Utrecht University

School of Economics

Lying in Two Dimensions

Diogo Geraldes, Franziska Heinicke &
Stephanie Rosenkranz

U.S.E. Working Paper Series nr: 21-01

Utrecht University **School of Economics**

U.S.E. Research Institute
Kriekenpitplein 21-22, 3584 EC Utrecht, The Netherlands Tel: +31 30 253 9800, e-mail: use.ri@uu.nl www.uu.nl/use/research

Utrecht University School of Economics

U.S.E. Research Institute
Working Paper Series 21-01
ISSN: 2666-8238

Lying in Two Dimensions

Diogo Geraldes*
Franziska Heinicke**
Stephanie Rosenkranz*
*Utrecht School of Economics Utrecht University
**Department of Economics, University of Mannheim

January 2021

Abstract
The expanding literature on lying has exclusively considered lying behavior within a one-dimensional context. While this has been an important first step, many real-world contexts involve the possibility of simultaneously lying in more than one dimension (e.g., reporting one's income and expenses in a tax declaration). In this paper, we experimentally investigate individual lying behavior in both one- and two-dimensional contexts to understand whether the multi-dimensionality of a decision affects lying behavior. In the one-dimensional treatment, participants are asked to roll two dice in one hand and to report the sum of both dice. In the two-dimensional treatment, participants are asked to roll two dice at the same time, but one in each hand, and to report the two dice separately. Our paper provides the first evidence regarding lying behavior in a multi-dimensional context. Using a two-dimensional die-roll task, we show that participants lie partially between dimensions, i.e., they distribute lying unevenly across dimensions, which results in greater over-reporting of the lower outcome die. These findings suggest a pertinent policy to tackle the infamous societal challenge of tax fraud: Tax report checks should focus on the item(s) for which a taxpayer profile hints at higher self-benefits in case of misreporting.

Keywords: lying, honesty, morals, multi-dimensional, lab experiment, lab-in-the-field experiment

JEL classification: C91; C93; D82; H26

Acknowledgements
The authors thank the participants of the EDS Seminar (Utrecht University), ESA World Meeting, TIBER Symposium (Tilburg University), 13th Nordic Conference on Behavioural and Experimental Economics, M-BEES, and the ZEW/University Mannheim Experimental Seminar for invaluable comments. The authors are particularly indebted for the comments of Jan Potters and Kai Barron. The authors gratefully acknowledges financial support from Utrecht School of Economics:

Comments welcomed to: d.geraldes @uu.nl

1 Introduction

In many real-world situations, people have the opportunity to lie for their own benefit. More often than not, these situations involve multi-dimensional decisions; that is, situations that involve the report of several outcomes in which the misreporting of each outcome yields self-benefits. In tax declarations, for example, people can under-report their income, under-report their wealth, and/or over-report their expenses in order to pay less in taxes. Likewise, students often have more than one assessment component in which they can cheat to improve the final grade of a course. In the field of science, for example, researchers have the option to manipulate the data collection and/or statistical analysis to obtain the desired result. While a myriad of studies has researched lying behavior in a one-dimensional context, it is unknown how people lie in a multi-dimensional context. In this paper, we study whether multi-dimensional decision-making affects lying behavior.

The seminal papers of Mazar et al. (2008) and Fischbacher and Föllmi-Heusi (2013) designed self-reporting tasks (the *Matrix-task* and *Die-roll task*, respectively), which became well-established paradigms to study lying behavior and spurred a rapidly growing literature. However, these two studies, as well as many follow-up studies on lying behavior, exclusively investigate one-dimensional contexts in which participants report a single outcome and can earn a higher payoff if they lie about the outcome. Even though these multiple studies provide valuable insights regarding dishonesty (for a review see, e.g., Abeler et al., 2019), the multi-dimensionality element inherent in many real-world situations in which lying is possible has not received attention in the literature.[1] Thus, a natural question to ask is whether a multi-dimensional context provides additional insight on lying behavior relative to a one-dimensional context. The objective of this paper is to tackle this question by investigating both one- and two-dimensional contexts.

To study the effect of dimensionality on lying behavior, we conduct a controlled laboratory experiment with two treatments: the *One-dimensional* (1D) treatment and the *Two-dimensional* (2D) treatment. More specifically, we implement a die-roll task with two dice, and, in both treatments, participants are paid according to the sum of two dice, which are rolled simultaneously. In the 1D treatment, we ask participants to report the sum of the two dice in a single report. In the 2D treatment, to induce two-dimensionality, we ask participants to report each die separately. Thus, the key distinction between the two treatments is the number of reports. In both treatments, participants repeat the die-roll task over ten rounds.

Moreover, in light of the growing demand for validation by combining different types of data in experimental work (e.g., Gneezy and Imas, 2017), we conduct a lab-in-the-field experiment to test the robustness of the results obtained in the laboratory. Specifically, we ask teenager students to perform the same die-roll task with two dice in their familiar school environment. Eliciting behavior from a different population and in a different setting makes the robustness check more stringent. These additional

[1] The recently created website http://www.preferencesfortruthtelling.com allows for visually identifying this limitation of the literature.

data are particularly important for checking the robustness of our laboratory experiment results because of the absence of previous findings regarding the novel setting that we explore. Thereby, we are also contributing to the stream of literature testing the generalizability of laboratory findings on lying behavior (Potters and Stoop, 2016).

To assess the impact of dimensionality of a setting on lying behavior, we fundamentally tackle two questions. Firstly, we analyse how people distribute lying between dimensions. Secondly, we assess whether the possibility of lying in more than one dimension affects the overall level of lying. By tackling these questions, we aim to understand: i. Whether we can extrapolate behavior on one dimension to behavior on a different dimension within the same setting; ii. Whether overall lying levels are comparable across settings with different dimensionality. That is, relative to a one-dimensional setting, does dimensionality increase, decrease, or not affect the overall level of lying?

Arguably, a multi-dimensional lying context might decrease the psychological costs involved with telling a lie because it opens the possibility of lying partially between dimensions. That is, if people combine truthful reports in one dimension with beneficial misreports in another dimension, they might believe that they are less likely to be seen as a liar and thus preserve a positive image of themselves (e.g., Mazar et al., 2008). Also, the truthful reports may serve as a justification of lies, which should further decrease the psychological cost of lying if one considers the concept of moral balancing (Merritt et al., 2010; Mullen and Monin, 2016). Therefore, when analyzing a single dimension in isolation, we are likely missing relevant traits of lying behavior. If so, to prevent drawing flawed conclusions, it is important to also investigate lying behavior in a multi-dimensional context.

Evidence suggests that there is a connection between the opportunity for moral balancing and lying behavior. Research in both economics and psychology has shown that people balance moral and immoral decisions (e.g., Dolan and Galizzi, 2015; Blanken et al., 2015; Mullen and Monin, 2016; Merritt et al., 2010). Examples of balancing behavior include declining to volunteer after having signed a petition (Guadagno et al., 2001), giving bad advice after disclosing a conflict of interest (Cain et al., 2005), or purchasing a luxury item after donating to charity (Khan and Dhar, 2006). Ploner and Regner (2013) find that people make larger donations after having cheated, and Cojoc and Stoian (2014) show that people lie more if they know that they can donate to charity later. Finally, Barron et al. (2019) combine a lying and a dictator game so that participants face two different moral motives; fairness and honesty. The authors find that people selectively adhere to the moral motive that provides the higher payoff. The theory underlying the balancing of moral and immoral actions is that such balancing allows an individual to maintain a positive self-image because the moral action counteracts the negative effect that an isolated immoral action would have on his or her self-image. Accordingly, in sequential decisions, we should expect immoral actions to be taken more frequently if they can be balanced with moral actions (Merritt et al., 2010).

