Reputation and lying aversion in the die roll paradigm: Reducing ambiguity fosters honest behavior

Ann-Kathrin Crede[1], Frauke von Bieberstein[2]

Abstract

This paper examines reputation as motive for lying aversion. In a control treatment, participants roll a six-sided die and report the outcome, which the experimenter cannot observe. In a digital die treatment, the outcome of the die roll is determined randomly on the computer. Contrary to prior literature, we reduce ambiguity in the digital die treatment by making observability common knowledge. We find that partial lying and full lying disappear when the experimenter can track participants' behavior. This result can be explained by reputational costs: Participants care about how they are viewed by the experimenter and abstain from lying.

[1]ann-kathrin.crede@iop.unibe.ch, Institute for Organization and Human Resource Management, University of Bern.

[2]frauke.vonbieberstein@iop.unibe.ch, Institute for Organization and Human Resource Management, University of Bern.

1 Introduction

In many economic settings, people have to decide whether to lie about their private information. The standard economic model predicts that an individual lies when the utility from the dishonest act outweighs the potential punishment in case of detection (Becker 1968). However, numerous empirical studies have found that not everyone behaves dishonestly even if dishonesty has no personal consequences, and people are rarely dishonest to the maximum extent (Jacobsen et al. 2018). For instance, managers do report truthfully (Evans III et al. 2001), people hand in full wallets to police stations (West 2005) and customers pay for a newspaper that they buy out of a box on the street (Pruckner and Sausgruber 2013). For the die roll paradigm by Fischbacher and Föllmi-Heusi (2013), the most widely used method to study lying behavior in the lab, Abeler et al. (2019) find in a meta-analysis that people forgo, on average, about three-fourth of the potential gains from lying.

The die roll paradigm studies lying behavior in a non-strategic environment with a simple experimental design: Participants privately roll a six-sided die and report the outcome, which determines their payoff. The higher the number they report, the higher the payoff, unless it is a six, which yields a payoff of zero. Because the experimenter cannot observe the outcome, participants have an incentive to report a higher number than they actually rolled (except six) to maximize their payoff. Although it is impossible for the experimenter to tell whether a particular individual lied or not, lying can be detected on an aggregate level, as the

underlying true distribution of a six-sided die roll is known. The authors identify three different types of behavior: There are participants who report truthfully, others who report the highest payoff, and those who lie partially, i.e. who report a high, but not the maximum, payoff. Whereas the first two types of behavior are easily explained with direct lying costs or standard preferences, respectively, partial lying is less straightforward. Fischbacher and Föllmi-Heusi (2013) suggest reputational concerns as an explanation for partial lying leading participants to disguise lies to an outside observer (the experimenter or a future self of the participant).

Recently, a number of theoretical contributions have started to systematically study the underlying motives for the aversion to lie. Gneezy et al. (2018), Abeler et al. (2019), and Khalmetski and Sliwka (2019) address reputational costs theoretically, and present models that incorporate the desire to appear honest in the utility function. In a different theoretical approach, Dufwenberg and Dufwenberg (2018) assume that individuals care about the beliefs others have about the *extent* of lying (instead of the chance of being a liar at all). Interestingly, they can account for the above mentioned three types of behavior (truth telling, partial lying, and full lying) without assuming heterogeneity in types and without imposing a direct lying cost. Differences in behavior emerge as part of a mixed strategy equilibrium in a psychological game (Geanakoplos et al. 1989, Battigalli and Dufwenberg 2009).[3] This result emerges when the experimenter cannot observe the true number. The authors propose as a testable implication of their

---

[3]The game is solved using psychological game theory because the utility of the decision maker depends on what she thinks the observer believes about the extent to which she lied.

model the case when the experimenter can observe the truth. In this case, predictions depend on how much participants care about the beliefs of the experimenter: if they care only little, full lying emerges, whereas if they care more strongly, all participants tell the truth.

To study the effect of the perception of honesty empirically, both Gneezy et al. (2018) and Abeler et al. (2019) independently from each other introduced such a treatment where the experimenter is able to observe the truth (a variant of the die roll is executed digitally).[4] Gneezy et al. (2018) find that partial lying, a behavior that could help the participant to appear more honest in the unobservable treatment, is strongly reduced when the truth is observable, while full lying prevails in both treatments. Abeler et al. (2019) find similar effects. In combination with their findings from other treatments, both authors conclude that people seem to have a preference for being honest and for being perceived as honest by other people. Importantly, in both studies, the instructions do not explicitly mention that the experimenter can verify the real outcome ex post.

