

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Barron, Kai

Working Paper Lying to appear honest

WZB Discussion Paper, No. SP II 2019-307

Provided in Cooperation with: WZB Berlin Social Science Center

Suggested Citation: Barron, Kai (2019) : Lying to appear honest, WZB Discussion Paper, No. SP II 2019-307, Wissenschaftszentrum Berlin für Sozialforschung (WZB), Berlin

This Version is available at: https://hdl.handle.net/10419/204826

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU





Kai Barron

Lying to appear honest

Discussion Paper

SP II 2019–307 October 2019

Research Area Markets and Choice

Research Unit **Economics of Change** Wissenschaftszentrum Berlin für Sozialforschung gGmbH Reichpietschufer 50 10785 Berlin Germany www.wzb.eu

Copyright remains with the authors.

Discussion papers of the WZB serve to disseminate the research results of work in progress prior to publication to encourage the exchange of ideas and academic debate. Inclusion of a paper in the discussion paper series does not constitute publication and should not limit publication in any other venue. The discussion papers published by the WZB represent the views of the respective author(s) and not of the institute as a whole.

Affiliation of the authors:

Kai Barron, WZB (kai.barron@wzb.eu) Abstract

Lying to appear honest*

This paper studies individual truth-telling behavior in the presence of multiple lying opportunities with heterogeneous stake sizes. The results show that individuals lie downwards (i.e. forgo money due to their lie) in low-stakes situations in order to signal honesty, and thereby mitigate the image repercussions of upward lying in high-stakes contexts. This constitutes first evidence of systematic downward lying in an unobserved lying game. The observed behavior is consistent with the spirit, but not the letter, of the prominent models of lying behavior. It therefore presents a challenge for these models.

Keywords: Private information, honesty, truth-telling, lying, image, reputation

JEL classification: C91; D83; H26; Z13

^{*} The author would like to thank Philipp Albert, Tilman Fries, Steffen Huck, Heather Jacklin, Joel Sobel, Agne Kajackaite, Robert Stüber, and Roel van Veldhuizen for helpful comments and suggestions.

1 Introduction

Opportunities to lie about one's private information for personal gain are pervasive in everyday life. Often these lying opportunities do not constitute once-off interactions, but rather one instance from many interactions with a particular audience. These interactions may vary widely in terms of the magnitude of the gains from lying. Therefore, individuals may be strategic in deciding precisely when and how to lie, resulting in patterns of behavior that differ from those observed in one-shot lying interactions.

In this project, I study individual lying behavior in the presence of multiple lying opportunities with heterogenous stake sizes. The central question asked in this paper is whether individuals give up money when making low stakes reports to try to signal honesty, with the goal of shrouding their (lucrative) lying in high stakes environments. This hypothesized behavior is in stark contrast to the predictions of the fully rational model, which predicts that individuals will maximize their earnings. Furthermore, while existing concepts, such as moral licensing (Monin and Miller, 2001), as well as the prominent lying models (e.g. Gneezy et al. (2018), Dufwenberg and Dufwenberg (2018), Abeler et al. (2019) and Khalmetski and Sliwka (2019)) can explain being honest in the low stakes scenarios along with upward lying in the high stakes scenarios, here I show that a substantial fraction of subjects actually lie downwards¹ in the low stakes scenarios, thereby giving up money to signal honesty.² This constitutes first evidence of substantial downward lying in the lying literature.³ While the observed behavior is consistent with the spirit of these prominent lying models (since reputation concerns play an important role), it is not easy to reconcile with the specific details of these models (in terms of the precise way in which reputation concerns operate).

This paper reports the results from an experiment that builds on the classic Fischbacher and Föllmi-Heusi (2013) die-rolling paradigm in which subjects roll a single die in private and can misreport the number they observe in order to increase their earnings. Here, I extend this paradigm by allowing for multiple lying opportunities with heterogeneous stake sizes. This is achieved by adding a second die. In this TwoDICE treatment, subjects roll two fair six-sided dice (a blue and a red die) simultaneously, and are instructed to report the number observed on each die. One report is a *High Stakes* report, which is multiplied by $\in 2$, and the other is a *Low Stakes* report, which is multiplied by

¹I follow the literature in using the concept *downward lying* to refer to a scenario where an individual makes a report that they know to be false which results in them receiving a lower payment than they would have received from reporting truthfully (see, e.g., Definition 4 in Abeler et al. (2019)).

²To be precise, some of these models (e.g. the Abeler et al. (2019) *Reputation for Honesty* + *LC utility* model) can sustain some downward lying in unobserved lying games, but not the to the degree observed in my data. This will be discussed in more detail below.

³More specifically, it constitutes first evidence of downward lying in an unobserved lying game – i.e. a lying game in which the experimenter does not observe the true state.

€ 0.1 in contributing to their payment. Prior to rolling the dice, subjects must choose whether the blue or the red die will be their *High Stakes* die. As in Fischbacher and Föllmi-Heusi (2013), a report of 6 always pays € 0, and in the text below I will refer to it as a report of 0.

