ECONOMIC RESEARCH REPORTS

Creating Culture in the Lab: Equilibrium Conventions in Inter – Generational Ultimatum Games

by

Andrew Schotter and **Barry Sopher**

RR#: 2000-07 **April 2000**

C.V. Starr Center for Applied Economics

Department of Economics

Faculty of Arts and Science

New York University

269 Mercer Street, 3rd Floor

New York, New York 10003-6687

Creating Culture in the Lab: Equilibrium Conventions in Inter-Generational Ultimatum Games

Andrew Schotter[*] and Barry Sopher[+]

April 10, 2000

Abstract

The Ultimatum Game and the experiments surrounding it, have presented economists with a puzzle that they have struggled to explain. But as Robert Aumann has pointed out, while there may be only one sub-game perfect equilibrium to the Ultimatum Game, there are an infinite number of Nash equilibria. All that is needed to maintain a non-sub-game perfect equilibrium is a set of Sender beliefs that the offer contemplated is the minimum that would be accepted and behavior on the part of the Receivers that confirms these beliefs. The only puzzle is how such a set of mutually consistent beliefs developed in the first place and how they are passed on from one generation of player to the next. Using an inter-generational game experimental setting, this paper investigates how "culture" serves as the selection mechanism which solves this puzzle. Culture is then simply a system of beliefs and self-confirming actions which support any one of these non-sub-game perfect Nash equilibria as the accepted solution to the game being played. The outcome is, as Robert Aumann has called it a "perfectly good" Nash equilibrium convention which is just not perfect.

JEL Classification: C91, C72

[*] Professor of Economics, New York University, [+] Associate Professor of Economics, Rutgers University. 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 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

The Ultimatum Game [1]and the experiments surrounding it, have presented economists with a puzzle that they have struggled to explain .(See Thaler (1988), Camerer and Thaler (1995), and Roth (1995) for surveys.) More precisely, given a theory that assumes that people are rational and that such rationality is commonly or at least mutually known, the puzzle of explaining why sub-game perfect equilibria do not appear experimentally has opened the door to a variety of explanations including inequality aversion, (Bolton (1991), Bolton and Ockenfels (1998, 2000 forthcoming), and Fehr and Schmidt (1999)), reciprocal fairness Rabin (1993), Dufwenberg and Kirchsteiger (1998), Blount (1995), "quasi-maximin preferences" (Charness and Rabin (2000)), etc. But as Robert Aumann (see, van Damme (1998)) has pointed out, while there may be only one sub-game perfect equilibrium to the Ultimatum Game, there are an in¯nite number of Nash equilibria. All that is needed to maintain a non-sub-game perfect equilibrium is a set of beliefs on the part of the Sender that the o®er contemplated is the minimum that would be accepted and behavior on the part of the Receivers that con¯rms these beliefs. Hence nothing strange is occurring if we observe one of these non-subgame-perfect equilibrium outcomes at the end of an experiment. We would simply be observing "perfectly good" Nash equilibria that are just not 'perfect' "[2]. The only puzzle is how such a set of mutually consistent beliefs developed in the ¯rst place and how they are passed on from one generation of player to the next.

We consider "culture" to be the selection mechanism which solves this puzzle. More precisely, if the same Ultimatum Game were played by di®erent sets of social agents, these players may create di®erent conventions of behavior for the play of the game because di®erent communities of agents are capable of creating di®erent "cultures" which select di®erent Nash equilibria as the solution to the Ultimatum Game. Culture is then simply a system of beliefs and self-con¯rming actions which support any one of these non-sub-game perfect Nash equilibria as the accepted solution to the game being played.

This idea of culturally supported Nash conventions is illustrated in an article by Alvin Roth, Vesna Prasnikar, Masahiro Okuna-Fujiwara and Shumuel Zamir (1991) who compare the behavior of subjects engaged in an Ultimatum Game across four countries: the United States, Japan, Israel, and Yugoslavia. At the end of their paper they conclude that the di®erence in the behavior they observe

[1] In an Ultimatum Game, there are two types of players: Senders and Receivers. The Sender is initially allocated a certain amount of money, say $10, that he or she has to divide into two amounts, $x and $10-$x. The amount $x is proposed to the Receiver as his portion which the Receiver could either accept or reject. If the Receiver accepts the proposal, the payo®s would be $x for the Receiver and $10-$x for the Sender. If the Receiver rejects the proposal, each subject's payo® would be zero.

Hence the Ultimatum Game is a two stage game. In Stage 1 the Sender makes an o®er and in Stage 2 the Receiver either accepts or rejects.

The sub-game perfect equilibrium requires that the Sender o®er the Receiver " and that the Receiver accepts this small o®er. Experimental evidence strongly contradicts this prediction with Senders o®ering amounts between $4 and $5 as their modal o®er.

[2] Comment by Robert Aumann. (See, van Damme (1998) interview of Robert Aumann).

is not the result of di®erences in the type of people inhabiting these countries (i.e. Israelis are not more aggressive than Americans) as much as a cultural di®erence that has emerged in these countries which leads them to a di®erent set of mutual expectations about what o®ers are acceptable.

> "This suggests that what varied between subject pools is not a property like aggressiveness or toughness, but rather the perception of what constitutes a reasonable o®er under the circumstances"
>
> (Roth et al., (1991, p. 1092).

The question of interest is to try to explain how such culturally determined outcomes ("reasonable o®ers") or conventions come about, how they are passed on from generation to generation of players in an intergenerational game, and how they change when circumstances dictate. It is our belief that we can observe such a culture in the laboratory with paid human subjects and observe the creation of such conventions ¯rst hand. This is what we attempt here.

In our experiments, the evolution of conventions is Lamarckian 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, once acquired, could be passed on in a non-genetic manner. We are interested in these transitions and the evolutionary dynamics they imply from an experimental perspective.[3]

This paper is the second of a series of papers on what we call "Inter-generational Games". (See Schotter and Sopher (2000)) for an analysis of Inter-generational Coordination Games.) In these games a sequence of non-overlapping "generations" of players play a stage game for a ¯nite number of periods and are then replaced by other agents who continue the game in their role for a similar length of time. Players in generation t are allowed to communicate with their successors in generation t+1 and advise them on how they should behave. In addition, they care about the succeeding generation in the sense that each generation's payo® is a function not only of the payo®s achieved during their generation but also of the payo®s achieved by each of their children in the game that is played after they retire. These types of games have proven to be very useful in describing the evolution of conventions of behavior in coordination games (see Schotter and Sopher (2000) and also provide many insights into how the terms of trade or sharing contracts emerge endogenously in the Ultimatum Game.

What we ¯nd (as summarized in the form of 10 Observations) is that advice is a key ingredient in explaining the behavior of subjects in our Inter-generational Ultimatum games. More precisely, when advice exists it tends to be followed in that it serves as the key variable explaining the o®ers sent by Senders . In addition, from examining the written advice o®ered from one generation of Sender

[3] Our research program is not very di®erent from that studied by Peyton Young (1998), (1993) who investigates an evolutionary model to explain the emergence of contractual conventions in which, over time, one conventional rate of compensation or sharing contract gets established to regulate the interaction of economic agents.

to the next, we conclude that arguments of fairness or backward induction are infrequently relied on by subjects in rationalizing the o®ers they suggest to their successors. What is relied on are arguments of expected payo® maximization. In fact, even when 50-50 splits, the hallmark of equity o®ers, are proposed, they are mostly proposed because the Sender perceives the probability of having lesser o®ers accepted to be unacceptably low. The advice of Receivers is di®erent, however, more often relying on fairness and spite arguments to justify behavior. Third, behavior is "more conventional" when advice is allowed in that the data of those experiments allowing advice is more easily organized if one posits that subjects were following conventions. Fourth, in our Intergenerational Ultimatum games Senders appear to "leave money on the table" in that they consistently make o®ers to their Receivers which are greater than the Receivers minimum acceptable o®ers. While this may appear to contradict the claim that the conventions we observe are equilibrium conventions, we demonstrate that our observed Sender behavior can be rationalized as being part of a Bayes-Nash equilibrium to a game with incomplete information played over time by Senders and Receivers. It can also be explained by observing the pessimistic nature of Sender beliefs which tends to lead them to o®er more than they, in fact, need to have their o®ers accepted.

We will proceed as follows. In Section 2 we will describe our experiment and experimental design. Section 3 reports our results, and Section 4 presents some conclusions.

2 The Experiment: Design and Procedures

2.1 General Features

Given our discussion above, it should be clear that any experiment on intergenerational 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-speci¯ed game repeatedly with the same opponent for a pre-speci¯ed 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 payo®s to any subject should be equal to the payo®s earned by that generation during their lifetime plus a discounted payo® which depends on the payo®s 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 e®ort 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 either at the Experimental 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 intergenerational games, a Battle of the Sexes Game (BOSG) (see Schotter and Sopher (1999) for a discussion of this game), an Ultimatum Game(UG) in which they were asked to divide 100 francs, and a Trust Game (TG) as de¯ned 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 di®erent opponent. For example, consider the following table:

Table 1:Rotation Scheme For Subjects

	Player		Game	
Period		Battle of Sexes	Ultimatum	Trust
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 ¯nished 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 subject has played each game against a di®erent opponent who has not played with any subject he has played with before. Each generation played the game once and only once and their payo® was equal to the payo® they received during their generation plus an amount equal to 1/2 of the payo® of their successor in the generation $t+1$ that followed them. (Payo®s 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 di®erent opponent in each period, however.

As a result of this design, when we were ¯nished 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 di®erent 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. [4]

In this paper we will report the results of only the Ultimatum Game played. In our Ultimatum Game, subjects were randomly assigned to role of Sender or Receiver. The Sender was initially allocated 100 units of a ¯ctitious laboratory currency called francs, which were later converted into dollars at the rate of 1 franc equals $.10. The task of the sender was to divide this 100 francs into two amounts, x and 100-x. The amount x was proposed to the Receiver as his portion which the receiver could either accept or reject. If the receiver accepted the proposal, the payo®s would be x for the Receiver and 100-x for the Sender. If the receiver rejected the proposal, each subject's payo® would be zero.

The procedures used in playing all games were basically the same. When subjects started to play any of the three games, after reading the speci¯c instructions for that game, they would see on the screen the advice given to them from the previous generation. In the Ultimatum Game Sender advice was in the form of a suggested amount that the previous Sender advised his or her successor to o®er. For the Receiver it was a suggested minimal accepted o®er that the previous Receiver suggested as the cut o® point for rejection. In addition, subjects were allowed to write free-form messages to their successors o®ering 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 subject could scroll through some subset of the previous history of the generations (perhaps all depending on the treatment) before it and see what each generational Sender o®ered and its acceptance or rejection. 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 their generational opponent was likely to do. A Receiver was also asked to state the minimum acceptable o®er that he or she agreed to accept if it were o®ered. To elicit beliefs we used a proper scoring rule which made truthful revelation optimal. The minimal acceptable o®er was not elicited in an incentive compatible manner, yet we are able to check if they are meaningful by observing if they are violated by the Receivers in their acceptance behavior. Our belief elicitation procedure worked as follows:

For the Receiver, we asked what they thought the probability was of receiving any amounts in the intervals 0-10, 10-20, 20-30, 30-40, 40-50, 50-60, 60-7-,

[4] As far as we know, this is the record for economic experiments.

70-80, 80-90, 90-100. In other words, we asked them to enter a vector $r = (r_1;r_2;r_3;r_4;r_5;r_6;r_7;r_8;r_9;r_{10})$, with $\sum_{k=1}^{10} r_k = 100$, indicating the probabilities de¯ned above.[5]. [6] Receivers were rewarded for their predictions in experimental points which were converted into dollars at the end of the experiment as follows:

Let $r = (r_1;r_2;r_3;r_4;r_5;r_6;r_7;r_8;r_9;r_{10})$ indicate the reported beliefs of the Receiver. Remember that these are the Receiver's belief that the amount sent will be contained in one of ten disjoint intervals. Since only one such amount will actually be sent, the payo® to player i (the Receiver) when an amount in interval l is chosen will be:

$$\pi_l = 20;000 - \left\{ ((100 - r_l)^2 + \sum_{k \neq l} (r_k)^2 \right\} : \qquad (1)$$

The payo®s 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 $r = (r_1;r_2;r_3;r_4;r_5;r_6;r_7;r_8;r_9;r_{10})$. If their opponent chooses to send an amount in interval l, then the subject would have been best o® if he or she had put all of their probability weight on l. The fact that he or she assigned it only r_l means that he or she has made a mistake. To penalize this mistake we subtract $(100 - r_l)^2$ from the subject's 20,000 point endowment. Further, the subject is also penalized for the amount he or she allocated to the other nine intervals , by subtracting $(r_k)^2$ from his or her 20,000 point endowment as well. The worst possible guess, i.e. putting all your probability mass on one interval only to have your opponent choose another, yields a payo® 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. [7]Telling the truth is optimal.

To elicit truthful beliefs from the Sender we do an equivalent procedure. The Sender is going to o®er an amount to the Receiver who is going to either accept or reject. hence, we ask the Sender to assign probabilities to the acceptance or rejection of any o®er in our ten intervals. For example, the Sender must enter 10 probability vectors describing the probability that he thinks the Receiver will accept or reject any o®er in this interval.. For example, let us index the intervals by k = 1,2,..., 10. Then the Sender would type ten probability vectors into the computer of the following form: $r_k = (\pi_a^k; \pi_r^k)$: Here π_a^k is the probability that if an amount in the k^{th} interval is sent it will be accepted while π_r^k is the complementary probability that the o®er will be rejected. From this point on the payo®s are identical to the ones de¯ned above but they are de¯ned conditional

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

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

[7] An identical elicitation procedure was used successfully by Nyarko and Schotter (1999) in their analysis of zero sum games and Schotter and Sopher (2000) in their investigation of inter-generational Battle of the Sexes games.

on the amount sent. For example, say that an amount in the k^{th} interval was sent, the Sender predicted that if he or she sent that amount it would be accepted with probability π_a^k, and it turns out that the o®er was accepted. Then that Sender's prediction payo® would be de¯ned as follows:

$$\Pi_k = 20{,}000 - \left\{(100 - \pi_a^k)^2 + (\pi_r^k)^2\right\}: \tag{2}$$

In other words, if the o®er was accepted but the Sender only predicted that it would be accepted with probability π_a^k;the payo® function penalizes him or her by subtracting $(100 - \pi_a^k)^2$ from his or her 20,000 point endowment. It also subtracts $(\pi_r^k)^2$ since that is the probability predicting that the o®er would be rejected which it was not. An analogous payo® can be de¯ned if the o®er was rejected. Note that since the Sender knows how much he or she will send before he makes his prediction, his reported probabilities are meaningful only for that interval since all the others have zero probability of being relevant. Hence, nothing guarantees that these reports are truthful for amounts in intervals not sent yet, the scoring function should be incentive compatible for the beliefs in the interval of actual amount sent. With this proviso, we will still refer to these "out of equilibrium beliefs" at various points and use them as truthful reports.[8] As you will see, however, none of our more important claims rely on this information.

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, the maximum earnings that could be earned in the prediction part of the Ultimatum Game was only $2.00 as opposed to the maximum payo® in the game itslef of $10.00). The fear here was that if more money could be earned by predicting well rather than playing well, then a Sender might want to o®er the full 100 points to the Receiver knowing that it will be accepted for sure and predict that outcome. This actually happened only once.

