



ELSEVIER

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/econbase



To deceive or not to deceive: The effect of deception on behavior in future laboratory experiments

Julian Jamison^a, Dean Karlan^b, Laura Schechter^{c,*}

^a Brain and Creativity Institute, University of Southern California 3620 McClintock Ave, Los Angeles, CA 90089, United States

^b Department of Economics, Yale University, 27 Hillhouse Avenue, New Haven, CT 06520, United States

^c Department of Agricultural and Applied Economics, University of Wisconsin at Madison, Taylor Hall, 427 Lorch Street, Madison, WI 53706, United States

ARTICLE INFO

Article history:

Received 17 January 2008

Received in revised form 6 September 2008

Accepted 7 September 2008

Available online 17 September 2008

JEL classification:

A12

C81

C90

Keywords:

Laboratory experimental methods

Experimental economics

Deception

Psychology and economics

Laboratory selection effects

ABSTRACT

Experimental economists believe (and enforce the idea) that researchers should not employ deception in the design of experiments. This rule exists in order to protect a public good: the ability of other researchers to conduct experiments and to have participants trust their instructions to be an accurate representation of the game being played. Yet other social sciences, particularly psychology, do not maintain such a rule. We examine whether such a public goods problem exists by purposefully deceiving some participants in one study, informing them of this fact, and then examining whether the deceived participants behave differently in a subsequent study. We find significant differences in the selection of individuals who return to play after being deceived as well as (to a lesser extent) the behavior in the subsequent games, thus providing qualified support for the proscription of deception. We discuss policy implications for the maintenance of separate participant pools.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

In two of the original experimental economics textbooks, Davis and Holt (1993) and Friedman and Sunder (1994), among others, proscribe the use of deception in experiments. The primary concern with deception is that many experimental laboratories use a common pool of participants. Thus, a public goods problem exists in which experiencing deception in one experiment may cause participants to react differently (and uncertainly) in future games with other researchers. Clearly, maintaining this “public good” involves trade-offs between benefits to the individual (ability to conduct experiments that require deception) and the group (maintaining a subject pool that is trained to believe that experiment instructions are truthful). As some research questions may be better answered by using deception, should we forego the knowledge that could be acquired through such experiments in order to maintain a common pool of deception-free participants? This concern warrants testing, and in this paper we determine the presence and extent of such sample contamination in a particular setting.

Although we focus here on the particular issue of deception, the overall methodological question of understanding what we mean by “control” in the laboratory is important more generally. Levitt and List (2007) discuss the tradeoffs between the

* Corresponding author.

E-mail address: lschechter@wisc.edu (L. Schechter).

laboratory and the field, specifically focusing on how the manner in which individuals are selected and watched may alter their behavior. Deception is just one example of how a subject pool may be altered that may influence interpretation of later experiments. If a subject pool is entirely confident that the environment is deception-free, this may heighten “control” at the expense of reality. On the other hand, in the field one may have participants who are suspicious of deception in unknown ways (thus exhibiting less control, but more reality). This paper does not compare the field to the laboratory with respect to behavior in economic experiments. Instead we examine an important selection effect: does a singular experience in the laboratory affect the likelihood that different individuals return in a future month for another experiment run by a different researcher?

Bonetti (1998) writes the first, and perhaps only, article arguing that deception should be allowed in experimental economics. He reviews the evidence from experiments in psychology and concludes that deception has a minimal effect on behavior. Two immediate, mainly philosophical, replies to Bonetti’s argument are Hey (1998) and McDaniel and Starmer (1998). Later, Ortmann and Hertwig (2002) conducted a thorough review of the evidence from experiments in psychology and concluded that the experience of deception does affect participants’ expectations, suspicions, and future behavior.

Here we give a brief overview of the different strands of the psychology literature looking either at the effects of suspicion of deception or the effects of past experiences with deception itself on different outcomes. A reader interested in more details should consult the review by Ortmann and Hertwig (2002). There is a rather large strand of literature in psychology looking at differences in behavior based on whether a participant admits to being suspicious of deception. Stricker et al. (1967), one of the seminal papers on this topic, find that subjects who admitted to suspecting deception in the experiment in which they participated conformed less (i.e. were less likely to agree to unreasonable propositions simply because other subjects agreed) than those who claimed not to suspect anything. The converse of this line of study is that which looks at whether subjects who have been deceived claim they will be suspicious of information given to them by experimenters in the future. Both Epley and Huff (1998) and Krupat and Garonzik (1994) find that deceived subjects anticipate being more suspicious in the future. There is a third line of literature looking at the effects on behavior of being warned before an experiment by the experimenter that the experiment “might” include misinformation or of being informed by a confederate before the game that it actually does involve deception. This type of forewarning tends not to lead to changes in behavior (see Wiener and Erker, 1986 for the first type of warning and Golding and Lichtenstein, 1970 for the second), although there is some mixed evidence.

The three areas of literature discussed above do not look at the direct effects of experiencing deception in one experiment on behavior in future experiments. These consequences are the principal motivation for the prohibition of deception in experimental economics and are the focus of this paper. There do exist a few experiments in psychology that attempt exactly this type of test, albeit with non-economic experiments such as personality tests and memory games.¹ Some of these papers find that experiencing deception leads to changes in behavior (Silverman et al., 1970; Christensen, 1977) while others do not (Fillenbaum, 1966). Some find that past experiences with deception lead to changes in behavior when the first and second set of experiments are noticeably similar, but not when the experiments are dissimilar (Brock and Becker, 1966; Cook et al., 1970).

There are many differences between the above psychological literature measuring the effects of experienced deception on play in future experiments and the setup in this paper. Some of the psychological experiments did not give the control group (those who were not deceived) any prior experimental experience (Silverman et al.). Other experiments subjected the control group to a very different experience than that experienced by the deceived (Christensen). In contrast, in this paper, the initial treatment was almost identical for those who were deceived and those who were not.

Additionally, in some of the psychological experiments, the first and second treatments were conducted immediately sequentially. After the participants took part in the first experiment, a different researcher walked into the room and claimed that he really needed more subjects and wondered if they might have time to participate in one more experiment being run next door (Brock and Becker, 1966; Cook et al., 1970; Christensen, 1977). This is not the conventional method of subject recruitment. In the current paper, the initial and secondary experiments were separated by approximately 3 weeks, and the recruitment mechanism for both experiments used e-mail, a conventional approach.