In this paper, we build on the studies of Gneezy et al. (2018) and Abeler et al. (2019) and take away part of the ambiguity that prevails in their observable treatments. In particular, even if the variant of the die roll is performed on the computer, some participants might be unsure whether the experimenter can and will check the true outcome. This reasoning could even be a form of advantageous "self-deception" to persuade oneself that misconduct cannot be tracked. Consequently, some participants might act as if they were in the unobserved treatment

_____

[4]Gneezy et al. (2018) use a single-blind procedure whereas Abeler et al. (2019) employ a double-blind procedure.

when in fact the outcome is observable. In contrast, we take away part of the ambiguity by explicitly mentioning that the outcome of the die roll is electronically stored, but not verified for the payment procedure.[5]

Whereas results in our control treatment with an analog die are in line with Fischbacher and Föllmi-Heusi (2013), Gneezy et al. (2018) and Abeler et al. (2019), partial lying and full lying both disappear in the unambiguous digital die treatment. Thus, whereas some participants might be able to persuade themselves that results cannot be checked when this possibility is not explicitly mentioned, in our unambiguous treatment this is no longer possible. Participants seem to care about how they are viewed by the experimenter, and abstain from lying when they are observed. The most likely underlying reason are reputational costs that can be due to an intrinsic preference for appearing honest, or due to the fear of actual consequences, such as being excluded from future experiments.

Our results fit very well with the model of Dufwenberg and Dufwenberg (2018).[6] In their model for unobservable outcomes, they present a "sailing-to-the-ceiling" equilibrium and apply it to the data presented in Fischbacher and Föllmi-Heusi (2013). For a value of $\theta = 3$, capturing how strongly an individual cares about the beliefs others have about the extent of lying, the predicted data from their model shows a striking similarity to the results in Fischbacher and Föllmi-Heusi (2013). Given that we find very similar behavior in our control treatment, we could assume that a value of approx. $\theta = 3$ also prevails in our subject pool. For this value (more exactly, for all $\theta > 1$), the model by Dufwenberg and Dufwenberg (2018)

---

[5]We do not address ambiguity regarding whether the outcome of the die roll will be verified at all at a later point.

[6]We are grateful to an anonymous referee for this observation.

exactly predicts that everybody tells the truth (no more partial or full lying) when the outcome is observable and everybody is aware of this fact.

The role of ambiguity is also covered in the models presented in Gneezy et al. (2018), Abeler et al. (2019), and Khalmetski and Sliwka (2019). For instance, Abeler et al. (2019) consider preferences that depend on the reputation for honesty. When the true state is observable, the probability of being a liar is either 0 or 1, whereas the probability can take intermediate values when the true state is unobservable. We argue that probabilities strictly between 0 and 1 are also possible in the observable state, if some participants are unsure and form beliefs about observability.[7]

Why do people experience reputational costs? The literature on *social image* suggests that many people care about how they are viewed by people in their surroundings, and that social image concerns can strongly affect economic behavior (see Bursztyn and Jensen 2017 for an overview). For example, people vote to be able to signal to others they voted (DellaVigna et al. 2016), work more for social recognition (Kosfeld and Neckermann 2011), and prefer donating in public

---

[7]Gneezy et al. (2018) consider the utility function to depend on direct monetary payoffs, a direct cost of lying, and a social identity cost. This social identity cost contains an argument for situations when the true state is publicly observed ($\lambda = 1$) that only depends on whether the participant lied or not, and an argument when only the participant knows the true state ($\lambda = 0$) that depends on the probability that a report is interpreted as being honest (see equations (2)-(5)). The authors discuss that a value of λ strictly between zero and one can be appropriate in the non-observed game, if the agent wrongly believes that the experimenter knows the true state. Here, we argue that a similar reasoning can apply to the observed treatment when participants are not sure of whether the experimenter can and will check the truthfulness of the report. Consequently, some participants might act as if they were in the non-observed treatment when using the digital die. Khalmetski and Sliwka (2019) discuss that when the subject is absolutely sure that her lying will be verified ex post by the experimenter, in their model the image loss from lying would be captured by the fixed cost of lying. We argue that not all subjects might be sure that their lying can and will be tracked, and thus there is not only the experimenter's perception of the subject's honesty but also the subject's perception of the experimenter's actions after the experiment that needs to be considered.

than in private (Ariely et al. 2009). With respect to the decision whether to lie in a laboratory experiment, participants' social image is more at risk when the experimenter can observe them lying compared to the case when the truth is not observable. Some studies support this idea, showing that people are more likely to behave according to a social norm when an audience is present that observes their choices (e.g., Kurzban et al. 2007, Andreoni and Bernheim 2009, Cohn et al. 2018).