While most of the literature discusses moral balancing regarding sequential decisions, there is also evidence that people exhibit moral balancing in simultaneous decisions. Engel and Szech (2017) find that participants balance between different ethical consumer labels in a willingness-to-pay experiment. Specifically, participants are willing to pay a premium for an organic product relative to a conventional product, but they are not willing to pay a significant premium for a product with both a fair-trade label and organic label compared to a product with the organic label only. Hence, since we implement two simultaneous lying decisions in the 2D treatment, we also analyze whether moral balancing emerges between dimensions in our simultaneous decisions.

This paper contributes to the literature on lying and honesty preferences; by expanding the context in which lying has been studied, we are able to detect systematic patterns of how people lie across dimensions. The fundamental result in this growing body of research is that people lie for their own benefit, but many do not take full advantage of the opportunity for lying (for a review of the literature and meta-studies see Rosenbaum et al., 2014; Jacobsen et al., 2017; Abeler et al., 2019; Gerlach et al., 2019). This finding is particularly remarkable because it has been found in experiments that offer participants the opportunity to lie to the full extent while preserving their anonymity and without imposing a negative externality on others. The actually observed pattern in these experiments basically consists of three types of people: i. Truth-tellers, who report truthfully regardless of the opportunity to earn a higher payoff if they lie; ii. Partial-liars, who do lie, but not to the full extent, i.e., they lie to increase their payoff but not to earn the maximum payoff possible; iii. Extreme-liars, who lie to the full extent possible and thus earn the maximum payoff possible (Fischbacher and Föllmi-Heusi, 2013). In our study, we also evaluate whether this pattern of lying types persists in a two-dimensional context.

Our paper provides the first evidence regarding lying behavior in a multi-dimensional context.[2] We find that people not only tell partial lies in each dimension but also distribute their lies unevenly across dimensions. More specifically, we show that participants lie partially between dimensions by favoring one dimension over the other for over-reporting. Notably, we detect that participants lie partially between dimensions based on the only observable difference between the two outcomes: the number rolled with each die. Namely, participants lie considerably more regarding the lower die. This result emerges not only in the laboratory experiment but also in the lab-in-the-field experiment, in which we test the finding in a different population.

Regarding the overall level of lying, we do not find a significant difference between the 1D and 2D treatments in both experiments. Thus, if the discussed decrease in the psychological cost of lying in a two-dimensional context is true, it does not necessarily lead to higher levels of over-reporting compared to a one-dimensional context. In other words, this study shows that with two simultaneous lying decisions, lying occurs to a greater extent in one decision than in the other, but this does not lead to more immoral behavior overall.

Our results have practical implications regarding policy design to cope with tax fraud. First, to detect fraud in tax reports, authorities should focus on the item(s)

[2] Shu et al. (2012) conduct an experiment in which people can lie in two different dimensions. However, the focus of their analysis is not on assessing how behavior is related between the two dimensions. And the behavior elicited on their experiment is not contrasted to a one-dimensional context either

for which the self-benefits of misreporting are higher, considering the taxpayer profile. Secondly, if fraud is detected in a specific item, authorities should not necessarily assume that the same pattern holds regarding other items.

The remainder of this paper is organized as follows. In Section 2, we describe the controlled laboratory experiment—which we denominate Experiment 1—, formulate hypotheses, and present the results of the laboratory experiment. In Section 3, we describe the lab-in-the-field experiment—which we denominate Experiment 2—and present its results. In Section 4, we provide the discussion and conclusion of our study.

2 Experiment 1: The effect of two-dimensionality

2.1 Experimental design

2.1.1 Treatments

Experiment 1 was designed with a two-fold purpose: i. Assess how participants distribute lying behavior between dimensions; ii. Compare the overall level of lying in a two-dimensional setting (2D treatment) with the level of lying in a one-dimensional setting (1D treatment). Regarding the first purpose, to carry out the assessment as clean as possible, we design the two dimensions to be identical. Moreover, to address the second purpose properly, we design the range of the the aggregate level of lying in the 2D treatment to be equivalent to the range of the level of lying in the 1D treatment.

More specifically, we used the well-established die-roll task (Fischbacher and Föllmi-Heusi, 2013), but with two dice. In the 1D treatment, participants were instructed to roll both dice in one hand and to report the sum of the two dice. In the 2D treatment, participants rolled the two dice simultaneously but in separate hands and reported the outcome of each die separately. For the sake of clarity, in each treatment, the instructions included a short video to illustrate how to roll the dice. A participant's payoff was determined according to the reported eyes, where 1 eye is worth €0.50. This means that payoffs could be between €1.00 and €6.00 in both treatments.

In both treatments, participants repeated the die-roll task over ten rounds, and one round was selected randomly for actual payment at the end of the experiment. The nine repetitions of the task were announced only after the first round. Therefore, the first round only allowed participants to engage in partial lying between dimensions, while the subsequent nine rounds also allowed participants to balance lying between rounds. The repetition of the task enabled us to test how robust partial lying between dimensions is if other balancing strategies are available.

2.1.2 Experimental procedure

The experiment was conducted at the Experimental Laboratory for Sociology and Economics in Utrecht (Netherlands). Participants were recruited using ORSEE (Greiner, 2015) among students of Utrecht University. In total, 139 participants took part in the experiment.[3] Participants were on average 23.6 years old and 74% were female. Demographic variables did not differ significantly between treatments. Upon arrival at the laboratory, participants received all instructions on-screen using oTree software (Chen et al., 2016) and were randomly assigned to the 1D treatment or the 2D treatment, resulting in group sizes of 68 and 71 participants, respectively. Including a participation fee of €2.00, the average payment per participant was €5.85. In total, the sessions lasted around 30 minutes.

2.2 Hypotheses

The focus of the literature on the determinants of preferences for honesty has been on explaining the heterogeneity in lying behavior observed in experimental studies. Thus, theoretical explanations on how and why people lie aim to accommodate the three distinct groups found in lying experiments: Truth-tellers, partial-liars, and extreme-liars.

In economic theory, the decision to lie has been modelled as a trade-off between the material gains from lying and the costs of telling a lie (caused by people's preference for honesty). Moreover, to accommodate the observed heterogeneity in lying behavior based on this trade-off, it further has to be assumed that the disutility from telling a lie consists of two types of costs (Abeler et al., 2019; Gneezy et al., 2018; Khalmetski and Sliwka, 2017). First, lying induces a direct cost due to the discomfort of violating a social or personal norm of honesty (Abeler et al., 2019). This direct cost can explain individual differences in the decision to lie because lying can be too costly for some people but affordable to others.

Besides direct lying costs, people can experience disutility from what Gneezy et al. (2018) coined the "likelihood dimension". The term refers to the notion that a lie is more costly the more likely it is to be identified as a lie. Being seen by others as a liar, or to think of oneself as a liar, negatively affects people's overall image of themselves, which includes their social-image if reports are observed by an audience, as well as their self-image as people also act as their own audience (Dufwenberg and Dufwenberg, 2018). Thus, the more obvious a lie is, the higher the costs involved in telling it (Mazar et al., 2008; Fischbacher and Föllmi-Heusi, 2013; Abeler et al., 2019; Dufwenberg and Dufwenberg, 2018). This indirect cost can explain partial-lying because partial-liars believe that partial lies appear more likely to be true statements, and so partial-liars are able to think of themselves in a more favorable light than when telling an extreme lie. Accordingly, this lying strategy allows partial-liars to maintain a positive image of themselves while still reaping some monetary benefits (Mazar et al., 2008; Gneezy et al., 2018).

