Big and Small Lies★

Diogo Geraldes[a], Franziska Heinicke[b] and Duk Gyoo Kim[b,*]

[a] Utrecht School of Economics, Utrecht University, the Netherlands
[b] Department of Economics, University of Mannheim, Germany

ARTICLE INFO

Keywords:
laboratory experiment
lying
luck
honesty

ABSTRACT

Lying involves many decisions yielding big or small benefits. Are big and small lies complementary or supplementary? In a laboratory experiment where the participants could simultaneously tell a big and a small lie, our study finds that lies are complementary. The participants who lie more in the big lie, also do so in the small lie and vice versa. Our study also finds that although replacing one dimension of the lying opportunities with a randomly determined prize does not affect the overall lying behavior, repeatedly being lucky on a high-stakes prize leads to less lying on the report of a low-stakes outcome.

1. Introduction

Lying behavior is pervasive in social, political, and economic life. Nevertheless, the lies that people tell are not all the same but rather differ considerably regarding the consequences they have. Some lies do not cause much harm relative to the counterfactual truth-telling situation. Other lies, however, have considerable consequences because they create a significant shift relative to the situation that would have occurred under truth-telling. For instance, deceiving an employer regarding oversleeping is likely innocuous to a corporation, whereas obtaining a job by misstating items in a resume will, in all likelihood, jeopardize it.

Big and small lying opportunities commonly occur together in the real world. A compelling example is the tax declaration because people have to report private information regarding several items, and the combination of all self-reports ultimately determines the tax liability. Thus, if people intend to adjust the total outcome in their favor, they can misreport all or just some of the individual items. However, misreports are not necessarily equal in size. That is, the consequence of a particular lie in the tax return depends on the misreported item. For example, misreporting the capital gains from assets held in a foreign country may lead to a considerable effect on the final tax payment, while overstating miscellaneous expenses will have a minor effect.

Although the recently growing literature on experimental economics has considerably advanced our understanding of the determinants of lying behavior, little is known about the interplay of lying behaviors. Our primary research questions consider the interaction between big and small lies in a *simultaneous* two-decision setting. When there are two lying opportunities of different sizes, how do people lie? Are big and small lies complementary? That is, do some people always lie for both big and small benefits? Or, do people supplement lies? If it is the latter, do people supplement one honest, less rewarding behavior with a big rewarding lie? Or do they lie more for petty outcomes because it is not a "big deal?"

To answer these questions, we conducted a laboratory experiment where a participant could tell big and small lies. Specifically, the participants of this experiment tossed a coin and rolled a dice, and they received a payoff based on the reports of both outcomes. The actual outcomes were unobservable by the experimenter. The reports on the coin toss represented the big lies because (a) we design higher-stakes payoffs associated with the coin than with the dice, and (b) the participants who lied about the coin accepted a larger deviation relative to the outcome under truthful reporting than the deviation possible on the dice.[1] Moreover, we compared the behavior in this main treatment—which we call the *Big and Small Lie* treatment (BSL)—with two control treatments. In the *Big Lie* treatment (BL), the participants reported only the coin while the dice was determined exogenously. In the *Small Lie* treatment (SL), the participants reported only the dice while the coin was determined exogenously. The two control treatments were designed to identify the effect of reporting, and possibly misreporting, two outcomes as compared to merely having an additional payment dimension.

Our results show that lying behavior is complementary when both big and small lying opportunities are available. That is, the participants who lie about the coin are also more likely to lie about the dice and vice versa. In absolute terms, we observe that more participants lied about the dice than about the coin. However, taking into account the size of the participants' lies, we observed that there was significantly more lying about the coin than about the dice.

★ We thank Collin Raymond and the participants of the HeiKaMaX Workshop, the ESA European Meeting 2019, and EEA Virtual Congress 2020 for invaluable comments and suggestions. Duk Gyoo Kim gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)–Project-ID 139943784–SFB 884.

* Corresponding author
diogogeraldes@gmail.com (D. Geraldes); f.heinicke@uni-mannheim.de (F. Heinicke); d.kim@uni-mannheim.de (D.G. Kim)

[1] We acknowledge that the outcome (and payoff) spaces are also different as the coin reports are binary while the dice reports vary from 1 to 6. As we discuss more in detail later, the largest possible lie in dice (reporting 6 when getting dice 1) is associated with a smaller payoff gain than the lie from a coin, so misreporting the coin outcome can be innocuously called a big lie.

The complementarity of lies does not lead to more lies when there are more lying opportunities. Neither the average coin reports in BL nor the average dice reports in SL differed from the corresponding reports in the main treatment. In other words, these findings indicate that lying is complementary when big and small lies can be reported jointly, but we do not observe more big (or small) lies when only one type of lying opportunity was available. In short, our results show that the participants behaved consistently across two lying opportunities.

Besides analyzing jointly reported big and small lies, we also tested the effect of observing a (un)favorable outcome on lying behavior. We found that the observation of an exogenous low-stakes outcome does not affect telling a big lie. Notably, however, we found that repeatedly observing a positive high-stakes exogenous outcome leads to decreased lying regarding the small lie.

This paper contributes to a growing body of research on lying behavior by exploring the interaction of lies of different sizes in a two-dimensional setting. In particular, it expands the understanding of lying behavior by analyzing the interaction of jointly told big and small lies while varying the size of a lie. Many economically relevant real-world settings present simultaneous lying opportunities. Thus, for deriving accurate policy implications, it is of utmost importance to examine the settings in the lab that resemble the vital features of the real-world settings.

The remainder of this paper is organized as follows. Section 2 provides an overview of the relevant literature. Section 3 outlines the experimental design, the hypotheses, and the procedures. Section 4 reports the results of the experiment, and Section 5 concludes.

2. Literature Review

Concerned with the adverse effects of self-serving dishonesty in economically relevant settings, the recently growing literature in experimental economics has been uncovering the determinants of lying behavior.[2] Traditionally, economists have considered lying to be inevitable as long as the material gains from a lie outweigh the risk and consequences of the lie being detected (Lewicki, 1983). However, the experimental literature on lying behavior has come to a different conclusion. Two recent meta-studies (Abeler et al., 2019; Gerlach et al., 2019), taking the impressively vast work conducted in the last few years altogether, have identified a considerable proportion of people that hold preferences for honesty. In particular, many people not only fear the material consequences of lying but also suffer moral costs from lying. Although this recent literature has largely enriched our understanding of the lying behavior, little is known regarding the interaction of two lies that vary by size. This paper aims to study the latter element by examining people's lying behavior when they are allowed to tell two lies with asymmetric consequences simultaneously.

Gneezy et al. (2018) identify the essential components of the size of a lie. The *stake size* component is the payoff that can be gained from lying. The *outcome* component is the deviation measure from the true state.[3] In our experiment, we manipulate the size of a lie considering both the outcome and the payoff components. Specifically, lying about a dice involved a small lie with several possible small deviations from the true outcome, while lying about a coin involved a big lie with one possible large deviation. In the following, we discuss the evidence regarding the effect of these two components.

The question of stake size has been discussed in the literature about lying from the beginning. In the experimental study by Mazar et al. (2008), the participants could get a higher payment by overreporting their performance of a real-effort task. To analyze the effect of the size of the payoffs on the lying, the authors vary the payoff incentives of the real-effort task. They find no significant differences in the lying even when the stake sizes quadrupled. Similarly, Fischbacher and Föllmi-Heusi (2013) conduct a low-stakes baseline treatment and a high-stakes treatment where the size of the stakes tripled. Belot and van de Ven (2019) conduct similar experiments to see whether dishonesty is persistent. Ruffle and Wilson (2018) examine how the participants respond to the different stakes concerning one visible characteristic, a tattoo. These studies find no significant effect of stake size on lying either.

