placed in Category 0. For example, the row player in generation 61 told his successor in generation 62 "May the force be with you". We placed this piece of advice in category 0.

Category 1 was reserved for all those pieces of advice which urged the subject to look at history but did not clearly specify how one should learn a lesson from

¹¹Such an exercise is very close to what H.P. Young (1998) does but his histories are generated from taking a k period random sample from a larger set of m periods. He then looks for the transitions between these k-period pairs of histories into themselves.

¹²This is true if we make the simplifying assumption that two four-period histories are equivalent as long as the fraction of 1 and 2 choices in them are identical even if their placement in the sequence differs. If this assumption fails, the number of states is much larger.

such history, For example Column subject 16 told his successor "Read the history before you play." while his successor in generation 17 told his successor in 18 "Read the history and do accordingly". Both of these messages are really devoid of normative content since they don't tell the subject how to draw a lesson from what he or she observes. If a message suggested looking at the history for a purpose like "You should choose column #1 since the history suggests a trend of cooperation" (Column subject 34) we did not place it in Category 1 but one of the remaining higher categories.

Category 2 was reserved for that advice which simply told the subject what to choose without an explanation. It was pure prescriptive advice. For example the Column subject in generation 35 simply said "Choose Column 1". Here a clear prescription is made without any justification. Other statements placed in this category may have been more elaborate. However, they all shared the property that they ultimately made pure suggestions without justification. Hence we placed it in Category 2.

Category 3 was reserved for all those comments in which a subject was told to choose a particular strategy because it was, in some sense, a best response to what they thought the other player would do. In this category we placed all advice which had a first order logic in which a subject was urged to do something because of an anticipated move by his opponent. For example, the row subject in generation 4 urged his successor to choose row 2 since "Your opponent is gong to choose the column with the highest payoff for themselves" . Clearly the action suggested is a best response to this belief.

Finally in Category 4 we place all of those pieces of advice which ask the next generation to think about what the other player is thinking about you, and then choose an action which is a best response to ones opponent's best response to you. For example, the column subject in generation 44 said the following: "Based on the history, I calculated the likelihood of choosing column 2 was 60%. If your opponent does the same, i.e., calculate from the statistics given, he'll think you will choose column 2 and the payoff is certainly the one you want. Well it worked for me". This quote illustrates a final principle we used which was that if a statement contained material from two different categories, we always chose the highest category to place it in. In this case, while the quote referred to history, since it gave a reason for suggesting what it did and since this reason involved a second order logic, we placed it in Category 4.

Following standard procedures for coding data each of the three people who coded the data did it independently and then the coding was compared. All

entries which had unanimous agreement we left alone. All others were reviewed and in almost all cases the entry was coded according to how the majority coded it.¹³

Tables 10a -10c present the results of our coding. In Table 10a we present the distribution of Row and Column advice across all five categories aggregated over time. Table 10b disaggregates this data by regime while Table 10c disaggregates the data by the state in which it was given.

Table 10a:Advice Data Coding:All Rounds

	Row	Column
Category 0	30	31
Category 1	2	7
Category 2	16	16
Category 3	27	22
Category 4	6	5

¹³The actual distribution of coding agreements was as follows: For the row players there were 42 quotes receiving a unanimous agreement on category, while for the column there were 58 quotes receiving unanimous agreement.

When disagreement occurred, the final states chosen for the Row and Column players were:

Row		Column	
Advice	#	Advice	#
0	5	0	4
1	2	1	2
2	15	2	10
3	12	3	5
4	4	4	1

Table 10b: Advice Data Coding by Regime

	Regime I			Regime II	
	Row	Column		Row	Column
Category 0	8	10	Category 0	9	8
Category 1	0	3	Category 1	1	0
Category 2	2	7	Category 2	4	2
Category 3	13	3	Category 3	4	9
Category 4	1	1	Category 4	3	2
	Regime III			Regime IV	
	Row	Column		Row	Column
Category 0	11	8	Category 0	2	5
Category 1	0	0	Category 1	1	4
Category 2	3	5	Category 2	7	2
Category 3	5	7	Category 3	5	3
Category 4	2	1	Category 4	0	1