In our 2D treatment, participants have the option to lie in two reports that jointly contribute to the final payoff. Thus, in contrast to the 1D treatment, partial lying in the 2D treatment can take two forms: i. Lie partially in each dimension; ii. Lie partially

[3] This sample size provides a power of 0.8 to detect an effect size comparable to moral licensing effects (Blanken et al., 2015) and a power > 0.8 to detect a medium effect size as suggested by Schindler and Pfattheicher (2017)

between dimensions. That is, a participant willing to lie can submit a combination of a a false and truthful report instead of a partial lie in each report. Thus, the latter lying strategy allows a participant to lie while avoiding the direct lying cost in one dimension. This means that a participant willing to lie will prefer to balance a lie and a truthful report between dimensions in case incurring a direct cost in only one dimension is less costly than splitting the direct cost between dimensions.

Moreover, if we consider image concerns, another motive for partial lying between dimensions is plausible. Since participants submit two separate reports in the 2D treatment, the image-related costs of lying will depend on participants' beliefs about whether the two reports taken together will be seen as truthful. Effron and Monin (2010) find that people are willing to forgive an ambiguous transgression if the person who is under suspicion also performed some good deed. Thus, when an ambiguous immoral act is accompanied by a moral act, people are more willing to assume that no transgression has taken place. In our 2D setting, this idea implies that coupling a lie in one dimension with a truthful report in the other dimension lowers the probability of being seen as a liar compared to having to report a lie in each report. Accordingly, by balancing between dimensions, a lying participant could reduce the likelihood of being seen as a liar.

Another motive to observe partial lying between dimensions in our setting might be that a truthful report in one dimension of the 2D treatment can act as a justification for lying in the other dimension. This conjecture is consistent with the literature on moral balancing (e.g., Blanken et al., 2015; Mullen and Monin, 2016). More specifically, according to this literature, we should expect that reporting truthfully in one dimension *licenses* participants to behave immorally in the other dimension.

Finally, we argue that if the discussed motives exist, and their magnitudes are substantial, partial lying between dimensions could cause overall lying in the 2D treatment to be higher than in the 1D treatment.

Based on this discussion, we formulate the following hypotheses:

Hypothesis 1. *The overall level of lying in the 2D treatment is higher than in the 1D treatment.*

Hypothesis 2. *In the 2D treatment, people lie partially between dimensions.*

2.3 Results

We start by analyzing the level of lying for the die-roll tasks in our setting. Figure 1 illustrates the distribution of the reported sums in each treatment across the ten rounds. We observe that both distributions are shifted towards higher reports. Sums on the left side of the distribution are reported less frequently while sums on the right side of the distribution are reported more frequently than expected under truth-telling, respectively.

Averaged over the ten rounds of the 1D treatment, participants reported an outcome of 7.3971, which indicates lying behavior. More specifically, the reports deviate significantly from the hypothetical outcome of two rolls of a fair die (p=0.0009,

Figure 1 Distribution of reported sums with the truthful distribution of rolling two fair dice as a reference

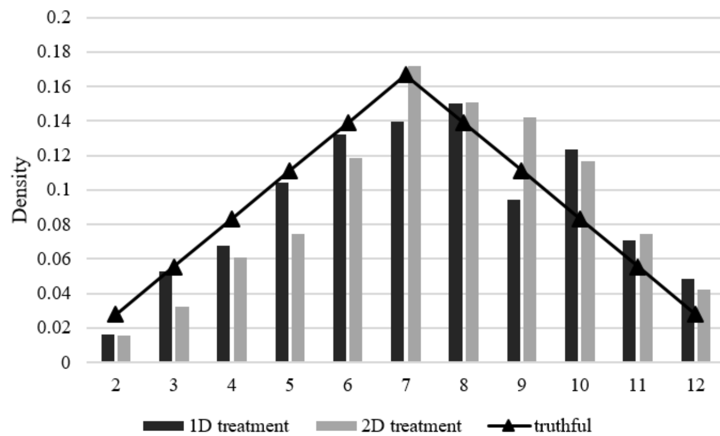

Kolmogorov–Smirnov one sample test (KS); p=0.0302, two-sided Wilcoxon signed-rank test (WSR)).[4] We reach the same conclusion when pooling the 10 rounds of the 2D treatment, in which participants reported an average sum of 7.6394 (p=0.0000, KS; p=0.0000, WSR).

2.3.1 Lying behavior: 1D vs. 2D

Our experiment allows contrasting the level of lying in a single-report task with the aggregate level of lying in a two-report task because we designed the range of the level of lying to be the same in both treatments.

Averaged over the ten rounds of the 2D treatment, participants' reports are higher than those of participants in the 1D treatment, but this difference is not significant (p=0.0666, Mann–Whitney (MW)). Hence, the results reject Hypothesis 1.

To further assess the level of lying, we run a regression analysis which is reported in Table 1. The dependent variable in this analysis is the average report over all rounds minus 7. Hence, the significant constant in Model (1a) indicates significant levels of over-reporting in the pooled data. In Model (1b), we run a bivariate regression with the 2D treatment dummy, which supports the earlier result that average reports do not differ significantly between treatments. This finding is robust when including control variables in Model (1c).

Result 1. *Reported sums in the 2D treatment are not significantly different than reported sums in the 1D treatment.*

[4]The distribution of two dice rolls has a bell-shaped distribution. Thus, as the KS is less sensitive to differences in the tail of distribution, we complement the test with the WSR test on the median.

Table 1 Linear regression analysis of the level of lying

Outcome variable: Average report - 7			
	(1a)	(1b)	(1c)
2D treatment		0.2424	0.3405
		(0.1870)	(0.1824)
constant	0.5209***	0.3971**	0.1683
	(0.0937)	(0.1336)	(0.2247)
Controls	No	No	Yes
N	139	139	137[a]
R2	0.0000	0.0121	0.0401

Standard errors in parentheses
Controls include gender, age and studying economics
[a]Two participants indicated gender=other.
* p < 0.05, ** p < 0.01, *** p < 0.001

2.3.2 Lying behavior between dimensions

The fundamental purpose of our experiment is to test whether lying balancing strategies exist between dimensions. More specifically, we hypothesized that people lie partially between dimensions in the 2D treatment.

To test lying balancing strategies, we perform two separate analyses. First, we analyze lying behavior between the two dimensions in the first round of the 2D treatment. Recall that participants learned about the subsequent nine rounds only after completing round 1. Thus, balancing between the two reports was the only available strategy in round 1. Second, we consider lying behavior across the subsequent nine rounds of the 2D treatment so that we test not only lying balancing between dimensions but also lying balancing between rounds. In the following, we analyze lying balancing between dimensions in round 1, as well as in the subsequent nine rounds. We test and discuss balancing between rounds in subsection 2.3.3.

In the 2D treatment, participants were instructed to roll one die in the left hand and the other one in the right hand. Thus, to analyze partial lying between dimensions, we examine the left and right dice. Reports on the right and left dice in the first round of the 2D treatment are not significantly different (3.6338 vs. 3.7183, p=0.7989, WSR). However, since the two dimensions are identical, participants had no reason to prefer lying about the left die over lying about the right one, or vice versa. Therefore, balancing between dimensions could be concealed by sometimes lying regarding the right die outcome and at other times lying about the left die outcome.

