ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 111

On the Persistence of Dishonesty

Stefania Bortolotti Felix Kölle Lukas Wenner

August 2021 www.econtribute.de

UNIVERSITÄT BONN Universität zu Köln

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866 is gratefully acknowledged.

On the Persistence of Dishonesty*

Stefania Bortolotti Felix Kölle Lukas Wenner

August 18, 2021

Abstract

In social and economic interactions, individuals often exploit informational asymmetries and behave dishonestly to pursue private ends. In many of these situations the costs and benefits from dishonest behavior do not accrue immediately and at the same time. In this paper, we experimentally investigate the role of time on dishonesty. Contrary to our predictions, we find that neither delaying the gains from cheating, nor increasing temporal engagement with one's own unethical behavior reduces the likelihood of cheating. Furthermore, allowing for a delay between the time when private information is obtained and when it is reported does not affect cheating in our experiment.

Keywords: Dishonesty, cheating, delay, discounting, experiment

JEL Classification Numbers: C91; D82; D91

*Corresponding author: Kölle, Department of Economics, University of Cologne, Albertus Magnus Platz, 50923 Cologne, Germany, e-mail: *felix.koelle@uni-koeln.de*. Bortolotti: Economics Department, University of Bologna & IZA, e-mail: *stefania.bortolotti@unibo.it*. Wenner: Department of Economics, University of Cologne, e-mail: *lukas.wenner@uni-koeln.de*. The authors thank Susanna Grundmann, Agne Kajackaite, and participants at various conferences for helpful comments and suggestions. We gratefully acknowledge the financial support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866. Significant parts of this research were developed while Bortolotti was at the Max Planck Institute for Research on Collective Goods, Bonn; she thanks this institutions for the financial and logistic support. Declarations of interest: none.

1 Introduction

Many economic interactions involve the reporting of private information to other agents. Often, there is a trade-off between truthful reporting, which may be morally desirable and/or socially optimal, and lying, which may be individually optimal (e.g., payoff-maximizing). Examples of such situations reach from making tax declarations or insurance claims to experts giving advice to their clients. A recent literature in economics has investigated a number of factors that influence (dis)honesty, including the possible gains from lying, the degree of social observability, and the ex-ante likelihood of different states of the world (see Abeler et al., 2019, for a an overview).

A common feature of the existing evidence is that these studies primarily focus on cases in which the consequences from lying occur immediately. For example, in the prominent "die-rolling task" (Fischbacher and Föllmi-Heusi, 2013), participants privately observe the outcome of a random variable, report the outcome, and receive a monetary payoff based on their report typically right at the end of the experiment. While attractive for its simplicity, such a design neglects that in many real-world situations in which individuals have to decide whether to engage in dishonest behavior or not, the costs and benefits of the agents' actions do not necessarily realize immediately and at the same time. Given the ample evidence of time discounting across a variety of domains such as education, health, or financial decision-making (see Frederick et al., 2002; Chabris et al., 2010, for overviews), this cast some doubts on the generalizability of previous results for cheating behavior in the field.

In this paper, we investigate how the timing of the reporting of private information and the timing of the realization of the benefits from lying affect dishonesty. We focus on three aspects. First, we consider the case in which potential psychological costs of lying realize immediately, but the benefits from such dishonest behavior are delayed. Misreporting income to tax authorities is a prototypical example as the intrinsic costs of such dishonest behavior occur immediately (i.e., at the time of declaring one's income), while the benefits, e.g., in form of a lower tax burden, are realized only in the future.[1] Second, we test if increased mental engagement with and temporal exposure to one's own dishonest behavior decrease dishonesty. When the benefits of a lie are accrued in the future, it is also likely that there will be reminders of the unethical behavior or the need to confirm the lie told in the first place. For example, employees misreporting their achievements in hope of a promotion can

[1] Of course, there might be other costs of dishonesty that also only accrue with delay. For example, the loss of one's reputation of being honest, or potential fines or penalties from court might also only occur in the future, once the dishonest behavior has been detected. For our experiment, to keep the setting focused on our question of interest, we abstract away from such effects.

be reminded of their unethical behavior every single day by personal notes or colleagues. Third, time might matter for dishonesty because in many situations the time at which private information are received and the time when they are reported is not the same. For example, when experts give advice to clients or when interacting in long-term relationships (either professionally or privately), the informed party typically has time to reflect upon the decision to be dishonest or not.

To study these three aspects, we designed an experiment in which participants have to flip, in private, a coin twice. Participants are told to truthfully report the outcome of the two coin flips, and that when they report *2 Heads* they receive an amount which is eight times higher than for any of the other two possible reports. To investigate the direct effect of a delayed payment, participants either receive their payment from the coin flip immediately thereafter or with a delay of one week. We implement an increased mental engagement by comparing a treatment where participants get paid for their report by simply returning to the experiment after a week, to a treatment where they have to return and correctly recall their previous report. Finally, to study how a delay in reporting affects dishonesty, we compare two reports made at the same point in time, but vary whether the underlying private information was generated one week or a few minutes before the reporting took place.

Overall, we find that 68% of the participants report having flipped *2 Heads*, which clearly exceeds the expected reporting frequency of 25% under full honesty. More strikingly, we find little evidence for differences in dishonesty levels across treatments. We find that neither delaying the gains from cheating, nor increasing the engagement with one's own unethical behavior reduces the likelihood of cheating. At the same time, we find that when providing individuals with some time to think about their decision and potentially forget about the outcome of their coin flip, this does not further increase dishonesty.