Table 10c: Advice by State

State	Row Category				Column Category			2	3	4
	0	1	2	3	4	0	1			
State (1,1)	5	0	3	6	2	3	0	3	10	0
State (1,2)	14	2	5	5	1	9	1	4	10	3
State (2,1)	4	0	1	1	1	6	0	0	1	0
State (2,2)	7	0	7	15	2	13	6	9	1	2

Tables 10a-10c tell an interesting story. As we see from our simple Battle of the Sexes stage game, the two pure strategy equilibria are asymmetric in the sense that at the (1,1) equilibrium the row player does significantly better than the column player while in the (2,2) equilibrium just the opposite is true. Calling the players doing well in an equilibrium the "haves" and the others the "have-nots" we see that when a convention is in place, the haves offer lower category advice (advice 0, 1, or 2) than do the have-nots who offer higher level advice (category 3 or 4). For example, as we see in Table 6b, in Regime I the column player offers advice which is category 2 or less in 20 of the 24 generations of that regime, while the Row players offer such low level advice only 10 times over those same generations. Furthermore, while the Row subjects offer category 3 advice 13 times, the column player only do so 3 times. Similarly, in Regime IV where we claim the (2,2) equilibrium re-emerges, the column subjects offers category 0, 1, advice 9 out of 15 generations while the row subject did so only 3 times. (The row subject did use category 2 advice here quite often, however).

The above observations are supported by a series of Wilcoxon Signed-Ranks tests. These tests reject for Regimes I and IV the hypotheses of equal distributions of advice in favor of the alternative that row players offer higher category advice. This is true at the 4% level ($z = 2.03$) in Regime I and at the 7% level in Regime IV ($z = 1.79$). The null hypothesis is not rejected for Regimes II and III ($z = -.56$ for Regime II and $z = -.52$ for Regime III).

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 3a and 3b where we present the probability that each generational subject felt his opponent would choose strategy 1.

[Figure 3a and 3b here]

Note that in Figures 3a and 3b 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 that 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 of 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: $z = -0.34$, 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.¹⁴

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

¹⁴This was done by forbidding them to write any instructions on the screen despite the fact that they were prompted to.

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 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 4a-4c plots the time series generated by these two treatments along with our original Baseline treatment.

[Figures 4a-4c 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.

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 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) 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 in the Appendix to this paper.

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 \hat{A}^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 calculate the test statistic,

$$\chi^2 = \sum_{i=1}^4 \sum_{j=1}^4 n_i \frac{(p_{ij}^T - p_{ij}^{P^k})^2}{p_{ij}^{P^k}}; \quad (4.1)$$

where n_i is the number of instances of state i in the data. This statistic is distributed as \hat{A}^2 with $4(3-d)$ degrees of freedom where d is the number of zeros in the P matrix and the summation in the above formula is taken over only those ij states for which $p_{ij}^{P^k} > 0$. Asymptotically this is equivalent to a likelihood ratio test based on the Neyman-Pearson lemma.

When these calculations are made 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 ($\hat{A}^2(12df) = 27.6521$, ($p = 0.000$)) or Treatment II ($\hat{A}^2(9df) = 59.4262$, ($p = 0.000$)). 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.

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.¹⁵

¹⁵There is no sense, then, that parents in our experiment are in any way "experts" as in the

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.

5. Some Simple Models

One might be tempted, in constructing a model to explain our data, to look for equilibria in inter-generational supergame strategies. Such strategies would specify a convention and a set of inter-generational punishments that go along with deviations from it.

We reject such an approach because we see no evidence that such strategies were used. More precisely, if such strategies were used we would expect to see evidence of them in the advice passed on from generation to generation since the strategy would have to be explained to successive generations. No such advice was ever observed. In fact, advice was overwhelmingly myopic and never forward looking.

The type of model we would ultimately want to construct to explain our data would have more in common with the type of model suggested by H.P. Young (1998). In his book, Young treats the evolution of conventions of behavior as a Markov chain defined on the state space of strategy pairs. (Actually his state space is the space of k -period histories where k is ...nite). The important features of the model is that behavior is myopic, in fact Markovian, and the process transits from state to state according to some stationary stochastic process. What is looked for

model of Ottaviani and Sorensen (1999).

in such models is not an equilibrium set of strategies but rather a long run steady state probability distribution defining what the likelihood is that the game will be in any state in the long run.

