Exp Econ (2006) 9: 123–145
DOI 10.1007/s10683-006-5386-y

# Trust and trustworthiness in games: An experimental study of intergenerational advice

**Andrew Schotter · Barry Sopher***

Received: 9 June 2004 / Revised: 22 August 2005 / Accepted: 22 August 2005


**Abstract** This paper investigates the development of conventions of trust in what we call *intergenerational games*, i.e., games played by a sequence of non-overplapping agents, who pass on advice on how to play the game across adjacent generations of players. Using the trust game of Berg et al. (1995) as our experimental decision problem, advice seems to decrease the amount of *trust* that evolves when this game in played in an inter-generational manner in that it decreases the amount of money sent from Senders to Returners. Ironically, advice increases *trustworthiness* in that Returners tend to send more back. Further, subjects appear to follows conventions of reciprocity in that they tend to Send more if they think the Returners acted in a "kind" manner, where kind means the Sender sent more money than the receiver expected. Finally, while we find a causal relationship running from trustworthiness to trust, the opposite can not be established. We note that many of our results can only be achieved using the tools offered by inter-generational games. The inter-generational advice offered provides information not available when games are played in their static form. Combining that information with elicited beliefs of the Senders and Returners adds even more information that can be used to investigate the motives that subjects have for doing what they do.

**Keywords** Trust · Trustworthiness · Advice · Intergenerational games

**JEL Classification** C91 · C72

*Resources for this research were provided by National Science Foundation grants SBR-9709962 and SBR-9709079 and by both the Center for Experimental Social Science and the C.V. Starr Center for Applied Economics at New York University. We would like to thank Shachar Kariv for both his comments and research assistance. We also thank Mikhael Shor and Judy Goldberg for research assistance, and Yevgeniy Tovshteyn for computer programming.

A. Schotter (✉)
Department of Economics New York University

B. Sopher
Department of Economics Rutgers University

## 1. Introduction

How do we learn to trust each other? Once a norm of trust is created how is it passed on from generation to generation? Does intergenerational communication increase or decrease trust? Does it increase or decrease trustworthiness? Is trust profitable? What is the causal relationship between trust and trustworthiness?

In this paper we aim to answer these questions by investigating a game of trust in an intergenerational setting where a game is played by a pair of players who subsequently are replaced by another pair, each replacement being a "descendent" of one of the original players and able to receive advice from their predecessor on how to behave in their play of the game. We analyze the impact of this intergenerational advice on behavior.

The trust game was first investigated experimentally by Berg et al. (1995). We consider the following version of the game: Player 1 moves first and can send Player 2 any amount of money $x \in [0, 100]$ or keep 100 for him or herself. Once $x$ is determined it gets multiplied by 3 and an amount $3x$ is received by Player 2. Player 2 can then decide how much of the $3x$ received, $y$, to send back to Player 1. The payoffs for the players are then $100 - x + y$ for Player 1 and $3x - y$ for Player 2. Note that this game is a game of trust since Player 1, by sending nothing, can elect to get a safe payoff of 100 for himself but if he sends any amount $x$ to Player 2, he places his fate in Player 2's hands and must trust him to reciprocate and send back at least $x$ to compensate him for his act of trust. Hence, Player 2 is trustworthy if he or she sends back an amount $y \geq x$ and is not trustworthy otherwise.

In our experiment this game is played as an intergenerational game. In these games a sequence of non-overlapping "generations" of players play a stage game (i.e., the static trust game) for a finite number of periods and are then replaced by other agents who continue the game in their role for an identical length of time. Players in generation $t$ are allowed to 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. In fact, the payoff to any subject in generation t will be equal to the payoff they receive in their lifetime plus an amount equal to $\frac{1}{2}$ of the payoff received by their successor in the game, so their one-period-ahead discount rate is $\delta = \frac{1}{2}$.

Some comments on our non-overlapping generations set up might be in order since some readers may feel that an overlapping generations model is more natural. First, in fact, our design is an overlapping generational experiment in the sense that the payoff functions people have span generations. That is, in our experiment the payoff to any one subject is equal to not only the payoff he gets in his lifetime but also $\frac{1}{2}$ of the payoff of his laboratory child. So people are linked inter-temporally via their payoff function but not strategically, as is usually the case in overlapping-generations models. Second, we wanted a non-overlapping structure for two other reasons. First is the practical reason that if we have overlapping generations and we did not finish all the experiments on one day, we would have to bring back the last generation again and have them play against the new young generation the next day or week. This creates some experimental design problems in addition to the problem of having to get the same people to show up.

More importantly, we feel that our set-up allows us to concentrate on the advice giving side of the problem as opposed to the strategic side. Let us explain. In order to get cooperation in an intergenerational game, be it overlapping or non-overlapping, we must have two things occurring. The players must all have the same equilibrium strategy in mind and, more importantly, they must coordinate on it and have it be commonly known that it is going to be used. We find it very difficult to believe that in games with multiple equilibrium subjects could find a way to pick out the same (complicated) dynamic strategy and coordinate on it without the use of some type of intergenerational advice. We use a non-overlapping set up to isolate the advice component and not confound it with the strategic element of overlapping generations. It is our feeling that while overlapping generations may lead to different types of equilibria, the transmission of these equilibria across generations is what advice is needed for and we concentrate on that.

Intergenerational games have proven to be very useful in describing the evolution of behavior in coordination (Battle of the Sexes), Ultimatum, and Minimum games (see Schotter and Sopher (2003a), Schotter and Sopher (2003b), and Chaudhuri et al. (2002), respectively). In the games studied here, each generation will play the trust game once and only once in a setting where $\delta = \frac{1}{2}$. Across treatments we vary the ability of subjects to receive advice and also the length of history they are able to look at before making their choice. More precisely, in the Baseline experiment, before a subject makes a choice, he or she is able to receive advice from his or her predecessor and can also view all of the choices made by both the Sender and Returner in all periods up to and including $t-1$. In the Advice-only Treatment the subject can get advice from his predecessor but can only view the history of choices made by his or her immediate predecessor. Finally, in the History–only Treatment subjects can view the entire history of the game before making a move but cannot give or receive advice. Note that advice is private and seen only by the person receiving it. Also, in no treatment is it possible to see any advice other than that given to you be your immediate predecessor so when people have access to history that is only the history of actions not advice.

What we demonstrate is that the presence or absence of advice in our experiments is a key factor determining the behavior and earnings of our subjects. For example, advice appears to decrease trust in our laboratory intergenerational games as measured by the mean amount of money sent from the Sender to the Returner. Specifically, while in the Baseline and Advice-only experiments the mean amount of money sent by the Sender was 25.9 and 28.1, respectively, it was 40.18 for the History-only Treatment where no advice exists. Ironically, the existence of advice appears to increase the trustworthiness of the Returner in the sense that they tend to return more than subjects who play the game without access to the advice of their predecessor. (We will explain this seeming contradiction by examining an asymmetry in the anchoring and adjusting heuristic that subjects seem to employ.) Despite, or perhaps because of this lack of trust, the payoffs to Senders in the Baseline and the Advice-only Treatment experiments where advice exists, (91.37 and 90.75 respectively) were higher than those of Senders in our no-advice (the History-only Treatment) experiment (83.84).

Later in this paper we investigate two different aspects of the trust problem. First, since our experiments are intergenerational as opposed to static, we can investigate the causality between trust and trustworthiness. For example, we can ask whether, in our data, an increase in the level of trust (measured as an increase in the amount of money sent by the Sender) calls forth higher levels trustworthiness (as measured by the amount returned) or whether the causality works in the reverse manner. What we find is that the causality is one way. If in the last generation a Returner returned more than the last Sender sent (i.e., there was reciprocation), the probability that the new Sender will increase the amount he or she sends over the last generation's Sender increases. However, if this generation's Sender increases the amount sent over last generation, the probability that the current Returner reciprocates does not significantly increase. Finally, we investigate whether our results can be explained by positing the existence of some type of interdependent utility function or reciprocal behavior as investigated by Rabin (1993) and Fehr and Schmidt (1999). Here we are fortunate to have elicited expectations of our Returners since we find that when Returners are sent an amount greater than what they expected to receive (i.e., when the Sender performed a "kind" act (as Rabin (1993) would call it)) they tended to reciprocate that kindness and increase the amount they sent back.

