DISCUSSION PAPER SERIES

No. 9384

DECEPTION DETECTION AND THE ROLE OF SELF-SELECTION

Kai A. Konrad, Tim Lohse and Salmai Qari

PUBLIC POLICY

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP9384.asp**

ISSN 0265-8003

DECEPTION DETECTION AND THE ROLE OF SELF-SELECTION

Kai A. Konrad, Max Planck Institut for Tax Law and Public Finance, and CEPR
Tim Lohse, Berlin School of Economics and Law, and Max Planck Institute for Tax Law and Public Finance
Salmai Qari, Max Planck Institute for Tax Law and Public Finance

Discussion Paper No. 9384
March 2013

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 9384

March 2013

ABSTRACT

Deception detection and the role of self-selection*

We consider a lie-catching experiment with 9240 judgements. A set of videotapes shows subjects participating in a tax compliance experiment. The subjects chose whether or not to misreport. Subjects knew that underreporters were chosen for an audit with some probability. An audit led to detection and to a punishment fee. This compliance framework induced only persons with high deceptive abilities to underreport and, so, caused self-selection. Among the students who judged these videos, we find that the deception detection rate was significantly below 50 percent and even lower if the self-selection pressure in the tax compliance experiment was higher. This suggests that, when subjects can choose whether to state the truth or to lie, there is a self-selection effect by which individuals with higher deceptive ability are more likely to lie.

JEL Classification: D83 and H26
Keywords: decision making, interpersonal interaction, judgment and perception

Kai A. Konrad
Max Planck Institute for Tax Law and Public Finance
Marstallplatz 1
80539 Munich
GERMANY

Email: kai.konrad@tax.mpg.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=116703

Tim Lohse
Berlin School of Economics and Law
Alt-Friedrichsfelde 60
10315 Berlin
GERMANY

Email: tim.lohse@hwr-berlin.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=176920

Salmai Qari
Max Planck Institute for Tax Law and
Public Finance
Marstallplatz 1
80539 Munich
GERMANY

Email: salmai.qari@tax.mpg.de

*We thank Hans Müller for programming the experimental environment and Nina Bonge for excellent research assistance. The usual caveat applies.

Submitted 04 March 2013

Research on lie-catching and deception detection has established important results in more than 200 experiments and has been well documented in a series of meta-studies.[1] One key finding of this large stock of evidence is that individuals have rather limited deception detection abilities overall. They typically do better than pure chance, albeit not by very much (Hartwig and Bond, 2011). Vrij (2008) reports accuracy estimates that range from about 42% to 68% and an average accuracy of around 54% among nonprofessionals. In their meta-study, DePaulo et al. (2003, p. 106) list three major types of critiques of the existing literature: (i) participants typically tell the truth or lie because they are told to do so; (ii) participants typically are not strongly motivated; (iii) participants typically do not get relevant feedback on their behavior. We report here about a lie-catching experiment with 9240 judgements. This experiment departs from earlier experiments along several dimensions and thereby addresses the three points of criticism raised by De Paulo et al. (2003) (see next section for details). Our subjects could self-select whether to lie or not. They did so under two different incentive regimes in case they were caught, one with low and one with high fine. This main innovation allows us to analyse whether or not choice and the implied selection matter. We ask: Is the share of subjects correctly classified as liars lower than the consensus estimates from the meta-studies? Is it even lower than 50 percent? And is the detection rate even lower for videoclips that show underreporters in a regime with high fine?

1 Method

Our experiment takes into account the critiques by DePaulo et al. (2003) and is designed as follows:

(i) The videotapes with honest and dishonest persons are generated in a tax compliance framework in which the decision whether to lie or not was the *endogenous* choice of the persons declaring their taxes. They faced a compliance decision in which underreporting had a monetary reward only if it remained undetected, and in which underreporting was punished if it was detected. Whether a subject chooses to underreport in this set-up should therefore depend on whether the subject assesses it to be likely that it will successfully deceive the tax officer, and, if subjects are correct in their beliefs, this should

