



Yale University
Department of Economics

Yale Working Papers on Economic Applications and Policy

Yale University
P.O. Box 208268
New Haven, CT 06520-8268
DISCUSSION PAPER NO. 18

**To Deceive or Not to Deceive: The Effect of Deception on Behavior in
Future Laboratory Experiments***

Julian Jamison
University of California, Berkeley

Dean Karlan
Yale University

Laura Schechter
University of Wisconsin, Madison

June 2006

*The authors thank Jennifer Alix-Garcia, James Andreoni, Colin Camerer, Martin Dufwenberg, Dan Friedman, John List, George Loewenstein, Richard Thaler and Roberto Weber for useful discussions, and the Yale University Institute for Social and Policy Studies for funding. Research assistance from Mark Borgschulte, Elysha Massatt, and Scott Nelson is appreciated. All errors remain our own.

This paper can be downloaded without charge from the Social Science Research Network
electronic library at: <http://ssrn.com/abstract=913057>

To Deceive or Not To Deceive: The Effect of Deception on Behavior in Future Laboratory Experiments

Julian Jamison
University of California, Berkeley
julison@berkeley.edu

Dean Karlan
Yale University
dean.karlan@yale.edu

Laura Schechter
University of Wisconsin, Madison
lschechter@wisc.edu

Abstract

Experimental economists believe (and enforce) that researchers should not employ deception in the design of experiments. The rule exists in order to protect a public good: the ability of other researchers to conduct experiments and 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, and then examining whether the deceived participants behave differently in a subsequent study. We find significant differences in both 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.

KEY WORDS: laboratory experimental methods; experimental economics; deception; psychology and economics; laboratory selection effects

JEL CLASSIFICATIONS: B40, C81, C90, C91, D80, D83

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 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 experiment instructions are truthful). As some research questions may be better answered by using deception, should we forgo the knowledge which could be acquired through such experiments in order to maintain a common pool of deception-free participants? This concern warrants testing to observe the presence of such sample contamination.

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) conduct a more thorough review of the evidence from experiments in psychology. They conclude that experiencing deception does affect participants’ expectations, suspicions, and future behavior.

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, Messick, and Jackson (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 than those who claimed not to suspect anything. There is another line of literature looking at effects on behavior of being warned before an experiment by the experimenter that the experiment “might” include misinformation, or 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.

These papers do not look at the direct effects of experiencing deception in a previous experiment on behavior in future experiments; these consequences are the principal motivation for the prohibition of deception in experimental economics. There are a few experiments in psychology which attempt exactly this type of test, albeit with non-economic experiments such as personality tests and memory games. Some find that experiencing deception leads to changes in behavior (Silverman, Shulman, and Wiesenthal 1970; Christensen 1977) while others do not (Fillenbaum 1966). Some find that past experiences with deception only 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 deception and those conducted in economics (as well as 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. 1970). Other experiments subjected these individuals to a very different experience than that experienced by the deceived (Christensen 1977). In contrast, in this paper, the initial treatment was almost identical for those who were

deceived and those who were not. In some of the psychological experiments the first and second treatments were conducted sequentially. After the participants participated 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 & Becker 1966, Cook et al. 1970, Christensen 1977). This is not the conventional method of subject recruitment. Here, the initial and secondary experiments were separated by approximately three weeks, and the recruitment mechanism for the second experiment 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.¹ 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 they are matched. Weimann (1994), Blount (1995), Scharlemann et al. (2002), Sanfey et al. (2003), and Winter and Zamir (2005) all tell players that their partners are humans, when in fact they are computers playing either predetermined or random strategies. 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 allows the researcher to test how subjects react to specific situations which may not arise naturally with high probability. In this paper, we test the effects of just this sort of deception on the sample selection of those who return to participate in further experiments, as well as the behavior in future experiments.

2 Experimental Procedures

2A 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, booths at student and staff activity fairs, and so forth. Subjects must actively choose to join the list, which then makes them eligible to receive announcements regarding (and to subsequently sign up for)

¹ For example, subjects were given two poems by Robert Frost and told that one was by Frost and one by a high school English teacher and were 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 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 that experiment was a sign of obsessive compulsive disorder. Then they were allowed to try again. After the experiment, they were debriefed again, being told that excellence in writing phone numbers was not actually a sign of obsessive compulsion; the experimenter wanted to see if the participant would write fewer phone numbers after hearing that misinformation.

² The psychology department maintains a separate subject list.

specific experimental sessions. There are approximately six to ten unique experiments (requiring the same number of email announcements requesting participants) in a given month.

Our first round involved 261 total subjects across ten sessions, each of which lasted for approximately one hour. 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 minutes for individual checks to be filled out and distributed. See below for specific descriptions of all games.