It is important to note that many of our findings can only be achieved using the tools offered by intergenerational games. The intergenerational structure (specifically, the historical information and the advice offered) provides information not available (to subjects or to researchers) when games are played in their static form. Combining this information with elicited beliefs of the Senders and Returners adds even more information that can be used to investigate the motives that subjects have for doing what they do.

We will proceed as follows. In Section 2 we present our experimental design. Sections 3, 4 and 5 presents our results while Section 6 offers some conclusions.

## 2. The experiment: Design and procedures

### 2.1. General features

In the intergenerational Trust Games reported on here subjects, once recruited, were ordered into generations in which each generation plays a Trust Game once and only once against an opponent. In the Baseline game subjects in generation $t$, after their participation in the game, are replaced by a next generation, $t + 1$, who will be able to view the history of what has transpired before them. Subjects in generation t are 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 percentage of the amount sent to return if they were the Returner). 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.

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 return strategy for the Returner 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 the entire history, depending on the treatment) before it and see what each generational Sender sent and the amount that was returned. 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 Returner using a proper scoring rule described in the Appendix. After the beliefs were elicited subjects played the Trust 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. After playing the game and learning their payoffs, subjects were then able to leave advice for their descendants. This advice took the form of a recommended action (how much to send for Senders, or the percentage to return, conditional on the amount sent, for Returners). (The procedures used are outlined in the Appendix.

The experiment was run either at the Experimental Economics Laboratory of the C.V. Starr Center for Applied Economics at New York University or at the Experimental Lab in the Department of Economics at Rutgers University. Subjects were recruited, typically in groups of 12, from undergraduate economics courses and divided into two groups of six with which they stayed for the entire experiment. During their time in the lab, for which they earned approximately an average of $26.10 for about $1\frac{1}{2}$ hours, they engaged in three separate intergenerational games, a Battle of the Sexes Game (BOSG) , an Ultimatum Game (UG) in which they were asked to divide 100 francs, and a Trust Game (TG) as defined by Berg et al. (1995). All instructions were presented on the computer screens and questions were answered as they arose. (There were relatively few questions so it appeared that the subjects had no problems understanding the games being played which purposefully were quite simple). All subjects were inexperienced in this experiment.

The experiment had three periods. In each period a subject would play one of the three games with a different opponent. For example, consider the following table:

**Table 1** Rotation scheme for subjects

| | | Game | | |
|---|---|---|---|---|
| | | BOS Game | Ultimatum | Trust |
| Period 1 | Row | Player 1 | Player 2 | Player 3 |
| | Column | Player 6 | Player 5 | Player 4 |
| Period 2 | Row | Player 2 | Player 3 | Player 1 |
| | Column | Player 4 | Player 6 | Player 5 |
| Period 3 | Row | Player 3 | Player 1 | Player 2 |
| | Column | Player 5 | Player 4 | Player 6 |

In this table we see six players performing our experiment in three periods. In period 1, Players 1 and 6 play the Battle of the Sexes Game while Players 2 and 5 play the Ultimatum Game and Players 3 and 4 play the Trust game. When they have finished their respective games, we rotate them in the next period so that in period 2 Players 2 and 4 play the Battle of the Sexes Game while Players 3 and 6 play the Ultimatum Game and Players 1 and 5 play the Trust game. The same type of rotation is carried out in period 3 so that at the end of the experiment each subjects has played each game against a different opponent who has not played with any subject he has played with before. Each generation played the game once and only once and their payoff was equal to the payoff they received during their generation plus an amount equal to 1/2 of the payoff of their successor in the generation $t + 1$ that followed them. (Payoffs were denominated in terms of experimental francs which were converted into U.S. dollars rates which varied according to the game played.) The design was common knowledge among the subjects except for the fact that the subjects did not know the precise rotation formula used. They did know they would face a different opponent in each period, however.

As a result of this design, when we were finished running one group of six subjects through the lab we generated three generations of data on each of our three games since, through rotation, each player played each game once and was therefore a member of some generation in each game. Thus for the set-up cost of one experiment we generated three generations worth of data on three different 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] We concentrate only on the results of our trust game here. The results of our other games can be found in Schotter and Sopher (2003a and 2003b)).

In our Trust Game, subjects were randomly assigned to role of Sender or Returner. The Sender was initially allocated 100 units of a fictitious laboratory currency called francs, which were later converted into dollars at the rate of 100 franc equals \$5. The task of the Sender was to send some amount of this 100 francs (say $x \in [0, 100]$) to the Returner. The amount $x$ was then multiplied by 3 and given to the Returner. The Returner could then either keep it all or send back any amount $y \in [0, 3x]$ to the Sender. Payoffs were $\pi_S = 100 - x + y$ for the Sender and $\pi_R = 3x - y$ for the Returner.

We made sure that the amount of money that could potentially be earned in the prediction part of the experiment was not large in comparison to the game being played. (In fact, over the entire experiment subjects earned, on average, \$26 while the most they could earn on all of their predictions was \$6.)

## 2.2. Parameter specification

The experiments performed can be characterized by four parameters. The first is the length of the history that each generation $t$ player is allowed to see. The second is the intergenerational 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 allowed between generations. In all of our experiments each generation lives for one period or plays 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 the are able to get advice from their predecessor or not. In the Baseline Trust 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. However, at period 52 we took the history of play and started two separate and independent new treatments at that point which generated a pair of new histories. In the Advice-only Treatment before any generation made its move it could see only the last generation's history and advice from the previous generation. This treatment captures the effect of advice on the play of the intergenerational game when history is not available. 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 captures the effect of history when advice is not available. This seemed to us the natural framework to use–to remove either advice or history, both of which would normally be present in an intergenerational setting–as a way of isolating the effect of the other type of information. The alternative of having neither advice nor history present in a baseline game, and then comparing play when advice, history, or both are added is less useful, as the baseline in this case is not, in any meaningful sense, intergenerational. The advice-only Treatment was run for an additional 78 generations while the History-only Treatment was run

[1] As far as we know, this is the record for economics experiments.

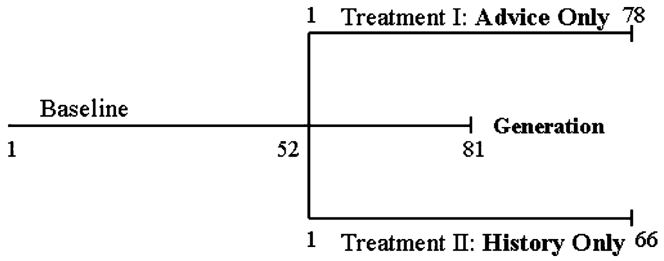


**Fig. 1** Experimental Design.

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 78 and the History-only Treatment was of length 66. Our experimental design can be represented by Figure 1.

As we can see from this diagram, we have only one time series for each treatment. The Baseline experiment lasts for 81 generations, the Advice-Only for 78 and the History-Only for 66 generations. In addition, the generations for the Advice-Only and History-Only treatments start at generation 52 of the Baseline. This was done to hold the past history constant for each treatment and have them start at a common point so as to isolate the impact of the treatment. Of course, in the Advice Only treatment generation 53 could only see the history of the last generation while the History-Only group could see the entire 52 period history of the Baseline.