To gain more insight, we sort the two reports in the 2D treatment into lower and higher reports to test whether balancing between the two dimensions occurred based on the observed outcomes.[5] Under truth-telling, the expected value for the lower of

[5]In this analysis, we assume that participants make lower reports for the die with the lower outcome than for the die with the higher outcome. This assumption is supported by the distributions of reports in round 1 (see Figure 2) where reports closely mirror the shape of the expected distributions under truth-telling. The corresponding distributions for the remaining nine rounds show the same pattern (see Appendix A1).

two dice rolls is 2.5278 because, among the 36 possible combinations, there are 11 combinations in which 1 is the lower outcome (e.g., reporting 1 for the left die and 3 on the right, or reporting 5 for the left die and 1 on the right die), 9 combinations in which 2 is the lower outcome, and so forth. The expected value for the higher of two dice rolls is 4.4722, where 6 is the most likely outcome with a probability of 11/36.[6] Importantly, in what follows, we consider deviations from the respective expected value to compare over-reporting between the two outcomes because the higher and lower dice have different expected values under truth-telling.

In the first round, for the ordered reports—which we designate the *higher die* and the *lower die*—participants lied significantly more regarding the outcome of the lower die. Specifically, the average deviation from the expected value under truthful reporting is higher for the lower report than for the higher report, and this difference is significant (0.0348 vs. 0.3173, p=0.0393, two-sided t-test (TT)).[7] This indicates that more over-reporting occurred with the lower die. Regarding the subsequent nine rounds, we observe the same behavioral pattern. The average values of the left and right dice did not differ significantly (3.9155 vs. 3.7559, p=0.0914, WSR). Nevertheless, the participants' average reports on the lower die show higher deviation from the expected value than the average reports on the higher die, and this difference is significant (0.2586 vs. 0.4128, p=0.0168, TT).

Higher levels of over-reporting on the lower die could occur because there is more room for over-reporting on the lower die since the higher outcome is bounded from above. Therefore, to test the robustness for partial lying between dimensions, we also analyze the distribution of the high and low reports. Figures 2a and 2b show the distribution of high and low reports in round 1. The corresponding distributions for the remaining nine rounds show the same pattern (see Appendix A1). In the first round, neither the lower die reports nor the higher die reports are significantly different from the respective truthful distribution (p=0.1626 and p=0.9999, respectively, KS). Notably, however, a one-sided binomial test (BT) shows significant under-reporting of outcome 1 for the lower die report (p=0.0284, BT). This means that the lower reports of participants in the first round include the outcome 1 significantly less than predicted under truth-telling. Regarding the subsequent nine rounds, reports of the lower and higher outcomes are both significantly different than expected under truth-telling (p=0.0000 and p=0.0001, respectively, KS).

To compare how far reports (r) differ from the truthful distribution (t), we use the Euclidean distance from the distribution.[8] Regarding the lower reports, the Eu-

[6]Equal reports on both dice are included in the analysis, with the same value being assigned to the higher and lower dice, and the expected values under truth-telling are derived under the same assumption. Alternatively, we can exclude ties, and the expected values for the higher and lower dice are 2.3333 and 4.6667, respectively. The results are qualitatively the same under that specification.

[7]Since both reports are centered on the expected mean, we use a t-test to compare the means of the deviations.

[8]We cannot directly compare the truthful distributions of the low and high outcomes because they are different, and the approaches of earlier literature to compare lying in different tasks are not applicable to this situation. The standardization method of Abeler et al. (2019) assumes symmetric distributions under truth-telling, and the lying calculator of Garbarino et al. (2016) is limited to binary outcomes. It

Figure 2 Distribution of lower and higher reports in round 1 with truthful distributions as a reference in the 2D treatment

(a) Distribution of higher report in round 1

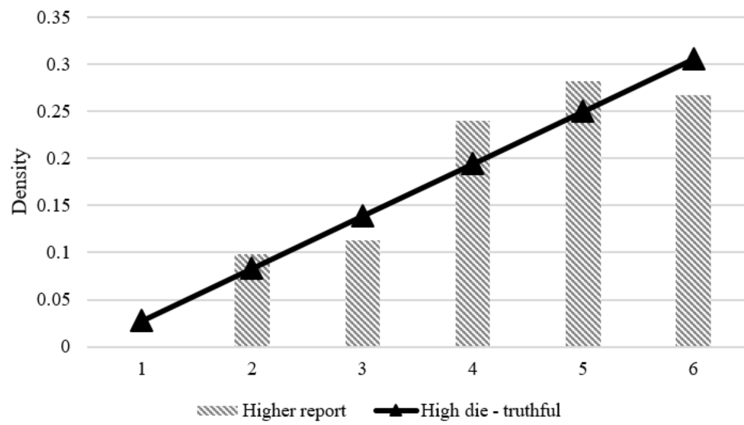

(b) Distribution of lower report in round 1

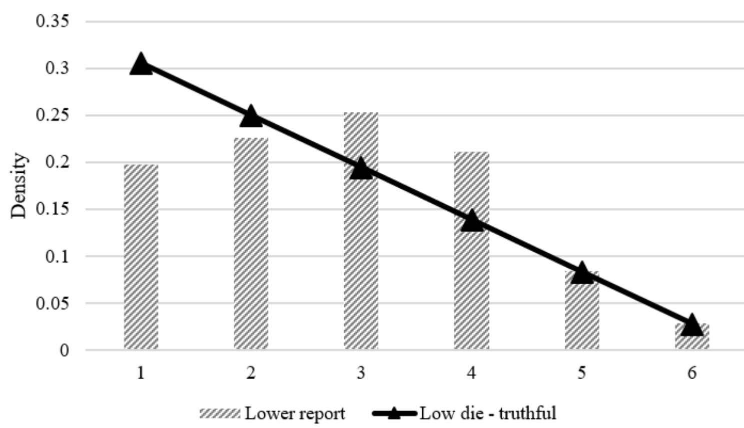

clidean distance to the truthful distribution is $d(r,t)_{low,1} = 0.1452$ in the first round and $d(r,t)_{low,2-10} = 0.1316$ in the subsequent nine rounds. Both these distances are significantly different from zero according to bootstrapped confidence intervals with 10000 replications. Regarding the higher reports, the analogous figures are $d(r,t)_{high,1} = 0.0785$ and $d(r,t)_{high,2-10} = 0.0847$. For the first round, the bootstrapped confidence interval of the Euclidean distance of the higher report includes zero, which means that we cannot reject the possibility that participants reported the higher outcome truthfully. For the subsequent nine rounds, zero is not included in the confidence interval, which indicates that participants also over-reported the higher outcome.

Finally, we assess the difference between the low and high reports. Regarding the

is possible to use the calculator for a task with multiple outcomes by bifurcating the reports into two payoff groups, but for the low die the low outcome would include reports of 1 and 2 while for the high die it would include 1, 2, 3, and 4. The Euclidean distance allows for a more symmetric treatment of both.

Euclidean distances, the distances are greater for the lower reports than for the higher reports in both the first round and the subsequent nine rounds. Econometric analysis supports the finding that people lie partially by lying more about the lower outcome. Comparing Models 1a and 1b in Table 2, we see that the lower die is more predictive of the difference between the two reports because the AIC of Model 1b is smaller. This indicates that over-reports on the lower die drive the difference between the two reports. Model 2a in Table 2 further shows that across the nine subsequent rounds, the average high report does not predict the difference between the two reports. Thus, if the difference between the two dice is not driven upwards for participants with high reports on the higher die, this necessarily implies that the report on the lower die was also high. Model 2b provides further evidence that participants adjusted their lower outcome more upwards than the high outcome: on average, high reports on the lower die significantly decrease the difference between the two reports, which indicates more over-reporting on the lower die. In short, the data support Hypothesis 2.