The experiment is designed to provide a simple environment for studying the relationship between decisions made under heterogeneous stake sizes – i.e. between big (high stakes) and little (low stakes) lying decisions with the same audience. In this context, I document two main findings. Firstly, I find evidence of downward lying in the *Low Stakes* reports, with 63% of subjects reporting having seen a low number (0, 1 or 2) on the *Low Stakes* die. Secondly, I show that high reports on the *High Stakes* die are observed in conjunction with low reports on the *Low Stakes* die more often than expected.⁴ More specifically, in terms of the joint distribution of reports on both dice, 50% of subjects reported having seen both a high number (3, 4 or 5) on the *High Stakes* die, and a low number (0, 1 or 2) on the *Low Stakes* die. According to the objective probability distribution, this (high,low) report should occur in 25% of cases. Furthermore, this (high,low) joint report is chosen almost twice as often as the higher-paying (high,high) report, which is chosen by 26% of subjects. This suggests that subjects are willing to give up money *and* lie to make a (high,low) report. Furthermore, this indicates that subjects preferred to choose a low *Low Stakes* report to simply telling the truth for the *Low Stakes* die.

This paper contributes to a large and fast-growing literature that studies when and why people choose to misreport their private information (see, e.g., Gneezy (2005), Mazar et al. (2008), Coricelli et al. (2010), Shalvi et al. (2011), Fischbacher and Föllmi-Heusi (2013), Abeler et al. (2014), Cohn et al. (2014), Kajackaite and Gneezy (2017), Gneezy et al. (2018), Abeler et al. (2019)). This literature has generated an impressive body of evidence, revealing systematic regularities in lying behavior. More recently, there has been an effort to synthesize the lessons learnt into a parsimonious theoretical framework. While there have been numerous contributions to this effort, arguably a small set of these models has risen to prominance, namely those by Gneezy et al. (2018), Dufwenberg and Dufwenberg (2018), Abeler et al. (2019) and Khalmetski and Sliwka (2019). Illustrating this point, Abeler et al. (2019) use a Popperian approach to argue that other previously proposed models are less compatible with the comprehensive dataset analyzed in their metastudy.⁵ These remaining four papers all posit that individuals experience disutility from being perceived as a liar. Three of these models, those offered by Gneezy

⁴Here "more often than expected" means that the combination of a high *High Stakes* report and low *Low Stakes* is observed: (i) more often than expected according to the objective distribution, but also (ii) more often than expected under commonly used models of lying aversion, which tend not to predict substantial downward lying. This point is discussed in more detail in the text below.

⁵Abeler et al. (2019) also mention in Appendix C2 that Dufwenberg and Dufwenberg (2018) cannot fully explain the data collected in their metastudy, but do still give the ideas discussed by Dufwenberg and Dufwenberg (2018) prominence throughout their paper.

et al. (2018), Abeler et al. (2019), and Khalmetski and Sliwka (2019), are very similar in spirit, with all three proposing that individuals desire both *to be honest*, and also *to appear honest* (i.e. they have an intrinsic cost of lying, as well as a reputational cost due to the inference about their honesty drawn from the observed action). In contrast, Dufwenberg and Dufwenberg (2018) propose a model that depends on the audience's perception regarding the degree of over-reporting (rather than just the perception that the individual is lying or not).

While the experiment under consideration here is not designed to test these models, reconciling the observed results with the models does present a challenge. The reason for this is that the models do not provide a motive for systematic downward lying on the Low Stakes report. To see this intuitively, consider the two lying aversion components of the models: (i) an intrinsic cost of lying, and (ii) a reputational cost. Regarding (i), taking a subject's High Stakes report as given, if she has some form of intrinsic lying cost, she can do weakly better by telling the truth on the Low Stakes die than by lying downwards (i.e. she earns more money, and tells fewer lies). Regarding (ii), the models typically use an equilibrium concept, where subjects are Bayesian and the reputational cost of lying is a function of the expected fraction of liars making a particular report. This expected fraction is accurate in equilibrium. Since twice as many subjects report (high,low) compared to (high,high) in my experiment, a subject could both reduce their reputational cost and increase their earnings by switching from (high,low) to (high,high). To see this in a different way, low *Low Stakes* reports are made more often than high *Low* Stakes reports. Therefore, according to the models, low Low Stakes reports incur greater reputational costs and lower earnings in comparison to high Low Stakes reports. Yet in my data they are observed more often.⁶

One potential reason for the imperfect fit between these theories and my data is that the theories have been designed to organize stylized facts generated predominantly by studies considering lying on a single dimension, or lying by individuals facing repeated homogeneous lying opportunities (see, e.g., Jiang (2013), Abeler et al. (2014), and Cohn et al. (2015)). The models appear to be highly successful in achieving this objective (as evidenced by the comprehensive analysis and discussion in Abeler et al. (2019)). In contrast, however, my study considers a slightly more complex environment, which alters the optimization problem.

The objective of this paper is to contribute evidence towards developing a better un-

⁶An alternative way to see this is by means of Figure 4 in Appendix C, which is a one-dimensional representation of the full distribution of reports in the TwoDICE treatment, with the states ordered such that they are increasing in payoffs. Given the uniform objective density over states, the models should predict an upward sloping probability density function. However, it is non-trivial to explain the downward slope observed within the two right-most sections of the figure, separated by the vertical dotted lines (reflecting the groupings of *High Stakes* 4's and *High Stakes* 5's.

derstanding of lying in such more complex environments, with the aim of facilitating the enrichment of these existing models. In line with these models, the evidence presented below suggests that reputation concerns play a highly important role in governing lying decisions.

The remainder of the paper proceeds as follows. Section 2 describes the experimental design. Section 3 presents the results. Section 4 discusses the relationship to the existing literature, and Section 5 concludes.

