

Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach

Andrew Schotter* and Barry Sopher⁺

Revision September 28, 2001

Abstract

In the real world, when people play games, they often receive advice from those that have played it before them. Such advice can facilitate the creation of a convention of behavior. This paper studies the impact of advice on the behavior subjects who engage in a non-overlapping generational Ultimatum game where after a subject plays he is replaced by another subject to whom he can offer advice.

Our results document the fact that allowing advice has a dramatic impact on the behavior of subjects. It diminishes the variance of offers made over time, lowers their mean, and causes Receivers to reject low offers with higher probability. In addition, by reading the advice offered we conclude that arguments of fairness are rarely used to justify the offers of Senders but are relied upon to justify rejections by Receivers.

JEL Classification: C91, C72

Acknowledgement 1 * *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 Drew Fudenberg for his valuable comments and the time he took to make them. We would also like to thank Sangeeta Pratap, Mikhael Shor and Judy Goldberg for their valuable research assistance, and Yevgeniy Towshteyn for writing the program upon which the experiments were run.*

1 Introduction

In many of the decisions we make we rely on the advice of others who have preceded us. For example, before we buy a car, choose a dentist, choose a spouse, find a school for our children, etc. we usually ask the advice of others who have experience with such decisions. The same is true when we play games. In international affairs before a current president or prime minister makes a decision in an important international situation he or she asks the advice of those who went before. In industry, when a C.E.O. retires he or she passes on the wisdom of his or her years to his or her successor. In this manner the conventions of behavior that have been established in the past are preserved and passed on from generation to generation.

For this reason it might be important when examining behavior in a game like the Ultimatum Game to play that game in the lab in a manner that mimics how such a game might be played in the real world. More precisely, the Ultimatum Game can be interpreted as a contracting game in which an offer must be made by one person and either accepted or rejected by the other. (A share-cropping game between landlord and farmer may be an example). If this game is played repeatedly over time by a sequence of generations, then we might likely expect previous generations to pass on advice to their successors as to how to play and also pass on to them whatever conventions of behavior pertain to the game such as what offers they expect will be accepted and what type will be rejected.

Such an explanation for conventionally-determined behavior is offered in an article by Alvin Roth, Vesna Prasnikar, Masahiro Okuna-Fujiwara and Shmuel

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 difference in the behavior they observe is not the result of differences in the type of people inhabiting these countries (i.e. Israelis are not more aggressive than Americans by nature) as much as a cultural difference that has emerged in these countries which leads them to a different set of mutual expectations about what offers are acceptable; i.e., a different convention.

“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 offer under the circumstances” (Roth et al., (1991, p. 1092).

Burke and Young (2001a and 2001b) use a similar type of argument to explain the emergence of share-cropping conventions.

In this paper we use the apparatus of what Schotter and Sopher (2000) call intergenerational games to investigate the impact of advice giving in intergenerational Ultimatum Games. In these games a sequence of non-overlapping “generations” of players play a stage game for a finite number of periods and are then replaced by other agents who continue the game in their role for 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 payoff is a function not only of the payoffs achieved during their generation but also of the payoffs 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)).

What we expected when we started this experiment was that over time, in our intergenerational Ultimatum Game, one offer would emerge as the “conventional” laboratory offer meaning that Senders would repeatedly make this offer to Receivers, advise their successors to do so also, and hold beliefs that supported this offer as payoff maximizing. For this reason in our experiment we not only solicited advice but elicited (using a proper scoring rule) the beliefs of the Senders and Receivers concerning the behavior of their opponents in an effort to see if we could detect not only an equilibrium in actions but also an equilibrium in beliefs.

What we find contradicts our initial expectations in that the time series of offers made in our experiment fails to converge to one conventional offer. Rather, throughout the length of the history of our game the variance of offers is positive and non-negligible. What we do find, however, is that there is a clear and dramatic impact of advice on behavior in our experiment. More precisely, we find that while the strong form of our expectations were not borne out, a weaker form was substantiated in that the variance of offers made is significantly less in those treatments where advice is available. In this sense, behavior is more “conventional” when advice exists. Another way to put this is to say that behavior is more erratic in intergenerational games played with only history to guide behavior than when people either only get advice or are allowed to receive

both advice and see the history of play before them. History alone is a poor guide to behavior.

