Available online at www.sciencedirect.com
ScienceDirect
Games and Economic Behavior 58 (2007) 365–393
GAMES and Economic Behavior
www.elsevier.com/locate/geb

# Advice and behavior in intergenerational ultimatum games: An experimental approach ☆

Andrew Schotter [a], Barry Sopher [b,*]

[a] *New York University*
[b] *Department of Economics, Rutgers University, 75 Hamilton Street, New Brunswick, NJ 08901, USA*

Received 22 September 2003
Available online 4 May 2006

**Abstract**

In the real world, when people play a game, 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 fosters the creation of a convention of behavior in Ultimatum games. 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

## 1. Introduction

In this paper we are interested in conventional behavior. What we want to know is what is conventional behavior, how does it come into existence, how is it passed on from one generation of social agent to the next and what are its properties? To investigate this phenomenon experimentally we use the Ultimatum game.

It is our claim here that one essential ingredient in the convention formation process is advice. Children learn to behave conventionally, to say please and thank you, to not talk when others

☆ This work was completed under N.S.F. grants SBR-9709962 and SBR-9709079.
* Corresponding author.
*E-mail address:* sopher@econ.rutgers.edu (B. Sopher).

0899-8256/$ – see front matter 
doi:10.1016/j.geb.2006.03.005

are talking, to open doors for each other etc., through a process of socialization communicated by parental advice. In our daily lives we learn that 15% is the right tipping rate because we are taught so by others while in business oligopolists learn not to engage in price wars because previous CEOs have learned that such behavior is dangerous and tell their successors to avoid it. In other words, conventions are part of the social learning we do as individuals. The thing that sets conventional behavior apart from other behavior is that it is automatic, non-thinking and socialized. Once we learn a convention we do it without thinking. In fact the reason the convention exists in the first place is to save us the cognitive load of solving repeated problems anew each time they arise.

Without the ability to obtain advice the only information available for us to learn from would be history. But in previous studies by these authors and others (see Schotter and Sopher, 2004, 2006; Iyengar and Schotter, 2005; Celen et al., 2004), it has been found that viewing previous history is a rather poor guide for current behavior. There are a number of reasons for this. First, people may not be especially able to digest and process a long time series. For example, giving an investor the quarterly time series of interest rates over the post World War II era may not help him invest in the bond market as well as a word from his neighbor or stock broker. We are trained from birth to listen to advice, not to interpret time series. Second, the process of advice giving and receiving forces us to think at a high cognitive level. Hence, advice may be worth listening to (see Iyengar and Schotter, 2005) because it is well reasoned and serious especially if the advice giver's interests are aligned with yours. For all of these reasons we expect to observe more conventional behavior when advice is present in a situation than when people have to rely on only history to guide their behavior. This logic furnishes the main hypothesis of this paper.

Bargaining is also a rather conventional endeavor. For example, in real estate negotiations it is common for the buyers to offer less than the asking price for the house they are interested in. Offering too little, however, can lead to the perception that the buyer is not serious and ruin the deal. Offering too much is wasteful, at least to the buyer. Depending on the community a convention exists that dictates which offer is acceptable and which is frivolous and real estate brokers are typically the repository of such wisdom.

The Ultimatum Game is an extreme type of bargaining game in which an offer is made by one person and either accepted or rejected by another. 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.[1]

Such an explanation for conventionally-determined behavior is offered in an article by Roth et al. (1991) who compare the behavior of subjects engaged in Ultimatum Games in four different countries: the United States, Japan, Israel, and Yugoslavia. At the end of their paper they conclude that the differences in the behavior they observe are 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.

[1] For example, Burke and Young (2001a, 2001b) investigate the creation of conventions determining the terms of sharecropping contracts in Southern Illinois and show that such terms appear to be "conventionally" determined in the sense that they are homogeneous within certain spatial areas but differ across space in a manner that can not be explained by economic fundamentals. In this paper we aim to present the results of a series of experiments whose purpose is to investigate how such conventions get established and passed on from generation to generation.

> "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).

Stated differently, what Roth et al. mean here is that while the Ultimatum Game may have only one subgame-perfect Nash equilibrium, it has many other Nash equilibria, each one of which defines an offer which is minimally acceptable to Receivers. Each one of these Nash equilibria could function as a convention of behavior dictating what offer should be sent and accepted. All that is needed to maintain a non-subgame perfect equilibrium is a set of beliefs on the part of the Sender that the offer contemplated is the minimum that would be accepted and behavior on the part of the Receivers that confirms 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 what Robert Aumann (see, van Damme, 1998) has called 'a "perfectly good" Nash equilibria that just isn't perfect'.[2]

In this paper we use the apparatus of what Schotter and Sopher (2004) call inter-generational games to investigate the impact of advice giving and convention creation in inter-generational 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, 2004) and games of trust (see Schotter and Sopher, 2006).

What we find is that behavior is more "conventional" in inter-generational Ultimatum games played with advice than without it. Another way to put this is to say that behavior is more erratic in inter-generational 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. In addition, the offers associated with these conventions are lower when advice exists than when it does not so not only does advice foster the creation of conventions but it lowers the offers associated with them.

In addition to observing the impact of advice on convention creation, our experiments turn up a number of other interesting findings. For example, we find 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.

Finally, using the written texts of the advice sent between generations we conclude that arguments of fairness or backward induction are infrequently relied on by subjects in rationalizing

[2] Comment by Robert Aumann.

the offers they suggest to their successors. What are 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.

We will proceed as follows. Section 2 describes our experiment and experimental design. In Section 3 we will present some definitions of what it means to behave in a conventional manner as well as a set of hypotheses to be tested. Section 4 reports the results of our hypothesis tests while Section 5 provides some additional results concerning the behavior of our subjects. Section 6 presents some conclusions.

## 2. The experiment: Design and procedures

### 2.1. General features

The general features of our inter-generational Ultimatum Game were as follows: Subjects once recruited were ordered into generations. Each generation played the game once and only once with an opponent. After their participation in the game, subjects in any generation $t$ are replaced by a next generation, $t+1$, who were able to view some or all of the history of what had transpired before them. Subjects in generation $t$ were able to give advice to their successors by suggesting an amount to offer or a minimum acceptable amount depending upon whether they were Senders or Receivers.

The payoffs to any subject in the experiment were equal to the payoffs earned during his or her lifetime plus a discounted payoff which depended on the payoffs achieved by his or her immediate successor in the game. Finally, during their participation in the game, subjects were asked to predict the actions taken by their opponent (using a mechanism which makes telling the truth a dominant strategy). This was done in an effort to gain insight into the beliefs existing at any time during the evolution of our experimental society. Each subject thus also received a payoff based on the accuracy of his or her beliefs, in light of what the subject's opponent actually did.

The experiment was run at both the Experimental Economics Laboratory of the C.V. Starr Center for Applied Economics at New York University or at the Experimental Lab in the Department of Economics at Rutgers University. Subjects were recruited, typically in groups of 12, from undergraduate economics courses and divided into two groups of six with which they stayed for the entire experimental session. During their time in the lab, for which they earned approximately an average of \$26.10 for about $1\frac{1}{2}$ hours, they engaged in three separate inter-generational games, a Battle of the Sexes Game, an Ultimatum Game in which they were asked to divide 100 francs, and a Trust Game as defined by Berg et al. (1995). All instructions were presented on the computer screens and questions were answered as they arose.[3] 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 the role of Sender or Receiver. The Sender was initially allocated 100 units of a laboratory currency called francs, which were later converted into

[3] There were relatively few questions so it appeared that the subjects had no problems understanding the games being played which purposefully were quite simple.

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.

The exact sequence of events was as follows. When a subject arrived at his or her terminal he or she received written on-screen instructions. After reading the instructions and having any questions answered, they were shown the advice offered by their predecessor. This advice had two parts: (i) A recommended action, which was either a suggested amount to offer by the Sender or a suggested minimally acceptable offer by the Receiver, and (ii) a free-form statement offering a justification for the proposed action. 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 generation (or generations, depending on the treatment) before it and see what the previous Sender (or Senders) offered and whether this (or these) offer(s) was (were) accepted or rejected. Subjects could not see, however, any of the advice given to their predecessors.

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

*2.1.1. Belief elicitation*

Our belief elicitation procedure worked as follows:

For the Receiver, we asked, in essence, 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. Specifically, 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 defined above.[4,5] Receivers were rewarded using a quadratic scoring rule described in Appendix A.