2 Experimental Design

The experiment involves a one-shot individual decision-making task. I build on the experimental paradigm popularized by Fischbacher and Föllmi-Heusi (2013) in which subjects privately observe the outcome of a random variable, and are asked to report this private information. When subjects make this report, they already know the mapping from reports to payoffs. This provides them with an opportunity to misreport their private information and increase their payoffs. For comparability, I follow Fischbacher and Föllmi-Heusi (2013) in using a six-sided die as the randomization device. One appealing feature of this paradigm is that it has been shown to correlate well with cheating behavior outside the lab (see, e.g., Hanna and Wang (2017), Dai et al. (2018), and Cohn and Maréchal (2018)).

Each subject in the study participated in one of two treatment conditions. The BASE-LINE condition is similar to the classic Fischbacher and Föllmi-Heusi (2013) game, providing a benchmark measure of lying behavior that can be compared to the existing literature. Subjects entered the lab and found the printed experimental instructions on their desk, along with two dice – a red die and a blue die. The reason for having two dice is to keep the two treatments as similar as possible. The instructions informed subjects that their payment would be determined by rolling one of the two dice, and writing down the number they saw.⁷ They were provided with a table that showed how the number they reported would contribute to their payment – i.e. they would receive $\in 2$ multiplied by the number reported, with the exception of a report of 6. I follow Fischbacher and Föllmi-Heusi (2013) in assigning a payoff of $\in 0$ to a 6. This implies that the range of possible payoffs for the task is between $\notin 0$ and $\notin 10$.

The TwoDICE treatment condition followed exactly the same experimental protocol, with the following exception. Instead of choosing one of the two dice for payment, subjects were told that they should privately choose one die to be their *High Stakes* die

⁷Subjects were told that they could choose either of the two dice.

and one to be their *Low Stakes* die.⁸ They were informed that the number reported on their *High Stakes* die would be multiplied by $\in 2$, and the number reported on their *Low Stakes* die would be multiplied by $\in 0.1$. Subjects rolled both dice simultaneously and were asked to write down the number shown on both their *High Stakes* and *Low Stakes* die. The *High Stakes* report is therefore payoff-equivalent to the report in the BASELINE condition. The difference is the addition of the *Low Stakes* die. Possible payoffs for the task range between $\in 0$ and $\in 10.5$ in this treatment.

The instructions for both treatment conditions are included in the Appendices (Sections F and G).⁹

2.1 Experimental Procedure

The experiment was conducted in October and November 2018 at the WZB-TU laboratory in Berlin. The sessions were implemented by providing subjects with printed instructions, and collecting the data using physical dice, pen and paper.¹⁰ Subjects were invited to participate in the experiment using ORSEE (Greiner, 2015). A total of 200 subjects participated in the experiment with 100 subjects in each treatment condition (i.e. five sessions of approximately 20 subjects). After completing the dice rolling task, subjects completed a very short post-experimental questionnaire that was comprised of questions about the subject's age, gender, and field of study. Each session lasted approximately 15 minutes. The average subject was 23 years old, and 49% of the subjects were female. The average earnings on the dice rolling task was \in 6.9. In addition, all subjects received a fixed payment of \in 8.

3 Results

3.1 The relationship between *High* and *Low Stakes* reports

This section asks whether there is a relationship between the *High* and *Low Stakes* reports in the TwoDICE treatment. Here I present evidence pertaining to the two main questions

⁸This feature of the experimental paradigm relates to some of the ideas discussed by Jiang (2013). However, the "mind game" element of the design is not of central interest in this study; rather it is a tool to provide participants with greater certainty that they are not being monitored. This is why two dice are used in both treatments.

⁹In the instructions, I used the word "dice" to refer both to the singular and plural form of dice. There were several reasons for this decision. Firstly, in modern English the term "dice" can be used to refer to both the singular and the plural. Secondly, my subjective belief is that the word "die" is more likely to be misinterpreted, especially by second language speakers. Thirdly, I endeavoured to make sure that the instructions were never ambiguous regarding whether I was referring to one or two dice. Subjects appeared to understand the task well, and no subject asked a question about this.

¹⁰The reason for this was to increase the simplicity and tangibility of the task, and also to avoid the disconnect with one's actions and one's physical identity sometimes experienced when using computers.

asked by this paper: (i) is there evidence of downward lying on the *Low Stakes* reports? and (ii) is there a higher-than-expected fraction of (high, low) reports in the TwoDICE treatment?

Figure 1 shows the binarized distribution of reports for the *High* and *Low Stakes* dice in the TwoDICE treatment (i.e. the outcome space is collapsed into a binary variable that equals 0 if the report is low (≤ 2) and equals 1 if the report is high (≥ 3)). The left-panel shows that there is upward lying for the *High Stakes* lying opportunity. However, the right-panel shows that subjects were more likely to report a low number (i.e. ≤ 2) than a high number (i.e. ≥ 3) for the *Low Stakes* die. Only 37% of subjects reported having seen a number ≥ 3 , which is significantly less than the expected 50% percentage [p=0.012, binomial test, two-sided].¹¹ This serves as evidence that we do observe downward lying in the *Low Stakes* reports.¹²

Finding 1. There is evidence of downward lying on the Low Stakes die.



Figure 1: Distributions of *High & Low Stakes* reports (binarized, marginal distributions)

The observation that there is upward lying on the *High Stakes* die, and downward lying on the *Low Stakes* die is not sufficient to conclude that there is a negative relationship between the reports made by individuals. Therefore, Table 1 and Figure 2 present evidence on this second question of interest by showing the joint distributions of the *High* and *Low Stakes* reports. Table 1 shows the joint distribution of the binarized variables and Figure 2 graphically depicts the full joint distribution.