The two aforementioned meta-studies confirm the null effect of the stake size. Considering 90 studies in which the participants could misreport a randomly generated outcome, Abeler et al. (2019) find that the stake size affects neither the average report nor the patterns of lying. For all stake sizes, outcomes at the lower end of the distribution are under-reported, and high outcomes are over-reported. Gerlach et al. (2019) compile data from 565 experiments, including those on lying about real-effort tasks and random outcomes and lying in sender-receiver games. They find that a higher maximum gain increases lying in the studies that used a coin toss to generate a random outcome, but this effect is not present when the analysis is limited to those studies that compare different stake sizes directly. For all other individual lying tasks, the stake size does not affect the reports. We should highlight that although the stakes are varied in these studies, each treatment corresponds to a single stake level,

[2] Abeler et al. (2014) and Arbel et al. (2014) investigate what individual characteristics that shape the costs of lying. Charness et al. (2019) show how the moral cost of lying may prevent lying in a loss frame. Hurkens and Kartik (2009), Houser et al. (2012), and Cojoc and Stoian (2014) study the relationship between social preferences and lying behaviors. Conrads et al. (2013) find that lying is more pronounced under team incentives. Cohn and Maréchal (2018), Dai et al. (2018), and Potters and Stoop (2016) examine whether and to what extent the laboratory measure of lying predicts misconduct in real situations. Similarly, Fosgaard (2020) analyzes how lying behavior differs between a student sample and a representative sample. Bucciol and Piovesan (2011) examine the lying behavior of children aged between 5 and 15. Djawadi and Fahr (2015) provide evidence of cheating outside of the laboratory.

[3] A third component is the *self-image concerns* of a lie, which relates to how blatant one's lie is. This component represents a challenge to maintain one's positive self-image. We do not manipulate the self-image concerns component in our experiment.

not to an interaction between two lying opportunities with different stakes.

The null empirical response to stake size could be due to two opposing directions of the effect. Stake size affects lying through two channels. First, a larger reward increases the marginal benefit of lying, which leads to more lying. The literature finds lying to be a trade-off between the monetary benefits from a lie and the psychological costs of lying, including the direct moral costs for breaking a moral norm and the reputational costs for possibly being considered as a liar by others (Gneezy et al., 2018; Abeler et al., 2019; Dufwenberg and Dufwenberg, 2018). Thus, increasing the benefit of a lie while keeping everything else equal will lead to more lying (Mazar et al., 2008; Hilbig and Thielmann, 2017). Second, the psychological or reputational cost of lying increases with the marginal benefit of lying, which leads to less lying. People are more likely to lie if they can justify it (Shalvi et al., 2015), and the lack of justification for lying increases the costs of lying for higher stakes (Mazar et al., 2008). Further, the reputational cost, incurred by the fear of being regarded as a liar, increases with the stake size of a lie (Kajackaite and Gneezy, 2017). These two opposing effects can explain the neutral effect of stake sizes on lying.

Moreover, evidence by Hilbig and Thielmann (2017) suggests that this effect is confounded by the fact that the stakes affect different types of liars differently. While people who are always lying (or always answering truthfully) are unaffected by the stakes, those who tell partial lies can be divided into two groups. The "corruptible" group reacts to high stakes by lying more, while the other group becomes more honest under the high stakes because this group is only willing to tell small lies.

The size of a lie due to the outcome component has received less attention in the literature so far. Gerlach et al. (2019) consider the dice roll and coin flip experiments separately in their meta-study. They argue that in the case of a binary task, people do not have options to lie partially. If they decide to lie, they have to do so to the full extent. Thus, in a task with more possible outcomes, people can tell smaller lies in the sense that they are closer to the true state. Gerlach et al. (2019) also find that coin and dice tasks do not differ concerning the average level of lying; however, more people lie on the dice tasks than on the coin tasks. The interpretation of this finding is that the people who tell small (partial) lies on the dice tasks will not lie on the coin tasks.

Closer studies to ours are Chowdhury et al. (forthcoming), Geraldes et al. (2019) and Barron (2019). Chowdhury et al. (forthcoming) also experimentally investigate the high- and low-stakes lies but in sequential order. Their focus and the relevant context are very different from our study. More specifically, they examine the effect of knowing (or not knowing) about a follow-up second-round lying opportunity when faced with the first-round lying opportunity. They show that people lie more on the first-round opportunity than on the second-round opportunity only when they are aware of the second round. In other words, lying in the first round increases if the participants could plan ahead of the second round. Using two dice, Geraldes et al. (2019) investigate individual lying behavior in a two-dimensional context to test whether multi-dimensional decision-making affects lying behavior. The participants' decision-making in their experiment is also simultaneous, but the stake and outcome components of two lying opportunities are kept constant. They find that the participants over-report significantly more on the lower outcome dice than on the high outcome dice. Barron (2019) also asks the participants of his experiment to report on two (one high-stakes and one low-stakes) dice rolls. He finds that compared to a uniform distribution, people over-report on the high-stakes dice but under-report on the low-stakes one.

3. Experimental design

3.1. Treatments

We design the laboratory experiments to observe how the participants behave when two misreporting opportunities differ in size. In the *Big and Small Lie* treatment (BSL), the participants were asked to toss a coin, roll a dice, and self-report the outcomes separately on the computer interface. This elicitation was repeated for ten rounds.

A participant's report, (C, D), where $C \in \{Head, Tail\}$ and $D \in \{1, 2, \dots, 6\}$, determines the points a participant earns in the following way:

$$\begin{cases} 15 + D & \text{if } C = Head \\ 7 + D & \text{if } C = Tail \end{cases}$$

At the end of the experiment, one round was selected randomly for actual payment to properly incentivize the participants (Azrieli et al., 2018). The conversion rate into euro was 1 point = 0.50 euro. Since the marginal benefit of lying about the outcome of the coin toss is larger than that of the dice roll, the coin toss is associated with a big lie. More specifically, although the small lie was scalable, the maximum benefit from the small lie was 5 (if reported 6 when the outcome of the dice was 1), which was strictly smaller than the marginal benefit of the big lie (reporting *Head* when the outcome of the coin was *Tail*). This treatment serves to capture the realm of real-life settings that involve big and small lies.

In the *Big Lie* treatment (BL), the participants were asked only to toss a coin and self-report its outcome and received a low-stakes prize based on the outcome of a dice roll, which was exogenously determined by the server computer. This elicitation was also repeated ten times. The *Small Lie* (SL) treatment was the reverse of BL. That is, the participants self-reported the outcome of a dice roll, and received a high-stakes prize based on the outcome of a coin toss, which was exogenously determined by the server computer. These two treatments serve to examine whether, and to what extent, an exogenous random event (relative to a self-reported random event) affects lying behavior on the other dimension. Table 1 summarizes the key differences in the three treatments.[4]

[4]The full experimental instructions for BSL are included in Appendix

Table 1
Experimental Treatments

	Coin ($\Delta_C = 8$)	Dice ($\Delta_D = 1$)
BSL	self-report	self-report
BL	self-report	exogenous
SL	exogenous	self-report

Δ_I is a benefit of marginally misreporting the outcome of $I \in \{Coin, Dice\}$.

3.2. Hypotheses

The primary purpose of our experiments is to enhance the understanding of individual lying behavior by observing the interactions between big and small lies. Thus, we do not posit "desired" conjectures before running the experiments. Although we learn many findings from the existing literature, we do not aim to verify or falsify extrapolated conjectures from previous studies. We do not aim to build an ex-post (or reverse-engineered) theoretical model either. Hence, we list "null hypotheses" as benchmarks for interpreting our observations.

Hypothesis 1 (Complete lying). *Every report, except for the exogenous outcomes, consists of Head and 6.*

If we assume self-interested rational individuals who maximize their monetary payoffs, then the theoretical predictions are trivial: Every participant reports *Head* and 6 for all of the ten rounds. If our data rejects this hypothesis, then our results support that the participants have the costs of lying (e.g., moral and/or psychological costs), as many previous studies have found (Abeler et al., 2019).

Hypothesis 2 (Complete honesty). *The distribution of the reports is uniform, that is, the probability of observing Head is 1/2, and the probability of any dice outcome is 1/6.*

If the costs of lying are sufficiently large, it is possible that the participants truthfully report the actual outcomes. Due to the law of large numbers, the empirical probability distribution of the repeated random draws with replacement will converge to the theoretical probability distribution. In addition, since we allow each participant to report their outcomes ten times, albeit the data would be noisy, we can also test this hypothesis within a participant. If our data rejects this hypothesis, then our results would indicate that the participants' costs of lying (assuming they exist) are not sufficient to overcome their temptation to make monetary gains.