Our paper contributes to the extensive literature investigating the motives and circumstances of dishonest behavior (Gneezy, 2005; Mazar et al., 2008; Shalvi et al., 2011; Erat and Gneezy, 2012; Fischbacher and Föllmi-Heusi, 2013; Conrads et al., 2013; Gneezy et al., 2013; Abeler et al., 2014; Kajackaite and Gneezy, 2017; Gneezy et al., 2018; Dai et al., 2018; Abeler et al., 2019; Charness et al., 2019; Dimant et al., 2020). In particular, we add to the scant evidence on the role of time on dishonest behavior. Most closely related to our paper is a recent study by Andersen et al. (2018) who investigate the effects of cooling-off periods on dishonesty. They find that giving participants one day to think about their decision to misreport private information had no effect on their degree of dishonesty. Here, we find a similar effect for a period of one week using a different task and subject pool. Unlike us, they do not study the effects of delayed benefits or the effect of increased mental engage-

ment over time. Somewhat related is the literature on the role time pressure and intuition on dishonesty. While the results are not clear cut, evidence seems to suggest that dishonesty is intuitive and that honesty requires deliberation and self-control (Shalvi et al., 2012; Köbis et al., 2019).

The remainder of the paper is organized as follows. In Section 2, we explain the design and procedures of our experiment, and derive our hypotheses. In Section 3 we present our main findings, and Section 4 concludes.

2 Experimental Design

2.1 General Setup

The experiment consists of two separate sessions conducted on Amazon Mechanical Turk (MTurk) – session 1 and session 2 – which took place exactly one week apart from each other. Each session comprised up to two parts:

- *Part 1 – Rating Task:* Participants are presented with 20 pairs of paintings, and for each pair they have to indicate their preferred one. We emphasized that we are interested in their subjective opinion and that there is no right or wrong answer (see Appendix B for a copy of the instructions). The purpose of this task is to conceal the main aim of the study and to make sure that in all treatments (see below) there is a meaningful task in both weeks;

- *Part 2 – Coin Task:* Participants are asked to flip a coin twice in private, and to report the outcome of the coin flips. This report then determines the bonus earned by the participant; if they report *2 Heads*, they receive a bonus payment of \$2.00, whereas any other report (*1 Heads and 1 Tails* or *2 Tails*) yields a bonus of \$0.25. While it is clearly stated that participants have to report the outcome of the coin flips truthfully, there is no monitoring and participants may choose to misreport the outcome in order to achieve a higher bonus.

2.2 Treatments and hypotheses

We implemented six between-subjects treatments in which we only varied the exact implementation of the coin task. More specifically, we varied when the bonus is paid, the degree to which participants are reminded of their potentially untruthful report, and the time when

Table 1: Experimental treatments

Treatments	**Session 1**		**Session 2**	
	Rating Task	Coin Task	Rating Task	Coin Task
Immediate	✓	flip, report, bonus	✓	—
PayDelay	✓	flip, report, bonus (delayed)	✓	—
PayDelay+ShowUp	✓	flip, report	✓	bonus
PayDelay+Recall	✓	flip, report	✓	recall, bonus
ImmediateWeek2	✓	—	✓	flip, report, bonus
ReportDelay	✓	flip	✓	report, bonus

the outcome is reported (see Table 1 for an overview). In the following, we explain each of the six treatments in detail and link each pair of treatments to one of our three research questions. We also provide a description of our hypotheses regarding the effects of time on dishonesty. All our hypotheses are based on recent models of dishonesty (Gneezy et al., 2018; Abeler et al., 2019; Khalmetski and Sliwka, 2019), which assume that an agent's behavior is determined by three components: (i) the monetary payoff; (ii) a direct intrinsic cost of lying; (iii) and social image concerns, i.e., the motivation to appear honest. All our treatments and hypotheses were pre-registered and all our main analysis were pre-specified (our pre-analysis plan can be found here: `https://www.socialscienceregistry.org/trials/3594`).

To answer our first research question of whether the timing of the benefits affects dishonesty, we implemented the following two treatments:

- *Immediate*: Participants report the outcome of the coin flips in week 1, and receive their corresponding bonus in week 1 (shortly after completing the session).
- *PayDelay*: Participants report the outcome of the coin flips in week 1, and receive their corresponding bonus with a delay of one week.

The two treatments are exactly identical, apart from the date at which the bonus from the coin task is paid. In both treatments, participation in session 2 – which comprised only the rating task – was not mandatory to receive the bonus from the coin task. Under the assumption that agents face a fixed intrinsic cost of lying, and that this cost realizes at the time of the untruthful report, delaying the bonus to the future should decrease dishonesty.

The reason is that while the intrinsic costs of lying are the same across treatments (in both cases they occur in the present), the gains from lying are smaller in *PayDelay* as agents discount future payments.[2] This leads to our first hypothesis:

Hypothesis 1 (**Delayed Payments**). *Delaying the benefits of dishonest behavior reduces misreporting: the share of 2 Heads is higher in Immediate than in PayDelay.*

Following our pre-analysis plan, we test this hypothesis by comparing the proportion of reports of *2 Heads* in session 1 between *Immediate* and *PayDelay*.