Because of the limitation on the size of our data set, we are limited in our ability to estimate such a model. However, we do investigate three extremely simple stochastic models which serve as a primitive first attempt at explaining our observed behavior. Despite their simplicity, two of them, The Pure Advice Model and the Pure Best-Response Model do appear to capture at least the qualitative nature of some aspects of the data.

5.1. The Pure Advice, Pure Best-Response and Mixed Strategy Models

In the Pure Advice Model, subjects simply follow the advice given to them by their parents. They completely ignore their beliefs and the history they have observed. In the Pure Best Response Model, subjects do the opposite. They completely ignore the advice they are given and best respond to their stated beliefs. In other words, in the Pure Advice Model $c_t^i = a_{t,i-1}^i$ where c_t^i is the strategy choice of generation subject of type i in generation t and $a_{t,i-1}^i$ is the advice that agent was sent. In the Pure Best Response Model, $c_t^i = \underset{c_i}{\operatorname{argmax}} E[\%_i(c_t)]$ where the expectation operator E is defined using the stated beliefs of the subject in generation t , and $\%_i(c_t)$ is the payoff function for agent i which depends on the vector of actions taken by both agents in generation t . These models are at the extreme ends of the spectrum of models that one could build using our data.

The Mixed-Strategy Model posits that generational subjects independently play the mixed strategy equilibrium to the stage game which involves the row player choosing a vector $p = (.25, .75)$ and the column player choosing $q = (.75, .25)$. We include this model in order to check that the stochastic movements we observe are not simply the result of independent stage-game play.

Note that, despite their deterministic behavior rules, both the Pure Advice Model and the Pure Best Response Model are stochastic models. In the Pure advice model, behavior is stochastic because, as we have seen, there is an underlying randomness in the advice that people are given conditional on any state reached in the game. In the Best Response Model, behavior is stochastic because new generations come to the experiment with new priors drawn from some underlying distribution which is unobserved. Hence, despite the fact that they best respond in a deterministic fashion, what each new generation is best responding to is stochastic. The Mixed-Strategy Model is stochastic for obvious reasons and all of

these models define a Markov chain in the state space of one-period histories.

5.2. Model Selection

We will select across these models in a very straight forward manner. First we will simply look at the time paths of actions that would be chosen by subjects had they followed either the Pure-Advice or Pure-Best-Response Models and compare the resulting time series to that which we observed in the Baseline and Treatment I experiments. (In Treatment II the Pure Advice Model can not be defined). Next we will look at the one-period transition matrices that are defined by these models and compare them to the one-period transition matrix defined by our experimental data. Using a chi-square test for goodness-of-fit, we will then choose among the models.

5.2.1. Time Series

To examine the Pure Advice and Pure Best-Response Models we present Figures 5a-5f which present the actual time series of outcomes in the Baseline and Treatment I experiments (Figures 5a and 5d) along with the hypothetical time series of outcomes that would have resulted in these two experiments if generational subjects had either simply best responded to their beliefs, (Figures 5c and 5f) or simply followed the advice of their parents (Figures 5b and 5e). (The Pure Advice Model can not be defined for Treatment II).

[Figure 5a-5f here]

A look at Figures 5a-5c (the Baseline Experiment) should be informative. First, it should be obvious that qualitatively the Pure Advice model captures the emergence of conventions of behavior more accurately than the Pure Best Response Model in the sense that it predicts the successive runs of repeated play of each stage-game equilibrium more accurately. For example, only the Pure Advice model predicts the nine successive plays of the (2,2) equilibrium during the first 13 generations. It is also successful in predicting the final five plays of the (2,2) equilibrium during generations 77-81. In addition, it was successful in predicting five out of ten occurrences of the (1,1) equilibrium between periods 25 and 45 while the Best Response Model predicted only three. Over all, there were 47 generations which chose equilibrium stage game actions over the 81 period history of the experiment. While the Advice model predicted 30 of them, and the Best Response Model predicted only 16 .