Besides reputational costs covered in recent models discussed above, Abeler et al. (2019) identify two additional classes of models: First, people experience direct lying costs; that is, not telling the truth involves a utility loss (e.g., Kartik 2009, Gibson et al. 2013, López-Pérez and Spiegelman 2013). Second, people's decision whether to behave dishonestly might be influenced by social norms and social comparisons (e.g., Weibull and Villa 2005, Rauhut 2013, Dieckmann et al. 2016). For instance, groups are more inclined to behave dishonestly than individuals, which is driven by communication and learning about norm compliance among group members (Kocher et al. 2017).

This paper contributes to the existing literature by taking away ambiguity when the experimenter can observe the true outcome in the die roll paradigm. Our findings have important implications: In an age of increased digitization, behavior can be tracked more and more easily, and people leave digital marks. By making this potential observability more salient, one could take advantage of the power of reputational costs to keep people from engaging in dishonest behavior.

The paper proceeds as follows: In section 2, we present the experimental design

7

and procedure. In section 3, we show the results. In section 4, we discuss the results and conclude.

2 Experimental Design

The experiment follows the design by Fischbacher and Föllmi-Heusi (2013) and consists of a short questionnaire to justify the payoff participants could earn (see Appendix A), the rolling of a die, and the (self-)reporting of the payoff. Participants learned that the payoff for completing the questionnaire would be different for every participant, and that a die roll served to determine the exact size of this payoff. Paying participants different amounts for filling out the same questionnaire might seem questionable. However, this way might be more credible than mentioning nothing about the purpose of the die roll to avoid an experimenter demand effect. The payoff amounted to € 1, 2, 3, 4, or 5 if participants rolled a 1, 2, 3, 4, or 5, respectively, and € 0 if they rolled a 6.

2.1 Treatments

In the control treatment, participants got a cup and a six-sided die, and were asked to roll the die to determine their payoff. The instructions (see Appendix A) explicitly explained that they should remember, and report, the outcome of the first die roll, but that they could roll the die several times to verify that the die was fair. After rolling the die once or more often, participants had to enter the outcome of their first die roll and the resulting payoff they would get on the

computer. This treatment is identical to Fischbacher and Föllmi-Heusi (2013) and very similar to the unobserved treatments in Gneezy et al. (2018) and Abeler et al. (2019), except the medium used to determine the random draw (a die instead of pieces of paper or chips).

In the digital die treatment, participants had to determine the outcome of a die roll on the computer. In the instructions, it was mentioned that the outcome of the die roll would be electronically saved, but not checked by the experimenter for the payment procedure after the experiment. This is different from Gneezy et al. (2018) and Abeler et al. (2019), who did not explicitly mention that the experimenter was able to track outcomes ex post. We added this sentence to prevent participants from forming different beliefs about this possibility. Participants saw a screen with a button labeled *Rolling die*. By clicking on this button, a number between 1 and 6 was randomly generated and displayed next to the button (see Appendix B for a screenshot). As in the control treatment, participants learned that they should remember, and enter, the outcome of the first die roll, but that they could click the button more often to verify that the random generator worked properly. Similar to Gneezy et al. (2018), but unlike Abeler et al. (2019), we did not employ a double-blind payment procedure for the digital die treatment.

2.2 Procedure

We conducted the experiment at the Business and Economic Research Laboratory at the University of Paderborn in December 2017.[8] The experiment was

---

[8]We preregistered the experiment in the American Economic Association's registry for randomized controlled trials. See https://www.socialscienceregistry.org/trials/2607/history/23742.

implemented in z-Tree (Fischbacher 2007), and participants were recruited with ORSEE (Greiner 2015). Given that the experiment took only a very short time, we added it to the sessions of another experiment. After participants finished the first experiment, they learned that there would be a second, very short experiment that was unrelated to the first one. At the end of the session, participants received their payoff from the first experiment and the payoff for this experiment. In total, 140 participants took part in this study. We ran two sessions with the control treatment ($n = 56$) and three sessions with the digital die treatment ($n = 84$). Because of potential spillover effects from hearing other participants rolling a real die, it was not possible to run a session that included both treatments. In our regression analysis, we control for sessions and for the amount of money gained in the preceding experiment, but do not find significant effects.