[1] The idea to detect deception on the basis of nonverbal expressions of emotion which are not as easily censored or disguised as the content of speech dates back to Charles Darwin (1872), Cesare Lombroso (1876) or Sigmund Freud (1959). A milestone in the experimental work on lie-detection is by Ekman and Friesen (1974). We do not survey the work here, but refer to surveys and metastudies by DePaulo et al. (2003), Bond and DePaulo (2008), Vrij (2008) and Hartwig and Bond (2011).

be reflected in the judgements. Due to this self-selection, the "quality of liars" is different as compared to the standard setting in the literature in which all subjects who lie are forced, or advised, to lie. This is the basis for our first hypothesis: our subjects should be less frequently correctly identified as liars than those emerging from the meta-studies. Moreover, the specific set-up of the compliance videoclips allows us to sharpen the prediction and formulate a second hypothesis. Our videoclips of persons who underreport come from two different compliance treatments, one with a high penalty and one with a low penalty for underreporters who are caught (see Konrad, Lohse and Qari 2011 for further details). The penal tax being low in one treatment and high in the other provides different incentives for self-selection. Average truthful reporting was about 35.75 percent in the low-fine treatment and about 66.69 percent in the high-fine treatment. Presumably, the group of tax payers who chose to underreport in the high-tax treatment is a more selective group and contains subjects who believe that they have particularly high deceptive abilities. Therefore, as a second hypothesis, we expect that the judges' ability to detect underreporters from this more selected group is even lower than for the group of tax payers who underreported in the low-penalty framework.

(ii) The persons who lie when making tax compliance decisions are motivated by monetary incentives. If they perform poorly (i.e., if their untruthful report is detected), they pay a fine. Also, the subjects who make judgements about the videotapes have incentives to try hard as they receive monetary rewards for correct assessments.

(iii) Participants in the compliance framework and subjects who judge the videotapes have several types of feedback. Participants in the compliance experiment who make a compliance decision get strong and direct feedback. They anticipate that their counterpart in the compliance interview rates their reply as more or less honest, and they learn indirectly about this assessment because the assessment affects their likelihood of receiving an audit. The subjects who judge the videos also receive a feedback because they receive a monetary reward for correct judgements.

Our data set is based on 2 sets of videotapes, each containing 40 individual declarations (produced in the context of an economic experiment on tax compliance at the MELESSA, the experimental laboratory at the University of Munich in 2012[2]), and 231 subjects (recruited from diverse fields of study at Technische Universität Berlin[3]) judging 40 videotapes each. A set of 40 compli-

[2]For details on this experiment see Konrad, Lohse and Qari (2011). All participant persons in Munich who acted as tax payers were recruited using the software ORSEE (Greiner, 2004).

[3]All participants of this assessment in Berlin were recruited using the software ORSEE (Greiner, 2004). Choosing from the subject pool in Munich for the tax-compliance game and

ance videos is composed of 20 videos showing truthful reports from tax payers and 20 with subjects who choose to underreport. Moreover, each subset of 20 videos is composed of 10 videos from the low-penal-tax regime and 10 from the high-penal-tax regime. The videos were shown in random order.[4]

The subjects in the lie-catching experiment, the "judges", were told that they will see a sequence of 40 clips with tax compliance dialogues on their computer screens and were informed roughly about how these videos were produced and what they show. They were told that the share of truthful reports among the 40 videos was roughly one half (in fact, it was *exactly* one half). After each video, they had to make the judgement on whether the tax payer was more likely to have made a truthful report or more likely to have made an untruthful report. They were not allowed to return to previous videos and alter their judgement. We recorded these judgements. They were rewarded for correct assessments.[5] This procedure led to the 9240 judgement data, which are the basis for our quantitative statistical analysis.

2 Results

Consider first the descriptive statistics. The fraction of correctly classified liars is equal to 47.66% in the low penalty condition and 45.71% in the high penalty condition. Since the two penalty conditions are equally likely, the overall detection rate is equal to 46.47%, the average of the two former rates. We also used generalized (probit) linear mixed-effects models (GLMM) to predict the probability that a video is classified as showing a lie.[6] We enter judge-specific and video-specific random effects into the models to account for the effects that the same judge assessed multiple videos, and that the same video was assessed by several judges. In Model 1, the main explanatory variable is a dummy indicating whether or not the videotaped subject was lying (i.e., had underreported). This model directly provides an estimate for the average detection rate of liars.

from the subject pool in Berlin for the deception detection experiment, we could reasonably rule out that subjects making judgements on the tax-payer clips would be acquainted with the tax payers they judge.