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. In the deception treatment (five sessions, 129 subjects), subjects were randomly assigned to roles, but in all cases their opponent's actions were determined by a computer program that simulated the human play.³ During the check-processing wait time, these subjects were informed that they had been deceived and had actually played against a computer rather than another human in the room. They were told that this was necessary for the research (but nothing more specific) and were asked to sign a second consent form allowing their data to be used.

Two to three weeks later, all subjects from the first round were sent a recruitment email for new experimental sessions (which were assigned a different researcher name from those in round one). 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. That in itself is not an uncommon occurrence, given various screening criteria used by other researchers at the lab.

142 people returned for one of the eight sessions that took place 3-4 weeks after the original sessions in round one. These lasted slightly under an hour, and each consisted of both the deceived and non-deceived subjects. A different 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 (see below for details).

Afterward, as they waited to be paid, they were informed (as per human subjects protocol requirements) 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 experiment.

2B. Specific games

In our first set of sessions, we ran a trust game (Berg, Dickhaut, and McCabe 1995) with a \$20 endowment to the trustor.⁴ 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.

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

⁴ 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 2006, Schechter 2006), we use the canonical terminology here of "trustor" and "trustee."

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 sometimes 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: if the trustee ever sent back less than what was invested, the trustor never again invested anything. The results section below outlines the summary statistics showing that the overall average play (for humans in each of the two roles) cannot be distinguished in the deception vs. the non-deception treatments.

In our second set of sessions, which occurred three to four 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, as 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 / 2x2 Matrix Games respectively.

The first 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 each subject ended up as only one or the other.

The second “game” was an ordered series of gambles (Holt and Laury, 2002), with ten 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% chance and \$8.80 with 90% chance versus a (risky) lottery that paid \$21.20 with 10% chance and \$0.55 with 90% chance. Here the first lottery was both risk-dominant and had a significantly higher expected value. As the choices progressed, the probability of the higher payoff in each lottery increased by 10%, 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 their earnings from the games (as promised in the recruitment email).

Our hypothesis was that players who had been deceived in the past would be suspicious in subsequent experiments that their partner was actually a computer, even when we told them

they were playing with a live partner. 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 (Abric and Kahan 1972, Kiesler et al. 1996, Rilling et al. 2004,). Thus, we suspected that those players who had been deceived in the past would cooperate less often in the Prisoners' Dilemma. We also suspected that players who had been deceived would suspect they were playing with a computer and keep more as dictator in the Dictator Game.⁵ We included the Holt-Laury risky lottery as a control, our hypothesis being that those who were deceived would not play any differently than those who were not deceived in the risky lottery 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

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. None of these differences is significant. 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. Males make up 37% of the deceived population and 42% of the non-deceived population. None of these differences is significant either. This suggests that the two groups were similar.⁶

3A Effect of Deception on Selection into Subsequent Experiment

The next step in our analysis involved looking at the effect of deception on return rates. All subjects were invited to participate in another round of experiments approximately three 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% of the original participants returned for the second round of experiments.

Table 1 shows that the same proportion of those who were deceived returned as those who were not. 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 are 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.

⁵ There is no previous evidence, of which we are aware, on this topic as it does not make very much sense to play a Dictator Game with a computer.

⁶ Psychologists believe there may be a "faithful subject" who attempts 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 "faithful."

The next few rows of Table 1 examine further selection effects. We divide the sample into four categories: “High Passers” and “Low Passers,” and Player A and Player B.⁷ We find no significant differences within three of these four groups. Only “Low Passers, 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% 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 B’s 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 the player also did not do well in the game. (Remember 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 randomly chose 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 return rates for those players for whom the last round determined their payout as opposed to those whose earlier rounds determined their payout. 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 and received a low payout. Perhaps these players felt that, not only did we deceive them in terms of the identity of their partner, but we may have also deceived them in terms of purposely choosing to pay them for the round with the lowest payoffs or programming the “trustor” computer to send them very little, which 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 feel that the procedures in the experiment were unjust, and not return for that reason. For more on this, refer to Thibaut and Walker (1975), whose seminal work

⁷ When categorizing players as high or low passers we ignore play in the last round of the trust game. Since many Player A’s 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. A player is a “High Passer” if he passed at least half of what he could pass, and a “Low Passer” otherwise.

on procedural justice finds that, holding the outcome constant, satisfaction depends upon the process by which outcomes are reached.

3B 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 practically differentiate 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% 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. Differences in behavior may also be due to the selection effect: a player who was deceived once may play no 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.

Table 3 presents the results. We measure four outcomes: (1) the amount kept in the dictator game; (2) risk aversion from the Holt-Laury 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 played as 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 which exhibits less selection.