3   Results

We first present the main variable of interest; that is, the payoff participants reported. Participants were asked to report a payoff of € 1, 2, 3, 4, or 5 if their die showed a 1, 2, 3, 4, or 5, respectively, and a payoff of € 0 if their die showed a 6.[9] Figure 1 shows the distribution of the reported payoffs in percentage in the control treatment.

---

[9]Participants reported both, the outcome of the die roll and the resulting payoff. They could only proceed to the next page once the reported outcome and payoff coincided.
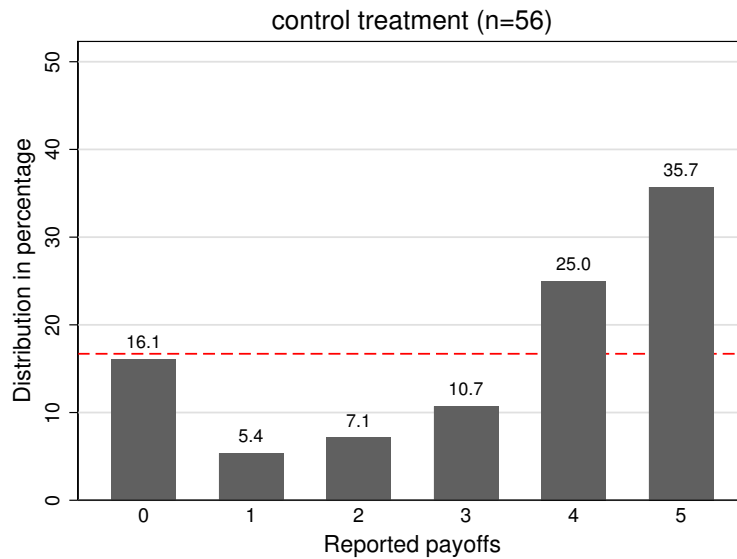
Figure 1: Distribution of reported payoffs in percentage in the control treatment. The red dashed line shows the uniform distribution.

As Figure 1 shows, the distribution is not uniform. We run two-sided binomial tests to check whether the observed percentage for each reported payoff differs from the theoretically predicted percentage of 16.7% (1/6) that would result in the case of full honesty (see the red dashed line). The observed percentage differs significantly from 16.7% for 1 (p=0.019) and 5 (p=0.001), differs weakly significantly for 2 (p=0.070) and 4 (p=0.106), and does not differ significantly for 0 (p=1.00) and 3 (p=0.284). These results are in line with the three types of behavior identified by Fischbacher and Föllmi-Heusi (2013). There are subjects who report truthfully (there is a positive fraction of participants who report 0), there are subjects who maximize their income (the fraction of participants reporting 5 is higher than 16.7%), and there are participants who lie partially (the fraction of participants reporting 4 is higher than 16.7%, though only weakly statistically significant). These results are in line with those of Gneezy et al. (2018) and Abeler

et al. (2019), who identify partial lying and full lying in their unobserved treatments. The main difference between our data and the results in Fischbacher and Föllmi-Heusi (2013) is that we find 16.1% of participants report 0, compared to 6.4% in their study. Thus, we find more honesty, and cannot fully replicate their monotonically increasing distribution.

