



Pecunia non olet: on the self-selection into (dis)honest earning opportunities

Kai A. Konrad¹ · Tim Lohse^{1,2,3} · Sven A. Simon¹

Received: 24 December 2019 / Revised: 25 September 2020 / Accepted: 19 November 2020 /
Published online: 11 January 2021
© The Author(s) 2021

Abstract

We study self-selection into earning money in an honest or dishonest fashion based on individuals' attitudes toward truthful reporting. We propose a decision-theoretic framework where individuals' willingness to pay for honest earnings is determined by their (behavioral) lying costs. Our laboratory experiment identifies lying costs as the decisive factor causing self-selection into honest earning opportunities for individuals with high costs and into cheating opportunities for those prepared to misreport. Our experimental setup allows us to recover individual lying costs and their distribution in the population.

Keywords Lying behavior · Lying costs · Misreporting · Honest earnings · Self-selection · Laboratory experiment

JEL Classification C91 · D81 · D91 · K42

1 Introduction

“Pecunia non olet” (money doesn't stink), as the Roman emperor Vespasian said to alleviate concerns that a presumably impure source of money would lower its value. One such impure source is dishonestly earned money. While some people are incorruptible and will suffer from lying, others follow Vespasian's adage and will cheat for personal gain without hesitation. This trait allows them to extract additional rents, and may make professions that allow for misreporting opportunities

Supplementary Information The online version contains supplementary material available at (<https://doi.org/10.1007/s10683-020-09691-7>).

✉ Sven A. Simon
sven.simon@tax.mpg.de

¹ Max Planck Institute for Tax Law and Public Finance, Marstallplatz 1, 80539 Munich, Germany

² Berlin School of Economics and Law, Badensche Straße 50-51, 10825 Berlin, Germany

³ CESifo, Poschingerstraße 5, 81679 Munich, Germany

particularly attractive to them. Indeed, survey evidence indicates large differences in trust in different professions. Typically, those professions with cheating opportunities, such as insurance agents, advertising executives but also politicians are the least trusted ones. In contrast, professions with little potential for manipulation, such as firefighters or teachers, are among the most trustworthy professions globally.¹ These opposites in the trustworthiness of different professions involve a deeper problem of causality. On the one hand, it is conceivable that the mere existence of a cheating opportunity in a certain profession might induce dishonesty. On the other hand, it is also conceivable that individuals self-select deliberately *ex ante* into different professions with different degrees of cheating opportunities. Such a self-selection could occur according to an individual's attitude toward honesty because money might stink for some, but not for others.

These considerations lead to empirical questions. Are individuals willing to pay to earn in an honest fashion? How badly do they suffer as a side effect of dishonest earnings? How is this lying cost related to their choice between an honest earning opportunity and one that allows them to make more money, but only if they are prepared to cheat? Finally, is dishonest behavior the outcome of the misreporting opportunity *per se*, or is it rather the outcome of deliberate self-selection into such opportunities? In other words, do certain professions attract dishonest people, or is it the cheating opportunity that makes people dishonest? We provide causal evidence of a self-selection effect, and rule out that the mere likelihood of facing the cheating opportunity induces dishonest behavior. Moreover, we show that the individual's lying costs determine the self-selection. Hence, abstracting from self-selection may lead to an underestimation of actual cheating behavior.

We generate these insights from a theory-guided experiment in the laboratory in which individuals choose between two binary lotteries in two different choice conditions. One lottery, let us call it the 'Bad Lottery,' has a zero outcome with high probability and a win outcome with a low probability. The other lottery yields the zero outcome with a low probability and the win outcome with a high probability. Let us refer to it as the 'Good Lottery.' The Bad Lottery is costless. To obtain the Good Lottery instead, individuals have to pay a price. We elicit individuals' willingness to pay (WTP) for the Good Lottery which has higher odds of having an opportunity to earn honestly instead of earning money through a misreporting opportunity. This WTP can reflect how badly different individuals suffer from making dishonest earnings. Therefore, this WTP can be seen as a subtle way to measure the heterogeneity in lying costs. Their lottery choice is recorded for two different conditions. One condition is truthful by design, and the lottery outcome is identical to the individual's final payout. Here, an individual's attitude toward truth-telling is irrelevant. The other condition allows for misreporting. In particular, the individual may claim the win outcome even if the true lottery outcome is zero. Observing individuals' actual reporting behavior when the misreporting opportunity emerges allows us to study

¹ See, e.g., the Ipsos (2019) 'Global Trust in Professions' survey, completed online by 19,587 adults in 23 countries or the Nuremberg Institute for Market Decisions (2018) 'Trust in Professions survey' based on 23,519 interviews in 20 countries.

how the WTP and the reporting behavior are related to each other. These findings give a good indication of whether individuals self-select into honest or dishonest earnings based on their lying costs.

Our results indicate that some, but not all individuals have a smaller WTP for the Good Lottery in the condition with the misreporting opportunity. Some individuals are not willing to spend money to obtain the Good Lottery. They also tend to misreport whenever the opportunity arises. Other individuals expend considerable resources to obtain the Good Lottery. These individuals with a higher WTP are also less likely to misreport if the misreporting opportunity emerges. Importantly, the finding does not depend on whether they are assigned to the preferred Good Lottery or whether they are exogenously assigned to the Bad Lottery instead. Their consistently honest behavior clearly indicates that (dis)honesty is determined by self-selection and not by the probability of facing the cheating opportunity.

Our work is related to the literature on cheating in general (as recently surveyed by Abeler et al. 2019 and Gerlach et al. 2019), and to two strands of this literature in particular. One literature strand touches on the issue of the self-selection of dishonest subjects into, or out of, lying opportunities. Several field experiments analyze the relationship between the occupational choice to work in corrupt sectors and the individual propensity to cheat. Departing from high corruption in the public sector of India, Banerjee et al. (2015) show that aspirants for public sector jobs engage in more corruptive behavior compared to aspirants for the private sector. Similarly, Hanna and Wang (2017) find a positive correlation between dishonest behavior in the laboratory and fraudulent behavior in the field for government nurses. A further dimension is the role of cheating opportunities for self-selection into competitive settings as documented by Faravelli et al. (2015) and Gino et al. (2015). However, clean and theory-guided evidence on whether subjects deliberately self-select into cheating opportunities based on their lying costs is limited so far. In Shalvi et al. (2011) the individuals are offered a fixed payoff as an exit option over participating in the dice-rolling game. In their framework the take-up of the exit option is not much affected by the possibility to lie, and lying behavior seems to be almost unrelated to the value of the exit option. We advance this literature since our design allows us to show that it is self-selection rather than the mere existence of a cheating opportunity which drives dishonest behavior. Additionally, by relating the WTP for honest earnings to reporting behavior, we identify an individual's lying costs as the determinant of self-selection. A potential implication is that screening procedures for applicants might be particularly promising for fighting dishonest behavior in certain professions.

Second, we contribute to the literature on the role of monetary incentives for lying and the measurement of lying costs. One approach is to vary the incentive size between treatments/experiments. Fischbacher and Föllmi-Heusi (2013) do not find a significant increase in lying when monetary incentives are tripled. Kajackaite and Gneezy (2017) come to a similar conclusion and show that the lying rate is largely insensitive to the multiplication of stakes. In meta-studies, Abeler et al. (2019) and Gerlach et al. (2019) confirm this pattern and document only little variation in dishonest behavior with

respect to the reward size. Hilbig and Thielmann (2017) use an alternative approach and vary the incentive size within-subjects for repeated cheating opportunities. They find evidence of ‘brazen liars’ and ‘incorruptible subjects,’ but also of subjects whose reporting decision is dependent on the stake size. However, the repeated setting might have an effect on cheating in itself. Finally, Gibson et al. (2013) and Vranceanu and Dubart (2019) elicit subjects’ reservation price to engage in fraudulent reporting and deceptive communication. Both find considerable heterogeneity in reservation prices, which suggests that misreporting is sensitive to monetary incentives. In contrast to most of the literature on dishonest decision-making, this approach implies that cheating is not a binary decision but rather there is a price for being dishonest by design. We propose a more subtle procedure to measure lying costs at the individual level. Our setting allows us to pin down a subject’s lying costs in a two-step procedure, but avoids implying that there is a price for misreporting by design. Hence, we complement the findings of the previous literature, and resolve some of the conflicting results on the heterogeneity of lying costs.

We proceed as follows. First, we sketch the backbones of the decision-theoretic framework which is closely mapped in our experimental design. We then derive predictions and provide a benchmark for the interpretation of the behavior in our laboratory experiment. We describe and discuss our experimental results before we finally conclude.