The top left quadrant of Table 3 shows the primary results for the full sample. As one can see, 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-Laury 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, eight 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-Laury 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 at one, and only one, point from the safe option to the risky one. This result suggests another potential test. We had hypothesized that the deceived would exhibit more self-interested behavior because they were less likely to believe that their partner actually existed. In addition, we might expect 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.

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, we find that the variance for the deceived is significantly higher than for the non-deceived at the 10% level.⁸ 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.⁹ Using the Bonferonni correction within panels of Table 3, the only finding which remains significant at the 10% level is that deceived females are more likely to play inconsistently in the Holt-Laury gamble.

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 were classified as Player A in Round One kept significantly more as dictator than do Player A’s who were not deceived. Further analysis shows that this result holds for both the inexperienced Player A’s (i.e. who have participated in fewer than 10 previous experiments in the Xlab) and female Player A’s. 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).

Looking only at the “lucky” players, we focus 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. This gives suggestive evidence that the difference in behavior is not due to selection, but rather, the effect of experiencing deception.

4 Conclusion

We find that deception may influence 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 (evidence, we suggest, of not taking the games seriously).¹⁰ We also find that Player A’s, especially those Player A’s who are either female or inexperienced, kept more of their money in the Dictator game in the subsequent round.

We have discussed these results with both psychologists and economists, and are struck by their reactions: both see our results as supporting their priors! Albeit this comes as no surprise, given what we know about confirmation bias (Lord, Ross, and Lepper, 1979). We put forth that whereas we do find differences, they are subject to interpretation as to their magnitude and economic (or psychological) importance. The irony is that further study of how deception

⁸ We carry out Levene’s test for equality of variances on risk aversion in the risk experiment and the amount kept in the dictator game. The variances are significantly different at the 10% level for risk aversion for all players, for all player B’s, and for player B’s who had participated in at least 10 experiments (at the 5% level). The variances of the amount kept in the dictator game are only significantly different at the 10% level in one comparison (all males) and in the wrong direction (the variance for the deceived is lower).

⁹ 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).

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.

influences behavior, both in the laboratory and in the real world, does require 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 real (and hence influences later behavior). An opponent of the no deception rule might argue that deception is no different than such other (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 the types of treatments that interest experimental economists. At the very least, since prior exposure to deception is potentially knowable information to researchers, it could be accounted for when possible.

Lastly, one further point should be made: all we have shown is that those who were deceived behave differently than those who were not. 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 generalizable 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 more similarly in a lab to how they would in their natural environment. These issues are drawn out in more detail in Levitt and List (2006), 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.

This debate, ironically, has been at a philosophical level (something to which economists are not accustomed), rather than pursued as a cost-benefit analysis (something to which economists are accustomed).¹¹ This paper identifies some of the costs, but does not address the benefits. The benefits clearly depend on the research questions that remain unanswered because of the experiment's design. The results presented here do not support an absolute rule against deception. Those who wish to deceive could compensate the lab by providing sufficient funds to maintain an appropriately-sized and "deception-free" participant list. Separate participant pools (as for the Xlab at Berkeley) can also be maintained. If a certain research question requires deceiving the "pure" pool, then individuals can be transferred, with compensation provided, to the "no-deception" lab for raising fees in order to increase the size of the participant pool.

Ortmann and Hertwig (2001) discuss four characteristics of experimental economics which 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 economists have maintained the reputation of eschewing deception, the costs of deception carried out by a few researchers in economics may be higher. The effect of deception on non-deceived participants is harder to measure; they may talk to deceived participants or hear about experiments using deception in

¹¹ 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 sub-field of one discipline to decide independently of the rest of academe. See Friedman and Sunder (1994) for more discussion on this.

their economics classes, thus tainting their perceptions of experimentation.¹² 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.

¹² Henrich (2001) makes the argument that this reputational spillover will be of lesser importance if experimentalists use subjects who are not students.

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(1), 161—188.
- 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(2), 131—144.
- Bonetti, S. (1998), 'Experimental economics and deception', *Journal of Economic Psychology* 19(3), 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. and Holt, C. A. (1993), *Experimental Economics*, Princeton University Press.
- Fillenbaum, S. (1966), "Prior deception and subsequent experimental performance: The 'faithful' subject", *Journal of Personality and Social Psychology* 4, 532—537.
- Friedman, D. and Sunder, S. (1994), *Experimental Methods: A Primer for Economists*, Cambridge University Press, Cambridge, England.
- Gibbons, R. & Van Boven, L. (2001), 'Contingent social utility in the prisoners' dilemma', *Journal of Economic Behavior and Organization* 45(1), 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(3), 397—401.