Another finding is that advice serves as the key variable explaining the offers sent by Senders. Further, not only is advice important but the type of advice given is also meaningful. For example, from examining the written advice offered from one generation of Sender to the next, we conclude that arguments of fairness or backward induction are infrequently relied on by subjects in rationalizing the offers they suggest to their successors. What is relied on are arguments of expected payoff maximization. In fact, even when 50-50 splits, the hallmark of equity offers, are proposed, they are mostly proposed because the Sender perceives the probability of having lesser offers accepted to be unacceptably low. The advice of Receivers is different, however, more often relying on fairness and spite arguments to justify behavior.

Finally, one of our most interesting observations deals with Receiver rejection behavior. What we find here is that the key element determining rejection behavior is the difference between the offer made and the expected offer of the Receivers. In other words, when an offer made is less than that expected by the Receiver (as defined by the Receiver's elicited beliefs), then it is almost always rejected even if that offer is above the stated minimally acceptable offer of the Receiver. Stated minimally-acceptable offers provide a poor guide to rejection behavior in that almost all offers above the minimally-acceptable ones are rejected if they are also below the offer the Receiver expects to receive. The distinction here is between the "hot" reaction that subjects seem to have when

offers arrive that are below their expectations and the “cool” response they give when hypothetically asked to state a minimum.

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

In the intergenerational Ultimatum Games reported on here subjects, once recruited, were ordered into generations in which each generation plays an Ultimatum Game once and only once against an opponent. After their participation in the game, subjects in any generation t are 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 are then required to give advice to their successors in the form of a strategy (a suggested amount to offer if they are a Sender) or a minimally acceptable offer (if they are a Receiver). They can also, if they wish, explain the reasons for their suggested strategy in the form of a free-form message. The payoffs to any subject are equal to the payoffs earned by that generation during their lifetime plus a discounted payoff which depends on the payoffs achieved by their immediate successors. Finally, during their participation in the game, subjects are asked to predict the actions taken by their opponent (using a mechanism which makes telling the truth a dominant

strategy).

The exact sequence of events is as follows. When a subject arrives at his or her terminal he or she receives written on-screen instructions. After reading the instructions and having any questions answered, they are shown the advice offered by their predecessor. This advice has two parts. A strategy which is a suggested amount to offer by the Sender and a suggested minimally acceptable offer by the Receiver and a free-form statement offering a justification for the proposed strategy. 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 offered and its acceptance or rejection. They could not see, however, any of the previous advice given to their predecessors.

After the advice was read, we elicited the beliefs of the Sender or Receiver using a proper scoring rule to be described in the Appendix. After the beliefs were elicited subjects played the Ultimatum Game and payoffs were determined. Their final payoff could only be determined after the next generation had finished, however, since their payoff depended on their actions.

Because running such intergenerational experiment can be very time consuming, given that we must do them sequentially, we decided on a slightly different design which allowed us to run three different intergenerational games in one session and hence generate observations of three different games.

More precisely, each experimental session started with the recruitment of 12

subjects who were randomly assigned to a group of six. During their experimental session each subject would play three different one-shot games with a different opponent. The three games were the Battle of the Sexes Game, the Ultimatum Game and the Game of Trust of Berg, Dickaut, and McCabe (1995). To play these games they were randomly matched with a different opponent in each period of the experiment. Hence, a subject might start the experimental session playing a one-shot Ultimatum Game with one of the other five subjects, then after that was over play a one-shot Battle of the Sexes Game with another subject, and finally the Trust Game with a third. In each game the subjects would read the instructions, see advice (or not, depending on the treatment) state their beliefs, and take an action. For example, an experimental session might consist of the following sequence of games played by our six subjects:

[Table 1 here]

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 once while Players 2 and 5 play the one-shot Ultimatum Game and Players 3 and 4 play the one-shot Trust game. When they have finished their respective games, we rotate them in the next period so that in period 2 Players 2 and 4 play the Battle of the Sexes Game while Players 3 and 6 play the Ultimatum Game and Players 1 and 5 play the Trust game. The same type of rotation is carried out in period 3 so that at the end of the experiment each subject has played each game once against a different opponent who has not played with any subject he has played with before. Each generation played the game once and only once and

their payoff was equal to the payoff they received during their generation plus an amount equal to $1/2$ of the payoff of their successor in the generation $t+1$ that followed them. (Payoffs were denominated in terms of experimental francs which were converted into U.S. dollars rates which varied according to the game played.) The design was common knowledge among the subjects except for the fact that the subjects did not know the precise rotation formula used. They did know they would face a different opponent in each period, however.