In almost all of these psychological experiments, the deception involved lying to the subjects about the purpose of the experiment. For example, subjects were given two poems by Robert Frost, were told that one was by Frost and one by a high school English teacher, and were then asked to rate the merits of the two poems. After the experiment they were debriefed and told that both poems were actually by Frost and that the experiment was looking at how beliefs regarding authorship affected the rating of the poems. In another experiment participants were told to try to write down as many phone numbers as they could as quickly as possible. After doing so, they were told that excellence in writing phone numbers quickly was a sign of obsessive compulsive disorder. Then they were told to try again. After the experiment, they were told that excellence in writing phone numbers was not actually a sign of obsessive compulsion; rather the experimenter wanted to see whether the participant would write fewer phone numbers after hearing that misinformation.

Experimental economists tend not to implement this particular form of deception. One of the more common forms of deception used in economic experiments is deceiving the player with regards to the partner with whom he is matched.

¹ Ortmann and Hertwig (2002) discuss in detail nine papers that look at the effects of past deception on behavior in future experiments in Section 3.2.2 of their paper. Here we focus on five of those studies.

Weimann (1994), Blount (1995), Scharlemann et al. (2001), Sanfey et al. (2003), and Winter and Zamir (2005) all told players that their partners are humans, when in fact they are computers playing either predetermined or random strategies. Kim and Walker (1984) told subjects that they were playing a public goods game with many more individuals than was in fact the case, using a stochastic variable to approximate the contributions from the nonexistent players. Ball et al. (2001) and Gibbons and Van Boven (2001) tell players that their partner is of a certain intelligence level or personality type. Although their partner is a human, not a computer, the type is generated artificially. This may not, as implemented, constitute deception per se, but it is similar to the type of experiment which does constitute unambiguous deception. These experiments allow the researcher to test how subjects react to specific situations that may not arise naturally with high probability.

Different types of deceit may lead to different contamination of subjects. In this paper, we focus on one specific type of deception, and the results may not be generalizable to all types of deception. Specifically, we deceived subjects as to the identity of their partner. We chose this type of deception because, from both conversations with experimental economists and the small literature in experimental economics using deception (mentioned in the previous paragraph), this seems to be the type of deception that economists are primarily interested in. We randomize subjects into a treatment group that experiences deception regarding the identity of their partner and a control group that participates in a similar experiment but with no deception. Then we measure the effects of deception on the sample selection of those who return to participate in further experiments, as well as on subjects' behavior in those future experiments.

2. Experimental procedures

2.1. Overall setup

The Xlab (experimental social science laboratory) of the University of California at Berkeley maintains two participant lists, consisting primarily of undergraduates but also a few graduate students as well as some staff members. One participant list does not maintain the no-deception rule; it is less active but is sometimes used by behavioralists in the business school.² The second participant list is used primarily by economists and maintains the no-deception rule. Occasionally, as in the present experiment, researchers are allowed to deceive subjects from the latter pool, who are then moved to the former.

In all cases, potential subjects are recruited using a variety of methods, including mass emails (to class lists or student organizations), flyers posted around campus, and booths at student and staff activity fairs. Subjects must actively choose to join the list, which then makes them eligible to receive announcements regarding (and to subsequently sign up for) specific experimental sessions. There are approximately six to ten unique experiments and concomitant e-mail announcements in a given month.

Our first round of experiments involved 261 subjects across ten sessions, each of which lasted for approximately 1 h. Five of these sessions involved deception (treatment) while five did not (control). The three initial sessions were necessarily control so that we could use that data to program the computer players, but after that sessions were randomized as to treatment or control. Subjects signed a consent form and were given a written instruction sheet (see Appendix A), which was then read aloud by the researcher. The experiment itself was conducted on laptop computers implementing a z-Tree (Zurich Toolbox for Readymade Economic Experiments) program. When all subjects had completed the experiment, payoffs were determined and subjects waited approximately 15 min for individual checks to be filled out and distributed.

In the non-deception treatment (five sessions, 132 subjects), subjects were randomly assigned to roles and matched with one another, and payoffs were determined according to actual play. (The actual games are described in detail below.) In the deception treatment (five sessions, 129 subjects), subjects were randomly assigned to roles, but in all cases their opponents' actions were determined by a computer program that simulated the human play.^{3,4} During the check-processing wait time, these subjects were informed that they had been deceived and that they had actually been partnered with a computer rather than with another human. They were told that this was necessary for the research (but no more specific details) and were asked to sign a second consent form allowing their data to be used. All subjects signed this consent form.

Two to three weeks later, all subjects from the first round were sent a recruitment email for new experimental sessions using the name of a researcher different from that in the first round. This email was identical to the standard Xlab recruitment email, except that it promised \$10 above and beyond the normal earnings in order to facilitate sufficient return rates. Subjects did not know that only they had received this particular email, although it is possible that some of them checked with friends and noticed that it had not gone out to the entire pool. However, that in itself is not an uncommon occurrence given various screening criteria used by other researchers at the lab.

In all, 142 people returned for one of the eight sessions that took place 3–4 weeks after the sessions in round one. These lasted slightly under an hour, and each consisted of a mixture of both deceived and non-deceived subjects. A different

² The psychology department maintains a separate pool, although it is entirely possible for a subject to be on both the Xlab list and the psychology list.

³ We explain this in more detail in the next section.

⁴ If we were only interested in the results from the first round of experiments, it may have been possible to use a technique suggested by Bardsley (2000) that avoids deception while at the same time maintains a setup in which players who think their partner is a human are actually partnered with a computer. Using this technique, though, would not allow us to measure the effects of being deceived on return rates and play in subsequent experiments.

researcher than in the first round was physically present in the room for these sessions. Subjects signed a standard consent form, were given the instructions, and then completed the experiment as three separate interactions on the VeconLab website.

Afterward, as they waited to be paid, they were informed that these sessions were in fact a continuation of the previous experiment, and they were asked to sign another consent form allowing use of their data. For those from the deception treatment, this was thus the fourth informed consent form that they signed in the course of the full protocol. None of the participants refused to give their consent at any point throughout the experiments.

2.2. Specific games

In our first set of sessions, we ran a trust game (Berg et al., 1995) with a \$20 endowment to the trustor.⁵ The trustor chose an amount (the ‘investment’) to send to the trustee, and this amount was then tripled. The trustee then returned as much or as little as desired. There were two practice rounds and four actual rounds. Anonymous partners were randomly chosen and maintained for all six rounds (and, like everything else above, subjects were told that this would be the case), but individuals were not told the identity of their partner at any point. No communication in the room was allowed.

One of the four actual rounds was randomly chosen to determine payoffs. All subjects were given an additional \$5 to make sure that no one received a final payout of \$0, so the potential individual range (for both roles) was \$5–65. In fact, players in both roles sometimes received \$5 (i.e. \$0 in the game itself), and some trustees received \$65.