- Holt, C. A. & Laury, S. K. (2002), 'Risk aversion and incentive effects', *American Economic Review* 92(5), 1644—1655.
- Karlan, D. (2005), 'Using experimental economics to measure social capital and predict financial decisions', *American Economic Review* 95(5), 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(1), 47—65.
- Levitt, S. and List, J. (2006) 'What do laboratory experiments tell us about the real world?' University of Chicago working paper
- Lord, C., Ross, L., and 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(3), 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(2), 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(4), 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(5626), 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(5), 617—640.
- Schechter, L., (2006), 'Traditional trust measurement and the risk confound: An experiment in rural Paraguay'. *Journal of Economic Behavior and Organization*, Forthcoming.
- Silverman, I., Shulman, A. D., & Wiesenthal, 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*. Hillsdale, NJ: Erlbaum.

Weimann, J. (1994), 'Individual behavior in a free riding experiment', *Journal of Public Economics* 54(2), 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(3), 363—385.

Appendix A

Experimental Instructions, Round One:

Welcome to this experiment on decision making and thank you for being here. You will be compensated for your participation in the experiment, though the exact amount you will receive will depend on the choices you and others make, and on random chance (as explained below). Please pay careful attention to these instructions, as a significant amount of money is at stake.

Information about the choices that you make during the experiment will be kept strictly confidential. Your name will appear only on the payment-receipt form and will not be linked to any specific choices you make. You will not be asked to reveal your identity or the content of your decisions to anyone else (either the experimenter or other participants) at any time during or after the experiment. In order to maintain privacy and confidentiality, please do not speak to anyone during the experiment and please do not discuss your choices with anyone even after the conclusion of the experiment.

In this experiment each of you will be paired with somebody else who is in this room. You will not be told who that person is either during or after the experiment. Half of you will be assigned to the role of Player A and half of you will be assigned to the role of Player B. Each person playing the role of A will be assigned \$20 to begin the experiment. Those assigned the role of A (henceforth Player A) will have the opportunity to send some, all, or none of their \$20 endowment to a player playing the role of B (henceforth Player B). Each dollar sent to Player B will be tripled. For example, if Player A sends \$4, Player B will receive \$12. If Player A sends \$18, Player B will receive \$54. Player B will then decide how much money (if any) to send back to Player A and how much money (if any) to keep.

You will play this game in its entirety six times (or rounds). You will remain with the same partner and role for all six rounds. The first two rounds are just for practice so that you learn the rules of the game. The next four rounds will be used to determine the payouts. In the end, one of these last four rounds will be chosen at random to determine you and your partner's payouts according to your payoff in that round.

Figure 1, below, illustrates the sequence.

PLAYER A => SEND between \$0 and \$20 to B.

PLAYER B => RECEIVE triple what A sent.

PLAYER B => SEND back some, none or all of new total to Player A.

Figure 1. One complete Round.

The earnings of A and B are given by:

A's earnings = 20 - [what (s)he gives to B] + [what (s)he receives from B]

B's earnings = [3 times what (s)he has been given by A] - [what (s)he returns to A]

Your role (A or B) will be told to you before the experiment starts. You will play this role in all subsequent rounds.

After all participants have made their choices in the first round, you will receive information about your earnings. Then you will play again and you will receive information about your earnings again. After the two practice rounds, you will play four more times. One of these last four rounds will be chosen at random to decide your actual payoffs. The dollar amount in that round, plus a \$5 show-up bonus, will be the amount you are paid for your participation in the experiment.

Once the experiment begins please remain silent and in your seat until it is over.

Computer Program Description

To make your decisions you will use the computer in front of you. Right now, you can see an initial waiting screen. The program will be activated when the instructions are finished.

Once the program is activated, you will go through two practice rounds to make certain that you correctly understand how the experiment will work. These practice rounds will be numbered “-1” and “0.” When we start, a new window will pop up and replace the initial waiting window. If you are Player A, the new window will resemble the one shown below. In this case, the upper-left of the screen informs you that you are in the first of 4 Rounds. The upper-right of the screen shows the time in which you are encouraged to make your decision. After you have completed the practice rounds, the experiment will begin. The rounds that determine your payoffs will be numbered “1” through “4.” The center of the screen will inform you that you are in round 1 and, if you are Player A, ask you for your decision. If you are Player B you will be asked to wait for the other player to make their choice. You may choose any discrete amount between \$0.00 and your full endowment (\$20.00 in the case of player A.) Your choice will be rounded to the nearest penny.

The screenshot shows a window titled "Round" with a light beige background. At the top left, it says "1 of 4". At the top right, it says "Time Remaining 21". The main area contains the following text: "You have been assigned the role of Player A", "You have an endowment of \$20.00", and "How much of that (if any) do you want to send to player B?". Below this text is a text input field containing the number "5". In the bottom right corner, there is a red "OK" button.