Result 2. *Partial lying between dimensions occurs, with more lying in the reported lower outcome.*

Table 2 Effect of higher and lower reports on the difference between the two reports in the 2D treatment

Round 1			Rounds 2–10		
Outcome variable	high die - low die		Outcome variable:	avg_highdie - avg_lowdie	
	(1a)	(1b)		(2a)	(2b)
highdie	0.3356***		avg_highdie	0.0923	
	(0.0778)			(0.1775)	
lowdie		-0.4063***	avg_lowdie		-0.5327***
		(0.0819)			(0.0517)
constant	0.1496	2.8179***	constant	1.3539	3.3568***
	(0.311)	(0.3019)		(0.8135)	(0.1686)
N	71	71	N	71	71
R2	0.1425	0.2338	R2	0.0075	0.4891
AIC	211.3222	203.3318	AIC	113.809	66.6595

Robust standard errors in parentheses
* p < 0.05, ** p < 0.01, *** p < 0.001

2.3.3 Order effects across rounds

In both treatments, repeating the task over ten rounds allowed participants to balance lying between rounds. Therefore, we analyze whether there are round effects, i.e., whether there is evidence of partial lying between rounds.

In Table 3, we see that the round effect is insignificant in both treatments, which indicates that participants did not change their behavior within rounds in a systematic fashion. Also, there is no evidence of participants balancing lying between rounds in the form of alternating a lie with a truthful report, which would have been the case if we had a significant negative coefficient of the lagged outcome variable. Figure 3 corroborates that participants made very similar lower reports across rounds in each treatment.

Result 3. *There is no evidence of partial lying between rounds.*

Figure 3 Reporting behavior over rounds

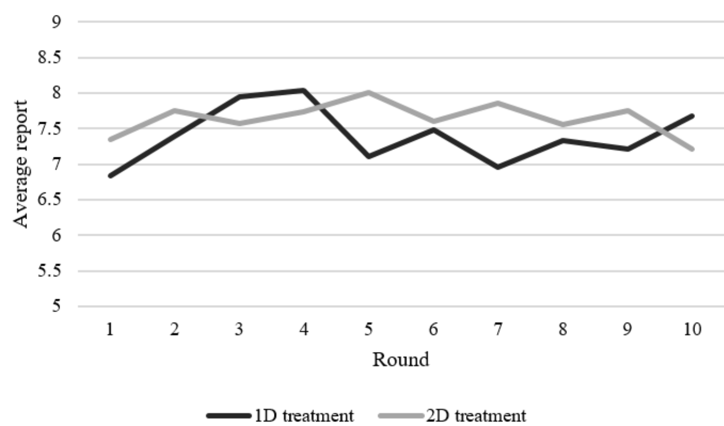

3 Experiment 2: Robustness check

3.1 Experimental design

To test the robustness of Experiment 1, we replicated the experiment with a different population. In Experiment 2, we collected data with adolescents in a German school.[9] Specifically, students completed exactly the same die-roll tasks as in Experiment 1 over ten rounds. The only distinctive design element was related to payoffs, which were determined as follows: (report-2)* €0.50. This means that the payoff could range between €0.00 and €5.00. We designed lower expected payoffs in Experiment 2 to account for the lower opportunity cost of adolescents.

3.2 Experimental procedure

The experiment was conducted on two consecutive days. Students were brought from their classroom to another room where computers were set up for the experiment. The workspaces were divided by screens such that students could not observe each other during the experiment. We minimized information transmission between students by conducting the experiment during lessons, with a maximum of two groups per class to prevent the students from talking to each other after completing the session. Upon

[9]We used a high school to implement the lab-in-the-field experiment because it offered the two features that we aimed for: i. A different population; ii. Ease of implementation.

Table 3 Two-step system GMM with lagged levels (t-2, t-3 and t-4) of the dependent variable as instruments

	1D		2D		2D		2D	
Outcome variable:	diesum		diesum		lowdie		highdie	
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
l.outcome	-0.028	-0.0218	0.0659	0.0527	-0.0284	-0.0752	0.1287*	0.1134*
	(0.0611)	(0.0569)	(0.0521)	(0.0525)	(0.0518)	(0.0501)	(0.0579)	(0.0568)
round		-0.042		-0.0329		-0.0036		-0.01
		(0.0360)		(0.0364)		(0.0241)		(0.0227)
constant	7.6168***	7.8212***	7.2170***	7.5132***	2.9902***	3.1680***	4.1773***	4.3047***
	(0.4855)	(0.4609)	(0.4181)	(0.4582)	(0.1821)	(0.2304)	(0.2711)	(0.2993)
Arellano–Bond serial	-0.24	-0.18	0.19	0.1	0.3	-0.14	-0.1	-0.18
correlation test [p-value]	[0.813]	[0.856]	[0.846]	[0.917]	[0.762]	[0.891]	[0.923]	[0.861]
Instruments	30	31	30	31	30	31	30	31
Sample size	612	612	639	639	639	639	639	639

Robust standard errors using Windmeijer correction in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

arrival in the room of the experiment, participants were randomly assigned to the 1D or 2D treatment, resulting in group sizes of 56 and 54 participants, respectively. The 110 participants (45% female, and average age 15.95, balanced over treatments) received all instructions on-screen using oTree software (Chen et al., 2016). The average payment per student was €2.8 and was paid out at the end of the school day. In total, the sessions lasted 10 minutes.

3.3 Results

Pooling all ten rounds, participants in Experiment 2 also showed significant lying behavior in both treatments. Figure 4 shows a tendency to under-report low values and to over-report high values in both treatments.[10] In the 1D treatment, participants reported an average outcome of 7.7875, and their reports are shifted significantly from the distribution expected under truth-telling (p=0.0000, KS; p=0.0000, WSR). In the 2D treatment, the average reports were 7.6852, which also differs significantly from the expectation under truth-telling (p=0.0000, KS; p=0.0001, WSR).

Figure 4 Distribution of reported sums with truthful distribution of rolling two fair dice as a reference in school sample

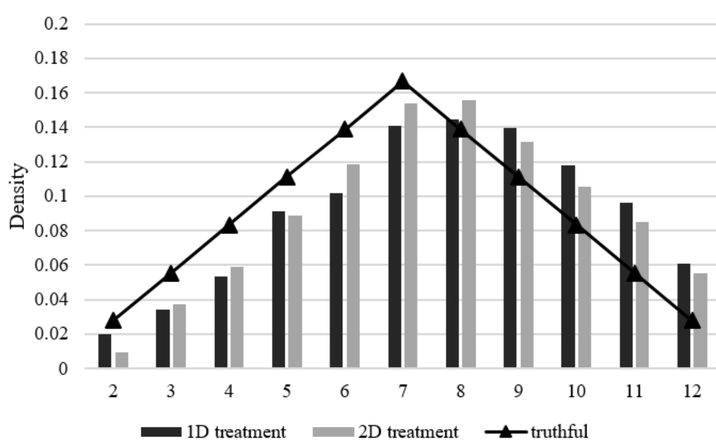

Results from the analysis of Experiment 2 are presented in Appendix A2. As in Experiment 1, there is no difference in the level of lying between the 1D and the 2D treatment (p=0.8225, MW test). The regression analysis with the average report-7 as the dependent variable confirms this result. In other words, Hypothesis 1 is again rejected in Experiment 2.