We ran three non-deception sessions first and used those to program the computer play in the deception sessions in order for it to match the human play as closely as possible. Since these were repeated games, we needed to try to mimic the entire strategy, not simply the observed play-paths. Obviously this was imperfect, but in the end we categorized trustors into five types and trustees into three types. A small percentage of trustors never invested more than \$5 of their endowment in any round, while a significant fraction invested \$20 every time. We also included a “trigger strategy” for all types of trustors (as observed in the real subjects): if the trustee ever sent back less than what was invested, the trustor never again invested anything. Section 3 outlines the summary statistics showing that the actual game play by participants in the deception and non-deception treatments is not significantly different.

In our second set of sessions, which occurred 3–4 weeks after the first session, we ran three different games with each subject. All games were run with subjects simultaneously connected to the VeconLab website at the University of Virginia. The three decision problems were run as independent sequential online sessions since the type of game was different in each case. All instructions were provided on the website under the game selection categories of Bargaining and Fairness/Bargaining Games, Individual Decision Problems/Lottery Choice Menu, and Game Theory Experiments/ 2×2 Matrix Games, respectively.

The first game was a dictator game with an endowment of \$20 (i.e. subjects chose how much of this amount to give away to an anonymous partner). All subjects played the role of sender (and therefore, in effect, also the role of potential receiver), but for payoffs only half of the matchings were consummated, so that each subject ended up as either sender or receiver but not both.

The second “game” was a series of gambles (Holt and Laury, 2002), with 10 ordered binary choices between two lotteries. The exact payoffs are shown in Appendix Table A. The first choice was between a (safe) lottery that paid \$11 with 10 percent chance and \$8.80 with 90 percent chance and a (risky) lottery that paid \$21.20 with 10 percent chance and \$0.55 with 90 percent chance. Here the first lottery was risk-dominant and also had a significantly higher expected value. As the choices progressed, the probability of the higher payoff in each lottery increased by 10 percent until the final choice was between \$11 (with certainty) versus \$21.20 (with certainty), so that the second lottery dominated in all respects.

In the third game, subjects played a prisoner’s dilemma with payoffs $(C,C) = (10,10)$, $(D,D) = (6,6)$, and $(D,C) = (15,1)$. They were randomly assigned to be the row or column player and were randomly matched with one another. No deception was involved for any of the games in these sessions. Finally, one of the three games was randomly chosen to determine payoffs. (In the risk game, which involved multiple decisions, one decision was randomly chosen.) We also paid each player \$10 in addition to the earnings from the games (as promised in the recruitment email).

These three games were chosen in part because they are so common in the field and laboratory and in part because previous research has suggested that subjects behave differently in the dictator game and in the prisoners’ dilemma when they think their partner is a human versus when they think their partner is a computer. Our hypothesis was that players who had been deceived in the past as to the identity of their partner would be suspicious in subsequent experiments that their partner was actually a computer, even when we told them they were playing with a human.

There is evidence that, in the prisoners’ dilemma, players cooperate more often when they are playing with another human than they do when playing with a computer (Abic and Kahan, 1972; Kiesler et al., 1996; Rilling et al., 2004). Thus, we hypothesized that those players who had been deceived in the past would suspect they might be playing against a computer and would thus cooperate less often in the prisoners’ dilemma. Frohlich et al. (2001) find that players in a dictator game who suspect that they were not paired with real people give less. Thus, we also hypothesized that players who had been deceived

⁵ Despite evidence suggesting that the first player in the Trust game is not merely acting due to “trust” but also due to a desire to take risks (Karlan, 2005; Schechter, 2006), we use the canonical terminology here of “trustor” and “trustee.”

Table 1

Selection effects in round two (based on player characteristics) mean, standard errors and number of observations.

	Total	Not deceived	Deceived	<i>t</i> -Stat: diff \neq 0
Proportion returned to the lab in round two	0.547 (0.031) 261	0.583 (0.043) 132	0.512 (0.044) 129	1.162
Proportion of males who returned to the lab in round two	0.558 (0.049) 104	0.482 (0.067) 56	0.646 (0.071) 48	−1.682*
Proportion of females who returned to the lab in round two	0.541 (0.040) 157	0.658 (0.055) 76	0.432 (0.055) 81	2.895***
Proportion of “High Passers” (Player A) who returned to the lab in round two	0.490 (0.050) 102	0.471 (0.071) 51	0.510 (0.071) 51	−0.393
Proportion of “Low Passers” (Player A) who returned to the lab in round two	0.724 (0.084)	0.800 (0.107) 34	0.643 (0.133) 47	0.927
Proportion of “High Passers” (Player B) who returned to the lab in round two	0.533 (0.075) 45	0.500 (0.096) 28	0.588 (0.123) 17	−0.564
Proportion of “Low Passers” (Player B) who returned to the lab in round two	0.538 (0.062) 65	0.655 (0.090) 29	0.444 (0.084) 36	1.706*

Standard errors are in parentheses. * significant at 10 percent; ** significant at 5 percent; *** significant at 1 percent. A player is a “High Passer” if they passed at least half of what they could pass in the Trust game, and a “Low Passer” otherwise.

would keep more as dictator in the dictator game. We included the Holt–Laury sequence of risky lotteries as a control, our hypothesis being that those who were deceived would not play any differently than those who were not deceived, as it did not involve a partner (unless they believed that part of the deception was our lying about the results of random draws).

3. Results

As there were 261 participants who participated in the first round and 142 in the second round, our sample size is somewhat small when compared with some studies of selection effects. This may make it more difficult to find statistically significant results. On the other hand, adding more participants would have meant deceiving even more subjects. Before conducting this research we did not know how large an effect deception might have on the participants and were thus conservative in our recruiting. A second potential drawback of the data is that in the Xlab no demographic information was universally collected other than sex and the number of experiments in which the participant had previously participated.

Before looking to see if being deceived in a previous experiment affects play in a subsequent experiment, we would like to make sure that both the randomization and the deception were successful. Remember, the first round of our experiment involved six rounds of a trust game. The senders who were deceived (not deceived) sent on average \$14.11 (\$12.88) in the first round, \$11.68 (\$9.95) in the last round, and an average of \$14.59 (\$13.33) in all six rounds. Using a *t*-test with equal variance, the *t*-values are 1.041, 1.160, and 1.281, respectively. The receivers who were deceived (not deceived) returned 0.435 (0.438) of what they had received in the first round, 0.186 (0.174) in the last round, and on average 0.413 (0.407) in all rounds. The *t*-values for these tests are 0.095, 0.261, and 0.318, respectively. Males make up 37 percent of the deceived population and 42 percent of the non-deceived population, yielding a *t*-value of −0.858. Since none of these differences is significant the randomization appears to have been successful.⁶

3.1. Effect of deception on selection into subsequent experiment

The next step in our analysis involves looking at the effect of deception on return rates. As described above, all subjects were invited to participate in another round of experiments approximately 3 weeks after the first round. Subjects were not led to believe that this subsequent experiment was in any way related to the previous experiments. Only 55 percent of the original participants returned for the second round of experiments.