In the next step, we consider how participants in the digital die treatment behaved. Figure 2 shows the distribution of reported payoffs in percentage by participants who rolled the die digitally on the computer.
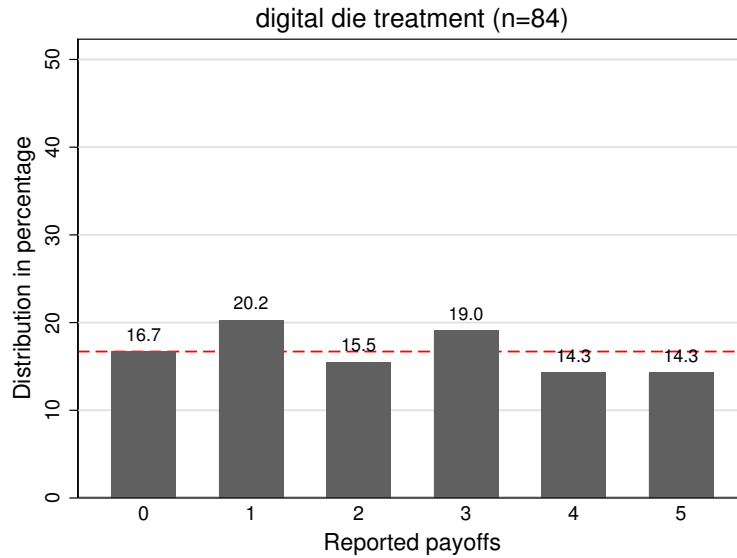

Figure 2: Distribution of reported payoffs in percentage in the digital die treatment. The red dashed line shows the uniform distribution.

As Figure 2 shows, the distribution is very close to uniform. Two-sided binomial tests for the observed percentage for each payoff reveal that none differs significantly from 16.7%. Thus, replacing a standard die with a digital die seems to prevent participants from lying, more specifically from partial lying and full lying (see Appendix C for a comparison of the distribution of reported payoffs in the

control treatment with the distribution of reported payoffs in the digital die treatment). Compared to Gneezy et al. (2018) and Abeler et al. (2019), we do not find that full lying remains when the outcome of the random draw can be tracked.

To confirm that participants abstain from lying when being observed, we can compare the outcome of the digital die roll and the reported payoff, as the outcome of the digital die was saved. Of the 84 participants in the digital die treatment, 78 participants (93%) truthfully reported the first number that the digital die showed. Correspondingly, only 6 participants entered a result that did not coincide with the true outcome of the digital die.[10] This finding supports the result that replacing the standard die with a digital die leads to more honesty, and not simply to a uniform distribution due to potentially other (opposing) changes in behavior.

Finally, we run an ordinary least squares (OLS) regression with the reported payoff as the dependent variable and the *digital die* treatment, the amount of *money* gained in the preceding experiment, the *session*, *age*, and being *male* as explanatory variables. Table 1 provides the results.

---

[10]Out of these 6 participants, 3 participants lied upward, and reported a higher outcome than the digital die actually showed (a 3 instead of a 1, a 5 instead of 2, and a 2 instead of a 1). Surprisingly, 3 participants entered a number that was lower than the actual outcome of the die roll. This might indicate downward lying. However, 2 of these 3 participants entered the outcome of the second die roll. Thus, it seems more plausible that they did not lie downward on purpose, but that they remembered only the second die roll.

Table 1: OLS regression with the reported payoff as the dependent variable

	Model 1	Model 2
digital die	-0.935***	-1.040***
	(0.306)	(0.308)
money		-0.104
		(0.083)
session		-0.099
		(0.107)
age		0.043
		(0.032)
male		0.280
		(0.318)
_cons	3.304***	4.877***
	(0.245)	(1.843)
N	140	140
R^2	0.065	0.102

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

As the results show, only the treatment dummy has a significant effect; that is, being in the digital die treatment compared to the control treatment reduces the payoff by € 1.04. The amount of money gained in the prior experiment, the session, age, and being male do not have a significant influence on the reported payoff.

4 Discussion and Conclusion

In this paper, we consider reputational concerns and lying aversion. In order to do so, we take away part of the ambiguity present in the observable treatments introduced by Gneezy et al. (2018) and Abeler et al. (2019): In the control treat-

ment, participants privately roll a six-sided die and report the outcome, which the experimenter cannot observe. In the digital die treatment, the die roll is executed on the computer, meaning that the experimenter can verify the outcome ex post. Importantly and contrary to the above two mentioned papers, we make it common knowledge to participants that the outcome of the random draw will be stored, but not checked for the payment procedure. We find that partial lying and full lying disappear when the experimenter is able to track participants' behavior in the digital die treatment. We suggest that this change in behavior can be explained by reputational concerns toward others, in this case toward the experimenter.