Hypothesis 3 (Lying costs unlinked to the lying devices). *Lying in the coin reports is the same as lying in the dice reports.*

If the participants' costs of lying are positively (or negatively) associated with the size of a lie, we may observe more (or less) lying in the dice reports than in the coin reports. Alternatively, if the costs of lying are unrelated to the size of a lie, we can expect more lying about the coin as the monetary gain from lying about the coin is larger. At a different angle, if the costs of lying are positively associated with the number of ways lying, then we can expect more lying about the dice. Thus, to assess lying behavior, we need to measure both the number of misreports and the extent of their lying. The latter aspect differs between a dice and a coin. Unlike the binary coin outcomes, the dice outcomes allow us to observe more deviations from the theoretical probability distribution.

Hypothesis 4 (No complementary lies). *An individual's distribution of coin reports and that of dice reports are unrelated.*

This hypothesis relates to our main research question. If lies are complementary, then a participant who reports *Head* more frequently will report higher dice outcomes more often. On the other hand, if a participant finds the two lying opportunities supplementary, then a right-skewed distribution on one dimension will be associated with a left-skewed distribution on the other.

Hypothesis 5 (No spillovers). *Having two lying opportunities does not make an individual lie more or less compared to having only one opportunity.*

This hypothesis concerns the comparison between BSL and BL and SL, respectively. If the empirical distributions from BSL are more skewed to the left (or right) than those from BL and SL, then it suggests that more lying opportunities facilitate more (or less) lying in each dimension.

Hypothesis 6 (No time-varying justification). *The realizations of the exogenous outcomes do not affect the distribution of the outcome reports.*

Our final hypothesis addresses the effect of the exogenous outcomes in BL and SL. If we find a negative correlation between the exogenous outcomes and reports, our results would suggest that the participants use the exogenous outcomes to justify their lies. Lying after receiving a low exogenous outcome would be justified because of the bad luck, while lying after a high exogenous outcome would lack this justification.

3.3. Procedures

The experimental sessions were conducted in English at the Mannheim Laboratory for Experimental Economics (mLab) of the University of Mannheim. The participants were drawn from the mLab participant pool. Four sessions were conducted for each treatment, and a total of 152 participants participated in one of the 12 ($= 3 \times 4$) sessions. The number of participants per session varied from 8 to 17 due to no-shows, but the number of participants per treatment varied to a small extent (from 48 in SL to 55 in BSL).

B. We also conducted additional BSLBL (BSL for ten rounds and BL for another ten rounds) and BSLSL (BSL for ten rounds and SL for another ten rounds) to examine within-subject changes. The results in these additional treatments are similar to the comparisons of BSL with BL and SL, respectively. Thus, for the sake of non-repetition, we do not report this additional data, but it is available upon request.

Table 2
Overview of results in BSL

	Average reports Mean (SD)	N
Head	0.6727 (0.2130)	55
Head for dice≤ 4	0.6646 (0.2491)	50
Head for dice≥ 5	0.6842 (0.2803)	55
Dice	4.0800 (0.8381)	55
Dice for *Tail*	3.9054 (1.0569)	44
Dice for *Head*	4.0757 (0.9036)	55

Note: Under truth-telling, the expected share of *Head* is 0.5 and the expected average of Dice is 3.5.

Python and its application Pygame were used to establish a server-client platform. After the participants were randomly assigned to separate computer cubicles, the experimenter read the general descriptions of the experiment out loud. The participants were asked to carefully read the instructions displayed on the monitor and to qualify a comprehension quiz. Importantly, we did not track where the participants were seated and emphasized to them that their decisions in this regard would stay anonymous.

In all the treatments, the participants were subsequently asked to fill out a survey asking about their basic demographic characteristics, risk preferences, and degree of familiarity with the experiment. The participants' risk preferences were measured by the dynamically optimized sequential experimentation (DOSE) method (Chapman et al., 2018; Imai and Camerer, 2018). The average payment per participant was 8.31 euro. The payments were made in private, and the participants were asked not to share their payment information. Each session lasted less than 25 minutes.

4. Results

In each treatment, the participants completed ten rounds of the same task, which means that we have ten observations per participant. When necessary, we call each of the observations as a single report. Unless indicated otherwise, the following analysis considers the aggregated observation at the participant level because the ten single reports from a participant are not independent of each other. Thus, in what follows, the term "average reports" refers to the average of aggregated participant reports.

This section consists of four subsections. Sections 4.1 to 4.3, respectively, summarize the findings from BSL, BL, and SL. Section 4.4 compares the BSL findings with those of BL and SL.

4.1. Big and small lies

4.1.1. Over-reporting on coin and dice

Table 2 summarizes the BSL reports. The reports of this treatment show that the share of *Head* reports is significantly larger than 50%, the expected share under truth-telling (two-sided Wilcoxon signed-rank test (WT), p<0.001).

In addition to the overall level of lying, we analyze the distribution of reports over the ten rounds as shown in Figure 1. This figure provides a better understanding of how people lied about the coin outcomes, i.e., whether the participants lied only in some rounds or whether lying occurred in all the beneficial cases. Compared to a binomial distribution with ten draws and probability 0.5, the reports significantly shifted toward higher numbers of *Head* reports (two-sided Kolmogorov-Smirnov test (KS), p<0.001). This shift is particularly driven by some participants who reported ten rounds of *Head*.

Figure 1: Distribution of the number of *Head* reported in BSL

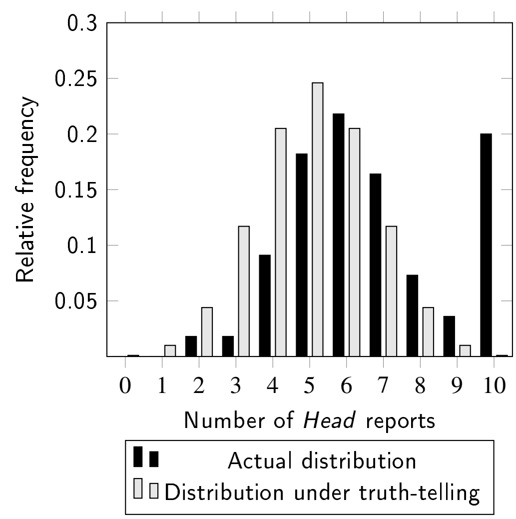

The distribution of the average reports on the dice roll in BSL significantly shifted from the value expected under truth-telling of 3.5 (p<0.001, WT). Further, we are interested in knowing how people lie on the dice, i.e., whether they only overstated the outcome of 6 or whether there was also over-reporting of the other outcomes. Figure 2 shows the distribution of single reports. The reports shifted toward higher outcomes, which resulted in a distribution significantly different from the uniform distribution expected from a fair dice (p<0.001, KS). More specifically, lying resulted in over-reporting of the outcomes 5 and 6 on the dice. Both of these outcomes were reported significantly more frequently than expected under truth-telling (Binomial test (BT), p=0.005 and p<0.001, respectively).

An analysis of the reports about the coin and the dice leads us to reject Hypotheses 1 and 2 and to conclude the following:

Result 1. *The participants lie significantly but not fully on the dice and the coin.*

4.1.2. Comparing lying about the coin and the dice

From Figure 2, we can conservatively estimate that the percentage of truthful reports regarding the dice roll is 57.6%.[5]

[5]Following Fischbacher and Föllmi-Heusi (2013), we assume that the participants do not report a lower number than observed, the percentage of reports of 1 is a conservative estimate of the true reports for each reported number (i.e., 9.6%×6). This is a conservative estimate since the subjects' truth-telling is likely to increase with the value of the observed outcome.

Figure 2: Distribution of single dice reports in BSL

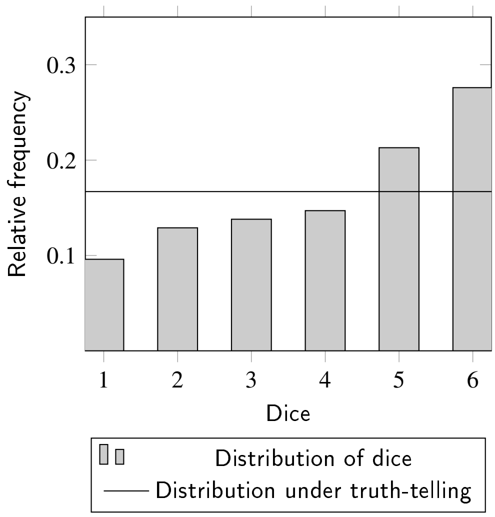

Figure 3: Estimates of lying and 95% confidence intervals based on Hugh-Jones (2019)

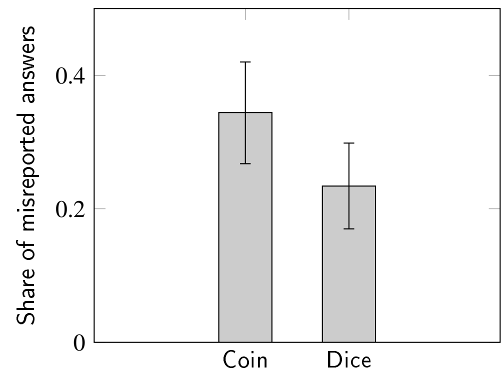

Regarding the coin, from Table 2, we can estimate that the percentage of truth-tellers is 65.4%.[6] These two estimates indicate that there were slightly more truth-tellers about the coin toss. However, that does not necessarily mean that the level of lying—which besides the *number* of lies, takes into account the *size* of the lies—is lower in the coin toss.