Table 1 shows that 51 percent of those who were deceived returned while 58 percent of those who were not deceived returned and that this difference is not statistically significant. While the subjects were not living in a vacuum during the 3 weeks between the first and second set of experiments, because of the randomization we only need to assume that these two groups had, on average, similar experiences in the time between the two experiments.

Interestingly, gender differences played a significant role in our analysis: females who have been deceived are significantly less likely to return than females in the control group. On the other hand, males who were deceived are significantly more likely to return than males in the control group. In results not shown here, we tested for return status based on the number of previous experiments at the Xlab in which the player participated. Although participants who have participated in other experiments in the past are significantly more likely to return in general, the deception effect does not differ for the more or less experienced.

⁶ Psychologists believe there may be “faithful subjects” who attempt to take the experiment at face value, even if he suspects deception. Hence, the fact that the deceived and non-deceived play no differently in the initial experiment could be a sign that the deceived do not suspect the deception or that they are merely being “faithful.”

Table 2

Selection effects in round two (based on player experience) mean, standard errors and number of observations.

	Total	Not deceived	Deceived	t-Stat: diff \neq 0
Proportion of “High Receivers” (Player A) who returned to the lab in round two	0.500 (0.063) 64	0.464 (0.096) 28	0.528 (0.084) 36	−0.497
Proportion of “Low Receivers” (Player A) who returned to the lab in round two	0.556 (0.068) 54	0.586 (0.093) 29	0.520 (0.102) 25	0.480
Proportion of “High Receivers” (Player B) who returned to the lab in round two	0.505 (0.052) 93	0.549 (0.070) 51	0.452 (0.078) 42	0.922
Proportion of “Low Receivers” (Player B) who returned to the lab in round two	0.676 (0.078) 37	0.867 (0.091) 15	0.545 (0.109) 22	2.117**
Proportion of players “whose actual payout was not determined by last round” who returned to the lab in round two	0.565 (0.037) 184	0.568 (0.051) 95	0.562 (0.053) 89	0.090
Proportion of players “whose actual payout determined by last round” who returned to the lab in round two	0.506 (0.057) 77	0.622 (0.081) 37	0.400 (0.078) 40	1.967*

Standard errors are in parentheses. * significant at 10 percent; ** significant at 5 percent; *** significant at 1 percent. Player A is referred to as a “High Receiver” if Player B sent back at least as half as much as he received in the Trust game. Player B is referred to as a “High Receiver” if their Player A passed at least half of what they could pass. The final two rows divide the sample into those whose payout was determined by the last round, and those whose payout was determined by any other round. Many passed zero back in the final round, since the final round was pre-announced. Thus, those whose last round was chosen as the actual round earned on average less money.

The next few rows of [Table 1](#) examine further selection effects. We look at return rates dividing the sample into four categories: “High Passers” and “Low Passers,” and Player As (trustors) and Player Bs (trustees).⁷ We find no significant differences within three of these four groups. Only “Low Passer, Player B” types seemed to be influenced by the deception in their decision to return. One possible explanation for this is that after the untrustworthy (i.e. the “Low Passers”) found out that they were deceived, they felt more self-conscious about their behavior.

One could worry that, in making multiple comparisons, a few of them are bound to appear significant when in fact there are no true differences. We employ a Bonferonni correction for multiple significance tests in order to correct the significance of independent comparisons and derive conservative estimates for these non-independent comparisons. Given our correlated outcomes, the Bonferonni correction may hold each test to an unreasonably high standard and increase the probability of a Type II error. Still, it is interesting to see which, if any, of our comparisons will stand up to the Bonferonni correction. In [Table 1](#), only one comparison remains significant (at the 5 percent significance level) after making the Bonferonni correction, and that is the fact that deceived females are less likely to return to participate in the second set of experiments.

In [Table 2](#), we look at selection effects not based on how the player himself played, but rather on how his partner treated him or on the round of the trust game that was randomly chosen to count for payoffs. When we compare return rates conditional on the play of the partner we find that, for the sub-sample of Player Bs who were sent very little by Player A, being deceived made them less likely to return. Thus, it seems that deception most influenced the individual’s decision to return to the laboratory when he also was not lucky in the game. (Recall that Player B had no endowment. If Player A sent him very little, then he went home with very little.) Clearly, then, it was not merely the experience of deception that altered the future decision, but rather deception coupled with low winnings.

This is further reinforced by the final two rows of [Table 2](#). Here we examine the effect of the deception treatment on those whose payoffs were determined by the final round of the Trust game. We had randomly chosen which round of the Trust game determined payoffs. Earnings in the final round, presumably due to backwards induction, were significantly lower than in the previous rounds. There is no difference in overall return rates for those players for whom the last round determined their earnings as opposed to those for whom an earlier round determined earnings. We find that players who were deceived and for whom the last round of the trust game counted were much less likely to return. (We should note that none of the comparisons in [Table 2](#) stand up to the Bonferonni correction.)

Summarizing [Table 2](#), being deceived made individuals less likely to return only when they were unlucky insofar as having received a low payout. Perhaps these players felt that not only did we deceive them in terms of the identity of their partner, but that we may have also deceived them either in terms of purposely choosing to pay them for the round with the lowest payoffs or by programming the “trustor” computer to send them very little. (This, of course, was not the case.) One could argue that it might not be the experience of deception *per se* that encouraged these ‘unlucky’ players not to return; they may have felt that the procedures in the experiment were unjust and not have returned for that reason. For more on this, refer to [Thibaut and Walker \(1975\)](#), whose seminal work on procedural justice finds that, holding the outcome constant, satisfaction depends upon the process by which outcomes are reached.

⁷ A player is a “High Passer” if he passed at least half of what he could pass, and a “Low Passer” otherwise. When categorizing players as high or low passers we ignore play in the last round of the trust game. Since many Player As passed no money in the last round, making it impossible to calculate the share returned by Player B in that round, this also allows us to include more observations.

