Do Participants Believe the Experimenter?*

Ro'i Zultan†

This version: September 5, 2015.

Abstract

A taboo on the deception of participants is a defining characteristic of experimental economics. Experimental economists avoid deception to maintain participants' trust in experimental instructions. Nonetheless, there is little experimental evidence on the extent to which this policy is successful, and on the factors influencing participants' trust. I develop and test a simple tool to estimate trust implicitly. Estimated trust is highly correlated with self reported trust. I find that participants are less likely to trust payoff-relevant instructions. Indirect exposure to deception, in the form of psychology studies, is not associated with higher suspiciousness. Direct exposure, in the form of past participation in a large number of psychology experiments is associated with an insignificant increase in self-reported trust, but a small significant reduction in estimated trust.

Keywords: experiments, methodology, deception, trust.

JEL classification: C90, C81

*I thank Bradley Ruffle and Michal Krawczyk for their thoughtful comments and Mor Uzan and Sagi Dekel for providing valuable research assistance. I thank Yoav Bar-Anan and Amos Schurr for providing participation statistics from their respective departments.

†Department of Economics, Ben-Gurion University, zultan@bgu.ac.il.

1 Introduction

A strict prohibition on deception of experimental participants is a defining characteristic of experimental economics (e.g., Davis and Holt, 1993; Hertwig and Ortmann, 2001). Proponents of this prohibition argue that deception erodes participants's trust, which is necessary to maintain experimental control and to draw valid inferences from observed behavior. Others have argued that participants do not view deceptive practices negatively, and that exposure to deception does not affect behavior in future experiments. After an initial interest in psychology starting with Kelman (1967), the advent of experimental economics has revived the debate, seeing overwhelming attention from experimental economists, social psychologists, and sociologists over the last couple of decades.[1]

Economic experiments typically create an economic microsystem, within which participants make choices given their preferences and endowments and the economic institutions (Smith, 1982). Clearly, trust in the instructions is necessary to maintain control over participants' beliefs, without which the hypotheses cannot be tested. Most importantly, the relation between participants' choices and the experimental outcomes must hold, as in Smith's (1982) precept of Saliency. If participants expect their payoff (or the payoffs of others) to be determined not according to the stated rules—for example, if everyone receives the same fixed payoff, or some random mechanism is manipulated in violation of the rules, as has been done in the past in psychology experiments—internal validity is compromised. Furthermore, administering different instructions, which may induce different levels of trust, in different experimental treatments creates an interaction between the independent variable and trust, undermining external validity (Cook and Campbell, 1979).

Whether deception erodes trust, and whether lack of trust in the experimenter compromises the experimental validity, are ultimately empirical questions. Several attempts have been made to test the compound question—whether deceptive practices affect behavior in future experiments—without explicitly looking at trust. For example, imagine that a participant is told in an experiment that she interacts with another person only to learn in a debriefing that the other person was a confederate or simulated by a computer. If this participant later takes part in an experiment that involves social preference—such as an ultimatum or a dictator game—she might question the existence of a human responder and behave more selfishly than she would in a real interaction. Jamison et al. (2008) did just that. In the first phase of their experiment, participants played a trust game, believing that they are playing with a human partner. Actually, half were paired with a computer designed to mimic the human participants' decisions, a fact that was only revealed to them in the post-experimental debriefing. Three to four weeks later, the same participants were invited back to the lab to play a dictator game, a risk elicitation task, and a

[1] The debate on deception is reviewed in the work by Ralph Hertwig and Andreas Ortmann (Hertwig and Ortmann, 2001, 2008a,b; Ortmann and Hertwig, 1997, 2002) Krawczyk (2013) reports results from a survey of experimental subjects and researchers' opinions on deceptive practices.

prisoner's dilemma game. Participants who played as a trustor and were deceived in the first phase kept slightly but significantly more in the second phase dictator game. This result, however, did not hold for those playing in the trustee role in the first phase, and was not replicated in a similar study by Barrera and Simpson (2012).