What could these reputational costs involve? First, participants could *intrinsically* dislike to be perceived as liars, and thus, abstain from misreporting. This is in line with studies on social image showing that people are more likely to behave according to a social norm when an audience is present. Second, participants could doubt that the outcome of the digital die is only electronically saved, but not checked before payment. Thus, they could be afraid of immediate real consequences in the sense that they forgo their payment or are confronted with their behavior when a potential lie is detected. Third, participants could be afraid of long-term consequences, such as being excluded from future experiments, when the experimenter detects misreporting ex post. At this point, we cannot disentangle the different explanations that we subsume under the term *reputational costs*. For further research, it would be interesting to employ a double-blind payment procedure to exclude that participants fear real consequences. In addition, a questionnaire after the experiment could help to learn more about participants' beliefs

regarding the consequences of being observed and their motives for the aversion to lie.

Comparing our results to the current experimental literature, we find similar patterns regarding partial lying. Compared to Gneezy et al. (2018) and Abeler et al. (2019), we also find that partial lying disappears when the experimenter is able to observe the outcome of the random draw. In contrast to their results, we additionaly find that full lying disappears when ambiguity regarding the storage of the data is removed.

Importantly, our results are in line with predictions arising from the model of Dufwenberg and Dufwenberg (2018). They show that their model produces very similar results to the data presented in Fischbacher and Föllmi-Heusi (2013) for a value of $\theta = 3$, capturing how strongly an individual cares about the beliefs others have about the extent of lying. Given the similarity of our data in the control treatment to Fischbacher and Föllmi-Heusi (2013), we can also assume a value of approx. $\theta = 3$ for our subject pool. For this value (more exactly, for values of $\theta > 1$), the model by Dufwenberg and Dufwenberg (2018) predicts our results with everybody telling the truth in an environment where the experimenter can observe the truth and where all subjects are aware of this. This is an important finding that shows the predictive power of the model by Dufwenberg and Dufwenberg (2018), because the design of the previous experiments did not allow to test this special case, as some ambiguity about the observability of the truth remained. Our design has the nice feature to exactly capture this case and yields that participants behave as the model predicts.

We contribute to the existing literature in several ways. First, we replicate the baseline treatment by Fischbacher and Föllmi-Heusi (2013), and find support for the three types of behavior identified in their study: There are participants who report truthfully, others who report the highest payoff, and those who lie partially. Second, we complement the experimental findings recently published by Gneezy et al. (2018) and Abeler et al. (2019) by taking away part of the ambiguity present in their treatments where outcomes are observable. Third, we test an important prediction from the model of Dufwenberg and Dufwenberg (2018) and find data in accordance with this prediction.

This study has several limitations: The subject pool consisted of only students. For further research, it would be interesting to run similar experiments with a mixed subject pool to verify the results. Abeler et al. (2014) investigated lying costs by calling a representative sample of the German population at home, and found that aggregate reporting behavior was close to the expected truthful distribution. Thus, investigating lying behavior in the laboratory might overestimate dishonesty, and calls for experiments conducted in an applied and representative context. Another drawback might be the task itself, as the die roll is artificial and not close to what people encounter in their daily life when it comes to the decision whether to lie or not. It would be interesting to study situations that are closer to the daily dilemmas people face. Next, we let participants make a one shot decision, and thus, could not study how behavior evolves over time. It would be interesting to implement a digital die treatment with repeated decisions to examine whether reputational concerns decline over time. Fischbacher and Föllmi-Heusi (2013) ran

a repeated version of their baseline die roll treatment, and found that participants reported higher payoffs when they participated a second time. This finding might suggest that reputational concerns extenuate over time. Finally, we focused on the observability by the experimenter, and thus, could investigate only reputational concerns toward him or her. In real life, people encounter observability by many different groups of people, such as family members, neighbors, or colleagues at work. It would be interesting to study the influence of the *type* of observer on lying behavior, and see whether reputational concerns differ depending on who is observing.

What are the policy implications? The results of the control treatment with unobservable outcomes generally seem promising, as they support the finding that people have an aversion to lie, and do not cheat as much as they can, even if they could do so without being detected. However, relying only on voluntary honesty might not be sufficient, as partial lying might still be a problem and full lying prevails. Based on the treatment with observable outcomes, we suggest taking into account that people care about how they are viewed by others. In the age of digitization, behavior can increasingly be observed and tracked, and people leave digital marks. By making this potential observability more salient, one could take advantage of the power of reputational costs to keep people from engaging in dishonest behavior. At the same time, being observed by humans still seems more powerful than being observed by machines: In a laboratory experiment, Cohn et al. (2018) find that participants cheat significantly more when interacting with a machine than with a person, suggesting that human interaction is key

to mitigating dishonesty. Thus, as our digital die treatment shows, making it unambigiously clear that a human being is able to check behavior as the outcome of the random draw is electronically stored, seems to be a promising road to mitigate lying behavior.