As a result of this design, when we were finished running one group of six subjects through the lab we generated three generations of data on each of our three games since, through rotation, each player played each game once and was therefore a member of some generation in each game. Thus for the set-up cost of one experiment we generated three generations worth of data on three different intergenerational 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.¹

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, they earned approximately an average of \$26.10 for about $1\frac{1}{2}$ hours. All instructions were presented on the computer

¹As far as we know, this is the record for economic experiments.

screens and questions were answered as they arose. (There were relatively few questions so it appeared that the subjects had no problems understanding the games being played which purposefully were quite simple). All subjects were inexperienced in this experiment.

In this paper we 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 fictitious 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 payoffs would be x for the Receiver and $100-x$ for the Sender. If the Receiver rejected the proposal, each subject's payoff would be zero.

2.1.1 Belief Elicitation

Our belief elicitation procedure (explained in more detail in the Appendix) worked as follows:

For the Receiver, we asked what they thought the probability was of receiving any amounts in the intervals 0-10, 11-20, 21-30, 31-40, 41-50, 51-60, 61-70, 71-80, 81-90, 91-100. 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 probabil-

ities defined above.² ³ Receivers were rewarded for their predictions using a quadratic scoring rule as described in the Appendix of this paper.

To elicit truthful beliefs from the Sender we do an equivalent procedure. The Sender is going to offer 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 offer in our ten intervals. More precisely, let us index the intervals by $k = 1, 2, \dots, 10$. 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 offer will be rejected. From this point on the payoffs are identical to the ones defined above but they are defined conditional on the amount sent. 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.⁴ 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

²In the instructions r_j is expressed as numbers in $[0,100]$, so we must divide by 100 to get probabilities.

³See Appendix 1 for the instructions concerning this part of the experiment.

⁴Obviously, there is no positive incentive to misrepresents beliefs in these intervals.

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 payoff in the game itself 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 offer 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 willingness to accept. By observing and coding the advice that is offered, 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.

2.2 Parameter Specification

The experiments can be characterized by four parameters. The first is the length of the history that each generation t player is allowed to see. The second is inter-generational discount rate indicating the fraction of the next generation’s payoff

to be added to any give generational players payoff. The third is the number of periods each generation lives for (i.e. the number of times they repeat the game) while the fourth indicates whether advice is allowable between generations. In all of our experiments each generation lives for one period or repeats the trust game only once and has a discount rate of $1/2$. Hence, they only differ on the basis of the length of history the subjects are allowed to view before playing and whether they are able to get advice from their predecessor or not. In the Baseline Ultimatum Game experiment subjects could pass advice to their successor and see the full history of all generations before them. This Baseline experiment was run for 81 generations. After we had run the Baseline experiment for 52 periods we started two separate and independent new treatments one which we call the Advice-only treatment (sometimes referred to as AO) and the History-only treatment (sometimes referred to as HO). In Advice-only Treatment before any generation made its move it could see only the last generation's history and nothing else. This treatment isolated the effect of advice on the play of the intergenerational game. The History-only Treatment was identical to the Baseline except for the fact that no generation was able to pass advice onto their successors. They could see the entire history, however, so that this treatment isolated the impact of history. Advice-only Treatment was run for an additional 80 generations while the History-only Treatment was run for an additional 66 generations, each starting after generation 52 was completed in the Baseline. Hence, our Baseline was of length 81, the Advice-only Treatment was of length

78⁵ and the History-only Treatment was 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.

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 offers of Sender subjects, the advice they were given, their beliefs, and the advice they offer their successors. We then proceed to look at the analogous behavior of Receivers.

3.2 Sender Behavior

3.2.1 Offers:

Observation 1: Advice Lowers the Variance of Offers.

Let M_B , M_{AO} , and M_{HO} be the mean offer in the Baseline, Advice-only Treatment and the History-only Treatment and let V_B , V_{AO} , and

⁵One generation was lost because of a computer crash. The lost generation was the third (last) period of a session. We were able to reconstruct the relevant data files

V_{HO} be their associated variances. Then $V_{AO} < V_B < V_{HO}$. In addition, in the Advice-only treatment, offers decrease as time progresses but this is not true in the Baseline or the History-only Treatment. Finally, if we look at the mean offers made during the last 40 generations, $M_{AO} < M_{HO} < M_B$.