An experiment by Alberti and Güth (2013) also aimed to study the effects of deception on behavior in a subsequent task. Unlike the previous studies, here the researchers attempted to avoid actually deceiving the participants by allowing a 'participant-experimenter' to do the deception. Deception had no effect on subsequent decisions. Note, however, the debriefing explicitly stated that deception by another participant will not be allowed in the second phase. In a recent paper, Krawczyk (2015) manipulated whether invitations to participate or the experimental instructions included explicit announcements that the experiment involves no deception. Self-reported trust was higher when the instructions explicitly stated that there is no deception, although behavior in tasks designed to be sensitive to trust issues was not similarly affected.

The current paper is aimed at developing and utilizing a way to measure trust directly yet implicitly. Previous studies measured behavior that is determined not only by the degree of trust placed in the experimental instructions, but also by social preferences and notions of responsibility. For example, it is possible that direct exposure to deception increases the belief in being deceived again, however participants may still place a non-negligible probability on not being deceived and are reluctant to run the risk of punishing another participant for a misbehavior on the part of the experimenter. This, of course, may be enough if we are only interested in the potential effects of deceptive practices on experimental behavior. Nonetheless, a deeper understanding of the direct effects of such practices on trust can enrich the debate, and is needed to generalize the results. For example, adopting relaxed guideline across the discipline may have a cumulative effect that is not observable following a single experience with deception.

Trust in the experimenter can be measured directly by simply asking the participants whether they believed the experimental instructions, as in Krawczyk (2015). However, participants may be reluctant to report suspicion, especially when the instructions explicitly state that no deception is involved. Some studies present evidence that such self reports are indeed not reliable and even biased. Newberry (1973) had a confederate reveal solutions to the experimental task to the participants. When asked directly whether they had any prior information about the experiment or not, only 4 of 20 participants admitted to having such knowledge. Closer to the current study, Taylor and Shepperd (1996) report on an experiment involving deception. The participants in a pilot, unaware that one of them was a graduate student called on to fill in for a student who did not show up, discussed their individual feedback and discovered that they were deceived. Nonetheless, when asked several times during the debriefing whether they suspected deception,

none of the participants divulged this information. The data presented in this paper also suggest that participants under-report suspicions. moreover, this tendency appears to vary across populations.

Thus, the aim of this paper is to measure trust in a direct yet unobtrusive way. I follow the suggestion outlined by McKenzie and Wixted (2001), and employed in McKenzie et al. (2004). Specifically, I estimate the degree of trust in the experimenter by measuring the extent of belief updating after providing participants with new information. Participants are asked to rate their belief in a set of separate statements, e.g., *The German language has 30 letters* and *Hugh Hefner was married three times*. Next, participants are presented with the same set of statements in pairs, and rated their beliefs again under the instructions that exactly one statement in each pair is correct. The belief that the instructions are truthful can then be estimated using a Bayesian model taking into account the prior and posterior beliefs reported by the participants. The procedure allows the estimation of the trust parameter at both the individual and group level.

The research findings are the following. First, the new measure is validated against self reports. The estimated trust parameter is highly predictive of participants' self reported suspicion. Nonetheless, a substantial share of the participants are estimated to place little trust in the instructions do not report being suspicious. This result suggests that participants are inclined to hide their distrust, as in Taylor and Shepperd (1996). Second, the trust measure is used to test the effects of direct exposure to deception—in the form of previous participation in psychology experiments—as well as to indirect exposure—in the form of psychology undergraduate studies. The results do not support indirect exposure effects, but provide some evidence for direct exposure effects, albeit of small size. A third question pertains to whether deception is payoff-relevant. Half of the participants were paid for their probability assessments based on accuracy, while the other half were paid a flat fee. My working hypothesis was that participants are less likely to expect deception when it is payoff relevant. Surprisingly, participants were more suspicious when the instructions were payoff relevant.