To answer our second research question of whether increased mental engagement with one's own actions affects dishonesty, we ran the following two treatments:

- *PayDelay+ShowUp*: Participants report the outcome of the coin flips in week 1, and receive their corresponding bonus in week 2. They receive this payment only if they take part in the second session.
- *PayDelay+Recall*: Participants report the outcome of the coin flips in week 1, and receive their corresponding bonus in week 2. They receive this payment only if they take part in the second session and correctly recall their report from week 1. If their reports from both weeks do not match, they receive no bonus.

The two treatments differ in the degree to which participants need to engage with their (potentially untruthful) report during the seven days between the two sessions. While in both cases participants need to complete the second session to receive the bonus payment from the coin task, in *PayDelay+ShowUp* they no longer need to engage with their report once they completed the first session as they can secure the bonus even if they forget about their report. In the *PayDelay+Recall* treatment, in contrast, in week 2 participants are required to correctly recall what they reported in the first session; if their are not able to recall correctly their initial report, they will not receive any bonus. This implies that all participants untruthfully reporting 2 Heads in the first week have to reiterate their false claim a second time. This also implies that agents need to remember their first report for an entire week, which might require them taking note of the result, hence enhancing their level

[2]One critical issue when studying intertemporal choices is to make all choices equivalent except for their timing. In particular, transaction costs associated with experimental payments must be constant across all treatments. This applies not only to physical transaction costs but also to the credibility of payments. In our setting, we achieve this by relying on the MTurk payment systems, which allows us to use exactly the same method for both immediate and delayed bonuses. Furthermore, given that participation in our experiment was restricted to individuals with a sufficient high overall approval rate (see below), all our participants had plenty of experience with this payment system.

of engagement. An agent anticipating that reporting dishonestly twice and being reminded of that comes at an increased cost (e.g., via reduced self-image), may decide to rather report honestly in the first place instead. This should result in a lower proportion of participants claiming the large bonus (*2 Heads*) in *PayDelay+Recall* than in *PayDelay+ShowUp*. This leads to our second hypothesis:

Hypothesis 2 (**Engagement Costs**)**.** *Increasing the extent to which participants have to engage with their dishonest behavior reduces dishonesty: the share of 2 Heads is higher in PayDelay+ShowUp than in PayDelay+Recall.*

As specified in the pre-analysis plan, to test our second hypothesis we rely on answers given in the first week, and we include in the analysis even those participants who did not show up in the second session. The reason is that attrition between the two weeks might be driven by different reasons across the two treatments. Specifically, participants in *PayDelay+Recall* may not return simply because they forgot what they reported. If forgetting is not anticipated but correlates with the reported outcome in session 1, then using the session 2 data would bias our results.

To answer our third research question of whether a delay between the time when private information is obtained and the time when it has to be reported affects dishonesty, we conducted the following two treatments:

- *ImmediateWeek2*: Participants are instructed about the coin task in week 2, when they have to flip the coin and report the result. The corresponding bonus is paid in week 2 (shortly after completing the session).
- *ReportDelay*: Participants are instructed to flip the coin in week 1, but are asked to report the outcome only in week 2. They receive the corresponding bonus in week 2 (shortly after completing the session).

While in *ImmediateWeek2* the coin task only takes place in week 2 (instructions, coin flips, and the report all happen in session 2), in *ReportDelay* the coins are flipped in week 1 but the outcome is reported only in week 2. In the latter case, agents may find it "convenient" to forget the outcome of the coin flip: in fact, if they are unable to recall the exact outcome, reporting *2 Heads* might come at a lower cost as they can maintain their self-image of being honest. This reasoning is in line with recent evidence by Zimmermann (2020) on the importance of motivated beliefs and by Exley and Kessler (2019) on motivated errors in decision making. Furthermore, failing to remember the outcome of the two coin flips might

serve as a justified excuse to act more unethically, as argued by Shalvi et al. (2012). Another potential channel through which misreporting could be affected across these two treatments is that while *ImmediateWeek2* participants need to report the outcome immediately after observing it, in *ReportDelay* participants have time to "cool down" and think about their decision. Previous literature on cooling down effects typically have found that agents become more rational and self-interested after some time has passed (see e.g., Xiao and Houser, 2005; Grimm and Mengel, 2011; Bolle et al., 2014; Dickinson and Masclet, 2015); even though a recent study by Andersen et al. (2018) shows that providing participants with an additional day to think about their decision did not affect the extent of cheating. If motivated beliefs/errors and cooling off are at work, one should expect the proportion of *2 Heads* to be higher in *ReportDelay* than in *ImmediateWeek2*. This leads to our third hypothesis:

Hypothesis 3 (**Motivated Recall**). *Providing individuals with an excuse to cheat by delaying the time of their report increases dishonesty: the share of 2 Heads is higher in ReportDelay than in ImmediateWeek2.*

Unlike for our tests of the previous hypotheses, here, again in accordance with our pre-analysis plan, we rely on data from participants who completed the second session, as this is the time when they report the outcome of the coin flips.

2.3 Procedures

The experiment was conducted online and participants were recruited via the online labor market platform MTurk (see Horton et al., 2011, for a detailed description). Participation was restricted to workers residing in the US and with a high approval rate.[3] We recruited a total of $n = 1235$ workers, and randomly assigned them to one of the six treatments (between-subjects design). The sample size was determined a priori using power calculations. Our variable of interest is the proportion of participants who report *2 Heads*. Based on a z-test for the difference of independent proportions, power calculations reveal that for a sample of $n = 200$ observations per treatment, we have a power of 80% to reject the equality of proportions hypothesis if the normalized effect size h is above 0.28, which is a small to medium effect size. Depending on the baseline proportion of reports of *2 Heads*, this corresponds to a treatment difference of 0.13 to 0.15.