It is interesting to note that our experiment provides a whole host of data and information that is missing in most if not all other studies of the ultimatum game. For example, since we elicit beliefs we are able to track the beliefs of generational agents over time. This is important since a convention of behavior depends very much on the underlying beliefs that people have about each other (what Schotter (1981) calls the "norms of society"). In addition, we are able to observe what the subjects report as their true willingnesses to accept. By observing and coding the advice that is o®ered, we are able get another insight into the thinking of our subjects that is not typically available. Hence, our data set involves actions, beliefs, and advice all of which we keep track of as our laboratory society evolves.

[8] Obviously, there is no positive incentive to misrepresents beliefs in these intervals.

2.2 Parameter Specification

The experiments performed can be characterized by a set of parameters $P = \{\Gamma; L_{h_t}; \delta; l, a\}$, where Γ 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, δ 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 Ultimatum Game experiment we set L_{h_t} = t-1, δ = 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 history 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 advice. Treatment I was run for an 81 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 77[9] and our Treatment II of length 66. Our experimental design can be represented by Figure 1:

[Figure 1 here]

3 Results:

In presenting our results we will proceed by presenting a set of observations which we hope to substantiate using the data generated. When we are finished with this exercise, we will turn our attention to the search for conventions of behavior in our Inter-generational Ultimatum Games and present evidence concerning the existence of what we call Weak and Strong (or Bayes-Nash) conventions.

[9] Due to a computer problem we lost one observation from Treatment I so there are only 77 observations instead of 78. .

3.1 Observations

In this section we will present a set of observations about our data and test a set of implied hypotheses which statistically substantiate these observations. We organize our presentation of the results by proceeding systemically and presenting a set of observations about the o®ers of Sender subjects, the advice they were given, their beliefs, and the advice they o®er their successors. We then proceed to look at the analogous behavior of Receivers.

3.2 Sender Behavior

3.2.1 O®ers:

Observation 1:Advice Alone Makes O®ers Less Variable and Lowers O®ers Over Time.

Let M_B; $M_{TI,}$ and M_{TII} be the mean o®er in the Baseline, Treatment I (No History) and Treatment II (No Advice) experiments and let V_B; $V_{TI,}$ and V_{TII} be their associated variances. Then $V_{TI} < V_B < V_{TII}$. In addition, in the Treatment I experiment, o®ers decrease as time progresses but this is not true in the Baseline or Treatment II. Finally, if we look at the mean o®ers made during the last 40 generations, $M_{TI} < M_{TII} < M_B$:

Substantiation

What Observation 1 says is that the variance of o®ers is least in Treatment I, where only advice is present, and greatest in Treatment II where there is no advice. This leads to the conclusion that advice is a key ingredient into making economic behavior in our experiments more orderly.

To explore o®er behavior more systematically, consider Table 2 which present some descriptive statistics about the o®er behavior of our subjects and Figures 2a-2c which presents a set of histograms of the o®ers in each experiment.

Table 2: O®ers of Subjects

	All Generations		Last 40 Generations	
	Mean	Variance	Mean	Variance
Experiment				
Baseline	44.70	223.58	45.66	252.98
Treatment I	37.16	166.35	33.68	183.22
Treatment II	42.45	482.28	43.90	386.59

[Figures 2a-2c Here]

There are some things to note. First note how variable o®ers are in Treatment II, the experiment without advice. In fact, the variance of o®ers is almost three times as great in Treatment II than in Treatment I where subjects have access exclusively to advice (except for a one period history). A series of one-tailed F-tests supports this observation for binary comparisons between with Treatment II and the Baseline ($F_{(65;80)} = 2.16$, $p = .00$) and Treatment II and Treatment I ($F_{(65;76)} = 2.90$, $p = .00$) . The same test found a di®erence between the variances of Treatments I and the Baseline at only the 10% level.

What this indicates is that history does not seem to supply a su±cient lesson for subjects to guide their behavior in a smooth and consistent manner. Advice seems to be needed.

With respect to time, it appears that only in Treatment I do o®ers change over time in a statistically signi¯cant (and negative) manner. To illustrate this point we ran a simple OLS regression of o®ers made on time. In all regressions, except the one run on Treatment I data, time was insigni¯cant at the 5% level. In the Treatment I regression, the coe±cient was negative and signi¯cant at the .003% level. [10] Looking at the mean o®ers in the last 40 generations we see that there is a statistically signi¯cant di®erence in the mean o®ers made between Treatments I and II using a Wilcoxon test ate the 2% level (z = -2.295, p = .02.). No such di®erence exists in the comparison between the Baseline and Treatment II.

It appears then that the inclusion of advice leads subjects to learn from advice that sending lower o®ers is a bene¯cial thing to do. Interestingly, this lesson seems to be a function of advice and disappears when subjects are allowed to view history even when advice is also allowed, as in the Baseline.

Our discussion of o®ers in Observation 1 suggests that we should investigate what factors are important in generating these o®ers. To pursue this question, we o®er the following Observation.

Observation 2: Advice Determines O®ers

Advice is the key determinant in deciding upon o®ers. In fact, subjects tend to follow the advice of their generational predecessor even when their own beliefs suggest that they would maximize their expected payo® by o®ering something else.

Substantiation

Before we present any statistical analysis to back up this observation, consider Figures 3a-3b which plot the times series of o®ers in each of our treatments involving advice against the advice the Sender received (in the Baseline and Treatment I) and also against their subjective payo® maximizing o®er. By subjective payo® maximizing o®er we mean that o®er which, given the elicited beliefs of the subjects, would maximize their expected payo® if sent. Remem-

---

10 Regressions of O®er on Time:

Baseline

	Coef.	Std. Err	t	P>jt
time	.0267615	.0714438	0.375	0.709
cons	43.60648	3.372018	12.932	0.000

$F_{1;79} = :14; p = 0:71$

Treatment I

	Coef.	Std. Err	t	P>jt
time	-.1954361	.0626288	-3.121	0.003
cons	44.79084	2.81133	15.932	0.000

$F_{1;75} = 9:74; p = 0:00$

Treatment II

	Coef.	Std. Err	t	P>jt
time	.0710156	.1427262	0.498	0.620
cons	40.07552	5.500362	7.286	0.000

$F_{1;64} = :25; p = 0:62$

ber, for each potential o®er in intervals 0-10, 10-20,....., 90-100, we have elicited the beliefs of the Sender subject as to the likelihood that such an o®er would be accepted. Hence, we can take an expected value by assuming an o®er at the midpoint of these intervals was sent and multiplying these o®ers by their elicited probabilities. This yields ten distinct values, each representing the expected payo® from sending an o®er in each interval where the expectation is taken over the subjects subjective elicited beliefs. We take the maximum of these ten values whose argmax can take one of the values 5, 15, 25, ...,or 95.

Figure 3a and 3b here

Note from Figures 3a and 3b the absolutely remarkable ¯t between the advice that Senders receive from their predecessors and the o®ers they make. This is true for both the Baseline experiment and Treatment I. Note also, however, that despite the fact that our payo®-maximizing o®er can only take on ten discrete values, they seem to ¯t the pattern of o®ers made reasonably well, though there are many exceptions.

To discriminate between these two variables, we ran a simple linear regression in which our dependent variable was the amount sent and the independent variables were the advice subjects were given and their subjective payo®-maximizing o®er. We ran this for both the Baseline and Treatment I experiments. (Obviously Treatment II did not have advice). These results are presented in Table 3.

Table 3: O®er Behavior in the Baseline and Treatment I

	Baseline			
	coe®	std. er.	t	P>jtj
maxexpec	.11	.11	1.05	0.30
adv sent	.26	.10	2.62	0.01
cons	27.42	6.64	4.13	0.00
$R^2 = :11$	F(2,77) = 4.64, Prob>F=.01			
obs = 81				

	Treatment I			
	coe®	std. er.	t	P>jtj
maxexpec	.08	.07	1.08	0.29
adv sent	.53	.10	5.11	0.00
cons	13.05	5.63	2.32	0.02
$R^2 = :27$	F(2,73) = 13.22, Prob>F = .00			
obs = 77				

These results once again indicate how important advice is for behavior in our experiments. Most striking is that fact that it seems to weigh more heavily in the minds of Senders than do their own beliefs in the sense that when the advice they get contradicts their best response predictions, they seem to opt for following advice rather than best responding to their beliefs.

The question that is raised by these results is how would subjects behave when no advice is given as was true in Treatment II. Would, under these circumstances, subjects concentrate on their best response o®er? Table 4 o®ers the answer to this question since it reports the results of a regression run on Treatment II data in which we regress the o®er made simply on the subjects subjective payo®-maximizing o®er.

Table 4: O®er Behavior in Treatment II

	coe®	std. er.	t	P>jtj
maxexpec	.16	.15	1.08	0.28
cons	34.56	7.78	4.44	0.00
R^2 = :02	F(1,64) = 1.17, Prob>F = .28			
obs = 66				

As Table 4 indicates, subjects do not appear to focus on their best response o®ers even in that experiment where they are not distracted by advice. If advice is so important, however, then it would be interesting to see how this advice varies across experiments which o®er subjects di®erent access to history of the generations before them.

Observation 3: Advice Alone Lowers Advice

The advice given by subjects to their successors is greater in the Baseline than in Treatment I.

Substantiation:

We substantiate this observation by presenting Table 5 which simply presents the mean, median and variance of advice o®ered by subjects in these two experiments along with the results of a simple Wilcoxon test run to test the null hypothesis that these two samples were drawn from the same population.

Table 5: Advice in the Baseline and Treatment I Experiments

Experiment	Mean	Median	Variance	Std. Dev.
Baseline	44.48	47	270.97	16.46
Treatment I	38.25	40	158.71	12.59

Wilcoxon Test: z = 3.12, p=.00

As we see, advice is lower in Treatment I and signi¯cantly so.

Observation 4: Pessimistic Beliefs

Beliefs of Senders tend to be overly pessimistic.

Substantiation

When we call beliefs overly pessimistic we mean the following. For each sub-interval 0-10, 10-20, 20-30, etc. we have elicited the belief of each of our generational Senders as to what they think the chances are that an o®er in this interval would be accepted. Hence, each Sender reports a vector of 10 such beliefs. Call these the subject's Stated Beliefs. At each generation we can also look at the history of play of the game and actually count the fraction of times o®ers in these intervals were accepted (assuming we have some observations in that interval). Call these fractions the subject's Historical Beliefs. By pessimistic we mean that the Stated beliefs of subjects are consistently below their Historical Beliefs.

To substantiate this observation we present Table 6 which provides a set of descriptive statistics to support our claim.

Table 6: Pessimistic Beliefs

Baseline

Interval	0 10	11 20	21 30	31 40	41 50	51 60	61 70	71 80	81 90	91 100
Stated	.10	.18	.28	.40	.55	.74	.78	.83	.82	.89
Historical	.28	.85	.41	.50	.94	1.0	1.0	-	-	1.0
N	3	4	5	13	42	11	1	0	0	2

Treatment I

Interval	0 10	11 20	21 30	31 40	41 50	51 60	61 70	71 80	81 90	91 100
Stated	.18	.24	.38	.52	.69	.83	.84	.86	.87	.88
Historical	0	.35	.34	.94	.96	1.0	-	-	-	-
N	5	5	10	26	30	1	0	0	0	0

Treatment II

Interval	0 10	11 20	21 30	31 40	41 50	51 60	61 70	71 80	81 90	91 100
Stated	.12	.18	.30	.44	.55	.70	.76	.82	.86	.91
Historical	.27	.42	.54	.51	.62	.85	-	1.0	.75	.89
N	4	4	16	12	18	4	0	3	2	3

What you see in this table is, for each treatment, the 10 intervals over which beliefs were elicited along with the average Stated and Historical beliefs of Senders for amounts in that interval. For example, take the interval 41-50 in the Baseline. In the row entitled Stated we have the average over all generations of the subjects' Stated Beliefs for that interval. As you see, on average, subjects felt that an offer in the 41-50 interval would be accepted with probability .55. In fact, if one looks historically at what actually happened when such offers were made (see the row entitled Historical) we find that on average, such offers were accepted with a probability of .94.(There were 42 such generations in which offers in the 41-50 interval were made).[11] Hence, subjects seemed, on average,

[11] A note of clarification here. This Historical beliefs probability is calculated by taking an average of the moving averages defining these historical belief. For example, assume that our experiment had only five periods and say that over those five periods there were four instances where offers in the interval 41-50 were sent (generations 1, 2, 3,and 5) and the Receivers decisions were Accept, Accept, Reject, and Accept. Then the historical beliefs at

to greatly under estimate the willingness of their opponents to accept o®ers in this interval.. The same pattern exists for all intervals and all treatments except in the intervals 81-90 and 91-100 for Treatment II where the opposite is true. Note the small number of observations here, however.[12]

There are some further aspects of Table 6 worth noting. First, note that all mean Stated beliefs are monotonically increasing in the interval so that, on average, subjects did feel that higher o®ers did have a higher probability of being accepted. This was not true for Historical beliefs, however. Also note that when we compare Stated beliefs across treatments, beliefs are always, (except for the comparison of beliefs in interval 91-100 between the Baseline and Treatment I) highest in Treatment I where no history is allowed. This leads to the impression that history tends to make people more pessimistic despite the fact that objectively it should make them more optimistic.

If beliefs are too pessimistic then o®ers would tend to be too high in the sense that Senders could in actuality lower their o®ers and increase their expected payo®s. This raises the question as to whether a signi¯cant portion of the Ultimatum Game puzzle, that subject do not send their sub-game perfect equilibrium o®er and tend to make o®ers in the middle of the allowable range (around 50) , is merely the result of misperceived probabilities. We are able to suggest that this may be true because we have elicited the beliefs of our subjects and are in a position to know what o®er was subjective payo® maximizing given Sender beliefs whereas such information was not available to previous investigators. In the last section of this paper we o®er an alternative equilibrium explanation for the puzzle along with qualitative evidence from the texts of advice indicating that equity can not explain the tendency for subjects to o®er "too much".

3.2.2 Receivers:

These ¯rst four observations explain the behavior of the Senders. The Receivers, however, also exhibited di®erences in their behavior depending upon which treatment they engaged in. The following two observations discuss some of these di®erences.

Observation 5: Advice Makes Receivers Tougher

De¯ning a low o®er as one below 25 and a "tough" Receiver as one who rejects low o®ers, the probability of having a low o®er accepted is lowest in Treatment I, second lowest in the Baseline and highest in Treatment II. In other words, the bigger the role allowed for advice (as in Treatment I where there is no history) the tougher are the Receivers.

---

these generations would be 1, 1, 2/3, 3/4 and the average of these would be .85 which is what we would report in this table.

[12] If we had time to present the full time series of these two belief series, the reader would see that this pattern is persistent over all generations and intervals and does not diminish toward the end of the experiment when there are relatively more observations, at least in some intervals.

Substantiation:

There are more conceptual di±culties involved in analyzing Receiver behavior than Sender behavior. For example, in analyzing the acceptance or rejection behavior of Receivers across treatments, we would ideally like to condition on the o®er made and see if, when identical o®ers are made, they are rejected or accepted with identical frequency across experiments. Unfortunately, the set of o®ers actually made may vary across experimental treatments and hence such a controlled comparison can not always be made.