This design differs from typical experimental design where an attempt is made to have many independent groups of subjects. In our intergenerational design this is not possible without facing a trade-off. Since generation $t + 1$ can not engage in the experiment until generation $t$ is done, our experiments are extremely time consuming. (In the typical experiment several groups can be run simultaneously rather than sequentially). This forces us to trade off the length on the time series we have against several independent trials of shorter lengths. Since we are interested in the evolution of trust conventions here, we chose to extend the length of the time series instead of having a larger number of independent groups each with shorter generational lengths. For most of the variables of concern in the paper we do have multiple observations, however. For example, for each treatment we have many observations on how subjects behave given advice, how they pass on advice given history, how they form beliefs, etc. Finally, we hasten to point out that many of our most cherished time series in economics, GNP, interest rates, etc. are unique time series and, as such, do not prevent statistical inference from being drawn from them. What we do have, however, is three different time series of considerable length run under different conditions—i.e., we are making use of the power afforded us by the experimental approach of generating data under alternative conditions.[2]

## 3. Aggregate results

In this section we summarize the aggregate empirical results of the experiment, including summary statistics on the amount sent, the amount returned, the earnings of Senders and

[2] There are many, say, GNP series (one for each country), but there are likely to be significant cross-correlations among them. They are, afterall, generated in the same "laboratory," i.e., the world. Our three series, by contrast, are substantially more independent of one another. The Advice-only and History-only groups do have the first 52 periods of the Baseline as a common heritage, but there is no obvious contemporaneous dependence among them, except as noted in the last paragraph.

**Table 1(a)** Summary statistics by game

| | Baseline | | Advice-only | | History-only | |
|---|---|---|---|---|---|---|
| Game | Mean | s.d. | Mean | s.d. | Mean | s.d. |
| Amount sent | 25.94 | 31.54 | 28.10 | 26.70 | 40.18 | 32.44 |
| Amount returned | 17.31 | 36.59 | 18.86 | 28.47 | 24.03 | 37.77 |
| Sender earnings | 91.37 | 26.54 | 90.76 | 21.09 | 83.85 | 32.79 |
| Returner earnings | 60.51 | 73.53 | 65.45 | 63.16 | 96.52 | 81.78 |
| Expected offer (Returner beliefs) | 23.34 | 18.12 | 27.81 | 18.52 | 32.38 | 17.41 |
| Expected return (Sender beliefs) | 35.04 | 48.25 | 36.62 | 43.28 | 61.26 | 56.26 |

Returners, and the expectations of Senders and Returners about what the other player would do in the game. Table 1(a) summarizes this information.

Since we analyze these results in detail in the paper let us pause only briefly to discuss this table. First note that while less is sent in treatments with advice more is returned when advice is absent (in the History-only treatment). Further, because of this behavior the earnings of Senders are greater than those of Returners. Finally, note that advice has the affect of dampening expectations. In both treatments where advice exists subjects both expect less will be sent and also expect less will be returned. These expectations turn out to be qualitatively correct while not precise quantitatively.

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 Trust Game last in a session, after playing the Ultimatum and 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 trust game first, one-third played the trust game second, and one-third played the trust game last in the sesssion. Denoting each group as a separate "cohort," we can test for possible differences between those who played the trust game with no experience, those who played with one previous 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. We conducted a series of tests using the Kruskal-Wallis test for differences in distributions, which are reported in the last column of Table 1(b). The testswere conducted for (i) amount

**Table 1(b)** Summary statistics by cohort

| | 1st cohort | | 2nd cohort | | 3rd cohort | | K-W test |
|---|---|---|---|---|---|---|---|
| Cohort | Mean | s.d | Mean | s.d. | Mean | s.d. | $X^2$ (prob) |
| Amount sent | 33.68 | 30.23 | 31.87 | 31.35 | 27.05 | 30.56 | 3.98 (.14) |
| Amount returned | 25.05 | 37.60 | 20.87 | 36.94 | 13.53 | 26.83 | 13.08 (.00) |
| Sender earnings | 91.37 | 29.63 | 89.00 | 26.75 | 86.48 | 24.43 | 1.03 (.60) |
| Returner earnings | 75.99 | 72.72 | 74.73 | 72.90 | 67.63 | 77.15 | 1.98 (.37) |
| Sender advice | 24.09 | 29.74 | 24.98 | 31.56 | 22.74 | 25.44 | .03 (.99) |
| Returner advice* | 15.01 | 39.67 | 6.09 | 13.04 | 13.55 | 47.49 | 5.14 (.08) |

*Advice received that is relevant to the amount returner has been sent.

**Table 1(c)** Relationship between Amount Returned (AR) and Amount Sent (AS), with cohort effects

| Dependent variable: AR | Coeff. | $t$-stat | Prob. > $\|t\|$ |
|---|---|---|---|
| Amount Sent (AS) | 1.11 | 7.93 | 0.00 |
| First cohort dummy | 13.74 | 1.47 | 0.14 |
| Last cohort dummy | −11.79 | −1.17 | 0.25 |
| (First cohort) × (AS) | −.16 | −0.80 | 0.43 |
| (Last cohort) × (AS) | −.05 | −0.23 | 0.82 |
| Constant | −26.46 | −3.84 | 0.00 |
| $N$ | 225 | | |
| # Left-censored | 88 | | |
| # Right-censored | 4 | | |
| Model Chi.Sq. (Prob. > Chi. Sq.) | 137.49 (.00) | | |

sent, (ii) amount returned, (iii) Sender earnings, (iv) Returner earnings, (v) expected offer (Returner beliefs), and (vi) expected return (Sender beliefs). We found a significant difference across cohorts only for the amount returned (at the 1% level of significance) and for Returner advice (at the 8% level of significance). Specifically, the amount returned declines over the three cohorts, from an average of 25 francs for those with no experience, to 21 for those with one previous interaction, to 13.5 for those with two previous interactions in the session. The fact that the earnings variables for both the Sender and the Returner do not differ across cohorts, though, seems to us more important. The Sender and Returner earnings variables range from about 86 to 91, and from about 68 to 76 francs. The return variable averages only about 19 francs, and a small (if significant) variation in it has a relatively small impact on earnings. Even though there are significant differences across cohorts in amount returned, these differences are not important, in that they do not affect (significantly) players's earnings. Similarly, the Returner advice variable, though it differs across cohorts, is evidently not decisively important.

It should be kept in mind, as well, that the amount returned is constrained by the amount sent. To illustrate this more clearly, we estimated the relationship between the amount returned with the amount sent as a dependent variable, allowing for cohort effects in both the constant and slope terms. The results of this regression are reported in Table 1(c). [3] The estimates show that the amount returned is largely a function of the amount sent. Controlling for the amount sent, neither the first nor the last cohort dummy variables are significantly different from the second cohort. However, the difference between the first and last cohort dummy variables is signficant ($F(1, 220) = 6.56$, Prob. $> F = .01$), indicating that significantly less is sent, on average, by the last cohort than by the first cohort. On the other hand, neither of the cross-slope terms are significant (nor is their difference significant), indicating no difference in the amount returned as a proportion of the amount sent. Overall, it does seem that there is a declining propensity to send in later cohorts within a session. We will see below in Section "Trustworthiness and returner behavior" that return behavior, as a function of the amount sent, differs more among treatments than among cohorts.

[3] The regression is a censored normal regression, a generalization of the Tobit model. See Section "Return behavior" below for more explanation.

## 4. Trust and sender behavior

The decision to send any money is an act of trust. Since keeping a dollar guarantees its consumption after the experiment, sending one and hoping for its return explicitly exposes the Sender to a loss unless he or she can trust that the Returner will reciprocate. In our experiment we find that the level of trust is closely related to the presence or absence of advice with the impact of advice being negative. That is, when advice exists subjects send less than they do in the experiment without it. This is not surprising since in general Senders tend to follow advice and the advice they receive is trust-decreasing.