The report about the coin is binary, while the report about the dice can take six different outcomes. Importantly, the latter report also has a lower marginal contribution to the overall payoff. To compare the lying behavior between the two dimensions, we standardize the reports (Abeler et al., 2019). Specifically, the reports are standardized such that the lowest possible report takes the value -1, and the highest possible report takes the value 1. For a dice roll with linear payments, the reports $[1, 2, 3, 4, 5, 6]$ become $[-1, -0.6, -0.2, 0.2, 0.6, 1]$. For a coin toss with *Head* paying the higher amount, *Head* will be evaluated at 1 and *Tail* at -1. In order to compare lying about the coin and about the dice, we calculate the average standardized report for each device. The average standardized report is 0.346 (SD 0.426) for the coin and 0.232 (SD 0.335) for the dice. The standardized single reports about the coin are significantly higher than on the dice, which indicates a higher level of lying about the coin (p=0.005, WT).

As a robustness check, we further use the Bayesian method proposed by Hugh-Jones (2019)[7] to compare lying about the coin and the dice. Since this method is designed for a binary event, we consider a report of 5 or 6 on the dice as the high outcome because Figure 2 shows that lying led to over-reporting of these two outcomes.[8] Figure 3 shows that the method proposed by Hugh-Jones (2019) corroborates that there is more lying about the coin than about the dice in BSL.[9] In other words, we reject hypothesis 3.

Result 2. *The participants lie to a greater extent on the coin than on the dice.*

4.1.3. Lying conditional on the second dimension

To tackle our main research question, we analyze the relationship between the reports on the coin and reports on the dice. To understand whether reporting on one dimension is complementary or supplementary, we conduct two analyses. First, we test whether reports on the coin (or dice) differ between rounds with a high and rounds with a low dice (or coin) report. A significant result for this test would yield confirmatory evidence for a supplementary relationship. Second, we test the correlation between the dice and coin reports. While a negative correlation would point towards a supplementary relationship, a positive correlation would provide evidence for a complementary relationship.

Reports conditional on the other dimension are in the overview in Table 2. These values are generated by calculating the average coin (or dice) report per participant for all of the rounds in which a high dice (or coin) value was reported. The sample size is lower for *Head* if the dice report was greater than 4 and lower for the dice if the coin was reported as *Tail* because the number of rounds used per participant to calculate the respective average varied. For a participant who only reported *Head*, no average dice report for *Tail* could be calculated.

[6] Assuming that the participants who report *Tail* are not lying, that is, assuming that no participant reports *Tail* when *Head* is observed, the percentage of *Tail* reports is an adequate estimate of the true reports for each reported outcome (i.e., 32.7%×2).

[7] The measure uses the total number of reports, the number of reports that indicate the high outcome, and the probability of receiving the low payoff outcome under truth-telling to update an initial prior and to calculate a point estimate of the share of misreported answers as well as the corresponding confidence intervals.

[8] If we instead split dice reports in the middle to differentiate between the high and low reports, the point estimate for the coin is still larger than for the dice but the confidence intervals for the dice (0.190–0.353) have a larger overlap with the confidence interval of the coin. However, since we observe no over-reporting of an outcome of 4, we consider 5 to be the more appropriate cutoff value.

[9] Alternatively, we can consider the confidence intervals resulting from the method by Garbarino et al. (2018). This method provides narrower confidence intervals than the method of Hugh-Jones (2019), which is more conservative in the estimation of confidence intervals and is more reliable for small sample sizes. This alternative method estimates a 95% confidence interval of 0.287–0.397 for the coin and of 0.191–0.280 for the dice.

Figure 4: Distribution of dice single reports for each coin outcome in BSL

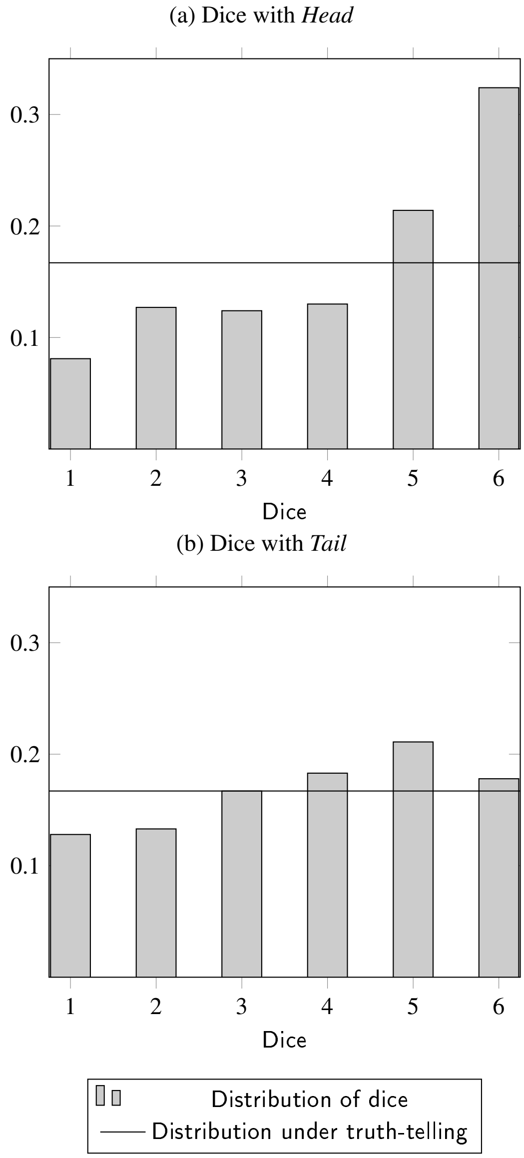

Since lying about the dice resulted in over-reporting of 5 and 6, we analyze whether *Head* was reported significantly more often when the dice report was 5 or 6. Table 2 indicates that only slightly higher shares of *Head* were reported for the higher reports about the dice. The difference is not significant (p=0.934, WT). Similarly, when *Head* was reported, the average report about the dice was slightly higher than when *Tail* was reported. This difference too is not significant (p=0.477, WT), and thus there is no indication of a supplementary relationship.

The average correlation between the coin and the dice reports takes a positive value of 0.643 and is significantly different from zero (Spearman correlation (SC), p<0.001). This indicates a complementary relationship between the two reports in the sense that the participants who lie for one outcome also tend to over-report the other. Figure 4, which shows the distributions of single dice reports for *Tail* and *Head*, corroborates the latter finding. Conditional on reporting *Head*, we see significant over-reporting of the outcomes 5 and 6 (p=0.018 and p<0.001, respectively, BT). Conditional on reporting *Tail*, we observe that reporting of 5 and 6 is more than 17%, respectively, but these inclinations are not significant (p=0.110 and p=0.689, respectively, BT). Moreover, the two distributions shown in Figure 4 are significantly different (p=0.009, KS). Notably, we cannot reject that the distribution of the dice reports conditional on reporting *Tail* is different from the expected distribution under truth-telling (p=0.304, KS).

Table 3
Linear regression analysis of dice reports in BSL

Dep. variable:	Average dice report (1)	(2)
shr_*Head*	2.891***	2.872***
	(0.367)	(0.387)
constant	2.135***	2.475***
	(0.259)	(0.407)
Controls	No	Yes
N	55	55
adj. R^2	0.531	0.536

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Finally, as a robustness check that allows us to control for the demographic variables, we conduct regression analysis. The dependent variable represents the average dice report across the ten rounds per participant, and the independent variable represents the share of *Head* reports across the ten rounds per participant.[10] Models (1) and (2) in Table 3 show that the participants who made many *Head* reports also made higher reports on the dice. The individual characteristics do not affect the reporting.[11] Overall, the regression analysis yields further evidence that lies on the two outcomes are complementary in the sense that the participants who tend to lie on one outcome also tend to lie on the other.