[4] These 80 videotapes were a selection from the whole set of compliance videosets generated by Konrad, Lohse and Qari (2011). Random selection was used to generate the different subsets of 20, 10 and 10. One set was shown to 120 subjects; the other set was shown to 111 subjects.

[5] Out of the 40 assessments by a judge, the computer randomly selected five rounds for payment. Judges were paid 5 EUR for each correct assessment. Judges also received a show-up fee of 5 EUR. Realizations of final payments were between 5 and 30 EUR with an average of 17.99 EUR (SD=6.23).

[6] We use the R Environment (R Core Team 2012) and, in particular, the lme4 package (Bates, Maechler and Bolker 2012) to fit the models.

Model 2 investigates whether the detection rate differs across the two penalty conditions. To this end, it extends Model 1 by entering a dummy indicating high penalty and a term interacting this dummy with the dummy for underreporting. As a further robustness check, we calculate, for each of the 231 judges, the share of liars that were detected for the two penalty conditions, collapsing essentially all observations for a given judge and a given penalty condition into one variable. We use linear mixed models (LMM) to analyse these shares and enter judge-specific random effects into the model. Model 3 simply predicts the average share of detected liars while Model 4 controls for penalty conditions.

The descriptive evidence is confirmed by the generalized linear mixed models. Model 1 shows that liars are successful since underreporting high-endowment subjects are more often classified as reporting truthfully compared to honestly reporting low-endowment subjects. In particular, the probability of being classified as dishonest is five percentage points smaller for the average underreporting high-endowment-subject (z=-1.738, $p < 0.1$) compared to subjects who complied honestly, yielding a probability of 46.37% (cf. the descriptives). Model 2 estimates that the probability of being classified as dishonest for the average underreporting subject is equal to 47.71% and 45.04% in the low and high penalty condition, respectively. Once again, these numbers match the descriptive statistics, and the differences in the model selection criteria indicate that Model 2 provides a significant improvement compared to Model 1 ($\chi^2 = 4.7279$, $p < 0.1$, $\text{AIC}_2 < \text{AIC}_1$). These results are robust to other model specifications. The linear mixed-effects model (Model 3) indicates that the fraction of liars detected overall is equal 46.47%, which is significantly smaller than that detected by random chance ($t = -4.9$, $p < 0.001$).[7] Model 4 indicates that the fraction of liars detected for low penalty videos is 47.66%, which is 2.34 percentage points below random chance. The deviation from 50% is significant ($p = 0.02$). The detection rate for high penalty videos is 2.38 percentage points smaller than the detection rate for the low penalty videos ($t = -1.69$, $p < 0.1$) and therefore 4.72 percentage points below chance.

Summarizing, all models indicate that the detection rate is significantly below 50% [8] and that the detection rate in the high penalty condition is smaller compared to the low penalty condition.

[7] As the degrees of freedom are sufficiently high, we use the normal distribution to calculate p-values from the t-statistics.

[8] This result is surprising from an economist's point of view. Students who made judgements here could have applied a simple but superior strategy: they could have used a random mechanism or simply have deterministically judged every second videoclip as honest etc. This strategy would have given them a higher expected return. But, in fact, they must have tried sincerely - and unsuccessfully - to extract information from the videoclips they saw to help them make their judgements.

3 Discussion

Our behavioral model suggests that individuals make choices about whether or not to tell the truth, and, in it, their decision to underreport (i.e., use deception) in a tax compliance decision is positively correlated with their self-assessed deceptive abilities. Individuals who choose to lie are a self-selected group. Their choice should be driven by their deceptive abilities and by the size of the fine in case their deception is detected. Higher fines make it attractive to choose deception only for individuals who think their deceptive abilities are very high. Our findings are in line with this hypothesis. We find that individuals who choose deception are successful on average. Their detection rate is lower than what could be achieved by a random audit device. Moreover, individuals who chose to underreport in the treatment regime with high penal taxes were even more successful on average. This is evidence in line with this behavioral theory.