We deliberately oversampled treatments *ReportDelay* and *ImmediateWeek2* because for

[3]Collecting data on Mturk has become increasingly popular in the social sciences, and recent studies have shown that behavior in a variety of games is consistent between representative samples on MTurk and laboratory participants (Arechar et al., 2018; Snowberg and Yariv, 2021).

those two treatments we have to rely on the data from participants who complete both sessions (and not just the first one as in the other treatments). To get an idea about the expected attrition rate between the two weeks, prior to our main experiment, we conducted a pilot study. The attrition rate in the pilot experiment was 10%, which is why we invited twenty additional participants in treatments *ReportDelay* and *ImmediateWeek2*.[4] Our final sample consists of $n = 1166$ individuals: $n = 198$ in *Immediate*, $n = 204$ in *PayDelay*, $n = 199$ in *PayDelay+ShowUp*, $n = 189$ in *PayDelay+Recall*, $n = 194$ in *ImmediateWeek2*, and $n = 182$ in *ReportDelay*. The attrition rate across weeks was fairly low: 14.5%, with no differences across treatments ($\chi^2(5) = 7.96, p = 0.159$).

Upon accepting our invitation for the first session, participants received a link that directed them to our experimental interface, which was programmed using Qualtrics. Participants then received instructions about the rules and nature of our study. After that, they were introduced to the rating task. Upon completion of this task, participants received instructions about the coin flip task; to ensure understanding of the incentives and the procedures of the coin flip task, all participants had to answer a set of control questions. Only participants who answered all control questions correctly (within two attempts) were allowed to participate in the study. Seven days after completion of the first session, we sent all participants a reminder via email, inviting them to participate in the second session. After receiving this email, participants had 24 hours to complete the study.

Final earnings are the sum of (i) a flat session payment, (ii) a variable bonus payment, and (iii) a flat completion payment. For the successful completion of each of the two sessions, participants received a flat payment of \$0.50, paid to them directly after completing the respective session. On top of that, depending on the reported outcome from the coin flip task, participants received a bonus payment of either \$2.00 (in case they reported *2 Heads*) or \$0.25 (in case they reported *1 Heads and 1 Tails* or *2 Tails*). To limit attrition across weeks, participants further received a completion payment of \$1.50 if they successfully completed both sessions. In order to not interfere with the incentives of the coin flip task, the completion payment was paid out three days after the second session (see Table A1 in Appendix A for an overview of the exact timeline). On average, participants earned a total of \$3.83, for a study that lasted around eight minutes (both sessions combined).

[4]The design of the pilot study was very similar to the design of our main study, except for the fact that participants did not have the possibility to behave dishonestly. Hence, no data on lying behavior was obtained.

3 Results

We start the discussion of our results by reporting the overall degree of dishonesty in our sample. If all participants would report the outcome of the coin task truthfully, the expected frequency of *2 Heads* would be equal to 25%. In stark contrast to that, we find that the fraction of participants reporting the high outcome is equal to 67.7%, which is significantly higher than the expected 25% under truthful reporting (Binomial test, $p < 0.001$). Under the assumption that nobody misreports to their disadvantage, we find about 43% of our participants to be dishonest – slightly more then reported in a recent meta-analysis of previous studies (Abeler et al., 2019), but close to the results by Garbarino et al. (2019) who us a task and subject pool similar to our study.

Next, we test each of our three hypotheses by comparing the fraction of participants reporting *2 Heads* across treatments. We start with the test of Hypothesis 1, stating that the number of *2 Heads* should be lower when the material benefits from misreporting are paid out later rather than immediately. Our evidence does not support this hypothesis. As displayed by the left panel of Figure 1, 67.7% of the participants in *Immediate* and 72.5% in *PayDelay* report *2 Heads*. If anything, dishonesty seems to be even higher under delayed payments, although the difference between treatments is not statistically significant ($\chi^2(1) = 1.14, p = 0.286$). This is further confirmed by logistic regressions (see Table A2 in Appendix A). We summarize this finding in our first result:

> **Result 1 (Delayed Payments):** *Delaying the material gains from dishonest behavior does not reduce dishonesty.*

A possible explanation for our finding is that participants in our experiment simply do not discount payments that are delayed by one week. To test for this possibility, we conducted a follow-up study with $n = 201$ new participants, recruited from the same subject pool on MTurk. Similarly to our main experiment, participants had to first complete a Rating Task. Afterwards, instead of reporting the outcomes of the coin flips, participants were asked to complete a standard price list design to elicit their time preferences (Coller and Williams, 1999; Harrison et al., 2002; Dohmen et al., 2010). Specifically, participants were asked to choose between thirteen smaller-sooner and larger-later rewards. The larger-later reward was a fixed payment of \$2.00, while the smaller-sooner reward decreased in increments of \$0.05 from \$2.10 in the first row to \$1.50 in the last row (see Figure A1 in Appendix A).

The results from this follow-up study are shown in Figure A2 in Appendix A, depicting the distribution of \$-amounts participants are willing to forgo to receive money today rather

Figure 1: Fraction of *2 Heads* by treatment.