This analysis rejects Hypothesis 4. We summarize below our third result.

Result 3. *Big and small lies are complementary.*

4.2. Big lies under an exogenous low-stakes prize

Table 4 gives an overview of the results in BL, where the participants reported the outcome of a coin toss while the outcome of a roll of dice was exogenously and randomly determined. In this treatment, we also find significant over-reporting about the coin (p<0.001, WT). In Figure 5, we can see the distribution of the participants' coin reports in the ten rounds of BL. The distribution is significantly different from a binomial distribution with ten draws and probability

[10]Since the participants report two outcomes simultaneously, this analysis does not capture any causal relationships. It only captures the correlation between the two reports.

[11]This result also holds when omitting the share of *Head*, which might already contain the effect of individual characteristics, from the model.

Table 4
Overview of results in BL

	Average reports Mean (SD)	N
Head	0.6633 (0.1933)	49
Head for dice≤ 4	0.6762 (0.2140)	49
Head for dice≥ 5	0.6281 (0.3349)	48

Note: Under truth-telling, the expected average of *Head* is 0.5.

0.5 (p<0.001, KS). This difference is driven by an increase in the frequency of 7,8,9, and, in particular, of 10 rounds of *Head*.

Figure 5: Distribution of the number of *Head* reported in BL

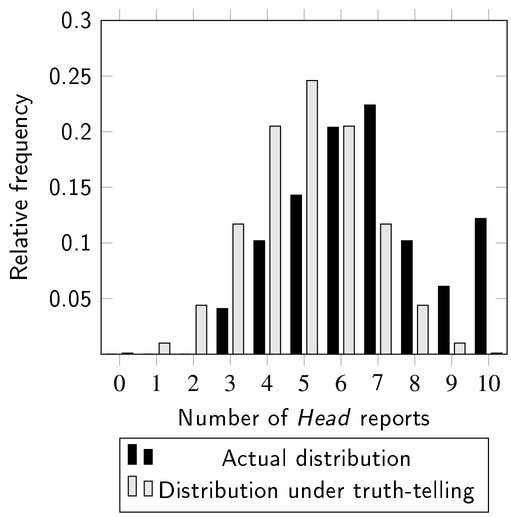

Regarding the effect of receiving an exogenously drawn low-stakes prize, we observe that there is less lying about the coin for the rounds in which the participants observed a high outcome on the dice, but the difference is not significant (p=0.474, WT).[12] The correlation between the observed dice and the reported coin is 0.0391 and not significant (p=0.790, SC) either. Hence, the evidence from BL indicates that reports on the coin were made independent of the exogenous dice.

Finally, we also conduct regression analysis for BL at the participant level. Specifically, we use a linear regression model in which the dependent variable is the average of *Head* that a participant reported over all the ten rounds, and the independent variable is the average dice report across the ten rounds per participant. Models (1) and (2) in Table 5 indicate that neither the average observed dice outcome nor the individual characteristics affect the number of rounds in which *Head* was reported. The low adjusted R^2 yields further support for the overall finding of this treatment that the exogenous small outcomes do not affect the reports of the big outcome.

[12]One participant only observed high values on the dice by chance and no average coin report could be calculated.

Table 5
Linear regression analysis of coin reports in BL

Dep. variable:	Share of *Head* (1)	(2)
avg_dice	0.0153	0.0075
	(0.052)	(0.055)
constant	0.609^{**}	0.751^{**}
	(0.186)	(0.240)
Controls	No	Yes
N	49	49
adj R^2	−0.019	−0.079

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6
Overview of results in SL

	Average reports Mean (SD)	N
Dice	4.1292 (0.9144)	48
Dice for *Tail*	4.0632 (1.0160)	47
Dice for *Head*	4.1731 (1.0084)	48

Note: Under truth-telling, the expected average of Dice is 3.5.

4.3. Small lies under an exogenous high-stakes prize

Table 6 summarizes the SL reports. Recall that the participants in this treatment only reported the outcome of the dice, while the outcome of the coin was exogenously and randomly determined. Reports on the dice significantly shifted from the expected value under truth-telling of 3.5 (p<0.001, WT). On the level of single reports, Figure 6 shows that reports shifted away from smaller outcomes toward the outcomes of 5 and 6, which are both reported significantly more frequently than expected (p=0.005 and p<0.001, respectively, BT).

Figure 6: Distribution of single dice reports in SL

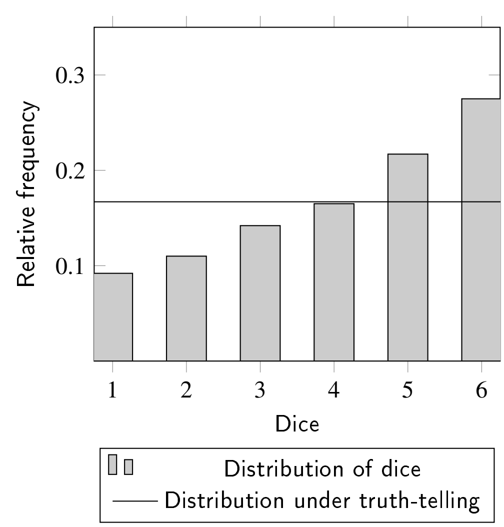

Regarding the effect of receiving an exogenously drawn

Figure 7: Distribution of single dice reports for each coin outcome in SL

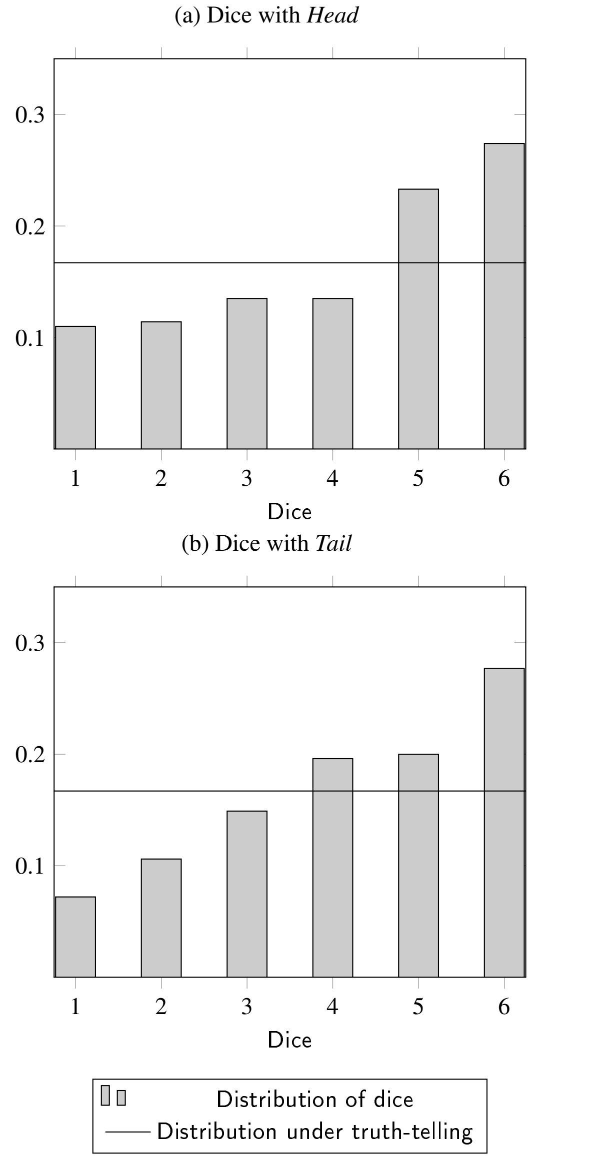

high-stakes prize, we observe that there is more over-reporting for the dice in the rounds in which the participants observed *Head* on the coin, but the difference is not significant (p=0.167, WT). A graphical analysis supports the latter finding. Figure 7 shows the distribution of single dice reports for each outcome of the coin. We observe that both distributions have a clear shift towards high outcomes and, most importantly, the two distributions are not significantly different (p=0.963, KS), which indicates that an exogenous outcome for the coin in a specific round does not affect reporting on the dice in the same round.