¹¹The full distribution of *Low Stakes* reports is portrayed in Figure 5 in Appendix E. The downward slope in reporting propensities contrasts with the upward slope commonly observed in this literature.

¹²It is worth noting that the amount of downward lying may be underestimated by looking at this figure. This would be the case if there is also any upward lying in the *Low Stakes* reports.

The conspicuous feature of Table 1 is that 50% of subjects make a high-report on the *High Stakes* die and a low-report on the *Low Stakes* die. This (high, low) joint report is chosen almost twice as often as the higher-paying (high, high) report (which is chosen by 26% of subjects) and approximately four times as often as (low,high) and (low,low).¹³ This suggests that subjects prefer to choose a combination of a high-report and a low-report in a way that allows them to earn a reasonably high payoff, but also (subjectively) signal honesty. This indicates that subjects are using the *Low Stakes* report to try to obscure their upward lying on the *High Stakes* report.

	Нідн Дісе				
	Low Report (≤ 2)	High Report (\geq 3)	Total		
Low Dice					
High Report (\geq 3)	11	26	37		
Low Report (≤ 2)	13	50	63		
Total	24	76	100		

Table 1: Observed binarized distribution of reports in the Two DICE game.

Finding 2. Half of the subjects make a (high,low) report, indicating an interaction between the High Stakes and Low Stakes reports.

Figure 2 reiterates this finding by showing that the majority of the density of the full distribution is below the diagonal, particularly in the bottom right quadrant.¹⁴

One potentially interesting ex post observation is that, of those subjects who made a report above the diagonal (i.e. reported having observed a higher number on the *Low Stakes* die than on the *High Stakes* die), 11/18 reported a 5 on the *Low Stakes* die (i.e. on the horizontal line). In comparison, only 7 subjects reported a number in the triangle in the north-west corner of Figure 2 enclosed by the two lines drawn onto the figure. One might speculate that some subjects viewed reporting a pair of numbers comprising a higher *Low Stakes* die value than *High Stakes* die value as another way to signal honesty. Therefore, by choosing the highest *Low Stakes* die value, namely 5, they could signal honesty in this way, and choose a reasonably high *High Stakes* die report. However, this observation involves only a small subset of the subjects in the experiment, so caution should be exercised when interpreting it.

¹³Testing whether each of the four cells is equal to the expected percentage of 25% by means of binomial tests indicates that (high,low) is observed statistically more often than expected, while (low, high) and (low, low) are observed less often than expected, with all three tests producing a p-value < 0.01 for a two-sided test. (high, high) is observed approximately as often as expected according to the objective distribution.

¹⁴Tables 4 and 5 in the Appendices report the exact frequencies for the distribution observed in the data, as well as the objective expected distribution respectively.



Figure 2: Full joint distribution High & Low Stakes reports.

3.2 Comparison with Fischbacher and Föllmi-Heusi (2013)

This section supplements the discussion of the main results above in two ways. Firstly, it provides a comparison of the behavior observed in my BASELINE treatment with Fischbacher and Föllmi-Heusi's (2013) results. This serves to alleviate concerns regarding the behavior described above in the TwoDICE treatment potentially being driven by the subject pool, or specific elements of the instructions, or the experimental protocol. Secondly, I compare the distribution of *High Stakes* reports observed in the TwoDICE and BASELINE treatment conditions. This treatment comparison provides evidence towards assessing whether the addition of a *Low Stakes* lying opportunity affects the *High Stakes* report, which could be the case if subjects are choosing both their reports in a strategic way to signal honesty.

The left panel of Figure 3 displays the distribution of reports in my BASELINE condition. The distribution is largely consistent with the pattern of reports commonly observed in the literature, with a higher frequency of reported 4's and 5's than expected under the objective distribution. Furthermore, Table 2 in Appendix A provides a direct comparison of my BASELINE treatment with the two comparable treatments from Fischbacher and Föllmi-Heusi (2013), showing that my results are very similar to theirs.

Turning to the comparison of behavior between my two treatments, it is worth noting that the monetary incentives that subjects face in choosing their report in BASELINE are identical to those for the *High Stakes* report in the TwoDICE treatment. The two panels of Figure 3 provide a graphical illustration of the treatment comparison. In contrast to BASELINE, in the TwoDICE condition, subjects report 3's more often, and 5's less often (see Table 3 in Appendix B for the statistical tests), with the implication that 4 becomes the modal report in the TwoDICE treatment. A comparison of the two distributions weakly

rejects the hypothesis that these two samples of reports are drawn from the same underlying population [p=0.087, Mann-Whitney]. This provides suggestive evidence in favor of the conclusion that the addition of the *Low Stakes* die is influencing subjects' choice of reports on the *High Stakes* die.



Figure 3: Reporting on High Stakes die across treatments (full distribution).

4 Discussion

4.1 The role of reputation concerns

The results described above suggest an important role for a motive other than a taste for money. The evidence points towards a prominent role for some form of reputation concerns, with subjects displaying a strong preference for making a (high, low) report – i.e. subjects appear to prefer making reports like (4,1) or (3,1), which contain a high and low number, rather than reports such as (4,4) or (3,4) which involve two high numbers. This behavior is not consistent with the equilibrium in models that propose that: (a) the perceived probability of lying after observing a report, r, is increasing in the frequency of observing report r in the population, and (b) that beliefs are correct in equilibrium (e.g. **Gneezy et al. (2018)**, Abeler et al. (2019), and Khalmetski and Sliwka (2019)). Under these models, subjects in the experiment could increase both their earnings and their reputation by switching to (high, high) reports that are observed less frequently (i.e. to the north-west quadrant of Figure 2).