### 4.1. Offers and advice: Does advice increase or decrease trust?

To demonstrate the relationship between the amount sent and the availability of advice, consider Figure 2 which present histograms of the amount sent in each treatment.

Note that in all treatments the amount sent is substantially above the zero predictions of the static sub-game perfect Nash equilibrium prediction. For example, in all of our treatment 84% of the subjects send something positive. The mean amount sent is 31, and the median amount sent is 25.

The presence of advice has a dramatic impact on sending behavior, however. As we can see in Figure 2, the amount sent is substantially higher in the History-only Treatment where there is no advice than in either the Baseline or Advice-only Treatment. For example, the mean (median) amount sent in the Baseline and Advice-only Treatment respectively is 25.94 (15) and 28.10 (25) while in the History-only Treatment where there is no advice, it was 40.18 (30). A set of two-sample Wilcoxon rank-sum tests indicate that while there is no

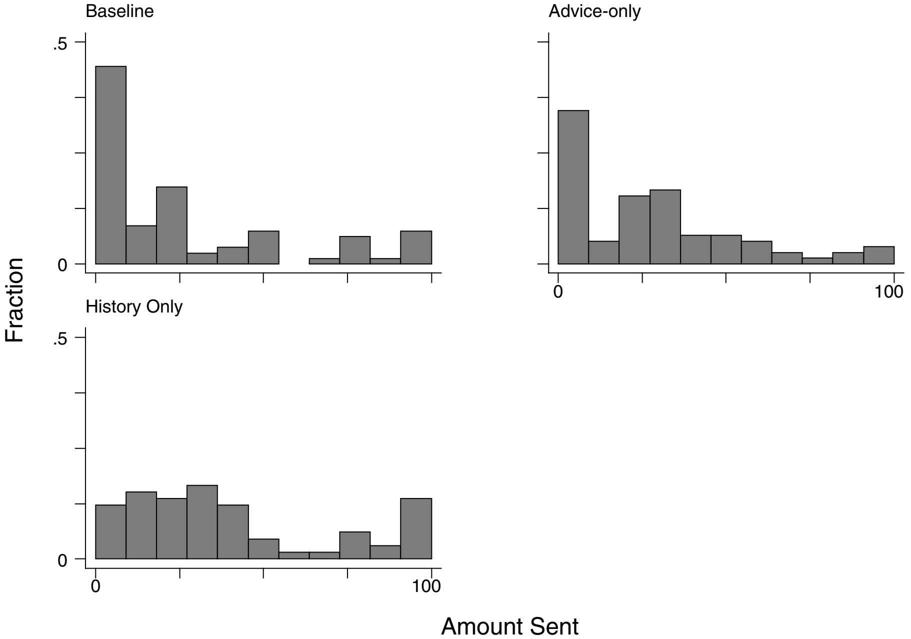


**Fig. 2** Histograms of amount sent by treatment

significant difference between the samples of Baseline and Advice-only Treatment offers ($z$-statistic $-1.24$, $p$-value .22) a significant difference did exist between the amounts sent in the History-only Treatment and both the Baseline and Advice-only Treatment (**Advice-only v. History-only** $z$-statistic $= -2.13$, $p$-value $= 03$, **Baseline v. History-only** $z$-statistic $= -3.03$, $p$-value $= 00$).

In addition, while the inter-quartile range of offers in the Baseline and the Advice-only Treatment were 1–40 and 5–40 respectively, the same range was 15–55 in the History-only Treatment. Another measure of trust can be gleaned from the upper end of the offer distribution. For instance, 10% of all offers in the Baseline and the Advice-only Treatment experiments were greater than 80 and 65 respectively, while 10% of all offers in the History-only Treatment were equal to 100 indicating an extreme willingness to "risk it all". Finally, to demonstrate the impact of advice on amounts sent we ran a linear regression of the amount sent on a $\{0,1\}$ dummy variable depicting whether or not advice was allowed in the experiment generating the observation. According to this regression we again observe a significant and negative relationship between the presence of advice and the amount sent. On the basis of these results we conclude that advice lowers the amount of trust in this game by lowering the amount of money sent.

### 4.2. Was advice followed?

What is interesting to note is that while advice lowers the amount of money sent, subjects tend to send more than they are advised to. For example, in the Baseline treatment the mean amount advised to be sent is 19.42 while the mean amount sent is 25.94, with 29 subjects sending more, and 24 sending less than they were advised. In the Advice-only treatment, while the mean amount sent tended to be no different from the amount advised, 28.10 versus 28.99, 36 subjects sent more than they were advised to while only 21 sent less. This difference is significant at the 11% level using a Wilcoxon matched pairs sign rank test ($z = 1.58$, $p > |z| = .11$), suggesting that the distribution of amount sent and advised amount to send may be different (pooled over both games where advice is allowed). The difference between what was sent and what was advised to be sent can be seen in Figure 3 where we plot for each generation the amount sent on the vertical axis and the corresponding advice on how much to send that the Sender received on the horizontal axis. The somewhat larger proportion of plotted points lying above the 45 degree line in the figure, especially for the Advice-only

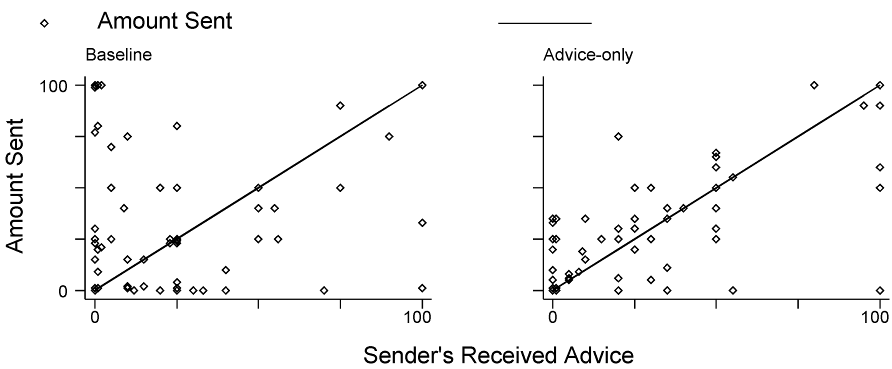


**Fig. 3** Amount sent vs. amount advised to send

game, indicates that there is some tendency for subjects to send more than they are advised to send.

What we think is going on here is that if left to their own devices, as in the History-only treatment, subjects tend to send a considerable amount. The amounts advised to be sent are so low, however, that it lowers the anchor in the mind of the subjects as to what is the appropriate amount to send. The subjects then adjust upwards but not enough to overcome the stingy advice. The net effect is to lower the amount sent. [4]

More generally, in the absence of advice, as Figure 2 showed, the amount sent is highly variable. Advice and history seem to reduce this variance considerably. The points in Figure 3 lie closer to the 45 degree line for the Advice-only treatment. In fact, the average difference between the amount sent and the advised amount to send is 6.71 for the Baseline treatment and −.65 for the Advice-only treatment. This suggests that advice is more influential without history—the amount sent is closer to the amount advised to send. When history is also present, as in the Baseline treatment, one may also, quite naturally, be influenced by the amounts that have been sent in the past.

### 4.3. Does advice increase payoffs?

The next question we ask is: Is having access to advice profitable in the sense that it increases the payoffs of Senders in the Baseline and Advice-only Treatment relative to those in the History-only Treatment? The answer is a qualified yes. The payoff to Senders in the experiments with advice (the Baseline and the Advice-only Treatment) were higher than those in the History-only Treatment where only history was available. For example, while the mean (median) payoffs to Senders in the Baseline and Advice-only Treatment were 91.37 (99) and 90.76 (95) respectively, those same payoffs were 83.85 (88) for the History-only Treatment. Two-sample Wilcoxon rank-sum tests indicate no difference between the treatments with advice (Baseline and Advice-only Treatment $z$-statistic .03, $p$-value .98) while there is some evidence, though not overwhelming, of a difference between the History-only Treatment and both the Baseline and Advice-only Treatment ($z$-statistic 1.55, $p$-value .12, and $z$-statistic 1.39, $p$-value .17, respectively).