Concerning Hypothesis 2, lying about the right and left dice does not differ significantly in the first round (3.963 vs. 3.7963, p=0.4773, WSR) nor in the nine later rounds (3.784 vs. 3.893, p=0.3679, WSR). Yet, as in Experiment 1, we observe more lying in the lower report. Deviations from the expected value in Round 1 are greater for

[10] Comparing Figure 4 to Figure 1, we also see that there was a greater extent of lying in Experiment 2 than in Experiment 1. This side result corroborates the existing evidence suggesting that honesty is increasing with age (Friesen and Gangadharan, 2013; Fosgaard, 2018).

the low reports than for the high reports, but this difference is not significant (0.2315 vs. 0.5278, p=0.0790, TT). The same holds for the nine subsequent rounds (0.2706 vs. 0.4064, p=0.0677, TT). Figure 5 compares high and low reports in round 1 to the expected distributions under truth-telling. For the low report, the observed distribution is significantly different from the truthful distribution (p<0.01, KS test), which is driven by significant under-reporting of outcome of 1 (p=0.0004, BT). High reports do not differ significantly from the truthful distribution (KS p>0.2).

The analysis of Euclidean distances shows that in the first round, lower reports ($d(r,t)_{low,1} = 0.2629$) have a greater distance from the truthful distribution than high reports ($d(r,t)_{high,1} = 0.1006$), and that the estimate of lower reports is not included in bootstrapped confidence intervals with 10000 repetitions of higher reports and vice versa. In the nine subsequent rounds, the Euclidean distance is greater for lower reports ($d(r,t)_{low,2-10} = 0.1176$) than for higher reports ($d(r,t)_{high,2-10} = 0.0884$), but estimates are included in each other's confidence intervals. Finally, as in Experiment 1, the regression analysis on the difference between the two reports reveals that the lower die was more predictive of this difference (see Appendix A2).

In Experiment 2, there was again no evidence of round effects or partial lying between rounds (see Appendix A2). Overall, Experiment 2 supports the results of Experiment 1.

Result 4. *The results of Experiment 1 are validated in a different population. Most importantly, partial lying between dimensions via more lying in the lower outcome reported is also observed in Experiment 2.*

4 Discussion and concluding remarks

This paper contributes to a growing body of literature on lying behavior. To the best of our knowledge, the present research is the first attempt in experimental economics to assess the effect of dimensionality on individual lying behavior.

The fundamental purpose of this study was to experimentally investigate multi-dimensional decision-making. More specifically, we analyzed a two-dimensional lying context to understand how the dimensionality of a decision affects lying behavior. Also, we aimed to contrast lying behavior in a two-dimensional context to a one-dimensional context.

First and foremost, our results uncover an interesting lying trait in the two-dimensional context: People lie partially between dimensions. More specifically, people systematically lie more in the dimension in which there is more room to lie. Importantly, this laboratory result is replicated in a lab-in-the-field experiment with participants possessing different characteristics. Additionally, we find no significant difference in the levels of lying in the 2D context compared to the 1D context.

An important implication of our results is that we should not extrapolate lying behavior from one dimension to another. In light of our results, the level of lying differs

Figure 5 Distribution of lower and higher reports in round 1 with truthful distributions as reference in the school sample

(a) Distribution of higher reports in round 1

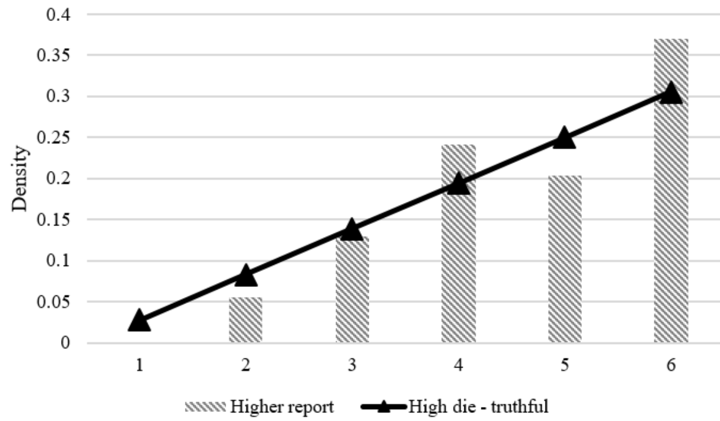

(b) Distribution of lower reports in round 1

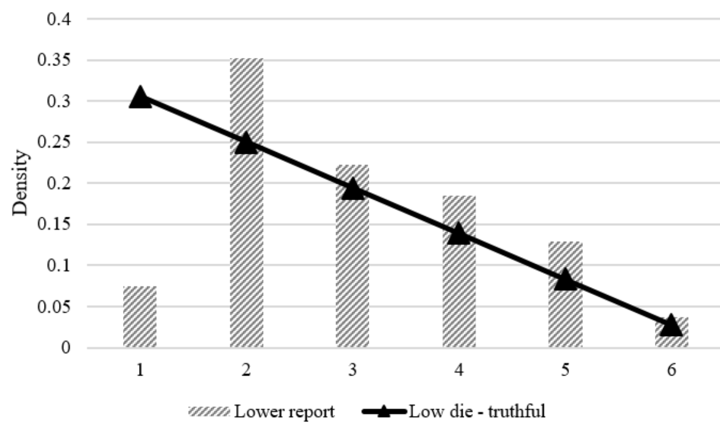

between dimensions because people lie partially between dimensions. Thus, extrapolating levels of lying from one dimension to another dimension might lead to inaccurate estimates of lying behavior. In the context of tax declarations, this would mean that if fraud is detected in a declared item, we cannot take that figure to estimate fraud related to other items or on the declaration as a whole.

While this paper is an important first step in understanding multi-dimensional lying, our research yields questions in need of further investigation. First, an evident follow-up is to test whether partial lying between dimensions holds in a setting with several dimensions. Second, we did not find support for the hypothesis that a 2D context increases overall lying behavior. However, a possible explanation for the observed similar level of lying in both contexts could be that the 1D and 2D contexts were too identical. In particular, we cannot rule out that participants in the 1D context also used a partial lying strategy since they also rolled two dice. Nevertheless, if the latter conjecture is true, that would be a further argument for the importance of studying a multi-dimensional reporting context: It uncovers traits of lying behavior that are concealed in a one-dimensional reporting context.[11] Third, from a laboratory experimental methodological perspective, it was advisable as a first step conducting a clean test of the effect of dimensionality on lying behavior. Yet, a natural and relevant extension to our work is to study lying in a multi-dimensional setting while making the dimensions dissimilar.

More research is needed to evaluate the external validity of our main finding. Nevertheless, if partial lying between dimensions proves robust in other multi-dimensional settings, our findings have a clear policy implication for detecting fraud in tax reports: When checking the correctness of tax reports, authorities should focus on item(s) for which the taxpayer profile hints at more tempting self-benefits.

In closing, we directed the discussion of policy implications to the tax evasion context. Unfortunately, however, fraud is ubiquitous. Hence, if future research corroborates that partial lying between dimensions is a robust lying trait in multi-dimensional decision-making, one could possibly derive policy implications for a number of other important societal challenges, such as school, science, and sports fraud.

[11]An alternative explanation for not observing a higher level of lying in the 2D context is that a costly moral action in one dimension does not justify an immoral action in the other dimension. Gneezy et al. (2012) show that the cost involved with acting morally requires high intrinsic motivation to be honest and that the emphasis on intrinsic motivation leads to consistently moral behavior.

References

Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4):1115–1153.