We can, however, estimate a conditional acceptance function by simply running a logit regression of the dichotomous acceptance variable against the amount o®ered in each of our three experiments and comparing the resulting acceptance functions.. We estimate the logistic relationship,

$$\Pr(\text{x accepted}) = \frac{e^{a+bx}}{1+e^{a+bx}};$$

where x is the amount o®ered and the left hand variable is a f0,1g variable taking a value of 1 if x is accepted and 0 otherwise. This would present us with an estimate of the conditional rejection behavior of subjects in our three experiments and we can use this as a basis of comparison.

The results of these estimations are presented in Figure 4 which plots the resulting estimated acceptance functions and superimposes them on the same graph.[13]

Figure 4 here

Table 6: Acceptance Behavior (Logit)

Table 6a: Baseline

Variable	coe±cient (Std. Err.)	(z	P>jzj
accept			
sent	.10(.03)	3.62	0:00
constant	-2.39 (1.07)	-2.24	0.03
obs = 81			
Pseudo R^2 = :24	LL =-29.62		

Table 6b: Treatment I

Varibale	coe±cient (Std. Err.)	(z	P>jzj
13 accept			
sent	.16 (.04)	4.10	0:00
constant	-4.20 (1.32)	-3.18	0.00
obs = 77			
Pseudo R^2 = :41	LL = 24.71		

Table 6c: Treatment II

Varibale	coe±cient (Std. Err.)	(z	P>jzj
accept			
sent	.022 (.01)	1.52	0:13
constant	{.048 (.61)	-0.08	0.94
obs = 66			
Pseudo R^2 = :03	LL = -39.16		

What we see in Figure 4 is that for low o®ers, the probability of acceptance is ordered in the manner described by the observation, i.e., they are least likely to be accepted when only advice exists (Treatment I) and most likely to be accepted when no advice is present but access to history is unlimited (Treatment II). The Baseline, in which both treatments exist simultaneously, is in between.

While Figure 4 presents a relationship between the likelihood of acceptance and the amount sent, it does not dig deeply into what motivates acceptance behavior. To investigate this, we ran a more elaborate logit estimation in which we tried to explain the dichotomous accept/reject behavior of subjects as a function of their stated minimum acceptable o®er, their expected o®er given their stated beliefs, the advice they received from their predecessors (in the Baseline and Treatment I), the o®er they received and appropriate di®erences among these variables. What we ¯nd is summarized in Observation 6 and substantiated below:

Observation 6: Unful¯lled Expectations Cause Rejection

While unful¯lled expectations about o®ers helps explains rejection behavior, they do not do well in explaining acceptance behavior. Just the opposite is true about a Receiver's stated minimum acceptable o®er, it does a good job at explaining acceptance but not rejection behavior. The advice a Receiver receives from his or her predecessor seems to o®er a compromise explaining both acceptance and rejection behavior fairly well.

Substantiation:

In our experiments we have elicited a great deal of information about Receivers which can be of great help in describing and explaining their rejection behavior. For example, we know what they stated as their ex ante minimum acceptable o®er, and we can calculate the o®er they expect to receive from the Sender using the beliefs elicited beliefs. In addition, we know what they have been advised to accept by their predecessor. By comparing the o®er received to these variables and observing rejection and acceptance behavior, we should be able to learn a great deal about how subjects decide to accept or reject an o®er.

In this section we will ¯rst present a table which records just these results and then explore this acceptance/rejection behavior more formally through a set of models all of which use these variables as inputs.

Table 7 describes the rejection and acceptance behavior of subjects on the basis of the di®erence between the o®er they receive and either their minimal acceptable, expected, or advised acceptable o®er .

Table 7: Rejection and Acceptance Behavior

	Di®erence of O®er and Variable										
	O®er-Exp.O®er				O®er - Min. Acc.				O®er - Advice		
Treatment	B	I	II	Total	B	I	II	Total	B	I	Total
¡Di®erence > 0¢ & Acceptance	33	29	19	81	62	59	43	164	50	41	91
¡Di®erence < 0¢ & Rejection	14	17	15	46	4	3	4	11	10	13	23
¡Di®erence < 0¢ & Acceptance	33	30	27	90	4	0	3	7	16	18	34
¡Di®erence > 0¢ & Rejection	1	1	5	7	11	15	16	42	5	5	10

A number of things are notable in this table. First, the di®erence between what a Receiver was o®ered and what they expected to receive is very good at correctly classifying rejections, but is terrible at classifying acceptances. For example, of the 15 rejections in the Baseline experiment, 14 occurred when the Receiver was not o®ered at least his expected amount. However, of the 66 acceptances in the Baseline, 33 occurred in instances where the amount o®ered was less than a Receiver's expectations. Similar patterns exist in the other treatments as well. This seems to imply that rejection behavior is a "hot" phenomenon perhaps triggered by a de°ation of expectations while stating a minimal acceptable o®er is more a more detached "cold" phenomenon.(See, Brandts J. and Charness, G., (2000)).

The di®erence between a Sender's o®er and a Receiver's stated minimum acceptable o®er has just the opposite e®ect; very good at classifying acceptances but bad at classifying rejections. For example, in the Baseline again, of the 66 acceptances 62 occurred when the o®er was greater than the stated minimum acceptable. (It is not surprising that the result here is stronger than that for the expected o®er since it is almost always the case that a Receivers expected o®er is greater than his or her stated minimum acceptable o®er). However, of the 15 rejections in the Baseline, 11 occurred when the o®er received was greater than the stated minimum.

The di®erence between the o®er and advice received variable is, perhaps, a good compromise, doing a reasonable, though not outstanding, job of classifying both acceptances and rejections. Hence one could state that advice is important for Receivers since it avoids the extremes exhibited by those other variables.

To delve deeper into the relationships of these variables, we consider four models containing four main explanatory variables in logistic regressions of the accept variable (1 if an o®er was accepted, 0 if it was rejected). The variables are, the o®er sent, the minimum acceptable o®er (which we elicited from subjects) the advice received from the previous generation, and ¯nally, Receiver's expected o®er. We estimated di®erent equations, each containing the o®er sent and one of the three variable mentioned: expected o®er, minimum acceptable o®er, and receiver's advice.

Tables 8a-8c contain logistic regression results for each of our three treatments. Two equations are estimated for each of the three explanatory variables noted above. While all models contain the o®er variable on the right hand side, for each treatment one equation contains it alone while another looks at the di®erence of the o®er and each of our three variables. Thus, there are a total of six estimated equations for the Baseline and Treatment I, (two equations for each of three explanatory variables, and four estimated equations for Treatment II (since there is no advice variable in that treatment). For each equation the estimated coe±cient is shown with the standard error of the estimate in parentheses. Signi¯cance levels of the coe±cients are indicated by * (1% level), # (5% level) and & (10% level), respectively. Also shown are a variety of measures of goodness of ¯t, including the log likelihood (LL), the Akaike Information Criterion (AIC) and the percentage of cases correctly classi¯ed by the model (%CC). A case is "correctly classi¯ed" if the predicted probability is

> =.5 (for an acceptance) or <.5 (for a rejection). Four additional diagnostic measures which give a more detailed picture of the performance of the model are also reported: "sensitivity" (probability that a true acceptance is classi¯ed as such by the model), "speci¯city" (probability that a true rejection is classi¯ed as such by the model), "positive predictive value" (probability that a case classi¯ed as an acceptance is a true acceptance), and "negative predictive value" (probability that a case classi¯ed as a rejection is a true rejection).

[Tables 8a-8c here]

For the Baseline game, the model with the o®er and the receiver's advice (column (3) of Table 8a) performs best according to the LL and the AIC. The model with the di®erence in the o®er and expected o®er (column (4)), does a bit better in 3 of the 4 diagnostic measures. In particular, it does a bit better at explaining rejection behavior as seen in Table 8a. For Treatment I, the model with the o®er and expected o®er separately (column (1) of Table 8b) is the best according to LL, AIC, and %CC, and it is also the overall best with respect to the diagnostic measures (two other models, in columns 5 and 6, do better in terms of sensitivity, but only at the expense of speci¯city-they do a very poor job at explaining rejections). For Treatment II, none of the models do very well, but the model with the di®erence in the o®er and the expected o®er (column (4), Table 8c) can be singled out as best. First, it is the only model which is not rejected by the chi-square test for the overall model. It also performs best according to the AIC. The models for Treatment II all do a very poor job of capturing rejection behavior accurately.

Overall, our analysis of receiver behavior shows a less dramatic role for advice than was true for Senders. Though the model with advice was, in some sense, the best model for the Baseline game, the alternative model with the di®erence in o®er and expected o®er seems to perform as well for acceptance behavior and slightly better for rejection behavior. In fact, the expected o®er is the only variable, aside from the o®er actually sent, that appears to be important in all three treatments.

3.3 Advice

While we have concentrated exclusively on the quantitative aspects of our data, we do have a plethora of qualitative data in the form of written advice from one generation to the next. These texts are a treasure trove of insight into what our subjects were thinking not only during their our experiment but, perhaps, even of what subjects think Ultimatum game experiments are about in general. Such data is obviously unique to our experiment.

More precisely, in rationalizing advice in our experiments, a subject might appeal to a number of di®erent motivations. For example, one might advise a particular split (say 50-50) on equity grounds. On the other hand, one might just as well rationalize a 50-50 split on payo® maximizing grounds if one thought that, given your subjective acceptance probabilities, such an o®er is a best response. Such a rationalization need not appeal to equity at all. Alternatively, one may support o®ering only 1 by appealing to the notion of backward induction as is expected of sub-game perfect equilibrium arguments. Backward induction arguments, however, need not only be used to support sending 1. One might advise one's successor that 10 is the best o®er to make because one thinks that there is a threshold below which one's opponent will reject any o®er but above which the o®er would be accepted. The argument here is identical to the sub-game perfect argument but the threshold is not zero. This is how a non-subgame perfect Nash convention can be established. Finally, one can refer to history and look for precedent in what to send or advise one's successor how to make predictions in the experiment since a subject's payo® was also a®ected by how well they predicted what their opponent would do.

In analyzing our advice data we proceeded as follows. First we read each Sender and Receiver comment. After doing this we broke down the Senders comments into 7 sub-groups: Best response Advice (BRA) which basically supports an o®er on the basis of expected payo® maximization, Backward induction advice with a threshold of zero (BI0), Backward induction advice with a strictly positive threshold (BI+), Fairness advice (FA), History-based advice (HBA) which refers to precedent or personal experience in the game, prediction advice (PA), which is advice informing one successor how to make a good prediction, and "other" (OA) which is advice that falls into none of the above categories. We di®erentiate payo® maximizing advice from Backward-Induction advice by noting that a Backward Induction type argument is one where a subject places himself in the position of the Receiver and asks what the minimum he would accept is if he were in those circumstances. If one is rational and cared only about money, the answer would be 0. Backward induction would then suggest sending ". If, for reasons of equity or justice or interpersonal utility, one would reject some positive o®ers, then Backward Induction would ask you to locate your opponents minimal acceptable threshold and make that your o®er. Payo® maximizing behavior looks at the entire distribution of rejection probability and o®ers a best response to it. Hence it approaches the problem di®erently.

For any text we simply recorded any and all types of advice it contained. For example, if a piece of advice contained references to fairness, backward induction, and payo® maximization, we counted all of them in our coding. Our point was not to de¯ne each piece of data as belonging to one and only one category, but rather to count all of the arguments used to bolster the advice given. Hence, in the Baseline where there were 81 generations there is likely to be more than 81 advice codings since the same text can be counted in many di®erent categories. For example, consider the following advice written by the Sender in generation 46 of Treatment I which includes elements of many di®erent types of advice in extremely pure form:

> "The guy before me thought I should send 50. Although, that would be fair, it's not going to maximize your payo®. I was greedy and o®ered 10, thinking that the other guy would accept anything he got, BUT that wasn't the case. They rejected. So my advice is to be a little more generous, so about 30 should do it. Good Luck"

This quote was coded as BRA, BI0, BI+, FA, and HA since it included elements of all of these.

Examples of a pure Backward Induction advice (BI0) were seen in the advice given by the Receivers in generations 34, and 35 of Treatment I who all told their successors to accept anything above 1 if it is o®ered with the following explanations:

> "accept any o®er that is o®ered to u because to reject means that you get nothing. (Generation 34)
>
> "De¯nitely accept anything, or else you get nothing". (Generation 35).

For the Receiver we proceeded as described above except that we changed the categories slightly given the di®ering roles of the subjects. We retained the codings BI0, BI+, FA, HA, PA and OA but dropped BRA since this was not appropriate to the context. We added a category SP (spite) for all those references which suggested retribution if the amount sent was too small and in doing so indicated that relative payo®s were important. Spite and fairness are very close to each other but we separated them because spite has a much more mean-spirited objective. You could lump them together if you wished, however.

A spite statement might read as did this one representing subject 45 in the Treatment I experiment who suggested a minimum acceptable o®er of 40:

> "you're pretty much at the mercy of the other person, if they try to screw you reject it and get them back, otherwise take the money and be happy"

Finally, we added a category PR for prescription which refers to statement that simply suggested a cut-o® point without any real justi¯cation. ("Don't take less than 40 { subject 47 of the Treatment I experiment). These statements are in fact close to BI+ statements and one might be tempted to lump them together, but they did not go all the way and remind their successor that 40 is better than nothing which is what we expect of backward induction thinking.

The results of this coding are presented in Table 9 which present the results of our coding for the Baseline and Treatment I.

Table 9: Coded Advice

Senders

Experiment	Type of Advice						
	BRA	BI0	BI+	FA	HA	OA	PA
Baseline	17	4	21	8	5	10	19
Treatment I	21	6	18	11	23	6	7

Receivers

	PR	BI0	BI+	FA	HA	OA	PA	Spite
Baseline	7	11	3	15	8	11	13	14
Treatment I	7	10	6	4	5	3	13	8

We summarize our ¯nding by the following two observations:

Observation 7: Sender and Receiver Advice Di®er

While the advice of Senders appears to be own payo® oriented and infrequently mentions fairness, Receiver advice re°ects a more inter-dependent utility orientation.

One of the most striking features of Table 9 is the relatively infrequent use by Senders of fairness considerations to support their prescriptions. For example, fairness was not a principle that was invoked often (only 8 times in the Baseline and 11 times in Treatment I). More interesting, however, is that fact that when 50-50 splits are suggested, they are most often supported by payo® maximizing arguments and not equity arguments. For example, in the Baseline, of the 24 cases in which a 50-50 split is suggested, only 7 are supported by references to fairness (a good number leave no written advice, however). In Treatment I, of the 15 times that a 50-50 split was suggested, only 3 were supported by fairness arguments. Hence, observing a 50-50 split does not appear to o®er proof of equity considerations.

Also notable in Table 9 is the infrequent use of pure backward induction arguments. For example, for Senders in the Baseline only four pieces of advice relied on sub-game perfect-like arguments while only six such pieces of advice relied on them in Treatment I. The overwhelming bulk of advice had Senders suggesting an o®er to their successor which, given their assessment of the probabilities of rejection, either maximized their expected payo® or constituted a best o®er given their assessment of the minimum acceptable o®er on the parts of Receivers. For example, there were 38 such pieces of advice in the Baseline and 39 in Treatment I. When backward induction is used, it is usually used to support sending a positive amount based on the assumption that anything less than that amount would be rejected for sure. Hence, backward induction-like arguments are used, but not to justify sending zero but rather to justify sending some positive amount.