References

Abeler, J., Becker, A., and Falk, A. (2014). Representative evidence on lying costs. *Journal of Public Economics*, 113:96–104.

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

Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, 77(5):1607–1636.

Ariely, D., Bracha, A., and Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1):544–555.

Battigalli, P. and Dufwenberg, M. (2009). Dynamic psychological games. *Journal of Economic Theory*, 144(1):1–35.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217.

Bursztyn, L. and Jensen, R. (2017). Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure. *Annual Review of Economics*, 9:131–153.

Cohn, A., Gesche, T., and Maréchal, M. A. (2018). Honesty in the digital age. *CESifo Working Paper Series No. 6996*.

DellaVigna, S., List, J. A., Malmendier, U., and Rao, G. (2016). Voting to tell others. *The Review of Economic Studies*, 84(1):143–181.

Dieckmann, A., Grimm, V., Unfried, M., Utikal, V., and Valmasoni, L. (2016). On trust in honesty and volunteering among europeans: Cross-country evidence on perceptions and behavior. *European Economic Review*, 90:225–253.

Dufwenberg, M. and Dufwenberg, M. A. (2018). Lies in disguise–a theoretical analysis of cheating. *Journal of Economic Theory*, 175:248–264.

Evans III, J. H., Hannan, R. L., Krishnan, R., and Moser, D. V. (2001). Honesty in managerial reporting. *The Accounting Review*, 76(4):537–559.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

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.

Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, 1(1):60–79.

Gerlitz, J.-Y. and Schupp, J. (2005). Zur Erhebung der Big-Five-basierten Persönlichkeitsmerkmale im SOEP. *DIW Research Notes*, 4.

Gibson, R., Tanner, C., and Wagner, A. F. (2013). Preferences for truthfulness: Heterogeneity among and within individuals. *American Economic Review*, 103(1):532–548.

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.

Hahn, E., Gottschling, J., and Spinath, F. M. (2012). Short measurements of personality–validity and reliability of the GSOEP Big Five Inventory (BFI-S). *Journal of Research in Personality*, 46(3):355–359.

Jacobsen, C., Fosgaard, T. R., and Pascual-Ezama, D. (2018). Why do we lie? A practical guide to the dishonesty literature. *Journal of Economic Surveys*, 32(2):357–387.

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies*, 76(4):1359–1395.

Khalmetski, K. and Sliwka, D. (2019). Disguising lies - image concerns and partial lying in cheating games. *American Economic Journal: Microeconomics (forthcoming)*.

Kocher, M. G., Schudy, S., and Spantig, L. (2017). I lie? We lie! Why? Experimental evidence on a dishonesty shift in groups. *Management Science*, 64(9):3995–4008.

Kosfeld, M. and Neckermann, S. (2011). Getting more work for nothing? Symbolic awards and worker performance. *American Economic Journal: Microeconomics*, 3(3):86–99.

Kurzban, R., DeScioli, P., and O'Brien, E. (2007). Audience effects on moralistic punishment. *Evolution and Human Behavior*, 28(2):75–84.

López-Pérez, R. and Spiegelman, E. (2013). Why do people tell the truth? Experimental evidence for pure lie aversion. *Experimental Economics*, 16(3):233–247.

Pruckner, G. J. and Sausgruber, R. (2013). Honesty on the streets: A field study on newspaper purchasing. *Journal of the European Economic Association*, 11(3):661–679.

Rauhut, H. (2013). Beliefs about lying and spreading of dishonesty: Undetected lies and their constructive and destructive social dynamics in dice experiments. *PloS One*, 8(11):1–8.

Weibull, J. W. and Villa, E. (2005). Crime, punishment and social norms. *SSE/EFI Working Paper Series in Economics and Finance No. 610*.

West, M. D. (2005). *Law in everyday Japan: Sex, Sumo, Suicide, and Statutes.* University of Chicago Press.

Appendix A: Instructions

Questionnaire

I see myself as someone who

1. ... tends to be lazy.