Substantiation

What Observation 1 says is that the variance of offers is least in Advice-only Treatment, where only advice is present, and greatest in the History-only Treatment 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 offer behavior more systematically, consider Table 2 which present some descriptive statistics about the offer behavior of our subjects and Figures 2a-2c which presents a set of histograms of the offers in each experiment.

[Table 2 here]

[Figure 2 here]

There are some things to note. First note that by comparing the offers made in Advice-only treatment to those of the History-only treatment we see that one impact of advice is to truncate the right tail of the offer distribution. In fact while only 10% of the offers in the Baseline and Advice-only treatment were above 58 and 50 respectively, in the History-only treatment 10% of the observations were above 80. Note also that the distribution is much flatter in the History-only treatment and that there is much less of a spike at the modal choice than in either of the other treatments. In fact, the standard deviation of

offers is almost twice as great in the History-only treatment than in the Advice-only treatment 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 the History-only treatment and the Baseline ($F_{(65,80)} = 2.16$, $p = .00$) and the History-only treatment and the Advice-only treatment ($F_{(65,76)} = 2.90$, $p = .00$). The same test found a difference between the variances of the Advice-only treatment and the Baseline at only the 10% level. What this indicates is that history does not seem to supply a sufficient lesson for subjects to guide their behavior in a smooth and consistent manner. Advice seems to be needed. Finally note that the mode does shift lower when we move from the Baseline to the Advice-only treatment.

With respect to time, it appears that only in the Advice-only treatment do offers change over time in a statistically significant (and negative) manner. To illustrate this point we ran a simple OLS regression of offers made on time. In all regressions, except the one run on the Advice-only treatment data, time was insignificant at the 5% level. In the the Advice-only treatment regression, the

coefficient was negative and significant at the .003% level. ⁶

Looking at the mean offers in the last 40 generations we see that there is a statistically significant difference in the mean offers made between the Advice-only and the History-only treatments using a Wilcoxon test at the 2% level ($z = -2.295$, $p = .02$). No such difference exists in the comparison between the Baseline and the History-only treatment.

It appears then that the inclusion of advice leads subjects to conclude that sending lower offers is a beneficial 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 offers in Observation 1 suggests that we should investigate what factors are important in generating these offers. To pursue this question, we offer the following observation.

Observation 2: Advice Determines Offers

Advice is the key determinant in deciding upon offers. In fact, subjects tend to follow the advice of their generational predecessor

Regressions of Offer on Time:					
Baseline					
	Coef.	Std. Err	t	P> t	
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					
⁶	Coef.	Std. Err	t	P> t	
	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> t	
	time	.0710156	.1427262	0.498	0.620
	cons	40.07552	5.500362	7.286	0.000
F _{1,64} = .25, p = 0.62					

even when their own beliefs suggest that they would maximize their expected payoff by offering something else.

Substantiation

Before we present any statistical analysis to back up this observation, consider Figures 3a-3b which plot the times series of offers in each of our treatments involving advice against the advice the Sender received (in the Baseline and the Advice-only treatment) and also against their subjective payoff maximizing offer. By subjective payoff maximizing offer we mean that offer which, given the elicited beliefs of the subjects, would maximize their expected payoff if sent. Remember, for each potential offer in intervals 0-10, 11-20,....., 91-100, we have elicited the beliefs of the Sender subject as to the likelihood that such an offer would be accepted. Hence, we can take an expected value by assuming an offer at the midpoint of these intervals was sent and multiplying these offers by their elicited probabilities. This yields ten distinct values, each representing the expected payoff from sending an offer 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 3 here]

Note from Figures 3a and 3b the close fit between the advice that Senders receive from their predecessors and the offers they make. This is true for both the Baseline experiment and the Advice-only treatment. Note also, however, that despite the fact that our payoff-maximizing offer can only take on ten discrete values, they seem to fit the pattern of offers 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 payoff-maximizing offer. We ran this for both the Baseline and the Advice-only treatments. (Obviously the History-only treatment did not have advice). These results are presented in Table 3.

[Table 3 here]

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 the History-only treatment. Would, under these circumstances, subjects concentrate on their best response offer? Table 4 offers the answer to this question since it reports the results of a regression run on the History-only treatment data in which we regress the offer made simply on the subjects subjective payoff-maximizing offer.