With respect to Receivers, the situation is di®erent. Here recommendations for behavior rely much more on fairness and spite-like arguments. For example, in the Baseline spite and fairness are referred to 29 times to support rejecting low o®ers while in Treatment I they are used 12 times. Note that pure backward induction arguments are more prevalent as well used 10 and 11 times for the Baseline and Treatment I. Here, being in the position of the Receiver probably makes it easier to see how accepting anything positive makes sense.

Observation 8: Subjects Create Oral History

When subjects do not have access to history but can pass on advice, they create an oral history through their messages which gets passed on from generation to generation.

Another interesting feature of the advice texts we read was the fact that in Treatment I, where subjects were denied access to any history other than their immediate predecessors, they included references to the meager history available to them far more often than in the Baseline where all subjects could scroll through the history of past generations. What we mean here is while in Treatment I subjects could not °ip through the past generations history and see what occurred, they were able to pass on their own experiences from one generation to the next. Hence, a subject could say that his predecessor told him that his predecessor made o®er x and it was accepted. In fact, it would be possible in such an experiment for all history to be passed on through the medium of advice. The problem, of course, is that if ever one generation fails to pass on a history, it is lost and the historical record must start again from scratch.

As we see, in Treatment I where no history was provided subjects made reference to either their own or their predecessors experience 23 times while they did so only 5 times when a full history was available in the Baseline. This oral history appeared to be an attempt to compensate for the otherwise meager historical setting of the experiment.

4 Searching For Conventions

In this section of the paper we commence our search for the existence of conventions of behavior in the data generated by our inter-generational Ultimatum game. Basically we will be looking for the existence of two di®erent types of conventions which we will call Weak and Strong Conventions. [14] A Weak Convention is an empirical artifact of the data and represents a regularity in the data in which the same or relatively the same o®er is repeatedly made over an interval of time in the experiment. We call these conventions Weak since they need not form a Nash or Bayes-Nash equilibrium to the inter-generational game being played. They are simply regularities in the data. A Strong Convention, on the other hand, is a set of o®ers and responses that do form a Bayes-Nash equilibrium to the game under investigation.

4.0.1 Weak Conventions

In discussing Weak Conventions we are interested in asking what our data would look like if it were generated by a set of economic agents whose behavior was "conventional", i.e. agents following a convention of behavior. Ideally, we would like to construct an index of conventionality for our data sets which would indicate how conventionally determined the data is and then use this index to compare data sets on the basis of their conventionality. Since deriving such an index axiomatically is beyond the scope of this paper, we will proceed here in a more empirical fashion. Our goal is to substantiate the following observation:

Observation 9 : Advice Fosters Convention Creation

In those experiments where advice was allowed, the data is better organized by assuming the existence of conventions of behavior.

To give the reader an insight into what we might mean here, remember that a convention is a regularity in the behavior of a group of subjects in which

[14] Actually if one considers the rule "do what your parents tell you" as a convention, then we have already established that such a conevntion existed in our experiments since advice was shown to be a key ingredient into Sender behavior.

one action is prescribed for behavior and this prescription is passed on from generation to generation. In the extreme, therefore, we would expect the time series of a perfect convention in our experiment to look like the one depicted in Figure 5.

Figure 5 here

Note that this time series is extreme in the sense that from the beginning to the end of the experiment the same o®er is prescribed and made in each generation. Two features of this time series are of note, however. One is the fact that one o®er, in this case 45, is made often (in this case always) over the entire course of the experiment. Such an o®er is "recurrent". The other is the fact that this recurrent o®er is persistent in the sense that the conditional probability of it being o®ered in period t+1 given that it was o®ered in period t is equal to 1.

In searching for Weak Conventional behavior, then, we are searching for o®ers that are recurrent and relatively persistent. However, it is unlikely that the data we will observe will be as orderly as the time series presented in Figure 5. One reason, obviously is that the variance around a recurrent o®er is unlikely to be zero as depicted in that time series. Hence let us call an o®er x an "(Â%) ¡ recurrent o®er over an interval of time T if it is one for which Â% of the o®ers are within an " ¡ neighborhood of x over that time period. The o®er x is "(') ¡ persistent over that time interval if the conditional probability of a generation t+1 o®er being within " of x is equal to '; given that the o®er in generation t was also within " of x is . An "(Â; ') ¡ convention over a time interval T is an "(Â%) ¡ recurrent o®er that is "(') ¡ persistent.

To illustrate the notion of an "(Â; ')-Convention consider the following two time series.

Figures 6a and 6b Here

In Figure 6a we see a hypothetical scatter of points representing the offers made over time in a ¯ctitious inter-generational game experiment. As you can see, the data is neatly organized by positing the existence of two "(Â; ') ¡ conventions one centered at 45 and one at 33. Note that in the ¯rst convention " = 2 while in the second " = 6. The break point between these two "(Â; ') - conventions occurs in generation 44 where there is a clear structural change in the data. Note that for this particular hypothetical time series all observations are included within the "-bands of these conventions, hence Â = 100% and ' = 1:

Figure 6b presents another such scatter diagram but this one is more messy and not as neatly described by "(Â; ')-conventions. In fact, in this Figure two "(Â; ') ¡ conventions are depicted but they do not characterize the data as neatly as the ¯rst. To begin, they exclude many observations and have broader " ¡ bands; the ¯rst band is of width 6 while the second is of width 8. In addition, ' < 1: Hence, because the second time series can only be organized by "(Â; ')-conventions with broader bands, and ones that, despite these bands, still do not include as much data as the first, we say that the time series in Figure 6a is more conventional than the time series in Figure 6b.

Restricting ourselves to using at most two $\varepsilon(\chi, \varphi)$–conventions per treatment, we can say that time series A is unambiguously more conventional than time series B if there exists a break point in time series A, t*, with $\varepsilon(\chi, \varphi)-$conventions to the right and left, such that these $\varepsilon(\chi, \varphi)$-conventions contain more observations and have smaller bands than any comparable break point and $\varepsilon(\chi, \varphi)$-conventions that can be found for time series B and have higher persistence. If we fix the number or fraction of the observations we want contained in the $\varepsilon(\chi, \varphi)-$conventions, (and allow at most one break point) then comparing data according to their conventionality would be equivalent to comparing the area of the bands (length $\times$ width) along with their persistence. For example, time series A would be called "more conventional" than time series B, at the 70% level, if we can find a break in the data at t* (which may occur at time period 0) and bands around the $\varepsilon(\chi, \varphi)$-conventional offers to the left and right of t*, x_l, x_r; such that we can pack 70% of the observations within the bands of time series A containing less area than those of time series B and these observation have greater persistence than do those contained in B. Obviously, the time series in Figure 5 illustrates the highest type of conventionality since the area contained in the bands covering 100% of the data is zero and the persistence is 1.

These considerations provide us a method of substantiating Observation 10.

Substantiation:

To substantiate the claim that advice fosters Weak Conventions we will look at the three time series of offers for our three experiments and search, using at most two $\varepsilon(\chi, \varphi)$-conventions, for the break point and $\varepsilon(\chi, \varphi)$-conventions which best organize the data. We will then compare these decompositions.

To do this we proceed in two steps. First we search over the data in any experiment for the optimal break-point using a switching regression technique. More precisely, we fit two regressions $o_l = c_l + \varepsilon_l$ and $o_r = c_r + \varepsilon_r$ on the data to the left and the right of a "break point" generation t*, where o_j, c_j and ε_j; $j = l, r$; are the offers, regression constant, and error terms of the regressions for the data to the left and right of t*. We then systematically search over all break points until we find that t* for which the sum of the squared residuals of the regressions is minimized. We then test if this break is significant. Such a break point would represent a behavioral shift in the offers made by our subjects. After finding the break points and constants, we systematically search for optimal $\varepsilon(\chi, \varphi)-$conventions around these offers which contain 70%, 80% and 90% of the observations, respectively.

To illustrate our methods, say in the 70% case, we are searching for a break point, along with an $\varepsilon(70\%, \varphi)$-convention to the left and one to the right, which contains 70% of the observed offers within the smallest area. Hence we are searching for a break point along with those minimum bands around the x_l, x_r which minimize the area contained in the bands. The resulting φ is then recorded and is not the object of maximization. We do the same exercise for the 80% and 90% cases.

Table 10 and Figures 7a - 9c present the results of this exercise on our three experiments.

Table 10: Conventionality Metric

Percentage	Baseline (Area, (Persistence))	Treatment I (Area, (Persistence))	Treatment II (Area, (Persistence))
70%	13.77 (.7142)	13.32 (.6086)	32.06 (.5777)
80%	22.66 (.7903)	20.49 (.8064)	39.45 (.6930)
90%	38.22 (.8750)	22.98 (.8358)	50.36 (.7241)

Table 10 presents our area metric applied to the time series of our three experiments at the 70%, 80% and 90% levels along with the associated persistence. The numbers in the cells are the actual areas contained in the bands encompassing di®erent amounts of data (i.e., 70%, 80%, or 90%), the numbers in parenthesis are persistence measures.

As we can see, according to our measure, behavior is most conventional when only advice is present in Treatment I as opposed to either just history, as in Treatment II, or where both advice and history are present, as in the Baseline. The Baseline and Treatment I are uniformly more conventional than Treatment II, where no advice was allowed, i.e. at all three levels, 70%, 80% and 90%, the area containing these amounts of data was smaller in the Baseline and Treatment I than in Treatment II and the persistence is greater. Hence, as Observation 10 implies, advice fosters conventionality.

Figures 7a-9c here

Figures 7a-9c illustrate these relationships. Three things are of interest in these ¯gures. First note generally how scattered are the o®ers in the Treatment II experiment in comparison to the Baseline and Treatment I experiments. It appears as if just having history at one's disposal does not provide as clear a guide to behavior as does having advice. In addition, note that in the Baseline experiment, over the ¯rst 25 generations, o®ers between 48 and 52 were made 17 times (11 times the o®er was exactly 50). In the 19 generations between generations 7 and 25 o®ers between 48 and 52 occurred 16 times. Hence over this time period behavior appeared to be quite neatly conventional. Looking at the graphs for the 70% bands, Figures 7a-7c, gives an indication of exactly how conventions help organize the data when advice is present.

4.1 Strong Conventions

While Weak Conventions are empirical constructs, Strong Conventions are more theory based in that they constitute equilibrium regularities of behavior. The actual de¯nition of a Strong Convention of behavior in our experiments depends upon how you model the inter-generational game being played by the subjects. If one assumes that the game played by our inter-generational agents is a game of complete information, then one would have to conclude that no equilibrium convention of behavior was created in any of our experiments.. The reason for saying this is simple. In order for our subjects to behave in a Nash-like manner we would have to observe o®ers being made which were equal to the minimum stated acceptable o®ers of Receivers. This is true because if the o®ers made were above the minimal acceptable o®ers of the receivers, then the Sender would be able to increase their payo® by sending less and they would be accepted. We have seen in Observation 6, however, that this was not the case since Senders repeatedly left money on the table in these experiments. Hence, if one models our inter-generational game in this fashion, one would have to conclude that the behavior we observed was not equilibrium behavior.

Alternatively, one could think of the game played by our subjects as a game of incomplete information played by a set of non-overlapping generations of players. More precisely, let us assume that the game is played as follows. At the beginning of each generation nature moves and draws a minimum acceptable o®er for the Receiver in that round. This minimum is private information known only to the Receiver. However, the Sender of that round knows that all such minima are drawn independently from an identical distribution $F^{®;¯}(m)$, with support [0,100]. Here $F^{®¯}(m)$ is one of a two-parameter family of such distributions. These minima drawn by generational Receivers are their "types". For Receivers in generation t, m_t de¯nes the true minimum acceptable o®er for that subject which is determined by a host of factors existing outside of the game, i.e., the subjects' essential sense of fairness, feelings of altruism, spite, etc.

After the Receiver draws his minimum, nature moves again and chooses a risk aversion coe±cient for the Sender which is the exponent in his utility function $U = (y)^r$; where y is the ¯nal net payo® of the Sender: In other words; each Sender will be assumed to be identical except for his exponent r. Each r implies a di®erent attitude toward risk and r is assumed to be drawn independently from an identical distribution. r is seen only by the Sender alive in that generation and is private information. The Receiver knows the distribution from which these r's are drawn although that will not matter for his strategy in the game.

After the Receiver views his m and the Sender his r, the Sender makes an o®er, x, to the Receiver who then accepts or rejects. If the Receiver accepts, the Sender gets 100-x and the Receiver gets x. In payo® terms, the Sender gets $U(100\text{-}x) = (100\text{-}x)^r$, while the Receiver gets x. (We can assume that the Receiver is risk neutral, all that matters is his minimum acceptable o®er).

The Bayes-Nash equilibrium for this game is particularly simple. For any generation t, the Receiver has a dominant strategy of accepting all o®ers above m_t and rejecting all o®ers below it. (Remember that m_t is the Receiver's true minimal acceptable o®er. Hence, by de¯nition, he has a dominant strategy). The Sender acts as follows. First he or she will look back at the history of the game up through generation t-1. This will present the Sender with a time series of o®ers made along with a reject or accept decision on the part of the generational Receivers alive at those times. An accept decision by generation t-k implies $m_{t¡k}$ **5** $x_{t¡k}$ while a reject decision implies $m_{t¡k}$ **=** $x_{t¡k}$:Using this information the Sender must estimate ® and ¯ to generate an estimate of $F_t^{®;¯}(m)$ (the generation t's best estimate of $F^{®;¯}(m)$), and then make an o®er that solves the following maximization problem:

$$\max_x E(\pi) = (100 - x)^r F_t^{\alpha,\beta}(x). \qquad (3)$$

If we call x* the argmax of this problem, it satisfies the following first order condition:

$$\frac{(100 - x^*)}{r} = \frac{F_t^{\alpha,\beta}(x^*)}{f_t^{\alpha,\beta}(x^*)}; \qquad (4)$$

where $f_t^{\alpha,\beta}(x^*)$ is the density function associated with $F_t^{\alpha,\beta}(x^*)$. In any generation t, if $x^*_t = m_t$, then the offer is accepted, if not it is rejected.

Before proceeding, let us pause to make some comments about our approach. First, we are not placing a great deal of reliance on the particular functional form assumed for the subjects' utility function. For example, we only introduced the random risk aversion coefficient in the utility function for the seller to allow us to provide some variability in the predictions of the model. Without a stochastic element for the preferences of the Senders, period-to-period offers would vary only slightly yet we observe quite a large amount of variability in the offers made from period to period. Other functional forms for the utility function would not change the qualitative features we care about, but some stochastic variability in the preferences of the Senders is necessary.

Second, we assumed that the true $F^{\alpha,\beta}(m)$ is one of a family of two-parameter functions since it will allow us to proceed empirically and estimate this model. More precisely, say that as time progressed in the experiment each sender repeatedly estimated the following logit accept/reject function using the time series of offers and acceptance decisions made up until generation t-1:

$$\text{Prob(accept } x) = \frac{e^{\alpha+\beta x}}{1 + e^{\alpha+\beta x}}; \qquad (5)$$

where, x is the amount sent and the left hand variable is a {0,1} variable taking on a value of 1 if the offer is accepted and zero if it is not. In other words, each generation learns by running a logit regression of the dichotomous accept/reject left-hand variable on the amount sent using all the data available. Knowing the Receiver's equilibrium strategy, the probability of having an offer of x^*_t accepted at generation t is equivalent to the probability that the Receiver has a minimum reservation value less than x^*_t since the two are informationally equivalent given the Receiver's strategy. Hence, at any time t the best estimate of $F_t^{\alpha,\beta}(x^*_t)$ is the updated estimate of the parameters in the logistic function $\frac{e^{\alpha+\beta x}}{1+e^{\alpha+\beta x}}$.