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 the ten intervals. More precisely, if we index the intervals by $k = 1, 2, \ldots, 10$, then the Sender would type ten probability vectors into the computer of the following form: $r_k = (\pi_a^k, \pi_r^k)$. Here $\pi_a^k$ is the probability that if an amount in the $k$th interval is sent it will be accepted while $\pi_r^k$ is the complementary probability that the offer will be rejected. Senders are rewarded using a quadratic scoring rule described in Appendix A, conditional on the amount actually sent. Note that since the Sender knows how much he or she will send before he states his beliefs, 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, i.e., the scoring function is incentive compatible only for the beliefs in the interval of the 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.[6] As you will see, however, none of our more important claims rely on this information.

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

[5] See Appendix A for the instructions concerning this part of the experiment.

[6] Still there seems to be no positive incentive to misrepresents beliefs in these intervals either.

We made sure that the amount of money that could potentially be earned in the belief elicitation 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 belief elicitation part of the Ultimatum Game was only $2.00 as opposed to the maximum payoff in the game itself of $10.00.) Our fear 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.

This scoring rule method has been used in other studies. For example, Nyarko and Schotter (2002) used it and found that it generates beliefs consistent with learning behavior in belief-learning models. Schotter and Sopher (2004, 2006) have used it to measure beliefs about opponent actions in other games. Other methods exist to elicit beliefs, for example, see Charness and Dufwenberg (2003) and Dufwenberg and Gneezy (2000). Some researchers argue that simpler, though technically not incentive compatible, procedures may be reliable. We prefer to use an incentive compatible method where possible since, in fact, there is no way to check whether reported beliefs are true beliefs. That is, one really only has the properties of the procedure as assurance that the elicited beliefs will be true beliefs. Note that we are not dogmatic on this point, as our willingness to use some of the sender beliefs (footnote 6) attests.

It is interesting to note that our experiment provides a whole host of data and information that has not been collected 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 Ultimatum 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 the are able to get advice from their predecessor or not. In the History-and-Advice treatment subjects could pass advice to their successor and see the full history of all generations before them. (They could not see any advice except for their predecessors.) This History-and-Advice treatment was run for 81 generations. After we had run the History-and-Advice treatment 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 the Advice-Only treatment, before any generation made its move, it could see only the last generation's actions and receive advice from them. So in this treatment people could get advice from the previous generation but could not see the history of any other previous generation nor get advice from any previous generation. (In any treatment with advice, subjects could only see the advice of

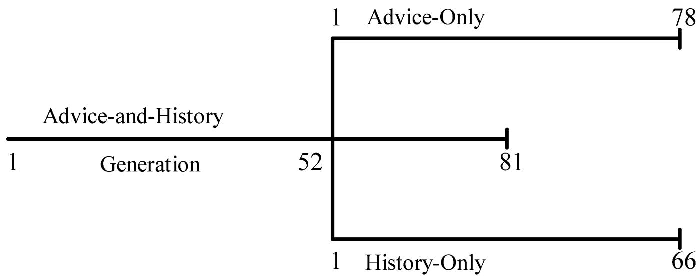


Fig. 1. Experimental design.

their immediate predecessor.) This treatment isolated the effect of advice on the play of the inter-generational game. We allow people to see the actions of their predecessor in the Advice-Only treatment since it is realistic to think that if you are around to get advice from someone you were likely to see how they behaved. The History-Only treatment was identical to the History-and-Advice treatment except for the fact that no generation was able to pass advice on to their successors. They could see the entire history, however, so that this treatment isolated the impact of history. The Advice-Only treatment was run for an additional 78 generations while the History-Only treatment was run for an additional 66 generations, each starting after generation 52 was completed in the History-and-Advice treatment. Hence, our History-and-Advice treatment was of length 81, the Advice-Only treatment was of length 78[7] and the History-Only treatment was of length 66. Our experimental design can be represented by Fig. 1.

## 3. Definitions and hypotheses

### *3.1. Definitions*

To answer the questions we are interested in this paper we need to be able to define what we mean when we say that behavior in society $A$ is more conventional than that of society $B$ when they each face the same situation or game. Put differently, say that you knew that two identical societies played the same game for an identical number of periods and you were presented with the time series of their interactions. What features of this data would lead you to say that behavior in society $A$ was more conventional than that of society $B$?

In order to answer these questions, we will need to define a number of concepts starting with that of a Strong Convention.

**Definition 1** *(A Strong Convention).* In an inter-generational Ultimatum game a "strong" convention of behavior consists of an offer $x^*$ (the "conventional offer"), a set of beliefs for the Senders and Receivers, and a pattern of advice such that:

(1) All Senders
  (a) expect offers less than $x^*$ to be rejected with probability 1,

[7] One generation was lost because of a computer crash. The lost generation was the third (last) period of a session. We were able to reconstruct the relevant data files, so the first generation of the next session received advice from the second generation of the previous session.

   (b) expect all offers at or above $x^*$ to be accepted with probability 1,
   (c) send $x^*$, and
   (d) advise their successors who play the game after them to send $x^*$;
(2) All Receivers
   (a) expect $x^*$ to be sent,
   (b) reject all offers below $x^*$,
   (c) accept all offers at or above $x^*$, and
   (d) advise their successors to do the same.

While such a strong convention may get established, it is more likely that behavior, while conventional, may not be so neatly so. Still, in this paper we will judge the conventionality of behavior by the extent to which it deviates from this strict or strong definition. Our claim in this paper is that behavior is more conventional in experiments with advice because the behavior we observe is closer to that predicted by the definition above in those experiments where advice is present.

To be more precise, let us assume that we performed an Ultimatum Game experiment and the results were in exact conformity to the definition above, i.e., our subjects established a strong convention. If this were so the data generated would have a very simple characteristic; all relevant variables in a given treatment would fit the same 0–1 step function. For example, consider the functions plotted in the top half of Fig. 2, with amounts of money on the horizontal axis and probability of acceptance on the vertical. These are the estimated logit acceptance function for each of our experimental treatments indicating the probability that a given offer is accepted by the Receivers.

If a strong convention existed, then, with respect to the amount of money sent, there would be some conventional amount, $x^*$, all of our subjects would have sent $x^*$ and all offers less than $x^*$ would be accepted with probability zero while all above would be accepted with probability 1. (In other words, the acceptance function would be a step function.) When looking at other variables, such as the senders' beliefs about what offer would be acceptable to the receivers, if a strong convention existed we would expect that all of our subjects would place all of their probability mass on the outcome that $x^*$ or greater would be acceptable. That is, the empirical distribution of such beliefs would define the same step function with the same break as above—the belief that an offer below $x^*$ would be acceptable would be zero while it would be equal to 1 for all offers greater than $x^*$. In fact, if a strong convention existed, the same 0–1 step function would arise for all of our relevant behavioral variables, i.e., amounts sent, Sender beliefs over acceptance, Receiver beliefs over amounts sent, as well as the advice offered to Sender and Receivers.

Hence, when a strong convention exists all of the data collected on all variables of interest (amounts sent, amounts accepted, beliefs, and advice) have what we will call the "step-function" property which is the defining characteristic of strongly conventional behavior. These considerations help us define the following concept.

**Definition 2** *(More Conventional Than).* The data on a particular behavioral variable in treatment $A$ is *more conventional than* the corresponding data in treatment $B$ if the empirical distribution function is closer to the 0–1 step function characterizing a strong convention.