[Table 4 here]

As Table 4 indicates, subjects do not appear to focus on their best response offers 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 offer subjects different access to history of the generations before them.

Observation 3: History Raises Advised Offers.

The advice given by subjects to their successors is greater in the Baseline than in the Advice-only treatment.

Substantiation:

We substantiate this observation by presenting Table 5 which simply presents the mean, median and variance of advice offered 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 here]

As we see, advice is lower in the Advice-only treatment and significantly so.

Observation 4: Pessimistic Beliefs.

Beliefs of Senders tend to be overly pessimistic, compared to what the history of the game implies their beliefs should be.

Substantiation

When we call beliefs overly pessimistic we mean the following. For each sub-interval 0-10, 11-20, 21-30, etc. we have elicited the belief of each of our generational Senders as to what they think the chances are that an offer 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 offers

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 here]

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).⁷ Hence, subjects seemed, on average, to greatly under estimate the willingness of their opponents to accept offers in this interval. The same pattern exists for all intervals and all treatments except in

⁷A 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 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.

the intervals 81-90 and 91-100 for the History-only treatment where the opposite is true. Note the small number of observations here, however.⁸

There are some further aspects of Table 6 worth noting. First, note that all mean Stated beliefs are monotonically increasing across intervals so that, on average, subjects did feel that higher offers had a higher probability of being accepted.⁹ 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 the Advice-only treatment) highest in the Advice-only treatment 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 have made them more optimistic.

If beliefs are too pessimistic then offers would tend to be too high in the sense that Senders could in actuality lower their offers and increase their expected payoffs. This raises the question as to whether a significant portion of the Ultimatum Game puzzle, that subjects do not send their sub-game perfect equilibrium offer and tend to make offers in the middle of the allowable range (around 50), is merely the result of misperceived probabilities. We are able to

⁸If we had room 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.

⁹In fact, 87% of subjects stated beliefs that were monotonically increasing in the offer amount. 6% stated beliefs that were "single-peaked," first increasing and then decreasing. Another 7% stated beliefs with at least 3 monotone segments. Restricting Table 6 to those with monotone stated beliefs would increase the average, especially in the 5 highest ranges of offers, but would not change any of the conclusions we draw below about the pessimism of stated beliefs.

suggest that this may be true because we have elicited the beliefs of our subjects and are in a position to know what offer was subjective payoff maximizing given Sender beliefs whereas such information was not available to previous investigators.

It should be pointed out, however, that for many offer ranges the number of observations is too small to be useful for comparison. For those situations where too little evidence exists, subjects are free to conjecture as they wish and one could make a case that some of the behavior we observe in this experiment might be consistent with the type of self-confirming equilibrium of the Fudenberg-Levine (1993) type where segments of the strategy space are left unexplored so that players are free to hold expectations about what would happen there that are unconstrained by experience.

3.2.2 Receivers:

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

Observation 5: Advice Makes Receivers Tougher

Defining a low offer as one below 25 and a "tough" Receiver as one who rejects low offers, the probability of having a low offer accepted is lowest in the Advice-only treatment, second lowest in the Baseline and highest in the History-only treatment. In other words, the bigger

the role allowed for advice (as in the Advice-only treatment where there is no history) the tougher are the Receivers.

Substantiation:

There are more conceptual difficulties 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 offer made and see if, when identical offers are made, they are rejected or accepted with identical frequency across treatments. Unfortunately, the set of offers 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 offered in each of our three treatments and comparing the resulting acceptance functions. We estimate the logistic relationship,

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

where x is the amount offered and the left hand variable is a $\{0,1\}$ 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 treatments 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.¹⁰

[Figure 4 here]

What we see in Figure 4 is that for low offers, 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 (the Advice-only treatment) and most likely to be accepted when no advice is present but access to history is unlimited (the History-only treatment). 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

Acceptance Behavior (Logit)				
Baseline				
Variable	coefficient (Std. Err.)	(z	P> z 	
accept				
sent	.10(.03)	3.62	0.00	
constant	-2.39 (1.07)	-2.24	0.03	
obs = 81				
Pseudo R² = .24	LL = -29.62			
Advice only				
Variable	coefficient (Std. Err.)	(z	P> z 	
accept				
sent	.16 (.04)	4.10	0.00	
constant	-4.20 (1.32)	-3.18	0.00	
obs = 77				
Pseudo R² = .41	LL = 24.71			
History only				
Variable	coefficient (Std. Err.)	(z	P> z 	
accept				
sent	.022 (.01)	1.52	0.13	
constant	-.048 (.61)	-0.08	0.94	
obs = 66				
Pseudo R² = .03	LL = -39.16			