Figure 4 present histograms of the Sender payoffs for each treatment. The tails of the distribution of payoffs for Senders in the History-only Treatment appear to be thicker than in either the Baseline or Advice-only Treatment. More precisely, the standard deviation of Sender payoffs in the Baseline and Advice-only Treatments respectively, are 26.54 and 21.09, while it is 32.79 in the History-only Treatment. $F$-tests indicate that there are significant differences for all pairwise comparisons of the treatments (Baseline v. Advice-only, .$F$-statistic 1.58, $p$-value .04, Advice-Only v. History-only, .$F$-statistic −2.42, $p$-value, .00, Baseline v. History-only, .$F$-statistic −1.53, $p$-value .08). Note that we do not have purely independent observations here, so the $F$-tests should be taken as suggestive. The observations are on different pairs of subjects, however, not repeated observations on the same subjects, so the observations can be reasonably thought to have some degree of independence.

In summary, there is some indication that the presence of advice is beneficial to the Sender but the significance levels for this result are not high. Advice does tend to lower the variance

[4] This behavior is not very different from that seen when Americans go to Europe where they do not know the tipping norms. After being advised that tipping means leaving the small change left on the plate for the server, most Americans will throw in some more, yet the effect is that they leave less than those uncoached Americans who think they need to tip as they do back home.

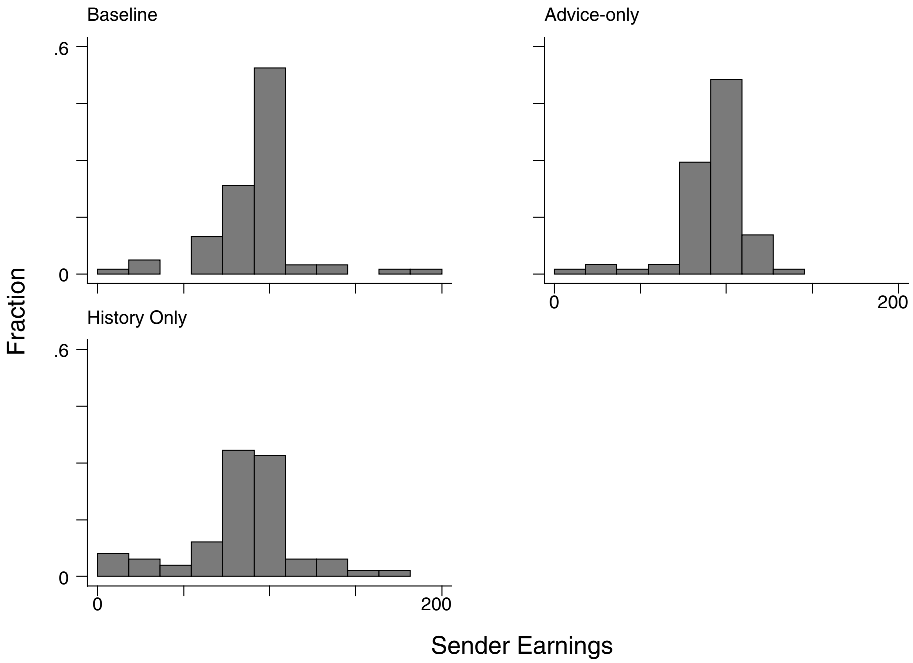


**Fig. 4** Sender earnings by treatment

in payoffs to Senders, while availability of advice tends to raise payoffs, with or without history.

## 5. Trustworthiness and returner behavior

### 5.1. Return behavior

While Senders exhibit trust in their sending behavior, Returners prove themselves to be trustworthy by what they do in return. Obviously, a Returner is trustworthy if he or she returns more than he or she received. On average, the number of times Returners reciprocated was rather high especially in comparison to the predictions of the static sub-game perfect equilibrium prediction of no reciprocity. In the Baseline experiment Returners returned an amount greater than or equal to the amount they received 40.7% of the time while in the Advice-only Treatment and the History-only Treatment they did so 38.5% and 37.9% of the time, respectively. These return ratios actually overstate the degree of reciprocation for a number of reasons. First, a large fraction of the time when money was sent it was very little. For example, if those cases where the amount sent was 0 (and thus, the amount returned was also 0) are not treated as reciprocation, then the rates are roughly halved, to 23%, 21% and 26% for the Baseline, Advice-only, and History-only treatments, respectively. As mentioned above, for the Baseline the median amount sent was 15 while in the Advice-only and History-only Treatments these median amounts were 25 and 30 respectively. Hence, in these circumstances, little trust was shown and little trustworthiness was shown in return. Further, reciprocation may not even be expected for such low amounts if subjects adhere

**Table 2** Average amount returned

| Amount sent | Baseline Amount returned (# of obs.) | Advice-only Amount returned (# of obs.) | History-only Amount returned (# of obs.) |
|---|---|---|---|
| 0 | 0 (14) | 0 (14) | 0 (8) |
| 1–10 | 1 (24) | 2.1 (15) | 6.5 (4) |
| 11–20 | 7.5 (6) | 11.2 (5) | 7.7 (12) |
| 21–30 | 14.7 (14) | 9 (13) | 14.3 (10) |
| 31–40 | 5.8 (4) | 26.1 (14) | 23 (10) |
| 41–50 | 25 (6) | 42 (5) | 21.4 (5) |
| 51–60 | – (0) | 50 (4) | 5 (1) |
| 61–70 | 30 (1) | 50 (2) | 0 (1) |
| 71–80 | 60 (5) | 100 (1) | 75 (4) |
| 81–90 | 20 (1) | 15 (2) | 62.5 (2) |
| 91–100 | 100.7 (6) | 86.7 (3) | 62 (9) |

to the fairness theories of scholars like Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) that predict that for low amounts sent it may be the case that fairness dictates that nothing be sent back and that the Sender would prefer that as well. For example, if 25 were the amount sent then keeping it would equalize the payoffs of the Sender and Returner in that generation at 75 each. Hence, amounts below 25 (what might be termed the *equal-payoff threshold*), even if kept by the Returner, still determine a distribution of payoffs equally favorable to the Sender and Returner, and hence keeping all could be justified on fairness grounds. That is, the resulting allocation will be more fair than if something were returned.

Our data offers an insight into whether the trust exhibited by Senders was reciprocated by Returners. To investigate this relationship see Table 2, which shows how much of the amounts sent in various intervals were reciprocated.

As we see in this table, the amounts returned vary depending on the amount sent. There is a roughly monotonic increasing pattern in the average amount returned (note that the number of observations upon which the average is based is often quite small). Overall, sending behavior seems to be unprofitable since the average amount returned is typically less than the mid-point of the interval of amounts sent. Figure 5 graphs the amount returned (on the vertical axis) against the amount sent (on the horizontal axis). Note that while there is a positive relationship between the amount sent and the amount returned, it is typical that less is returned than was sent, since most of the data lies below the 45 degree line in the figure. On average, sending anything is a losing proposition.