2 The decision-theoretic framework

Consider an individual that maximizes its expected monetary payoff. The individual chooses between two binary lotteries (labeled ‘Good Lottery’ and ‘Bad Lottery’). Each lottery yields either a high payout x_H (denoted as ‘win outcome’) or a low payout x_L (denoted as ‘zero outcome’), with $x_H > x_L$. The Bad Lottery B yields the high outcome x_H with probability p_B and the zero outcome x_L with probability $1 - p_B$. The Good Lottery G yields x_H with probability p_G and x_L with probability $1 - p_G$. Lottery G is the better lottery because $p_G > p_B$.

2.1 The Honest Condition: choosing in the absence of a misreporting opportunity

An individual can choose the Bad Lottery B for a zero price. For a price $z \geq 0$ she can choose the Good Lottery G instead. We denote the two choice alternatives by $(B, 0)$ and (G, z) . Once this choice is made, the individual learns the lottery outcome and is rewarded with this true outcome as her payoff. The expected payoffs from $(B, 0)$ and (G, z) are $\pi_B = p_B x_H + (1 - p_B)x_L$ and $\pi_G = p_G x_H + (1 - p_G)x_L - z$. The indifference price

$$z^* = (p_G - p_B)(x_H - x_L) \quad (1)$$

gives the individual the same expected payoff from $(B, 0)$ and from (G, z^*) . For a price $z \leq z^*$ the purchase of the Good Lottery G leads to an equal or higher expected




Fig. 1 Predicted fraction of individuals purchasing the Good Lottery G as a function of its price z

payoff. Hence, the expected payoff-maximizing individual chooses (G, z) for $z \leq z^*$ and $(B, 0)$ for $z > z^*$. The fraction of individuals who choose (G, z) as a function of z is depicted as a function $\gamma_T(z)$ in Fig. 1: $\gamma_T(z) = 1$ for $z \leq (p_G - p_B)(x_H - x_L)$ and $\gamma_T(z) = 0$ for any higher price.

2.2 The Untruthful Reporting Condition: choosing with a misreporting opportunity

Let us change one assumption in the setup. As before, the individual chooses between $(B, 0)$ and (G, z) and then learns the true lottery outcome. But then the individual is asked to report the lottery outcome. This report does not need to be truthful, i.e., the individual can choose any report $\xi \in \{x_H, x_L\}$. She receives exactly this reported value ξ as her payoff.

Suppose that the individual has a (behavioral) cost of reporting untruthfully. Let this cost be equal to $\theta \geq 0$ and a draw from a random variable with cumulative distribution function $F(\theta)$ with support $[0, \infty)$.² This θ might be interpreted as the sum of internal lying costs, including (self-) image concerns and the costs of violating (self-imposed) social norms, as modeled explicitly e.g., in Gneezy et al. (2018) or Abeler et al. (2019). For brevity, we refer to θ as the individual’s lying costs.

Based on these primitives the optimal reporting behavior of an individual with lying costs θ is the following. A win outcome x_H is always reported honestly, $\xi(x_H) = x_H$. In contrast, whether the individual reports truthfully if the lottery outcome is x_L depends on θ : $\xi(x_L) = x_L$ if $\theta \geq (x_H - x_L)$ and $\xi(x_L) = x_H$ if $\theta < (x_H - x_L)$.

² We assume that an individual’s θ for reporting $\xi(x_L) = x_H$ is the same in both lotteries. In the experimental setting, we are able to address whether this is a valid assumption.

The individual misreports on x_L if and only if her lying costs θ fall short of the monetary benefits of misreporting. Taking this reporting behavior into account, the expected payoffs for the choices $(B, 0)$ and (G, z) are

$$\pi_B(\theta) = \begin{cases} x_H - (1 - p_B)\theta & \text{if } \theta < x_H - x_L \\ p_B x_H + (1 - p_B)x_L & \text{if } \theta \geq x_H - x_L \end{cases} \tag{2}$$

and

$$\pi_G(\theta) = \begin{cases} x_H - (1 - p_G)\theta - z & \text{if } \theta < x_H - x_L \\ p_G x_H + (1 - p_G)x_L - z & \text{if } \theta \geq x_H - x_L \end{cases} \tag{3}$$

This makes the lottery choice straightforward: an individual pays z and chooses (G, z) if $\pi_G(\theta) - \pi_B(\theta) \geq 0$, and goes for $(B, 0)$ if $\pi_G(\theta) - \pi_B(\theta) < 0$, where

$$\pi_G(\theta) - \pi_B(\theta) = \begin{cases} (p_G - p_B)\theta - z & \text{if } \theta < x_H - x_L \\ (p_G - p_B)(x_H - x_L) - z & \text{if } \theta \geq x_H - x_L \end{cases} \tag{4}$$

For the case of indifference, we assume that the individual chooses (G, z) . The fraction of individuals who purchase (G, z) as a function of z is depicted as $\gamma_M(z)$ in Fig. 1.

For a price $z > z^* = (p_G - p_B)(x_H - x_L)$ the individual will never purchase (G, z) . This behavior is independent of the individual's θ : for $\theta \geq x_H - x_L$ the individual reports truthfully, but we have $\pi_G(\theta) - \pi_B(\theta) = (p_G - p_B)(x_H - x_L) - z < 0$. For $\theta < x_H - x_L$ the individual reports $\xi = x_H$ for any true lottery outcome, and we have $\pi_G(\theta) - \pi_B(\theta) = (p_G - p_B)\theta - z < (p_G - p_B)(x_H - x_L) - z < 0$. Hence, $\gamma_M(z) = \gamma_T(z) = 0$ for all $z > z^*$.

For $z \leq z^*$, individuals with $\theta \geq x_H - x_L$ report truthfully and choose (G, z) , as purchasing (G, z) gives them a higher expected payoff than $(B, 0)$. In contrast, individuals with $\theta < x_H - x_L$ report $\xi = x_H$ for any true lottery outcome. For these individuals, $\pi_G(\theta) - \pi_B(\theta)$ is weakly increasing in θ : the higher the individual's lying costs, the worse the zero outcome is compared to the win outcome, and the higher the individual's WTP for the Good Lottery. An individual is just indifferent between $(B, 0)$ and (G, z) for $(p_G - p_B)\theta - z = 0$, or $\theta = z/(p_G - p_B)$.³ Individuals with a lying cost higher than (or equal to) $z/(p_G - p_B)$ purchase (G, z) , while individuals with smaller lying costs choose $(B, 0)$. Accordingly, for a given z the fraction of individuals who choose (G, z) is equal to the probability that $\theta \geq z/(p_G - p_B)$, or,

$$\gamma_M(z) = 1 - F\left(\frac{z}{p_G - p_B}\right) \text{ for } z \leq z^* \tag{5}$$

³ This relationship allows for an estimation of an individual's lying costs based on her articulated WTP. For a detailed discussion of the estimation procedure and results from the experiment, refer to Sect. 4.3.

3 Experimental analysis

3.1 Design

3.1.1 Overview

We conducted our laboratory experiment in September and October 2018 at the econlab Munich.⁴ Each subject made decisions in two conditions in a sequential order.⁵ In the ‘Honest Condition’ (HC), the individual chooses one of two lotteries and receives the true payout of this lottery. In the ‘Untruthful Reporting Condition’ (URC) the lottery choice is followed by a reporting stage with an opportunity to misreport and to potentially earn a dishonest income by reporting a win outcome although the lottery yielded the zero outcome. As both conditions were fully computerized, dishonest reporting is identifiable at the individual level in our data.⁶

3.1.2 Lotteries

In both conditions there are two lotteries, the Bad Lottery and the Good Lottery. The lotteries have the same binary outcomes. The zero outcome has a payout x_L of EUR 0 and the win outcome has a payout x_H of EUR 12. For the Bad Lottery, the probability p_B of the win outcome x_H is 1/6, and the probability $(1 - p_B)$ of the zero outcome x_L is 5/6. For the Good Lottery, the probability p_G of the win outcome x_H is 4/6 and the probability $(1 - p_G)$ of the zero outcome x_L is 2/6.⁷ Evidently, the Good Lottery is better than the Bad Lottery by first-order stochastic dominance.⁸

⁴ A total of 308 subjects participated in 18 sessions of the experiment (average length of a session 75–90 min). These were predominately local university students (average age 24; 55.5 percent female students). Subjects were recruited from the laboratory’s subject pool using ORSEE (Greiner 2015). The experiment was programmed and conducted with z-Tree (Fischbacher 2007). The average payoff from the experiment was EUR 17.90 (min. EUR 6.50, max. EUR 27.50) including a show-up fee of EUR 9 and earnings up to EUR 2.90 from the post-tests.

⁵ We randomly varied the order of conditions. Half the subjects started with the URC and the other half with the HC. The behavior of subjects in both groups is comparable, and we do not find any major order effects. Hence, we abstract from order effects in the following analysis except for controlling for the order in our regression analyses. Moreover, in line with standard experimental procedures, we use the condition on order first for the between-subjects comparison of the WTP for the Good Lottery. To exclude any income effects or hedging behavior, only one of the two conditions was chosen at random for payout at the end of the experiment.