2 The trust model

For each statement in a pair $i \in \{1, 2\}$, write p_i^1 for the prior probabilities and p_i^2 for the posterior probabilities assigned to statement i. If the announcement that exactly one statement is correct is fully believed, by Bayes' rule,

$$\begin{aligned} p_1^2 &= \frac{p_1^1(1-p_2^1)}{p_1^1(1-p_2^1)+p_2^1(1-p_1^1)}, \\ p_2^2 &= \frac{p_2^1(1-p_1^1)}{p_1^1(1-p_2^1)+p_2^1(1-p_1^1)}. \end{aligned} \tag{1}$$

Now add a trust parameter r, denoting the probability assigned to 'any information provided by the experimenter is correct'. Note that ignoring the announcement completely implies $r = 0.5$. Bayes' rule now gives

$$\begin{aligned} p_1^2 &= \frac{r[p_1^1(1-p_2^1)] + (1-r)(p_1^1 p_2^1)}{r[p_1^1(1-p_2^1) + p_2^1(1-p_1^1)] + (1-r)[p_1^1 p_2^1 + (1-p_1^1)(1-p_2^1)]}, \\ p_2^2 &= \frac{r[p_2^1(1-p_1^1)] + (1-r)(p_1^1 p_2^1)}{r[p_1^1(1-p_2^1) + p_2^1(1-p_1^1)] + (1-r)[p_1^1 p_2^1 + (1-p_1^1)(1-p_2^1)]}. \end{aligned} \tag{2}$$

In words, statement 1 can be true either if the announcement is truthful, statement 1 is true, and statement 2 is false; or if the announcement is untruthful and both statements are true, and correspondingly for statement 2. To illustrate, consider the prior probabilities $p_1^1 = 0.8$ and $p_2^1 = 0.4$ (the median probability pair obtained in the experiment). Applying equations (1), we obtain the posterior probabilities $p_1^2 = 0.857$ and $p_2^2 = 0.143$. By (2), if the probability assigned to exactly one statement being true is $r = 0.9$, we obtain $p_1^2 = 0.847$ and $p_2^2 = 0.190$. For $r = 0.8$, we obtain $p_1^2 = 0.836$ and $p_2^2 = 0.239$.

Naturally, with empirical data it is very unlikely that any value of r will perfectly balance both equations in (2). The parameter is therefore estimated parametrically using a maximum-likelihood non-linear regression that takes the functional form of (2) simultaneously on all statement pairs, 14 in number in the current experiment.

3 Experimental procedure and design

The experiment was conducted at the Negev Experimental Economics Laboratory at Ben-Gurion University of the Negev. The probability assessment task was completed as part of an unrelated large scale experiment in collaboration with psychology researchers. This experiment involved a large recruiting effort from the subject pools at the psychology and management departments, which ensured large proportions within the participants of psychology students and participants with experience in psychology experiments, allowing the test of the effects of both direct and indirect exposure to deception.

In total, 146 males participated in the experiment.[2] For exactly half of the participants, this was the first time participating in an experiment at the Experimental Economics Laboratory. These participants were required to read and sign the laboratory rules during the experiment (in addition to accepting the rules online when registering to the subject pool). The rules include the following item, which may have contributed to the generally high level of trust we observe in the experiment:

[2] Females were excluded because the main experiment involved the nasal administration of the neuropeptide oxytocin.

Table 1: Statements used in the experiment.

1	Archimedes was born before Sophocles
2	The first Harry Potter book was translated to more than 60 languages
3	The Titanic's disaster more than 1400 people were killed
4	The planet Saturn is closer to the sun than the planet Jupiter
5	Princess Diana was killed at age 34
6	Simon and Garfunkel were born on the same year
7	The Russian novel "Anna Karenina" was published before 1860
8	The Eiffel tower has more than 1600 stairs
9	Madonna's first album was released in 1980
10	A rat can survive longer than a camel without water
11	The most expensive painting in the world was painted by Van Gogh
12	The shortest period of presidency by a president of the U.S.A. was one month
13	The German alphabet has 30 letters
14	Hugh Hefner has been married three times

> The rules of the experiment, including the rules for determining the payoff for participants, will be presented to participants in the experimental instructions. The rules are always accurate and reflect the true procedure of the experiment. The laboratory directors make sure that the instructions are never misleading nor present the experimental procedure in a deceptive way. If it appears to any participant that the instructions are misleading in any way, he or she are asked to report it to the laboratory team directly or by mail at experiments.bgu@gmail.com.