3.2. Effect of deception on play in a subsequent experiment

Next, we analyze whether deception has an effect on play in future experiments. We preface these results with the caveat that we cannot differentiate in practice between any indirect selection effects and direct deception effects. If all participants came back to play in the second round of experiments, we could claim that any differences in play were due to the direct effect of having experienced deception. Since only 55 percent of the participants came back for the second round, any differences in play in the second set of experiments may be due to a change in behavior at the individual level after experiencing deception: a player who was deceived once begins to second-guess the experimenter and uses a different strategy than he would otherwise. However, differences in behavior may also be due to the selection effect: a player who was deceived once may not play differently in future experiments, but only certain types of players may decide to return after having been deceived. Although we cannot separate the indirect and direct effects, the combination of the two makes up the true effect on subsequent experimental outcomes, which is the practical outcome of interest.

Table 3 presents these results. We measure four outcomes: (1) the amount kept in the dictator game, (2) risk aversion from the Holt-Laurry gambles (if consistent), (3) a binary variable if responses to the risk aversion questions were inconsistent, and (4) defection in the prisoners' dilemma. We analyze the results for the sample as a whole as well as for (a) inexperienced versus experienced participants, (b) males versus females, and (c) those who took on the role of Player A versus Player B in the first round.

Lastly, we make an attempt at controlling for selection. We saw previously that those players who were deceived and unlucky (either because their payoffs were decided by the last round of the trust game or because they were a trustee who received very little from their trustor) were less likely to return than those who were not deceived and unlucky. By limiting our analysis to only 'lucky' players, we can look specifically at the group of players exhibiting less selection.

The top left quadrant of Table 3 shows the primary results for the full sample. We show via a *t*-test that mean behavior of the deceived is different from the non-deceived. We also carry out Levene's test comparing the variance of behavior of the deceived to that of the non-deceived. For the results on variance we look only at the amount kept in the dictator game and at the number of safe choices in the Holt-Laurry risk game. We do not compare the variances of the two binary outcome variables. One might imagine there to be more noise in the actions of the deceived because they may believe that their actions are less likely to map into payoffs.

As can be seen, we find that the deceived individuals were more likely to behave inconsistently in the risky gambles, and we find no significant differences in the other three outcomes. For the Holt-Laurry risk lotteries, we find no differences in risk aversion either in the full sample or in any sub-sample analysis (except for those who were 'lucky'). However, 8 percent of the players did not make consistent choices and were categorized as "switchers," and we do find a difference for these individuals. Remember, the Holt-Laurry experiment involves a series of choices between a risky and safe option. As one moves along in the series, the risky option becomes more and more appealing. A rational player should not switch back and forth between the risky and safe options, but rather switch from the safe option to the risky one at one and only one point.

If we compare the risk aversion of the players who played consistently, there is no difference, as we expected, between those who were deceived and those who were not. The measure of risk-aversion in Table 3 is the number of safe choices, which could hypothetically range from 0 to 10, with higher values representing more risk aversion. Using Levene's test for equality of variance, in the full sample we find that the variance for the deceived is significantly higher than for the non-deceived at the 5 percent level. The variances are also significantly different for all Player Bs at the 10 percent level, for all players and Player Bs who had participated in at least 10 previous experiments (at the 5 percent level), and for all females (at the 10 percent level). The variances of the amount kept in the dictator game are not significantly different in any of the comparisons.

We also find that significantly more of the deceived players are "switchers"; it is possible that they no longer take the research seriously and thus switch back and forth between the risky and safe choices. This result is quite significant among the general population of players, as well as among females and the more inexperienced participants.⁸ Using the Bonferonni correction within panels of Table 3, the only finding that remains significant at the 10 percent level is that deceived females are more likely to play inconsistently in the Holt-Laurry gamble.⁹

In the dictator game, we find no significant effect of previous deception on the amount kept in the pooled sample. We do find, however, that those who played as trustor in the first round and were deceived kept significantly more as dictator than did Player As who were not deceived. Further analysis shows that this result holds for both the inexperienced Player As (i.e. who had participated in fewer than 10 previous experiments in the Xlab) and female Player As. For the prisoners' dilemma, we find no difference in behavior among those who were deceived versus those who were not (except that females were almost significantly more likely to defect if deceived).

⁸ There is also one player who chose all 10 safe lotteries. As the tenth lottery gives 100% probability to the higher of the two payoffs, and the high payoff is higher for the risky choice, no rational player should choose all 10 safe lotteries. The player who did this had also been deceived (although we do not consider this player to be a switcher).

⁹ Note that this effect, albeit significant statistically, is small in terms of the proportion of individuals who answer inconsistently in either the deceived or undeceived pool.

Table 3
Effect of deception in round one on behavior in round two games mean and standard errors.