10

function of their stated minimum acceptable offer, their expected offer given their stated beliefs, the advice they received from their predecessors (in the Baseline and the Advice-only treatment), the offer they received and appropriate differences among these variables. What we find is summarized in Observation 6 and substantiated below:

Observation 6: Unfulfilled Expectations Cause Rejection

The key variable influencing rejection behavior appears to be whether or not the offer received by the Receiver is above or below that which he expects to receive and not the relationship between the offer and the stated minimal acceptable offer of the receiver.

Substantiation:

In our experiment we have elicited a great deal of information about Receivers which can be of great help in describing rejection behavior. For example, we know what they stated as their ex ante minimum acceptable offer, and we can calculate the offer 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 offer 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 offer.

Table 7 describes the rejection and acceptance behavior of subjects on the basis of the difference between the offer they receive and either their minimal acceptable, expected, or advised acceptable offer. Note that if any one of these three variables explains either acceptance or rejection behavior it must be such

that whenever the offer exceeds any one of them it is accepted while when it is below it is rejected. For example, if expectations matter for behavior, then we would expect any offer below a subject's expectations would be rejected while any offer above would be accepted. Table 7 presents our results:

[Table 7 here]

A number of things are notable in this table. First, the difference between what a Receiver was offered and what they expected to receive is very good at correctly classifying rejections, but is bad at classifying acceptances. For example, of the 15 rejections in the Baseline experiment, 14 occurred when the Receiver was not offered at least his expected amount. However, of the 66 acceptances in the Baseline, 33 occurred in instances where the amount offered was less than a Receiver's expectations. Similar patterns exist in the other treatments as well.

The difference between a Sender's offer and a Receiver's stated minimum acceptable offer has just the opposite effect; very good at classifying acceptances but bad at classifying rejections. For example, in the Baseline again, of the 66 acceptances 62 occurred when the offer was greater than the stated minimum acceptable. (It is not surprising that the result here is stronger than that for the expected offer since it is almost always the case that a Receiver's expected offer is greater than his or her stated minimum acceptable offer). However, of the 15 rejections in the Baseline, 11 occurred when the offer received was greater than the stated minimum. This seems to imply that rejection behavior is a "hot" phenomenon perhaps triggered for some subjects by a deflation of

expectations, while stating a minimal acceptable offer is more a more detached “cold” phenomenon.(See, Brandts J. and Charness, G., (2000)).

The difference between the offer 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. Overall, however, our analysis of Receiver behavior shows a less dramatic role for advice than was true for Senders.

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 and the results we reach are summarized by the following observation.

Observation 7: Sender and Receiver Advice Differ

While the advice of Senders appears to be own payoff oriented and infrequently mentions fairness, Receiver advice reflects a more inter-dependent utility orientation.

Substantiation.

In rationalizing advice in our experiment, a subject might appeal to a number

of different 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 payoff maximizing grounds if one thought that, given your subjective acceptance probabilities, such an offer is a best response. Such a rationalization need not appeal to equity at all. Alternatively, one may support offering 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 offer to make because one thinks that there is a threshold below which one's opponent will reject any offer but above which the offer 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 payoff was also affected 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 6 sub-groups: Best response Advice (BRA) which basically supports an offer on the basis of expected payoff maximization, Backward induction advice (BI), which is the type of advice consistent with subgame perfection in that it posits that the Receiver will accept any offer and then advises the Sender to send as little as possible given that expectation, 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.

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 payoff maximization, we counted all of them in our coding. Our point was not to define 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 different categories. For example, consider the following advice written by the Sender in generation 46 of the Advice-only treatment which includes elements of many different 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 payoff. I was greedy and offered 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, BI, FA, and HA since it included elements of all of these.

For the Receiver we proceeded as described above except that we changed the categories slightly given the differing roles of the subjects. We retained the codings 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 payoffs were important and also BI+ which is basically advice that says accept anything above a strictly positive threshold. Spite and fairness are very close to each other so we have merged them in the table below, but we point out that spite has a much more mean-spirited objective.