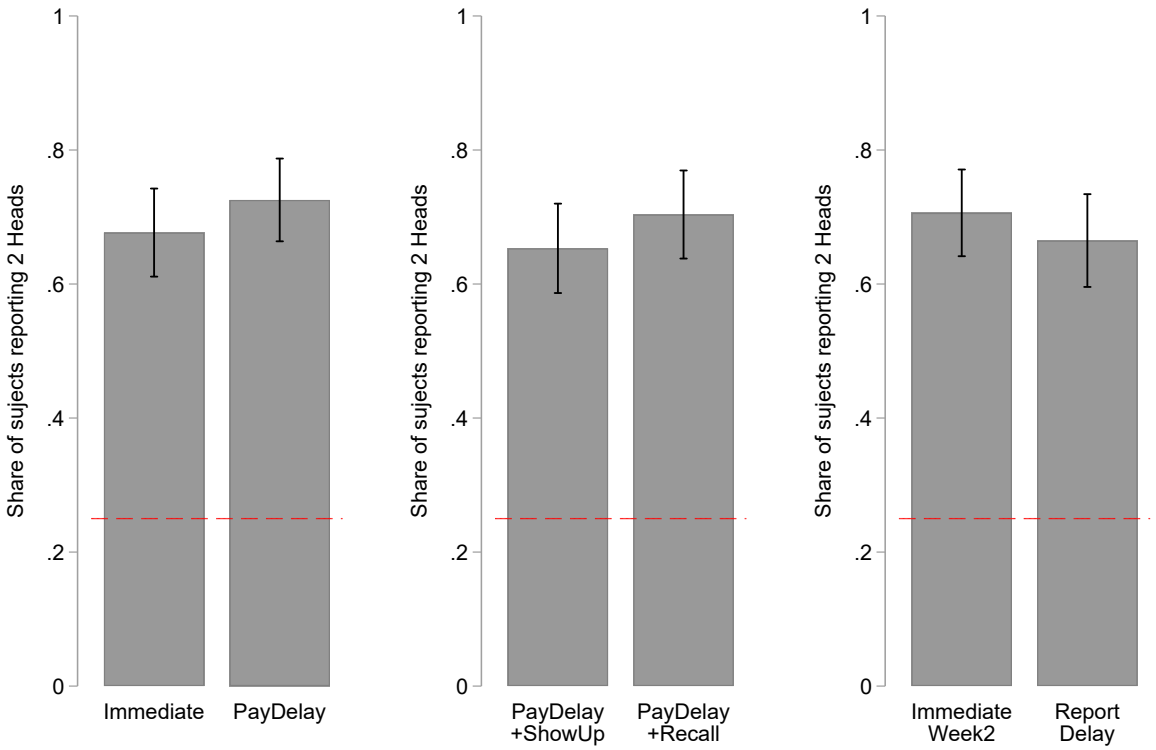

Notes. Solid lines indicate 95% confidence intervals. Dashed lines indicate the expected frequency if everyone would report truthfully.

than \$2.00 in one week. The results reveal that while around 43% of participants always choose the larger monetary amount irrespective of the timing, the majority of participants display some degree of discounting. Among the latter, on average participants are willing to forgo \$0.32 (or 16%) in order to advance the payment by one week. Furthermore, around one quarter of our participants (24%) are willing to forgo at least \$0.50 (or 25%), indicating a substantial degree of impatience. Overall, these results reveal that the fact that we do not observe a lower rate of dishonesty in *PayDelay* compared to *Immediate* is not due to low time discounting over the period of time considered here (one week).

Next, we turn to the comparison of *PayDelay+ShowUp* and *PayDelay+Recall*, to test if forcing participants to recall and reiterate their initial report reduces dishonesty. We do not find support for this hypothesis. As shown in the middle panel of Figure 1, 65.3% of participants report *2 Heads* in *PayDelay+ShowUp*, compared to 70.4% in *PayDelay+Recall*; the difference between treatments is not significant ($\chi^2(1) = 1.13, p = 0.288$; see also column (2) in Table A2 in Appendix A). Results are quantitatively and qualitatively similar if we

consider the subsample of participants who completed both sessions. For this subsample, the fraction of *2 Heads* is 64.0% in *PayDelay+ShowUp* and 69.8% in *PayDelay+Recall*, and the difference is not statistically significant ($\chi^2(1) = 1.31, p = 0.252$). This is reassuring, as the lack of significant differences thus cannot be driven by participants willingly not engaging with their lie. As for the recall rate, only 3.5% (6 out of 172) of the participants in our *PayDelay+Recall* failed to correctly recall their initial report. We do not find any systematic relationship between the initial report and the recall rate.[5] While we do not have any information about how and by which means participants managed to recall their initial report (e.g., by taking a photo or screenshot, writing it down, keeping it in mind), the fact that 96.5% of participants correctly recalled their report clearly suggests that they engaged with it. We summarize these findings in our second result:

> **Result 2 (Engagement Costs):** *Increasing the extent to which participants have to engage with their dishonest behavior does not reduce dishonesty.*

Finally, we test our third hypothesis, on motivated recall. To this end, we compare the fraction *2 Heads* in *ImmediateWeek2* and *ReportDelay*: the results are shown in the right panel of Figure 1. The fraction *2 Heads* is 70.6% in *ImmediateWeek2* and 66.5% in *ReportDelay*, and the difference is not statistically significant ($\chi^2(1) = 0.75, p = 0.388$; see also column (3) in Table A2 in Appendix A). In fact, if anything, adding a cooling-off period by delaying the report by one week led to less misreporting, although the differences are small and not significant. This constitutes our third result:

> **Result 3 (Motivated Recall):** *Allowing individuals to think about their decision and providing them with an excuse to cheat by delaying the time of their report does not increase dishonesty.*

4 Discussion & Conclusion

Given the high costs of dishonesty for society, understanding the motives underlying such behavior is of great importance. A relatively recent body of studies in economics has started to investigate the circumstances under which people engage in dishonest behavior (see

[5] Among the six participants who did not correctly recall their initial report, one reported *2 Tails* but recalled *1 Heads and 1 Tails*, one reported *1 Heads and 1 Tails* but recalled *2 Heads*, two reported *1 Heads and 1 Tails* but recalled *2 Tails*, and two reported *2 Heads* but recalled *2 Tails*. In addition, we also do not find any systematic relationship between participants' initial report and their likelihood of showing up and completing the second session ($\chi^2(1) = 0.33, p = 0.564$). The overall completion rate is thereby not significantly different from the one in *PayDelay+ShowUp* ($\chi^2(1) = 2.02, p = 0.156$).

Abeler et al., 2019, for an overview). We extend this literature by showing that, contrary to our pre-specified hypotheses, dishonesty is persistent over time: neither delaying the gains from cheating, nor increasing the engagement with one's own unethical behavior reduces the likelihood of cheating. At the same time, we show that allowing individuals to think about their decision and providing them with an excuse to cheat does not further increase dishonesty.

While the latter result is in fact good news for those concerned about the costs of dishonesty, our first two results provide a more pessimistic message. The fact that delaying the benefits of dishonesty does not, despite substantial time discounting, decrease the likelihood of individuals engaging in such dishonest behavior suggests that simple policy solutions aimed at reducing misconduct by shifting incentives to the future might not be very effective. Making dishonesty more salient by forcing individuals to recall their own misreporting behavior seemed, a priori, another useful strategy to reduce misconduct. Here, we find no evidence for the efficacy of such a policy.

One possible explanation for the absence of any delayed payments effect is that misreporting in general seems to be only marginally affected by the potential gains from lying (see Abeler et al., 2019, for a discussion, p. 1123). As such, one might argue that our manipulation of the incentives to lie might have been too weak, and that only extremely long delays might have a sizable impact on cheating. The null effect of increased engagement could instead be explained by very low costs of cheating in our setting. Specifically, given that in our experiment participants were completely anonymous and detection of misreporting was ruled out by design, the costs of cheating might have been so low that even when having to reiterate the misreport a second time (and engaging with it in between), the benefits from cheating might have still outweighed total costs. Future research should therefore test the robustness of our results, for example by investigating longer time horizons or by considering decision environments in which the costs of lying are increased, e.g., situations in which social image concerns are more prominent or where agents have to lie about their performance rather then luck (see Kajackaite, 2018).

References

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

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

Andersen, S., U. Gneezy, A. Kajackaite, and J. Marx (2018): "Allowing for reflection time does not change behavior in dictator and cheating games," *Journal of Economic Behavior & Organization*, 145, 24–33.

Arechar, A. A., S. Gächter, and L. Molleman (2018): "Conducting interactive experiments online," *Experimental Economics*, 21, 99–131.

Bolle, F., J. H. Tan, and D. J. Zizzo (2014): "Vendettas," *American Economic Journal: Microeconomics*, 6, 93–130.

Chabris, C. F., D. I. Laibson, and J. P. Schuldt (2010): "Intertemporal choice," in *Behavioural and Experimental Economics*, Springer, 168–177.

Charness, G., C. Blanco-Jimenez, L. Ezquerra, and I. Rodriguez-Lara (2019): "Cheating, incentives, and money manipulation," *Experimental Economics*, 22, 155–177.

Coller, M. and M. B. Williams (1999): "Eliciting individual discount rates," *Experimental Economics*, 2, 107–127.

Conrads, J., B. Irlenbusch, R. M. Rilke, and G. Walkowitz (2013): "Lying and team incentives," *Journal of Economic Psychology*, 34, 1–7.

Dai, Z., F. Galeotti, and M. C. Villeval (2018): "Cheating in the lab predicts fraud in the field: An experiment in public transportation," *Management Science*, 64, 1081–1100.

Dickinson, D. L. and D. Masclet (2015): "Emotion venting and punishment in public good experiments," *Journal of Public Economics*, 122, 55–67.

Dimant, E., G. A. Van Kleef, and S. Shalvi (2020): "Requiem for a nudge: Framing effects in nudging honesty," *Journal of Economic Behavior & Organization*, 172, 247–266.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010): "Are risk aversion and impatience related to cognitive ability?" *American Economic Review*, 100, 1238–60.

Erat, S. and U. Gneezy (2012): "White lies," *Management Science*, 58, 723–733.

Exley, C. L. and J. B. Kessler (2019): "Motivated errors," Tech. rep., National Bureau of Economic Research.

Fischbacher, U. and F. Föllmi-Heusi (2013): "Lies in disguise—an experimental study on cheating," *Journal of the European Economic Association*, 11, 525–547.

Frederick, S., G. F. Loewenstein, and T. O'Donoghue (2002): "Time discounting and time preference: A critical review," *Journal of Economic Literature*, 40, 351–401.

Garbarino, E., R. Slonim, and M. C. Villeval (2019): "Loss aversion and lying behavior," *Journal of Economic Behavior & Organization*, 158, 379–393.

Gneezy, U. (2005): "Deception: The role of consequences," *American Economic Review*, 95, 384–394.