However, we find a significant negative correlation between the average dice reports and the average observed coin outcomes of −0.410 (p<0.01, SC). In Models (1) and (2) of Table 7, we report the results of a linear regression analysis at the participant level where the dependent variable is

Table 7
Linear regression analysis of dice reports in SL

Dep. variable:	Average dice report (1)	(2)
shr_*Head*	-2.481^{**}	-2.166^{*}
	(0.889)	(0.943)
constant	5.396^{***}	5.010^{***}
	(0.470)	(0.655)
Controls	No	Yes
N	48	48
adj. R^2	0.126	0.111

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the average dice report a participant made over ten rounds. These models' results confirm that the number of rounds of *Head* is negatively correlated with the average dice report. In other words, this analysis unveils that the participants who observed more rounds of *Head* (i.e., had more luck) report, on average, lower outcomes for the dice. Individual characteristics have no significant effect on this reporting.

4.4. Comparison of treatments

4.4.1. BSL vs. BL: Comparison of self-reports on the coin

First, we test whether receiving a low-stakes prize had a significant effect on reports on the coin outcomes, i.e., whether the overall level of reporting about the coin is different between BSL and BL. We find that the share of *Head* does not differ significantly between BSL and BL (p=0.929, WT). Thus, the overall level of lying about the coin is similar, regardless of whether the outcome of the dice is reported by the participant or exogenously determined.

Second, we assess how the coin reports depend on the dice outcomes, which are either self-reported or exogenously determined. While in BSL, *Head* is reported more frequently together with high reports on the dice, the pattern is reversed in BL, i.e., *Head* is reported less frequently when the participants observe a high dice outcome. The difference in the patterns for coin reports conditional on reporting (BSL) or observing (BL) high outcomes on the dice is not significant (Mann-Whitney U test (MW), p=0.503). The same holds for coin reports, conditional either on reporting (BSL) or observing (BL) low outcomes on the dice (p=0.709, MW).

In Section 4.1, we have shown a positive correlation between the share of *Head* and the average dice report, whereas in Section 4.2 we have shown no significant effect of the observed average dice roll on the amount of *Head* reports. Figure 8 illustrates this difference between treatments. In BSL, we observe a positive correlation between the two reports. In BL, however, we observe evidence of no correlation between the observed dice and the reported coin.[13]

[13]Table A.1 in Appendix A provides a quantitative analysis of the difference in the correlations between BSL and BL.

Figure 8: Relation of average of *Head* and average dice in BSL and BL

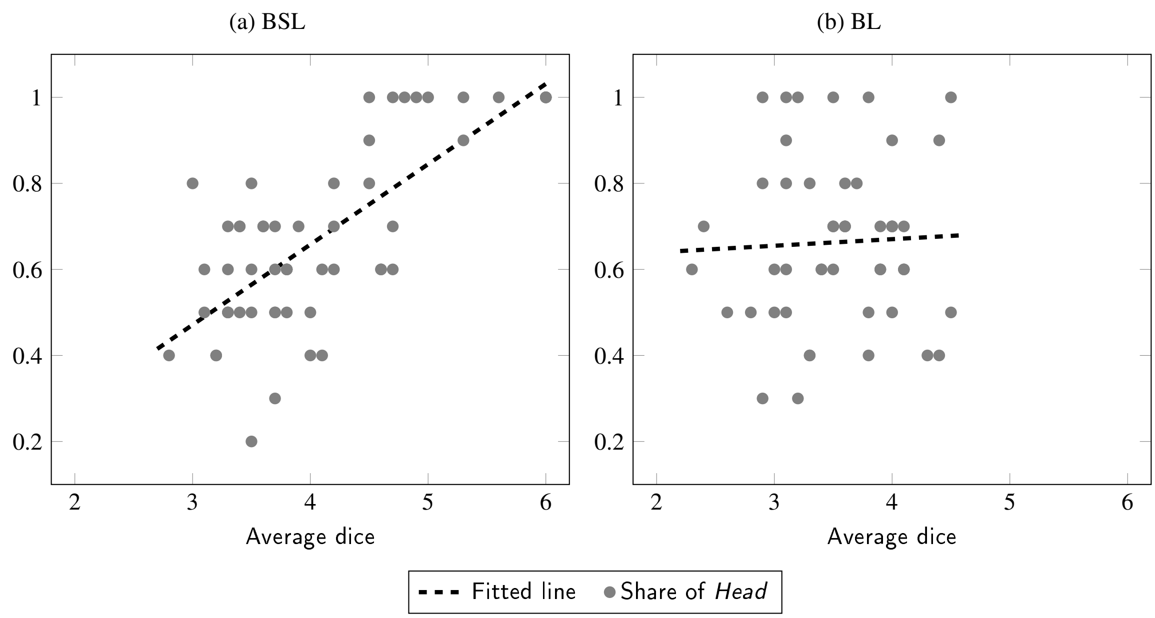

4.4.2. BSL vs. SL: Comparison of self-reports on the dice

First, we test whether receiving a high-stakes prize had a significant effect on the reports on the dice outcome, i.e., whether the overall level of reporting on the dice is different between BSL and SL. We find that lying about the dice is not significantly different between BSL and SL (p=0.999, KS; p=0.629, MW). Thus, whether there was a possibility to lie on the coin, or not, did not significantly affect the dice reports.

For the participants who self-report (BSL) or observe (SL) *Head*, the reports on the dice do not differ significantly (p=0.840, KS; p=0.481, MW). The average reports on the dice in *Tail* rounds do not differ significantly between the two treatments (p=0.259, KS; p=0.290, MW).

From Sections 4.1 and 4.3, we know that the effect of the total number of *Head* reported, or observed, over the ten rounds on the average dice report is different in the two treatments. In BSL, the participants who reported a higher number of *Head* made higher dice reports. In SL, the participants who observed a higher number of *Head* drawn by the computer made, on average, lower reports on the dice. Figure 9 illustrates the difference between the treatments. We see a positive correlation between the participant reports on the coin and those on the dice in BSL, whereas the correlation between the observed coin and dice reports is negative in SL.[14]

The comparison of the treatments leads us to not reject Hypothesis 5.

Result 4. *Having two lying opportunities does not make the participants lie more or less, compared to having one opportunity.*

Regarding Hypothesis 6, we cannot reject it for the reports on the coin. Notably, however, we reject Hypothesis 6 for the aggregate dice reports.

Result 5. *The realizations of the exogenous outcomes do not affect the distribution of reports on the coin. However, being repeatedly lucky on the coin—i.e., observing Head several times—significantly decreases the reports on the dice.*

The latter result essentially indicates that being unlucky on a high-stakes prize leads to justifying more lying in the report of a lower-stakes outcome.

5. Discussion and Conclusion

The purpose of the present research was to assess the relationship of jointly-reported high- and low-stakes lies. As far as we know, few studies in experimental economics investigate the interaction of two different lies reported simultaneously.[15]

To this end, we experimentally investigate how asymmetry in the size of a lie affects lying behavior within a two-dimensional context across ten rounds. More specifically, we elicit the participants' reports in a setting where they could simultaneously tell a big and a small lie. Further, we also investigate how the report of a big (or small) lie is affected when we replace the small (or big) opportunity by an exogenously determined small-stakes (or high-stakes) prize.

We show that lies are complementary, i.e., people who are more willing to tell a big lie are also more willing to tell a small lie. Further, we find that, when asked to report jointly,

[14] Table A.2 in Appendix A provides a quantitative analysis of the difference in the correlations between BSL and SL.

[15] The contemporaneous studies of Barron (2019) and Geraldes et al. (2019) also study two different simultaneous lying opportunities.