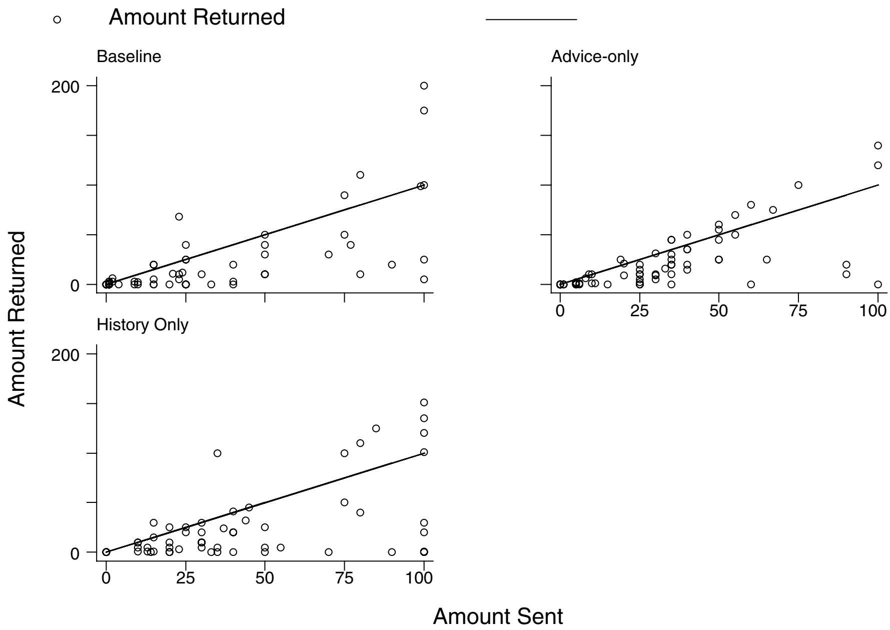


**Fig. 5** Amount returned vs. amount sent

To investigate the relationship between the amount returned and the amount sent in a more systematic manner, we use a censored normal regression, a generalization of the Tobit model that allows for variable threshold levels. The upper censoring threshold level varies because the amount returned has as an upper limit equal to three times the amount sent. We assume that the amount returned is a reflection of an underlying latent variable, determined by factors which affect trustworthiness. A low level of trustworthiness leads to nothing being returned, while a high level of trustworthiness leads to the return of everything (3 times the amount sent). Intermediate levels of trustworthiness result in a positive amount being returned. The use of the censored regression is important here, both conceptually and practically.

Just as Tobin (1958) conceived of expenditures on consumer durables as a measure of desire for such goods, with this measure censored when expenditures were zero (i.e., he conceived of negative desire for such goods), we can conceive of the amount returned as a measure of trustworthiness. Trustworthiness thus conceived is censored at zero, because some individuals returning zero might be near the threshold of giving something back (if, say, only a litte more had been sent to them), while others could be quite far away (angry, say, that the other person only sent 1 franc). The coefficient on the Amount Sent will be biased downward if this censoring is not accounted for in the estimation. Finally, note that the amount returned is also censored at the upper end by three times the Amount Sent: In principle, one might wish to return more than three times the Amount Sent (i.e., the entire amount received from the Sender), and if one in fact returns the full amount, then one could conceive of the possibility that he or she would have returned more, if this were possible. There are far fewer cases in the data that are at the upper limit than at the lower limit of zero, and this case is in fact less plausible, and as a practical matter less important. Ignoring the right censoring does not change any of the estimates appreciably.

**Table 3** Relationship between Amount Returned (AR) and Amount Sent (AS)

| Dependent variable: AR | Coeff. | $t$-stat | Prob. > $|t|$ |
|---|---|---|---|
| Amount Sent (AS) | 1.19 | 8.70 | 0.00 |
| Advice-only Dummy | 6.28 | 0.67 | 0.50 |
| History-only Dummy | 14.39 | 1.45 | 0.15 |
| (Advice-only) × (AS) | −.06 | −0.27 | 0.79 |
| (History-only) × (AS) | −.37 | −1.88 | 0.06 |
| Constant | −31.70 | −4.77 | 0.00 |
| $N$ | 225 | | |
| # Left-censored | 88 | | |
| # Right-censored | 4 | | |
| Model Chi.Sq. (Prob. > Chi. Sq.) | 129.48 (.00) | | |

The regression results shown in Table 3 indicate that the relationship between the amount returned (AR) and the amount sent (AS) is a significantly positive one. The Advice-only Treatment appears not to be significantly different than the Baseline, but there is some evidence that the History-only Treatment is different. Though the relationship is positive, it is insufficient to ensure reciprocation. The fitted values of the estimated equation, as well as the actual amount returned, are plotted against the Amout Sent in Figure 6. Note in Figure 6 that the actual Amount Returned values equal to zero, or very small, are concentrated to the left of Amount Sent equal to 25, while the relationship is more clearly a positive for values of Amount Sent greater than 25.

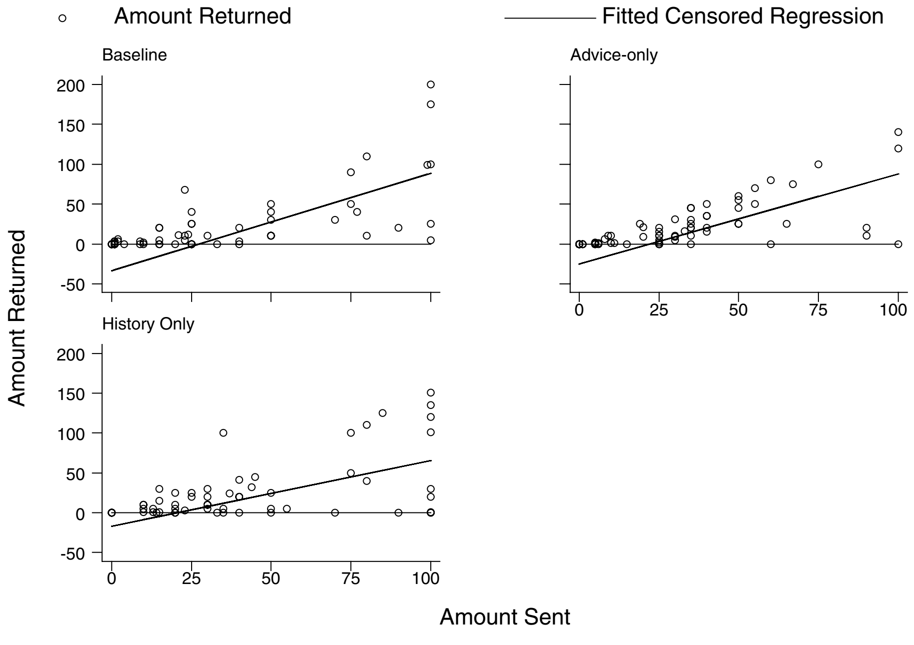


**Fig. 6** Fitted and actual amounts returned vs. sent

The coefficients imply that the Sender gets a positive return only when he or she sends an amount above a threshold that is approximately equal to the equal payoff-threshold of 25. The fitted values imply thresholds of 26.64, 22.50, and 21.11 for the Baseline, Advice-only and History-only games. The hypothesis that the threshold equals 25 cannot be rejected for any treatment. The fitted values of the slope coefficient on the Amount Sent are 1.19, 1.13, and .82 for the Baseline, Advice-only and History-only games, respectively. The hypothesis that the slope coefficient is at least one cannot be rejected for any of the games, however. The fitted values are larger for the advice games, suggesting that advice makes reciprocation more likely. This reciprocation seems to be of a qualified sort, however, in that it is evidently conditional on some basic level of equity being achieved. As we will see in our next section, this relationship can be interpreted as one in which the advice offered functions as a lower bound on the Returner's actions by locating a salient anchor from which return payments are adjusted upwards.

### 5.2. Does advice make returners more or less trustworthy?

Interestingly, advice appears to lead to higher return rates than those existing in the History-Only treatment. For example, consider Table 4 where we list the mean and median differences between the amount sent and returned in all three experiments. (Note a smaller number implies a greater returned amount, relative to the amount sent).