⁶ Instructions stated that the data is analyzed anonymously, and the lottery outcome was unknown to other participants or the (implementing) laboratory staff. We address the potential effects of observability with a post-experimental control question. See the section ‘Robustness discussion’ below for details.

⁷ Probabilities of the outcomes were visualized as drawing numbers from ‘1’–‘6’ on a computerized wheel of fortune. Each number represented either the zero or the win outcome of the respective lottery. The wheel of fortune determined the true lottery outcome for the subject.

⁸ Important to our design, subjects are not given the possibility to opt out completely, but only to improve their chances of an honest win outcome. Uncertainty over the outcome and the potential temptation to lie are not removed completely. The WTP is therefore more likely to be due to the reduction of the likelihood of facing temptation, and not due to completely avoiding an uncertain outcome.

3.1.3 Elicitation method

In both conditions the subject chooses between the two lotteries. We use the strategy method (Selten 1967) to elicit an individual's reservation price for choosing the Good Lottery over the Bad Lottery. Subjects make a series of choices between the Good and the Bad Lottery for 18 different possible price levels of the Good Lottery. The price for the Bad Lottery was EUR 0 for each choice, while the price z of the Good Lottery declined in steps of EUR 0.50 from EUR 8.5 to EUR 0.⁹ Accordingly, a subject's maximum WTP z_i^{\max} is defined by the highest price for which the subject chooses the Good Lottery over the Bad Lottery.

3.1.4 Implemented choice

When making their choices, subjects know that the laboratory had pre-determined a seller price z_0 that would apply, and that would be announced by the laboratory once they made their binary buying decisions for all possible values of z . The seller price z_0 was set to EUR 4 in the Honest Condition and to EUR 1.5 in the Untruthful Reporting Condition.¹⁰ As subjects only receive the information on the price z_0 once they have made their choices, it was in their best interest to consider their choices between $(B, 0)$ and (G, z) as potentially payoff-relevant for all possible $(B, 0)$ -versus- (G, z) comparisons.¹¹

3.1.5 Lottery allocation

If a subject chooses the Bad Lottery for z_0 (i.e., $z_i^{\max} < z_0$), the cost-free Bad Lottery at a price of EUR 0 is implemented ('Bad Lottery-self-selected'). If a subject expressed a WTP that (weakly) exceeded z_0 (i.e., $z_i^{\max} \geq z_0$), then, with probability 2/3, the subject was assigned to the Good Lottery at a price of z_0 ('Good Lottery-self-selected'). With the remaining 1/3 probability the subject was assigned to the Bad Lottery and did not have to pay z_0 ('Bad Lottery-assigned'). This procedure provides us with information on how individuals make their reporting choices if they have a high WTP for the Good Lottery but have received the Bad Lottery.¹²

⁹ The multiple price list employed here is a variant of the Becker-DeGroot-Marschak mechanism (see Becker et al. 1964) to elicit a subject's WTP in an incentive-compatible manner.

¹⁰ The price EUR 1.50 in the URC was chosen to achieve a balanced number of subjects who face the misreporting opportunity in the three possible lottery allocations. We based our estimation on the results of a structurally related pilot with different parameterizations.

¹¹ We addressed the potential concern that subjects might be suspicious about whether the seller price z_0 was predetermined or not. In their cubicle, we placed a sealed envelope which they were only allowed to open at the end of the experiment. The envelope contained information on the pre-determined price z_0 .

¹² It also accounts for the potential effect of loss aversion, as buying the Good Lottery is costly but does not imply an honest success with certainty. The literature typically finds more cheating in the loss domain than in the gain domain (e.g., Grolleau et al. 2016 or Schindler and Pfattheicher 2017), with an increasing propensity to cheat the less likely the bad outcome is (Garbarino et al. 2019). Therefore, the lying rate in the group 'Good Lottery-self-selected' may have an upward bias (that runs against our prediction). If so, focusing on the lying rate in the group 'Bad Lottery-assigned' instead allows for a clean estimation of the selection effect.

3.1.6 The conditions

In the Honest Condition, the lottery outcome of the allocated lottery is directly payoff-relevant and ends the condition. In the Untruthful Reporting Condition, the procedure continues. The subject learns the respective lottery outcome of the allocated lottery and is asked to make a report on this outcome. In turn, this report determines the payoff in the URC. The subject may report untruthfully, i.e., she might choose to report $\xi(x_L) = x_H$. There are no audits or formal sanctions for misreporting. Our design allows both for a comparison of the willingness to pay between-subjects (decision in the condition which was first in order) as well as within-subjects (decision in both conditions). Moreover, we are able to identify dishonest subjects in the URC by comparing their report to their true lottery outcome which are both observable in the data.

3.1.7 Post-tests and payoffs

The experiment concluded with an incentivized elicitation of beliefs about the choice and the reporting behavior of other participants in the respective session.¹³ The Post-Test section also included the Cognitive Reflection Test (Frederick 2005), Murphy et al.'s (2011) test for social value orientation and a variant of the Ellsberg paradox (Ellsberg 1961) to test for ambiguity aversion. Finally, we conducted a brief socio-economic questionnaire. Before subjects were paid out, they received a detailed overview of their earnings from the experiment.¹⁴

3.2 Hypotheses

We develop two sets of hypotheses along our decision-theoretic framework. The first set of hypotheses addresses the impact of the misreporting opportunity on individuals' WTP for choosing (G, z) over $(B, 0)$.

Using the actual numbers from the experimental setup, risk neutral, payoff-maximizing subjects have a predicted indifference price z^* of EUR 6 in the Honest Condition.¹⁵ For all prices $z > z^*$, subjects are predicted to choose the Bad Lottery. For

¹³ In contrast to honest subjects, dishonest subjects correctly predict the gap in the WTP between both conditions. Beliefs on the reporting behavior are biased toward the own reporting choice, and only the belief for the own lottery allocation is by and large accurate. We abstract from a detailed discussion of beliefs but results are available from the authors upon request.

¹⁴ We ran additional robustness sessions to account for potential confounders from the interaction effects of payment schemes or risk preferences (see Appendix A.1 in the Supplementary Information for a detailed discussion).

¹⁵ Our decision-theoretic framework assumes risk neutral decision-makers which is covered by our laboratory measure of risk preferences. Around 49 percent of subjects are broadly risk-neutral, 23 percent are risk-averse, 21 percent are risk-seeking, and 7 percent do not allow for a risk classification (proportions refer to subjects with a consistent statement of their WTP, numbers in the complete sample are comparable). Empirically, we do not find any strong relationship of risk preferences and the WTP for the Good Lottery. This makes sense, as individuals' choice to pay for the Good Lottery is not equivalent to reducing risk. Such payments not only affect the probability, but they also change the final monetary outcomes from x_L and x_H into $x_L - z$ and $x_H - z$. It makes the win outcome more likely, but it also reduces the zero

all prices $z \leq z^*$, subjects are predicted to choose the Good Lottery. In the Untruthful Reporting Condition, subjects with high lying costs $\theta \geq 12(= x_H - x_L)$ face a decision problem that is exactly equivalent to the Honest Condition. For these subjects, the misreporting opportunity does not matter and the choice behavior should be the same in both conditions. But subjects with low or intermediate lying costs $\theta < 12$ are likely to misreport on the zero outcome. Ex ante, because the misreporting opportunity makes the zero outcome less bad for them, this leads to a smaller WTP z_i^{\max} for the Good Lottery in the URC. However, the higher the individual's lying costs, the worse the zero outcome is and the closer z_i^{\max} is to z^* .

Hence, the comparison of $\gamma_T(z)$ and $\gamma_M(z)$ in Fig. 1 suggests that, for a given $F(\theta)$, the distribution of the WTP in the Honest Condition first-order stochastically dominates the WTP in the Untruthful Reporting Condition. For prices $z \in (z^*, \infty)$, the comparison is trivial as $\gamma_T(z)$ and $\gamma_M(z)$ coincide and all individuals choose $(B, 0)$ in both conditions. For prices $z \in [0, z^*)$, the fraction of individuals who purchase G for a given z deviate from each other. All types of individuals choose (G, z) in the HC. Their maximum WTP for the Good Lottery G is z^* . In contrast, the fraction of subjects choosing (G, z) in the URC is $1 - F(\frac{z}{p_G - p_B})$ for a given $z \in (0, z^*)$. This share is monotonically declining in z . At $z = z^*$, only individuals with high lying costs $\theta \in [x_H - x_L, \infty)$ would choose (G, z) . Hence, $1 - F(\frac{z^*}{p_G - p_B})$ is strictly smaller than 1 at $z = z^*$.