After Player A makes his (or her) decision, Player B will be given the option of sending some, all or none of his (or her) money back to Player A. The computer will display the screen below:

Round

1 of 4

Time Remaining 18

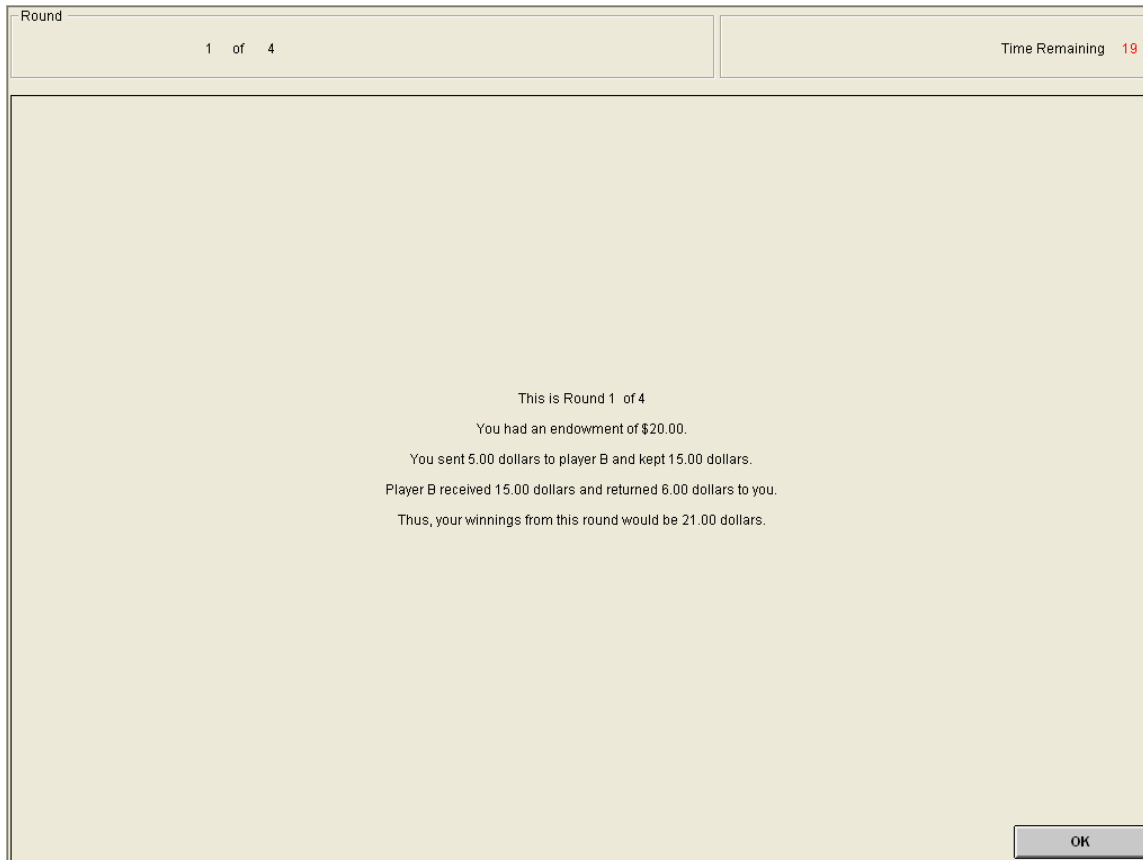
Player A received an endowment of \$20.00.

Player A sent you 5.00 dollars. This was tripled so you have received 15.00 dollars.

How much of that (if any) do you want to send back to player A?

OK

In this case, Player B would be able to send back any amount (s)he wants to Player A. Player B can choose to send "\$0.00," "\$15.00" or any number in between. Player A will view a waiting screen while Player B makes their decision.



Finally, you will be shown the results of the round, including your payout and a summary of the choices you and the other player made. Player A's screen is shown above. Player B will view a similar screen. You can click OK, or wait until the "Time Remaining" expires, at which point you will automatically progress to the next screen.

Then, the next round will start. You will continue to play the role of Player A or B, and your partner will not change. Your final payout for your time here will be chosen randomly from the four non-practice rounds, so you will want to pay attention throughout the experiment.

After the last round is finished, please sit quietly while we process the results and issue your payment. You may go to the bathroom, work on homework or read during this period, which should last between 10 and 20 minutes.

At no point in the experiment will you be told who your partner is, nor will we identify you to your partner. There are no right or wrong choices in this experiment.

Example

Player A starts with \$20. She sends \$15.

Player B receives \$45, and sends back \$30.

HOW MUCH does each player have at the end of the round?

Rules

Please do not talk with anyone during the experiment. We ask everyone to remain silent through the end of the last round.

Your participation in the experiment and any information about your earnings will be used solely for research purposes. Your name and association to your decisions will be kept strictly confidential. If you have any remaining questions about the confidentiality agreement you will be given a chance to ask in just a minute.