Note however, that the Best response model does a far better job at predicting the disequilibrium states than does the Pure Advice Model. For example, the Pure Best Response Model predicts the disequilibrium (1,2) state in 20 out of 27 instances while it predicts the (2,1) state in 2 out of 7 instances. The Pure Advice Model predicts these states successfully only 7 out of 26 and 3 out of 7 times respectively.

A similar pattern emerges from our Treatment I experiment where we have advice but no history. Here once more the Pure Advice Model does a better job of predicting the equilibrium states in the actual data (11 out of 23 times for state 1 and 5 out of 15 times for state 4) while the Pure Best Response Model does a better job of predicting the dis-equilibrium states. This may be a result of our previously noted over-confidence bias in the belief data which leads each type of subject to best respond with that action which is associated with the equilibrium which is best for him or her and thereby leading to an abundance of (1,2) states (41) while there were only 24 such (1,2) states in the Pure Advice Model. Because subjects happen to use these strategies most often, we see once again that the Pure Best Response Model does a better job of explaining the disequilibrium states than does the Pure Advice Model.

5.2.2. Transition Matrices

To further compare the performance of our three models we look at the state-to-state transition matrices they define and test whether such transitions could have been generated by the same process that generated the transitions found in our data. We do this for all treatments and all models. (The estimated transition matrices can be found in the Appendix to this paper).

To test if the actual transition probabilities are generated by a process equivalent to the one described by our three models we use a \hat{A}^2 goodness of fit test described above.

The results of these chi-square tests are presented in Table 11:

Table 11:
Goodness-of- Fit Tests

Experiment	Model		
	PBR	PAM	MSM
Baseline	42.95 (.00) ⁹	37.37(.00) ¹²	57.56 (.014) ¹²
Treatment I	12.34 (.19) ⁹	6.67 (.67) ⁹	16.90 (.15) ¹²
Treatment II	8.17 (.42) ⁸	NA	4.87 (.96) ¹²

Format: \hat{A}^2 (p value)^{degrees of freedom}

These results point up some interesting features of our data. First note that the Mixed Strategy Model is rejected in the Baseline (at the 1.4% level). In Treatment I, while we could not reject the hypothesis that the data observed there was consistent with iid mixed strategy play, the p-value was .15 as opposed the Treatment II where we also could not reject the mixed strategy model but with greater authority ($p = .96$). Still, in the Baseline none of our three simple models fit the transition data well. This is true because while the Pure Advice Model does a good job of explaining equilibrium behavior in the game and the Pure Best Response Model does a good job of explaining disequilibrium behavior, a model capable of explaining the observed data might need to be a hybrid of these two extreme models if it is going to provide a sufficiently good fit.

5.2.3. Steady State Probabilities

The estimated transition matrices defining the evolution of our generational states for each of our treatments form an irreducible, a-periodic Markov chain. For such chains the stationary distribution of the transition process can be found by solving the system of equations $xT = x$ for the vector x , where T is the appropriate transition probability matrix. (See Appendix B for the actual matrices).

When we solve this system for our three experiments using the actual transition matrices estimated using the data generated, we find that the long-run probabilities of visiting any of our four states are:

Table 12:
Long -Run Transition Probabilities

Experiment	State			
	1	2	3	4
Baseline	.1975	.3230	.0848	.3945
Treatment I	.2837	.35424	.1670	.1984
Treatment II	.2002	.5652	.1452	.0891
Mixed Strategies	.1875	.5625	.0625	.1875

There are some interesting characteristics of these three vectors. For instance, although the Baseline long run probabilities do not appear to be like those in our No History experiment, they both share the property that in each experiment, in the long run, the fraction of time spent in equilibrium (i.e. State 1 or 4) are high relative to the time to be spent in the dis-equilibrium states (States 2 and 3 - -.594 for the Baseline and .4821 for the No History (but advice) experiment). For the No Advice experiment, on the other hand, the fraction of time expected to be spent in an equilibrium state is only .28939 substantiating our conjectures above that advice is a needed component for successful coordination but that advice plus access to history is best. Further, note that the long run behavior of Treatment II is similar to that which would evolve if subjects used iid mixed strategies at each point in time (the last row of the table). The similarity is especially strong in the probability both vectors place on State 2. Finally, note that the mixed strategy model severely under-predicts the use of equilibrium strategies in both the Baseline and Treatment I Experiments but over-predicts it in the case of Treatment II.