To illustrate what we mean here consider the other line (not the step function) plotted in each graph in the upper half of Fig. 2. As stated above these are the estimated acceptance functions. The step function super-imposed in each plot is that 0–1 step function that best matches the con-

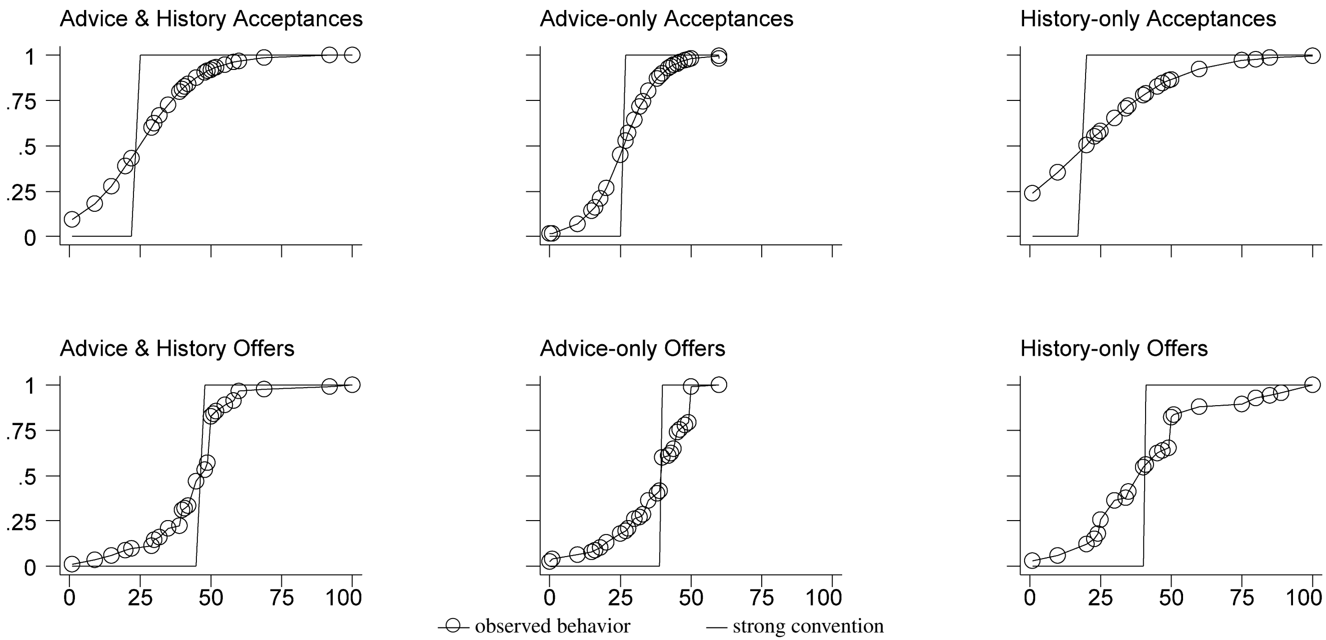


Fig. 2. Probability of acceptance and offers.

tinuous empirical function, in the sense of minimizing the area between the empirical function and the step function. It can be thought of as the strong convention that most closely approximates actual behavior in the treatment. Note that in each figure there is, in general, a different break point, $x^*$, in the step function, implying that different conventions were established in each treatment.

To measure whether behavior is more conventional in one treatment than the other we will measure the area between the estimated acceptance function and its associated 0–1 step function. If this area is smaller in one treatment than another, we call behavior in that treatment "more conventional" than that in another.[8] For example, as can be seen by comparing the acceptance functions of the Advice-Only and History-Only treatments, the acceptance function more closely approximates the idealized 0–1 step function in the Advice-Only treatment. The difference between the step and estimated functions is measured by the sum of areas between the empirical function and the step function in each diagram. That sum is 8.58 in the Advice-Only treatment while it is 18.13 (or more than twice as large) in the History-Only treatment. Further, as can be seen by the shape of the acceptance function for the Advice-Only treatment in Fig. 2, there is a much sharper change in the probability of acceptance on either side of $x^*$. In other words, it is much more clear cut in the Advice-Only treatment what an acceptable offer is since the acceptance function is steeper than in the History-Only treatment. Further, low offers, have a minimal probability of acceptance in the Advice-Only experiment while they remain viable in the History-Only treatment.

Note that a similar result holds when we compare the cumulative frequency of the amounts sent in the Advice-Only and History-Only treatments in the lower half of Fig. 2. Here our area metric again indicates that the offer function is more closely step-like in the Advice-Only treatment, where the area metric equals 8.12, than in the History-Only treatment, where the area metric equals 16.12. So for our two key variables, amount sent and probability of acceptance, it is rather clear that behavior was more conventional in our Advice-Only treatment than in our History-Only treatment. This result holds in general for the other variables (see Figs. 3 and 4). In Section 4, where we discuss our results, we will use this area metric as our measure of conventionality to provide a formal test between the treatments.

### *3.2. Hypotheses*

When a strong convention exists it specifies what actions we should expect people will take, what their expectations should be about the actions of others, and what advice they should offer their successors. The definitions above help us to frame a set of hypotheses which all concern the impact of advice on our key behavioral variables. As we indicated in the Introduction to this paper, we expect that advice is an essential ingredient for effective convention formation. People seem to have more trouble interpreting a historical time series than assimilating advice and either following it entirely or modifying their actions because of it. These assertions are borne out in our previous experience with inter-generational games (see Schotter and Sopher, 2004, 2006) where it is clear that advice fosters convention creation. As a result we will investigate the following hypothesis.

[8] More precisely, we choose the one-step function, increasing from 0 to 1 at the step, to minimize $\sum |f_i - s_i|$, where the sum is over the discrete support of the offer distribution, i.e., from 0 to 100, and where $f_i$ and $s_i$ are the empirical and step function values, respectively, at offer $= i$. It is easy to see that the step must occur at a point where the empirical function takes on the value 0.5.

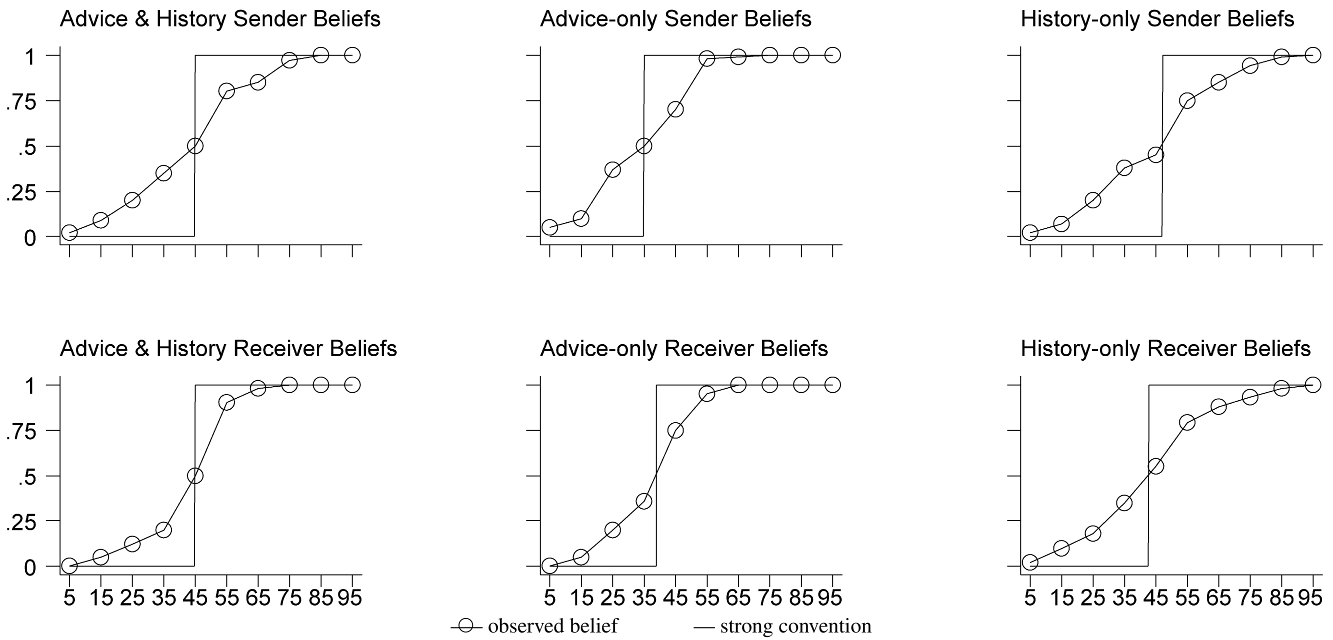


Fig. 3. Beliefs: amount sent and acceptable offers.