We ask that you assist us in maintaining the integrity of the experiment. If at all possible, please do not discuss the experiment, your choices in it, or the outcomes with anyone, either during the experiment, or after you leave the Xlab.

Questions

Any clarification questions should be asked at this time. Please raise your hand and wait for a researcher to come to your desk. If you need to use the bathroom, please do so before the start of the experiment.

Experimental Instructions, Round Two

Welcome to this experiment, and thank you for agreeing to participate. You will be compensated for taking part, but the exact amount that you receive will depend on the choices that you and others make (and possibly also on random chance), as will be explained below.

The experiment will consist of a total of three question problems. At the end of the experiment, we will randomly choose (with a computer) exactly one question problem. For that one, the computer will look at your choices (confidentially) and determine your outcome. You will then be paid (in real dollars) the amount determined by your choices plus \$10 just for participating.

When it says "session name," type in zeam2037q, and submit. You should now be on the login page that includes a series of instructions.

Pay careful attention to the instructions on the log in page and be sure you understand them before you log in. When you are ready to log in, you must use your computer number as your codename.

Please let us know (by raising your hand quietly) if you have any questions at any point during the experiment.

Instructions (ID =), Page 1 of 3

- **Single Decision:** The experiment consists of a single **round** in which you are matched with another randomly selected person. **Note:** You will only make a single decision in this experiment.
- **Interdependence:** The decisions that you and the other person make will determine the amounts earned by each of you.
- **Decisions:** One of you will be designated as a **Proposer** (decision-maker) who will begin the round by deciding on a division of an amount of money, **\$20.00**. The other person, designated as **responder**, is told how the money is divided between the two people and has no choice to make.

Continue with Instructions

Vecon Lab - June 12, 2006

Instructions (ID = , Proposer) Page 2 of 3

- **Role:** You have been randomly assigned to be a **Proposer** (or initial decision-maker) in this process, and you will decide on amounts of money for each of you that sum to **\$20.00**.
OR **Role:** You have been randomly assigned to be a **Responder** in this process. The other person (proposer/decision-maker) will begin by deciding on amounts of money for each of you that sum to **\$20.00**.
- **Earnings:** Then each person earns their part of the division.

Vecon Lab - June 12, 2006

Instructions Summary (ID = , Proposer)

- **One Decision Only:** Please remember that you are randomly matched with another participant, and that this process will not be repeated, you make only one decision.
- **Earnings:** The proposer (decision-maker) in each pair decides on a division of the **\$20.00**; this decision determines earnings for the round.
- **Finish:** After you and the person you are matched with have made your decisions, earnings will be determined and announced. There will not be another round of bargaining with this person or with anyone else.

Finished with Instructions

Vecon Lab - June 12, 2006

Instructions (ID =), Page 1 of 6

You will be making choices between two lotteries, such as those represented as "Option A" and "Option B" below. The money prizes are determined by the computer equivalent of throwing a ten-sided die. Each outcome, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, is equally likely. If you choose Option A in the row shown below, you will have a 1 in 10 chance of earning \$11.00 and a 9 in 10 chance of earning \$8.80. Similarly, Option B offers a 1 in 10 chance of earning \$21.20 and a 9 in 10 chance of earning \$0.55.

Decision	Option A	Option B	Your Choice
1st	\$11.00 if the die is 1	\$21.20 if the die is 1	A: <input type="radio"/> or B: <input type="radio"/>
.	\$8.80 if the die is 2 - 10	\$0.55 if the die is 2 - 10	
.			

Continue with Instructions

Vecon Lab - June 13, 2006

Instructions (ID =), Page 2 of 6

- Each row of the decision table contains a pair of choices between **Option A** and **Option B**.
- You make your choice by clicking on the "A" or "B" buttons on the right. Only one option in each row can be selected, and you may change your decision as you wish.
- **Note:** Note, try clicking on one of the radio buttons, then change by clicking on the other one.

Decision	Option A	Option B	Your Choice
1st	\$11.00 if the die is 1	\$21.20 if the die is 1	A: <input type="radio"/> or B: <input type="radio"/>
.	\$8.80 if the die is 2 - 10	\$0.55 if the die is 2 - 10	
.			

Continue

Vecon Lab - June 13, 2006

Instructions (ID =), Page 3 of 6

Even though you will make ten decisions, **only one** of these will end up being used. The selection of the one to be used depends on the "throw of the die" that is determined by the computer's random number generator. No decision is any more likely to be used than any other, and you will **not** know in advance which one will be selected, so please think about each one carefully. This random selection of a decision fixes the row (i.e. the Decision) that will be used. For example, suppose that you make all ten decisions and the throw of the die is 9, then your choice, A or B, for decision 9 below would be used and the other decisions would not be used.