We summarize these considerations in the following

Hypothesis 1 The willingness to pay for the Good Lottery in the Untruthful Reporting Condition is

- (1) between-subjects and, on average, smaller than in the Honest Condition;
- (2) within-subjects, smaller than or equal to the Honest Condition.

Our second set of hypotheses focuses on individuals with a lottery outcome x_L in the Untruthful Reporting Condition. The decision-theoretic framework suggests a close relationship between the articulated WTP for the Good Lottery and the subsequent reporting behavior in this condition. As reporting behavior is observable in our setting, we are able to investigate this relationship by analyzing the WTP conditional on reporting behavior, and vice versa.

In the Untruthful Reporting Condition, misreporting only pays off for sufficiently small lying costs $\theta < x_H - x_L$. Accordingly, individuals base their WTP for the Good Lottery on their anticipated reporting behavior. Honest individuals disregard the misreporting opportunity and should therefore not alter their WTP

Footnote 15 (continued)

and win outcome to a lower level. Hence, paying for the Good Lottery does not transform the Bad Lottery into a lottery that is less risky (in the sense of Rothschild and Stiglitz 1970). Rather, it is an activity similar to what Ehrlich and Becker (1972) called 'self-protection.' Dionne and Eeckhoudt (1985) show that the decision-maker with a higher Arrow-Pratt measure of risk aversion is not necessarily inclined to pay more for the better lottery, and McGuire et al.'s (1991) conclusion is negative on general predictions about how risk aversion affects the willingness to pay for a better lottery.

$z^* = (p_G - p_B)(x_H - x_L)$. Their WTP for the Good Lottery reflects the increased chances of an honest win outcome. In contrast, as the misreporting opportunity makes dishonest individuals better off, dishonest individuals should express a smaller maximum WTP $z_i^{\max} = (p_G - p_B)\theta < (p_G - p_B)(x_H - x_L) = z^*$. Their WTP for the Good Lottery reflects their expected reduction of lying costs. Hence, dishonest individuals have a smaller WTP than honest individuals in the URC. As a consequence, individuals with $z_i^{\max} < z^*$ are expected to misreport $\xi(x_L) = x_G$ and individuals with $z_i^{\max} \geq z^*$ are expected to report truthfully $\xi(x_L) = x_L$.¹⁶

In the Honest Condition an individual’s readiness to misreport is irrelevant by design. In this condition individuals with different attitudes toward misreporting should articulate the same maximum WTP: $z^* = (p_G - p_B)(x_H - x_L)$. A comparison of both conditions reveals that dishonest individuals (as identified by their reporting behavior in the URC) have a smaller WTP in the URC than in the HC. In contrast, the condition has no influence on the WTP of honest individuals.

Finally, consider the reporting behavior conditional on the lottery allocation. Let the seller price of the Good Lottery be $z_0 \in (0, z^*)$. Individuals with a WTP smaller than z_0 should self-select into $(B, 0)$. Since $z_0 < z^*$, these individuals should misreport if such an opportunity emerges. Individuals with a WTP (weakly) higher than z_0 self-select into (G, z_0) . This group comprises two types of individuals: individuals with $\theta \geq \frac{z^*}{p_G - p_B}$ who will not misreport and individuals with $\theta \in [\frac{z_0}{p_G - p_B}, \frac{z^*}{p_G - p_B})$ who will misreport $\xi(x_L) = x_H$. In the aggregate, the fraction of dishonest reports in the group of individuals who select into $(B, 0)$ should therefore be higher than in the group who end up with (G, z_0) .

These considerations are collected as

Hypothesis 2 Consider the group of individuals who have the zero outcome x_L in the Untruthful Reporting Condition. In the Untruthful Reporting Condition

- (1) individuals who misreported $\xi(x_L) = x_H$ articulate a smaller WTP than truthfully reporting individuals;
- (2) individuals who misreported $\xi(x_L) = x_H$ articulate a smaller WTP than they do in the Honest Condition;
- (3) a higher fraction of individuals who select into $(B, 0)$ misreport as compared to individuals who select into (G, z_0) .

¹⁶ For many reasons (including risk aversion, lack of attention, calculation errors), the observed reporting behavior may deviate from this sharp prediction. In particular, for individuals with θ close to $\theta = \frac{z^*}{p_G - p_B}$, the difference between the payoff from lying, $x_H - \theta$, and the payoff from reporting truthfully, x_L , is small. Some of the individuals with a maximum willingness to pay above z^* might misreport, and some of the individuals with a maximum WTP below z^* may report honestly in the URC. A ‘softer’ prediction that accounts for these effects is that the propensity to lie is a declining function of the maximum articulated WTP.

4 Results

The analysis of the experimental data evolves along the lines of our predictions. We start with a comparison of the WTP in the Untruthful Reporting Condition and the Honest Condition. We show that the WTP in the URC is significantly smaller. Then, we investigate the relationship of the WTP in the URC and the subsequent reporting behavior. We find evidence that the WTP is indicative of the reporting behavior and that subjects self-select based on their lying costs. For the most part of the analysis, we focus on subjects with a WTP smaller or equal to EUR 6 in both conditions and who make a consistent statement of the WTP.¹⁷

4.1 The willingness to pay for the Good Lottery

4.1.1 The demand for the Good Lottery

We start with a (pseudo) between-subjects comparison of the WTP of the Good Lottery in the two conditions, focusing on the condition that was first in order. Figure 2 displays the aggregate distribution of subjects willing to buy the Good Lottery over the full support of possible prices. It has a similar interpretation as the theoretical demand curves $\gamma_T(z)$ and $\gamma_M(z)$ in Fig. 1. The solid line illustrates the empirical demand curve $\gamma_T(z)$ for the Good Lottery in the Honest Condition, and the dashed line the empirical demand curve $\gamma_M(z)$ in the Untruthful Reporting Condition.

Comparing the WTP in both conditions, Fig. 2 suggests that the WTP in the HC first-order stochastically dominates the WTP in the URC. A stable gap between both conditions of around 20 percentage points emerges up to a WTP of EUR 6. Econometric tests for this range of prices confirm this observation ($N = 86$ in the URC and $N = 74$ in the HC). The average WTP of 2.45 EUR in the URC is significantly smaller than the average WTP of 3.88 EUR in the HC (Wilcoxon-Mann-Whitney test: $p < 0.01$). A Kolmogorov-Smirnov test also rejects the null hypothesis of equal distributions ($p < 0.01$). Hence, subjects seem to take the misreporting opportunity into account when making their lottery choice, and they reduce their WTP accordingly. However, some subjects value the higher likelihood of an honest win outcome even in the presence of the misreporting opportunity. Only 14 percent have a non-positive WTP in the URC, and the median WTP in the URC is significantly larger than zero (one-sample median test: $p < 0.01$).¹⁸ For prices higher than EUR 6, both empirical demand curves seem to converge, which may indicate that the decision is

¹⁷ A higher valuation for the Good Lottery may not be explained by lying costs and, hence, is beyond the scope of our decision-theoretic framework. The consistent statement of the WTP is defined as choosing either the Bad Lottery or the Good Lottery for all price levels, or as switching only once from the Good Lottery to the Bad Lottery at z_i^{\max} . In a robustness estimation, we use the complete sample without any restrictions. The detailed discussion and analysis is provided in Appendix A.2 in the Supplementary Information.

¹⁸ A non-positive WTP in the URC may indicate subjects with lying costs of zero ($\theta = 0$) who are exactly indifferent between both lotteries.




Fig. 2 Between-subjects comparison of the willingness to pay for the Good Lottery in the Untruthful Reporting Condition and the Honest Condition. The gray area between EUR 6 and EUR 6.5 represents the theoretical indifference price z^* as measurement of the WTP was implemented in intervals of EUR 0.50. Figure 2 displays the data without any restrictions to the level of the WTP in both conditions (URC: $N = 115$; HC: $N = 116$)

not related to lying costs in this range. Consistent with *Hypothesis 1(1)*, we summarize our findings in

Result 1 In a between-subjects comparison, the WTP in the Untruthful Reporting Condition is smaller than in the Honest Condition.

4.1.2 The within-subjects differences in the willingness to pay

Since all subjects state their WTP in the HC as well as in the URC, we are also able to make a within-subjects comparison of the valuation of the Good Lottery in both conditions. Figure 3 illustrates the within-subjects difference between the WTP in the Honest Condition and the Untruthful Reporting Condition, i.e., $z_{i,HC}^{max} - z_{i,URC}^{max}$. In addition to the restrictions mentioned above, we only consider subjects who allow for a meaningful computation of the difference. This leaves us with 121 subjects here.