6. 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 stochastic one (a Markov chain) in which non-overlapping generations of players create and pass on conventions of behavior in a Lamarckian fashion from generation to generation. 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. In addition, this paper helps make a case for the use of Lamarckian, as opposed to Darwinian models, in analyzing social evolution. These models, we feel, give greater scope to the abilities of humans to think creatively and socialize their offspring, thereby avoiding being trapped in an unsatisfactory, but perhaps evolutionarily stable, equilibria.

Appendix:

Transition Matrices

Transition matrix Baseline

	State			
State	1	2	3	4
1	.187	.5	.187	.125
2	.296	.259	.148	.296
3	.142	.574	0	.285
4	.133	.233	0	.633

Transition matrix Treatment I: No History

	State			
State	1	2	3	4
1	.541	.333	.083	.041
2	.148	.333	.296	.222
3	.230	.538	.076	.158
4	.200	.266	.133	.400

Transition matrix Treatment II: No Advice

	State			
State	1	2	3	4
1	.384	.538	.0	.076
2	.135	.540	.216	.108
3	.200	.800	.000	.000
4	.200	.400	.200	.200

Transition Matrices: Pure Advice Model

Transition Matrix

Pure Advice Model: Baseline

	State			
State	1	2	3	4
1	.454	.136	.136	.272
2	.384	.153	.230	.230
3	.333	.444	.111	.111
4	.114	.085	.057	.742

Transition Matrix

Pure Advice Model: Treatment I

	State			
State	1	2	3	4
1	.384	.230	.153	.230
2	.260	.304	.000	.434
3	.857	.000	.000	.142
4	.136	.500	.136	.227

Transition Matrices: Pure Best-Response Model

Transition Matrix

Pure Best Response Model:
Baseline

	State			
State	1	2	3	4
1	.000	.714	.000	.285
2	.113	.545	.068	.272
3	.166	.500	.000	.333
4	.043	.521	.130	.304

Transition Matrix

Pure Best-Response Model:
Treatment I

	State			
State	1	2	3	4
1	.277	.500	.055	.166
2	.205	.435	.051	.307
3	.000	.666	.000	.333
4	.262	.578	.000	.157

Transition Matrix

Pure Best Response Model:
Treatment I I

	State			
State	1	2	3	4
1	.250	.500	.083	.166
2	.142	.485	.028	.342
3	.000	1.00	.000	.000
4	.184	.523	.030	.261

7. 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] Ullman-Margalit, E.,(1977), *The Emergence of Norms*, Oxford University Press, Oxford, England.
- [23] VanHuyck, J., Battalio, R., and Beil, R., (1990), "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure", *American Economic Review*, Vol. 80, pp. 234-248.
- [24] Vega-Redondo, F. (1996), *Evolution, Games and Economic Behavior*, Oxford University Press, Oxford, England.
- [25] Weibull, J., (1995), *Evolutionary Game Theory*, MIT Press, Cambridge, Massachusetts..
- [26] Young, H.P., (1993), "The Evolution of Conventions", *Econometrica*, vol. 61, No. 1, pp. 57-84.

- [27] Young, H.P., (1996) , "The Economics of Conventions" , Journal of Economic Perspectives, vol. 10, No. 2, pp. 105-122
- [28] Young, H.P., (1998), Individual Strategy and Social Structure, Princeton University Press, Princeton New Jersey.