A spite statement might read as did this one representing subject 45 in the the Advice-only treatment who suggested a minimum acceptable offer 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.”

Examples of a pure Backward Induction advice (BI) were seen in the advice given by the Receivers in generations 34 and 35 of the Advice-only treatment who all told their successors to accept anything above 1 if it is offered with the following explanations: “accept any offer that is offered to you because to reject means that you get nothing. (Generation 34), “Definitely accept anything, or else you get nothing”. (Generation 35).

Finally, we added a category PR for prescription which refers to statement that simply suggested a cut-off point without any real justification. (“Don’t take less than 40 – subject 47 of the the Advice-only treatment). 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 8 which present the results of our coding for the Baseline and the Advice-only treatment.

[Table 8 here]

One of the most striking features of Table 8 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 the Advice-only treatment). More interesting, however, is that fact that when 50-50 splits are suggested, they are most often supported by payoff 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 the Advice-only treatment, 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 offer proof of equity considerations.

Also notable in Table 8 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 the Advice-only treatment. The overwhelming bulk of advice had Senders suggesting an offer to their successor which, given their assessment of the probabilities of rejection, either maximized their expected payoff or constituted a best offer given their assessment of the minimum acceptable offer on the parts of Receivers. For example, there were 38 such pieces of advice in the Baseline and 21 in the Advice-only treatment. 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 different. 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 offers while in the Advice-only treatment 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 the Advice-only treatment. 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.

Substantiation.

Another interesting feature of the advice texts we read was the fact that in the Advice-only treatment, 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 the Advice-only treatment subjects could not flip 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 offer 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 the Advice-only treatment 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 Conclusions

This paper has studied the impact of advice in intergenerational Ultimatum Games. What our results demonstrate is the overwhelming influence 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 offers when it is present. Hence, games played with advice generate behavior which is more “conventional” than those where advice is absent. Advice is also important for Receivers affecting both their rejection and acceptance behavior. However, for Receivers it appears as if rejection behavior is most affected by a deflation of their expectations since most rejections occur when they receive an offer that was lower than what they were expecting even if that offer is above their stated

minimal accepted offer.

5 Appendix

Elicitation procedure for Receivers:

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 0-10, 11-20, 21-30, 31-40, 41-50, 51-60, 61-70, 71-80, 81-90, 91-100. Since only one such amount will actually be sent, the payoff to player i (the Receiver) when an amount in interval l is chosen will be:

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

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

Note what this function says. A subject starts out with 20,000 points and states a belief vector $r = (r_1, r_2, r_3, r_4, r_5, r_6, r_7, r_8, r_9, r_{10})$. If their opponent chooses to send an amount in interval l , then the subject would have been best off 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 payoff of 0. It can easily be demonstrated that this reward function provides an incentive for subjects to reveal their true beliefs about the actions of their opponents.¹¹ Telling the truth is optimal.

Elicitation Procedure for the Sender:

As indicated above the Sender types ten probability vectors into the computer of the following form: $r_k = (\pi_a^k, \pi_r^k)$, where k is the index of one of the 10 intervals between 0 and 100. Hence, 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 offer will be rejected. From this point on the payoffs are determined by a quadratic scoring rule. For example, say that an amount in the k^{th} interval was sent, the Sender predicted that if he or she sent that amount

¹¹An 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.

it would be accepted with probability π_a^k , and it turns out that the offer was accepted. Then that Sender's prediction payoff would be defined as follows:

$$\Pi_k = 20,000 - \{((100 - \pi_a^k)^2 + (\pi_r^k)^2)\}. \quad (2)$$

In other words, if the offer was accepted but the Sender only predicted that it would be accepted with probability π_a^k , the payoff 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 offer would be rejected which it was not. An analogous payoff can be defined if the offer was rejected.