In Fig. 3, we mainly observe two groups of subjects. There is a large group with a difference in their WTP of around zero (between EUR -1 and EUR 1, 61 percent), indicated by the gray interval. These are subjects with a WTP which is broadly similar in both conditions.¹⁹ Our decision-theoretic framework suggests that subjects in this group may have high lying costs ($\theta \geq 12$). In this case, subjects forgo

¹⁹ To account for small mistakes and fuzziness in choice behavior, we define the group that has a similar WTP in both conditions as having an absolute difference smaller or equal to EUR 1. A more narrow definition of this group (absolute difference smaller or equal to EUR 0.5) does not qualitatively change the results. Around 10 percent have a negative difference, 50 percent have a similar WTP in both conditions, and 40 percent have a positive difference.

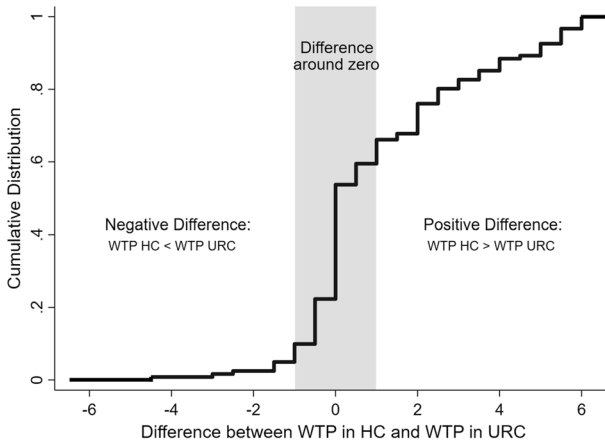


Fig. 3 Within-subjects differences in the WTP between the Honest Condition and the Untruthful Reporting Condition

the misreporting opportunity and both conditions are payoff-equivalent. Second, a group of 34 percent of subjects has a positive difference between EUR 1 and EUR 6, i.e., a higher WTP in the HC as compared to the URC. This group presumably consists of subjects with zero, small or medium lying costs $\theta \in [0, 12)$. These subjects should be willing to misreport for the zero outcome. Their WTP for the Good Lottery is determined by the desire to avoid the costs of lying. The smaller the lying costs, the smaller the WTP in the URC and the larger the difference between both conditions. The prevalence of these two groups of subjects with distinct differences in their WTP is in line with the prediction from our decision-theoretic framework. Finally, 5 percent of subjects have a negative difference between EUR -4.5 and EUR -1. Our framework does not accommodate an explanation for this third group, but its share is small in comparison to the other groups. The average within-subjects difference over all three groups is EUR 1.08 between both conditions, which is significantly larger than zero (Wilcoxon signed-rank test: $p < 0.01$). In accordance with *Hypothesis 1(2)*, we state:

Result 2 In a within-subjects comparison, the difference between the WTP in the Honest Condition and the Untruthful Reporting Condition is either zero or positive for most subjects.

4.2 Reporting behavior

The analysis so far has revealed that the misreporting opportunity in the URC decreases the demand for the Good Lottery. Now, we investigate the relationship of the WTP in the URC and the subsequent reporting behavior. In our decision-theoretic framework, both are determined by an individual's lying costs θ . Subjects may hence self-select based on their WTP, and report accordingly when facing the misreporting opportunity. To shed light on this, we focus on subjects with a zero outcome




Fig. 4 WTP in the Untruthful Reporting Condition and Honest Condition conditional on reporting behavior in the Untruthful Reporting Condition, *** $p < 0.01$

(for the respective lottery). These subjects had a monetary incentive to lie and to report the win outcome instead.²⁰

4.2.1 The willingness to pay conditional on reporting behavior

We start by examining the WTP conditional on the observed reporting behavior in the URC ($N = 71$). Misreporting subjects are classified as dishonest subjects, and those who made a truthful report as honest. Figure 4 displays the average WTP in both conditions for honest and dishonest subjects, respectively. The between-subjects comparisons of the WTP reveal that dishonest subjects in the URC are willing to pay less than a third of what honest subjects would pay for the Good Lottery. This difference is highly significant (Wilcoxon-Mann-Whitney test: $p < 0.01$). In contrast, there is no significant difference between honest and dishonest subjects in the HC (Wilcoxon-Mann-Whitney test: $p = 0.74$). The within-subjects comparisons across conditions corroborate these findings. Dishonest subjects reduce their WTP for the Good Lottery by more than 70 percent when they have the opportunity to misreport (Wilcoxon signed-rank test: $p < 0.01$). For honest subjects, we find no such difference in the WTP between the conditions (Wilcoxon signed-rank test: $p = 0.45$).

²⁰ For subjects with an honest win outcome ($N = 131$ in the complete sample), we cannot observe whether they had lied or not, and therefore they are excluded here. However, none of these subjects lied downward by stating a lower number than actually drawn.

Table 1 Multivariate analysis of the WTP in the Untruthful Reporting Condition

Variables	(1A) Honest	(1B) Dishonest	(2A) Honest	(2B) Dishonest
WTP Honest Condition	0.820*** (0.112)	0.162 (0.120)	0.649*** (0.127)	0.232 (0.147)
Female	- 0.227 (0.443)	- 0.652 (0.665)	0.331 (0.423)	- 0.830 (0.720)
Age	- 0.027 (0.073)	0.078 (0.091)	- 0.047 (0.067)	0.099 (0.094)
Socio-economic controls	Yes	Yes	Yes	Yes
Procedural controls	Yes	Yes	Yes	Yes
Post-tests	No	No	Yes	Yes
Constant	0.637 (1.993)	3.871 (2.900)	- 0.946 (1.774)	4.058 (2.915)
Observations	98	55	98	55

The table presents results of an interval regression. The dependent variable is the WTP in EUR in the Untruthful Reporting Condition. Specification “1A/2A” includes subjects that reported honestly a zero outcome and specification “1B/2B” includes subjects that dishonestly reported a win outcome and had a WTP smaller than or equal to EUR 6 in the URC. “WTP Honest Condition” is a subject’s WTP for the Good Lottery in the Honest Condition in EUR, “Female” is a dummy for female subjects, and “Age” is the age of the respective subject. “Socio-economic controls” include a dummy for business/economics students and a dummy for subjects with non-German mother tongue, “Procedural controls” include a dummy for consistent choice behavior and a control for order effects on both conditions, and “Post-tests” include the number of correct answers (0–3) in the Cognitive Reflection Test (Frederick 2005), Murphy et al.’s (2011) measure of distributional preferences (– 16.26 to 63.39), a dummy for ambiguity aversion (Ellsberg 1961) and an indicator for an individual’s risk preferences. Robust standard errors in parentheses.

*** $p < 0.01$

In summary, as our decision-theoretic framework suggests, only in the URC is there a difference between dishonest and honest subjects, and only dishonest subjects have a smaller WTP in the URC than in the HC. Hence, our results indicate that dishonest subjects deliberately reduce their WTP in the presence of the misreporting opportunity. In contrast, both conditions are payoff-equivalent for honest subjects as they tend to state a similar WTP in the URC and HC. In line with *Hypothesis 2(1) and 2(2)*, we summarize these findings in

Result 3 Dishonest subjects have a smaller WTP in the Untruthful Reporting Condition

- (1) as compared to honest subjects;
- (2) as compared to their WTP in the Honest Condition.

We confirm these findings in an interval regression analysis on the WTP for the Good Lottery in the Untruthful Reporting Condition (Table 1).²¹ The main result of

²¹ Here, we control for a consistent statement of the WTP in ‘Procedural controls’ directly. Moreover, as the misreporting opportunity is not relevant for honest subjects, we also include honest subjects with a WTP larger than EUR 6 in the URC (specification 1A/2A). Imposing the restriction of a WTP smaller or equal to EUR 6 for honest subjects does not qualitatively change the results in specification 1A/2A.

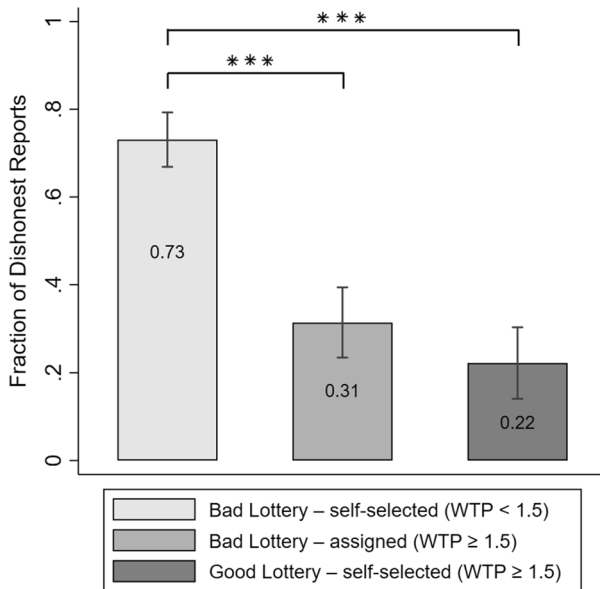


Fig. 5 Dishonest reporting conditional on the assigned lottery, *** $p < 0.01$

interest is the coefficient on the WTP in the Honest Condition. For the subgroup of honest subjects (specifications (1A)/(2A)), we find a positive and highly significant relationship between the WTP in both conditions. The point estimate implies that an increase of EUR 1 in the WTP in the HC is associated with an increase of EUR 0.65 to EUR 0.82 in the WTP in the URC. As both conditions are payoff-equivalent for honest subjects, this is in line with our theory framework. In contrast, for the subgroup of dishonest subjects (specifications (1B)/(2B)), the coefficient on the WTP in the HC is small and not significantly different from zero. Hence, the WTP in the URC seems to be based on a different motivation than in the HC, namely the subject's lying costs.