2. ... values aesthetic experiences.

3. ... tends to be disorganized.

4. ... is reserved, quiet.[11]

Answers given on a 7-point Likert scale: 1 = totally disagree, ..., 7 = totally agree

Die roll game

Control treatment

For the questionnaire, you will receive a small additional payoff. However, this payoff is not the same for every participant. You determine your payoff by rolling the die twice as soon as you are asked to. The first roll determines how much you receive. If you roll a 1, 2, 3, 4, or 5, you receive one euro per number rolled (i.e. if you roll a 1, you receive € 1 if you roll a 2, you receive € 2, etc.). If you roll a 6, you receive nothing. The second die roll only serves to make sure that the die is working properly. If you wish, you can roll the die more than twice. However, only the first die roll counts.

If you have any questions, please raise your hand.

If you are ready, please press OK.

---

[11]Questions are taken from the German Socio Economic Panel (GSOEP) Big Five Inventory (see e.g. Gerlitz and Schupp 2005, Hahn et al. 2012).

Digital die treatment

For the questionnaire, you will receive a small additional payoff. However, this payoff is not the same for every participant. You determine your payoff by rolling the digital die on your computer twice as soon you are asked to. Therefore, you click the button "Rolling die" and will see the result of your die roll. The result of your die roll is electronically stored, but will not be verified for the payment procedure. The first roll determines how much you receive. If you roll a 1, 2, 3, 4, or 5, you receive one euro per number rolled (i.e. if you roll a 1, you receive € 1, if you roll a 2, you receive € 2, etc.). If you roll a 6, you receive nothing. The second die roll only serves to make sure that the die is working properly. If you wish, you can roll the die more than twice. However, only the first die roll counts.

If you have any questions, please raise your hand.

If you are ready, please press OK.

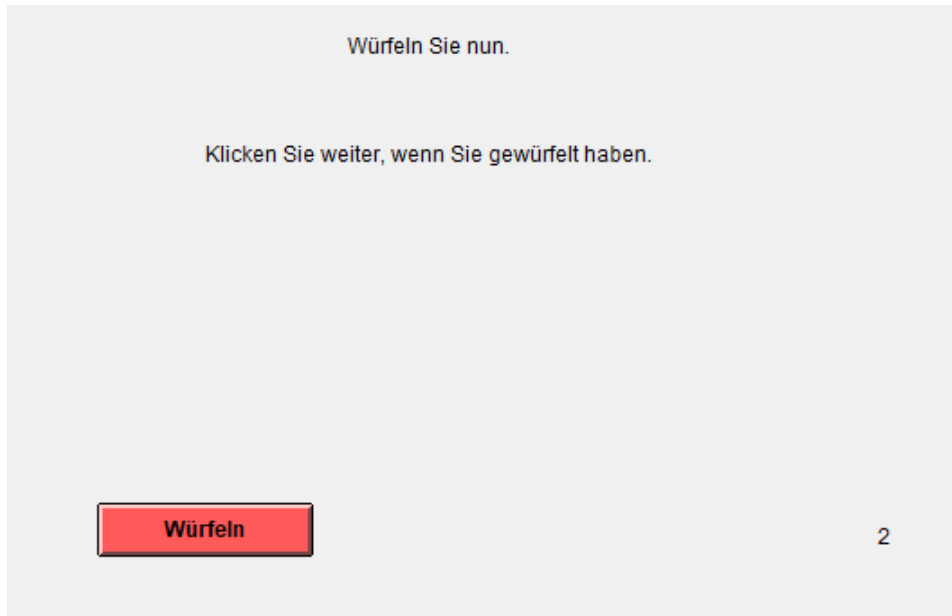# Appendix B: Screenshot digital die treatment

Figure 3: Implementation of the digital die in z-Tree with the button "Rolling die" on the left side and the outcome of the die roll on the right side. Participants could click on the button several times to check that the digital die works properly.

Appendix C: Comparison of treatments

Table 2: Shares of participants in percentage reporting the corresponding payoff in the control treatment and the digital die treatment. A Fisher exact test comparing the two distributions reveals a statistically significant difference (p=0.004).

	0	1	2	3	4	5	Fisher exact test
Control treatment (n=56)	16.1	5.4	7.1	10.7	25	35.7	p=0.004
Digital die treatment (n=84)	16.7	20.2	15.5	19	14.3	14.3	