Given this functional form, in any generation t the optimal offer can be defined by solving the following maximization problem:

$$\max_x E(\pi) = (100)^r \cdot \left(\frac{e^{\alpha+\beta x}}{1 + e^{\alpha+\beta x}}\right); \qquad (6)$$

where the associated first order condition is,

$$\frac{(100 - x)}{r} = \frac{1 + e^{\alpha + \beta x}}{\beta}. \qquad (7)$$

Hence, given the generation t estimate of α_t and β_t, and knowing r, a generation t subject can define his or her optimal bid. Note also that in generation t we can calculate the r_t implied by an offer of x by solving,

$$\frac{(100 - x_t)\beta}{1 + e^{\alpha + \beta x_t}} = r_t. \qquad (8)$$

This structure will allow us to explain all of the salient features of our data. For example, a successful explanation for what happened in our experiments would have to be able to explain how our subjects could be in equilibrium and still have the offers made by the Senders be so much above the minimal acceptable offers of the Receivers. In addition, a successful explanation must also be able to explain why Sender offers varied so much from period to period despite the fact the information varied so little between adjacent generations. Such facts are not surprising in our inter-generational game since each generation involves a new player whose preferences, i.e. risk aversion parameter is different. Hence, even on the basis of relatively identical information, adjacent generations may make very different offers. In addition, given the incomplete information in the problem, it is not surprising that the arg max of (6) is above the unknown minimal acceptable offer of Receiver.

Observation 10: Behavior is Bayes-Nash Consistent

The behavior of subjects in our Baseline and Treatment II experiments is consistent with a Bayes-Nash equilibrium to the game of incomplete information described above if the risk aversion parameters of our Senders are drawn from the distributions presented below.

Substantiation:

Looking at the first order condition presented in (7) we see that if subjects are risk neutral, i.e., r=1 all t, then using our functional form assumption for F(x), (7) can be rewritten as

$$(100 - x) = \frac{F_t(x)}{f_t(x)} = \frac{1 + e^{\alpha + \beta x}}{\beta}; \qquad (9)$$

and the optimal offers for the Baseline and Treatment II are presented in Figures 10a and 10b.

Figures 10a and 10b here

Note that one feature of these figures is that the period-to-period variability in the optimal risk-neutral offers is rather small. This is true because the subjects updated estimates of α and β in their logit functions can not vary much from period to period since they di®er only on the basis of one observation. Note, however, that the time series of actual o®ers does vary substantially from period to period. It is to explain this di®erence that we have introduced the possibility that Senders have utility functions drawn from a one-parameter family of such functions characterized by the risk parameter r, i.e. $U(y) = y^r$: For each o®er x_t made in period t, given the updated estimates for ®$_t$ and ¯$_t$, we can solve (8) for that r_t which would make the observed o®er optimal. These r's are presented as histograms in Figures 11a and 11b and completely rationalize the behavior of our subject Senders.

Figures 11a and 11b here

Note that the overwhelming majority of subjects act as if they were risk averse with r's less than 1. For example, in the Baseline over 75% of the subjects acted as if they had r's less than .704 while in Treatment II 75% of the subjects acted as if they had r's less than 1.073. (The mean and median r's in each experiment were .73 and .36 for the Baseline and .79 and .58 for Treatment II). This indicates that they tend to send more than the risk neutral Bayes Nash optimal o®er.

5 Conclusions:

This paper has attempted to provide an explanation for non-subgame equilibrium behavior in an inter-generational ultimatum game experiment. We posit that what allows such equilibria to exist is a set of conventions of behavior which are supported by beliefs which lead Senders to make strictly positive o®ers and Receivers to reject o®ers that are positive but "too low". These conventions are transmitted from generation to generation through the socializing in°uence of advice o®ered by one generation to the next. When such advice is absent, the associated conventions are harder to establish and the data less able to be organized by conventions.

What our results demonstrate is the overwhelming in°uence of advice on the behavior of our subjects. As we have seen, advice tends to be followed closely by Senders and dramatically lowers the variability of o®ers when it is present. Advice is also important for Receivers a®ecting both their rejection and acceptance behavior. However, for Receivers it appears as if rejection behavior is most a®ected by a de°ation of their expectations since most rejections occur when they receive an o®er that was lower than what they were expecting even if that o®er is above their stated minimal accepted o®er.

In the process of our analysis, we have presented a model in which a Bayes-Nash equilibrium convention is established and have estimated the parameters of one such possible model. While we recognize that all such models are in some sense arbitrary, we do feel that our model is successful in capturing the stochastic nature of the data and providing a plausible framework within which to organize the data generated. Clearly, more work along these lines is needed until we are able to furnish a complete model of culture creation and evolution.

6 Bibliography

References

[1] Blount, S. (1995), "When Social Outcomes Are'n Fair: The E®ect of Causal Attributions on Preferences", Organizational Behavior and Human Decision Processes, Vol. 63, pp. 131-144.

[2] Bolton, G., (1991), "A Comparative Model of Bargaining: Theory and Evidence", American Economic Review", Vol. 81, pp. 1096-1136.

[3] Bolton, G., and Ockenfels, A., (2000, forthcoming), "ERC: A Theory of Equity, Reciprocity, and Competition", forthcoming , American Economic Review.

[4] Bolton, G., and Ockenfels, A., (1998), "Strategy and Equity: An ERC Analysis of the Guth-van Damme Game, Journal of Mathematical Psychology, Vol. 42, pp. 215-226.

[5] Brandts., J., and Charness, G., (2000, forthcoming) ,"Hot vs Cold: Sequential Responses and Preference Stability in Simple Experimental Games", Experimental Economics

[6] Camerer, C, and Thaler, R.(1995), 'Ultimatums, Dictators and Manners", Economic Perspectives, Vol. 9, pp. 209-219.

[7] Charness, G. and Rabin, M., "SOcial Preferences: SOme SImple Tests of a New Model, mimeo, University of California - Berkeley, (2000),

[8] Dufwenberg, M. and Kirchsteiger, G., (1998) "A Theory of Sequential Reciprocity", mimeo.

[9] Fehr, E., and Schmidt, K., (1999) "A Theory of Fairness, Competition, and Cooperation", Quarterly Journal of Economics, Vol. 16, pp. 324-351.

[10] 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.

[11] Rabin, M. (1993), "Incorporating Fariness into Game Theory and Economics", American Economic Review, vol. 83, pp. 1281-1302.

[12] Roth, A.E.(1995), "Bargaining Experiments", in J. Kagel and A. Roth (eds.), Handbook of Experimental Economics, (Princeton University Press, Princeton ,New Jersey).

[13] Roth, A., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S., "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study", American Economic Review, vol. 81, pp. 1068-1095.

[14] Schotter, A. (1981), The Economic Theory of Social Institutions, Cambridge University Press, Cambridge England.

[15] Schotter, A. and Sopher, B., (2000), "Social Learning and Coordination Conventions in Inter-Generational Games: An Experiment in Lamarckian Evolutionary Dynamics" Russell Sage Working Paper #155, Russell Sage Foundation, New York, New York.

[16] Thaler, R. (1988) , "The Ultimatum Game", Journal of Economic Perspectives, Vol. 2, Autumn, pp. 195-296.

[17] van Damme, E.(1998), "On the State of the Art in Game Theory: An Interview with Robert Aumann", Games and Economic Behavior, Vol. 24, pp. 181-200.

[18] Young, H.P., (1993a), "An Evolutionary Model of Bargaining", Journal of Economic Theory, Vol. 59, pp. 145-168.

[19] Young, H.P., (1993b), "The Evolution of Conventions", Econometrica, vol. 61, No. 1, pp. 57-84.

[20] Young, H.P., (1998), Individual Strategy and Social Structure, Princeton University Press, Princeton New Jersey.

