It's time to cheat!

It's Time to Cheat!*

Alessandro Bucciol†
University of Verona

Simona Cicognani
Leiden University

Natalia Montinari
University of Bologna

Abstract

We study the correlation between time preferences and cheating. In our experiment, cheating increases the earnings of those who commit it and only entails a moral cost. We are the first to measure both (a proxy for) the propensity to cheat and time preferences at the individual level, determining whether cheaters are more likely to be more present-biased or to have a higher discount factor. We observe widespread cheating, which prevails among subjects with present bias and overconfidence.

Keywords: Cheating; Time Discounting; Quasi-hyperbolic preferences.
JEL Classification: D91; C91; D81.
Declarations of interest: None.

---

* The present study was approved by the Bioethics Committee of the University of Bologna (Italy), prot. #127883, 2/7/2020. We are grateful to participants at the 2019 IMEBESS in Utrecht and at the 2019 NCBEE in Kiel, as well as to seminar participants at the Universities of Alicante, Bologna, Bozen, Innsbruk, Parma, Siena and Verona. A special thanks goes to Bradley Ruffle, Loukas Balafoutas, Valerio Capraro, David Hugh-Jones, Martin Kocher, Maria Vittoria Levati, Giovanni Ponti, Shaul Shalvi, Charlie Sprenger, Toke Reinholt Fosgaard and Luca Zarri for useful comments.

† Corresponding author. Postal address: University of Verona, Dept. of Economics, Via Cantarane 24, 37129 Verona, Italy. Email: alessandro.bucciol@univr.it. ORCID: 0000-0001-8163-7617.

1. Introduction

Understanding the determinants of dishonesty is the prerequisite for designing effective policies to limit the adverse effects of misbehavior in our societies (Bucciol & Montinari, 2019). Many of the daily cheating behaviors undertaken involve an intertemporal dimension. The act of cheating and its benefits usually do not arise at the same moment, and when it is the case, individuals who cheat may still face the risk of being caught and sanctioned later. Consider, for instance, everyday dishonesty such as intellectual property theft (music, film, and software piracy), omitting income, inflating deductions, overstating the value of insurance claims and overcharging consumers (Jacobsen et al., 2018; Vranka et al., 2019). In these unethical acts, benefits occur (almost) immediately, incorporating a component associated with self-control ability. In contrast, costs driven by being caught are uncertain and, if any, only occur in the future.

Our study focuses on both the temporal and risky components of cheating. We analyze the correlation between the tendency to cheat and time preferences (split in discount factor and present bias), controlling for individual risk attitude. Being able to connect these two domains of behavior could endow the policymaker with the possibility to customize a set of incentives that, building on the time preferences of the agents (more easily observable), can limit the emergence and the magnitude of cheating (a trait that individuals are often reluctant to reveal). Collecting clean evidence on this phenomenon from observational data is challenging since people frequently try to hide their dishonest behavior. Therefore, an alternative approach relies on the methodology of tightly controlled laboratory experiments, where participants are confronted with (cheating) games characterized by the possibility of acting dishonestly (typically at the expense of the experimenters).

Our goal is to investigate two main questions: is there a positive correlation between the likelihood of cheating and: i) time discounting; ii) present bias? We conjecture that the tendency to cheat is more frequent among individuals who attribute more importance to the present, displaying a lower discount factor and exhibiting a more substantial present bias. Specifically, the discount factor can better explain planned acts of cheating, while the present bias can better explain impulsive cheating behavior. In the first case, individuals may strategically act dishonestly to gain an advantage in the future by setting up detailed and meticulous plans; think, for example, of cases where the main cheaters have been politicians or bankers. In other cases, the acts of cheating appear dominated by impulsivity, such as, for example, everyday non-planned actions of stealing.

We devise a novel task that builds on existing literature to identify a proxy for dishonesty at the individual level. Participants are confronted with 10 multiple-choice trivia questions and receive a bonus only if they report solving at least 8 correctly. Participants self-report their performance after having checked the solutions available on the back of the printed instruction sheet. The main driver for choosing this cheating task rather than, for instance, the coin toss or the die roll tasks was that we wanted to infer cheating behavior at the individual level with higher precision compared to these other tasks. As in the other cheating tasks used in the literature, the risk of being caught cheating is null, the cost of cheating is only moral, and payments occur immediately. Importantly, to be able to make cheating claims at the individual level, we validate our cheating task in two pilot studies. The validation allows us to safely state that for participants reporting 8 or more correct answers, we have a good proxy for cheating. To estimate time preferences, we rely on the Convex Time Budget (CTB) method developed by Andreoni and Sprenger (2012), which uses a single, relatively simple instrument to capture risk aversion, long-run discounting, and present bias in a utility function with quasi-hyperbolic discounting (Cohen et al., 2020; Laibson, 1997; O'Donoghue & Rabin, 1999).

To this aim, we implemented a laboratory experiment with two separate tasks to elicit cheating and time preferences. Our design also varies the order in which the two tasks are presented to account for possible spillover effects from one task to the other. This way we introduce an unavoidable element of delay in payments. Specifically, students who first play the cheating game must wait for the completion of the time preference task before receiving the final payment. This may make the (potential) monetary consequences associated with cheating less salient compared to the situation in which the cheating game is played as a second task. However, our results suggest that this feature of the experimental design has no impact on the decision to cheat between treatments (i.e., depending on the sequence of tasks faced). In addition, we made it very clear in the instructions that all decisions are anonymous and that the experimental team makes no link between individual decisions and individual identity, this way limiting the concern about possible delayed consequences of cheating behavior.

The decision to implement two separate tasks (i.e., the time preferences elicitation and the cheating elicitation) stems from our goal of measuring the behavior in two different domains and connecting them at the individual level. The alternative would have been introducing a time component in the cheating task or, conversely, a cheating component in the time preference elicitation task. However, this design choice comes with a limitation: each elicitation (e.g., cheating) task would correspond to a specific configuration of the other parameter of interest (e.g., time preferences). Our primary goal, though, is to establish connections between two distinct domains of behavior: moral behavior underlying cheating behavior and time preferences, rather than fixating on individual parameter choices. In this sense, our contribution differs from previous studies analyzing cheating by varying the timing of receiving the prize at the treatment level (e.g., Ruffle and Tobol, 2014) since it allows us to consider individuals' time preferences without restrictions or manipulations on specific time preference parameters.

Our contribution to the literature is threefold. First, by presenting the first study that measures both the propensity to cheat (using a proxy at the individual level) and time preferences, we can determine whether cheaters are more likely to be more present-biased or to have a higher discount factor. Second, compared to previous studies that associate self-control depletion with cheating (e.g., Gino et al., 2011; Mead et al., 2009), we estimate an essential proxy for self-control (i.e., the present bias parameter) and connect it to the likelihood of cheating. Third, we are among the first to investigate the association between cheating and self-confidence at the individual level. Differently from Adams et al. (2018), who focus on different mechanisms of how cheating can affect confidence, we analyze whether a nexus exists between cheating and overconfidence without manipulating any of these two behavioral domains.