Figure 9: Relation of average dice and share of *Head* in BSL and SL

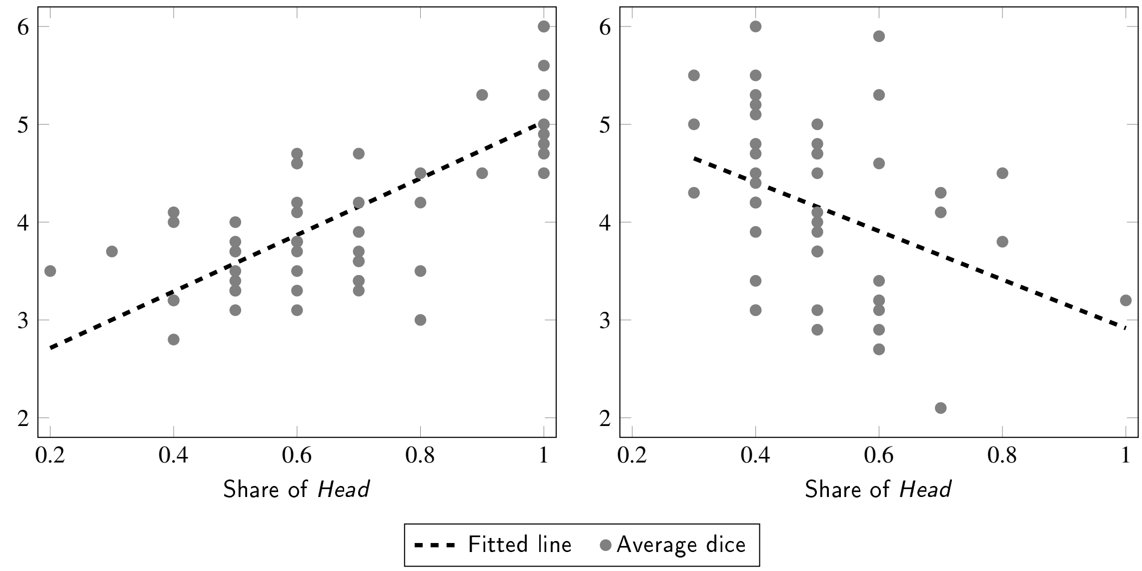

people lie more for the larger lying opportunity. We also find that offering more than one lying option neither facilitates nor suppresses lying behavior. In light of this evidence, we conclude that people behave consistently across asymmetric lying options.

Our experiment also allows us to test the effect of observing an exogenous payoff-relevant outcome on lying behavior. We find that observing an exogenous prize does not immediately affect the same-round lying behavior. Notably, however, when taking the dynamics across rounds into account, we find that being repetitively lucky in high-stakes outcomes decreases reports of small lies. In contrast, when the exogenous income involves small stakes, observing multiple lucky low-stakes outcomes in the previous rounds does not affect reports of big lies.

A possible explanation for the latter unanticipated result is that being lucky (or unlucky) on a high-stakes prize removes (or provides) the justification for telling small lies. To the best of our knowledge, the asymmetric effect of different stakes of luck on the lying behavior is a novel finding. The sparse literature on luck and lying has essentially compared lying about luck to lying about performance[16], whereas our finding provides a size effect of luck on the lying behavior.

The findings of this study have practical implications for tax compliance. Firstly, since lies are complementary, a tax authority should extend the check to all items when misreporting is detected in a given declared item. Secondly, in light of our findings regarding the effect of an exogenous outcome on lying, tax authorities should not be concerned about potential shifts in lying when strengthening the audit of a specific item. To put it differently, costly investigation of tax compliance on one item should not lead to lesser compliance on the other items.

In closing, a few words of caution are in order. This study offers the additional insight into distinguishing a big lie from a small one. However, even though it is transparent how we define a big and small lie, we are agnostic about the relative importance of the outcome and the stake components. In other words, our data cannot determine how these two components of a lie separately affect our results. Future research is needed to disentangle the relevance of these two channels. Notwithstanding this incompleteness and its exploratory nature, this study offers an important first insight into understanding decision-making in the presence of simultaneous asymmetric lying opportunities.

[16]The evidence is mixed. Kajackaite (2018) reports that lying about luck is intrinsically less costly than lying about performance, whereas Gravert (2013) shows that the participants who earned a performance-based payoff were three times more likely to take the undeserved maximum payoff than the participants with a randomly allocated payoff.

References

Abeler, J., Becker, A., Falk, A., 2014. Representative evidence on lying costs. Journal of Public Economics 113, 96–104. doi:10.1016/j.jpubeco.2014.01.005.

Abeler, J., Nosenzo, D., Raymond, C., 2019. Preferences for truth-telling. Econometrica 87, 1115–1153. doi:10.3982/ECTA14673.

Arbel, Y., Bar-El, R., Siniver, E., Tobol, Y., 2014. Roll a die and tell a lie—what affects honesty? Journal of Economic Behavior & Organization 107, 153–172. doi:10.1016/j.jebo.2014.08.009.

Azrieli, Y., Chambers, C., Healy, P.J., 2018. Incentives in experiments: A theoretical analysis. Journal of Political Economy 126, 1472–1503. doi:10.1086/698136.

Barron, K., 2019. Lying to appear honest. Discussion Paper SP II 2019–307. WZB. URL: http://hdl.handle.net/10419/204826.

Belot, M., van de Ven, J., 2019. Is dishonesty persistent? Journal of Behavioral and Experimental Economics 83, 101464. doi:10.1016/j.socec.2019.101464.

Bucciol, A., Piovesan, M., 2011. Luck or cheating? a field experiment on honesty with children. Journal of Economic Psychology 32, 73–78. doi:10.1016/j.joep.2010.12.001.

Chapman, J., Snowberg, E., Wang, S., Camerer, C., 2018. Loss Attitudes in the U.S. Population: Evidence from Dynamically Optimized Sequential Experimentation (DOSE). Working Paper 25072. National Bureau of Economic Research. doi:10.3386/w25072.

Charness, G., Blanco-Jimenez, C., Ezquerra, L., Rodriguez-Lara, I., 2019. Cheating, incentives, and money manipulation. Experimental Economics 22, 155–177. doi:10.1007/s10683-018-9584-1.

Chowdhury, S.M., Kim, C., Kim, S.H., forthcoming. Pre-planning and its effects on repeated dishonest behavior: An experiment. Bulletin of Economic Research doi:10.1111/boer.12238.

Cohn, A., Maréchal, M.A., 2018. Laboratory measure of cheating predicts school misconduct. The Economic Journal 128, 2743–2754. doi:10.1111/ecoj.12572.

Cojoc, D., Stoian, A., 2014. Dishonesty and charitable behavior. Experimental Economics 17, 717–732. doi:10.1007/s10683-014-9391-2.

Conrads, J., Irlenbusch, B., Rilke, R.M., Walkowitz, G., 2013. Lying and team incentives. Journal of Economic Psychology 34, 1–7. doi:10.1016/j.joep.2012.10.011.

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, 1081–1100. doi:10.1287/mnsc.2016.2616.

Djawadi, B.M., Fahr, R., 2015. "...and they are really lying": Clean evidence on the pervasiveness of cheating in professional contexts from a field experiment. Journal of Economic Psychology 48, 48–59. doi:10.1016/j.joep.2015.03.002.

Dufwenberg, M., Dufwenberg, M.A., 2018. Lies in disguise—a theoretical analysis of cheating. Journal of Economic Theory 175, 248–264. doi:10.1016/j.jet.2018.01.013.

Fischbacher, U., Föllmi-Heusi, F., 2013. Lies in disguise—an experimental study on cheating. Journal of the European Economic Association 11, 525–547. doi:10.1111/jeea.12014.

Fosgaard, T.R., 2020. Students cheat more: Comparing the dishonesty of a student sample and a representative sample in the laboratory. The Scandinavian Journal of Economics 122, 257–279. doi:10.1111/sjoe.12326.

Garbarino, E., Slonim, R., Villeval, M.C., 2018. A method to estimate mean lying rates and their full distribution. Journal of the Economic Science Association 4, 136–150. doi:10.1007/s40881-018-0055-4.

Geraldes, D., Heinicke, F., Rosenkranz, S., 2019. Lying in two dimensions and moral spillovers. Working Paper. SSRN. URL: https://ssrn.com/abstract=3381277.

Gerlach, P., Teodorescu, K., Hertwig, R., 2019. The truth about lies: A meta-analysis on dishonest behavior. Psychological bulletin 145, 1–44. doi:10.1037/bul0000174.

Gneezy, U., Kajackaite, A., Sobel, J., 2018. Lying aversion and the size of the lie. American Economic Review 108, 419–53. doi:10.1257/aer.20161553.

Gravert, C., 2013. How luck and performance affect stealing. Journal of Economic Behavior & Organization 93, 301–304. doi:10.1016/j.jebo.2013.03.026.

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. doi:10.1016/j.cognition.2017.02.011.

Houser, D., Vetter, S., Winter, J., 2012. Fairness and cheating. European Economic Review 56, 1645–1655. doi:10.1016/j.euroecorev.2012.08.001.