Result 4 The WTP in the Honest Condition does not predict the WTP in the Untruthful Reporting Condition for dishonest subjects, which indicates that subjects who misreported based their WTP on their lying costs instead.

4.2.2 Reporting behavior conditional on the willingness to pay

As a next step, we analyze the reporting behavior based on a subject's WTP in the Untruthful Reporting Condition. The sample is again restricted to the subgroup of subjects with the zero outcome ($N = 114$).²² Figure 5 displays the

²² As the WTP is not the focus of the analysis, we do not impose the restriction of making a consistent statement of the WTP here. Results are robust to this selection choice.

fraction of dishonest reports for the three possible lottery allocations. A comparison between the group ‘Bad Lottery–self-selected’ ($N = 52$) and the group ‘Bad Lottery–assigned’ ($N = 35$) allows us to investigate the consequences of self-selection for dishonest behavior. Subjects in both groups take part in the cost-free Bad Lottery, but have a different WTP. In the former group, the WTP is below the pre-determined price of EUR 1.5 for the Good Lottery, while subjects in the latter group were willing to acquire the Good Lottery but were randomly assigned to the Bad Lottery (prob. 1/3). Our decision-theoretic framework suggests that the first group consists predominantly of dishonest subjects and the second group of a mixture of honest and dishonest subjects. We find evidence in line with this prediction. Subjects with a small WTP are more than twice as likely to give a dishonest report as subjects with a high WTP. The difference of more than 40 percentage points is highly significant (χ^2 -test: $p < 0.01$).

We also compare the first two groups to the group of subjects whose choice of the Good Lottery was implemented (‘Good Lottery–self-selected’, $N = 27$, prob. 2/3). It is important to note that there are two core differences between the lotteries. The Good Lottery has a smaller probability of the zero outcome, and subjects pay EUR 1.50 for the Good Lottery independent of the outcome. Both aspects might affect the reporting behavior of subjects but are outside the scope of our model. The comparison of subjects who self-selected into the Bad Lottery with subjects who self-selected into the Good Lottery reveals that the latter are significantly less likely to use the misreporting opportunity (χ^2 -test: $p < 0.01$). In contrast, we do not find a significant difference in the reporting behavior between subjects with a high WTP that were involuntarily assigned to the Bad Lottery and those who were assigned to their preferred Good Lottery (χ^2 -test: $p = 0.42$).

These results provide evidence that the lying costs do not depend on the probability distribution of the respective lottery. This suggests that subjects do not find it easier to misreport in one of the two lotteries. Moreover, subjects with the zero outcome in the Good Lottery also seem to abstain from recovering their loss of EUR 1.5 by misreporting. Hence, self-selection into the Good Lottery reflects the preference for honest earnings. Subjects with a high WTP in the Untruthful Reporting Condition have a lower propensity to misreport. In accordance with *Hypothesis 2(3)*, we state

Result 5 The fraction of dishonest subjects is significantly lower among subjects willing to self-select into the Good Lottery than among subjects who self-select into the Bad Lottery.

Table 2 Multivariate analysis of misreporting behavior

Variables	(1)	(2)	(3)	(4)
WTP Untruthful Rep. Cond.	− 0.300*** (0.061)	− 0.312*** (0.062)	− 0.364*** (0.069)	− 0.329*** (0.070)
Female	− 0.704** (0.276)	− 0.684** (0.277)	− 0.606** (0.282)	− 0.541* (0.321)
Age	0.024 (0.039)	0.017 (0.039)	0.018 (0.042)	0.038 (0.033)
WTP Honest Condition			0.139** (0.058)	
Socio-economic controls	Yes	Yes	Yes	Yes
Procedural controls	No	Yes	Yes	Yes
Post-tests	No	No	No	Yes
Constant	0.430 (0.981)	− 0.0407 (1.085)	− 0.548 (1.199)	− 1.404 (1.056)
Observations	114	114	114	114

The table presents the results of probit specifications. The dependent variable is dishonest reporting (dummy variable). The dataset is restricted to subjects with a zero outcome and a WTP smaller than or equal to EUR 6 in the URC. “WTP Untruthful Rep. Cond.” is a subject’s WTP in the Untruthful Reporting Condition in EUR, “Female” is a dummy for female subjects, “Age” is the age of the respective subject, and “WTP Honest Condition” is a subject’s WTP in the Honest Condition in EUR. “Socio-economic controls” include a dummy for business/economics students and a dummy for subjects with non-German mother tongue, “Procedural controls” include a dummy for consistent choice behavior and a control for order effects on both conditions and “Post-tests” include the number of correct answers (0–3) in the Cognitive Reflection Test (Frederick 2005), Murphy et al. (2011) measure of distributional preferences (− 16.26 to 63.39), a dummy for ambiguity aversion (Ellsberg 1961) and an indicator for an individual’s risk preferences. Robust standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2 reports the results of a probit regression model with dishonest reporting as the dependent variable. Our model predicts a binary relationship between the WTP in the Untruthful Reporting Condition and reporting behavior. Here, we investigate a softer prediction, namely that a higher WTP in the URC is related to a smaller probability of misreporting. We find evidence for this prediction. The marginal effect of an increase of EUR 1 in an individual’s WTP in the URC corresponds to a 8–10 percentage-point decrease in the probability of misreporting. Both the sign and the magnitude are robust to the inclusion of diverse control variables in the alternative specifications (2)–(4). Female subjects are approx. 13–20 percentage points less likely to misreport, while we do not find a significant effect on the age of subjects. A higher WTP in the Honest Condition significantly increases the probability of giving a dishonest report (note that our theoretic framework does not make a prediction on the coefficient). Finally, subjects with a higher cognitive ability (Frederick 2005) are significantly more likely to misreport, while none of the other post-tests (social value orientation, ambiguity aversion and risk preferences) has a significant effect.




Fig. 6 Estimated distribution of lying costs based on observed reporting behavior and the willingness to pay in the Untruthful Reporting Condition

Result 6 In the Untruthful Reporting Condition subjects self-select based on their WTP. Subjects with a higher WTP are less likely to misreport.

4.3 Estimation of lying costs

The final part of our analysis focuses on an estimation of the individuals' lying costs based on the WTP in the Untruthful Reporting Condition. This is a subtle way of measuring an individual's lying costs. First, subjects state their willingness to pay to improve their chances of drawing the win outcome honestly. Only then do some subjects face the actual decision of whether to misreport or not. Importantly, the decision in both steps is determined by the individual's lying costs.

As lying is observable, the estimation procedure is based on the reporting behavior of subjects with the zero outcome.²³ Setting the left side of Eq. (4) equal to zero and solving for the lying costs θ , we recover a subject's lying costs as

$$\begin{aligned} \theta &= z/(p_G - p_B) \text{ if } \xi(x_L) = x_H \\ \theta &\geq x_H - x_L \text{ if } \xi(x_L) = x_L \end{aligned} \quad (6)$$

Figure 6 displays the estimated distribution of the lying costs ($N = 114$). For honest subjects, the costs of lying outweigh monetary benefits, and hence, lying

²³ Our procedure may also allow the estimation of lying costs if the reporting behavior is not observable. The key challenge here is to separate honest subjects with prohibitively high lying costs but a smaller than predicted WTP (for behavioral reasons, such as risk preferences or loss aversion), from subjects that are willing to lie but have strictly positive lying costs. Both groups have a WTP above zero but below EUR 6 in the Untruthful Reporting Condition. However, since only dishonest subjects have a smaller WTP in the URC as compared to the HC, a similar WTP in both conditions may indicate honest subjects with prohibitively high lying costs. Classifying the latter group of subjects as honest allows for a separation of both groups and an estimation as proposed in Eq. (6).

costs have to be equal to or higher than EUR 12. This is a lower bound of the true lying costs, and the proportion of these subjects is captured by the bar indicating lying costs of EUR 14. For dishonest subjects, the costs of lying are smaller than the monetary benefits, and lying costs are recovered as of the WTP in the Untruthful Reporting Condition. The result of this estimation suggests an almost binary distribution of the lying costs. Roughly 30 percent of subjects have zero lying costs, and around 50 percent of subjects have prohibitively high lying costs. A remaining 20 percent of subjects have intermediate lying costs in the interval EUR (0, 12). This is in line with survey evidence that subjects' reporting behavior is by and large insensitive to increased incentives in cheating games (e.g., Abeler et al. 2019 and Gerlach et al. 2019).