There are different possible explanations for why the observed behavior differs from the predictions of the models. Firstly, subjects may have a different objective function to the one considered in these models (e.g. a different form of reputation concern). For example, subjects may view a (high, high) report as appearing less honest than a (high, low) report, independently of the relative frequencies of observing the two reports in the population. Secondly, subjects may hold reputation concerns in line with those described by these models, but have inaccurate subjective beliefs about the behavior of other individuals. Put simply, if a subject (incorrectly) believes that most other subjects will make a (high, high) report, then a (high, low) report is consistent with these models. The first explanation operates via preferences and the second via beliefs. In this paper, I do not distinguish between them, but this appears to be a fruitful avenue for future research. Importantly, both explanations point towards reputation concerns playing an important role in motivating subjects' decisions. The results here show that the nature of this reputation concern is more nuanced than indicated by the previous evidence.

This paper joins other nascent work that also explores the intricate role of reputational concerns. In a market setting, Tergiman and Villeval (2019) show that when the role of reputation is magnified, subjects drastically reduce their use of detectable lies, but continue to use lies where they can maintain plausable deniability. In another interesting contribution that shares some similarity in its fundamental essence to the current paper (i.e. doing something bad to appear good), Soraperra et al. (2019) provide evidence for the seemingly contradictory result that individuals are willing to take a socially harmful action to improve their social image. However, different from the current paper, in both these studies an individual's reputation can lead to real monetary consequences (see, also, Ely and Välimäki (2003) and Grosskopf and Sarin (2010) for important early work on this topic). My results add to this by providing evidence that social image alone is sufficient to generate this type of nuanced effect of reputational concerns on behavior.¹⁵

4.2 Increasing the dimensionality of lying games

Much of the experimental economics literature on lying involves studying situations where individuals make a single decision. This decision typically involves an explicit tradeoff between telling the truth and obtaining more money. This scientific approach of simplifying a decision context until the psychological processes of interest can be studied in isolation is both elegant and highly informative. However, it is important to then build on this foundational knowledge by adding complexity to the decision environment, thereby providing stress tests for the application of our models to more realistic everyday life situations. The current paper contributes to this effort by considering the interaction

¹⁵In an interesting recent contribution, Crede and von Bieberstein (2019) demonstrate just how strong reputation effects can be. By making subjects explicitly aware that the experimenter can observe the true state of the world, the authors essentially make partial and full lying disappear – subjects in this treatment abstain from lying.

between heterogeneous lying opportunities. However, there are also several other ongoing strands of research gathering evidence regarding the influence of increasing the dimensionality of the decision problem in other directions.

The first strand of research increases the temporal dimensionality of the decision problem by allowing subjects to make repeated homogeneous lying decisions (see, e.g., Jiang (2013), Abeler et al. (2014), Cohn et al. (2015)). In their meta-analysis, using a cross-study comparison (with all the qualifiers that brings), Abeler et al. (2019) find evidence of significantly lower lying rates when faced with repeated lying opportunities.¹⁶ An important caveat is that these studies were typically not designed primarily to test between lying behavior in one-shot and repeated settings. In recent work, Casal and Filippin (2018) do explicitly compare one-shot reporting with a repeated reporting task and observe similar average levels of cheating across treatments. Belot and van de Ven (2018) also study behavior when facing repeated lying opportunities, and similar to the current paper, they consider heterogeneous stakes. The authors implement a treatment in which subjects are first exposed to several rounds of a high-stakes sender-receiver game and then to a low-stakes sender-receiver game. They compare this with a treatment where the order of high and low stakes incentives is reversed. Belot and van de Ven (2018) hypothesize that subjects facing high incentives first will continue to lie at high rates when the incentives are reduced due to some form of habituation effect. However, instead they observe a marked drop after the incentive reduction. I view this observation as being consistent with my results, albeit in a very different context, and in an experiment designed to test a different hypothesis.

The second strand of research considers non-temporal increases in the dimensionality of the decision problem. In work related to the current paper, Geraldes et al. (2019) study situations in which there are multiple homogeneous lying opportunities. The paper tests whether the increased dimensionality of the decision problem leads to an increase in lying, but finds no evidence in support of this hypothesis. There is also a set of papers that introduce the possibility for additional non-monetary motives (e.g. social preferences) to influence the decision of whether to lie, studying the interaction between motives. In a seminal contribution, Gneezy (2005) discusses the role of consequences in lying decisions, classifying lies as being white, altruistic, spiteful and selfish lies. The results demonstrate a sensitivity in making the decision to lie to the implications for both one's own and others' payoffs.¹⁷ Barron et al. (2019) also study the interaction of social preferences a tension between two conflicting moral motives, truth-telling and fairness. The paper

¹⁶Although, they also observe a small upward slope in the trend across rounds towards higher lying rates.

¹⁷Alempaki et al. (2018) build on this work by studying the responsiveness of spiteful and selfish lies to earlier unkind behavior (i.e. deception due to reciprocity).

evaluates how subjects resolve this tension, providing evidence that subjects are more likely to adhere to the moral motive that is more in line with their private self-interest. In a related vein, Danilov and Saccardo (2019) argue that their subjects use a truth-telling norm as an excuse to justify discrimination.