Twenty four (16.44%) of the participants were listed in the subject pool as psychology majors. Fifty eight participants (39.73%) were registered to the management department subject pool. Fifty one participants (34.93%) participated in at least one experiment in the psychology department. Of those, 47 participants were registered for the paid experiments pool, and the other four for the course credit pool.[3]

[3] The faculty of management does not keep actual participation statistics. For the department of psychology I matched subjects on both name and email and ignored partial matches. The results are qualitatively robust for using matches based only on one of the two criteria. Registration to the two subject pools was not significantly correlated.

The experiment was comprised of two stages. In the first stage, the participants saw the statements presented in Table 1 on a single screen, and were asked to indicate the probability they assign for each statement to be true.[4] Next, the participants repeated the task by pairs of statements under the instruction that exactly one sentence in each pair is true. Finally, the participants indicated whether they...[5]

1. believed that exactly one statement is true in each pair (trust).
2. were not certain that exactly one sentence is true in each pair (suspicion).
3. did not believe that exactly one statement is true in each pair (distrust).

Participants were randomly allocated into two payment treatments. The payment in the flat fee treatment was 10 NIS (approximately $2.60 at the time of the experiment) regardless of performance. Payment in the per performance treatment was determined by the quadratic scoring rule $1-(t-p)^2$, where $t \in \{0,1\}$ indicates whether the statement is actually true or false. The mean payoff in this treatment was 7.87 NIS, with a standard deviation of 1.39.

4 Results

Responses to the probability assessment tasks had enough variance to estimate the trust parameter. The standard deviation of the prior probabilities overall was 32.55%, with 74.11% of reported probabilities strictly between 0 and 1. If the participants believe and incorporate the information provided, then the posteriors in each pair should sum to 1. Participants clearly responded to the information, as posteriors sum to 1 in 82.1% of all pairs (compared to 18.2% of prior pairs). Another indication of correct updating is that the mean absolute difference between the two probabilities increase from 34.3 percentage points for the prior to 66.8 percentage points for the posteriors, in line with Bayesian updating regardless of trust. Accuracy payment did not improve performance in the first stage. Applying the CLP framework of experimental incentives (Camerer and Hogarth, 1999), this is taken to mean that—given the task properties and participants' abilities—incentivized effort does not improve performance. Therefore the measurement of trust can be assumed to not be confounded with motivation or deliberation. I thus proceed to discussing the results of the trust estimation.

The first row in Table 2 reports estimates of overall trust levels. A global maximum-likelihood non-linear regression estimated the trust parameter at 0.947. Such a high value indicates both that the participants trusted the instructions and were able to

[4] The 14 statements were selected out of a longer list piloted to assure that most respondents are uncertain of their veracity, i.e. provide internal probabilities.

[5] The debriefing made it clear that no deception was actually used in the experiment.

Table 2: Reported and estimated trust.

Reported trust	Estimated trust	95% confidence interval	Mean estimated trust	95% confidence interval
All	0.947	[0.934 – 0.960]	0.880	[0.850 – 0.910]
	(0.007)		(0.015)	
	N = 144		N = 137	
Trust	0.986	[0.976 – 0.996]	0.914	[0.881 – 0.946]
	(0.005)		(0.017)	
	N = 103		N = 98	
Suspect	0.868	[0.826 – 0.909]	0.836	[0.779 – 0.894]
	(0.021)		(0.029)	
	N = 34		N = 32	
Distrust	0.598	[0.500 – 0.696]	0.591	[0.469 – 0.714]
	(0.050)		(0.062)	
	N = 7		N = 7	

Standard errors in parentheses. Standard errors for mean estimated trust in the second to fourth rows are obtained by regressing the individually estimated trust on dummy variables for self-reported trust. N indicates the number of participants.