	Full sample				Player A in round one				Player B in round one			
	Not deceived	Deceived	t-Stat: mean diff \neq 0	Levene test f-stat: S.D. ratio \neq 1	Not deceived	Deceived	t-Stat: diff \neq 0	Levene test f-stat: S.D. ratio \neq 1	Not deceived	Deceived	t-Stat: diff \neq 0	Levene test f-stat: S.D. ratio \neq 1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
<i>Panel A: Full Sample</i>												
Amount kept in dictator game	14.584 (0.484)	14.923 (0.518)	-0.477	1.032	13.472 (0.676)	15.441 (0.697)	-2.029**	0.996	15.561 (0.658)	14.355 (0.771)	1.193	0.963
Risk aversion (Holt-Laury gambles if consistent)	6.067 (0.156)	5.804 (0.236)	0.968	0.586**	6.171 (0.254)	5.867 (0.338)	0.732	0.660	5.975 (0.191)	5.731 (0.331)	0.685	0.514*
Risk aversion (Holt-Laury gambles switchers)	0.026 (0.018)	0.138 (0.043)	-2.537**		0.028 (0.028)	0.118 (0.056)	-1.461		0.024 (0.024)	0.161 (0.067)	-2.117**	
Defected in prisoner's dilemma	0.653 (0.055)	0.677 (0.058)	-0.293		0.611 (0.082)	0.686 (0.080)	-0.651		0.692 (0.075)	0.667 (0.088)	0.223	
Number of observations	77	65			36	34			41	31		
<i>Panel B: participants played at least 10 prior games in laboratory</i>												
Amount kept in dictator game	15.179 (0.687)	15.032 (0.734)	0.146	1.103	13.941 (0.979)	14.882 (0.996)	-0.674	0.967	16.136 (0.922)	15.214 (1.125)	0.630	1.055
Risk aversion (Holt-Laury gambles if consistent)	6.211 (0.233)	5.923 (0.404)	0.659	0.487**	6.411 (0.429)	6.143 (0.592)	0.376	0.639	6.048 (0.244)	5.667 (0.555)	0.724	0.338**
Risk aversion (Holt-Laury gambles switchers)	0.026 (0.026)	0.133 (0.063)	-1.723*		0.000 (0.000)	0.125 (0.085)	-1.510		0.045 (0.045)	0.143 (0.097)	-1.017	
Defected in prisoner's dilemma	0.684 (0.076)	0.700 (0.085)	-0.138		0.824 (0.095)	0.706 (0.114)	0.792		0.571 (0.111)	0.692 (0.133)	-0.689	
Number of observations	39	31			17	17			22	14		
<i>Panel C: participants played fewer than 10 prior games in laboratory</i>												
Amount kept in dictator game	13.974 (0.675)	14.824 (0.740)	-0.850	0.930	13.053 (0.947)	16.000 (0.985)	-2.154**	1.034	14.895 (0.939)	13.647 (1.057)	0.886	0.883
Risk aversion (Holt-Laury gambles if consistent)	5.919 (0.206)	5.700 (0.272)	0.653	0.711	5.944 (0.286)	5.625 (0.375)	0.687	0.652	5.895 (0.305)	5.786 (0.408)	0.219	0.756
Risk aversion (Holt-Laury gambles switchers)	0.026 (0.026)	0.143 (0.060)	-1.828*		0.053 (0.053)	0.111 (0.076)	-0.637		0.000 (0.000)	0.176 (0.095)	-1.961*	
Defected in prisoner's dilemma	0.622 (0.081)	0.657 (0.081)	-0.310		0.421 (0.116)	0.667 (0.114)	-1.504		0.833 (0.090)	0.647 (0.119)	1.252	
Number of observations	38	35			19	18			19	17		
<i>Panel D: males</i>												
Amount kept in dictator game	14.111 (0.902)	14.516 (0.746)	-0.349	1.272	13.818 (1.394)	14.214 (1.065)	-0.230	1.346	14.313 (1.217)	14.765 (1.066)	-0.280	1.227
Risk aversion (Holt-Laury gambles if consistent)	5.703 (0.291)	5.393 (0.339)	0.694	0.714	5.545 (0.511)	5.545 (0.593)	0.000	0.742	5.813 (0.356)	5.294 (0.418)	0.938	0.683
Risk aversion (Holt-Laury gambles switchers)	0.000 (0.000)	0.097 (0.054)	-1.671		0.000 (0.000)	0.214 (0.122)	-1.661		0.000 (0.000)	0.000 (0.000)	NA	
Defected in prisoner's dilemma	0.769 (0.084)	0.581 (0.090)	1.508		0.818 (0.122)	0.643 (0.133)	0.948		0.733 (0.118)	0.529 (0.125)	1.178	
Number of observations	27	31			11	14			16	17		
<i>Panel E: females</i>												
Amount kept in dictator game	14.840 (0.567)	15.294 (0.725)	-0.498	0.902	13.320 (0.774)	16.300 (0.936)	-2.532**	0.942	16.360 (0.723)	13.857 (1.143)	1.942*	0.715
Risk aversion (Holt-Laury gambles if consistent)	6.271 (0.175)	6.214 (0.314)	0.170	0.534*	6.458 (0.276)	6.053 (0.415)	0.842	0.557	6.083 (0.216)	6.556 (0.444)	-1.063	0.632
Risk aversion (Holt-Laury gambles switchers)	0.040 (0.028)	0.176 (0.066)	-2.122**		0.040 (0.040)	0.050 (0.050)	-0.158		0.040 (0.040)	0.357 (0.133)	-2.829***	
Defected in Prisoner's dilemma	0.592 (0.071)	0.765 (0.074)	-1.645		0.520 (0.102)	0.714 (0.101)	-1.342		0.667 (0.098)	0.846 (0.104)	-1.163	
Number of observations	50	34			25	20			25	14		
<i>Panel F: non-'unlucky'</i>												
Amount kept in dictator game	13.867 (0.604)	14.675 (0.649)	-0.912	0.976	13.593 (0.779)	15.043 (0.870)	-1.245	0.942	14.278 (0.976)	14.176 (0.990)	0.073	1.03
Risk aversion (Holt-Laury gambles if consistent)	6.091 (0.213)	5.480 (0.265)	1.803*	0.858	6.192 (0.304)	5.667 (0.388)	1.079	0.888	5.944 (0.286)	5.267 (0.358)	1.499	0.763
Risk aversion (Holt-Laury gambles switchers)	0.022 (0.022)	0.154 (0.059)	-2.214**		0.037 (0.037)	0.182 (0.084)	-1.679*		0.000 (0.000)	0.118 (0.081)	-1.504	

Table 3 (Continued.)

	Full sample				Player A in round one				Player B in round one			
	Not deceived	Deceived	t-Stat: mean diff ≠ 0	Levene test f-stat: S.D. ratio ≠ 1	Not deceived	Deceived	t-Stat: diff ≠ 0	Levene test f-stat: S.D. ratio ≠ 1	Not deceived	Deceived	t-Stat: diff ≠ 0	Levene test f-stat: S.D. ratio ≠ 1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Defected in prisoner's dilemma	0.659 (0.072)	0.700 (0.073)	-0.397		0.630 (0.095)	0.652 (0.102)	-0.162		0.706 (0.114)	0.765 (0.106)	-0.378	
Number of observations	45	40			27	23			18	17		

Standard deviations are in parentheses. * significant at 10 percent; ** significant at 5 percent; *** significant at 1 percent. Some observation numbers are different than those reported. For the dependent variable 'proportion kept in dictator game', the observation numbers change to the following: For Panel C, there are 34 for 'not deceived' in the Full Sample; 17 for 'not deceived' in Player A. For the dependent variable 'Risk Aversion if Consistent', the observation numbers change to the following: For Panel A there are 75 'not deceived' and 56 'deceived' in the Full Sample; 35 for 'not deceived' and 30 for 'deceived' in Player A and 40 for 'not deceived' and 26 for 'deceived' in Player B. For Panel B, there are 38 'not deceived' and 26 'deceived' in the Full Sample; 14 'deceived' in Player A; 21 'not deceived' and 12 'deceived' in Player B. For Panel C, there are 37 'not deceived' and 30 'deceived' in the Full Sample; 18 for 'not deceived' and 16 for 'deceived' in Player A; 14 for 'deceived' in Player B. For Panel D there are 28 'deceived' in the Full Sample; 11 for 'deceived' in Player A. For Panel E there are 48 'not deceived' and 28 'deceived' in the Full Sample; 24 for 'not deceived' and 19 for 'deceived' in Player A; 24 for 'not deceived' and 9 for 'deceived' in Player B. For Panel F there are 44 'not deceived' and 33 'deceived' in the Full Sample; 26 for 'not deceived' and 18 for 'deceived' in Player A; 15 for 'deceived' in Player B. For the dependent variable 'Risk Aversion, Switchers', the observation numbers change to the following: For Panel B there are 30 'deceived' in the Full Sample, 16 'deceived' in Player A, and 21 'not deceived' and 12 'deceived' in Player B. For Panel C there are 17 'deceived' in Player B. For Panel F there are 39 'deceived' in the Full Sample; 22 'deceived' in Player A. For the dependent variable 'defected in prisoner's dilemma', the observation numbers change to the following: For Panel A, there are 75 for 'not deceived' in the Full Sample; 35 for 'deceived' in Player A; 39 for 'not deceived' and 30 for 'deceived' in Player B. For Panel B, there are 38 for 'not deceived' and 30 for 'deceived' in the Full Sample; 21 for 'not deceived' and 13 for 'deceived' in Player B. For Panel C, there are 37 for 'deceived' in the Full Sample; 18 for 'not deceived' in Player B. For Panel D there are 26 for 'not deceived' in the Full Sample; 15 for 'not deceived' in Player B. For Panel E there are 49 'not deceived' in the Full Sample; 22 for 'not deceived' in Player A; 24 for 'not deceived' and 13 for 'deceived' in Player B. For Panel F there are 44 'not deceived' in the Full Sample; 17 for 'not deceived' in Player B.