Finally, this paper also relates to the literatures studying moral licensing (see, e.g., Monin and Miller (2001), Cain et al. (2005), Merritt et al. (2010), Mullen and Monin (2016), and Engel and Szech (2017)) and conscience accounting (Gneezy et al., 2014). However, the results I observe do not seem to be explained by these motives, since they predict that subjects would offset lying on one die (i.e. the *High Stakes* die) with telling the truth on the other die (e.g. the *Low Stakes* die). This is not what I observe.¹⁸

5 Conclusion

While lying has been studied extensively in the behavioral economics literature, much of this work has studied situations in which individuals can tell a single lie (or multiple homogeneous lies) to receive more money. However, lying opportunities in everyday life often do not involve a simple tradeoff between telling a single lie and earning money. Rather, individuals often face multiple interactions with a given observer. The evidence reported here suggests that individuals try to strategically build a reputation for honesty in interactions where the potential monetary benefits of being deceptive are low, thereby facilitating the exploitation of more lucrative lying opportunities without appearing completely dishonest – i.e. they try to use little lies to cover up their bigger lies.

This behavior indicates that subjects face a tension between a strong concern regarding the reputational influence of their decisions and their monetary self-interest. This strong concern for one's reputation is consistent with the spirit of the prominent lying models. However, in contrast to my findings, these models would not predict systematic downward lying in low stakes interactions. This observed pattern of behavior suggests that reputational concerns might play a more nuanced role in influencing when and how individuals choose to lie. Essentially, it suggests that individuals are sometimes willing to take a *bad* (or immoral) action if they perceive it as signaling that they are doing something *good* (or moral). The implication is that sometimes the motivation to appear good is stronger than the motivation to be good.

¹⁸Incidentally, in terms of the dichotomy of dynamic moral behavior according to *consistency* and *licensing* as discussed by Mullen and Monin (2016), my results are not consistent with either *licensing* (see the main text) or *consistency*, since this would imply that subjects who lie upward on the *High Stakes* die are more likely to lie upward on the *Low Stakes* die, which is also not what I observe.

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(4), 1115–1153.
- Alempaki, D., G. Doğan, and S. Saccardo (2018). Deception and reciprocity. *Experimental Economics (forthcoming)*.
- Barron, K., R. Stüber, and R. van Veldhuizen (2019). Motivated motive selection in the lying-dictator game. *WZB Discussion Paper SP II 2019-303*.
- Belot, M. and J. van de Ven (2018). The persistence of lying. Working Paper.
- Cain, D. M., G. Loewenstein, and D. A. Moore (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies* 34(1), 1–25.
- Casal, S. and A. Filippin (2018). The effect of reporting multiple private information outcomes on the inclination to cheat. *Working Paper*.
- Cohn, A., E. Fehr, and M. A. Maréchal (2014). Business culture and dishonesty in the banking industry. *Nature 516*, 86 89.
- Cohn, A. and M. A. Maréchal (2018). Laboratory measure of cheating predicts school misconduct. *The Economic Journal 128*, 2743–2754.
- Cohn, A., M. A. Maréchal, and T. Noll (2015). Bad boys: How criminal identity salience affects rule violation. *The Review of Economic Studies* 82(4), 1289–1308.
- Coricelli, G., M. Joffily, C. Montmarquette, and M. C. Villeval (2010). Cheating, emotions, and rationality: an experiment on tax evasion. *Experimental Economics* 13(2), 226–247.
- Crede, A.-K. and F. von Bieberstein (2019). Reputation and lying aversion in the die roll paradigm: Reducing ambiguity fosters honest behavior. *Working Paper*.
- 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(3), 1081–1100.
- Danilov, A. and S. Saccardo (2019). Disguised discrimination. Mimeo.
- Dufwenberg, M. and M. A. Dufwenberg (2018). Lies in disguise–a theoretical analysis of cheating. *Journal of Economic Theory* 175, 248–264.

- Ely, J. C. and J. Välimäki (2003). Bad reputation. *The Quarterly Journal of Economics* 118(3), 785–814.
- Engel, J. and N. Szech (2017). A little good is good enough: Ethical consumption, cheap excuses, and moral self-licensing. *CESifo Working Paper Series*.
- Fischbacher, U. and F. Föllmi-Heusi (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association* 11(3), 525–547.
- Geraldes, D., F. Heinicke, and S. Rosenkranz (2019). Lying on two dimensions and moral spillovers. *Working Paper*.
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review 95*(1), 384–394.
- Gneezy, U., A. Imas, and K. Madarász (2014). Conscience accounting: Emotion dynamics and social behavior. *Management Science* 60(11), 2645–2658.
- Gneezy, U., A. Kajackaite, and J. Sobel (2018). Lying aversion and the size of the lie. *American Economic Review 108*(2), 419–53.
- Greiner, B. (2015, Jul). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association* 1(1), 114–125.
- Grosskopf, B. and R. Sarin (2010). Is reputation good or bad? an experiment. *American Economic Review 100*(5), 2187–2204.
- Hanna, R. and S.-Y. Wang (2017). Dishonesty and selection into public service: Evidence from india. *American Economic Journal: Economic Policy* 9(3), 262–90.
- Jiang, T. (2013). Cheating in mind games: The subtlety of rules matters. *Journal of Economic Behavior & Organization 93*, 328–336.
- 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. *AEJ:Micro [forthcoming]*.
- Mazar, N., O. Amir, and D. Ariely (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of marketing research* 45(6), 633–644.
- Merritt, A. C., D. A. Effron, and B. Monin (2010). Moral self-licensing: When being good frees us to be bad. *Social and personality psychology compass* 4(5), 344–357.

- Monin, B. and D. T. Miller (2001). Moral credentials and the expression of prejudice. *Journal of personality and social psychology 81*(1), 33.
- Mullen, E. and B. Monin (2016). Consistency versus licensing effects of past moral behavior. *Annual review of psychology 67*, 363–385.
- 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(2), 181–190.
- Soraperra, I., A. Suvorov, J. Van De Ven, and M. C. Villeval (2019). Doing bad to look good: Negative consequences of image concerns on pro-social behavior. *Working Paper*.

Tergiman, C. and M. C. Villeval (2019). The way people lie in markets. *Working Paper*.

APPENDICES

A Comparison of BASELINE with FFH

Table 2: Comparison of BASELINE reports with Fischbacher and Föllmi-Heusi (2013)

	Нідн Дісе								
	0	1	2	3	4	5	Ν		
This paper									
BASELINE [1]	9	9	6	9	31	36	100		
FFH									
BASELINE [2]	6.4	7.2	11.6	12.6	27.2	35.0	389		
HIGH STAKES [3]	11.25	5	15**	8.75	27.5	32.5	80		

Notes: (i) The rows display the observed percentage of each die report, (ii) in the rows labeled [2] and [3] corresponding to treatments from Fischbacher and Föllmi-Heusi (2013), stars display a test of whether the report percentage observed in the relevant FFH treatment differs significantly from the percentage observed in my BASELINE treatment (i.e. row [1]) using a two-sided test of proportions; * 10%, ** 5%, *** 1%.

B Comparison of reported and objective distributions

	Нідн Дісе						
	0	1	2	3	4	5	Ν
BASELINE [1]	9	9	6	9	31	36	100
TwoDice [2]	10	6	8	18	37	21	100
Expected Dist. [3]	16.7	16.7	16.7	16.7	16.7	16.7	
Stat. Comparisons							
[2] - [1]	1	-3	2	9*	6	-15**	
[3] - [1]	7.7**	7.7**	10.7***	7.7**	-14.3***	-19.3***	
[3] - [2]	6.7*	10.7***	8.7***	-1.3	-20.3***	-4.3	

Table 3: Comparison of reported distributions between treatments

Notes: (i) The first three rows display the percentage of each die report, (ii) the bottom three rows display the percentage point differences between the first three rows, considering binary comparisons, (iii) for the fifth and sixth rows (i.e. [3]-[1] and [3]-[2]), stars display the significance level of a two-sided binomial test that the observed percentage differs from 16.7%; * 10%, ** 5%, *** 1%, (iv) for the fourth row (i.e. [2]-[1]), stars display the significance level of a two-sided test of proportions; * 10%, ** 5%, *** 1%. The significance levels of the fourth row are unchanged when instead considering a two-sided Fisher exact test.

C One-dimensional depiction of distribution in TwoDICE



Figure 4: Report distribution across all 36 states, ordered by earnings

Note: The x-axis reflects the 36 possible states, ordered in increasing value. For example, "5/4" refers to a report of 5 on the high value dice, and a report of 4 on the low value dice. The vertical dotted lines divide the state-space according to the 6 high dice values. The horizontal dashed line reflects the objective probability of each state.

D Theoretical and reported joint distributions in TwoDICE

	Нідн Дісе							
	0	1	2	3	4	5	Total	
Low Dice								
5	1	3	1	3	3	2	13	
4	0	0	1	0	3	3	7	
3	4	0	1	3	6	3	17	
2	0	1	1	2	9	2	15	
1	0	1	4	5	8	5	23	
0	5	1	0	5	8	6	25	
Total	10	6	8	18	37	21	100	

Table 4: Observed distribution of reports in the Two DICE game.

	Нідн Дісе						
	0	1	2	3	4	5	Total
Low Dice							
5	2.8	2.8	2.8	2.8	2.8	2.8	16.7
4	2.8	2.8	2.8	2.8	2.8	2.8	16.7
3	2.8	2.8	2.8	2.8	2.8	2.8	16.7
2	2.8	2.8	2.8	2.8	2.8	2.8	16.7
1	2.8	2.8	2.8	2.8	2.8	2.8	16.7
0	2.8	2.8	2.8	2.8	2.8	2.8	16.7
	16.7	16.7	16.7	16.7	16.7	16.7	100

Table 5: Expected ojective distribution of die rolls in the Two DICE game.

E Marginal Low- and High-Stakes distributions in TwoDICE

Figure 5: Reported distribution on *Low-Stakes* die in TwoDICE treatment.



Figure 6: Reported distribution on *High-Stakes* die in TwoDICE treatment.



F Instructions: BASELINE

Every participant will receive a 5 Euro show-up fee for participating in the experiment, which will be paid out independently of the decisions made in the experiment.

Furthermore, you will be able to earn additional money. At the end of the experiment, you will receive the income which you earned over the course of the experiment plus the 5 Euro show-up fee in cash.

In front of you, there are two dice – one red dice, and one blue dice. Your payment for today's experiment will be determined by rolling one of these two dice. You will need to decide which dice you want to roll. After you have decided, you will roll that dice and report the number you see. For your payment, the table below shows how the dice roll will contribute to your payoff.

Payoff contribution from the dice roll

Number thrown	1	2	3	4	5	6
Resulting payoff	€2	€4	€6	€8	€10	€0

After rolling the dice, you should write down the number shown on your dice in the box below. After doing this, you will be asked to complete a short questionnaire. You will be paid a \in 3 fixed fee for completing the questionnaire. When you have finished the questionnaire, you can also roll the dice more times to check that they are working properly and are fair dice (they are).

Your total earnings = $\in 5 + \in 3 + DICE$ payoff

After everyone has completed this dice rolling task, and the questionnaire, the experiment will be over and we will proceed to payment.

Please now decide which one of the two dice you would like to roll. When you have done this, please roll the dice and report the number on the dice.

Number on DICE	

After you have completed this task, please wait for the experimenter to hand out the questionnaire.

G Instructions: TwoDice

Every participant will receive a 5 Euro show-up fee for participating in the experiment, which will be paid out independently of the decisions made in the experiment.

Furthermore, you will be able to earn additional money. At the end of the experiment, you will receive the income which you earned over the course of the experiment plus the 5 Euro show-up fee in cash.

In front of you, there are two dice – one red dice, and one blue dice. Your payment for today's experiment will be determined by rolling these two dice. You will need to decide which dice you want to be your HIGH-STAKES-DICE, and which you want to be your Low-STAKES-DICE. After you have decided, you will roll both dice at the same time and report the number you see on each dice. For your payment, the tables below shows how the two dice rolls will contribute to your payoff.

Number thrown	1	2	3	4	5	6
Resulting payoff	€2	€4	€6	€8	€10	€0

Payoff contribution from the HIGH-STAKES-DICE

Payoff contribution from the Low-STAKES-DICE

Number thrown	1	2	3	4	5	6
Resulting payoff	€0.1	€0.2	€0.3	€0.4	€0.5	€0

After rolling the two dice, you should write down the number shown on your HIGH-STAKES-DICE and your Low-STAKES-DICE in the boxes below. After doing this, you will be asked to complete a short questionnaire. You will be paid $a \in 3$ fixed fee for completing the questionnaire. When you have finished the questionnaire, you can also roll the dice more times to check that they are working properly and are fair dice (they are).

Your total earnings = € 5+€ 3+HIGH-STAKES-DICE payoff+Low-STAKES-DICE payoff

After everyone has completed this dice rolling task, and the questionnaire, the experiment will be over and we will proceed to payment.

Please now decide which is your HIGH-STAKES-DICE and which is your Low-STAKES-DICE. When you have done this, please roll both dice and report the number on each of the two dice.



After you have completed this task, please wait for the experimenter to hand out the questionnaire.

Discussion Papers of the Research Area Markets and Choice 2019

Research Unit: Market Behavior	
Azar Abizada, Inácio Bó Hiring from a pool of workers	SP II 2019-201
Philipp Albert, Dorothea Kübler, Juliana Silva-Goncalves Peer effects of ambition	SP II 2019-202
Yves Breitmoser, Sebastian Schweighofer-Kodritsch Obviousness around the clock	SP II 2019-203
Tobias König, Sebastian Schweighofer-Kodritsch, Georg Weizsäcker Beliefs as a means of self-control? Evidence from a dynamic student survey	SP II 2019-204
Rustamdjan Hakimov, Dorothea Kübler Experiments on matching markets: A survey	SP II 2019-205
Puja Bhattacharya , Jeevant Rampal Contests within and between groups	SP II 2019-206
Kirby Nielsen, Puja Bhattacharya, John H. Kagel, Arjun Sengupta Teams promise but do not deliver	SP II 2019-207
Julien Grenet, Yinghua He, Dorothea Kübler Decentralizing centralized matching markets: Implications from early offers in university admissions	SP II 2019-208
Joerg Oechssler, Andreas Reischmann, Andis Sofianos The conditional contribution mechanism for repeated public goods – the general case	SP II 2019-209
Rustamdjan Hakimov, CPhilipp Heller, Dorothea Kübler, Morimitsu	SP II 2019-210
How to avoid black markets for appointments with online booking systems	
Thibaud Pierrot Negotiation under the curse of knowledge	SP II 2019-211
Research Unit: Economics of Change	
Kai Barron, Steffen Huck, Philippe Jehiel Everyday econometricians: Selection neglect and overoptimism when learning from others	SP II 2019-301
Marta Serra-Garcia, Nora Szech The (in)elasticity of moral ignorance	SP II 2019-302

Kai Barron, Robert Stüber, Roel van Veldhuizen Motivated motive selection in the lying-dictator game	SP II 2019-303
Maja Adena, Steffen Huck Can mass fundraising harm your core business? A field experiment on how fundraising affects ticket sales	SP II 2019-304
Maja Adena, Rustamdjan Hakimov, Steffen Huck Charitable giving by the poor: A field experiment on matching and distance to charitable output in Kyrgyzstan	SP II 2019-305
Maja Adena, Steffen Huck Personalized fundraising: a field experiment on threshold matching of donations	SP II 2019-306
Kai Barron Lying to appear honest	SP II 2019-307

Research Unit: Ethics and Behavioral Economics

Daniel Parra, Manuel Muñoz-Herrera, Luis Palacio	SP II 2019-401
The limits of transparency as a means of reducing corruption	