Figure 1: Experimental Design

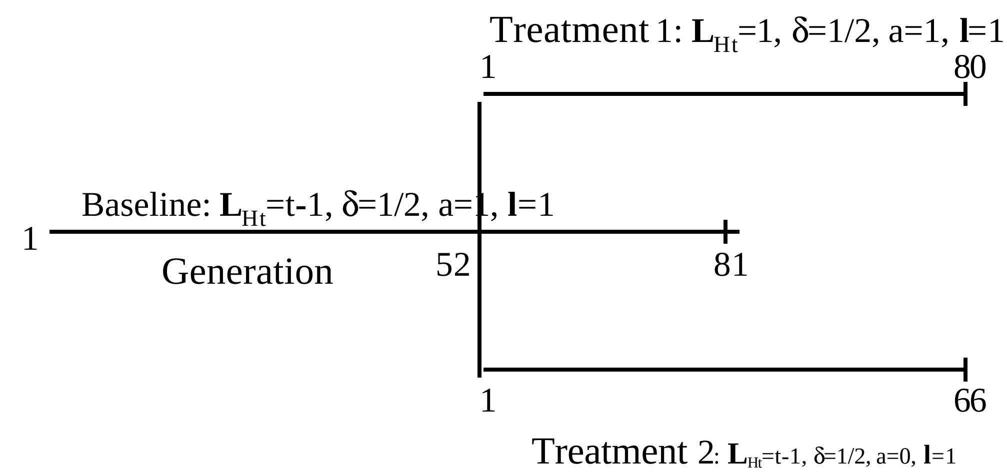

N:\schotter\exp\analysis\sopher\game1\figures.ppt

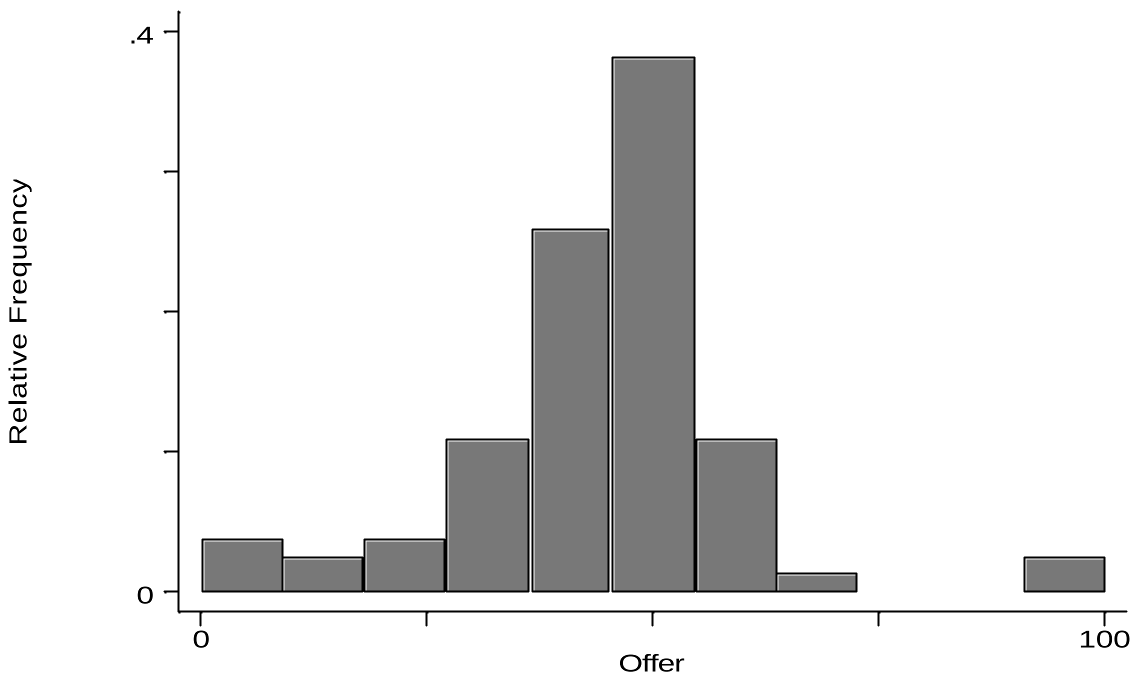

Figure 2a: Offers, Baseline

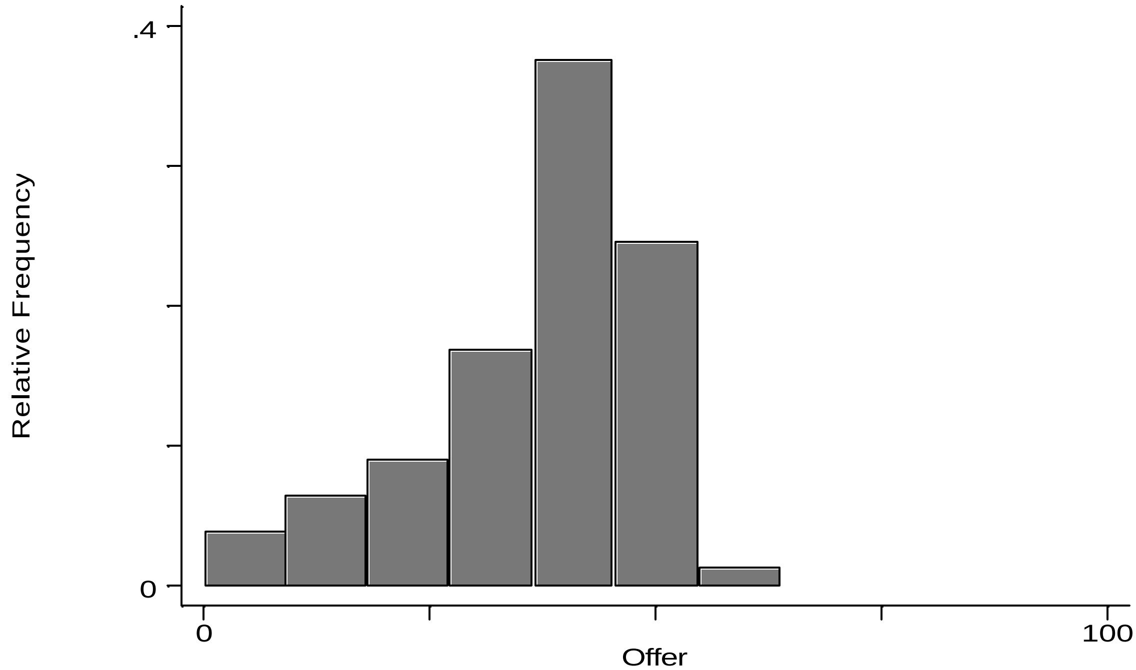

Figure 2b: Offers, Treatment I

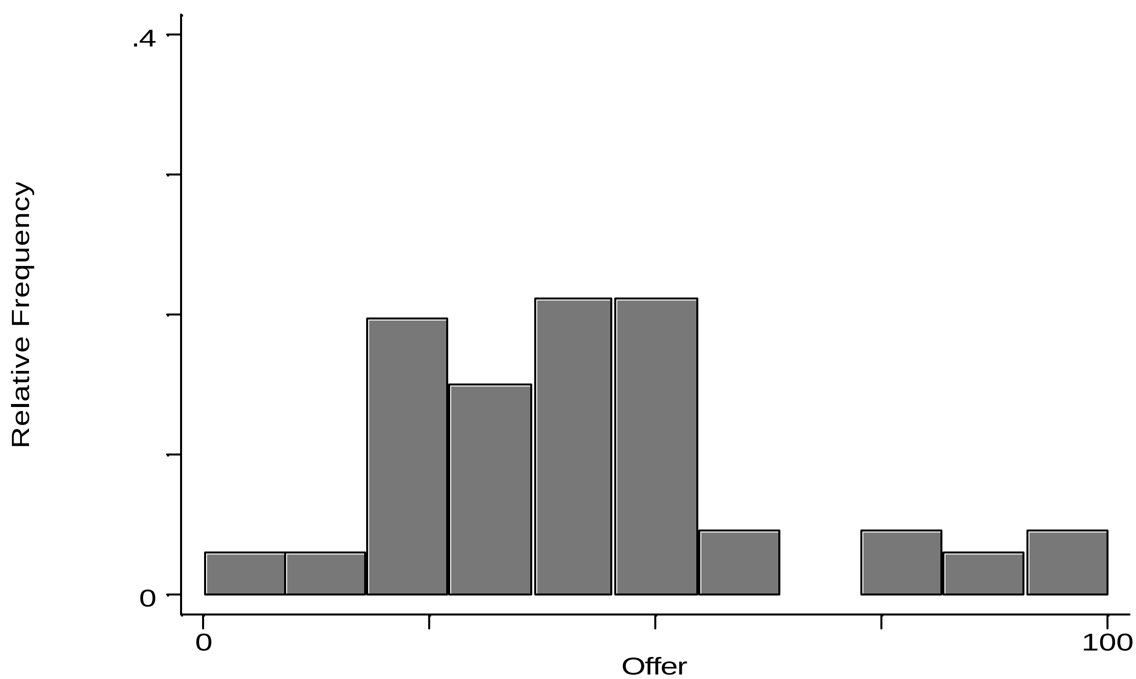

Figure 2c: Offers, Treatment II

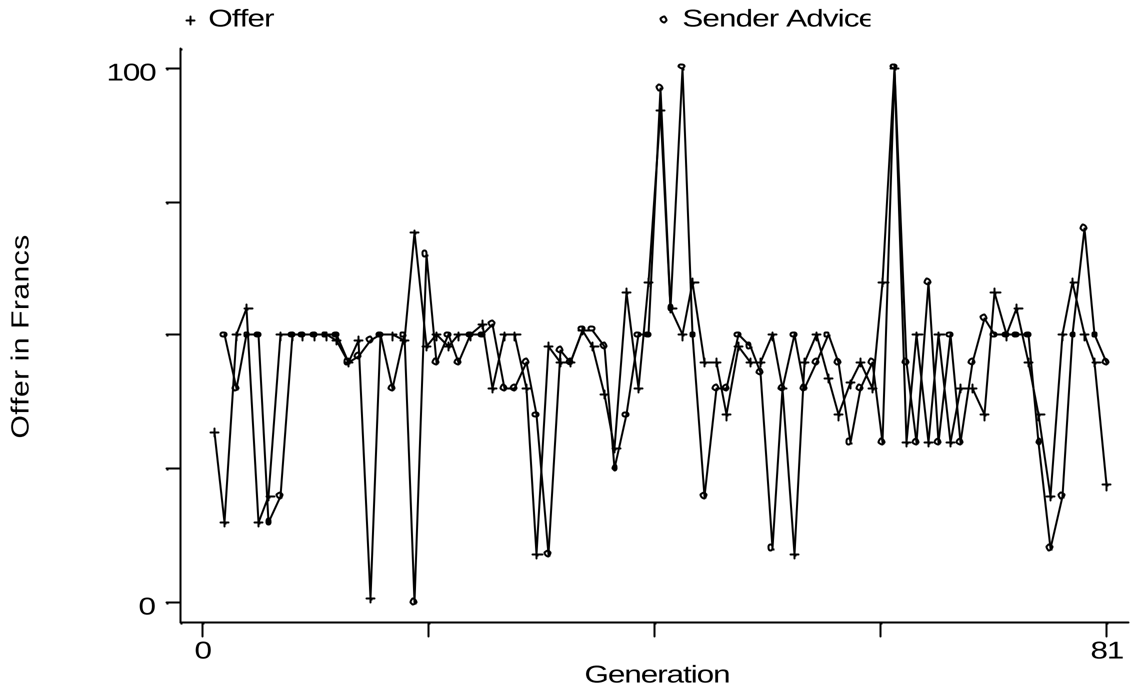

Figure 3a.i: Offer and Advice, Baseline

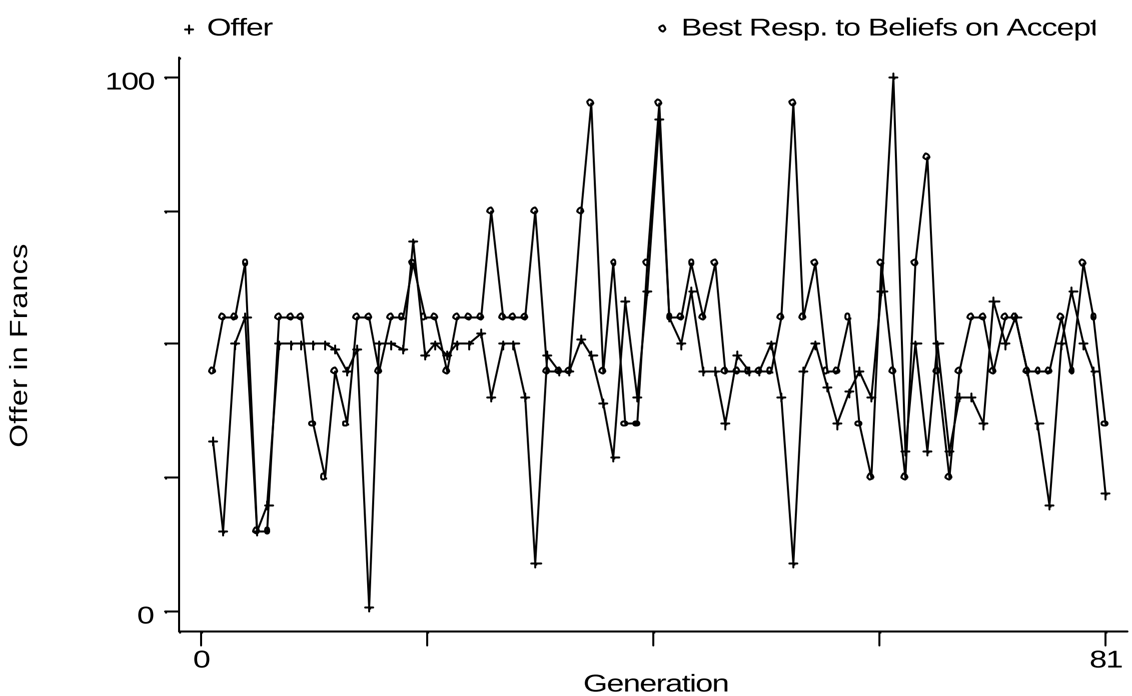

Figure 3a.ii: Offer and Payoff Maximizing Offer, Baseline

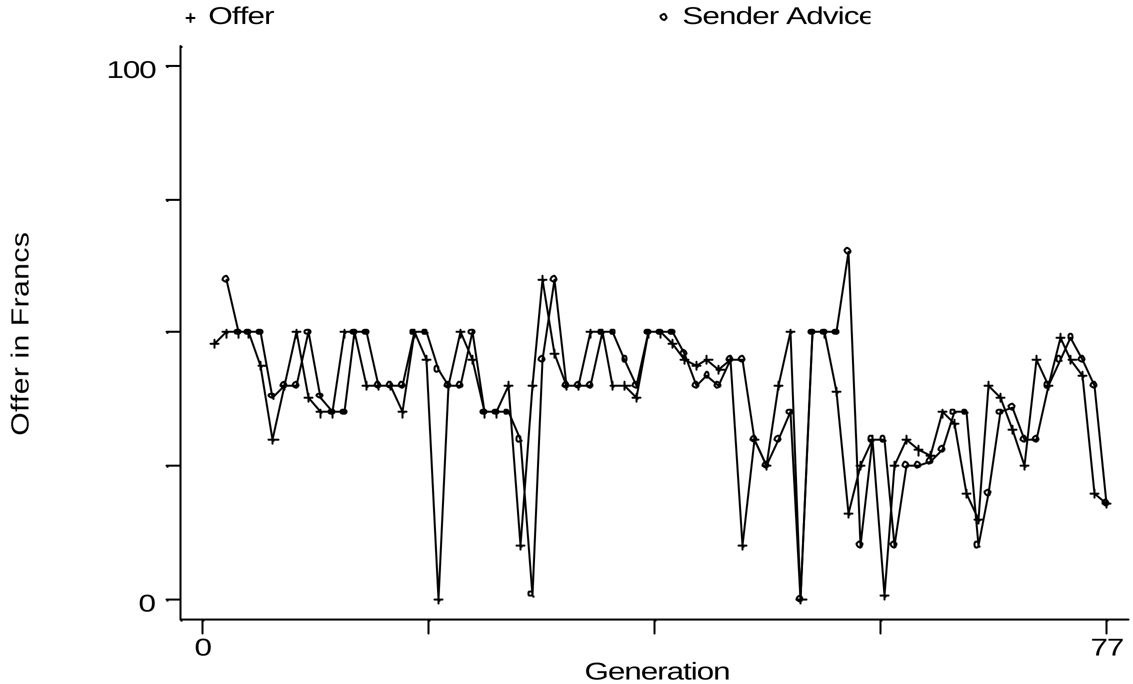

Figure 3b.i: Offer and Advice, Treatment I

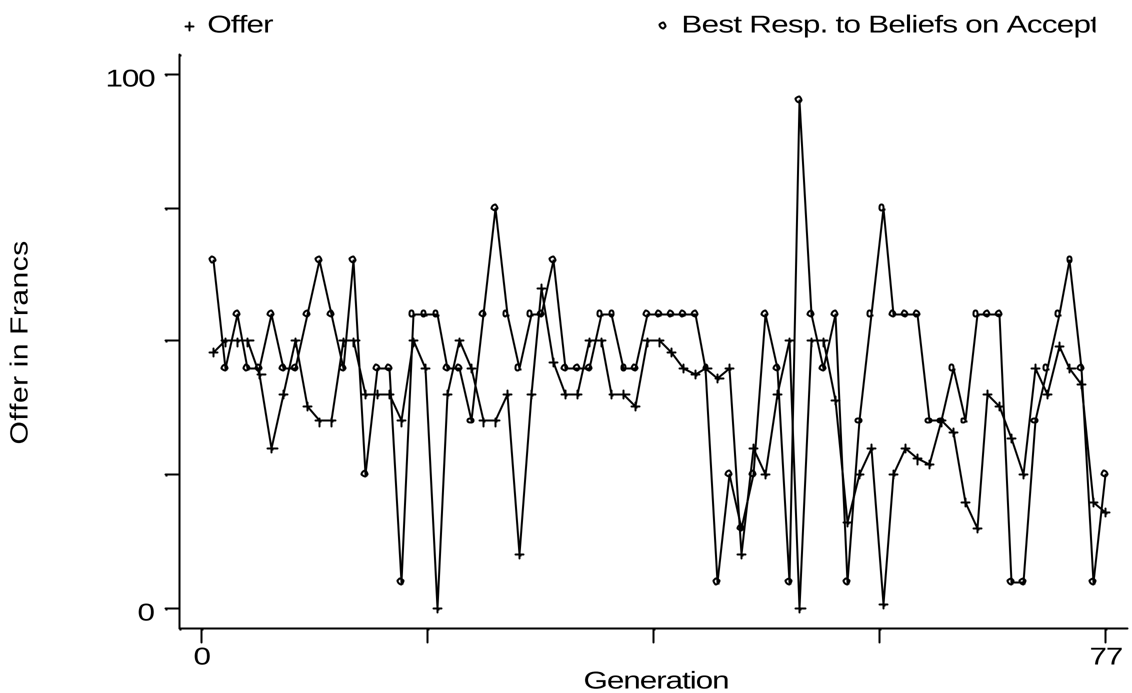

Figure 3b.ii: Offer and Payoff Maximizing Offer, Treatment I

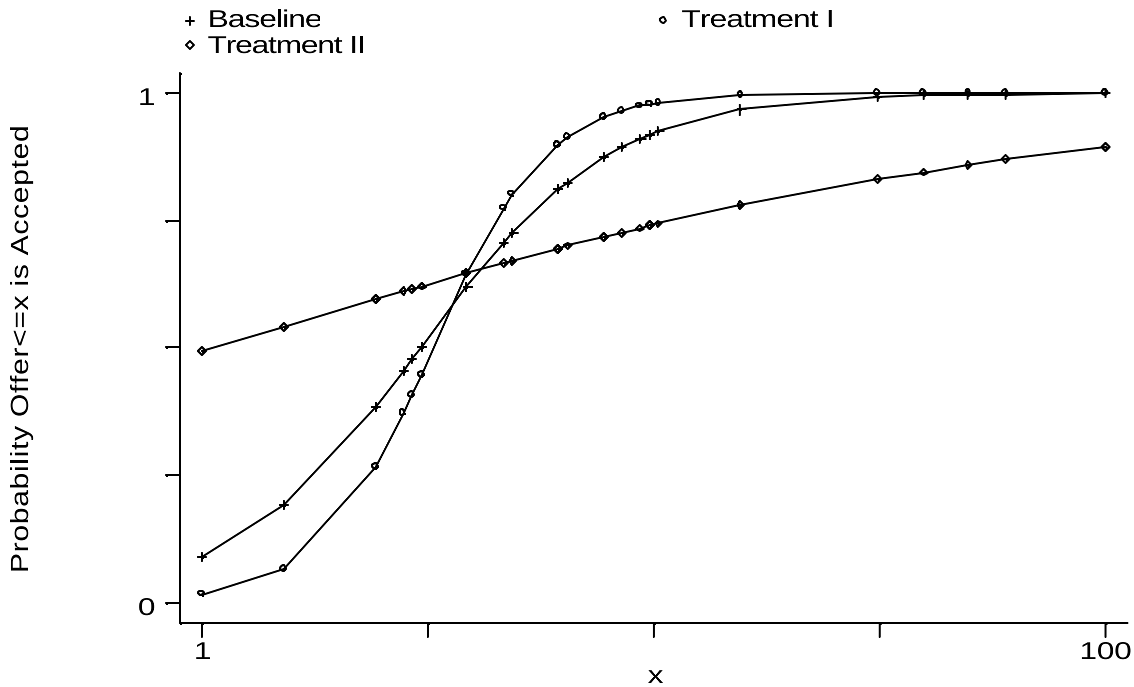

Figure 4: Estimated Probability of Acceptance, by Treatment

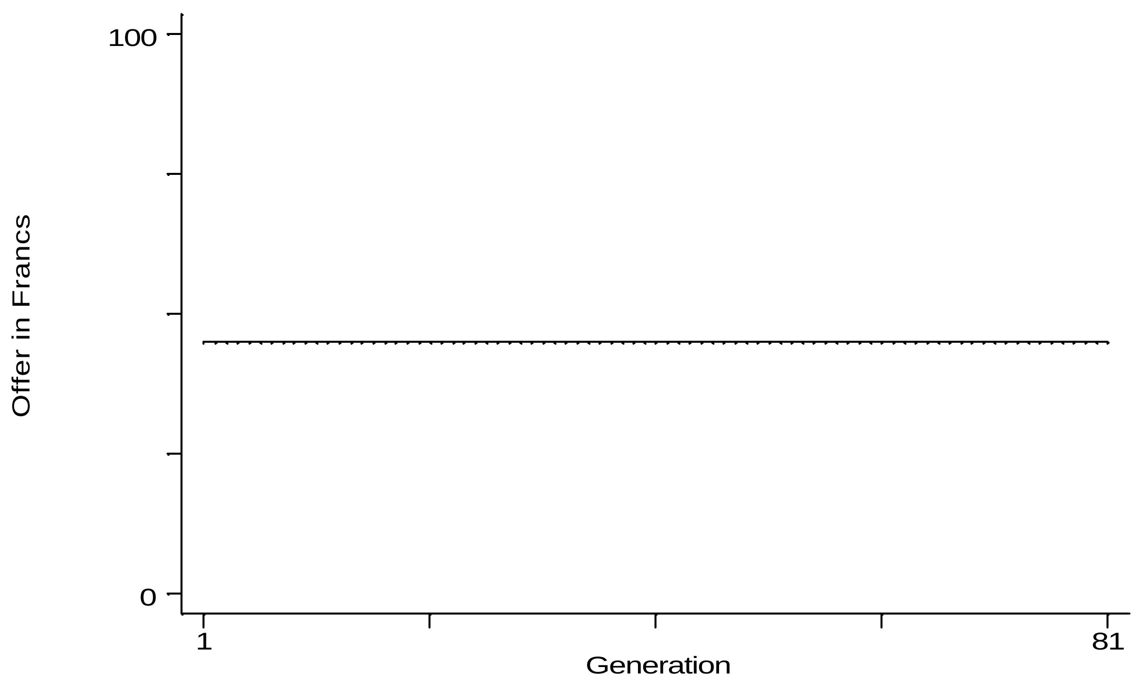

Figure 5: A Perfect Convention

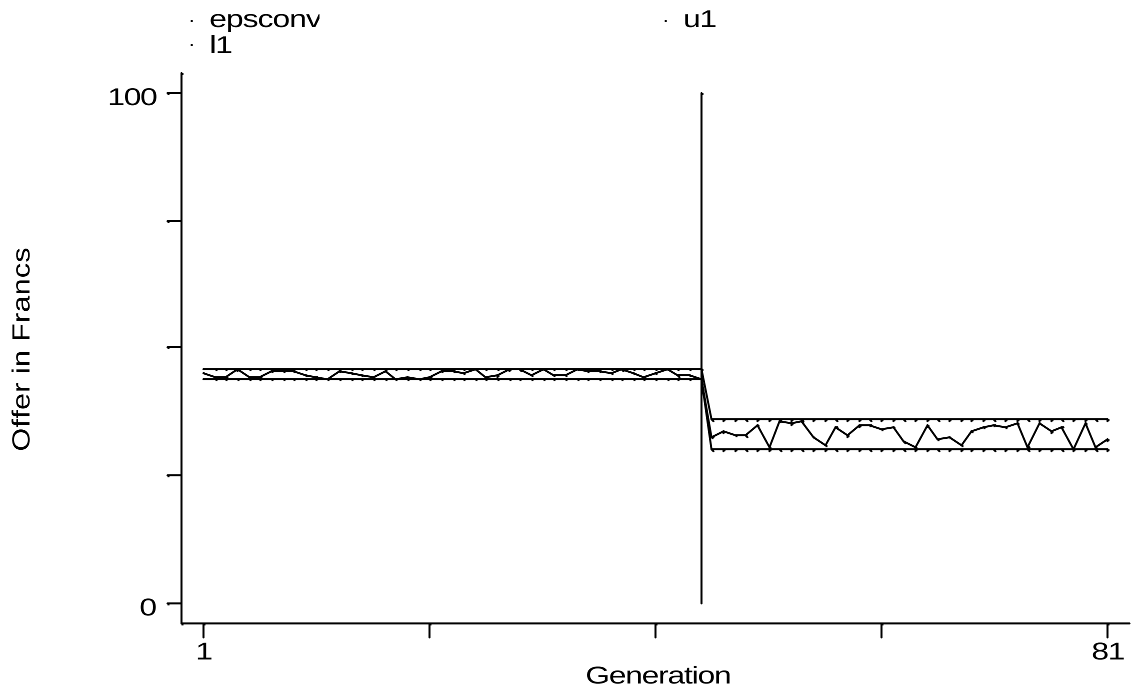

Figure 6a: An Epsilon Convention

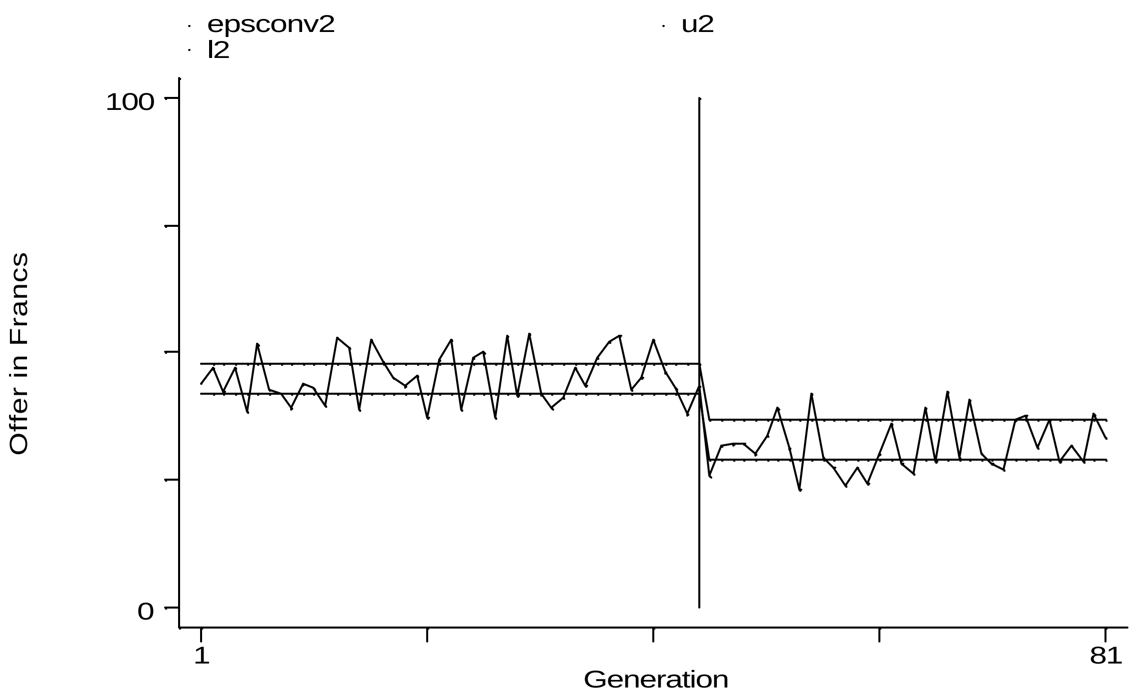

Figure 6b: A Less Conventional Epsilon Convention

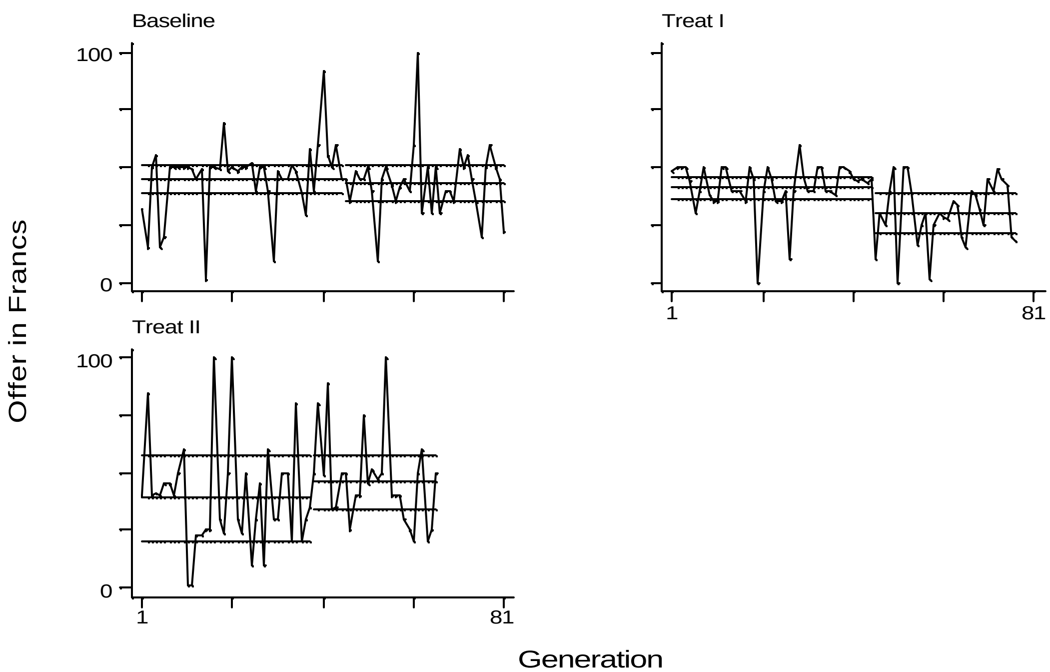

Figure 7: 70% Conventionality Bands for Offers

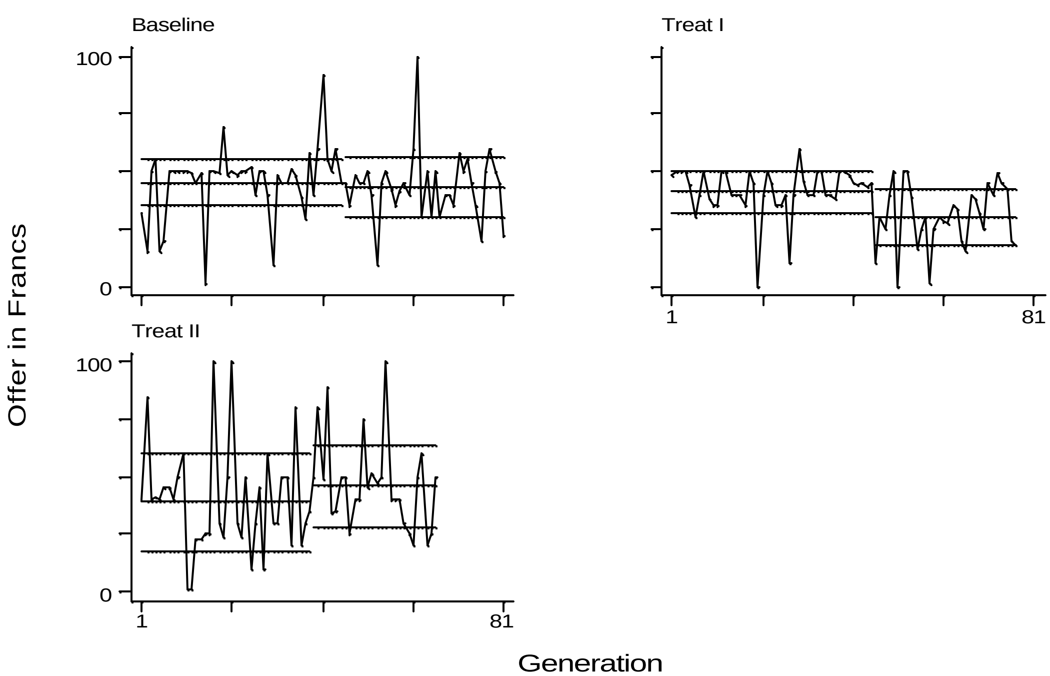

Figure 8: 80% Conventionality Bands for Offers

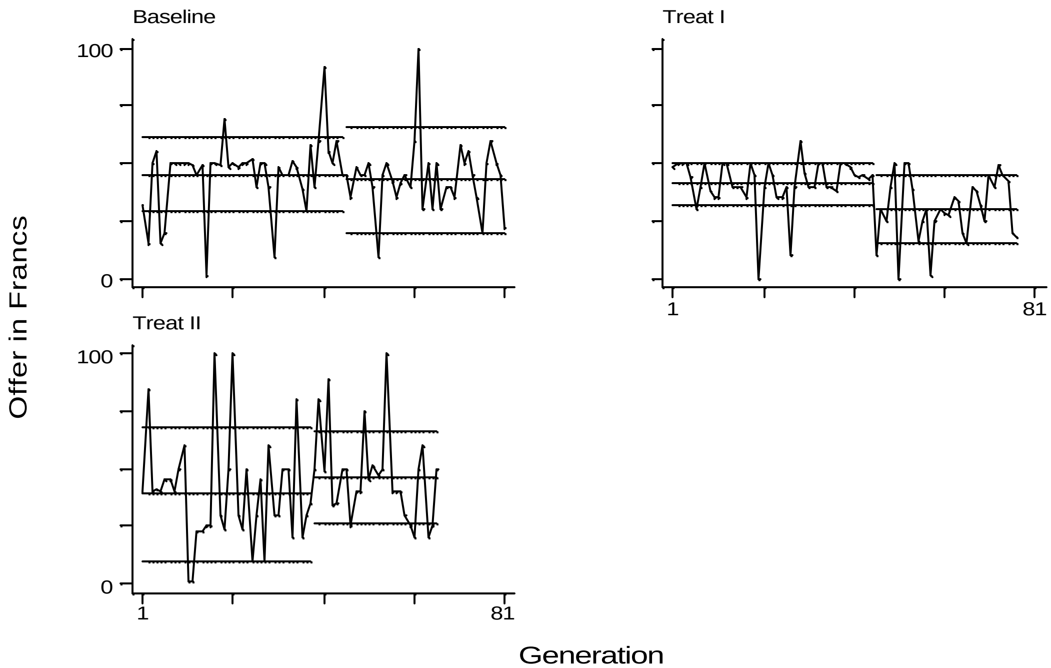

Figure 9: 90% Conventionality Bands for Offers

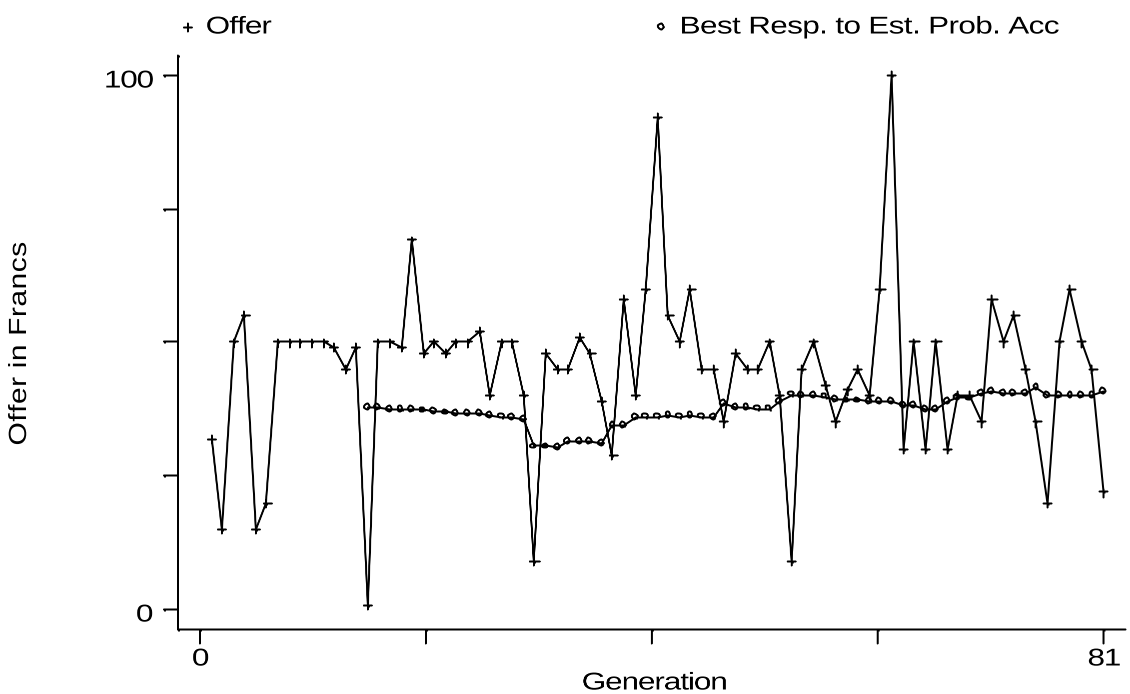

Figure 10a: Offers vs. Optimal Offer, Baseline

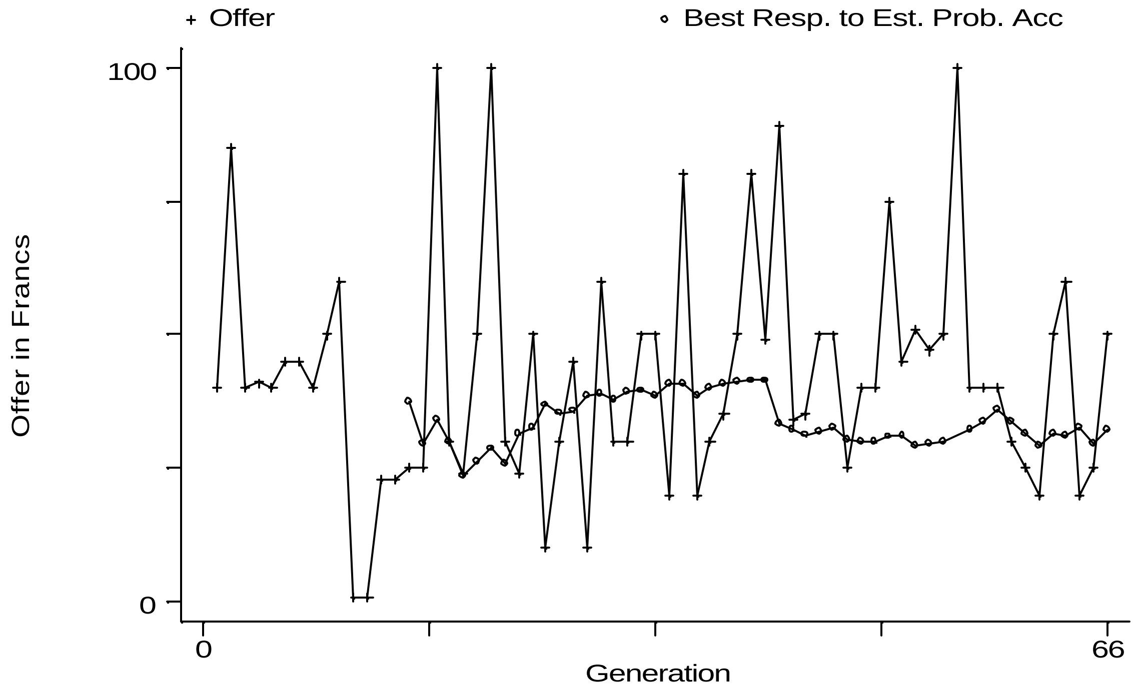

Figure 10b: Offers vs. Optimal Offer, Treatment II

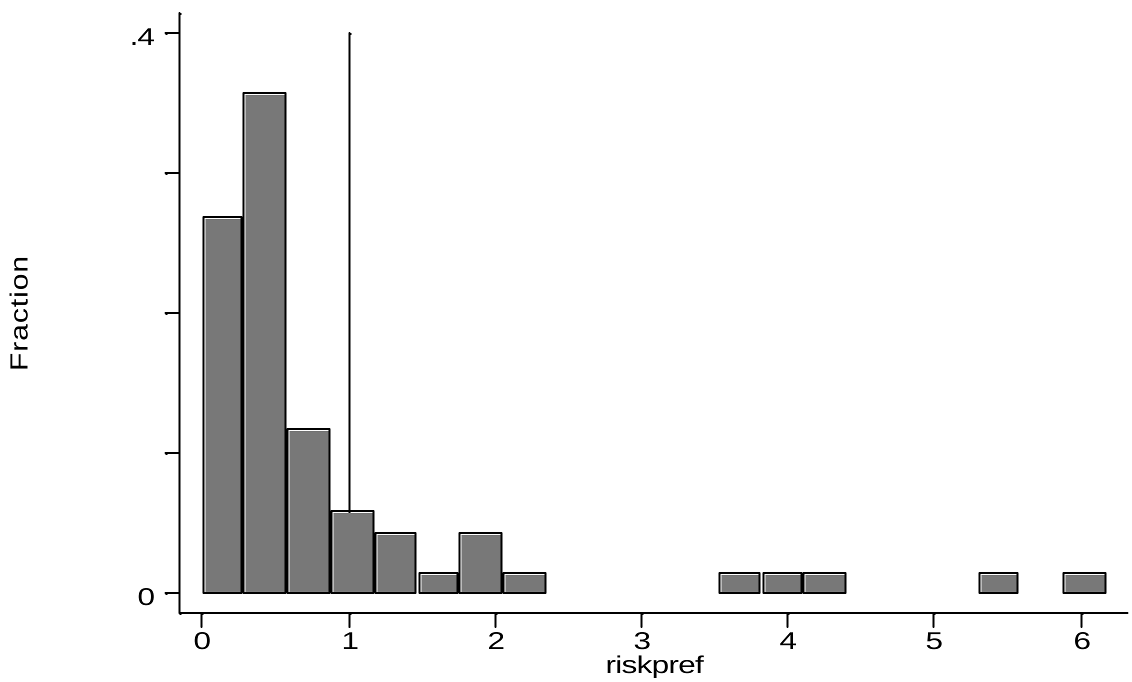

Figure 11a: Risk Aversion Parameter, Baseline

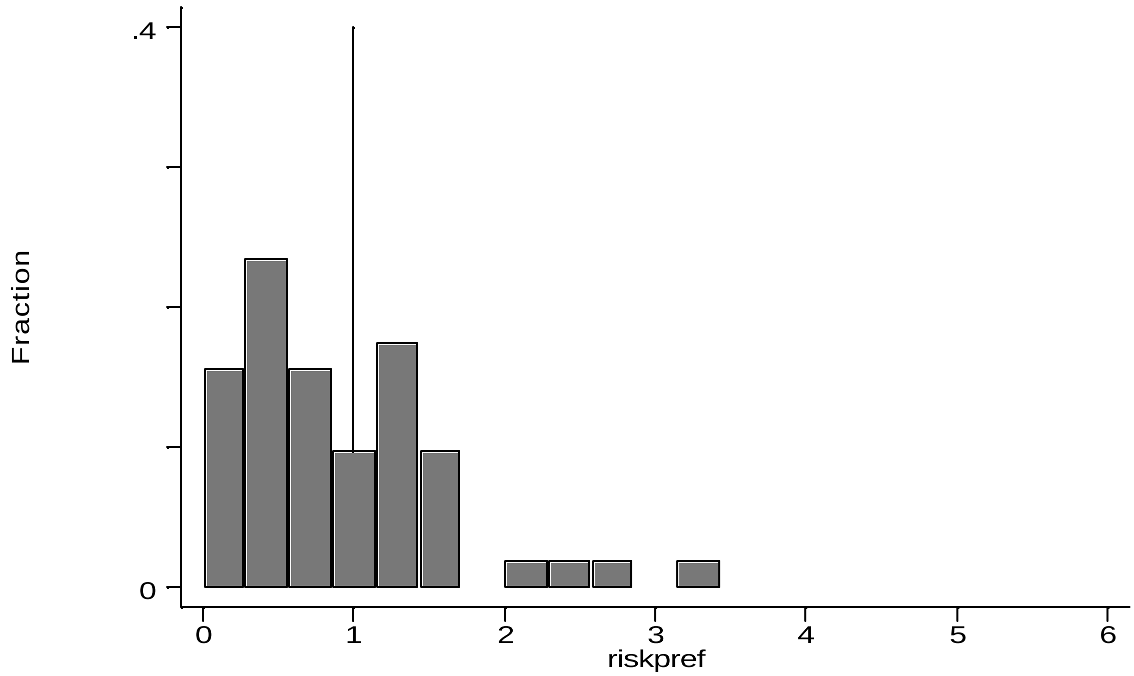

Figure 11b: Risk Aversion Parameter, Treatment II

APPENDIX : INSTRUCTIONS

The following are the instructions to the Ultimatum 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 or she played).

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

General Instructions For the Experiment

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 taken 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.

Specific Instructions For Ultimatum Game are Presented Below:

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 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 replacement 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 = $%e.

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 sender must decide how much of a given amount of francs, 100, to send to the receiver. The receiver, after receiving this offer from the sender, must then decide whether to accept or reject the offer. If the receiver accepts the offer, then the receiver gets a payoff equal to the offer and the sender gets a payoff equal to 100 minus the offer. If the receiver rejects the offer, then both the sender and the receiver will get a payoff of zero. For example, say the Sender chooses to offer the receiver x francs out of the hundred. Then, if the offer is accepted, the Sender=s payoff will be 100-x and the receiver=s will be x. If it is rejected both the Sender and the Receiver will get 0.

To make your decisions you will use a computer. If you are the sender there will be an Offer Box on your screen into which you can type your offer for the receiver. To do this simply use the mouse to click in the space provided for you in the offer box and type your offer. After you do so, the computer will ask you to confirm your choice by stating:

<<Are you sure you want to want to offer x francs? >>

