Journal of Public Economics 113 (2014) 96–104

Contents lists available at ScienceDirect

Journal of Public Economics

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

Representative evidence on lying costs ☆

Johannes Abeler [a,*], Anke Becker [b], Armin Falk [b]

[a] University of Oxford, United Kingdom
[b] University of Bonn, Germany

ARTICLE INFO

Article history:
Received 6 December 2012
Received in revised form 2 January 2014
Accepted 7 January 2014
Available online 30 January 2014

JEL classification:
C93
D01
D82
H26
H51

Keywords:
Private information
Lying costs
Tax morale
Representative experiment

ABSTRACT

A central assumption in economics is that people misreport their private information if this is to their material benefit. Several recent models depart from this assumption and posit that some people do not lie or at least do not lie maximally. These models invoke many different underlying motives including intrinsic lying costs, altruism, efficiency concerns, or conditional cooperation. To provide an empirically-validated microfoundation for these models, it is crucial to understand the relevance of the different potential motives. We measure the extent of lying costs among a representative sample of the German population by calling them at home. In our setup, participants have a clear monetary incentive to misreport, misreporting cannot be detected, reputational concerns are negligible and altruism, efficiency concerns or conditional cooperation cannot play a role. Yet, we find that aggregate reporting behavior is close to the expected truthful distribution suggesting that lying costs are large and widespread. Further lab experiments show that this result is not driven by the mode of communication.

1. Introduction

Situations with asymmetric information are ubiquitous. Most of economic theory assumes that people misreport their private information if this is to their material benefit; behavior is only determined by the trade-off between financial gains from misreporting and monetary fines when misreporting is detected.[1] In contrast, many recent models in various domains of Public Economics (and in Economics more generally) rely on the assumption that people can experience a psychological disutility which holds them back from misreporting, at least to some extent. These models invoke different underlying motives. Kartik et al. (2014), for instance, assume that people face an intrinsic lying cost and show that in this case the social planner can fully implement a much wider range of social choice rules compared to the standard Maskin (1977) case without lying costs (see, e.g., Matsushima (2008) and Dutta and Sen (2011) for similar assumptions). Many studies about incentive systems for doctors assume that doctors are altruistic towards their patients and thus do not always state the profit-maximizing diagnosis but rather treat patients honestly (e.g., Ellis and McGuire, 1986; Chalkley and Malcomson, 1998). The large literature on "tax morale" (e.g., Lewis, 1982; Cowell, 1990; Andreoni et al., 1998; Slemrod, 2007; Torgler, 2007) demonstrates that many tax payers misreport their income only a little bit or not at all. This literature is usually agnostic about the exact underlying motives but some studies cite efficiency concerns (e.g., Alm et al., 1992), patriotism (Konrad and Qari, 2012), religiosity (Torgler, 2006), fairness (Bordignon, 1993), conditional cooperation (Traxler, 2010) or honesty (Erard and Feinstein, 1994).

To further improve these models and to provide an empirically-validated microfoundation, it is crucial to understand the relevance of the different potential motives. Additionally, understanding these motives could inform the design of more psychologically-realistic policies, e.g., in the area of tax enforcement, that have a higher potential of being successful. In this paper, we focus on intrinsic lying costs and investigate how widespread and how large lying costs are. The ideal data set to answer these questions would allow studying lying costs for a representative sample of the population and in an environment without the confounding effects of strategic interaction (including the

☆ Financial support from the German Science Foundation (SFB/TR 15), the European Research Council (Starting Grant) and the Economic and Social Research Council (ES/K001558/1) is gratefully acknowledged. We thank Steffen Altmann, Vince Crawford, Alain Cohn, Eddie Dekel, Holger Gerhardt, David Huffman, Sebastian Kube, Michel Maréchal, Theo Offerman, Collin Raymond, Paul Schempp, Christian Traxler, Matthias Wibral, and Florian Zimmermann for helpful discussions. Valuable comments were also received from numerous seminar and conference participants.

* Corresponding author.
E-mail addresses: johannes.abeler@economics.ox.ac.uk (J. Abeler), anke.becker@uni-bonn.de (A. Becker), armin.falk@uni-bonn.de (A. Falk).

[1] See, e.g., Allingham and Sandmo (1972) on tax evasion, Falkinger (1991) on public good provision, Pitchik and Schotter (1987) on credence goods, along with the seminal Becker (1968) on crime.

http://dx.doi.org/10.1016/j.jpubeco.2014.01.005

levy of fines), reputational or efficiency concerns, or altruism. So far, the best evidence on lying costs comes from experiments conducted in tightly controlled laboratory situations. A robust result is that many subjects misreport their private information to their own advantage but that a substantial share of subjects refrains from reporting the payoff-maximizing type and that some are fully honest (e.g., Gneezy, 2005; Charness and Dufwenberg, 2006; Fischbacher and Föllmi-Heusi, forthcoming; de Haan et al., 2011; Houser et al., 2012; Shalvi et al., 2011; Wibral et al., 2012; Serra-Garcia et al., 2013). These studies are a strong first indicator that lying costs influence behavior. However, lab experiments do not allow for inferences with respect to the prevalence of lying costs in the overall population since they have been conducted almost exclusively with student samples (DellaVigna, 2009; Falk and Heckman, 2009). Also, decision making took place in an austere laboratory environment which might trigger behavior representative only of certain non-lab situations. It could thus be that there are systematic differences between behavior of students in the laboratory and behavior of non-student subjects outside the lab.