Decision	Option A	Option B	Your Choice
· · 9th	\$11.00 if the die is 1 - 9 \$8.80 if the die is 10	\$21.20 if the die is 1 - 9 \$0.55 if the die is 10	A: <input type="radio"/> or B: <input type="radio"/>

Continue

Vecon Lab - June 13, 2006

Instructions (ID =), Page 4 of 6

After the random die throw fixes the Decision row that will be used, we need to obtain a second random number that determines the earnings for the Option you chose for that row. In Decision 9 below, for example, a throw of 1, 2, 3, 4, 5, 6, 7, 8, or 9 will result in the higher payoff for the option you chose, and a throw of 10 will result in the lower payoff.

Decision	Option A	Option B	Your Choice
9th	\$11.00 if the die is 1 - 9 \$8.80 if the die is 10	\$21.20 if the die is 1 - 9 \$0.55 if the die is 10	A: <input type="radio"/> or B: <input type="radio"/>
10th	\$11.00 if the die is 1 -10	\$21.20 if the die is 1 - 10	A: <input type="radio"/> or B: <input type="radio"/>

For decision 10, the random die throw will not be needed, since the choice is between amounts of money that are fixed: \$11.00 for Option A and \$21.20 for Option B.

Vecon Lab - June 13, 2006

Instructions (ID =), Page 5 of 6

- **Making Ten Decisions:** After you finish these instructions, you will see a table with 10 decisions in 10 separate rows, and you choose by clicking on the buttons on the right, option A or option B, for each of the 10 rows. You may make these choices in any order and change them as much as you wish until you press the Submit button at the bottom.
- **The Relevant Decision:** One of the rows is then selected at random, and the Option (A or B) that you chose in that row will be used to determine your earnings. **Note:** Please think about each decision carefully, since each row is equally likely to end up being the one that is used to determine payoffs.
- **Determining the Payoff:** After one of the decisions has been randomly selected, the computer will generate another random number that corresponds to the throw of a ten sided die. The number is equally likely to be 1, 2, 3, ... 10. This random number determines your earnings for the Option (A or B) that you previously selected for the decision being used.

Continue

Vecon Lab - June 12, 2006

Instructions Summary (ID =)

- To summarize, you will indicate an option, A or B, for each of the rows by clicking on the "radio buttons" on the right side of the table.
- Then a random number fixes which row of the table (i.e. which decision) is relevant for your earnings.
- In that row, your decision fixed the choice for that row, Option A or Option B, and a final random number will determine the money payoff for the decision you made.
- Payoffs will be made in cash, unless your instructor indicates otherwise.

Finished with Instructions

Vecon Lab - June 12, 2006

Instructions

- **Rounds and Matchings:** The experiment consists of a number of **rounds**. **Note:** You will be matched with the **same** person in all rounds. The number of rounds in this part of the experiment is: **1**.
- **Interdependence:** Your earnings are determined by the decisions that you and the other person make.
- **Roles:** In each pair of people, one person will be designated as the "row" player and the other will be the "column" player. You will be a **column player** (or) **row player** in all rounds.
- **Decisions:** On the next page, you will see a table of payoffs. The row player will choose between the Top and Bottom rows, and the column player will choose between the Left and Right columns. The intersection of the designated row and column will determine the earnings for each person.

Continue with Instructions

Vecon Lab - June 12, 2006

Instructions (Role and ID)

- The column player will press either the **Left** or the **Right** button. The row player will choose **Top** or **Bottom**. These choices determine which part of the matrix is relevant (Top Left, Top Right, Bottom Left, Bottom Right). In each cell, the row player's payoff is shown in blue and the column player's payoff is shown in red.

Payoff Matrix (Row, Column)

	Left	Right
Top	\$10.00, \$10.00	\$1.00, \$15.00
Bottom	\$15.00, \$1.00	\$6.00, \$6.00

- If you are a row player, your decision buttons will be on the left side of the payoff table, and if you are a column player, your decision buttons will be above the table. Please press one of your decision buttons to go to the next page.

Continue

Instructions Summary

- **Matchings:** Please remember that you will be matched with the **same** person in all rounds.
- **Earnings:** Your earnings are determined by the choices that you and the other person make in the round. You begin with a fixed payment of **\$0.00**, and earnings will be added to this amount (losses, if the game has negative payoffs, will be subtracted). Your total earnings will be displayed in a cumulative earnings column on the page that follows.
- **Rounds:** There will be **1 round** in this part of the experiment.

Finished with Instructions

Vecon Lab - June 12, 2006

Table 1: Selection Effects in Round Two (Based on Player Characteristics)
Mean, Standard Errors and Number of Observations

	Total	Not Deceived	Deceived	t-stat: diff≠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%; ** significant at 5%; *** significant at 1%. 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.

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≠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%; ** significant at 5%; *** significant at 1%. 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.