If the answer is yes, click on the yes button and your decision will be entered in to the computer. If you would like to change the offer you have made, click on the No button . In this case you will be allowed to enter another amount.

If you are the receiver, the offer from the sender will appear on your screen in the Accept/Reject box. To accept the offer, simply click your mouse on the YES button.. If you wish to reject the offer, click the NO button. You will then be asked to confirm your choice and after you do your payoff will be reported to you in the payoff box.

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.

Predicitng Other people=s Actions

Before you make your decisions, you will be given an opportunity to earn additional money by predicting the choices of your opponent in the decision problem. To make a prediction click on the **Prediction** button. At this time, if your are a Sender you will see on the screen a prediction form as follows:

Sender Prediction Form here

Offer	Chance of Acceptance	Chance of Rejection = 100 - Chance of Acceptance
0-10		

11-20		
21-30		
31-40		
41-50		
51-60		
61-70		
71-80		
81-90		
91-100		

If you are a sender, you will be making make predictions about how likely the receiver is to accept or reject an offer from 0 to 10 francs, from 10 to 20 francs, etc. For example, suppose you are a sender and you think there is a 40% chance that the receiver will accept an offer between 0 and 10 francs if you happened to send one (and hence a 60% chance of it being rejected), that there is a 75% chance that the receiver will accept an offer between 11 and 20 francs, that there is a 80% chance that your opponent will accept any offer between 21 and 30, and that there is a 100% chance that your opponent will accept any offer from 31 to 100. If these are your beliefs about the likely decision of the receiver, then click in the space next to each possible range of offers that you might make, and type the numbers corresponding to your belief about the chances of an offer in that range being accepted. Note that when you type in the percentage chance of an offer in any given range being accepted, the corresponding percentage chance of the offer being rejected automatically appears in the next column. That is, the sum of the chances of an offer in any given range being accepted and the chances of such an offer being rejected must sum to 100.

We will then determine your prediction payoff as follows: Suppose you actually decide to send 27 francs and you predicted that your opponent will accept with a 60% chance and reject with a 40% chance. In that case you will place 60 next to the accept box of the 21-30 interval and 40 next to the reject box. Suppose now that your pair member actually rejects the 27 francs offered. In that case your payoff will be **Prediction Payoff = 20,000-$(100-40)^2$ - $(60)^2$** = 20,000 -3600 - 3600 = 12,800. 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 innacurate 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 a certain action is going to be taken with a 100% chance 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 you opponent will do for each amount sent to him/her. Any other prediction will decrease the amount you can expect to earn as a prediction payoff..

Receiver=s Predictions

If you are a Receiver to make a prediction click on the **Prediction** button. At this time, if you are a Receiver you will see on the screen a prediction form as follows:

Receiver Prediction Form

Offer	Chance of Receiving
0-10	
11-20	
21-30	
31-40	
41-50	
51-60	
61-70	
71-80	
81-90	
91-100	