The median difference between the amounts sent and returned in the History-only treatment is 12 while in the Baseline and Advice only treatment this difference is 0 and 5 respectively. Our explanation for this is similar to the one we used to explain the impact of advice on the amount sent. Here we assume that the advice received by one's predecessor acts like a floor or anchor from which one revises or adjusts his or her estimate of what is the right amount to return. Without this advice, the floor may be 0 and hence returned amounts are lower. (A similar result was found in Issac and Plott (1981) where the existence of a non-binding price floor raised prices in a double oral auction by functioning as a psychological anchoring point for the bidding behavior of subjects). So the existence of advice has an asymmetrical impact on Sender and Returner behavior. For the Sender, advice on what to send creates an anchor that is below what subjects would normally send and hence even when the amount sent is adjusted upwards it remains below what subjects send without advice. For Returners just the opposite is true. If left to their own devices subjects might return zero, but the advice they receive is strictly positive and when they adjust the amount they send upward

**Table 4** Trustworthiness (Send-return)

| Treatment | Mean | Median | Std. dev. |
|---|---|---|---|
| Baseline | 8.63 | 0 | 26.54 |
| Advice-only | 9.24 | 5 | 21.09 |
| History-only | 16.15 | 12 | 32.79 |

Test for Difference in Means of (Send-Return)

| Difference Tested | $t$-stat | Prob. $> \|t\|$ |
|---|---|---|
| Baseline v. Advice-only | −.16 | .87 |
| Baseline v. History-only | −1.54 | .13 |
| Advice-only v. History-only | −1.53 | .13 |

it leads to returning more than subjects do in the History-Only experiment. These differences between treatments are not statistically significant, however.

To get a better idea of the advice function (i.e. the relationship between the amount of money that one is advised to return out of what is sent) consider, the results of a censored normal regression of the advised amount to return on the amount sent. We assume that the amount that the Returner is advised to return by his or her predecessor is a reflection of an underlying latent variable, determined by factors which affect trustworthiness. The idea is that if one has a sufficiently low level of trustworthiness, then one is likely to advise one's successor to return nothing, while if one has a higher level of trustworthiness, then one is more likely to advise that one's successor return a large proportion. The dependent variable is constructed from the advice, which was in the form of a percentage of the amount sent to return, and it is regressed on the midpoint of the relevant range of possible amounts that one might be sent. We also include a dummy variable indicating whether advice was present in the experiment generating the observation.

Letting ADR stand for the advised amount to return and AS the (hypothetical) amount sent, these regressions yield the relationships contained in Table 5.

Comparing these regression results with those for the return-send relationship described in Table 3 above, we confirm our observation above that, typically, more is returned than the subject is advised to return. In particular, the regressions for the games with advice have essentially the same constant term, but the slope coefficient on amount sent is greater than 1 for the actual amount returned, less than 1 for the advised amount to return. Thus, on average, subjects advise the next generation to return a smaller proportion of the amount they receive than, on average, subjects actually return.

## 5.3. Fairness and reciprocity

### *5.3.1. Fairness*

Do the regressions presented above give us any insight into whether our subjects are selfish or fair? In some sense they do. For example, these functions are conditional advised-return functions since they indicate how much a Returner is advised to return conditional on receiv-

**Table 5** Relationship Between Amount Advised to Return (ADR) and Amount Sent (AS)

| Dependent variable: ADR | Coeff. | $t$-stat | Prob. $> \|t\|$ |
|---|---|---|---|
| Amount Sent (AS) | 0.87 | 12.31 | 0.00 |
| Advice-only Dummy (AO) | −1.97 | −0.33 | 0.74 |
| (AO) × (ADR) | −0.12 | −1.22 | 0.22 |
| Constant | −30.57 | −7.28 | 0.00 |
| $N$ | 1590 | | |
| # Left-censored at AS <= 0 | 608 | | |
| # Right-censored at AS >= 100 | 17 | | |
| Model Chi.Sq. (Prob. > Chi. Sq.) | 239.06 (.00) | | |
| LL | −5687.40 | | |
| Pseudo-Rsq | 0.02 | | |

ing any amount, AS. As we see, there is a positive (and significant) relationship between the amount advised to be returned and the amount sent. Returners suggest that their successors return positive amounts but only after the amount sent was above some threshold. Hence, it would appear that our subjects were not selfish maximizers since this behavior is not consistent with selfish preferences. [5](The advice function for a selfish Returner would be a function with a zero slope and zero intercept.) Returners are not on average advised to send back more than the Sender sent however. So while Returner subjects are not perfectly selfish, that fact would not be consolation to a selfish Senders who typically fail to get back what they sent.

However, the fact that Senders do not fully receive back what they send does not mean that they are always unhappy. To know their true utility level we would need to know their intentions in sending money in the first place and the beliefs of the Returners about their intentions. For example, say a Sender was a true egalitarian. He might send 25 and hope that the Returner keep it all. If the Returner did so, then the payoff to both of them would be 75 which would be exactly what the Sender wanted. (25 is the smallest amount sent by the Sender which is capable of generating equal payoffs for the Sender and Returner. Larger amounts could also be sent which, if the proper amount was returned would also generate equal payoffs, but they subject the Sender to greater risks.) Hence, if a Returner assumed that his opponent was an egalitarian, returning nothing would not be evidence of selfish behavior but could very well be a Pareto optimal action. In fact, returning part of the 75 received by the Returner would, ironically, diminish the satisfaction of his egalitarian cohort Sender.

If we were to posit some identifying assumptions we might be able to estimate or calculate the implied coefficients of a Fehr-Schmidt (1999) or Bolton-Ockenfels (2000) utility function for each type of subject. (See also Kockesen and Ok (2000) for an analysis of negatively dependent preferences). We have decided not to pursue that direction since doing so might be quite complex, as there are a number of people whose utility a Returner might care about, i.e. his contemporaneous Sender, his parental Returner, his opponent's parent etc., and such empathetic preference might be hard to map. Still, it should be clear that fairness considerations cannot be ruled out in analyzing our data while, in fact, selfish behavior on the part of Returners can be. Senders could well be selfish, however, since they may well be hoping to get back more than they sent.

### *5.3.2. Reciprocity*

Selfish maximizing behavior requires people to maximize their own monetary payoff without regard for the payoffs or actions of others. Alternatively, scholars have come to think that behavior in trust situations is determined by judging the intentions of others and responding or reciprocating to "kind actions" in a positive manner and to "nasty" actions in kind. (see Rabin (1993)).

When it comes to looking for evidence of reciprocal behavior, we are fortunate in having data on expectations. More precisely, remember that reciprocal behavior posits that when someone does something kind (unkind) you are likely to respond likewise. For our purposes here we will define a kind act by a Sender as one of sending more than the Returner expected to receive.[6] Let DAS denote the difference between the amount sent and the expected amount

[5] Note that if people care about the payoffs of people two generations removed from them, i.e. their "grandchildren", then they might still be selfish maximizers and tell their children to return positive amounts hoping that they would then suggest returning nothing to their children.

[6] Note that given our belief elicitation procedure, we are able to define a quantitative measure of kindness which is less arbitrary than that offerd by Rabin (1993).

**Table 6** Relationship between Amount Returned and difference in Amount Sent (AS) and Expected Amount Sent (EAS)

| Dependent Variable: AR | Coeff. | $t$-stat | Prob. $> \|t\|$ |
|---|---|---|---|
| AS-EAS (DAS) | 0.94 | 7.02 | .00 |
| Advice-only Dummy (AO) | 3.75 | .85 | .40 |
| History-only Dummy (HO) | 3.30 | 0.70 | .48 |
| (AO) × (DAS) | −.09 | −0.44 | .66 |
| (HO) × (DAS) | −.19 | −1.07 | .29 |
| Constant | 14.86 | 4.79 | .00 |
| $N$ | 225 | | |
| Model $F(5, 219)$ (Prob. $> F$) | 24.88(.00) | | |
| $R$-squared | 0.36 | | |

sent to a Returner in period t. When a Returner receives an amount greater than his expectations we would expect more to be sent back than otherwise. The results of a linear regression of the amount returned on the difference between the amount expected and actually sent are reported in Table 6.