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

Gneezy, U., B. Rockenbach, and M. Serra-Garcia (2013): "Measuring lying aversion," *Journal of Economic Behavior & Organization*, 93, 293–300.

Grimm, V. and F. Mengel (2011): "Let me sleep on it: Delay reduces rejection rates in ultimatum games," *Economics Letters*, 111, 113–115.

Harrison, G. W., M. I. Lau, and M. B. Williams (2002): "Estimating individual discount rates in Denmark: A field experiment," *American Economic Review*, 92, 1606–1617.

Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011): "The online laboratory: Conducting experiments in a real labor market," *Experimental Economics*, 14, 399–425.

Kajackaite, A. (2018): "Lying about luck versus lying about performance," *Journal of Economic Behavior & Organization*, 153, 194–199.

Kajackaite, A. and U. Gneezy (2017): "Incentives and cheating," *Games and Economic Behavior*, 102, 433–444.

Khalmetski, K. and D. Sliwka (2019): "Disguising lies—Image concerns and partial lying in cheating games," *American Economic Journal: Microeconomics*, 11, 79–110.

Köbis, N. C., B. Verschuere, Y. Bereby-Meyer, D. Rand, and S. Shalvi (2019): "Intuitive honesty versus dishonesty: Meta-analytic evidence," *Perspectives on Psychological Science*, 14, 778–796.

Mazar, N., O. Amir, and D. Ariely (2008): "The dishonesty of honest people: A theory of self-concept maintenance," *Journal of Marketing Research*, 45, 633–644.

Shalvi, S., J. Dana, M. J. Handgraaf, and C. K. De Dreu (2011): "Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior," *Organizational Behavior and Human Decision Processes*, 115, 181–190.

Shalvi, S., O. Eldar, and Y. Bereby-Meyer (2012): "Honesty requires time (and lack of justifications)," *Psychological Science*, 23, 1264–1270.

Snowberg, E. and L. Yariv (2021): "Testing the Waters: Behavior across Participant Pools," *American Economic Review*, 111, 687–719.

Xiao, E. and D. Houser (2005): "Emotion expression in human punishment behavior," *Proceedings of the National Academy of Sciences*, 102, 7398–7401.

Zimmermann, F. (2020): "The dynamics of motivated beliefs," *American Economic Review*, 110, 337–61.

A Additional Figures and Tables

Figure A1: Screenshot of decision screen for the elicitation of time preferences

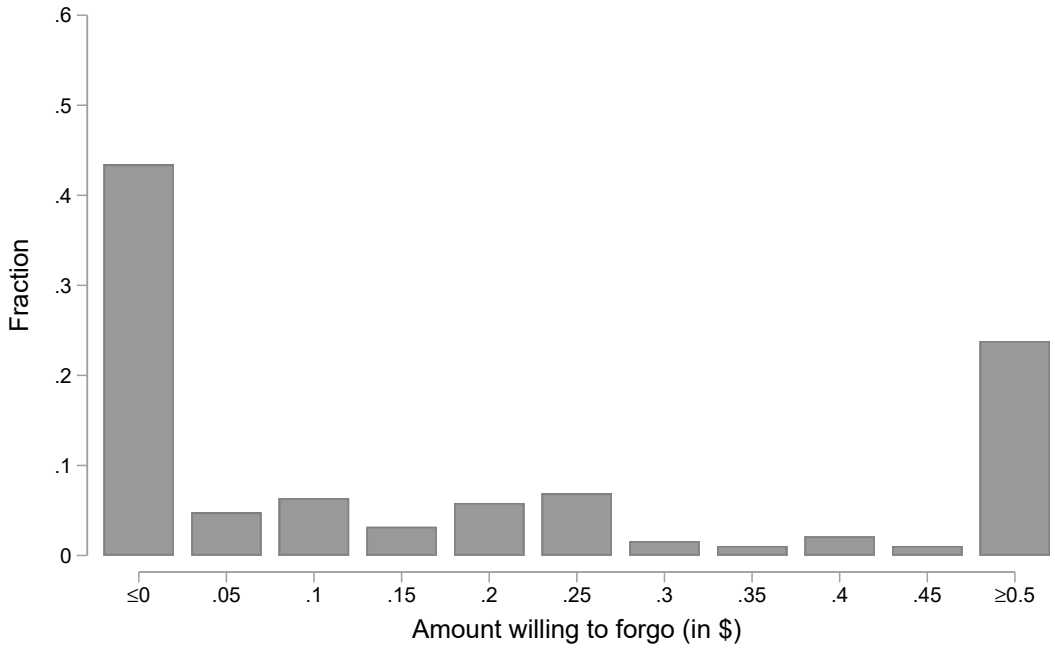

Figure A2: Distribution of $-amounts participants are willing to forgo in order to receive the money today, rather than 2$ in one week (excluding $n = 12$ participants with multiple switching points).

	Week 1 (11/27)			Week 2 (12/04)					12/07
	10am (ET)		from 10am onwards	4am (ET)	6am (ET)	6pm (ET)		from 6am onwards	4am (ET)
Immediate	invite	flip + report	base ($0.50) + bonus		invite	reminder	—	base ($0.50)	completion ($1.50)
PayDelay	invite	flip + report	base ($0.50)	bonus	invite	reminder	—	base ($0.50)	completion ($1.50)
PayDelay+ShowUp	invite	flip + report	base ($0.50)		invite	reminder	—	base ($0.50) + bonus	completion ($1.50)
PayDelay+Recall	invite	flip + report	base ($0.50)		invite	reminder	recall	base ($0.50) + bonus	completion ($1.50)
ImmediateWeek2	invite	—	base ($0.50)		invite	reminder	flip + report	base ($0.50) + bonus	completion ($1.50)
ReportDelay	invite	flip	base ($0.50)		invite	reminder	report	base ($0.50) + bonus	completion ($1.50)

Note: "—" indicates that no action with respect to the bonus needs to be taken.