As a Receiver you will be asked to make predictions about how likely the receiver is to make you an offer from 0 to 10 francs, from 11 to 20 francs, etc. For example, suppose you are a receiver and your beliefs are as shown in the following table:

Receiver Prediction Form

Offer	**Chance of Receiving**
0-10	30
11-20	5
21-30	5
31-40	5
41-50	5
51-60	5
61-70	5
71-80	5
81-90	5
91-100	30

If these are your beliefs about the likely decision of the sender, then click in the space next to each possible range of offers that might be made to you, and type the numbers corresponding to your belief about the chances of an offer in that range being sent to you. Note that the total of the chances you type in must sum to 100.

In addition say that the Sender actually sends you 27 francs. We will then determine your prediction payoff as follows: **Prediction Payoff = $20{,}000 - (100\text{-}5)^2 - (30)^2 - (5)^2 - (5)^2 - (5)^2 - (5)^2 - (5)^2 - (5)^2 - (5)^2 - (30)^2 = 20{,}000 - 9025 - 900 - 25 - 25 - 25 - 25 - 25 - 25 - 25 - 900 = 9000$**. 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 innacurate your prediction was. To do this when we find out what choice your pair member has made (i.e. in this case they sent you 27 which is between 21 and 30 francs), take the number you assigned to that choice, in this case 5, subtract it from 100 and square it. Then take the numbers you assigned to the offer ranges that did not contain the actual choice made by your pair member, square them, and add them up. . These squared numbers will then be subtracted from the 20,000 points we initially gave you to determine your final point payoff. Your point payoff will 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 a certain action is going to be taken with a 100% chance 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 you pair member will do for each amount sent to him/her. Any other prediction will decrease the amount you can expecte to earn as a prediction payoff.

Advice to 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. If you are a sender, then you can suggest the amount that your successor should send to his/her receiver. If you are a receiver, then you can suggest the minimum offer that your successor should accept. That is, any offer less than the Aminimum offer@ should be rejected by your successor, and any offer at or above the Aminimum offer@ should be accepted. You are also provided with a space where you can write any comments you have for them about the choice they should make. In addtion, in this space, if you like, you can tell your successor the advice given to you by your predecessor and the history of your predecessor which you saw and your successor did not see.

To give advice, click on the **Advice** button. You will then see on the screen the following advice form which provides you an opportunity to give advice to

Sender Advice form

Write Specific Advice Here
You Should Send _______ francs

Write General Advice Here

History Button

Receiver Advice Form

Write Specifi Advice Here
You Should not accept any offer less than _____ francs.

Write General Advice Here

History Button

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 the previous %k 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 %k predecessors participated in. You will not, however, see the advice given to the person you are paired with by his/her predecessor.

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 %k 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 you 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.