Barron, K., Stüber, R., and van Veldhuizen, R. (2019). Motivated motive selection in the lying-dictator game. Discussion Papers, Research Unit: Economics of Change SP II 2019-303, WZB Berlin Social Science Center.

Blanken, I., van de Ven, N., and Zeelenberg, M. (2015). A Meta-Analytic Review of Moral Licensing. *Personality and Social Psychology Bulletin*, 41(4):540–558.

Cain, D. M., Loewenstein, G., and Moore, D. A. (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies*, 34(1):1–25.

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree — An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Cojoc, D. and Stoian, A. (2014). Dishonesty and charitable behavior. *Experimental Economics*, 17(4):717–732.

Dolan, P. and Galizzi, M. M. (2015). Like ripples on a pond: Behavioral spillovers and their implications for research and policy. *Journal of Economic Psychology*, 47(Supplement C):1–16.

Dufwenberg, M. and Dufwenberg, M. (2018). Lies in Disguise - A Theoretical Analysis of Cheating. *Journal of Behavioral and Experimental Finance*, 175:248–264.

Effron, D. A. and Monin, B. (2010). Letting people off the hook: When do good deeds excuse transgressions? *Personality and Social Psychology Bulletin*, 36(12):1618–1634.

Engel, J. and Szech, N. (2017). A little good is good enough: Ethical consumption, cheap excuses, and moral self-licensing. CESifo Working Paper Series 6434.

Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in Disguise—an Experimental Study on Cheating. *Journal of the European Economic Association*, 11(3):525–547.

Fosgaard, R. T. (2018). Students cheat more: Comparing dishonesty of a student and a representative sample in the laboratory. *The Scandinavian Journal of Economics*, forthcoming.

Friesen, L. and Gangadharan, L. (2013). Designing self-reporting regimes to encourage truth telling: An experimental study. *Journal of Economic Behavior & Organization*, 94:90–102.

Garbarino, E., Slonim, R., and Villeval, M. C. (2016). Loss Aversion and lying behavior: Theory, estimation and empirical evidence. GATE Working Paper 1631.

Gerlach, P., Teodorescu, K., and Hertwig, R. (2019). The truth about lies: A meta-analysis on dishonest behavior. *Psychological bulletin*, 145(1):1.

Gneezy, A., Imas, A., Brown, A., Nelson, L. D., and Norton, M. I. (2012). Paying to be nice: Consistency and costly prosocial behavior. *Management Science*, 58(1):179–187.

Gneezy, U. and Imas, A. (2017). Lab in the field: Measuring preferences in the wild. In *Handbook of economic field experiments*, volume 1, pages 439–464. Elsevier.

Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying Aversion and the Size of the Lie. *American Economic Review*, 108(2):419–53.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.

Guadagno, R. E., Asher, T., Demaine, L. J., and Cialdini, R. B. (2001). When saying yes leads to saying no: Preference for consistency and the reverse foot-in-the-door effect. *Personality and Social Psychology Bulletin*, 27(7):859–867.

Jacobsen, C., Fosgaard, T. R., and Pascual-Ezama, D. (2017). Why Do We Lie? A Practical Guide to the Dishonesty Literature. *Journal of Economic Surveys*, 32(2):357–387.

Khalmetski, K. and Sliwka, D. (2017). Disguising Lies - Image Concerns and Partial Lying in Cheating Games. CESifo Working Paper Series 6347.

Khan, U. and Dhar, R. (2006). Licensing Effect in Consumer Choice. *Journal of Marketing Research*, 43(2):259–266.

Mazar, N., Amir, O., and Ariely, D. (2008). The Dishonesty of Honest People: A Theory of Self-Concept Maintenance. *Journal of Marketing Research*, 45(6):633–644.

Merritt, A. C., Effron, D. A., and Monin, B. (2010). Moral self-licensing: When being good frees us to be bad. *Social and personality psychology compass*, 4(5):344–357.

Mullen, E. and Monin, B. (2016). Consistency Versus Licensing Effects of Past Moral Behavior. *Annual Review of Psychology*, 67(1):363–385.

Ploner, M. and Regner, T. (2013). Self-image and moral balancing: An experimental analysis. *Journal of Economic Behavior & Organization*, 93(Supplement C):374–383.

Potters, J. and Stoop, J. (2016). Do cheaters in the lab also cheat in the field? *European Economic Review*, 87:26–33.

Rosenbaum, S. M., Billinger, S., and Stieglitz, N. (2014). Let's be honest: A review of experimental evidence of honesty and truth-telling. *Journal of Economic Psychology*, 45:181–196.

Schindler, S. and Pfattheicher, S. (2017). The frame of the game: Loss-framing increases dishonest behavior. *Journal of Experimental Social Psychology*, 69:172–177.

Shu, L. L., Mazar, N., Gino, F., Ariely, D., and Bazerman, M. H. (2012). Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. *Proceedings of the National Academy of Sciences*, 109(38):15197–15200.

Appendix

Appendix A1: Additional results Study 1

Figure A1.1 Distribution of lower and higher reports in rounds 2–10 with truthful distributions as reference

(a) Distribution of higher reports in rounds 2–10

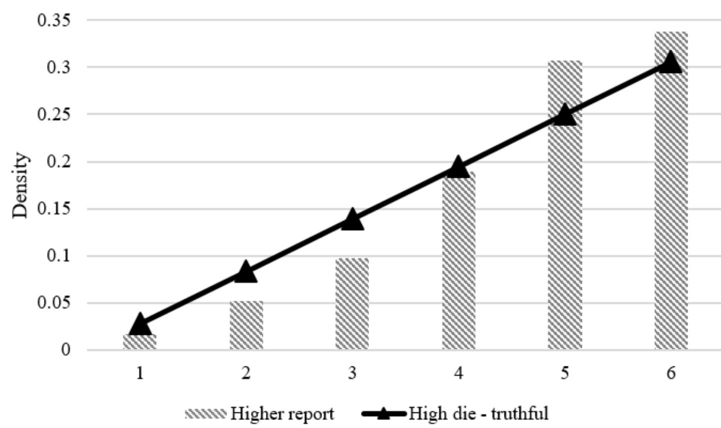

(b) Distribution of lower reports in rounds 2–10

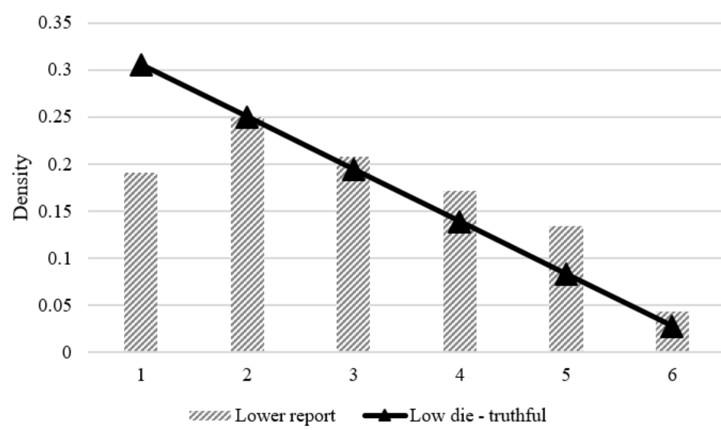

Appendix A2: Additional results Study 2

Table A2.1 Linear regression analysis of the level of lying in Experiment 2

Outcome variable: Average report - 7			
	(1a)	(1b)	(1c)
2D treatment		-0.1023	-0.1108
		0.2105	0.2145
constant	0.7373***	0.7875***	0.7750***
	0.1049	0.1475	0.1627
Controls	No	No	Yes
N	110	110	110
R2	0	0.0022	0.0049