The paper is organized as follows. Section 2 reviews the related literature on cheating and time preferences and discusses our study's main elements of novelty. Section 3 describes the experimental design, formulates our research hypotheses, and details the procedures and summary statistics. Section 4 reports and discusses our results, while Section 5 concludes. Finally, the Appendix reports some robustness checks and more details on the experiment and its design.

2. Related Literature

Our cheating task is inspired by Nagin and Pogarsky (2003) and Hugh-Jones (2016). As in our experiment, the authors deal with a set of trivia questions that are intended to be so complex that it would be unlikely for participants to know a given number of correct answers (that allows them to obtain a monetary bonus) and impossible to earn a bonus by just guessing. We depart from these studies in two important ways. Contrary to Nagin and Pogarsky (2003), we preserve the anonymity of participants' choices, which is especially relevant when studying dishonesty. Moreover, contrary to both works, we first validate our questions in two pilot studies, which allows us to safely state that participants who get the monetary bonus are cheating. Importantly, Hugh-Jones (2016) finds a very high (and positive) correlation between this kind of task and a coin flip task (as in Bucciol and Piovesan, 2011) implemented on the same participants.

Regarding time preferences, the papers most closely related to ours are Andreoni and Sprenger (2012) and Andreoni et al. (2015), which introduce the Convex Time Budget (CTB, henceforth) method, which we also use in our experiment. CTB is a widely recognized task in the time preferences literature (Cohen et al., 2020) in which quasi-hyperbolic time preferences are elicited in combination with constant relative risk aversion from variations in linear budget constraints over early and later incomes.

There is a small set of articles connecting cheating to time manipulation. Nagin and Pogarsky (2003) are the first to investigate the nexus between cheating and time discounting. They consider a measure of present orientation by asking participants a hypothetical question. However, the answer to this question is non-incentivized, and the authors cannot obtain participants' time preferences.

Ruffle and Tobol (2014) add a time preferences component to the "standard" die roll cheating game designed by Fischbacher and Föllmi-Heusi (2013), which detects cheating at the group level. They find that temporally distancing the decision task from the payment of the reward increases honest behavior. This is similar to what Bortolotti et al. (2022) do in an online experiment, in which they manipulate the timing of reporting of private information and the timing of the realization of the benefits from lying, to examine whether manipulations affect the dishonesty observed in a coin flip task. They do find minimal differences across treatments. However, neither of these two works estimates time preferences, which we do.

Alan et al. (2020) study cheating behavior among elementary school children. In their study, participants are confronted with a creative performance task where it is easy to imitate others' work. In this context, cheating is defined as presenting output that is not one's own. In contrast, time preferences are elicited using an allocation task based on a simplified version of Andreoni and Sprenger (2012). They find that children with higher IQ and socio-economic status are more likely to cheat, while risk and time preferences have no robust direct effect.

Time preferences and self-control are intertwined (Hoch & Loewenstein, 1991). Often, the urge to decide (for instance, to buy a product) involves overriding long-term preferences. Self-control is the psychological capacity that enables people to enact behaviors consistent with their long-term goals (e.g., being an ethical person) and refrain from engaging in behaviors driven by short-term, selfish reasons. In economics, self-control problems have been modeled as time-inconsistent, present-biased preferences, where present bias identifies the tendency of people to attach more substantial weight to payoffs closer to the current time when considering trade-offs between two future moments (O'Donoghue & Rabin, 1999). For example, Mead et al. (2009) and Gino et al. (2011) find that dishonesty increases when an initial act depletes people's self-control resources.

Time delays in cheating opportunities are also connected to self-control. In other words, what happens to cheating if individuals are endowed with extra time to reflect on their choices (e.g., the cooling-off effect on dishonest decisions)? For instance, Andersen et al. (2018) compare the results in cheating games when choices are immediate vs. when choices are taken after an extra day of reflection. They find that allowing for extra reflection time does not impact cheating.

Since both time discounting and self-control may matter as determinants of cheating, we consider a framework allowing us to elicit both dimensions in our analysis.

3. Experimental Design and Hypotheses

Our experiment comprises two main parts: a paper-based survey with a cheating opportunity (C) and a computerized task to measure time preferences (T), allowing us to also elicit risk preference parameters. We collect data on two groups of subjects, varying the order of the two tasks. The rationale is to check for order effects. One could argue that, besides possible spillovers from one task to the other, what changes across sequences is how soon, after the cheating task, subjects (may) receive the payouts from cheating. The Appendix reports the experimental instructions in the original Italian language, the English translation, and all the materials used.

3.1. Survey with cheating opportunity

3.1.1. The survey

Participants receive an 8 EUR flat payment to complete the survey. The survey, shown in the Appendix, is paper-based, identified by an ID number assigned to each participant, and made of three sections. The first one aims to collect personal information about the participants (such as the number of siblings, gender, age, field of study, and city of origin). The second one aims at gathering information on the socio-economic status of the participant's household (such as the average after-tax income of parents, education and employment status of both parents). The third section measures individual attitudes and personality traits using validated scales from economics and psychology. Specifically, we include a 9-item measure of fluid intelligence based on a selection of Raven's matrices (IQ test, henceforth) elaborated by Bilker et al. (2012). The 9 items selected by Bilker et al. (2012) guarantee administration time savings and, at the same time, high predictive power similar to the original 60-item Raven scale. We also include a measure of self-confidence, obtained as the difference between the expected and actual number of correct answers in the IQ test.

Given that the IQ test is conducted in paper and pencil mode and the self-confidence task is part of the same paper-based survey, we decided not to incentivize it. Despite this self-confidence measure, our results in terms of overconfidence do not deviate from prior (incentivized) studies that report males as more overconfident than women.

3.1.2. The cheating task within the survey

At the end of the survey, participants are confronted with ten multiple-choice questions on a separate sheet with no ID. Subjects are requested first to answer the questions and then check the number of correct answers on the back of the instruction sheet to report it on the computer screen to speed up calculating the earnings for that task. Letting subjects self-assess their performance has been used in other studies, such as Mazar et al. (2008) and Gino et al. (2009). Subjects are told they will earn an additional 8 EUR only if they report eight or more correct answers out of ten questions. The experimenter collects the paper-based survey at the end of the task, while the sheet with the trivia quiz is left on each participant's desk to get the instructions. Participants are told that they can shred their answer sheet and that we will do it for them if they do not. The time employed should guarantee good quality of the answers.

Participants are given 30 minutes to answer the survey and the trivia quiz. The time assigned was enough since all participants inserted the number of correct answers (together with the personal ID) on the computer in due time.

There are several ways to cheat in this task. One could try to answer, check which answers are correct, and then decide to report eight or more correct answers. One could do the same without the intermediate step of checking the correct answers. It could also be possible to check the correct answers and then report them in the answer sheet or to skip any attempt to answer the questions. We have no control over the way subjects cheat. Whatever the chosen method, however, the purpose is identical: making a false claim, which is the focus of our study. It is also likely that those who report seven or fewer correct answers are also cheating. The main driver for choosing this cheating task in our analysis was that we wanted to be able to infer cheating behavior at the individual level with higher precision compared to other popular tasks, such as the coin toss or the die roll tasks, where rewarding outcomes may also arise because of luck, in a fashion similar to what proposed by Yaniv et al. (2019). A notable exception that cleanly measured cheating at the individual level is represented by Dai et al. (2018), who devise a ticket fraud game which they then link to whether the subject is a fare dodger to study the external validity of their lab measure. However, no time dimension of cheating is designed in their study.

3.1.3. Validation of the cheating task

To validate our cheating task, we conducted two pilot studies.[1]

Preliminary pilot study (non-incentivized). We tested 20 challenging questions in a preliminary pilot study run at the University of Bologna on 125 students (more details are in the Appendix). The study was run in a classroom, and subjects were not incentivized. We picked the 10 questions that were less frequently answered. Specifically, nobody in the first pilot study responded correctly to 8 of these 10 questions (average number of correct answers: 2.06).

Second pilot study (incentivized). Later, we invited to the BLESS experimental laboratory of the University of Bologna 96 subjects who faced the same incentive structure as the main experiment but were given no cheating opportunity. Participants reported on average 2.63 correct answers (maximum 6) to the same 10 questions of the final experiment (which proved to be the 10 most difficult questions in the preliminary pilot) and 4.76 correct answers to the other 10 questions (the least difficult 10 questions of the pilot). Correct answers showed no significant correlation with gender and age. We took care of excluding students in the pilots from the pool of subjects who got the invitation to the experiment.

The probability (from a binomial distribution) of correctly answering by chance eight or more of these questions, where each question presents four alternatives, is tiny and equal to 0.000416. Based on the results of the pilot studies, even if each subject is assumed to know 3 correct answers, the probability of randomly answering the other questions and earning the bonus is as tiny as 0.0129. Because the questions are tough, and it is virtually impossible to answer them by chance correctly, we assume that those who report eight or more correct answers are reasonably cheating.

[1] Data from both pilot studies are not included in this manuscript but are available upon request.

3.2. Time preferences elicitation

To elicit time preferences, we implement a computer-based task drawn from Andreoni et al. (2015) and inspired by Andreoni and Sprenger (2012): the so-called "Convex Time Budget" (CTB) task. This task elicits quasi-hyperbolic time preferences with constant relative risk aversion from variations in linear budget constraints over early and later incomes. Precisely, we empirically estimate the discount factor parameter δ, the present bias parameter β, and the risk aversion parameter α with a non-linear least squares method from the following intertemporal utility function $U(\ldots)$ with argument x at time t and time $t+k$:

$$U\left(x_t, x_{t+k}\right) = \begin{cases} \dfrac{1}{1-\alpha}\left(x_t^{(1-\alpha)} + \beta\delta^k x_{t+k}^{(1-\alpha)}\right) & \text{if } t = 0 \\ \dfrac{1}{1-\alpha}\left(x_t^{(1-\alpha)} + \delta^k x_{t+k}^{(1-\alpha)}\right) & \text{if } t > 0 \end{cases}$$

The CTB in Andreoni et al. (2015) consists of 24 choices, with the last 12 choices. A choice involves 6 options, each made of two amounts paid at two different points in time; the last 12 choices regard points farther in time. A choice may be inconsistent when allocating a smaller budget share to the future payment date than the previous choice (Giné et al., 2018). In our experiment, we rescale the payoffs of Andreoni et al. (2015), reducing them by 50%; therefore, our participants were confronted with allocations of 10 EUR. An example of the decision screen is reproduced in Appendix Figures A1a and A1b.

3.3. Hypotheses

Based on the evidence discussed in Section 2, we formulate two hypotheses on the association between the likelihood of cheating, the time discount factor, and the present bias.

Hypothesis 1. *There is a positive correlation between time discounting and the likelihood of cheating at the individual level.*

Hypothesis 2. *There is a positive correlation between present bias and the likelihood of cheating at the individual level.*

Hp.1 and Hp.2 are based on the idea that individuals who give more importance to the present than the future are more likely to cheat. Note that a higher weight given to present outcomes and costs corresponds to a lower discount factor parameter, and a higher present bias corresponds to a lower present bias parameter. Stated differently, we expect a negative correlation between the discount factor parameter δ, the present bias parameter β, and the number of reported correct answers, respectively.

However, the two hypotheses refer to different dimensions of individual time preferences. Specifically, Hp.1 postulates an effect of the discount factor placed on returns receivable (or costs payable) in the future. It is consistent with (time-constant) exponential discounting models of time preferences. The classical discounted utility model implies that choices are time consistent; individuals will make the same utility trade-off between two periods regardless of when they make the allocation (Strotz, 1955). According to this hypothesis, the tendency to act dishonestly should not necessarily be associated with a lack of self-control (or higher present bias). Still, it could reflect a different weight given to rewards and costs occurring at different points in time, i.e., a difference in the discount factor, with a higher factor corresponding to higher patience, as documented by Ruffle and Tobol (2014) and Nagin and Pogarsky (2003).

Hp.2 explicitly refers to the role of the present bias component, which describes inconsistencies (or preferences reversal) arising when comparing short-term preferences with long-term preferences only considering the discount rate. At the same time, it becomes consistent when (quasi) hyperbolic discounting models of time preferences are used. According to this hypothesis, individuals who have a stronger present bias (i.e., a lower present bias parameter) are more likely to encounter self-control problems and should also display a higher likelihood to cheat, in line with the results reported by Mead et al. (2009) and Gino et al. (2011). An example of inconsistency coherent with present bias is observed when, for example, a subject prefers 10 USD now rather than 12 USD in a day, but he/she prefers 12 USD in a year plus a day rather than 10 USD in a year. When there is inconsistency, the present bias parameter is below one. The lower the present bias parameter, the higher the bias toward the present.

Note that the two behavioral hypotheses rely on two assumptions regarding individual rationality. While acts of cheating resulting from a preference for the present compared to the future do not entail a failure of individual rationality and may respond to standard monetary incentives, acts of cheating resulting from a present bias imply that the individual is likely to cheat, deviating from the original plans and make the incentives designed for time-consistent individuals ineffective (O'Donoghue & Rabin, 2006). Moreover, sophisticated individuals (i.e., those who are – to some extent – aware of their self-control problems) may be more willing to adopt strategies to reduce the opportunities in which cheating is an option or to buy commitment devices to help them to stick to their original plans (O'Donoghue and Rabin, 1999). If cheating and present bias are correlated, the policymaker will deal with individuals i) who cannot correctly predict their future behavior and ii) for which ex-ante preferences do not correspond to their normative preferences and would require quite a different approach compared to the case in which cheating is associated to time consistent behavior (O'Donoghue & Rabin, 2006).

Besides behavioral predictions on the two main explanatory variables of the analysis, we propose conjectures on both participants' IQ and self-confidence. Specifically, in line with prior studies (e.g., Alan et al., 2020), we expect subjects with a higher IQ to undertake, on average, more cheating. Moreover, we expect overconfidence (excessive self-confidence) to correlate positively with cheating. For example, in our task, an overconfident subject may think of being above average in correctly answering to the trivia quiz and be more tempted to think that he/she knew the answer in the moment of self-assessment. In this sense, overconfident individuals use cheating in the attempt of maintaining a positive self-image.

3.4. Experimental procedures

We conducted the experiment using z-Tree (Fischbacher, 2007) at the BLESS experimental laboratory of the University of Bologna, Italy. Subjects were undergraduate students from different fields of study recruited using the ORSEE software (Greiner, 2015). In June 2018, 178 subjects participated, divided into 6 sessions of up to 30 subjects per session. All sessions were run in a between-subjects design. Subjects were randomly assigned to the sessions. Each session was about 60 minutes, and the average and median payment was 17 EUR, ranging from 10.5 to 26.5 EUR, including a show-up and participation fee of 2.5 EUR.

One part between the survey (C) and the time preference elicitation task (T) was randomly selected for payment at the end of the experiment to avoid wealth effects. If the elicitation task was chosen for payment, one of the 24 choices was randomly selected, and subjects received the preferred option. This means they may receive immediate payments (by cash), delayed payments (by wire transfer), or a combination of the two. One word of caution should be used in this respect. All students must have a bank account to interact with the university. Usually, they also have familiarity with online banking and debit cards, which makes a bank account as liquid as cash. Therefore, the difference between the two available payment methods (cash and wire transfers) may be weak for them. However, we know that a few subjects may choose between the two options based on their preference for the payment method.

We minimized any contact with the experimenter, which fully preserved anonymity. For this purpose, we used the double-anonymous payout system, as indicated in Barmettler et al. (2012): subjects were associated with an ID that they found on a sealed envelope when entering the laboratory and which was different from the number of the desk they used in the lab. The use of this ID mediated all their interactions. A new set of IDs was generated at the beginning of each session and then destroyed after the payment had been completed so that it was impossible to trace back the identity of participants and the ID used. Participants received their earnings in a sealed envelope identified by the ID if the survey was selected for payment. Even in the case of wire transfer, our experimental procedures allow us to preserve complete anonymity, given that, after the session, the experimenters could not merge the ID assigned and the identity of the participants. Notice that our experimental procedures made it very clear that the contact with the experimenter was minimized, this way limiting concern about delayed consequences associated with their cheating behavior.

3.5. Summary statistics

One of the 178 potential participants skipped some answers in the paper-and-pencil questionnaire. Moreover, we do not have estimates on time and risk preferences for seven subjects. These individuals reported inconsistent responses in the time discounting task, indicating the most present-oriented option almost constantly and the most future-oriented option occasionally. We also exclude five observations reporting outlier values for one or more preference parameters: first, we remove one with a massive risk aversion value (above 20); second, we remove four observations with present bias more than two standard deviations higher than the maximum value. Episodes of inconsistent or weird answers commonly arise in tasks that measure time preferences (Andreoni and Sprenger, 2012). Our final sample consists of 165 observations that we use throughout the analysis.

Table 1 reports summary statistics on the variables under investigation. The critical variable is the self-reported number of correct answers that, on average, is 7.54 out of 10 (with a median of 8, i.e., the minimum level necessary to get the reward). We split the remaining variables into explanatory and control variables. The explanatory variables are time and risk preferences, IQ test, and overconfidence, i.e., the difference between the participant's predicted and actual outcome in the IQ test).

Table 1. Summary statistics (*N*=165)

	Mean	**Std. Dev.**	**Min.**	**Max.**
Correct answers	7.539	2.199	0	10
Discount factor	0.930	0.250	0.000	1.054
Present bias parameter	0.801	0.328	0.000	1.565
Risk aversion parameter	0.912	3.091	0.000	12.423
IQ test	5.788	1.814	0	9
Overconfidence	0.806	1.984	-4	7
Female (d)	0.545	0.499	0	1
Age	23.303	2.473	19	36
Born abroad (d)	0.103	0.305	0	1
Siblings (d)	0.915	0.280	0	1
Low income (d)	0.200	0.401	0	1
High income (d)	0.200	0.401	0	1
Cheating task first (d)	0.497	0.502	0	1

Note. (d) indicates that the variable is a dummy.

We observe, on average, a risk aversion parameter of 0.91, a discount factor of 0.93, significantly different from 1 (t-test: -3.60; p-value <0.001), and a present bias parameter equal to 0.801, also significantly different from 1 (t-test: -7.77; p-value <0.001), with 78% of the subjects identified as present-biased (that is, they are associated with a quasi-hyperbolic discount factor below 1).[2] Average statistics align with Andreoni and Sprenger (2012), except

[2] Repeating our benchmark analysis with a dummy equal to one if the present bias parameter is below one, rather than the point estimate, confirms our findings. Results are available upon request.

that they find no present bias with monetary rewards.[3] Similar results are documented by Augenblick et al. (2015) with effort choices. Our results align with the micro-based literature on risk aversion and time preferences (for a review, see, respectively, Dohmen et al., 2011; Cohen et al., 2020). Moreover, on average, subjects respond correctly to 5.79 out of the 9 questions in the IQ test, and they tend to overestimate their performance by 0.81 more correct answers. We see this variable as a proxy for overconfidence.

The control variables are standard socio-demographics. The average subject is female (in 55% of the cases), 23 years old (with age comprised between 19 and 36), born in Italy, has siblings (in 92% of the cases), and comes from a middle-income household. We define low-income and high-income households as earning respectively up to 1,500 euros and at least 4,000 euros per month, net of taxes. These two categories absorb the top and bottom 20% of the subjects. As a further control variable, we consider the order of the tasks in the experiment, with a dummy equal to one when the survey with cheating opportunity comes first (sequence CT) and zero after the time discounting task (sequence TC).

Figure 1 plots the distribution of the self-reported number of correct answers. It turns out that 69% of the participants declare 8 or more correct answers, this way earning the monetary reward. The distributions for both the non-incentivized and the incentivized pilot studies are shown in Appendix Figures A2a and A2b. They are different from the one reproduced in Figure 1. This evidence is statistically supported by a Kolmogorov-Smirnov test of equality of the distributions.[4]

3 More specifically, their aggregate annual discount rate is 0.30, which they acknowledge is lower than the one estimated by most other researchers, and their aggregate curvature is estimated at 0.92. As evidenced by Cohen et al. (2020), time preference estimates are characterized by high variance.

4 The Kolmogorov-Smirnov test reports the following values: Distributions in Figure A2a and Figure 1: 0.827 (p-value <0.001). Distributions in Figure A2b and Figure 1: 0.776 (p-value <0.001). Distributions in Figure A2a and Figure A2b: 0.148 (p-value 0.189).

Figure 1. Distribution of correct answers (*N*=165)

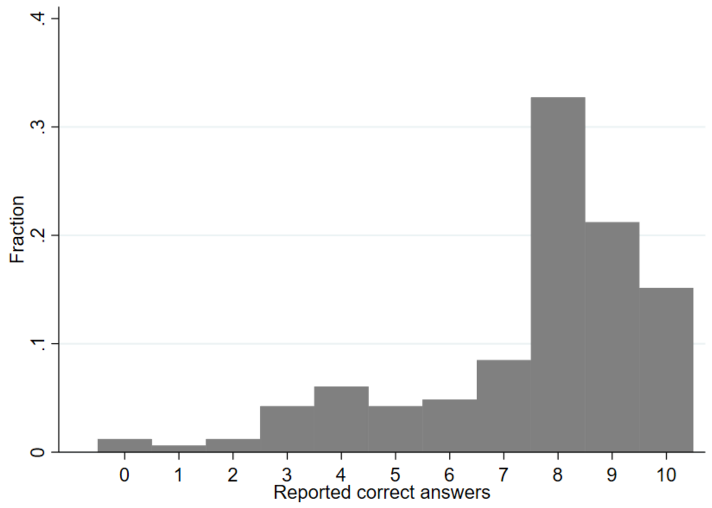

We label as "no earners" those who declared answers between 0 and 7 and as "earners" those who declared answers between 8 and 10. Following the validation of our cheating task, we refer to earning as cheating. Based on our preliminary evidence, it is doubtful that subjects who reported 7 or 6 did not cheat and answered correctly so many questions. We believe there are two main reasons for cheating in this setting: the *monetary reward* (I get extra money if I report 8 or more) and the *intrinsic or self-image motivation* (I feel better if I report a higher number and I want to maintain a positive image of myself). One may argue that cheating in our setting is due to liquidity constraints. However, this is unlikely given the high sample homogeneity in terms of liquidity (University students) and since we control for self-reported income in the econometric analysis (the income coefficient is never significant). Subjects who report 8 or more may have a combination of the two reasons for cheating, while subjects who report 7 or fewer cheat only for their intrinsic motivation and to signal that they did not want to cheat at the expense of the community. Our setting makes only the monetary reward salient, and consistently, we pay attention to the cheating made for monetary rewards. However, intrinsic motivation may explain why some subjects report 9 or 10. These subjects may display a more sophisticated cheating pattern. For instance, one may argue that reporting 8 may appear suspicious, representing the threshold for earning extra money. Therefore, reporting a 9 or 10 might be an attempt to disguise the lie and set apart from those more likely to be considered cheaters.[5]

Table 2 compares the average of the observed explanatory and control variables in the two groups of earners and non-earners; the last column shows results from a Wilcoxon non-parametric rank-sum test on the equality of the averages in the two groups. The two groups are very similar, as we find evidence of a difference significant at 5% only in overconfidence, which is higher among earners. Notably, the two groups do not differ in terms of socio-demographics and, overall, there is no effect of the sequence in which the two tasks are presented, that is, of the delayed consequences associated with the decision to cheat. In Section 4, we control for other characteristics of subjects in a multivariate setting.

Table 2. Earners and non-earners (means)

	Non-earners	**Earners**	**Test**
Discount factor	0.959	0.917	1.760
Present bias parameter	0.865	0.773	1.679
Risk aversion parameter	0.559	1.070	-1.055
IQ test	5.882	5.746	0.145
Overconfidence	0.275	1.044	-2.119*
Female (d)	0.569	0.535	0.399
Age	23.314	23.298	-0.162
Born abroad (d)	0.078	0.114	-0.693
Siblings (d)	0.882	0.930	-1.008
Low income (d)	0.235	0.184	0.756
High income (d)	0.176	0.211	-0.504
Cheating task first (d)	0.549	0.474	0.892
Observations	51	114	

*Note. (Non-)Earners are those individuals who report an outcome of 8 or (less)more correct answers in the TRIVIA quiz. The test is a Wilcoxon rank-sum test on the equality of the mean in the two groups. (d) indicates that the variable is a dummy. *p<0.05.*

5 We also believe the case of those who reported 7 correct answers in our incentivized task is particularly interesting. They probably lied, but they restrained themselves from lying more and earning monetary rewards. In their case, guilt may have prevailed over the potential benefit of declaring 8 or more correct answers.

4. Results

This section performs multivariate non-linear regression analyses to study the correlation between the self-reported number of correct answers and our explanatory variables. All the models use standard errors clustered by experimental session to account for potential session-specific features.[6] To ease the interpretation of the coefficients, the three preference parameters (i.e., the only continuous variables in the model) are included in the specifications in a standardized version (that is, with mean 0 and standard deviation 1). Their marginal effects then measure variations of the probability of cheating resulting from a standard deviation change in the preference.

Table 3 shows average marginal effects from probit regressions where the dependent variable is a dummy equal to one if the self-reported number of correct answers is 8 or more and zero otherwise. In so doing, we compare subjects who cheated for monetary incentives with those who did not cheat.

Column 1 includes only the preference parameters in the specification. We find a significantly negative effect for the present bias parameter, supporting Hypothesis 2. For this reason, we do not discuss it further in this section. The negative effect of the present bias parameter means that subjects with a more pronounced bias toward the present are more likely to cheat. Specifically, one standard deviation drop in the present bias parameter reduces the probability by 6.7%. Interestingly, we find no significant effect for the discount factor (this way failing to support Hypothesis 1), suggesting that short-sighted or forward-looking does not correlate with the probability of cheating. This result may be partially explained by the fact that our experimental task seems better suited to elicit impulsive cheating, as discussed in the Conclusion.

[6] For instance, subjects may self-select into the sessions. In our data we find significant differences in terms of some demographic variables (age, born abroad, siblings, high income). Our findings are preserved if we alternatively use standard errors clustered at the experiment day level. Results are available upon request.

Table 3. Probability to cheat (average marginal effects)

Dependent variable	(1) **Prob. 8-10**	(2) **Prob. 8-10**	(3) **Prob. 8-10**	(4) **Prob. 8-10**
Discount factor	-0.430	-0.472	-0.692	
	(0.683)	(0.727)	(0.634)	
Present bias parameter	-0.067***	-0.054*	-0.065*	
	(0.021)	(0.023)	(0.027)	
Risk aversion parameter	-0.414	-0.442	-0.640	
	(0.707)	(0.747)	(0.642)	
IQ test		0.027*	0.016	
		(0.013)	(0.014)	
Overconfidence		0.053***	0.060***	
		(0.012)	(0.015)	
Discount factor, males				-1.717*
				(0.839)
Present bias, males				-0.103
				(0.060)
Risk aversion, males				-1.682*
				(0.845)
IQ test, males				0.042*
				(0.020)
Overconfidence, males				0.068*
				(0.024)
Discount factor, females				0.250
				(0.583)
Present bias, females				-0.066*
				(0.026)
Risk aversion, females				0.443
				(0.479)
IQ test, females				-0.005
				(0.013)
Overconfidence, females				0.057*
				(0.027)
Female			0.020	-7.730
			(0.071)	(4.637)
Age			0.006	0.009
			(0.019)	(0.019)
Born abroad			0.170***	0.186***
			(0.050)	(0.054)
Siblings			0.181	0.180
			(0.115)	(0.116)
Low income			-0.151	-0.140
			(0.122)	(0.139)
High income			0.046	0.060
			(0.111)	(0.100)
Cheating task first			-0.104	-0.101
			(0.073)	(0.068)
Log-likelihood	-100.066	-96.946	-92.766	-89.282
Pseudo R^2	0.019	0.050	0.091	0.125
Observations	165	165	165	165

*Note. Probit model. Standard errors clustered by experimental session in parentheses. *** p<0.001, ** p<0.01, * p<0.05.*

Column 2 adds IQ test and overconfidence to the specification. This does not modify our previous findings, and in addition, we now obtain positive effects on IQ and being overconfident. To the best of our knowledge, no evidence supports an association between overconfidence and cheating, except for Adams et al. (2018), who find that cheaters have higher stated beliefs in their ability. Note that our measure of overconfidence is obtained from the IQ test. Still, overconfidence may explain cheating in our experiment since overconfident participants may – in the moment of the self-assessment – convince themselves that they knew the right answers, and therefore ex-post they tend to report more correct answers to maintain their self-image as high performers.

Column 3 adds control variables to the specification. Previous results are mostly confirmed, except for the IQ test, which is no longer significant. The control variables are never significant, apart from being born abroad, which shows a positive effect (+17%). It is noteworthy that, among the control variables, the coefficient on the ordering of the tasks in the experiment is not significant. This means that the decision to report a high number of correct answers is influenced neither by an earlier task focusing on time preferences nor by monetary rewards obtained in earlier stages. Since we find no order effect, we conclude that what we observe is likely to capture the tendency of impatient people to cheat more.[7]

To summarize, we find that cheating correlates with overconfidence and present bias, suggesting that individuals more likely to exhibit self-control problems are also more likely to cheat. We thus find support for Hypothesis 2, while we find no support for Hypothesis 1 on the effect of the discount factor.

4.1 The moderating role of gender and further robustness checks

In an ex-post analysis, we explore the moderating role of gender on our key dimensions. The motivation is that we find large differences in the descriptive statistics, with females being more future-oriented (in terms of both discount factor and present bias) and less self-confident. Therefore, even if we did not find a significant effect for being female in the previous analyses, it could be that gender plays a role through other dimensions. For this reason, in Column 4 of

[7] We wish to thank an anonymous referee for highlighting this issue.

Table 3, we repeat the analysis in Column 3 but replace the variables on preferences, IQ test, and overconfidence with their interactions with gender. We introduce one set of interactions with being male and one set of interactions with being female. This specification is as informative as the one in Column 3 with the inclusion of one single set of interactions but has the advantage of directly providing separate effects on males and females.

Estimates show interesting results, which we discuss considering the existing evidence. First, the present bias parameter acts only through females displaying a negative association with the likelihood of cheating. This difference may be related to the overall finding that females show greater patience than males (McLeish & Oxoby, 2007; Castillo et al., 2011), with impatient females representing an exception concerning the average trait in their group. Second, the other two preference parameters (the discount factor and risk aversion) negatively affect only the males' probability of cheating. This evidence, which was not observed in the general models pooling the effects of males and females, may reflect our male sample's larger variability of preferences.[8] In the literature, Bucciol et al. (2013) and Hübler et al. (2018) already found a negative association between cheating and risk aversion, irrespective of gender.

In general, our benchmark results of Table 3 hold when using linear probability models (OLS) instead of probit models (see Appendix Table A1) and when considering three variants of the regression specification. Specifically, in one case we add the score of the self-control scale (Tangney et al. 2004); see Appendix Table A2. The score displays no significant coefficient, possibly because our sample of university students shows a slight variation in this measure.

In the other two cases, we change the definition of the preference variables. We first replace the preference variables with ordinal variables equal to 1, 2, …, and 10 to indicate the decile of their sample distribution; see Appendix Table A3. We run this exercise because preference variables may be estimated with error. We then replace the preference variables with two non-parametric measures derived from the CTB task; see Appendix Table A4. One measure, which we label “beta”, is constructed as the average later reward in choices 13-24 of the task, while

[8] The standard deviation of both parameters is 72% higher in males than in females. Discount factor: 0.310 for males and 0.180 for females; Risk aversion: 3.838 for males and 2.233 for females.

the other, that we label "delta", is constructed as the difference between the average later reward in choices 1-12 of the task and the average later reward in choices 13-24. This approach allows us to account for the observations that were excluded from the analysis because of inconsistent or outlying behavior.

4.2 Definition of cheating

Table 4 replicates the specification in Column 3 of Table 3 using an alternative dependent variable. The new variable is ordinal and distinguishes the group of subjects who did not cheat for the monetary reward (by reporting less than 8 correct answers) and, separately, each of the groups reporting 8, 9, and 10 correct answers. The purpose of looking at these four groups rather than just two and combining all those individuals receiving the monetary reward is to pay special attention to acts of cheating that are not solely motivated by monetary gains. This represents an ex-post analysis, for which we did not have a priori hypotheses but that we deem worth investigating further, given our results.

The monetary reward is granted whenever the subject reports at least 8 correct answers. There is no further reward for reporting 9 or 10 correct answers, presumably farther from the truth. Still, many individuals report 9 and 10. Table 4 then shows results from an ordered probit regression, where the three columns indicate average marginal effects on the probability that the self-reported outcome is 8, 9, or 10, respectively. Compared to Table 3, we find fewer significant effects because each outcome of the dependent variable is associated with a smaller number of observations. However, we still find significant present bias and overconfidence effects, but only when the outcome is 9 or more. The effect is also growing with the self-reported outcome. For instance, one standard deviation rise in present bias increases the probability of reporting 9 by 3.59% and 10 by 5.85%. It then seems that only those who commit severe cheating differ from the others. This is interesting evidence, as the monetary incentives of reporting 8 or more than 8 answers are identical in our setting. It could be that those reporting 9 or 10 correct answers cheat more sophisticatedly than those reporting 8 correct answers, i.e., the threshold to receive the incentive. Their higher reporting might be an attempt to disguise the lie and separate them from those more likely to be considered cheaters. Unfortunately, we have no data to dig deeper in this direction, leaving it to future research.

In a robustness check reported in Appendix Table A5, we consider further definitions of cheating. We consider two alternative cases. In one case, we still take a dummy variable equal to 1 when the subject is considered a cheater, but we define cheaters as all those reporting 6 or more correct answers or 7 or more correct answers. One could view these two variables as trying to identify all likely cheaters (no matter for monetary rewards or not). In the other case, inspired by Rahwan et al. (2018) and our pilot studies without cheating, we group individuals into three categories based on their answers: certainly honest (0-4 reported correct answers), likely dishonest (5-7 answers) and certainly dishonest (8-10 answers). The idea is that, in addition to those who received a monetary reward, another set of individuals reported a relatively high number of correct answers, likely cheating for other reasons (e.g., self-image; see on this Thielmann and Hilbig, 2019).

In the models of Appendix Table A5, we no longer observe a significant effect of the present bias. In contrast, we systematically find strong effects of overconfidence which – interestingly – seems to have a different impact depending on whether cheating is also motivated by monetary rewards. Specifically, overconfidence negatively impacts the likelihood of cheating for those classified as likely dishonest. However, the opposite is true for those individuals classified as certainly dishonest, with overconfident being positive and significant. Overconfidence thus seems to play a crucial role in explaining dishonest behavior.

Based on this evidence, we make the following speculations on cheating not necessarily being driven by monetary rewards. First, the effect of the present bias is not as straightforward since it seems highly dependent on whether the cheater is aiming at obtaining monetary rewards. Our results show that there is not only one single motivation for the observed cheating pattern. Maintaining a favorable self-concept seems to be one of the relevant motivations underlying the evidence stemming from our experiment. In line with Fischbacher and Föllmi-Heusi (2013, p.542), "the phenomenon of partial lying is robust" and deserves more attention in future studies.

Table 4. Probability to cheat by intensity of cheating (average marginal effects x100)

Dependent variable	(1) **Prob =8**	(2) **Prob =9**	(3) **Prob =10**
Discount factor	4.902	-29.340	-47.882
	(4.541)	(19.127)	(33.522)
Present bias parameter	0.599	-3.586*	-5.852**
	(0.592)	(1.560)	(2.129)
Risk aversion parameter	4.665	-27.921	-45.567
	(4.308)	(19.026)	(32.542)
IQ test	0.024	-0.141	-0.230
	(0.052)	(0.374)	(0.568)
Overconfidence	-0.189	1.129**	1.842***
	(0.154)	(0.392)	(0.511)
Female	0.292	-1.750	-2.856
	(0.490)	(2.670)	(4.459)
Age	-0.105	0.629	1.026
	(0.099)	(0.394)	(0.729)
Born abroad	-0.888	5.316*	8.675**
	(0.629)	(2.691)	(3.314)
Siblings	-1.280	7.6<61	12.502
	(1.195)	(3.986)	(9.242)
Low income	1.079	-6.460*	-10.542
	(0.662)	(2.789)	(5.578)
High income	0.169	-1.011	-1.649
	(0.754)	(4.114)	(7.186)
Cheating task first	0.809	-4.843*	-7.904
	(0.975)	(2.090)	(4.428)
Log-likelihood	-208.928	-208.928	-208.928
Pseudo R^2	0.057	0.057	0.057
Observations	165	165	165

*Note. Ordered probit model. Standard errors clustered by experimental session in parentheses. *** p<0.001, ** p<0.01, * p<0.05.*

5. Conclusions

Research on cheating is growing fast, but much still has to be learned about the mechanisms behind cheating and the relationship between cheating and other domains of individual preferences, such as the moral one. Our research contributes to the literature by studying the correlation with time preferences. Specifically, we ran a laboratory experiment to explore the correlation between cheating and time preferences. In doing this, building on previous literature, we devised and validated a cheating task that allows us to obtain a proxy for cheating at the individual level with higher precision than in previous studies. Moreover, the fact that cheating in our experiment arises from self-reporting contributes to enhancing the external validity of our results. Indeed, many cheating opportunities in everyday life are also connected to self-reporting circumstances, characterized by a very low risk of getting caught cheating (one can think for instance of tax declarations, health insurance questionnaires and migration forms, among others). We found evidence that cheating is more frequent among individuals who exhibit a present bias. In addition, cheating is more likely among overconfident individuals. Policymakers then must carefully design incentives to discourage cheating, considering that it seems associated with time inconsistency.

We see three main limitations to this study, which call for future research. First, in our task, individuals may cheat to maintain a positive self-image because the performance in the task is driven by ability, i.e., knowledge or culture, rather than luck, as in the coin flip task. While we cannot exclude this possibility, we expect the confounding factor represented by self-image to be negligible based on Hugh-Jones (2016), who compared our task with a coin flip task on the same subject pool, finding that the two tasks are highly correlated. Second, we cannot control whether our cheating task frustrated the participants, who may realize it is tough to answer our multiple-response questions. Some may be irritated or infer that we want them to cheat, eventually deciding to cheat. Future research could disentangle "genuine" cheating from cheating due to this frustration, for instance, by asking an ad-hoc question in an ex-post questionnaire or eliciting incentivized beliefs about the average number of questions other participants can answer correctly. Third, our study – based on two separate tasks for revealing cheating and time preferences – is correlational. Future research may aim to unify the tasks to investigate a causal link between present bias and cheating and whether individuals aware of their self-control problems would be more likely to adopt strategic behaviors to avoid situations in which they may be tempted to cheat.

Since our results support only correlational rather than causality claims, policy implications cannot be conclusive. However, based on our results showing that cheating is more widespread among individuals who attribute more importance to the present, we argue that general goals such as increasing the severity of punishments, magnifying the future costs of cheating, shortening the length of trials or discovering crimes faster may be ineffective. The reason is that individuals with present bias might still be more influenced by the immediate benefits of cheating. Better strategies should focus on preventive measures and interventions targeting the underlying biases and motivations, implementing educational programs to raise awareness about the consequences of cheating, and fostering a culture of integrity and ethical behavior. By understanding and addressing the root causes of present bias, it would be possible to create more meaningful and impactful strategies to deter cheating in the long run. For example, Alan et al. (2021) show that a school intervention promoting a forward-looking perspective in children boosts their self-control. According to our findings, we conjecture that the effect of the school intervention might also spill over to reduced cheating rates among these children.

Ex-post analysis on gender shows that the present bias seems to matter on both males and females. In contrast, only males are influenced by the discount factor (negatively), risk aversion (negatively), and IQ test (positively). This evidence is exciting and opens new avenues for research on gender and cheating, where it is frequently found that females cheat less regularly (e.g., see the review in Jacobsen et al., 2018). We also found evidence that the correlation between cheating, overconfidence, and the present bias is more pronounced when the intensity of cheating is higher. This evidence, however, needs to be explored more in detail, employing an experimental design specifically aimed at detecting these mechanisms. We leave it as an attractive future avenue of research. Future research could also explore the role of beliefs about the task's difficulty, for instance, by eliciting the average number of correct answers to the trivia quiz from other participants. In such a manner, it would be possible to control for the presence of frustration or anger mechanisms; if one considers the trivia quiz too difficult, that may somewhat motivate cheating.

References

Adams, N.R., Kuhn, M.A., & Waddell, G.R. (2018). Confidence and contrition: Is cheating internalized in performance assessments? Mimeo, University of Oregon.

Alan, S., Ertac, S., & Gumren, M. (2020). Cheating and incentives in a performance context: Evidence from a field experiment on children. *Journal of Economic Behavior & Organization*, *179*, 681-701.

Alan, S., Ertac, S., & Gumus, I. (2021). Does a forward-looking perspective affect self-control and the demand for commitment? Results from an educational intervention. *Economic Inquiry*, *59*(4), 1533-1546.

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

Andreoni, J., Kuhn, M.A., & Sprenger, C. (2015). Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior & Organization*, *116*, 451-464.

Andreoni, J., & Sprenger, C. (2012). Estimating time preferences from convex budgets. *American Economic Review, 102*(7), 3333-56.

Augenblick, N., Niederle, M., & Sprenger, C. (2015). Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics*, *130*(3), 1067-1115.

Barmettler, F., Fehr, E., & Zehnder, C. (2012). Big experimenter is watching you! Anonymity and prosocial behavior in the laboratory. *Games and Economic Behavior*, *75*(1), 17-34.

Bilker, W.B., Hansen, J.A., Brensinger, C.M., Richard, J., Gur, R.E., & Gur, R.C. (2012). Development of abbreviated nine-item forms of the Raven's standard progressive matrices test. *Assessment*, *19*(3), 354-369.

Bortolotti, S., Kölle, F., & Wenner, L. (2022). On the Persistence of Dishonesty. *Journal of Economic Behavior & Organization, 200*, 1053-1065.

Bucciol, A., & Montinari, N. (2019). *Dishonesty in Behavioral Economics*, Amsterdam: Elsevier.

Bucciol, A., & Piovesan, M. (2011). Luck or cheating? A field experiment on honesty with children. *Journal of Economic Psychology*, *32*(1), 73-78.

Bucciol, A., Landini, F., & Piovesan, M. (2013). Unethical behavior in the field: Demographic characteristics and beliefs of the cheater. *Journal of Economic Behavior & Organization*, *93*, 248-257.

Castillo, M., Ferraro, P.J., Jordan, J.L., & Petrie, R. (2011). The today and tomorrow of kids: Time preferences and educational outcomes of children. *Journal of Public Economics*, *95*(11-12), 1377-1385.

Cohen, J., Ericson, K. M., Laibson, D., & White, J. M. (2020). Measuring time preferences. *Journal of Economic Literature*, *58*(2), 299-347.

Dai, Z., Galeotti, F., & Villeval, M. C. (2018). Cheating in the lab predicts fraud in the field: An experiment in public transportation. *Management Science*, *64*(3), 1081-1100.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. (2011). Individual risk attitudes: Measurement, determinants and behavioral consequences. *Journal of the European Economic Association*, *9*, 522-550.

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

Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, *11*(3), 525-547.

Giné, X., Goldberg, J., Silverman, D., & Yang, D. (2018). Revising commitments: Field evidence on the adjustment of prior choices. *The Economic Journal*, *128*(608), 159-188.

Gino, F., Gu, J., & Zhong, C.B. (2009). Contagion or restitution? When bad apples can motivate ethical behavior. *Journal of Experimental Social Psychology*, *45*(6), 1299-1302.

Gino, F., Schweitzer, M.E., Mead, N.L., & Ariely, D. (2011). Unable to resist temptation: How self-control depletion promotes unethical behavior. *Organizational Behavior and Human Decision Processes*, *115*(2), 191-203.

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

Hoch, S.J., & Loewenstein, G.F. (1991). Time-inconsistent preferences and consumer self-control. *Journal of Consumer Research*, *17*(4), 492-507.

Hübler, O., Menkhoff, L., & Schmidt, U. (2018). Who is cheating? The role of attendants, risk aversion, and affluence. DIW Berlin Discussion Paper No. 1736.

Hugh-Jones, D. (2016). Honesty, beliefs about honesty, and economic growth in 15 countries. *Journal of Economic Behavior & Organization*, *127*, 99-114.

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

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, *112*(2), 443-478.

Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, *45*(6), 633-644.

McLeish, K.N., & Oxoby, R.J. (2007). Gender, affect and intertemporal consistency: An experimental approach. IZA Discussion Paper No. 2663.

Mead, N.L., Baumeister, R.F., Gino, F., Schweitzer, M.E., & Ariely, D. (2009). Too tired to tell the truth: Self-control resource depletion and dishonesty. *Journal of Experimental Social Psychology*, *45*(3), 594-597.

Nagin, D.S., & Pogarsky, G. (2003). An experimental investigation of deterrence: Cheating, self-serving bias, and impulsivity. *Criminology*, *41*(1), 167-194.

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review*, *89*(1), 103-124.

O'Donoghue, T., & Rabin, M. (2006). Incentives and self-control. In Blundell, R., Newey, W., & Persson, T. [Eds.], *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress* (Econometric Society Monographs, 215-245). Cambridge: Cambridge University Press.

Rahwan, Z., Hauser, O. P., Kochanowska, E., & Fasolo, B. (2018). High stakes: A little more cheating, a lot less charity. *Journal of Economic Behavior & Organization*, *152*, 276-295.

Ruffle, B.J., & Tobol, Y. (2014). Honest on Mondays: Honesty and the temporal separation between decisions and payoffs. *European Economic Review*, *65*, 126-135.

Strotz, R.H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, *23*(3), 165-180.

Tangney, J.P., Baumeister, R.F. & Boone, A.L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of Personality, 72*(2), 271-324.

Thielmann, I., & Hilbig, B. E. (2019). No gain without pain: The psychological costs of dishonesty. *Journal of Economic Psychology*, *71*, 126-137.

Vranka, M., Frollová, N., Pour, M., Novakova, J., & Houdek, P. (2019). Cheating customers in grocery stores: A field study on dishonesty. *Journal of Behavioral and Experimental Economics*, 83, 101484.

Yaniv, G., Greenberg, D., & Siniver, E. (2019). Telling an impossible lie: Detecting individual cheating in a die-under-the-cup task. *Review of Behavioral Economics*, *6*(2), 119-131.