This evidence leads to the question about what the basis on which subjects may judge their own deceptive abilities is. The question is related to a key question in the context of the lie-catching literature that tries to trace what the cues that an individual may use when trying to detect deceptive behavior are.[9] Our research result does not speak directly to the research that tries to identify and trace cues, and our findings are compatible both with the existence or the non-existence of such informative cues. Our experiment suggests that subjects make a judgement about their own deceptive abilities, but the experiment is silent about the basis of the subjects' judgements about their own deceptive abilities. For a subject to have an assessment of their own deceptive abilities that is informative, it is not required that the subject knows or understands the process or the cues by which their judges assess them. Likewise, information about their ability to control their own performance is not required. A person's assessment of own deceptive ability may simply be based on their own personal life experience and their previous successes and failures in deceiving other people.

4 Conclusion

In a compliance decision, a person must decide whether to report truthfully or to lie and misreport. Motivational factors for this decision should be the benefit of successful deception and the cost associated with a detected deception attempt, in comparison to the outcome in the case of truthful report, and the likelihood

[9] In fact, much of the research on deception detection can be seen as a search for what exactly the cues that are used for lie-catching are and whether these cues have explanatory power (see, for instance, the surveys by DePaulo et al. (2003) and Hartwig and Bond (2011)).

of detection of successful deception. Persons who feel confident that they are successful liars should, hence, be less likely to tell the truth in a tax compliance experiment. This reasoning suggests that choice and the self-selection implied are important aspects for lie-catching in a natural compliance environment: the set of liars may consist of persons who think that they have good deceptive abilities, and increasingly so if the fine in case of detection is higher.

Our experiment tests an implication of this behavioral model. Using a large data set, we find that liars whose deception is the outcome of their own choice are less frequently detected than liars in standard experiments. This effect is stronger in our data if the incentives for tax payers to report truthfully are stronger. The shares of correctly classified liars are 47.66% and 45.28% in the low- and high-fine treatments, respectively, and are significantly below random chance (50%). This is support for our theory.

References

[1] Bates, Douglas, Martin Maechler and Ben Bolker (2012). lme4: Linear mixed-effects models using S4 classes. R package version 0.999999-0. http://CRAN.R-project.org/package=lme4.

[2] Bond, Charles F., and Bella M. DePaulo (2008). Individual differences in judging deception: Accuracy and bias, *Psychological Bulletin* 134(4): 477–492.

[3] Darwin, Charles (1872). *The expression of the emotions in man and animals,* London: John Murray, Albemarle Street.

[4] DePaulo, Bella M., James J. Lindsay, Brian E. Malone, Laura Muhlenbruck, Kelly Charlton, and Harris Cooper (2003). Cues to deception, *Psychological Bulletin* 129(1): 74–118.

[5] Ekman, Paul, and Wallace V. Friesen (1974). Detecting Deception from Body or Face, *Journal of Personality and Social Psychology* 29(3): 288–298.

[6] Freud, Sigmund (1959). Fragment of an analysis of a case of hysteria. In: Collected papers, 3, New York: Basic Books.

[7] Greiner, Ben (2004). An online recruitment system for economic experiments, University Library of Munich, MPRA Paper Nr. 13513.

[8] Hartwig, Maria, and Charles F. Bond (2011). Why Do Lie-Catchers Fail? A Lens Model Meta-Analysis of Human Lie Judgments, *Psychological Bulletin* 137(4): 643–659.

[9] Konrad, Kai A., Tim Lohse, and Salmai Qari (2011). Compliance and the Subjective Audit Probability, Max Planck Institute for Tax Law and Public Finance Working Paper 2011–18.

[10] Lambroso, Cesare (1876). *L'uomo delinquente. In rapporto all'antropologia, alla giurisprudenza ed alle discipline carcerarie*, Turin: Bocca.

[11] R Core Team (2012). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org/.

[12] Vrij, Albert (2008). *Detecting Lies and Deceit*, 2nd ed. Chichester: Wiley.