Standard errors in parentheses
Controls include gender and age
* p < 0.05, ** p < 0.01, *** p < 0.001

Table A2.2 Effect of higher and lower reports on the difference between the two reports in Study 2

Round 1			Rounds 2–10		
Outcome variable: high die - low die			Outcome variable: avg_highdie - avg_lowdie		
	(1a)	(1b)		(2a)	(2b)
highdie	0.4248***		avg_highdie	0.1969	
	(0.1165)			(0.1451)	
lowdie		-0.4728***	avg_lowdie		-0.5297***
		(0.1022)			(0.0781)
cons	-0.3501	3.0928***	cons	0.8746	3.3629***
	(0.4647)	(0.4257)		(0.6603)	(0.2489)
N	54	54	N	54	54
R2	0.1919	0.2593	R2	0.0352	0.435
AIC	165.8164	161.1148	AIC	86.7503	57.8568

Robust standard errors in parentheses
* p < 0.05, ** p < 0.01, *** p < 0.001

Table A2.3 Two-step system GMM with lagged levels (t-2, t-3 and t-4) of the dependent variable as instruments in Study 2

Outcome variable:	1D diesum		2D diesum		2D lowdie		2D highdie	
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
l.outcome	0.0884	0.1164	-0.0312	-0.0039	-0.0265	-0.0053	-0.0247	-0.0278
	(0.0683)	(0.0681)	(0.0685)	(0.0642)	(0.0654)	(0.0625)	(0.0676)	(0.0665)
round		0.0239		0.0282		0.0196		0.0068
		(0.0420)		(0.0422)		(0.0298)		(0.0243)
constant	7.0358***	6.6754***	7.8512***	7.4618***	2.946***	2.756***	4.848***	4.8202***
	(0.5884)	(0.5654)	(0.5738)	(0.4826)	(0.2215)	(0.2173)	(0.3532)	(0.3413)
Arellano-Bond serial	-0.02	0.14	0.16	0.33	-0.74	-0.57	0.26	0.24
correlation test [p-value]	[0.987]	[0.3960]	[0.877]	[0.739]	[0.461]	[0.570]	[0.794]	[0.811]
Instruments	30	31	30	31	30	31	30	31
Sample size	504	504	486	486	486	486	486	486

Robust standard errors using Windmeijer correction in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure A2.1 Distribution of lower and higher reports in rounds 2–10 with truthful distributions as reference in the school sample

(a) Distribution of higher report in rounds 2–10

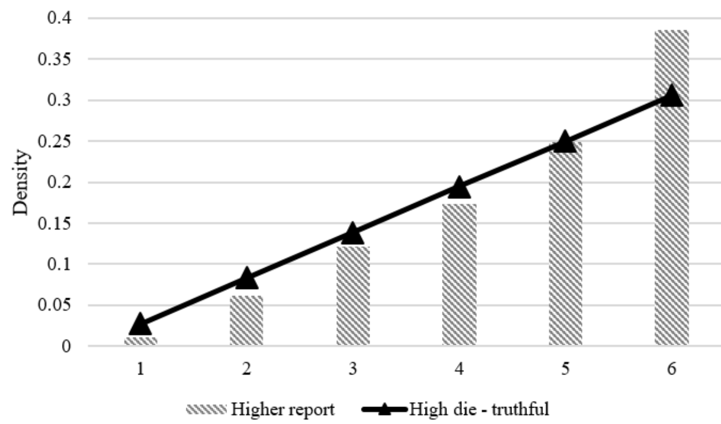

(b) Distribution of lower report in rounds 2–10

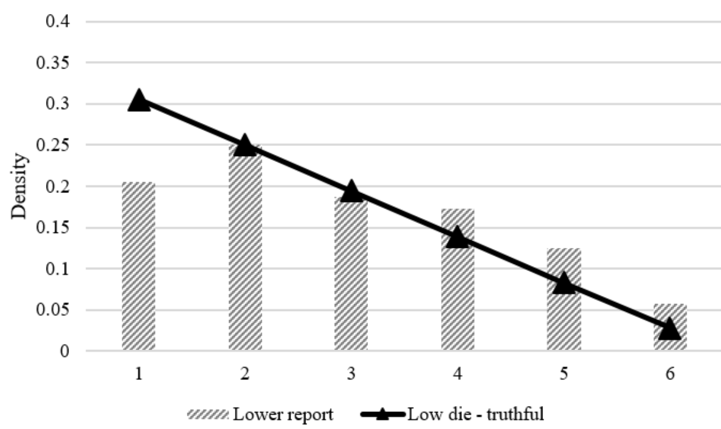

Appendix B: Experimental instructions

The following shows the instructions for the 1D (2D) treatment. Instructions were displayed on-screen and in English in Experiment 1 and in German in Experiment 2.

Welcome
Thank you for participating in this experiment. All instructions will be given on-screen. You are now taking part in an experiment in decision making. For showing up on time today, you will be paid €2. In addition, you can earn money with the decisions you make. Hence, it is important that you fully understand the instructions that follow. Please read them carefully.
Please enter the number of your computer so that your earnings at the end of the experiment can the matched to your computer cubicle. We use this number only to determine your payment. This means that the number will not be linked to your name for the data analysis. Your anonymity is therefore secured.
Please raise your hand if you have a question at any point of the experiment.

Instructions
In this experiment you will receive money based on the outcome of a rolling two dice. You can find two dice in front of you on the table. You can inspect them now to see that they are regular 6-sided dice.
On the following screen you will be asked to roll both dice at the same time and to report the sum of eyes that you see. (On the following screen you will be asked to roll two dice, which we label "right die" and "left die", respectively, at the same time and to report the outcome of each die roll separately.) Your payment will be:
Sum of both dice*€0.50
(Outcome "right die"*€0.50 + Outcome "left die"*€0.50)
Example 1: If the sum of the two dice is 12 (If the outcome of both dice is 6), you will receive €6.
Example 2: If the sum of the two dice is 2 (If both dice show a 1), you will receive €1.

Figure B.1 Reporting screen 1D treatment

Report 1

Please now take both dice in one hand and roll them as shown in the video below:

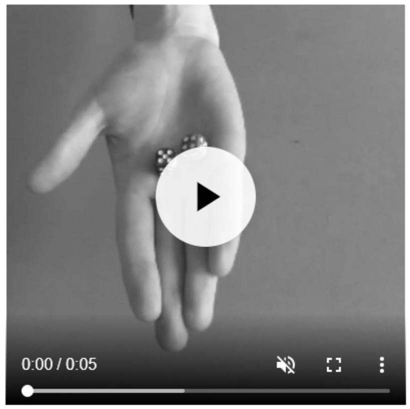

What is the sum of the two dice?

Next

Figure B.2 Reporting screen 2D treatment

Report 1

Please now take one die in your right hand and one in your left hand and roll both of them as shown in the video below:

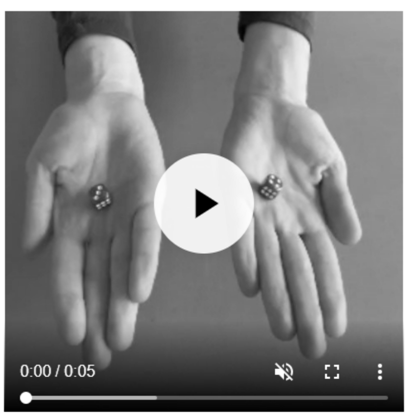

What is the outcome of the right die?

What is the outcome of the left die?

Next