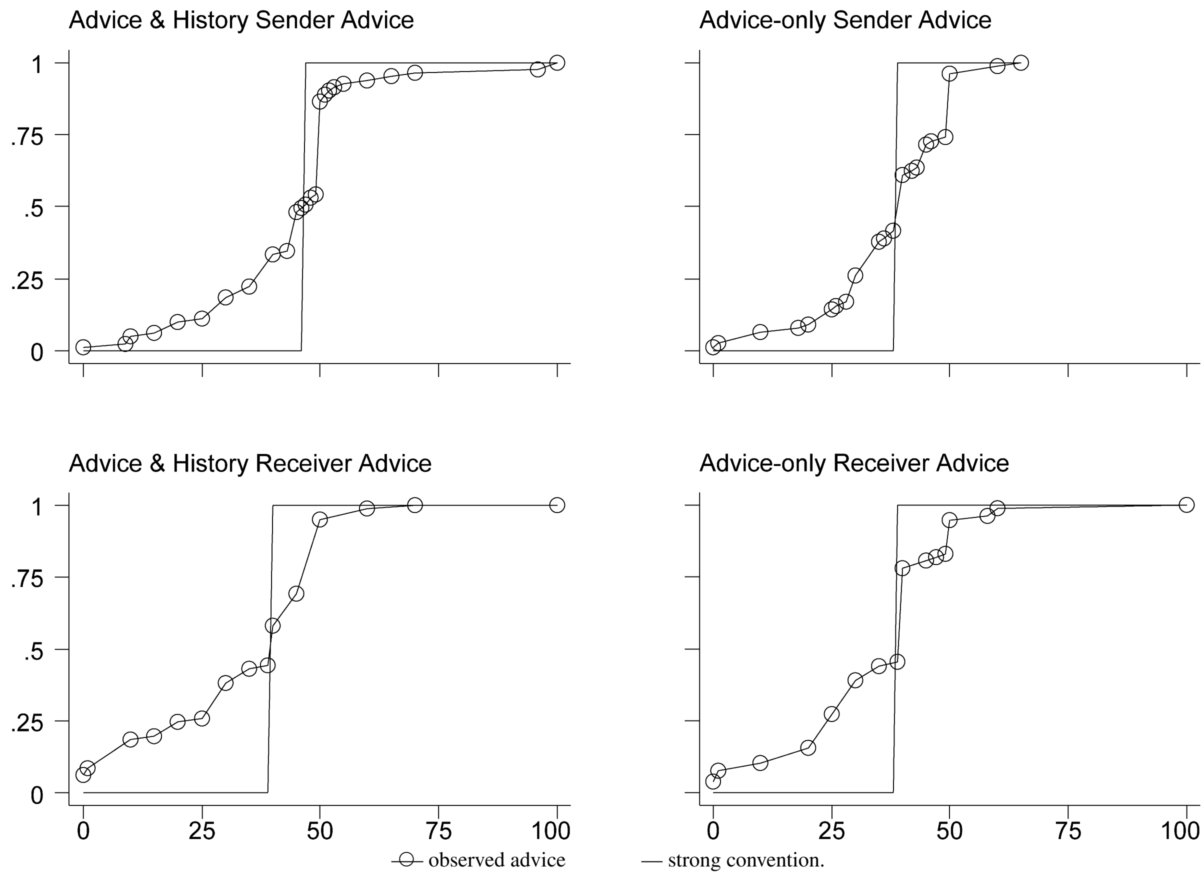


Fig. 4. Advice: amount sent and minimum acceptable offers.

**Hypothesis 1.** In the three Inter-generational Ultimatum treatments, we expect behavior to be more conventional in the treatments with advice (History-and-Advice and Advice-Only) than in the treatment without advice (History-Only).

What this means is that for each treatment and each of our six key behavioral variables, actions (amounts sent and accepted), beliefs (of Senders and Receivers) and advice (of Senders and Receivers), we will take the empirical distribution and compare it to the best fitting 0–1 step function posited by the strong convention concept using our area metric. We will call the behavior of that variable "more conventional" in one treatment than another if the area associated with it is smaller.

In addition to the simple between-treatment measure of conventionality that we are considering, it is useful to consider as well the within-treatment consistency of the measure. When a strong convention is in place the beliefs of the Senders should be consistent with those of the Receivers, which should be consistent with what is actually sent and accepted, etc. The overall consistency among all of the variables within a treatment can be measured by comparing the implied breakpoints of the strong convention associated with each variable. To measure consistency among variables, let us call $x^*_{ij}$ the break point estimated for variable $i$ in treatment $j$, $i = 1, 2, 3, 4, 5, 6$. If a strong convention existed then $x^*_{1j} = x^*_{2j} = x^*_{3j} = x^*_{4j} = x^*_{5j} = x^*_{6j} = x^*$. In other words, there must be agreement across variables in a treatment as to what $x^*$ is for a strong convention to hold. A measure of the departure from such agreement is the variance of these $x_{ij}$s ($\sigma^2_j$). These variances are presented in Table 4 and constitute our measure of within-treatment consistency.

**Definition 3** *(More Consistent Than).* We call behavior in experimental treatment $j$ *more consistent than* behavior in experimental treatment $k$ if $\sigma^2_j < \sigma^2_k$.

**Hypothesis 2.** In the three Inter-generational Ultimatum treatments, we expect behavior to be more consistent in the treatments with advice (History-and-Advice and Advice-Only) than in the treatment without advice (History-Only).

## 4. Results

### *4.1. Descriptive statistics—Sender and Receiver behavior*

Before we consider our convention creation hypotheses, we first look at raw data: the sending and rejecting behavior of our subjects—in each of our three treatments. We also check that there is no evidence that experience in the other games that were played by subjects in the same session systematically affected the results.

To explore offer behavior more systematically, consider Table 1a which presents some descriptive statistics about the offer behavior of our subjects and Fig. 5 which shows histograms of the offers made for each treatment.

The main empirical regularities evident are that when only advice is present (i.e., not history) the offers are lower and their variance is smaller. To demonstrate this 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 Advice-Only treatment were above 50, in the History-and-Advice treatment 17% were above 50 and in the History-Only treatment 18% of the observations were above 50. Note

Table 1a
Offers by senders

| Treatment | All generations | | Last 40 generations | |
|---|---|---|---|---|
| | Mean | S.d. | Mean | S.d. |
| Advice-and-History | 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 1b
Summary statistics by cohort on key variables

| Cohort: | 1st cohort | | 2nd cohort | | 3rd cohort | | K.–W. test | |
|---|---|---|---|---|---|---|---|---|
| | Mean | S.d. | Mean | S.d. | Mean | S.d. | $X^2$ | (prob) |
| Amount sent | 40.65 | 27.21 | 43.17 | 18.78 | 40.51 | 16.20 | 0.21 | (0.90) |
| Acceptance | 0.71 | 0.46 | 0.83 | 0.38 | 0.76 | 0.43 | 1.63 | (0.22) |
| Sender earnings | 39.19 | 27.24 | 43.65 | 25.48 | 42.95 | 27.22 | 1.31 | (0.52) |
| Receiver earnings | 31.48 | 22.64 | 39.01 | 28.82 | 32.73 | 22.11 | 2.14 | (0.34) |
| Sender advice | 59.96 | 28.32 | 57.19 | 29.66 | 57.32 | 30.36 | 0.82 | (0.65) |
| Returner advice[a] | 49.93 | 28.92 | 44.60 | 31.10 | 51.30 | 32.77 | 1.66 | (0.44) |

[a] Advice received that is relevant to the amount returner has been sent.

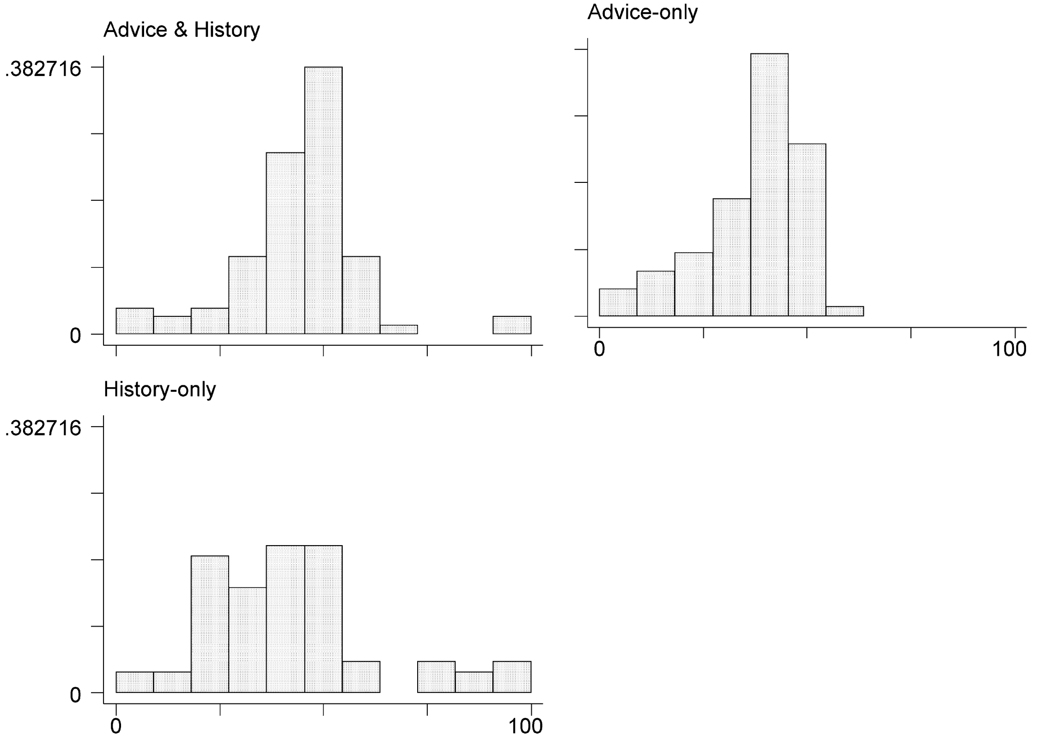


Fig. 5. Amount sent.