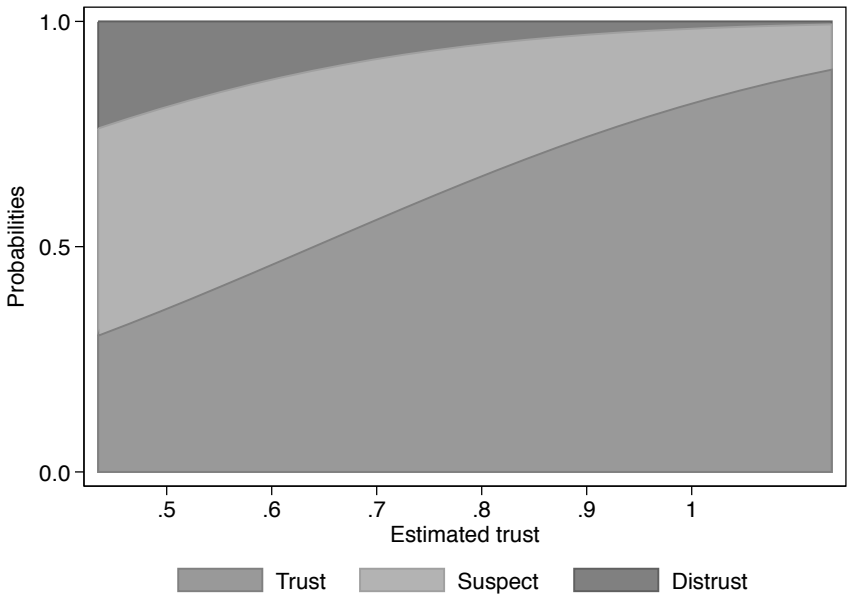

Figure 1: Predicted reported trust based on estimated trust.

approximate Bayesian updating. Next, the trust parameter was estimated for each participant individually, based on the fourteen equations implied by the seven statement pairs.[6] The mean estimated trust thus obtained is 0.880.

The rest of Table 2 reports the estimated trust separately by the three self-reported trust categories. Participants who reported not being certain in the veracity of the instructions are indeed estimated to be less trusting. The difference to the estimated trust for participants who reported explicit distrust is even more pronounced, with trust levels less than 0.6, and not clearly significantly above the disbelief baseline 0.5.

Figure 1 plots the predictions from an ordered probit regression of the self-reported trust on estimated trust. The model predicts that, as the estimated trust increases, the probability of reporting trust decreases while the probabilities of reporting suspicion and distrust decrease. The predicted probability of reporting trust increases from 36.2% when the estimated trust is 0.5 to 81.8% when the estimated trust is 1. Similarly, the predicted probabilities of reporting suspicion and distrust decrease from 44.8% and 19.0% to only 16.6% and 1.6%, respectively.

[6] Estimation failed for 7 participant who indicated the same posterior in all statement pairs, leaving 139 observations. Of these, 21 (15%) have $r > 1$. I refrained from censoring at 1 under the assumption that there is a symmetric measurement error (26 participants have $0.99 < r < 1$), and censoring would introduce a systematic bias.

The correspondence of the estimated and self-reported trust measures validate the new measure. Furthermore, the results are consistent with the conjecture, following Newberry (1973) and Taylor and Shepperd (1996), that participants under report their disbelief, as many participants who report no suspicion ignore the new information. This interpretation should be taken cautiously, however, as (part of) this result can be explained by partial updating or lack of motivation.

4.1 Correlates of trust

Table 3 reports regressions of estimated trust on the different predictors. The first column only includes the constant term, which is identical to that in the first row of Table 2. The second column adds the treatment variable indicating whether participants were paid according to accuracy. In contrast to the ex-ante hypothesis, we see that participants are approximately 9 percentage points less likely to believe a statement if it is payoff-relevant.[7]

Thus, participants may suspect that the experimenter is more likely to deceive to save money—the major motivation for, e.g., using bogus partners—or possibly participants are more alert, and therefore more likely to be sceptic about new information, when it comes to the determination of their payoff.