As we see from these regressions, it appears that subject Returners behaved reciprocally in the sense that the coefficient in front of the DAS variable is positive and significantly different from zero. This indicates that when subjects received amounts greater than their expected amount, i.e. when their Sender acted kindly, they returned more as a result. It is interesting to note, however, that the mean amount expected by the Returners was 27.54, 23.33, and 27.80 for the Baseline, Advice-Only and History-Only treatments which are amounts strikingly close to the equal share amount of 25 (the equal-payoff threshold amount). On average when this amount is sent, Returners tended to send something back instead of keeping it all, but as more and more was sent they responded positively although never with a coefficient greater than 1.

## 6. Do increases in trust lead to trustworthiness or does trustworthiness call forth increased trust?

In our final section we try to disentangle a simple riddle: do increases in trust in society lead people to be more trustworthy or does the proven trustworthiness of people lead others to increase their willingness to trust. Put differently, is a person more willing to trust others when they see that others have acted in a trustworthy (reciprocating) manner in the past, or does the causality work the other way, i.e., do increases in the level of trust in the past lead others to be trustworthy today. Note that this question is well suited to being answered by our intergenerational setup.

The idea of trust leading to trustworthiness is, to start, just simple reciprocation from one person to another. The idea of trustworthiness leading to an increase in trust relies more heavily on there being a convention of reciprocation "in the air". That is, for one to be willing to send even more than was sent in the past, one has to believe that this trust will be reciprocated. In our experiment, reciprocation can occur intragenerationally, but the converse–an increase in trust due to past reciprocation–can only occur intergenerationally.

We define an increase in trust as an *increase* from one generation to the next in the amount sent. We further define trustworthiness as the result that a Returner returns at least as much as what the Sender expected (i.e., the expectations of Senders are fulfilled). We then can ask

**Table 7** Relationship between trust and trustworthiness: Does trustworthiness increase trust?

| Marginal effects on the probability that amount sent increases | | | |
|---|---|---|---|
| Variable | dy/dx | $z$-stat. | Prob. > $z$ |
| Past trustworthiness | .31 | 3.85 | .00 |
| Advice-only dummy | .12 | 1.51 | .13 |
| History-only dummy | .07 | 0.81 | .42 |
| $N$ | 222 | | |
| Model Chi.Sq. (Prob. > Chi. Sq.) | 17.55 | (.00) | |
| Log likelihood | −141.48 | | |
| Pseudo-$R$-square | .06 | | |

whether the probability of an increase in trust is influenced by trustworthiness in the past, and vice-versa. A significant (and, in our interpretation, positive) estimated coefficient on the independent variable is interpreted as "causing" the dependent variable, in the sense that the former increases the probability of the latter.

The results of this exercise are that there is causation from trustworthiness to trust, but not the other way around. Specifically, if in the last generation the Returner returned at least as much as the Sender in the current generation would have expected (applying his or her beliefs about the behavior of Returners to the amount sent in the previous generation), the probability that more is sent this generation than last is .31 higher.

On the other hand, if the Sender in this generation sends more than the Sender last period, this does not increase the probability that the Returner in this generation will reciprocate with at least as much as the Sender expected.

This suggests that it is possible for past reciprocity to lead to increased trust intergenerationally. On the other hand, increased trust in the current generation tends not to engender reciprocity in the current generation. Table 8 shows that increasing the amount sent in the current generation over what was sent in the previous generation decreases the probability that expectations are met for the Sender in the current generation. Thus, it is not hard to see why the "conventional" level of trust shown by subjects (Senders) in our experiment is so low, on average. It is not sufficient that one be trusting today to engender more trust tomorrow. One must be trusting today, and others today must be trustworthy, in order for

**Table 8** Relationship between trustworthiness and trust: Does increased trust increase the probability of trustworthiness?

| Marginal effects on prob (Amount Returned > Expected Amount Returned) | | | |
|---|---|---|---|
| Variable | dy/dx | $z$-stat. | Prob. > $z$ |
| Increase in trust | −.15 | −2.59 | .01 |
| Advice-only dummy | .02 | 0.27 | .79 |
| History-only dummy | −.18 | −2.79 | .01 |
| $N$ | 222 | | |
| Model Chi.Sq. (Prob. > Chi. Sq.) | 14.52 | (.00) | |
| LL | −125.16 | | |
| Pseudo-$R$-square | .05 | | |

trust tomorrow to increase. Of course, one may be trustworthy even in the face of little trust, but the "target" one must hit to be considered trustworthy is not too difficult to achieve as long as expectations are low.

## 7. Conclusions

It was our belief that how one behaves when one faces a situation where trust is involved is influenced by what advice you are given by those who have interacted in that situation before you. What we find is that this conjecture is true in the sense that behavior in intergenerational trust games with advice is quite different than behavior in such games played without it.

More precisely, using the Berg et al. (1995) trust game as our experimental decision problem, advice seems to decrease the amount of trust that evolves when this game is played in an intergenerational manner in that it decreases the amount of money sent from Senders to Returners. Ironically, advice increases trustworthiness in that Returners tend to send more back. We explain this contradiction by appealing to the asymmetrical impact that advice has when serving as an anchor from which Sending and Returning behavior is adjusted. However, in no case, on average, does it pay to send any money. Further, we have discovered that subjects appear to follows conventions of reciprocity in that they tend to Send more if they think the Returners acted in a "kind" manner where kind means that the Sender sent more money than the Returner expected. Finally, while we find a causal relationship between trustworthiness and trust, the opposite cannot be established.

We feel that many of the questions addressed here can most easily (or only) be addressed using our intergenerational game setup. The use of advice as a behavioral explanatory variable is unique to this work, as is the elicitation of beliefs. Providing subjects with histories of past interactions is also a feature of this experiment that has not been used in previous static trust game experiments. Our History-only game is comparable to Berg, et al (1995) in terms of the level of giving by Senders, though not in return behavior. It should be noted that our game is not strictly comparable to Berg et al., in that our Senders and Returners had asymmetric initial endowments (100 and 0 francs, respectively) while their's had symmetric initial endowments of $10 each. Our advice games are dramatically different from the History-only game.

## References

Berg, J., Dickhaut J., & McCabe, K. (1995). Trust, reciprocity, and social history, *Games and Economic Behavior. 10*, 122–142

Bolton, G., & Ockenfels, Axel. (2000) ERC: A theory of equity, reciprocity, and competition. *American Economic Review, 90*(1), 166–93

Chaudhuri, A., Schotter, A., & Sopher, B. (2002). Talking ourselves to efficiency: Coordination in intergenerational minimum games with private. *Almost Common, and Common Knowledge of Advice*, March 2002.

Fehr, E., & Schmidt. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*(3), 817–868

Isaac, R. Mark, & Plott Charles, R. (1981). Price controls and the behavior of auction markets: An experimental examination. *American Economic Review, 71*(3), 448–459

Kockesen, L., & Ok, E. (2000). Negatively inter-dependent preferences. *Social Choice and Welfare, 3*, 533–558.

Nyarko, Y., & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica, 70*(3), 971–1005.

Rabin, M. (1993). Incorporating fairness into game theory and economics, *American Economic Review, 83*, 1281–1302

Schotter, A., & Sopher, B. (2003a) Social learning and coordination conventions in intergenerational games: An experimental study. *Journal of Political Economy, 111*(3), 498–529

Schotter, A., & Sopher, B. (2003b). Advice and behavior in intergenerational ultimatum games: An experimental approach, *Games and Economic Behaviour*, fourthcoming

Tobin, J. (1958) Estimation of relationships for limited dependent variables. *Econometrica, 26*, 24–36.