Result 7 The majority of subjects have either zero or prohibitively high lying costs, and only 20 percent of subjects fall in the intermediate range.

5 Robustness discussion

In this section, we discuss several features of our experimental design. A main point is the observability of lying at the individual level in our data. It allows for a one-to-one analysis of the relationship between reporting behavior and individual characteristics, and for a direct comparison of dishonest behavior between the three lottery allocations.²⁴ The computerized implementation of the URC also minimizes differences in the implementation of the HC. However, the observability may have consequences for subjects' reporting behavior. On the one hand, it may increase subjects' lying costs, for example, through more pronounced social image concerns. In turn, subjects may lie less as compared to unobservable settings, leading to a lower bound for dishonesty. On the other hand, subjects may anticipate the purpose of the experiment due to the observability of true outcomes, or even consider lying as an implicit rule. This raises concerns about an experimenter demand effect, which may imply a larger propensity to cheat. Moreover, subjects' behavior may reverse outside the laboratory, which could challenge the external validity of our findings. In the following, we discuss these concerns, and provide robustness analyses with our post-experimental control questions. We also address how lying under passive observability in the data is different from active audits or being caught with probability one.

The literature on dishonesty has used paradigms with different degrees of observability, ranging from lying in one's mind (Jiang 2013) over lying on an unobservable dice roll (Fischbacher and Föllmi-Heusi 2013) to computerized settings such as ours (e.g., Kocher et al. 2018). Differences in the lying behavior between these paradigms

²⁴ The Bad Lottery and the Good Lottery have different probabilities of the win outcome, and a simple comparison of the frequency of subjects reporting the win outcome would be misleading. While the correction for true win outcomes is straightforward in observable settings, unobservable settings have to correct for the *expected* number of stochastic true win outcomes (cf. Abeler et al. 2019).

indicate that changes to the observability of outcomes may have an impact on a subject's lying cost. For example, in Jiang's mind game (2013) subjects tend to cheat more as compared to other settings (Kajackaite and Gneezy 2017). Some contributions address the role of observability directly, for example, van de Ven and Villeva (2015) show that lying behavior is largely insensitive to privacy interventions, while Mol et al. (2020) conclude that only active as opposed to passive observation deters cheating behavior. An underlying question is whether dishonesty under observability is lying at all, and how it affects the structure of lying costs. Social norms may give insights to the first part of the question. Simon (2020) finds that misreporting on a computerized lottery outcome is regarded as dishonest behavior by the majority of subjects. Finally, there are typically lower levels of dishonesty as well as no partial lying in observable as compared to 'unobservable' settings, but lying behavior is still substantial (Gneezy et al. 2018; Abeler et al. 2019). Importantly, none of the mentioned paradigms completely rule out feelings of being observed. At least subjects themselves, and depending on their religious attitudes a metaphysical third party, observe their dishonesty. Therefore, considerations about who is able to observe one's lie are present in all settings, and the fear of being regarded as a liar is an essential component for the trade-off on whether to be dishonest or not. Observability becomes a matter of how much, rather than whether or not. The observability in the data may increase the weight of this component, but does not change the qualitative nature of the lying problem.²⁵

Complementary to this discussion, we assess the effect of observability empirically. One of our post-experimental control questions asked whether subjects felt observed at any point in time (please also refer to Appendix A.3 in the Supplementary Information). Overall, a large group of 78 percent of subjects stated not having felt observed. Only a small group of 22 percent of subjects were concerned about being observed.²⁶ This indicates that the computerized setting may have had only a limited impact on the reporting behavior. More importantly, the behavioral differences between the two groups are in line with our quantitative interpretation of higher lying costs under passive observability. Subjects who felt observed are marginally less likely to lie than their counterparts (35 percent vs. 53 percent, χ^2 -test:

²⁵ This becomes evident from a comparison to experiments on (tax) compliance. Even though the observability of the true outcome as in our setting is a shared component with compliance problems, the important difference is that lying is not uncovered by an audit. Consequently, it cannot be sanctioned in monetary terms with fines or in social terms with public shaming. To conclude, the observability may increase subjects' lying costs ($\theta_{\text{observable}} > \theta_{\text{unobservable}}$), but it is different from the cost of being caught (with certainty). Hence, our experiment reflects reporting situations where lying might be identifiable and the delinquent cannot be certain to keep the secret forever—a feature that arguably applies to many reporting situations.

²⁶ Note that this question not only refers to observability in the data, but also to concerns on observability for various reasons. As a precautionary check, we also compare the answers of subjects who had no incentive to misreport (since they drew the win outcome) to those who faced the trade-off between dishonestly reporting the win outcome or not (since they drew the zero outcome). In the former group, 82/77 percent of subjects did not feel observed, while the share is 75/73 percent in the latter group (main analysis/complete sample, respectively). The difference is not significant (χ^2 -test: $p = 0.32$ and $p = 0.40$), and subjects response to the exit question seems to be unbiased. Hence, we conclude that most subjects actually did not feel observed.

$p = 0.10$) and have larger estimated empirical lying costs (EUR 10.29 vs. EUR 7.16, Wilcoxon-Mann-Whitney test: $p = 0.02$). This fits well into our theoretic framework. The subjects who felt observed may be more concerned about their social image and have higher lying costs. In contrast, an experimenter demand effect would predict that these subjects are more likely to lie and would have smaller estimated empirical lying costs. In a second step, we repeat our analysis in section 4 for both subgroups separately. Due to the small sample of subjects who felt observed as well as due to the theoretical reasoning of larger lying costs, results are stronger in the group of subjects who did not feel observed but we confirm our results in both subgroups. Figure 8 in the Appendix A.3 in the Supplementary Information illustrates the WTP for the Good Lottery in the HC and the URC (compare to Fig. 2). The gap in the WTP the HC and URC is larger for subjects who did not feel observed, but subjects who felt observed also exhibit a difference in the WTP between both conditions. In summary, these results dampen concerns of an experimenter demand effect. Both from a theoretical and empirical perspective, it rather seems that the observability increased subjects' lying costs.

As a further robustness analysis, we compare the reporting behavior of experienced to unexperienced subjects. Unexperienced subjects might be less prone to consider cheating as part of the game. We calculate an experience score based on self-reported previous knowledge of our post-tests. Even though unexperienced subjects are less likely to lie (34 percent vs. 67 percent, χ^2 -test: $p < 0.01$) and have higher estimated lying costs (EUR 9.72 vs. EUR 5.73, Wilcoxon-Mann-Whitney test: $p < 0.01$), we confirm our findings in a subsample analysis for unexperienced subjects.²⁷

Finally, in light of the initial example of self-selection into different occupations, the question of the external validity of our results arises. There is convincing evidence of a positive correlation between cheating in the lab and dishonest behavior in the field (e.g., Hanna and Wang 2017 for skipping work, Potters and Stoop 2016 for the return of overpayment, Dai et al. 2018 for fare dodging in public transport, Cohn and Maréchal 2018 for school misconduct, and Cingl and Korbel 2020 for adolescent delinquents). This also applies to the related setting of tax compliance experiments (Alm et al. 2015). A corresponding field experiment is beyond the scope of this paper. However, the findings on the external validity of other forms of cheating behavior suggest that our results may allow for conclusions on the behavior outside the laboratory.

6 Conclusion

The paper studies differences between individuals who seemingly have no reservations of earning money by lying, and those who incur behavioral costs of misreporting and have a preference for honest money. Whether the former group of

²⁷ Detailed results are available from the authors upon request.

individuals are also more likely to deliberately select into misreporting opportunities is a crucial resulting question.

In our theory-guided experiment, we assess individuals' willingness to pay to earn money in an honest fashion. We experimentally investigate whether subjects self-select into honest and dishonest earning opportunities based on their lying costs, and whether this self-selection is predictive of their reporting behavior. Individuals choose between two binary lotteries, a Bad one and a Good one. The Bad Lottery has a zero outcome with high probability. In contrast, the Good Lottery yields a zero outcome with a low probability. We elicit an individual's willingness to pay to choose the Good Lottery over the Bad Lottery in two different conditions. One condition is truthful by design. The other condition allows for misreporting. Specifically, individuals may claim the win outcome even if their true lottery outcome is zero. The choice of the Bad Lottery confronts individuals with a high chance of the misreporting opportunity whereas the Good Lottery involves having a high chance of earning the win outcome honestly. Hence, the WTP for the Good Lottery can reflect how badly different individuals suffer from making dishonest earnings, and is a subtle way of measuring subjects' lying costs.