6 Bibliography

References

- [1] Berg, J., Dickhaut, J. and McCabe, K. (1995), "Trust, Reciprocity, and Social History.," *Games & Economic Behavior*, 10(1) 122-142
- [2] Brandts., J., and Charness, G., (2000, forthcoming) , "Hot vs Cold: Sequential Responses and Preference Stability in Simple Experimental Games", *Experimental Economics*.
- [3] Burke, Mary A. and Young, H.P., (2001a), "Competition and Custom in Economic Contracts: A Case Study in Illinois Agriculture," *American Economic Review*, 91, 551-573.
- [4] Burke, Mary A. and Young, H.P., (2001b), "Contract Uniformity and Factor Returns in Agriculture," May 2001.
- [5] Fudenberg, D. and Levine, D. (1993) "Self-Confirming Equilibria" , *Econometrica* 61, 1993, pp. 523-546.
- [6] Nyarko, Y. and Schotter, A., (1999), "An Experimental Study of Belief Learning Using Real Beliefs" Economic Research Report 98-39, C. V. Starr Center for Applied Economics, New York University.
- [7] 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.
- [8] Schotter, A. (1981), *The Economic Theory of Social Institutions*, Cambridge University Press, Cambridge England.
- [9] Schotter, A. and Sopher, B., (2000), "Social Learning and Coordination Conventions in Intergenerational Games: An Experiment in Lamarckian Evolutionary Dynamics" Russell Sage Working Paper #155, Russell Sage Foundation, New York, New York.

Table 1: Rotation Scheme For Subjects

Period		Game		
		Battle of Sexes	Ultimatum	Trust
Period 1	Subjects 1	1	2	3
	Subjects 6	6	5	4
Period 2	Subjects 2	2	3	1
	Subjects 4	4	6	5
Period 3	Subjects 3	3	1	2
	Subjects 5	5	4	6

Table 2: Offers by Senders

Treatment	All generations		Last 40 generations	
	Mean	s.d.	Mean	s.d.
Baseline	44.70	14.95	45.66	15.95
Advice only	37.16	12.89	33.68	13.53
History only	42.45	21.96	43.90	19.66

Table 3: Offer Behavior in the Baseline and the Advice-only treatment

	Baseline			
	coeff.	s.e.	t	P> t
Payoff-max. offer	.11	.11	1.05	0.30
Advice sent	.26	.10	2.62	0.01
constant	27.42	6.64	4.13	0.00
R ² = .11	F(2,77) = 4.64, Prob>F=.01			
N = 81				

	Advice Only			
	coeff.	s.e.	t	P> t
Payoff-max. offer	.08	.07	1.08	0.29
Advice sent	.53	.10	5.11	0.00
constant	13.05	5.63	2.32	0.02
R ² = .27	F(2,73) = 13.22, Prob>F = .00			
N = 77				

Table 4: Offer Behavior in the History-only treatment

	coeff	std. er.	t	P> t
	Payoff-max. offer	.16	.15	1.08
constant	34.56	7.78	4.44	0.00
R ² = .02	F(1,64) = 1.17, Prob>F = .28			
N = 66				

Table 5: Advice in the Baseline and the Advice-only treatment

Treatment	Mean	Median	Variance	Std. Dev.
Baseline	44.48	47	270.97	16.46
Advice-only	38.25	40	158.71	12.59

Wilcoxon Test: $z = 3.12$, $p = .00$

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
	Advice Only 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
	History Only 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

Table 7: Rejection and Acceptance Behavior
Variable: Offer-Expected Offer

	Treatment		
Prediction	B	AO	HO
Acceptance: Offer \geq Expectation	33/66	29/59	19/46
Rejection: Offer $<$ Expectation	14/15	17/18	15/20

Variable: Offer-Minimum Acceptable Offer

	Treatment		
Prediction	B	AO	HO
Acceptance: Offer \geq Minimum	62/66	59/59	43/46
Rejection: Offer $<$ Minimum	4/15	3/18	4/20

Variable: Offer - Advice


	Treatment		
Prediction	B	AO	HO
Acceptance: Offer \geq Advice	50/66	41/59	NA
Rejection: Offer $<$ Advice	10/15	13/18	NA

Table 8: Coded Advice

Experiment	Senders					
	Type of Advice					
	BRA	BI	FA	HA	OA	PA
Baseline	38	4	8	5	10	19
Advice-only	21	6	11	23	6	7

	Receivers						
	PR	BI	BI+	$\frac{FA}{Spite}$	HA	OA	PA
Baseline	7	11	3	29	8	11	13
Advice-only	7	10	6	12	5	3	13

Figure 1. Experimental Design



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

Figure 1:




Figure 2: Offers

Figure 2:




Figure 3: Advice, Offer and Payoff Maximizing Offer

Figure 3:




Figure 4: Estimated Probability of Acceptance, by Treatment

Figure 4: