On the persistence of dishonesty

Published Version:
Bortolotti, S., Kölle, F., Wenner, L. (2022). On the persistence of dishonesty. JOURNAL OF ECONOMIC BEHAVIOR & ORGANIZATION, 200, 1053-1065 [10.1016/j.jebo.2022.07.007].

Availability:
This version is available at: https://hdl.handle.net/11585/891099 since: 2022-07-21

Published:
DOI: http://doi.org/10.1016/j.jebo.2022.07.007

On the Persistence of Dishonesty*

Stefania Bortolotti Felix Kölle Lukas Wenner

July 4, 2022

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, providing individuals with an excuse to lie by inserting a delay between the time when private information is obtained and when it is reported does not affect cheating.

Keywords: Dishonesty, cheating, delay, discounting, experiment

JEL Classification Numbers: C91; D82; D91

*Corresponding author: Bortolotti: Economics Department, University of Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy & IZA, e-mail: *stefania.bortolotti@unibo.it*. Kölle, Department of Economics, University of Cologne, Germany, e-mail: *felix.koelle@uni-koeln.de*. 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 institution 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 several 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 an overview).

A common feature of the existing evidence is that it is based on studies that primarily focus on cases in which both the gains and the costs 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 within a few minutes. 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. To shed light on this, in the current paper we provide first clean evidence on dishonest behavior in an intertemporal context.

There are several ways in which time can affect the incentives to cheat, which can either lead to an increase or decrease in dishonesty. Dishonesty might be reduced in situations in which psychological costs of lying realize immediately but benefits accrue delayed and hence are discounted. Misreporting income to tax authorities is a prototypical example of such a situation as the (psychological) costs of misreporting occur immediately (i.e., at the time of declaring one's income), while the benefits of a lower tax burden are realized only in the future.[1] Alternatively, the intertemporal context might reduce cheating via increased costs. This can be the case when it is likely that the costs of dishonesty accrue at several points in time, for instance, when being reminded of own misbehavior or when there is a need to

[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. In our experiment, to keep the setting focused on our question of interest, we abstract away from such effects.

reiterate a lie in the future. Job candidates faking their résumé to get invited for an interview is an example of such a situation since candidates need to potentially reiterate their false statements several times during the application process. In other circumstances, time may facilitate cheating by reducing the psychological costs of dishonesty. This is particularly relevant in situations in which private information do not have to be reported immediately after being obtained. As shown by a growing literature on motivated beliefs and motivated errors (Epley and Gilovich, 2016; Exley and Kessler, 2019), people tend to "conveniently" forget negative feedback, for instance, to maintain a positive self-image (Zimmermann, 2020). While it is difficult to forget an unlucky draw when it needs to be reported right away, people may conveniently forget the unfavorable outcome when there is a delay between the realization of the state of the world and the time one needs to report it. In the latter case, people can exploit the "moral wriggle room" of manipulating their recollection of the events, thereby lowering their psychological costs from lying.

To study these aspects, we design an experiment in which participants have to privately flip 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 eight times higher than for any of the other two possible reports. To investigate the direct effect of delayed payment, participants either receive their payment from the coin flips 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 one 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.

In line with the evidence from previous literature, we find substantial but less than maximum degrees of dishonesty: 68% of our participants report having flipped *2 Heads*, which exceeds the expected reporting frequency of 25% under full honesty but falls short the 100% one should observe if everyone would simply maximize material payoffs. More strikingly, we observe no systematic differences in dishonesty across treatments. Specifically, 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 flips does not further increase dishonesty. Overall, our results indicate that dishonesty is little affected by the factor time.

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; Kocher et al., 2018; Abeler et al., 2019; Charness et al., 2019; Dimant et al., 2020). We add to this literature by investigating the role of time on dishonest behavior, an exercise which, to the best of our knowledge, has not yet been made. 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 did not affect 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, however, Andersen et al. (2018) do not study the effects of delayed benefits or the effect of increased mental engagement over time.[2]

The remainder of the paper is organized as follows. In Section 2, we propose a simple theoretical framework that models preferences for honesty in an intertemporal context. In Section 3, we explain the design and procedures of our experiment, and derive our hypotheses. The results of our experiment are presented in Section 4. Section 5 provides a discussion of our findings and concludes.

2 Theoretical Framework

To provide some intuition about how time can affect dishonesty, in the following we provide a simple theoretical framework that captures preferences for truth-telling in an intertemporal context. Our framework is based on previous studies assuming that agents are motivated by material payoffs, and that they derive some dis-utility from lying (see, e.g., Gneezy et al., 2018; Abeler et al., 2019; Khalmetski and Sliwka, 2019). We consider a setting in which there are two possible states of the world, L and H. The agent privately observes the true state of the world, $\omega \in [L, H]$, and subsequently has to make a report $r \in [L, H]$. Reporting the high state of the world yields a high payoff, while reporting the low state of the world yields a low payoff, i.e., $\pi(H) > \pi(L)$. We first consider the "standard" case where there is only one period t in which the state of the world is observed, the report about the state of the world is made, and in which all costs and benefits are realized. We assume the utility of the agent to be given by

[2] A more distant related literature investigates the role of time pressure and intuition on dishonesty. The evidence from this literature is inconclusive. While some studies report evidence that seems to suggest that dishonesty increases with time pressure (Shalvi et al., 2012; Köbis et al., 2019), others fail to replicate such an effect (Van der Cruyssen et al., 2020).

$$\pi_t(r) - c\mathbf{1}_{\omega \neq r} \tag{1}$$

where $\pi_t(r)$ denotes the material payoff obtained in period t from reporting r, and $c \geq 0$ denotes the agent's costs of lying that occur when not reporting the true state of the world, i.e., if $\omega \neq r$.[3] We allow agents to differ in their individual costs of lying by assuming that c is drawn from some continuous distribution on $[0, \bar{c}]$.

In the following, we extend this framework in two different directions, in each case explicitly removing the simultaneity of events which is present in the static model above. First, we consider the case where in period t the state of the world is realized and a report is being made, but payoffs are only obtained in period $t+1$:

$$\delta\pi_{t+1}(r) - c\mathbf{1}_{\omega \neq r} - a\delta c\mathbf{1}_{\omega \neq r} \tag{2}$$

Assuming standard exponential discounting (see, e.g., Frederick et al., 2002), payoffs in period $t+1$ are, from the perspective of period t, discounted by $\delta \leq 1$. In addition, we assume that lying costs not only accrue immediately when a false report is being made (the second term in the expression above) but, on top of that, also allow for the possibility that an agent incurs additional costs in the future. As discussed in the introduction, this might be the case when being required to reiterate a false claim a second time (e.g., when being asked about a false statement in the resume during a job interview) or when being reminded about a past misreport. We model such a situation by assuming that the costs of lying not only occur in period t but also in period $t+1$. It is conceivable that the costs of reiterating or being reminded of a false claim are lower than the (discounted) costs associated with the initial misreport, which is why we introduce an additional weighting factor $0 < a < 1$.

The second extension considers a situation where the report and its associated costs and benefits happen in period t, while the realization of the true state of the world already happens in period $t-1$. This extension is relevant in light of the recent literature on motivated beliefs and motivated errors showing how agents tend to conveniently forget negative information (see, e.g., Epley and Gilovich, 2016; Exley and Kessler, 2019; Zimmermann, 2020). We assume that in this case an agent's utility is given by:

$$\pi_t(r) - \eta^L_{t,t-1}c\mathbf{1}_{\omega \neq r} \tag{3}$$

[3]Here, we do not delve into a deeper discussion about the source of these costs. As argued in the literature, lying costs might arise from moral or religious reasons, or from self- or social image concerns (see, for instance, the discussion in Abeler et al., 2019).

Here, $\eta^L_{t,t-1}$ denotes the probability that an agent who obtained $\omega = L$ in period $t-1$ actually remembers this fact in period t. We make the simplifying assumptions that with probability $1 - \eta^L_{t,t-1}$ she (wrongly) believes that she obtained the high state of the world, while there is perfect recall for the realization of $\omega = H$, that is $\eta^H_{t,t-1} = 1$.[4]

As discussed in the introduction, we view these two extensions as natural ones to study once one considers a more dynamic framework of dishonest behavior. Below, we describe in detail how we can use our experiment to test the importance of these dynamic components in the agent's decision-making. In particular, our experiment allows us to study delayed payments (Hypothesis 1 below), increased engagement costs of a lie (Hypothesis 2 below, varying a) and motivated recall (Hypothesis 3 below, investigating the role of $\eta^L_{t,t-1}$).

3 The Experiment

3.1 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.[5] Each session is comprised of 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. The purpose of this task is to conceal the main aim of the study and to make sure that in

[4] This reasoning is also in line with the argument by Shalvi et al. (2012) stating that failing to remember the outcome of a coin flip might serve as a justified excuse to act more unethically. However, we show below that the qualitative predictions are unchanged if we instead assume that an agent who forgets the state believes that she obtained the high state with the objective probability, or if we instead assume that $\eta^H_{t,t-1} < 1$. In fact our prediction only hinges on the assumption that agents are more likely to forget a low realization than a high one, i.e., $\eta^H_{t,t-1} > \eta^L_{t,t-1}$.

[5] 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 (Paolacci et al., 2010; Arechar et al., 2018; Snowberg and Yariv, 2021). With regard to lying, recent evidence by Pascual-Ezama et al. (2020) reveals no differences in lying behavior between laboratory and MTurk participants in a coin-flipping task. Given this evidence, the simplicity of our design, and the fact that participation in our experiment was restricted to experienced subjects with a high approval rate (above 90%), we feel confident about the quality of our data. The relatively short delay of one week was chosen to keep the attrition rate between the two sessions at relatively low levels. It was further informed by the evidence from the literature on intertemporal choice showing that people exhibit substantial degrees of impatience even over very short delays (see Frederick et al., 2002; Ericson and Laibson, 2019, for overviews). The results from our follow-up experiment (see Figure A2) confirm that the majority of participants in our sample are willing to forgo substantial amounts of money to advance payoffs by one week.

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 to achieve a higher bonus.[6]

The general structure of the experiment, as well as the content of each of the two sessions, was explained to the participants at the beginning of the first session (see Appendix B for a copy of the instructions).

3.2 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 the outcome is reported (see also Table 1). 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 treatments and hypotheses were pre-registered and all our main analyses 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 immediately (within three hours) 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.

[6] Recent papers that study intertemporal preferences increasingly focus on the effort rather than monetary domain, especially when testing for present bias. Given that our main contribution is to the literature on dishonesty (rather than on time preferences), here we consider lying behavior in the monetary domain, which maximizes the comparability of our results with those of previous studies that almost exclusively have focused on monetary incentives (for a notable exception see Kajackaite, 2018).

Table 1: Overview of the experimental treatments

Treatments	**Session 1**		**Session 2**	
	Rating Task	Coin Task	Rating Task	Coin Task
Immediate	✓	flip, report, bonus (immediate)	✓	—
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 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. Delaying the bonus to the future will thus reduce any gains from cheating, due to the discounting of future payments, while costs of lying in both cases occur in the present – the gains from lying are smaller in *PayDelay* as agents discount future payments. This leads to our first hypothesis:

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

Proof. According to our theoretical framework, in *Immediate* an agent misreports if $\pi_t(H) - \pi_t(L) \geq c$, while in *PayDelay* she does so if $\delta(\pi_{t+1}(H) - \pi_{t+1}(L)) \geq c$. For $\delta < 1$, there thus exist values of c for which an agent cheats in *Immediate* but not in *PayDelay*, while the reverse is not true.[7] □

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*. In both treatments, session 2 only comprised the rating task. While session 2 is, in fact, not needed to answer our

[7] Note that this argument essentially assumes that for these two treatments $a = 0$, i.e., the agent does not experiences additional lying costs when participating in session 2. To us, this seems plausible because agents only complete the rating task in this session. However, the prediction is unchanged if we allow for $a > 0$, as long as a is not larger in *Immediate* than in *PayDelay*, which seems a reasonable assumption; in fact we would not expect a to differ across these treatments at all.

first research question, to keep instructions as similar as possible across all our treatments (see below), and to avoid differential selection effects, we also included the second session in this set of treatments.

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 are identical, except for the degree to which participants need to engage with their (potentially untruthful) report. 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 they are not able to recall correctly their initial report, they will not receive any bonus. Importantly, all of this was explained to participants already in week 1.[8] This implies that all participants untruthfully reporting *2 Heads* in the first week knew that they had to reiterate (and thus remember) their false claim in the second week.[9] This leads to our second hypothesis:

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

Proof. For *PayDelay+Recall*, we can use our theoretical model to show that an agent misreports if $\frac{\delta(\pi_{t+1}(H)-\pi_{t+1}(L))}{1+a\delta} \geq c$. As the left-hand side of this inequality increases as a decreases

[8] Specifically, in week 1 participants were told the following: "In HIT 2, seven days from today, we will ask you to recall your report from today. You will only receive your bonus payment if you can correctly recall the outcome of today's coin flip. If instead, your report in HIT 2 does not match today's report then you receive no bonus payment." To ensure participants' understanding of this procedure, they had to answer several control questions regarding their bonus in case their two reports do or do not match.

[9] We note, however, that given the online nature of our experiment, we do not observe how strongly participants engaged with their own behavior and by which means (e.g., memorizing, written note, taking a picture or screenshot) participants aimed at recalling their initial report in week 2. We will come back to this issue when discussing our results in Section 5.

and the premise of our experimental design is that the engagement costs with the lie are lower in *PayDelay+ShowUp* (maybe even zero) than in *PayDelay+Recall*, there thus exist values of c for which an agent cheats in *PayDelay+ShowUp* but not in *PayDelay+Recall*, while the reverse is not true. □

As specified in the pre-analysis plan, to test our second hypothesis we rely on the reports from 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 correlated with the reported outcome in week 1, then using the data from the second week 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 (within three hours 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 (within three hours after completing the session).

While in *ImmediateWeek2* the coin task is introduced and takes place only in week 2, in *ReportDelay* it is already introduced in week 1, and participants are instructed to flip the coins in week 1 but to report the outcome only in week 2.[10] As we have argued in Section 2, in the latter case agents may be more likely to forget the result of their coinflip, or might find it "convenient" to do so, in case the bad state of the world realized.[11] This leads to our third hypothesis:

[10]The exact instructions were as follows. In the first session, participants in *ReportDelay* were told: "After you have answered the questions correctly, you will have to flip a coin twice. In HIT 2, seven days from today, you have to report the outcome of the coin flips, by ticking the corresponding box on the screen." In the second session, it was stated that "After you have answered the questions correctly, you will have to report the outcome of the coin flips from HIT 1 by ticking the corresponding box on the screen." So while the instructions on when to flip the coins and when to report the outcome were clear, we note that given that participants flip the coins in private, we cannot observe when they actually flip the coins. We will come back to this issue in Section 5 when we discuss our results.

[11]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

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.*

Proof. For *ImmediateWeek2*, an agent's decision whether to misreport is formally the same as in *Immediate*, i.e., an agent lies and reports the high state if $(\pi_t(H) - \pi_t(L)) \geq c$. In contrast, in *ReportDelay* an agent finds it optimal to lie if $\frac{(\pi_t(H)-\pi_t(L))}{\eta^L_{t-1,t}} \geq c$. If motivated recall plays a role, i.e., $\eta^L_{t-1,t} < 1$, there exist values of c for which an agent cheats in *ReportDelay* but not in *ImmediateWeek2*, while the reverse is not true. □

Here, 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 (in accordance to the pre-analysis plan).

3.3 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 of more than 90%. 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 considered a small to medium effect size (see, e.g., Cohen, 1988; Schäfer and Schwarz, 2019).[12] Depending on the baseline proportion of reports of *2 Heads*, this corresponds to a treatment difference of 0.13 to 0.15.

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.

[12]While effect sizes can vary greatly across studies, a number of highly influential experiments on cheating behavior have reported standardized effect sizes two to three times larger than 0.28. Abeler et al. (2019) report a standardized effect size of 0.77 when analyzing the effect of observability on lying (Observable vs. Unobservable treatments). Similarly, Weisel and Shalvi (2015) find an effect size of 0.91 when testing the role of collaboration in shaping corruption (Aligned Outcomes vs. Individuals treatments). Slightly smaller effect sizes, ranging from 0.14 to 0.57, are reported by Kajackaite and Gneezy (2017) when comparing treatments with and without possible exposure of one's lie (Cheating Game vs. Mind Game, four different stake sizes). However, we cannot rely on any previous study directly testing the role of time on cheating, hence calling for the need of independent repetitions of the same experiment, as discussed in the conclusions

As specified in our pre-analysis plan, we deliberately over-sampled observations in *ReportDelay* and *ImmediateWeek2* because for 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*.[13] 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 overall attrition rate across weeks was fairly low, 14.5%, and, more importantly, not significantly different across treatments ($\chi^2(5) = 7.96, p = 0.159$).[14]

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 were determined by 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 completed both sessions. 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

[13]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.

[14]The attrition rates for each treatment are as follows: *Immediate*: 18.7%, *PayDelay*: 14.2%, *PayDelay+ShowUp*: 13.6%, *PayDelay+Recall*: 9.0%, *ImmediateWeek2*: 15.3%, *ReportDelay*: 15.7%.

of \$3.83, for a study that lasted around eight minutes (both sessions combined).[15]

4 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 68.9%, which is significantly higher than the expected 25% under truthful reporting (Binomial test, $p < 0.001$) but significantly lower than the expected 100% if everyone would maximize monetary payoffs (Binomial test, $p < 0.001$). Under the assumption that nobody misreports to their disadvantage, we can infer that about 43% of our participants reported dishonestly. To put these results into perspective, we can compare them to the findings of the meta analysis by Abeler et al. (2019). For studies using a (non-repeated) coin-flip task, they find a standardized report – measuring how much money participants leave on the table by not lying to the fullest extent – of 0.325, which is very similar to the 0.377 we observe in 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 ($\chi^2(1) = 1.14, p = 0.286$), a result that is 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.*

What can explain this surprising result? One possible explanation is that participants in our experiment simply do not discount payments that are delayed by one week. To test

[15] 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 the exact same method for both immediate and delayed bonuses. Furthermore, given that participation in our experiment was restricted to individuals with a sufficiently high overall approval rate (see above), all our participants had plenty of experience with this payment system.

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

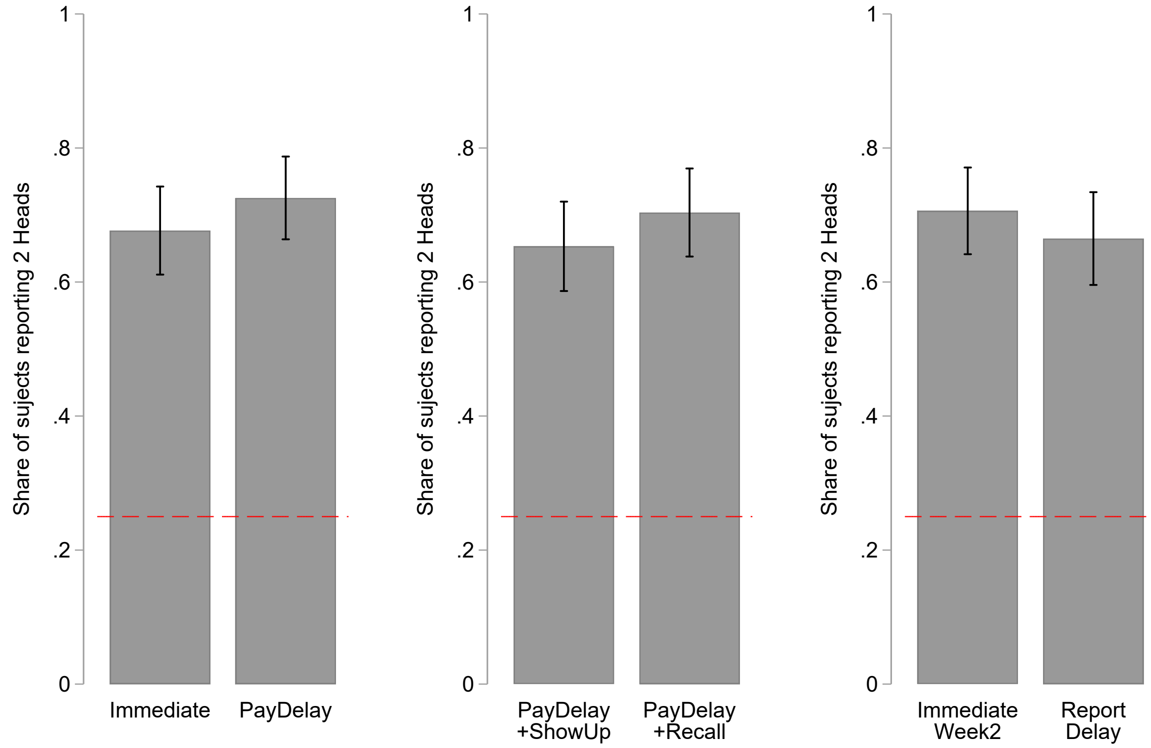

Notes. Solid lines indicate 95% confidence intervals. Dashed lines indicate the expected frequency if everyone would report truthfully.

for this possibility, we conducted a follow-up study with $n = 201$ new participants, recruited from the same subject pool on MTurk. Similar to our main experiment, participants in the follow-up study had to first complete a Rating Task. Afterwards, instead of reporting the outcome of the coin flips, they were asked to complete a standard price list eliciting 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 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%) 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). In Section 5 we discuss some alternative explanations.

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 ($\chi^2(1) = 1.13, p = 0.288$), a finding that is again confirmed by regression analysis (see column (2) in Table A2 in Appendix A). The 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 by not showing up to the second session. Furthermore, we find that only 3.5% (6 out of 172) of the participants in our *PayDelay+Recall* failed to correctly recall their initial report.[16] 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 in some way. 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 stating that introducing a delay between the time when private information are received and when they are shared increases dishonesty. To this end, we compare the fraction *2 Heads* in *ImmediateWeek2* and *ReportDelay*. The results are shown in the right panel of Figure 1. Contrary to our hypothesis, the figure reveals a

[16]We do not find any systematic relationship between the initial report and the recall rate. 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 ($\chi^2(2) = 0.84, p = 0.656$). The overall completion rate is thereby not significantly different from the one in *PayDelay+ShowUp* ($\chi^2(1) = 2.02, p = 0.156$).

very similar fraction of *2 Heads* across treatments: 70.6% in *ImmediateWeek2* and 66.5% in *ReportDelay*. 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 ($\chi^2(1) = 0.75, p = 0.388$; see also column (3) in Table A2 in Appendix A). 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.*

5 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 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 find 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 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 payment effect is that misreporting in general seems to be only marginally affected by the size of the potential gains from lying (see Abeler et al., 2019, for a discussion, p. 1123). As such, even when discounting future payoffs very strongly this might have only a minor impact on dishonesty. Another explanation might be rooted in insights from multi-attribute choice theory (see, e.g., Houston and Sherman, 1995), which assumes that the decision weight put on a particular attribute is inversely related to how many attributes differ between choice options. Some evidence along these lines comes from recent studies showing that time matters less when decisions involve more than one dimension. For instance, Cubitt et al. (2018) show that people are considerably more patient when comparing options of different types (e.g., apples now versus oranges later) compared to decision-situations that only involve options of the same type (e.g., apples now versus apples later). Similar evidence is provided by Deck and Jahedi (2015), Ederer and Schneider (2019), and Kölle and Wenner (2021) who find that the timing of consequences has little impact in situations in which choices affect both oneself and others. In a similar vein, it might be that in our context in which individuals have to evaluate multiple attributes – time and lying – at the same time, they simply ignore (or put very little weight on) the factor time.

Our study set out to study dishonesty in an intertemporal context in a very simple and minimalist environment. While we judge this the right approach for an initial analysis, we believe that further research is needed to gain a more complete understanding of dishonesty and its underlying mechanisms in such contexts. We see several fruitful avenues for future research. For instance, given the anonymous online nature of our experiment, in *PayDelay+Recall* we could not control how subjects tried to recall their report from one week to the other and if that had an effect on their report in the first session. Similarly, in *ReportDelay* we did not have control over the time at which participants actually flipped the coin.[17] Future studies could build on this by considering settings in which these mechanisms are better observable, which, in turn, could shed additional light on the size of the engagement and recollection costs in such contexts. Such studies could also focus on investigating individual-level heterogeneity, e.g., by employing a design that elicits time preferences and lying behavior within subjects.

Future research should also replicate our results, especially given that based on our ex-ante hypotheses our findings can be considered as surprising. As argued by Maniadis et al. (2014), already a few independent replications of an initial finding can improve the probability that a declaration of a research finding is true.[18] Complementary to this approach, the robustness of our findings could be tested by employing new research designs aimed at testing our hypotheses, for instance, by considering decision environments in which the costs of lying are increased (e.g., situations in which social image concerns are more prominent)

[17] Most studies do not have control over this dimension and cannot check if a person decided to flip the coin at all. For a notable exception and evidence on this point, see Pascual-Ezama et al. (2020).

[18] Maniadis et al. (2014) refer to this as the Post-Study Probability (PSP), which depends on the prior probability that a certain phenomenon is true, the power of the study, and the number of independent researchers that test the phenomenon. To determine whether an initial finding is indeed true, one needs to make an assumption about the prior. Based on the benchmarking exercise of Maniadis et al. (2014, see Table 4 on p. 288), given our power of 0.8 and given that there are no replications of our results yet, the PSP for a prior of 0.05 would be 0.12, for a prior of 0.20 it would be 0.38, and for a prior of 0.35 it would be 0.57. Based on these estimates it is clear that more research is needed before drawing final conclusions.

or in which decisions have non-monetary consequences.

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.

Cohen, J. (1988): *Statistical power analysis for the behavioral sciences*, NY: Academic Press.

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.

Cubitt, R., R. McDonald, and D. Read (2018): "Time matters less when outcomes differ: Unimodal vs. cross-modal comparisons in intertemporal choice," *Management Science*, 64, 873–887.

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.

Deck, C. and S. Jahedi (2015): "Time discounting in strategic contests," *Journal of Economics & Management Strategy*, 24, 151–164.

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.

Ederer, F. and F. Schneider (2019): "Trust and promises over time," *American Economic Journal: Microeconomics*.

Epley, N. and T. Gilovich (2016): "The mechanics of motivated reasoning," *Journal of Economic Perspectives*, 30, 133–40.

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

Ericson, K. M. and D. Laibson (2019): "Intertemporal choice," in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 2, 1–67.

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.

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.

Houston, D. A. and S. J. Sherman (1995): "Cancellation and focus: The role of shared and unique features in the choice process," *Journal of Experimental Social Psychology*, 31, 357–378.

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.

Kocher, M. G., S. Schudy, and L. Spantig (2018): "I Lie? We Lie! Why? Experimental Evidence on a Dishonesty Shift in Groups," *Management Science*, 64, 3995–4008.

Kölle, F. and L. Wenner (2021): "Is Generosity Time-Inconsistent? Present Bias across Individual and Social Contexts," *The Review of Economics and Statistics*, 1–45.

Maniadis, Z., F. Tufano, and J. A. List (2014): "One swallow doesn't make a summer: New evidence on anchoring effects," *American Economic Review*, 104, 277–90.

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.

Paolacci, G., J. Chandler, and P. G. Ipeirotis (2010): "Running experiments on amazon mechanical turk," *Judgment and Decision Making*, 5, 411–419.

Pascual-Ezama, D., D. Prelec, A. Muñoz, and B. G.-G. de Liaño (2020): "Cheaters, Liars, or Both? A New Classification of Dishonesty Profiles," *Psychological Science*, 31, 1097–1106.

Schäfer, T. and M. A. Schwarz (2019): "The meaningfulness of effect sizes in psychological research: Differences between sub-disciplines and the impact of potential biases," *Frontiers in Psychology*, 10, 813.

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.

Van der Cruyssen, I., J. D'hondt, E. Meijer, and B. Verschuere (2020): "Does honesty require time? Two preregistered direct replications of Experiment 2 of Shalvi, Eldar, and Bereby-Meyer (2012)," *Psychological Science*, 31, 460–467.

Weisel, O. and S. Shalvi (2015): "The collaborative roots of corruption," *Proceedings of the National Academy of Sciences*, 112, 10651–10656.

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

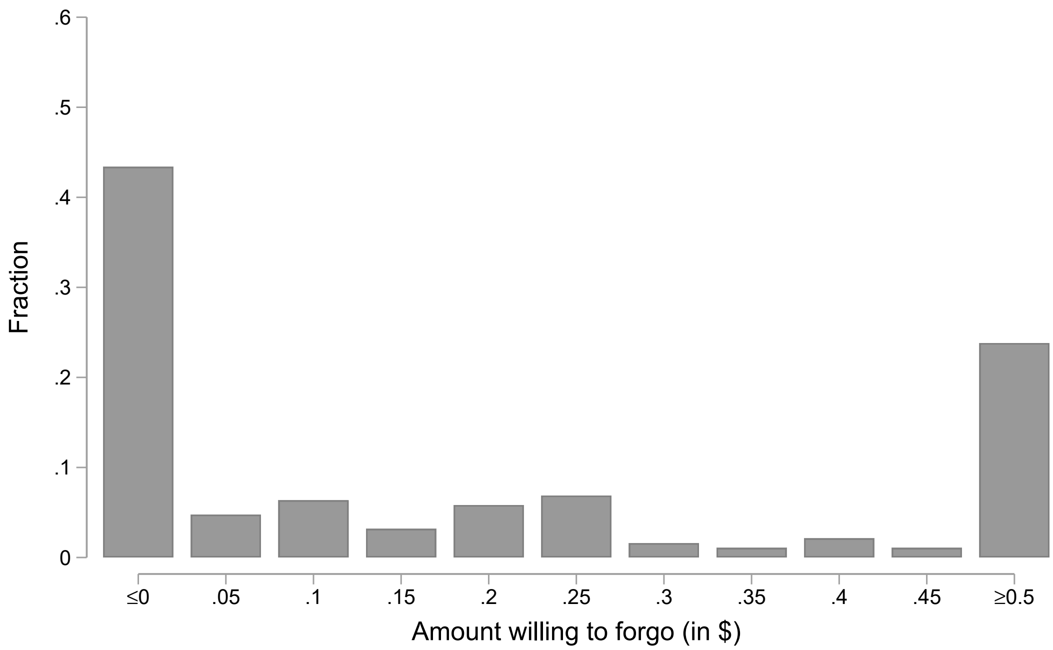

Figure A2: Distribution of $-amounts participants are willing to forgo in order to receive the money today, rather than 2$ in one week (excluding $n = 12$ participants with multiple switching points).

	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)	0.233		
	(0.219)		
PayDelay+Recall (d)		0.232	
		(0.218)	
ReportDelay (d)			-0.192
			(0.223)
Constant	0.739***	0.633***	0.877***
	(0.152)	(0.149)	(0.158)
Observations	402	388	376

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.

→