The third row tests the effects of indirect exposure to deceptive practices through psychology studies. Psychology students are 8% less likely to be suspicious in their first year, but 3.5% more likely to be suspicious later on. However, these effects and their interaction are not close to being significant. The results, therefore, do not support the hypothesis that indirect exposure undermines trust.

The rest of Table 3 relates to direct exposure. Deceptive practices are sometimes employed in both management and psychology departments of the university. Nonetheless, being registered to the subject pool of neither department is correlated with trust. The number of psychology experiments participated in is, however, significantly correlated with trust, although the effect size is very small. The pattern of the data suggests that the direction of the correlation flips, as is illustrated in Columns (6)–(7), which restrict the analysis to participants with a history of less than 20 or more than 10 psychology experiments, respectively. A piecewise linear regression tested this hypothesis formally. The model included an intercept, a cut-off point, and slopes below and above the cut-off point. The results estimated the intercept at 0.875 ($SE = 0.018$), the initial slope at 0.010 ($p = 0.339$), the cut-off point at 8 experiments ($p = 0.254$), and the slope above the cut-off point at -0.003 ($p < 0.01$). Thus, it appears that participation in a few experiments at the psychology department does not decrease—and perhaps even increases—trust,

[7] Note that, as trust and accuracy are somewhat confounded, paying for accuracy may increase estimated trust by increasing motivation rather than by reducing suspicion. Paying for accuracy, however, did not improve and even harmed the mean performance in the first stage ($p < 0.05$).

Table 3: OLS regressions on correlates of trust.

	(1) Baseline	(2) Payoff	(3) Studies	(4) Experiment history	(5) Experiment history	(6) Experience < 20	(7) Experience > 10
Trust	0.880	0.924	0.875	0.876	0.892	0.876	0.968
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Payoff relevant		-0.089					
		(0.003)					
Psychology major			-0.035				
			(0.546)				
First year			-0.005				
			(0.878)				
Psychology × First year			0.115				
			(0.162)				
Management				0.006			
				(0.849)			
Psychology				0.005			
				(0.867)			
# Experiments					-0.001	0.008	-0.003
					(0.053)	(0.066)	(0.030)
Observations	139	139	132	139	139	120	26

Note: p-values in parentheses.

whereas participation in many experiments slightly decreases trust. This may be due to the low share of experiments that employ deception at the psychology department.

Interestingly, the correlation with past participation in psychology experiments does not exist for self-reported trust. If at all, participants with more extensive experience with psychology experiments are *less* likely to report suspicion. Whereas 30 out of the 93 (32.3%) participants who did not participate in any experiment at the psychology department reported suspicion or distrust, only 11 (21.6%)of the remaining 51 participants and 4 of the 25 (16.0%) who participated in more than 12 psychology experiments did so. Although these differences are not significant and may be confounded with selection effects, clearly the self reports are less sensitive than the implicit measure. Furthermore, when controlling for the self report, the correlation between estimated trust and past participation in psychology experiments is highly significant ($p < 0.01$).

5 Conclusion

This paper presents a simple exercise in estimating trust using a Bayesian method. The measure is highly correlated with self reports, yet is free of problems of social desirability and under reporting. It is based on a short questionnaire that can be administered in less than five minutes, in the laboratory or online,[8] and can generate estimates of trust on both individual and group levels. Thus, the main contribution of this paper is to provide experimentalists with a simple tool to study issues of deception and trust.

I tested whether participants are less suspicious of experimental instructions if these are payoff relevant. Surprisingly, the result is the opposite. Possibly participants are weary of being fooled, but are willing to believe information when not exposing themselves to loss by believing.

I also use the data to explore issues of direct and indirect exposure to deception. I find no evidence for the effect of indirect exposure, and a mild effect for direct exposure. In view of the latter result, the lack of consistent findings in the literature can be understood. The effect of deception appears to be small and cumulative, hence detecting the effect of a single experience, as attempted, e.g., by Jamison et al. (2008) and Barrera and Simpson (2012), is difficult.[9]

When considering the policy of no deception in experimental economics, the results provide justification for the policy used by the American Psychological Association, which states that deception should be used as a last resort, only if "ef-

[8] Care should be taken that respondents do not search for the answers, which was blocked on the laboratory computers during the experiment.