Figure 1: Experimental Design




Figure 2: Baseline Outcomes

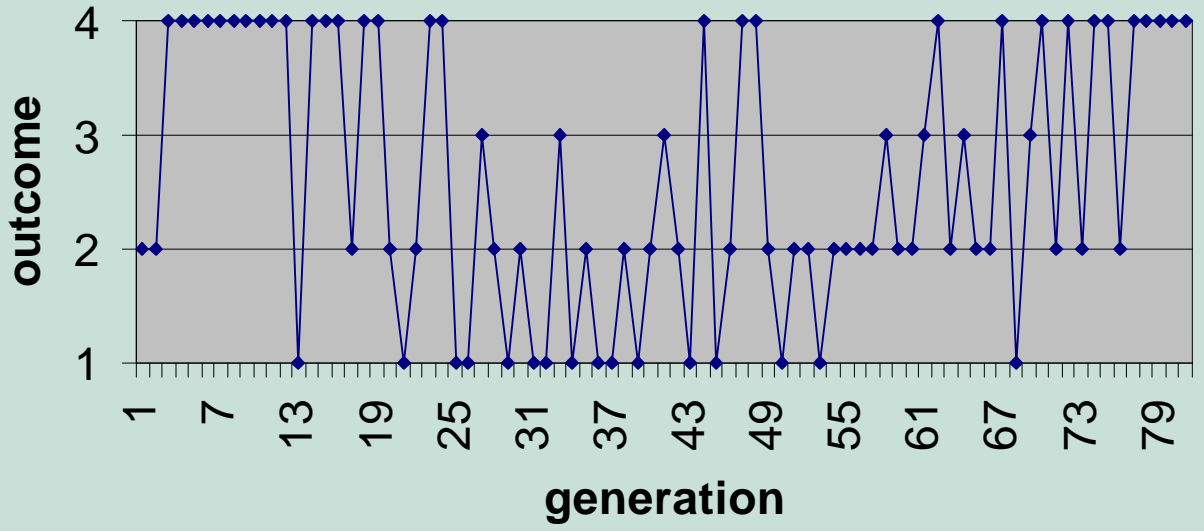


Figure 3a: Row's Beliefs About Column - Probability Column Chooses 1




Figure 3b: Column's Beliefs About Row, Baseline




Figure 4a: Baseline Outcomes




Figure 4b: Treatment I Outcomes





Figure 4c: Treatment II Outcomes




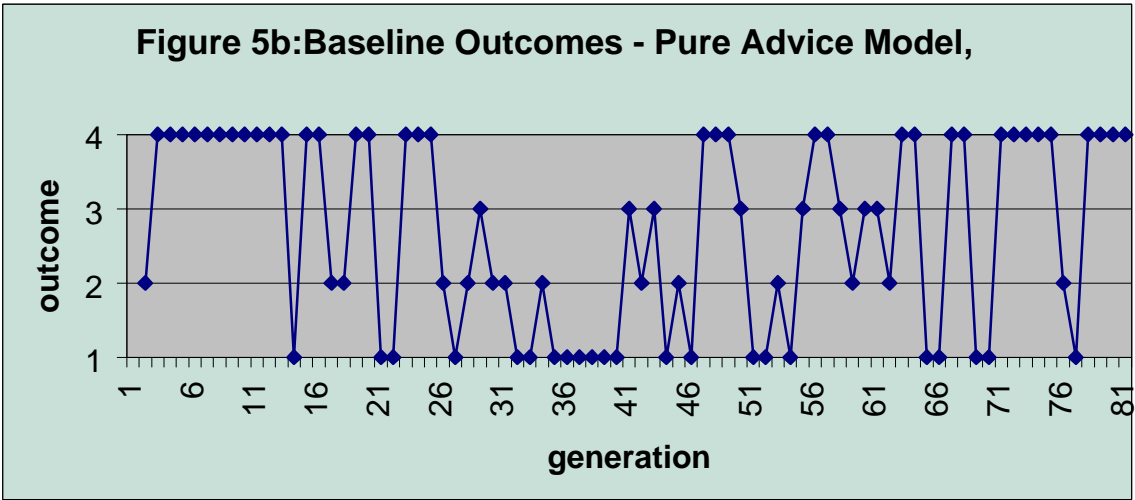



Figure 5d: Treatment I Outcomes




Figure 5e: Treatment I Outcomes - Pure Advice Model





Figure 5f: Treatment I Outcomes - Pure Best-Response Model



APPENDIX : 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 like conversion rate of experimental currency to dollars have been left blank.