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).

These observations are supported by statistical tests. Concerning the location of the distribution, the Kruskal–Wallis test yields a chi-squared statistic equal to 2.59 with 2 degrees of freedom, implying a significance level of 0.002. This means that the hypothesis that the samples of offers are drawn from the same distribution is rejected. A series of multiple pairwise comparison tests indicate that there is a significant difference (at the 5% level) between the Advice-and-History treatment and the Advice-Only treatment. There is not a significant difference between the Advice-Only and History-Only treatments, while there is a less significant (at the 10% level) difference between the Advice-and-History and History-Only treatments. These results imply that the offers are significantly lower in the Advice-Only game than the Advice-and-History. Primarily, it would seem, because of the high variance in the History-Only treatment, it is harder to determine its location, relative to the treatments with advice.

A note about our statistical tests are in order. The statistical tests and estimated regressions reported in this section require independence of observations. Our position is that since each sender makes only one offer and each receiver makes only one accept/reject decision, the observations (of offers or accept/reject decisions) may reasonably be thought of as independent, at least under the null hypothesis. If the null of no difference is rejected, then the advice or history evidently does have an effect. However, it is not clear to us that this should be interpreted as implying that the observations are not independent. The fact that advice and/or history is present may simply alter the environment in which the decision is made in a more or less similar fashion for all subjects in that treatment. This seems especially true for Advice: since one only sees the advice of one's immediate predecessor, it seems unreasonable to suppose that subjects who are not in adjacent positions in the game are not to some degree independent of one another. If the observations were repeated observations on the same subjects, then there would be a clear lack of independence, and the corrections for this are well known and easy to apply. We grant that there may be some lack of independence, but it is less obvious what the nature of the dependence is and it is not clear how one should correct for it.

A series of one-tailed standard deviation tests for binary comparisons between treatments addresses the question of variance. There is a significant difference between the History-Only treatment and the Advice-and-History ($F_{(65,80)} = 2.16$, $p = 0.00$) and the History-Only treatment and the Advice-Only treatment ($F_{(65,76)} = 2.90$, $p = 0.00$). The same test found a difference between the variances of the Advice-Only treatment and the Advice-and-History at only the 10% level. So the main result here is that both of the treatments with advice have lower variance of offers than the one without history. 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 and, in fact, absence of history appears to be helpful in lowering the variance. The lower offer variance in the treatments with advice indicates that behavior is "more conventional" there, since it indicates that there is a greater consensus about what are acceptable offers when advice is present. We will substantiate this assertion in greater detail in our next section on convention creation.

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 a time trend, including treatment variables for the constant and slope terms. The constant and the time trend for the Advice-Only treatment are the only

Table 2
Relationship between offers and time

| Dependent variable: Offer | Coeff. | $t$-stat | Prob. > $\|t\|$ |
|---|---|---|---|
| Time | 0.03 | 0.34 | 0.74 |
| Constant | 43.61 | 3.73 | 0.00 |
| Advice-Only constant (AOC) | 1.18 | 0.22 | 0.83 |
| History-Only constant (HOC) | −3.53 | −0.63 | 0.53 |
| AOC × time | −0.22 | −1.91 | 0.06 |
| HOC × time | 0.04 | 0.33 | 0.74 |
| $N$ | 224 | | |
| $F$(3,218) | 2.85 | | |
| Prob. > $F$ | 0.02 | | |
| Adj. $R$-squared | 0.04 | | |

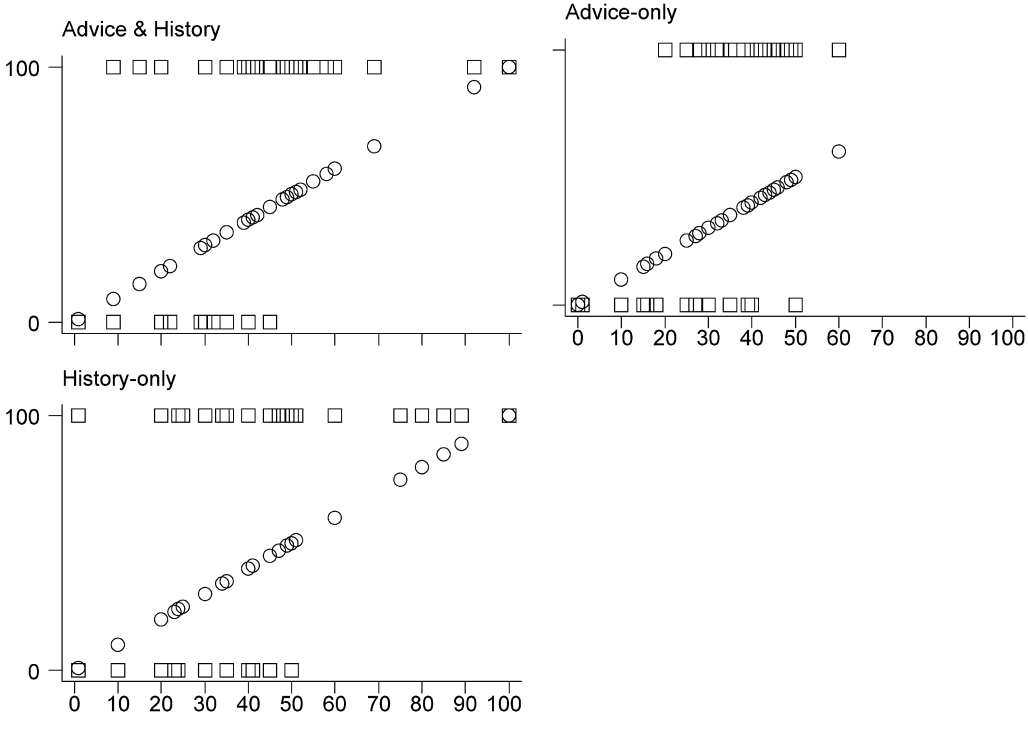


Fig. 6. Amount sent and acceptances. Raw data.

significant coefficients. The coefficient on time was negative and significant at the 0.06 level. See Table 2.

With respect to rejection behavior it appears as if the presence of advice makes Receivers tougher in the sense that they are more likely to reject low offers. Figure 6 illustrates the raw data—amount sent and whether accepted—for each treatment. Squares indicate offers made (on horizontal axis) and whether accepted or rejected (100 or 0 on vertical axis, respectively). Circles indicate the 45 degree line. Note that, besides the greater concentration of offers in the Advice-Only treatment, as we already saw in Fig. 5, there is also a greater concentration of acceptances.

To illustrate this fact more clearly, we estimate a logistic acceptance function for Receivers of the following form,

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

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 provides us with a continuous function showing the conditional rejection behavior of subjects in our three treatments and we can use this as a basis of comparison. The estimated functions are presented in the upper half of Fig. 2, which plots the fitted acceptance functions for each treatment. See Table 3.

What we see in Fig. 2 is that for low offers, the probability of acceptance is ordered in the manner described by the observation, i.e., lower offers 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), and vice-versa for higher offers. The Advice-and-History, in which both treatments exist simultaneously, is in between. In other words, the Advice-Only treatment exhibits the steepest acceptance function, and the History-Only game the flattest acceptance function.