Our results indicate that the presence of the misreporting opportunity leads to a significant and sizable reduction in the WTP for the Good Lottery. The demand for the Good Lottery in the truthful condition first-order stochastically dominates its demand in the condition with a misreporting opportunity. However, this finding is not unique and is not the same for all subjects. Subjects seemingly anticipate their ability to withstand or surrender to the temptation of dishonest earnings and select the lottery accordingly. Some subjects are not willing to spend money to obtain the Good Lottery, and tend to misreport whenever the opportunity arises. Other subjects expend considerable resources on making their earnings honestly. In turn, these subjects are also less likely to misreport and to earn money by lying. We interpret these findings as evidence that subjects self-select into misreporting opportunities based on their individual lying costs. The experiment thus shows that subjects select themselves when choosing between profitable ways to earn money. People who hesitate to lie choose situations in which they can earn their money predominantly honestly, in which the opportunity of earning money by lying is rare, and they do not make use of these opportunities should they arise. People who find it easy to lie choose situations in which lying is a more frequent way of making money and, if the opportunity arises, they make use of it.

Our findings may have implications outside the lab. While some situations in real life leave subjects with little influence on the opportunity to misreport, for instance, because the opportunity arises by chance or by complete surprise, others allow for an anticipation of the cheating opportunity and to pre-plan how to behave optimally. An example might be the choice of profession. Our results suggest that dishonest subjects are more likely to self-select into environments with cheating opportunities, and consequences for truth-telling may be detrimental if such self-selection is possible.

Acknowledgements Open Access funding enabled and organized by Projekt DEAL. We thank Loukas Balafoutas, Jana Cahlíková, Nicolas Fugger, Vitali Gretschko, Werner Güth, Martin Kocher, Florian

Morath, Tobias Riehm, Raisa Sherif, Matthias Sutter, participants of the 2019 ESA European Meeting in Dijon, workshops in Innsbruck, at the ZEW in Mannheim and at the LUISS in Rome, and our colleagues at the MPI Munich for helpful comments. Additionally, we thank the editor John Duffy and two anonymous referees for valuable comments and suggestions. Moreover, we thank the econlab in Munich for providing the laboratory resources. At conferences and seminars, we presented a preliminary version of this paper entitled ‘Better to Win Honestly Than to Get Rich by Lying.’

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abeler, J., Nosenzo, D., & Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4), 1115–1153.
- Alm, J., Bloomquist, K. M., & McKee, M. (2015). On the external validity of laboratory tax compliance experiments. *Economic Inquiry*, 53(2), 1170–1186.
- Banerjee, R., Baul, T., & Rosenblat, T. (2015). On self selection of the corrupt into the public sector. *Economics Letters*, 127, 43–46.
- Becker, G., DeGroot, M., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9, 226–232.
- Cingl, L., & Korbøl, V. (2020). External validity of a laboratory measure of cheating: Evidence from Czech juvenile detention centers. *Economics Letters*, 191, 109094.
- Cohn, A., & Maréchal, M. A. (2018). Laboratory measure of cheating predicts school misconduct. *Economic Journal*, 128, 2743–2754.
- Dai, Z., Galeotti, F., & Villeval, M. C. (2018). Cheating in the lab predicts fraud in the field: An experiment in public transportation. *Management Science*, 64(3), 1081–1100.
- Dionne, G., & Eeckhoudt, L. (1985). Self-insurance, self-protection and increased risk aversion. *Economics Letters*, 17, 39–42.
- Ehrlich, I., & Becker, G. S. (1972). Market insurance, self-insurance and self-protection. *Journal of Political Economy*, 80(4), 623–648.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75(4), 643–669.
- Faravelli, M., Friesen, L., & Gangadharan, L. (2015). Selection, tournaments, and dishonesty. *Journal of Economic Behavior & Organization*, 110, 160–175.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise—An experimental study on cheating. *Journal of the European Economic Association*, 11(3), 525–547.
- Fosgaard, T. R., Hansen, L. G., & Piovesan, M. (2013). Separating will from grace: An experiment on conformity and awareness in cheating. *Journal of Economic Behavior & Organization*, 93, 279–284.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25–42.
- Garbarino, E., Slonim, R., & Villeval, M. C. (2019). Loss aversion and lying behavior. *Journal of Economic Behavior & Organization*, 158, 379–393.
- Gerlach, P., Teodorescu, K., & Hertwig, R. (2019). The truth about lies: A meta-analysis on dishonest behavior. *Psychological Bulletin*, 145(1), 1–44.

- Gibson, R., Tanner, C., & Wagner, A. F. (2013). Preferences for truthfulness: Heterogeneity among and within individuals. *American Economic Review*, *103*(1), 532–548.
- Gino, F., Krupka, E. L., & Weber, R. A. (2015). License to cheat: Voluntary regulation and ethical behavior. *Management Science*, *59*(10), 2187–2203.
- Gneezy, U., Kajackaite, A., & Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, *108*(2), 419–453.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, *1*(1), 114–125.
- Grolleau, G., Kocher, M. G., & Sutan, A. (2016). Cheating and loss aversion: Do people cheat more to avoid a loss? *Management Science*, *62*(12), 3428–3438.
- Hanna, R., & Wang, S. (2017). Dishonesty and selection into public service: Evidence from India. *American Economic Journal: Economic Policy*, *9*(3), 262–290.
- Hilbig, B. E., & Thielmann, I. (2017). Does everyone have a price? On the role of payoff magnitude for ethical decision making. *Cognition*, *163*, 15–25.
- Ipsos, (2019). Global trust in professions: Who do global citizens trust? Retrieved Sept 2nd, 2020 from: <https://www.ipsos.com/en/its-fact-scientists-are-most-trusted-people-world>.
- Jiang, T. (2013). Cheating in mind games: The subtlety of rules matters. *Journal of Economic Behavior & Organization*, *93*, 328–336.
- Kajackaite, A., & Gneezy, U. (2017). Incentives and cheating. *Games and Economic Behavior*, *102*, 433–444.
- Kocher, M. G., Schudy, S., & Spantig, L. (2018). I lie? We lie! Why? Experimental evidence on a dishonesty shift in groups. *Management Science*, *64*(9), 3995–4008.
- Lohse, T., Simon, S. A., & Konrad, K. A. (2018). Deception under time pressure: Conscious decision or a problem of awareness? *Journal of Economic Behavior & Organization*, *146*, 31–42.
- McGuire, M., Pratt, J., & Zeckhauser, R. (1991). Paying to improve your chances: Gambling or insurance? *Journal of Risk and Uncertainty*, *4*, 329–338.
- Mol, J. M., van der Heijden, E. C. M., & Potters, J. J. M. (2020). (Not) alone in the world: Cheating in the presence of a virtual observer. *Experimental Economics*, *23*, 961–978.
- Murphy, R. O., Ackermann, K. A., & Handgraaf, M. J. J. (2011). Measuring social value orientation. *Judgement and Decision Making*, *6*(8), 771–781.
- Nuremberg Institute for Market Decisions. (2018). Trust in Professions 2018. Retrieved Sept 2nd, 2020 from: https://www.nim.org/sites/default/files/medien/135/dokumente/2018_-_trust_in_profession_s_-_englisch.pdf.
- Potters, J., & Stoop, J. (2016). Do cheaters in the lab also cheat in the field? *European Economic Review*, *87*, 26–33.
- Rothschild, M., & Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, *2*(3), 225–243.
- Schindler, S., & Pfattheicher, S. (2017). The frame of the game: Loss-framing increases dishonest behavior. *Journal of Experimental Social Psychology*, *69*, 172–177.
- Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperimentes. In H. Sauermann (Ed.), *Beiträge zur experimentellen Wirtschaftsforschung* (pp. 136–168). Tübingen: J.C.B. Mohr (Paul Siebeck).
- Shalvi, S., Handgraaf, M. J. J., & De Dreu, C. K. W. (2011). Ethical manoeuvring: Why people avoid both major and minor lies. *British Journal of Management*, *22*, 16–27.
- Simon, S. A. (2020). Reporting under Ignorance – Is it a lie if I don't know?, Working Paper of the Max Planck Institute for Tax Law and Public Finance No. 2020-12.
- van de Ven, J., & Villeval, M. C. (2015). Dishonesty under scrutiny. *Journal of the Economic Science Association*, *1*, 86–99.
- Vranceanu, R., & Dubart, D. (2019). Deceitful communication in a sender-receiver experiment: Does everyone have a price? *Journal of Behavioral and Experimental Economics*, *79*, 43–52.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.