[9] The current study is lacking compared to the previous studies in that it utilizes existing variation in exposure to deception rather than experimentally manipulating it—which would violate the laboratory rules.

fective nondeceptive alternative procedures are not feasible" (American Psychological Association, 2002, article 8.07). If deception is kept rare and far in between, participants can learn that experimental instructions are typical valid, and the pollution of the public good that is trust is negligible. Extensive use of deception, as developed in social psychology in the 1960's and 1970's, may well harm the experimental control over participants' beliefs, as the cumulative effects of deception come into play. Nonetheless, I would not advocate relieving the existing taboo on deception in experimental economics without placing practical restrictions on gratuitous deception—after all, deception in psychology was widespread until recent years despite the APA guidelines—and without collecting more evidence on the moderators of trust. The tool presented in this paper provides a way to further our understanding of which experimental practices can promote participants' trust and which practices substantially undermine it.

References

Alberti, F. and W. Güth (2013). Studying deception without deceiving participants: An experiment of deception experiments. *Journal of Economic Behavior & Organization 93*, 196–204.

American Psychological Association (2002). Ethical principles of psychologists and code of conduct. *American psychologist 57*(12), 1060–1073.

Barrera, D. and B. Simpson (2012). Much ado about deception consequences of deceiving research participants in the social sciences. *Sociological Methods & Research 41*(3), 383–413.

Camerer, C. F. and H. M. Hogarth (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty 19*(1), 7–42.

Cook, T. D. and D. T. Campbell (1979). *Quasi-experimentation: Design & analysis issues for field settings*. Houghton Mifflin Boston.

Davis, D. D. and C. A. Holt (1993). *Experimental economics*. New Jersey: Princeton university press.

Hertwig, R. and A. Ortmann (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences 24*(03), 383–403.

Hertwig, R. and A. Ortmann (2008a). Deception in experiments: Revisiting the arguments in its defense. *Ethics & Behavior 18*(1), 59–92.

Hertwig, R. and A. Ortmann (2008b). Deception in social psychological experiments: Two misconceptions and a research agenda. *Social Psychology Quarterly 71*(3), 222–227.

Jamison, J., D. Karlan, and L. Schechter (2008). To deceive or not to deceive: The effect of deception on behavior in future laboratory experiments. *Journal of Economic Behavior & Organization 68*(3), 477–488.

Kelman, H. C. (1967). Human use of human subjects: the problem of deception in social psychological experiments. *Psychological Bulletin 67*(1), 1.

Krawczyk, M. (2013). Delineating deception in experimental economics: Researchers' and subjects' views. *mimeo*.

Krawczyk, M. (2015). "trust me, i am an economist." a note on suspiciousness in laboratory experiments. *Journal of Behavioral and Experimental Economics 55*, 103–107.

McKenzie, C. R. and J. T. Wixted (2001). Participant skepticism: If you can't beat it, model it. *Behavioral and Brain Sciences 24*(03), 424–425.

McKenzie, C. R., J. T. Wixted, and D. C. Noelle (2004). Explaining purportedly irrational behavior by modeling skepticism in task parameters: an example examining confidence in forced-choice tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition 30*(5), 947.

Newberry, B. (1973). Truth telling in subjects with information about experiments: Who is being deceived? *Journal of Personality and Social Psychology 25*, 369–374.

Ortmann, A. and R. Hertwig (1997). Is deception acceptable? *American Psychologist 52*, 746–747.

Ortmann, A. and R. Hertwig (2002). The costs of deception: Evidence from psychology. *Experimental Economics 5*(2), 111–131.

Smith, V. L. (1982). Microeconomic systems as an experimental science. *American Economic Review 72*(5), 923–955.

Taylor, K. and J. Shepperd (1996). Probing suspicion among participants in deception research. *American Psychologist 51*, 886–887.