Given our design one might be concerned that the behavior of subjects could be colored by their experience with other games played in the same session. For example, if one plays the Ultimatum Game last in a session, after playing the Trust Game and the Battle of Sexes Game, then the experience a subject has giving and receiving advice in this game may be affected by experience in previous interactions. Recall that in each session, subjects played three different games (with different opponents). One third of subjects played the Ultimatum Game first, one third played the Ultimatum Game second, and one third played the Ultimatum Game last in the session. Denoting each group as a separate "cohort," we can test for possible differences between those who played the Ultimatum Game with no experience, those who played with one previous

Table 3
Acceptance behavior (logit)

| | Coefficient (std. err.) | $z$ | $P > \lvert z \rvert$ |
|---|---|---|---|
| *Advice & History variable* | | | |
| Accept | | | |
| Sent | 0.10 (0.03) | 3.62 | 0.00 |
| Constant | −2.39 (1.07) | −2.24 | 0.03 |
| Obs. = 81 | | | |
| Pseudo $R^2 = 0.24$ | $LL = -29.62$ | | |
| *Advice-Only variable* | | | |
| Accept | | | |
| Sent | 0.16 (0.04) | 4.10 | 0.00 |
| Constant | −4.20 (1.32) | −3.18 | 0.00 |
| Obs. = 77 | | | |
| Pseudo $R^2 = 0.41$ | $LL = 24.71$ | | |
| *History-Only variable* | | | |
| Accept | | | |
| Sent | 0.022 (0.01) | 1.52 | 0.13 |
| Constant | −0.048 (0.61) | −0.08 | 0.94 |
| Obs. = 66 | | | |
| Pseudo $R^2 = 0.03$ | $LL = -39.16$ | | |

interaction in a session and those who played in two previous interactions in the session. If we find evidence of differences in key behavioral variables across the cohorts, then there could be cause for concern. Means and standard deviations by cohort for each of six variables, along with the Kruskal–Wallis test for differences in distributions, are reported in Table 1b. The variables are (i) amount sent, (ii) accept/reject decision, (iii) Sender earnings, (iv) Receiver earnings, (v) expected offer (Receiver beliefs), and (vi) expected return (Sender beliefs). We found no significant difference across cohorts for any of these variables. We thus conclude that the playing of other games in the same experimental session had no important effect, and now proceed to examine our hypotheses in more detail.

### 4.2. *Convention creation—hypothesis tests*

We will now systematically analyze the data, making use of both the conventionality measure and the consistency measure. Our general claim is that behavior was more conventional when advice existed. To illustrate this consider again Figs. 2–4. These figures present the behavior of both Senders and Receivers in each treatment for one of the three types of variables we are analyzing: actions, beliefs and advice. That is, amounts sent and the corresponding accept/reject decisions, beliefs of the Senders about rejection behavior and of Receivers about sending behavior, and the advice of Senders and Receivers to their successors.

For each of the variables in each figure, we have computed a measure of conventionality—the area between the relevant empirical function and the closest idealized (strongly conventional) step function. For each of our relevant variables, we have computed a goodness-of-fit area measure—the distance between the implied breakpoints in the associated step functions for each variable.

The same analysis is repeated and portrayed for each variable illustrated in Figs. 2–4. The top half of Fig. 2 plots the estimated acceptance functions while the bottom half of Fig. 2 plots the amount sent using the cumulative distribution of offers. Figure 3 plots the median belief of Senders (upper half of figure) and Receivers (lower half of figure), respectively, for amounts sent in each of ten ranges (for values of the amount sent from 0 to 10, 11 to 20, etc.). The horizontal axis indicates the midpoint of each range. The Sender belief functions (about likelihood of acceptance) are uncumulated (i.e., they indicate what the Sender believes is the likelihood of acceptance for any amount sent in any of our ten ranges), while the Receivers' belief functions (over the amount sent) are cumulated, since we elicited the probability that an amount sent in each range would arise and therefore could define for each amount $x$ what the probability is that an amount less than that $x$ would be sent.[9] In Fig. 4, the Sender and Receiver advice functions are presented. Both of these are cumulative functions representing the frequency with which advice to send or accept amounts less than a specific amount $x$ is observed. Of course, there are no graphs for advice for the History-Only treatment. The calculated areas are contained in Table 4. Also shown are the break points where the step occurs in the benchmark step functions.

While our area metric serves as our measure of conventionality for each variable, the behavior of our six key variables must be internally consistent if a meaningful convention is to exist. One could hardly say that a meaningful convention existed in treatment j if Senders thought that the conventional amount to send ($x^*_{\text{send } j}$) was, say, 25, while the Receivers were instructed to accept

[9] We use the median rather than the mean as it reduces the biases that naturally arise for the lower and higher ranges, where it is not possible to err very far in one direction. For the cumulated Receiver belief functions, we use the median of the cumulated function, not the cumulation of the medians.

Table 4
Best fitting step functions for action, belief, and advice distributions

| Variable | Advice-and-History | | Advice-Only | | History-Only | |
|---|---|---|---|---|---|---|
| | Break | Area | Break | Area | Break | Area |
| Amount sent | 46 | 10.14 | 40 | 8.12 | 40 | 16.21 |
| Acceptance function | 25 | 13.48 | 26 | 8.58 | 19 | 18.13 |
| Receiver beliefs (over amt. sent) | 45 | 6.39 | 39 | 6.32 | 42.5 | 9.47 |
| Sender beliefs (over acceptance) | 45 | 9.75 | 35 | 9.80 | 47 | 9.84 |
| Sender advice (amt. sent) | 47 | 11.62 | 39 | 9.48 | NA | NA |
| Receiver advice (min. acceptable offer) | 40 | 14.04 | 39 | 13.62 | NA | NA |
| Mean | 41.33 | 10.90 | 36.33 | 9.32 | 37.13 | 13.41 |
| Variance | 69.87 | 7.85 | 28.67 | 5.95 | 154.40 | 19.46 |

nothing less than 40 ($x^*_{\text{accept } j} = 40$) and the beliefs of Receivers had a step at 15, etc. Again, when a strong convention exists such inconsistencies cannot happen. To measure consistency we calculate the variance of these $x_{ij}$s within any treatment. These variances are presented in Table 4 as well.

Using the information contained in Table 4 we can consider a test of our two hypotheses. Since in each treatment we only have one time series, albeit they are of considerable length, we have a small number of data points on the area metric and consistency variables with which to work. Specifically, we have, for each treatment, only the 6 measures (4 for the History-Only treatment), 3 from the senders, and 3 from the receivers (2 each for the History-Only treatment). Moreover, the 3 Sender variables are not independent of one another, nor are the 3 Receiver variables independent. Still Table 4 does offer information that allows us to contemplate each of our hypotheses. Further, by appropriately combining the information from the different player types in each game and treating them as separate observations, we are able to conduct a valid nonparametric test that allows us to test Hypothesis 1.

#### *4.2.1. Hypothesis 1: conventionality*

A careful look at Figs. 2–4 provides support for Hypothesis 1. For example, the two most important variables constituting a convention of behavior in any Ultimatum game are the offer and acceptance behavior. If a convention of behavior is going to be established then it must be that the acceptance (or rejection) behavior indicates a clear demarcation between what is an acceptable amount to be offered while the Senders must have a clear indication of what they should send. Figure 2 shows that while such clarity about what an acceptable offer is might have been observed in the Advice-Only treatment, it certainly was not in the History-Only treatment. As was stated before, this is indicated by the difference in the area metrics (8.58 as opposed to 18.13) as well as the shape of the functions. The same can be said about the offer behavior of Senders in the lower half of Fig. 2. Here again, the area metric is much smaller in the Advice-Only treatment as opposed to the History-Only treatment (8.12 versus 16.21). Note also the vertical portion of the offer function at the step of 40 in the Advice-Only treatment. This implies that there is a considerable probability mass exactly at the implied step of 40. Such a mass point is not evident at the step of the History-Only offer function.

Further evidence of a convention of behavior in treatments with advice can be seen in the Sender and Receiver advice functions displayed in Fig. 4. Here again we see a set of functions that very closely approximate a step function. For example, in all of these figures the empirical function presented have mass points at or very near the break point of the step function. Again this indicates that there was a fair consensus about what was a conventional offered and this convention was evidently passed on from generation to generation.

Differences between the beliefs of Senders and Receivers across treatments was less clearly differentiated and tended to be less step function-like (see Fig. 3). This is may be because beliefs were measured in a less precise manner than other variables by asking questions about ranges rather than points and because beliefs are less concrete to subjects than, say, offers or acceptances, which can be made very precise.