In panel F we look only at the ‘lucky’ players, focusing attention on the group of players that exhibits less selection. Within this group of players, we still find that players who experienced deception are more likely to play inconsistently in the risk game. On the other hand, there are no differences in play in the dictator game or the prisoners’ dilemma within this group of people. This gives suggestive evidence that the difference in behavior in the risk game, which does not involve a partner, is not due to selection, but rather to the effect of experiencing deception. On the other hand, any differences in behavior in the dictator game and prisoners’ dilemma, games that do involve partners, may be due to selection.

4. Conclusion

We find that deception influences both the selection of experiment participants as well as their behavior. Specifically, we find that females are less likely to return after being deceived and that those who fare badly due to luck *and* are deceived are less likely to return than those who fare badly due to luck and are not deceived. Regarding behavior for those who return, we find an increase in the likelihood of answering risk aversion questions inconsistently and an increase in the variance of the number of safe gambles chosen for those who do answer consistently. We suggest that this is evidence of not taking the games seriously.¹⁰ We also find that Player As, especially those who are either female or inexperienced, kept more of their money in the dictator game in the second round.

We have discussed these results with both psychologists and economists and are struck by their reactions: both see the data as supporting their priors! Perhaps this should come as no surprise, given what we know about confirmation bias (Lord et al., 1979). We fully understand that although we do find clear differences in behavior, they are subject to interpretation as to their economic (or psychological) importance, as well as to further refinement regarding their magnitude and generalizability. The irony is that further study of how deception influences behavior, both in the laboratory and in the real world, requires relaxing the no-deception rule.

In one sense, all we have shown is that prior experiences in life influence the way individuals play games in a laboratory and that experience in the laboratory is a part of real life (and hence influences later behavior). An opponent of the no-deception rule might argue that deception is no different than other such (uncontrollable) experiences, which must simply be assumed to be orthogonal to the treatments of interest. However, experimental deception is in fact controllable and, as we have demonstrated, has some non-random effects on behavior in the types of treatments that interest experimental economists. At the very least, since prior exposure to deception is potentially knowable information to researchers, it should be accounted for when possible.

Another possibility is that it may be necessary to deceive subjects in order to convince them that they have not been deceived.¹¹ This is perhaps most evident when dealing with probabilities, where people tend to have strong and incorrect preconceived notions. For instance, it is well known that individuals underestimate the number of long strings of either heads or tails in a sequence of random coin-flips, so when presented with a truly random sequence they may well become suspicious upon observing an extended sequence of one outcome. Slightly more subtly, if an experiment involves a very small probability of some extreme event, subjects may doubt the researcher if it ever does occur (especially if this entails a negative outcome for the subject). This provides an incentive to ensure that such extreme events never happen.

Lastly, one further point should be made: all we have shown is that those who were deceived behave differently than those who were not, primarily in terms of return rates. If deception is deemed rampant in the real world, are we better off testing behavior in a sterile, deception-free environment? Or would we get more widely relevant results if individuals were suspicious of the administrators, just as they may be in the real world with respect to economic transactions? To the extent that we want to use laboratory experiments to infer real world behavior, merely finding that deceived individuals behave differently than non-deceived individuals does not tell us which pool behaves in a lab more as they would in their natural environment. These issues are drawn out in more detail in Levitt and List, which discusses the link between the laboratory and the real world and how issues such as those discussed here influence the interpretability of laboratory experiments.

The debate in the deception literature thus far has been pursued at a philosophical level (something to which economists are not often accustomed), rather than as a cost-benefit analysis (something to which economists are accustomed).¹² This paper identifies some of the costs, but does not address the benefits. The benefits clearly accrue from the research questions that would remain unanswered because of a proscription on deception. Maintaining separate subject pools, conceivably with a higher rate paid for access to unadulterated subjects, might be one way to address this issue.

Ortmann and Hertwig (2001) discuss four characteristics of experimental economics that are not used by experimental psychologists, one of which is the proscription of deception. They argue that psychologists can learn from economists’ practices. Roth (2001) replies that the public costs to deception may be lower for psychologists than they are for economists. Even if psychologists were to stop using deception as an experimental tool today, students would continue to be taught about research that used deception. Thus, participants in psychological experiments may remain suspicious for many years. Since

¹⁰ An alternative explanation is selection: those who did not fully understand the deception and why it mattered are more likely to return, and such individuals are also less likely to understand the risk aversion questions.

¹¹ Thanks go to Matthew Rabin for suggesting this line of reasoning.

¹² For instance, some academics pose ethical considerations as reason for this rule. We suggest this is the role for the human subjects review boards at universities, not for one subfield of one discipline to decide independently of the rest of academe. See Friedman and Sunder for more discussion on this.

economists have maintained the reputation of eschewing deception, the costs of deception carried out by a few researchers in economics may be higher for the profession as a whole. The effect of deception on non-deceived participants is even harder to measure. Those who have not yet been deceived may talk to deceived participants or hear about experiments using deception in their economics classes, thus tainting their perceptions of experimentation.¹³ If these spillover effects are truly large, it may suggest that the maintenance of a “deception-free” participant list is not enough to limit the effects of deception on experimental outcomes.