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: diff≠0	Not Deceived	Deceived	t-stat: diff≠0	Not Deceived	Deceived	t-stat: diff≠0
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Full Sample									
Amount Kept in Dictator Game	14.584 (0.484)	14.923 (0.518)	-0.477	13.472 (0.676)	15.441 (0.697)	-2.029**	15.561 (0.658)	14.355 (0.771)	1.193
Risk Aversion (Holt-Laury Gambles if Consistent)	6.067 (0.156)	5.804 (0.236)	0.968	6.171 (0.254)	5.867 (0.338)	0.732	5.975 (0.191)	5.731 (0.331)	0.685
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	
Panel B: Participants Played At Least 10 Prior Games in Laboratory									
Amount Kept in Dictator Game	15.179 (0.687)	15.032 (0.734)	0.146	13.941 (0.979)	14.882 (0.996)	-0.674	16.136 (0.922)	15.214 (1.125)	0.630
Risk Aversion (Holt-Laury Gambles if Consistent)	6.211 (0.233)	5.923 (0.404)	0.659	6.411 (0.429)	6.143 (0.592)	0.376	6.048 (0.244)	5.667 (0.555)	0.724
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	
Panel C: Participants Played Fewer Than 10 Prior Games in Laboratory									
Amount Kept in Dictator Game	13.974 (0.675)	14.824 (0.740)	-0.850	13.053 (0.947)	16.000 (0.985)	-2.154**	14.895 (0.939)	13.647 (1.057)	0.886
Risk Aversion (Holt-Laury Gambles if Consistent)	5.919 (0.206)	5.700 (0.272)	0.653	5.944 (0.286)	5.625 (0.375)	0.687	5.895 (0.305)	5.786 (0.408)	0.219
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	
Panel D: Males									
Amount Kept in Dictator Game	14.111 (0.902)	14.516 (0.746)	-0.349	13.818 (1.394)	14.214 (1.065)	-0.230	14.313 (1.217)	14.765 (1.066)	-0.280
Risk Aversion (Holt-Laury Gambles if Consistent)	5.703 (0.291)	5.393 (0.339)	0.694	5.545 (0.511)	5.545 (0.593)	0.000	5.813 (0.356)	5.294 (0.418)	0.938
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	
Panel E: Females									
Amount Kept in Dictator Game	14.840 (0.567)	15.294 (0.725)	-0.498	13.320 (0.774)	16.300 (0.936)	-2.532**	16.360 (0.723)	13.857 (1.143)	1.942*
Risk Aversion (Holt-Laury Gambles if Consistent)	6.271 (0.175)	6.214 (0.314)	0.170	6.458 (0.276)	6.053 (0.415)	0.842	6.083 (0.216)	6.556 (0.444)	-1.063
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	
Panel F: Non-'unlucky'									
Amount Kept in Dictator Game	13.867 (0.604)	14.675 (0.649)	-0.912	13.593 (0.779)	15.043 (0.870)	-1.245	14.278 (0.976)	14.176 (0.990)	0.073
Risk Aversion (Holt-Laury Gambles if Consistent)	6.091 (0.213)	5.480 (0.265)	1.803*	6.192 (0.304)	5.667 (0.388)	1.079	5.944 (0.286)	5.267 (0.358)	1.499
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
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%; ** significant at 5%; *** significant at 1%. 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 A, 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.

Appendix Table A: Holt-Laury Risky Lottery

Screen	Option A		Option B	
1	9/10 of \$8.80	1/10 of \$11.00	9/10 of \$0.55	1/10 of \$21.20
2	8/10 of \$8.80	2/10 of \$11.00	8/10 of \$0.55	2/10 of \$21.20
3	7/10 of \$8.80	3/10 of \$11.00	7/10 of \$0.55	3/10 of \$21.20
4	6/10 of \$8.80	4/10 of \$11.00	6/10 of \$0.55	4/10 of \$21.20
5	5/10 of \$8.80	5/10 of \$11.00	5/10 of \$0.55	5/10 of \$21.20
6	4/10 of \$8.80	6/10 of \$11.00	4/10 of \$0.55	6/10 of \$21.20
7	3/10 of \$8.80	7/10 of \$11.00	3/10 of \$0.55	7/10 of \$21.20
8	2/10 of \$8.80	8/10 of \$11.00	2/10 of \$0.55	8/10 of \$21.20
9	1/10 of \$8.80	9/10 of \$11.00	1/10 of \$0.55	9/10 of \$21.20
10	0/10 of \$8.80	10/10 of \$11.00	0/10 of \$0.55	10/10 of \$21.20

The respondents were presented with the above choices in sequence, starting in Screen 1 with the first choice above, and ending in Screen 10 with the final choice. We record in our analysis (a) whether individuals were consistent in their answers; and (b) if they were consistent, at what point they switch from Option A to Option B.