Further evidence in support of Hypothesis 1 can be seen in Table 4. Here we see the area metric for all relevant variables tend to be smallest in the Advice-Only treatment than in either of the other two. For example, of the six variables of interest, the area metric is lower in the Advice-Only treatment for five of the six variables when comparisons are made between the Advice-Only and Advice-and-History treatments, while for the four variables that are common between the History-Only and the Advice-Only treatments the area metrics are smaller for the Advice-Only treatments in all four. The Advice-and-History treatment (where both advice and history are available) area measures are uniformly smaller than those in the History-Only treatment.

We employ the Kruskal–Wallis test, a nonparametric test for differences in distributions, to test Hypothesis 1. We combine the information gathered from Senders (Amount sent, Sender's Beliefs) and the information gathered from Receivers (Acceptance behavior, Receiver Beliefs) separately for each treatment and test for differences between the treatments.[10] Specifically, we use the average of the area metric for the two cumulated Sender variables and for the two cumulated Receiver variables. We thus have six observations in all, treating each treatment as a single session in which two pieces of data are gathered, one from the Senders and one from the Receivers. This test yields a test statistic of 4.57, which is significant at the 10% level, indicating that there is some difference between at least two of the treatments.[11] Multiple pairwise comparisons between the treatments using the Mann-Whitney two-sample test indicate a significant difference, at the 5% level, between the Advice-Only and History-Only treatments, with the area metric significantly smaller in the Advice-Only treatment. We interpret this as evidence that behavior in the Advice-Only treatment is more conventional than in the History-Only treatment.

#### *4.2.2. Hypothesis 2: consistency*

Hypothesis 2 states that there should be more consistency between beliefs and actions in the treatments with advice (Advice-and-History and Advice-Only) than in the treatment without advice (History-Only) with respect to what to offer and accept. We measure this consistency by the variance of the break-points in the 0–1 step functions. As we see in Table 2, the variance across our variables is by far smallest in the Advice-Only treatment (28.67) followed by the Advice-and-History (69.87) and the History-Only treatment (154.40). There appears to be more of a consensus about what to do, what to believe and what to advise others in the Advice-Only treatment than in the others. The conventional thing to do is more obvious.

As for our test of Hypothesis 1, we can combine information gathered from Senders in each treatment and from Receivers in each treatment to form a total of six observations. Specifically, we use the average of the breakpoints that minimize the area metric for Amount Sent and Sender

[10] We decided not to include the area metric for the advice variables for Senders and Receivers in the treatments with advice, as this could potentially influence the test, since there is not a comparable measure in the History-Only treatment. Any effect of advice should, in fact, be reflected in the actions and beliefs of the players.

[11] There is, in fact, no higher significance level possible for this test when there are only two observations per sample. See Siegel and Castellan (1988) for details.

Beliefs for Senders, and for Acceptance behavior and Receiver Beliefs for the Receivers. A formal test of Hypothesis 2 would require more independent data than we have at our disposal, since the hypothesis is about variance, and we have only two observations for each player type. But a comparison of the *difference* between these variables across treatments is suggestive of what such a test might yield. These differences (average Sender variables − average Receiver variables) are 10.1 for the Advice-and-History, 5.0 for the Advice-Only treatment and 12.75 for the History-Only treatment. Sender behavior is, on average, consistent with a higher breakpoint than is Receiver behavior for all treatments, but the *difference* is smallest for the Advice-Only treatment.

## 5. Other findings

While we have concentrated so far mostly on the impact of advice on the creation of conventions of behavior, we have in our data a wealth of information on a variety of other behaviors which we will highlight in this section.

### 5.1. Advice determines offers

It has become obvious from both this experiment and others that we have done using this inter-generational design (Schotter and Sopher, 2004, 2006; Chaudhuri et al., 2005), that advice has a big impact of behavior. To support this observation we present Fig. 7 which contains plots of the times series of the difference between offers and advice in each of the treatments involving advice in the Advice-and-History and the Advice-Only treatments.

Note from Fig. 7 the close fit between the advice that Senders receive from their predecessors and the offers they make. A simple yet effective model describing the offer behavior of Senders would simply be to explain amount sent as a function of the advice received. In other words, senders tend to follow advice. This can also be seen by comparing the relevant plots in Figs. 2 and 4 for the Advice-Only treatment. As one can see, there is a remarkable similarity between the two indicating that the offer behavior of Senders is closely related to the advice they receive.

### 5.2. Fairness

A great deal of attention has been given in the experimental literature to the impact of fairness considerations on behavior. In fact, the original debate about the Ultimatum game was whether

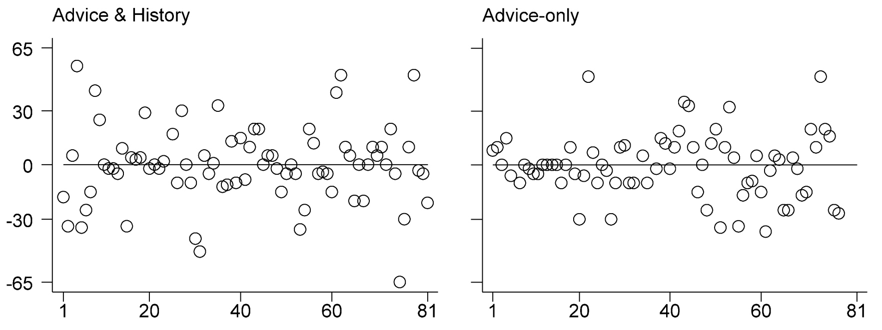


Fig. 7. Deviation of amount sent from advice.

what was observed was the result of fairness considerations or some artifacts of the experimental design. (See Camerer, 2003 for a review of these possible factors.) In addition, this simple game has led to the construction of new interdependent utility functions for economic agents which take the payoffs of others into account in defining utility (see Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002) and to the use of procedural fairness concepts (see Rabin, 1993).

While we have concentrated so far 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 the experiment but, perhaps, even of what subjects think Ultimatum Game experiments are about in general. Such data is unique to our experiment and the results we reach are summarized by the following observation.

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 one would do to justify the sub-game perfect equilibrium choice. 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 Advice-and-History where there were 81 generations there are 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 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 a statement that simply suggested a cut-off point without any real justification. ("Don't take less than 40"—subject 47 of 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 5 which present the results of our coding for the Advice-and-History and the Advice-Only treatment. (Note that because we allow a piece of advice to be fit (coded) into many categories, the total number of pieces of advice can vary between Senders and Receivers.)

One of the most striking features of Table 5 is the relatively infrequent use by Senders of fairness considerations to support their prescriptions. For example, fairness was not a principle

Table 5
Advice coding[a]

| Experiment | Type of advice | | | | | | |
|---|---|---|---|---|---|---|---|
| | | | | *Senders* | | | |
| | BRA | BI | FA | HA | OA | PA | |
| Advice-and-History | 38 | 4 | 8 | 5 | 10 | 19 | |
| Advice-Only | 21 | 6 | 11 | 23 | 6 | 7 | |
| | | | | *Receivers* | | | |
| | PR | BI | BI+ | $\frac{\text{FA}}{\text{Spite}}$ | HA | OA | PA |
| Advice-and-History | 7 | 11 | 3 | 29 | 8 | 11 | 13 |
| Advice-Only | 7 | 10 | 6 | 12 | 5 | 3 | 13 |

[a] See Section 5.2, *Fairness*, for definitions of codes.

that was invoked often (only 8 times in the Advice-and-History 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 Advice-and-History treatment, of the 24 cases in which a 50-50 split is suggested, only 7 were 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.

It is important to note however, that the fact that Senders did not justify their offers to each other in terms of fairness does not mean that fairness was absent from their thoughts. For instance, they certainly may have feared, and rightly so, that the Receivers did care about fairness and hence they had better offer more than they would like. Our point here is not to argue that fairness was not a factor in the minds of Senders but rather that their written advice indicates that they did not have a preference for fairness in that they seemed not to have the Receiver's payoff as an argument in their utility function. The only motive for offering anything was simply the fear of rejection and not inter-dependent preferences.

Also notable in Table 5 is the infrequent use of pure backward induction arguments. For example, for Senders in the Advice-and-History 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 Advice-and-History 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, not to justify sending zero but rather to justify sending some positive amount.

All of these observations lead us to conclude that we must reject the fairness hypothesis as the primary explanation of behavior, at least for Senders.

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 Advice-and-History 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 Advice-and-History 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.

### *5.3. Unfulfilled expectations cause rejections*

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 may be able to learn more about how subjects decide to accept or reject an offer.