Table A1: Timeline of the experiment

Dependent variable:	Report of 2 Heads		
	(1)	(2)	(3)
PayDelay (d)	0.212		
	(0.347)		
PayDelay+Recall (d)		0.356	
		(0.325)	
ReportDelay (d)			0.023
			(0.387)
Constant	2.186***	1.895***	2.462***
	(0.236)	(0.211)	(0.269)
Observations	402	388	373

Notes: Logistic regressions. The dependent variable takes the value one if the report was *2 Heads* and zero otherwise. The baseline category in model (1) is the treatment *Immediate*, in model (2) it is *PayDelay+ShowUp*, and in model (3) it is *ImmediateWeek2*. Numbers in parentheses indicate robust standard errors. Significance levels * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: The effects of our treatment manipulations on reporting behavior.

B Experimental Instructions

Note: The following screenshots are taken from the PayDelay treatment. The instructions for the other treatments are very similar and available upon request.

Week 1

General Instructions

Welcome! This study includes two HITs that have to be taken on two separate days. Each HIT will last about 6 minutes.

- **HIT 1 -- Today:** this HIT consists of two tasks:
  - a rating task and;
  - a coin flip task.

  If you complete both tasks, you will earn a fixed amount of **$0.50**;
- **HIT 2 -- in seven days from today (2018/12/04):** this HIT consists of one task:
  - a rating task similar to the one of HIT 1.

  If you complete it, you will earn a fixed payment of **$0.50**;
- **If you complete both HIT 1 and HIT 2:** you will earn a completion payment of **$1.50** in ten days from today (2018/12/07).

You will be eligible to participate in HIT 2 only if you complete HIT 1 today. If you complete HIT 1, in seven days (2018/12/04) you will receive an email from us, containing the link for HIT 2. Once you receive the link, you will have 24 hours to complete the HIT. **Only workers who complete both HITs are eligible for the completion payment of $1.50.**

On top of these payments, you may earn a bonus of up to **$2.00**. You will get to know the details of how this works after the first task of today.

After the completion of each HIT, you will receive a code to collect your payment via MTurk. Please note that the bonus and the fixed payment will be paid separately.

If you want to participate, please **enter your MTurk worker ID** and proceed to the instructions.

Instructions for Rating Task

In this task you will have to rate paintings.

There will be 20 pairs of paintings. For each pair, we would like to know which of the two you like better.

Please indicate your choice by clicking on your preferred painting.

There is no right answer, we are simply interested in your opinion.

→

1

Painting A

Painting B

→

Instructions for Coin Flip Task

Please read the following instructions carefully. After you have done this, we will ask you a few questions to ensure that you have understood the procedures of this task.

After you have answered the questions correctly, you will have to flip a coin twice and to report the outcome of the coin flips, by ticking the corresponding box on the screen. Please notice that you have to report the outcome of the coin flips truthfully. *(If you do not have a coin with you, you can use an online website, such as http://justflipacoin.com/.)*

Your bonus depends on these two coin flips as follows:

- If you get **two Heads**, you will be paid a **$2.00 bonus**;
- If you get **one Heads and one Tails,** you will be paid a **$0.25 bonus;**
- If you get **two Tails,** you will be paid a **$0.25 bonus.**

You will receive the **bonus payment** based on your report **seven days from today, a few hours before you will receive the invitation for HIT 2**.

→

Please answer the following questions.

(if you fail a question twice, your HIT will not be accepted)

You can access the instructions here: link to instructions

Question 1.

If the outcome of the two coin flips is 2 Heads, what is your bonus payment?

- $0.00
- $0.25
- $2.00

Question 2.

If the outcome of the two coin flips is 1 Heads and 1 Tails, what is your bonus payment?

- $0.00
- $0.25
- $2.00

Question 3.

If the outcome of the two coin flips is 2 Tails, what is your bonus payment?

- $0.00
- $0.25
- $2.00

Question 4.

When will you be paid the bonus?

- Today, within a few hours after you completed this HIT
- In seven days from today, r a few hours before you will receive the invitation for HIT 2
- On 2018/12/07, together with the completion payment of $1.50

→

Please flip a coin twice and report the result of your coin flips below. In case you don't have a coin, you may also use http://justflipacoin.com/ or a similar website to virtually flip a coin.

2 HEADS (=bonus $2.00)

1 HEADS and 1 TAILS (=bonus $0.25)

2 TAILS (=bonus $0.25)

You will receive the **bonus** payment based on your report **seven days from today, a few hours before you will receive the invitation for HIT 2.**

→

Week 2

General Instructions

Welcome! This is the second and last HIT of a study you started a week ago. Only workers that have completed HIT 1 and have received the link to this HIT via email are eligible to participate.

- **HIT 2 -- today:** this HIT consists one task:
  - a rating task similar to the one of HIT 1.

  If you complete it you will earn a fixed payment of **$0.50** paid today, and a completion payment of **$1.50** paid on 2018/12/07.

After the completion of the HIT, you will receive a code to collect your payment via MTurk.

→