General Instructions

Introduction

You are about to participate in an experiment in the economics of decision making. Various research foundations have provided the money to conduct this research. If you follow the instructions and make careful decisions, you might earn a considerable amount of money.

Currency

The currency used in this experiment is francs. All monetary amounts will be denominated in this currency. Your earnings in francs will be converted into U.S. Dollars at an exchange rate to be described later. Details of how you will make decisions and earn money, and how you will be paid, will be provided below.

The Decision Problem

In this experiment, you will participate in three distinct decision problems. In each problem, you will be paired with another person and you will each make decisions. The monetary payoff that you receive depends upon the decisions that you make and upon the decisions that the person you are paired with makes.

After you have played the first decision problem, you will then be paired with another person, different from the one you were first paired with, to play a second game. Again, your payoff in this second decision problem will depend upon the decisions that you make and upon the decisions that the person you are paired with makes.

After you have participated in the second decision problem, you will once more be paired with another person, different from either of the people you were paired with in the first two decision problems. Your payoff in this third decision problem will, again, depend upon the decisions that you make and upon the decisions that the person you are paired with makes.

You will never be informed of the identity of any of the people you are paired with, nor will any of them be informed of your identity.

The details of the three different decision problems that you will participate in will be briefly described to you just prior to each decision problem. What follows here is a general description of the structure of the decision problems and of the procedures that will be followed for each decision problem.

General Structure

In general, you and the person you are paired with will not be the first pair who has participated in a particular decision problem. That is, in general, other pairs will have participated before you, either earlier today, or on previous days. Further, you and the person you are paired with will not be the last pair to participate in the decision problem. That is, other pairs will participate in the decision problem after you, either later today or on later days.

Roles

In each decision problem, you will be replacing a person who has participated before you. In each decision problem there are two decision makers, A and B, and you will be assigned the role of either A or B.

Payoffs

In each decision problem, you will make a decision and the person you are paired with will make a decision, and these decisions will determine your payoff from playing the decision problem. In addition, you will also receive a payment equal to a fraction of the earnings made by your replacement when he/she takes your place. (Your predecessor will also be earning a payment equal to a fraction of what you earn). Thus, a player's total payoff from any particular decision problem is the sum of the earnings from the decision problem one plays with the person one is paired with, plus a payment equal to a fraction of the earnings from the decision problem one's successor plays with the successor of the person one is paired with in the decision problem.

Advice

Since, in general, your total payoff depends on your own decision and on the decision of the person who succeeds you in your role in a decision problem, you will be allowed to pass on advice on what action to take in the decision problem to your successor. The person you are paired with will also be allowed to pass on advice to his/ her successor. The person who was in your role when the last decision problem was played will be able to leave you advice on what action to take in the decision problem. Similarly, the person who was in the role of the person you are paired with when the decision problem was last played will have left him/ her advice on what action to take in the decision problem.

History

Since others have participated in a decision problem before you, you will be able to see some part of the history of the actions taken in the decision problem before you. Specifically, you and the person you are paired with will be able to see the decisions made by all previous pairs in this decision problem.

Predictions

At various points in the decision problem, prior to making a decision, you will be asked how likely you believe it is that your opponent is going to take any given action in the decision problem. To give you the incentive to state your beliefs as accurately as possible, you will be compensated according to how accurate your stated beliefs are, in light of what your opponent ends up doing. The details of how you will be compensated will depend on which decision problem you are participating in. Details of how you will be compensated will thus be deferred until the specific instructions for the different decision problems.

How you get paid

You will receive \$5 simply for showing up today and completing the experiment. You will receive, in addition, a payment today based on the outcome of the three decision problems you participate in. A second payment, based on the outcome of the three decision problems of your successors, will be available at a later time. You will be notified when your later payment is ready for you to pick up.

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 r round(s). In each round, every participant will engage in the following decision problem where you will either play the role of the Asender@ or Areceiver@. (Which type you are will be told to you before your participation in the decision problem begins):

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 $(1/2)$ 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 you 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 accept or reject), take the number you assigned to that choice, in this case 40 on reject, 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 accept, 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 = %f 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 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.