Hugh-Jones, D., 2019. True lies: Comment on garbarino, slonim and villeval (2018). Journal of the Economic Science Association 5, 255–268. doi:10.1007/s40881-019-00069-x.

Hurkens, S., Kartik, N., 2009. Would i lie to you? on social preferences and lying aversion. Experimental Economics 12, 180–192. doi:10.1007/s10683-008-9208-2.

Imai, T., Camerer, C., 2018. Estimating Time Preferences from Budget Set Choices Using Optimal Adaptive Design. Working Paper. LMU.

Kajackaite, A., 2018. Lying about luck versus lying about performance. Journal of Economic Behavior & Organization 153, 194–199. doi:10.1016/j.jebo.2018.07.010.

Kajackaite, A., Gneezy, U., 2017. Incentives and cheating. Games and Economic Behavior 102, 433–444. doi:10.1016/j.geb.2017.01.015.

Lewicki, R.J., 1983. Lying and deception: A behavioral model, in: Bazerman, M.H., Lewicki, R.J. (Eds.), Negotiation in Organizations, Sage Publications. pp. 68–90.

Mazar, N., Amir, O., Ariely, D., 2008. The dishonesty of honest people: A theory of self-concept maintenance. Journal of Marketing Research 45, 633–644. doi:10.1509/jmkr.45.6.633.

Potters, J., Stoop, J., 2016. Do cheaters in the lab also cheat in the field? European Economic Review 87, 26–33. doi:10.1016/j.euroecorev.2016.03.004.

Ruffle, B.J., Wilson, A.E., 2018. The truth about tattoos. Economics Letters 172, 143–147. doi:10.1016/j.econlet.2018.08.024.

Shalvi, S., Gino, F., Barkan, R., Ayal, S., 2015. Self-serving justifications: Doing wrong and feeling moral. Current Directions in Psychological Science 24, 125–130. doi:10.1177/0963721414553264.

Table A.1
Linear regression analysis of coin reports at the participant level in BSL and BL

Dep. variable:	Share of *Head*	
	(1)	(2)
avg_dice	0.140^{***}	0.187^{***}
	(0.025)	(0.028)
BL	0.068	0.698^{***}
	(0.038)	(0.200)
avg_dice*BL		-0.171^{**}
		(0.053)
constant	0.100	−0.089
	(0.104)	(0.115)
N	104	104
adj. R^2	0.227	0.292

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.2
Linear regression analysis of dice reports at the participant level in BSL and SL

Dep. variable:	Average dice report	
	(1)	(2)
shr_*Head*	1.4182^{**}	2.891^{***}
	(0.500)	(0.399)
SL	0.279	3.260^{***}
	(0.181)	(0.481)
shr_*Head**SL		-5.372^{***}
		(0.885)
constant	3.126^{***}	2.135^{***}
	(0.313)	(0.263)
N	103	103
adj. R^2	0.089	0.339

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A. Appendix - Additional results

Table A.1 shows the results of a linear regression analysis of the coin reports between BSL and BL. Note that in this regression, the dice variable is endogenous in BSL but exogenous in BL. The models show no causal relationships. The large significant coefficient for BL has to be interpreted jointly with the interaction term. For an average dice of 3.5, participant reports on the coin in BL were, on average, $0.698 + 3.5 * (-0.171) = 0.0985$ points higher than in BSL.

Table A.2 shows the results of a linear regression analalysis of the dice reports between BSL and SL. Note that in this regression, the coin variable is endogenous in BSL but exogenous in SL. The models show no causal relationships. The coefficient of SL has to be interpreted together with the interaction term. For five rounds of *Head*, the average dice in SL is $3.26 + 0.5 * (-5.372) = 0.574$ points higher than in BSL.

B. Appendix - Experimental Instructions

[*Note: Phrases in curly brackets are for the BSL treatment only. Corresponding phrases for the BL and SL treatments are followed in gray.]

Welcome to this experiment. Please read these instructions carefully.

Overview: The experiment consists of 10 rounds. In each round, {your task is to toss a coin, roll a dice, and report the outcomes.} Your cash payment will be based on your reports. The details follow.
{BL: your computer rolls a fair dice, and your task is to toss a coin and report the outcome.}
{SL: your computer tosses a fair coin, and your task is to roll a dice and report the outcome.}

Your task: {You can find one coin and one dice in front of you on the table. Please inspect them to verify that they are fair. In each round, toss a coin and roll a dice. Report the outcomes on the computer.}
{BL: You can find one coin in front of you on the table. Please inspect it to verify that it is fair. In each round, your computer rolls a fair dice, and you toss a coin. Report the outcome on the computer.}
{SL: You can find one dice in front of you on the table. Please inspect it to verify that it is fair. In each round, your computer tosses a fair coin, and you roll a dice. Report the outcome on the computer.}

You repeat this procedure for 10 rounds. In each round, your points will be determined as follows:

- 15+[outcome of the dice] if the coin lands Head
- 7+[outcome of the dice] if the coin lands Tail

For example, if you report (Head, 4), your points will be 19 (=15+4). If you report (Tail, 6), your points will be 13 (=7+6).

Payment: The server computer will randomly select one round, and your points in that round will be paid. This means that each round has an equal chance to be selected for the final cash payment. Thus, it is in your best interest to take each round equally seriously. Your points will be converted into Euros at the exchange rate of 2 points = 1 euro.

Anonymity: Your choices and answers will be linked with a computer number of your seat. We will never link your identity with your responses in any way. Your personal information provided for your payments will never be stored nor used for any research. In addition, since we do not track where you seat, we cannot match you with your reports, although we match the reports with the computer number.

Quiz: To ensure your understanding of the instructions, we will provide you with a quiz. If you have one or more wrong answers, you have to re-take the quiz. This quiz is only intended to check your understanding of the instructions. It will not affect your earnings.
Q1. If the coin lands head, and the outcome of the dice is 4, how many points do you receive?

Q2. If the coin lands tail, and the outcome of the dice is 1, how many points do you receive?
Q3. If the coin lands head, and the outcome of the dice is 6, how many points do you receive?
Q4. If the coin lands tail, and the outcome of the dice is 5, how many points do you receive?

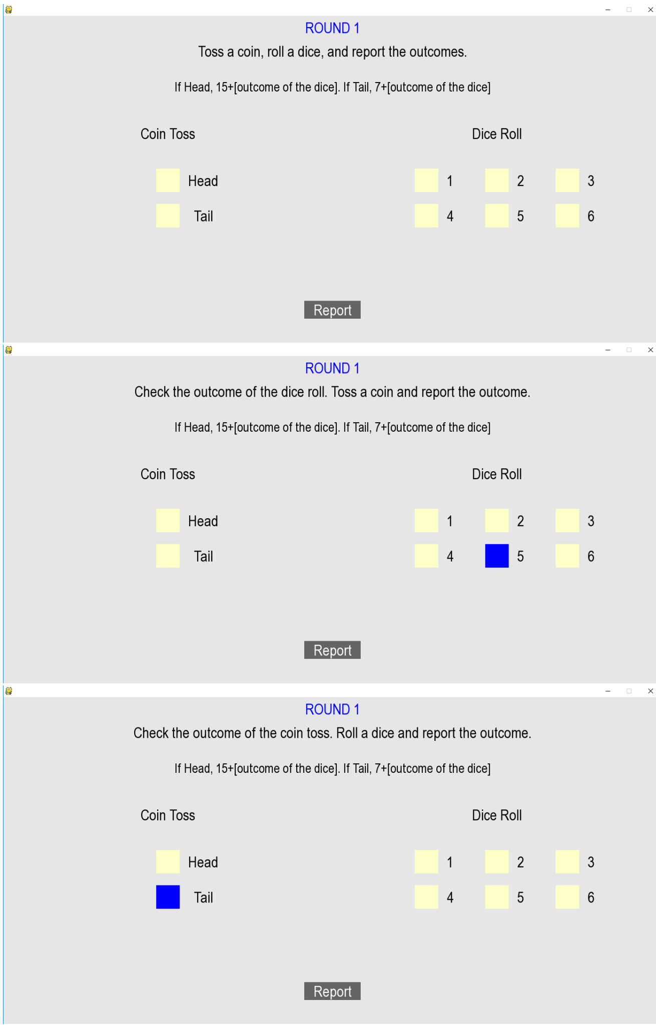

Figure B.1: Screenshots: BSL (top); BL (middle); SL (bottom)