To circumvent these limitations, we measure how people report their private information outside the laboratory by calling participants on the phone at their home. Participants were drawn randomly from the German population, yielding a representative sample. An incentivized experiment was embedded in the interview. The experimental setup is related to the design of Fischbacher and Föllmi-Heusi (forthcoming) and is extremely simple: participants were asked to toss a coin and report their type, i.e., either "heads" or "tails". Reporting tails yielded a payoff of 15 euros, which participants could choose to receive in cash or as an Amazon gift certificate, while reporting heads yielded a payoff of zero. Participants thus had a clear monetary incentive to report tails regardless of their true type. It was obvious that the true outcome was only known to the participants, as they tossed the coin privately at home. In this setup, we cannot draw reliable conclusions about the truthfulness of any individual report. But we can learn about aggregate behavior by comparing the distribution of reports to the true distribution of a fair coin (50% tails) and to the payoff-maximizing distribution (100% tails). This indirect observation therefore allows us to study the behavior of subjects in a situation in which private information is kept truly private and in which subjects do not face any risk of detection.[2] Moreover, the decision is non-strategic; altruism does not play a role as the money is not taken from any individual person; and reputational concerns are minimized since the interviewer is a stranger with whom no future interaction can be expected.

If all our participants were rational money maximizers, we would expect that all of them reported tails. If behavior on the phone was similar to previous comparable laboratory experiments (e.g., Houser et al., 2012), we would expect about 75% of subjects reporting tails.

In contrast to these predictions, observed behavior does not statistically differ from everybody reporting honestly. If anything, participants report the payoff maximizing outcome less often than expected under truthful reporting. This latter effect, however, is small and disappears in a second treatment in which participants were asked to report the total number of tails in four consecutive coin tosses and received 5 euros times the number of reported tails. The resulting distribution of reports in the 4-Coin Treatment is indistinguishable from the distribution under complete truth-telling. Moreover, while previous studies (e.g., Dreber and Johannesson 2008) have found correlations between individual characteristics, like gender, and truth-telling, we do not find any robust correlations between individual characteristics and reporting behavior. This is not surprising if almost all participants report truthfully. Reports are solely determined by chance, namely the coin toss, which cannot be related to any individual characteristic. Our results thus show that lying costs are pervasive and are influencing behavior regardless of gender, religious beliefs, education, or age.

We complement our telephone study with two additional control treatments in the laboratory to better understand what shapes lying costs, in particular the effect of the mode of communication. In both lab treatments subjects reported the outcomes of four consecutive coin tosses. Incentives were the same as in the 4-Coin Treatment in the telephone study: 5 euros times the numbers of tails reported. In the first lab treatment, subjects had to report the outcome directly to an interviewer via the phone, mirroring our telephone study. We observe the same pattern of behavior as in previous lab experiment: subjects lie much more than in the telephone study. In the second control treatment, subjects reported the outcomes by clicking a number between 0 and 4 on the computer screen as in most previous lab experiments. We find that subjects who enter their report by clicking report slightly higher numbers but this difference is not statistically significant. The difference to the telephone study persists: the average report in each lab treatment is higher than in the telephone study. This shows that the mode of communication does not systematically influence reporting behavior strongly and is not driving the widespread truth-telling in our telephone study. We also elicit beliefs about the behavior of other participants and find in all four treatments that participants believe others to lie more than they actually do. Older participants (correctly) believe that lying is less prevalent. In the lab, higher beliefs are correlated with higher own reports. We find no evidence that being a student has a significant impact on behavior, or that the perceived time pressure on the telephone or the limited experience of the survey participants with the abstract design of economics experiments played a role.

Our paper adds to the nascent literature studying lying outside the lab. Previous studies focused on particular groups: Bucciol and Piovesan (2011) study a sample of children and find that many of them lie, unless they are reminded to be honest; Cohn et al. (2013) study prisoners and find that they become less honest when reminded of their criminal identity; and Utikal and Fischbacher (2013) ask a small sample of nuns to report the roll of a dice and find significant downward lying. Studies looking at unethical behavior in less abstract environments include Azar et al. (2013) who find that the majority of customers in a restaurant do not return excessive change. Similarly, Bucciol et al. (2013) study free-riding in public transportation in Italy and find that 43% of passengers evade the fare. We add two features: we study a representative sample and we can investigate the underlying motives by conducting additional lab experiments using the same well-defined decision.

Taken together, our results strengthen the doubts that previous lab experiments have cast on the assumption of zero lying costs: we find evidence for even higher lying costs in the telephone study. This suggests that studying the theoretical implications of such costs (e.g., Kartik et al., 2007, 2014; Doerrenberg et al., 2013) is a promising research avenue. At the same time, it is very likely that altruism, efficiency concerns, etc. are also important factors in the decision to pay taxes or how to treat patients, for example. Future research would need to investigate the relative importance of different motives that hold people back from misreporting and the interactions between motives. Our results also do not mean that lab experiments are uninformative about non-laboratory settings. However, the difference in behavior between our telephone study and our previous lab experiments rather shows how malleable reporting behavior can be. This opens many new questions

[2] In other studies concerning how people report their private information (e.g., Gneezy, 2005; Charness and Dufwenberg, 2006), the experimenter knows or will later know the subject's true type (and the subject is aware of this) and can thus judge whether an individual was honest or not. In our experiment, only the participant knows his or her private information. Our setup is thus closer to situations in which information is truly private and only known by the individual, while Gneezy's and Charness & Dufwenberg's setup is more representative of situations in which the private information is known by more than one person, e.g., when filing a joint tax declaration. These papers are also interested in the interaction between sender and receiver, from which we abstract. (See, however, the recent paper by Deck et al. (2013) who do not find an additional effect of promises on cooperation in single-blind and double-blind conditions.)