Acknowledgements

The authors thank Jennifer Alix-Garcia, James Andreoni, Colin Camerer, Martin Dufwenberg, Dan Friedman, John List, George Loewenstein, Andreas Ortmann, Matthew Rabin, Uri Simonsohn, Richard Thaler, Roberto Weber, and seminar participants at UC Berkeley and the ESA 2006 meetings for useful discussions, and the Yale University Institute for Social and Policy Studies for funding. Research assistance from Mark Borgschulze, Elysha Massatt, Scott Nelson, and David Owens is appreciated.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2008.09.002.

References

- Abric, J.C., Kahan, J., 1972. The effects of representations and behavior in experimental games. *European Journal of Social Psychology* 2, 129–144.
- Ball, S., Eckel, C., Grossman, P.J., Zame, W., 2001. Status in markets. *Quarterly Journal of Economics* 116, 161–188.
- Bardsley, N., 2000. Control without deception: individual behaviour in free-riding experiments revisited. *Experimental Economics* 3, 215–240.
- Berg, J., Dickhaut, J., McCabe, K., 1995. Trust, reciprocity, and social history. *Games and Economic Behavior* 10, 122–142.
- Blount, S., 1995. When social outcomes aren't fair: the effect of causal attributions on preferences. *Organizational Behavior and Human Decision Processes* 63, 131–144.
- Bonetti, S., 1998. Experimental economics and deception. *Journal of Economic Psychology* 19, 377–395.
- Brock, T.C., Becker, L.A., 1966. Debriefing and susceptibility to subsequent experimental manipulations. *Journal of Experimental Social Psychology* 2, 314–323.
- Christensen, L., 1977. The negative subject: myth, reality, or a prior experimental experience effect? *Journal of Personality and Social Psychology* 35, 392–400.
- Cook, T.D., Bean, J.R., Calder, B.J., Frey, R., Krovetz, M.L., Reisman, S.R., 1970. Demand characteristics and three conceptions of the frequently deceived subject. *Journal of Personality and Social Psychology* 14, 185–194.
- Davis, D.D., Holt, C.A., 1993. *Experimental Economics*. Princeton University Press, Princeton.
- Epley, N., Huff, C., 1998. Suspicion, affective response, and educational benefit as a result of deception in psychology research. *Personality and Social Psychology Bulletin* 24, 759–768.
- Fillenbaum, S., 1966. Prior deception and subsequent experimental performance: the ‘faithful’ subject’. *Journal of Personality and Social Psychology* 4, 532–537.
- Friedman, D., Sunder, S., 1994. *Experimental Methods: A Primer for Economists*. Cambridge University Press, Cambridge, UK.
- Frohlich, N., Oppenheimer, J., Moore, J.B., 2001. Some doubts about measuring self-interest using dictator experiments. *Journal of Economic Behavior and Organization* 46, 271–290.
- Gibbons, R., Van Boven, L., 2001. Contingent social utility in the prisoners’ dilemma. *Journal of Economic Behavior and Organization* 45, 1–17.
- Golding, S.L., Lichtenstein, E., 1970. Confession of awareness and prior knowledge of deception as a function of interview set and approval motivation. *Journal of Personality and Social Psychology* 14, 213–223.
- Henrich, J., 2001. Challenges for everyone: real people, deception, one-shot games, social learning, and computers. *Behavioral and Brain Sciences* 24, 414–415.
- Hey, J.D., 1998. Experimental economics and deception: a comment. *Journal of Economic Psychology* 19, 397–401.
- Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. *American Economic Review* 92, 1644–1655.
- Karlan, D., 2005. Using experimental economics to measure social capital and predict financial decisions. *American Economic Review* 95, 1688–1699.
- Kiesler, S., Sproull, L., Waters, K., 1996. A prisoner’s dilemma experiment on cooperation with people and human-like computers. *Journal of Personality and Social Psychology* 70, 47–65.
- Kim, O., Walker, M., 1984. The free-rider problem: experimental evidence. *Public Choice* 43, 3–24.
- Krupat, E., Garonzik, R., 1994. Subjects’ expectations and the search for alternatives to deception in social psychology. *British Journal of Social Psychology* 33, 211–222.
- Levitt, S., List, J., 2007. What do laboratory experiments measuring social preferences tell us about the real world? *Journal of Economic Perspectives* 21 (2), 153–174.
- Lord, C., Ross, L., Lepper, M., 1979. Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37, 2098–2109.
- McDaniel, T., Starmer, C., 1998. Experimental economics and deception: a comment. *Journal of Economic Psychology* 19, 403–409.
- Ortmann, A., Hertwig, R., 2001. Experimental practices in economics: a methodological challenge for psychologists? *Behavioral and Brain Sciences* 24, 383–403.
- Ortmann, A., Hertwig, R., 2002. The costs of deception: evidence from psychology. *Experimental Economics* 5, 111–131.
- Rilling, J.K., Sanfey, A.K., Aronson, J.A., Nystrom, L.E., Cohen, J.D., 2004. The neural correlates of theory of mind within interpersonal interactions. *NeuroImage* 22, 1694–1703.
- Roth, A.E., 2001. Form and function in experimental design. *Behavioral and Brain Sciences* 24, 427–428.
- Sanfey, A.G., Rilling, J.K., Aronson, J.A., Nystrom, L.E., Cohen, J.D., 2003. The neural basis of economic decision-making in the ultimatum game. *Science* 300, 1755–1758.
- Scharlemann, J.P.W., Eckel, C.C., Kacelnik, A., Wilson, R.K., 2001. The value of a smile: game theory with a human face. *Journal of Economic Psychology* 22, 617–640.
- Schechter, L., 2006. Traditional trust measurement and the risk confound: an experiment in rural Paraguay. *Journal of Economic Behavior and Organization* 62, 272–292.

¹³ Henrich (2001) makes the argument that this reputational spillover will be of lesser importance if experimentalists use subjects who are not students.

- Silverman, I., Shulman, A.D., Wiesenhal, D.L., 1970. Effects of deceiving and debriefing psychological subjects on performance in later experiments. *Journal of Personality and Social Psychology* 14, 203–212.
- Stricker, L.J., Messick, S., Jackson, D.N., 1967. Suspicion of deception: implications for conformity research. *Journal of Personality and Social Psychology* 5, 379–389.
- Thibaut, J., Walker, L., 1975. *Procedural Justice: A Psychological Analysis*. Erlbaum, Hillsdale, NJ.
- Weimann, J., 1994. Individual behavior in a free riding experiment. *Journal of Public Economics* 54, 185–200.
- Wiener, R.L., Erker, P.V., 1986. The effects of prebriefing misinformed research participants on their attributions of responsibility. *Journal of Psychology* 120, 397–410.
- Winter, E., Zamir, S., 2005. An experiment on the ultimatum bargaining in a changing environment. *Japanese Economic Review* 56, 363–385.