Table 6
Rejection and acceptance behavior

| Prediction | Treatment | | |
|---|---|---|---|
| | A&H | AO | HO |
| *Variable: offer vs. expected offer* | | | |
| Acceptance: Offer $\geqslant$ Expectation | 33/66 (0.50) | 29/59 (0.49) | 19/46 (0.41) |
| Rejection: Offer < Expectation | 14/15 (0.93) | 17/18 (0.94) | 15/20 (0.75) |
| *Variable: offer vs. minimum acceptable offer* | | | |
| Acceptance: Offer $\geqslant$ Minimum | 62/66 (0.94) | 59/59 (1.00) | 43/46 (0.93) |
| Rejection: Offer < Minimum | 4/15 (0.027) | 3/18 (0.17) | 4/20 (0.25) |
| *Variable: offer vs. advised offer* | | | |
| Acceptance: Offer $\geqslant$ Advice | 50/66 (0.76) | 41/59 (0.69) | NA |
| Rejection: Offer < Advice | 10/15 (0.66) | 13/18 (0.72) | NA |

Table 6 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. The table shows the number of cases for which the difference between the offer and the other variables (expectation, minimum acceptable offer, or advice) correctly predicts acceptance or rejection behavior, over the number of possible cases:

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 Advice-and-History treatment, 14 occurred when the Receiver was not offered at least his expected amount. However, of the 66 acceptances in the History-and-Advice treatment, 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 is, in contrast, very good at classifying acceptances but bad at classifying rejections. For example, in the History-and-Advice treatment 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 History-and-Advice treatment, 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 a more detached "cold" phenomenon. (See Brandts and Charness, 2000.) Rejection behavior when expectations are unfulfilled is also consistent with Rabin's (1993) idea that kind and unkind acts are defined by the beliefs that people hold about each other which are used to measure their behavior.

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.

## 6. Conclusions

This paper has studied the impact of advice in inter-generational 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.

We find that the presence of advice makes a big difference in the Ultimatum Game, leading to more conventional behavior. The presence of history in combination with advice, however, does not seem to help. This is in contrast to our findings in a related study of behavior in an inter-generational coordination game (the Battle of the Sexes Game) in Schotter and Sopher (2004), where the most conventional behavior was observed in the Advice and History treatment. Advice alone in that study did lead to more coordination than did history alone. In yet another study, this of an inter-generational Trust Game in Schotter and Sopher (2006), we found that the presence of advice led to behavior more consistent with a backward induction solution to the game being studied than did the presence of history alone. The presence or absence of history in combination with advice, however, did not have a clear effect in that study.

It is thus a fair question to ask why the presence of history combined with advice sometimes leads to a higher, sometimes leads to a lower, and sometimes has an ambiguous effect on the degree of conventional play in different games. We believe the structure of the particular game in question is important in this regard. The Battle of the Sexes Game involves only two possible choices for each player, and there are two pure strategy equilibria. The presence of history combined with advice may help in coordinating on one of these equilibria in that the advice can refer to specific past instances of play, especially when there has been a string of play of one equilibria, which often was the case in that study. The Ultimatum Game, in contrast, allows the sender many more possible choices (101, to be precise), and we typically did not observe clear sequences of the same choices. In this instance, then, the presence of history may make it more difficult to converge to a particular solution to the game, as the history may be very diffuse in the guidance it provides. The Advice-Only treatment, then, may have a better chance of converging to a particular outcome, in that only the more recent history (that of the advice givers) is in view when Senders are making their choices.

Finally, the asymmetric nature of the Ultimatum game, and the considerations of fairness and equity that are often discussed in the context of experimental results, deserves some discussion. In contrast to the interpretation of many other researchers, we are not inclined to attribute the failure of the subgame perfect equilibrium prediction strictly to fairness in the sense of players having interdependent preferences. The Senders in our experiment, in particular, exhibit little evidence in their written advice to indicate that they regard the fairness of an allocation as the primary issue when recommending an action to a successor. Receivers, on the other hand, are more likely to express their unhappiness with what they regard as unacceptable offers in terms of the relative (un)fairness of the offer. We hazard the conjecture that what Receivers really find unfair is their relative powerlessness in the game, at least insofar as they are able to effect a nontrivial allocation of the available resources (they control only the degenerate allocation that gives zero to both parties). We believe that it is probably not reasonable to infer too much about behavior, in general, from the behavior of Receivers in this very artificial situation. Put differently, the behavior of Senders, and the way that they are affected (and helped) by the presence of advice is the most important finding here. The Senders have the more interesting and challenging task, in trying to find a proposal that will result in a profitable outcome, and the fact that advice helps in this endeavor and leads to more conventional behavior is our main finding, and the result that we think is most likely to have more general relevance in other settings.

## Acknowledgments

The financial support of the Russell Sage Foundation and the C.V. Starr Center for Applied Economics and the Center for Experimental Social Science 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 and the comments of two anonymous referees. We would also 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.

## Appendix A

### *A.1. Belief 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:

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

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.[12] Telling the truth is optimal.

[12] An identical elicitation procedure was used successfully by Nyarko and Schotter (2002) in their analysis of zero sum games and Schotter and Sopher (2004) in their investigation of inter-generational Battle of the Sexes games.

### A.2. Belief 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, $\pi_a^k$ is the probability that if an amount in the $k$th interval is sent it will be accepted while $\pi_r^k$ 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 it would be accepted with probability $\pi_a^k$, and it turns out that the offer was accepted. Then that Sender's prediction payoff would be defined as follows:

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

In other words, if the offer was accepted but the Sender only predicted that it would be accepted with probability $\pi_a^k$, 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.

## References


Berg, J., Dickhaut, J., McCabe, K., 1995. Trust, reciprocity, and social history. Games Econ. Behavior 10, 122–142.

Bolton, G., Ockenfels, A., 2000. ERC: A theory of equity, reciprocity, and competition. Amer. Econ. Rev. 90, 166–193.

Brandts, J., Charness, G., 2000. Hot vs. cold: Sequential responses and preference stability in simple experimental games. Exper. Econ. 2, 227–238.

Burke, M.A., Young, H.P., 2001a. Competition and custom in economic contracts: A case study in Illinois agriculture. Amer. Econ. Rev. 91, 551–573.

Burke, M.A., Young, H.P., 2001b. Contract uniformity and factor returns in agriculture, May.

Camerer, C., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton Univ. Press, Princeton.

Celen, B., Kariv, S., Schotter, A., 2004. The advice puzzle: An experimental study of social learning where words speak louder than actions. Mimeo. NYU, Version: May 27.

Charness, G., Dufwenberg, M., 2003. Promises & partnership. Research papers in economics 2003:3. Stockholm University, Department of Economics.

Charness, G., Rabin, M., 2002. Understanding social preferences with simple tests. Quart. J. Econ. 117, 817–869.

Chaudhuri, A., Schotter, A., Sopher, B., 2005. Talking ourselves to efficiency: Coordination in inter-generational minimum effort games with private, almost common, and common knowledge of advice. October.

Dufwenberg, M., Gneezy, U., 2000. Measuring beliefs in an experimental lost wallet game. Games Econ. Behav. 30, 163–182.

Fehr, E., Schmidt, K., 1999. A Theory of Fairness, Competition, and Cooperation. Quarterly Journal of Economics 114, 817–868.

Iyengar, R., Schotter, A., 2005. Learning with a meddlesome boss: An experiment in sloppy principal–agent relationships. Mimeo. New York University.

Nyarko, Y., Schotter, A., 2002. An experimental study of belief learning using elicited beliefs. Econometrica 70, 971–1006.

Rabin, M., 1993. Incorporating fairness into game theory and economics. Amer. Econ. Rev. 83, 1281–1302.

Roth, A., Prasnikar, V., Okuno-Fujiwara, M., Zamir, S., 1991. Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. Amer. Econ. Rev. 81, 1068–1095.

Schotter, A., 1981. The Economic Theory of Social Institutions. Cambridge Univ. Press, Cambridge, England.

Schotter, A., Sopher, B., 2004. Social learning and coordination conventions in inter-generational games: An experimental study. J. Polit. Economy 111, 498–529.

Schotter, A., Sopher, B., 2006. Trust and trustworthiness in intergenerational games: An experimental study of intergenerational advice. Exper. Econ., in press.

Siegel, S., Castellan, N.J., 1988. Nonparametric Statistics for the Behavioral Sciences, second ed. McGraw–Hill, Boston.

van Damme, E., 1998. On the state of the art in game theory: An interview with Robert Aumann. Games Econ. Behav. 24, 181–200.