

# Observability and Lying

Tilman Fries, Uri Gneezy, Agne Kajackaite, and Daniel Parra\*

May 18, 2021

## Abstract

Experimental participants in a cheating game draw a random number and then report any number they wish, receiving a monetary payoff based only on the report. We study how these reports depend on the level of observability of both the random draw and the report by the experimenter. Our results show that whereas increasing the observability of the random draw decreases cheating, increasing the anonymity of the reports does not affect average reports.

Key Words: Lying, Cheating, Observability, Social Image, Double Blind, Laboratory Experiment.

JEL Classification: C91, D82, D91.

\* Fries: WZB Berlin Social Science Center, Berlin, D-10785, Germany (email: [tilman.fries@wzb.eu](mailto:tilman.fries@wzb.eu)); Gneezy: Rady School of Management, University of California, San Diego, La Jolla, CA 92093, and CREED, University of Amsterdam (email: [ugneezy@ucsd.edu](mailto:ugneezy@ucsd.edu)); Kajackaite: WZB Berlin Social Science Center, Berlin, D-10785, Germany (email: [agne.kajackaite@wzb.eu](mailto:agne.kajackaite@wzb.eu)); Parra: WZB Berlin Social Science Center, Berlin, D-10785, Germany (email: [daniel.parra@wzb.eu](mailto:daniel.parra@wzb.eu)).

# 1 Introduction

Lying is central to economic interactions with asymmetric information. The classic modeling in economics makes strong simplifying assumptions, such as selfishness and no cost of lying, and derives the theoretical benchmark according to which people would lie whenever the strategic cost of lying is lower than the benefit. A growing literature in experimental economics, however, shows that the simplifying assumptions are rejected when we observe behavior in the laboratory, suggesting some psychological cost of lying: All else equal, a large fraction of people prefers not to lie (e.g., Gneezy, 2005; Fischbacher and Föllmi-Heusi, 2013). Recent models try to incorporate such costs, relaxing some of the simplifying assumptions (Gneezy et al., 2018; Dufwenberg and Dufwenberg, 2018; Abeler et al., 2019; Khalmetski and Sliwka, 2019).

Experimental economics takes two popular approaches to study lying. The first uses the “deception game” (Gneezy, 2005), which is a two-player game in which the sender is informed about the state of the world. She sends a message regarding payoffs to the receiver who is not informed about the state of the world, and the receiver chooses an outcome. This decision determines the payoff of both the sender and the receiver. In this game, the sender may choose to lie in her message if she expects to benefit from it, and her lie is observed by the experimenter. The results from experiments using the deception game show that some participants choose to send an honest message even when doing so is costly (e.g., Gneezy, 2005; Dreber and Johannesson, 2008; Sutter, 2009; Rode, 2010; Wang et al., 2010; Cappelen et al., 2013).

The fact that, in the deception game, the state of the world is also observed by the experimenter may deter the sender from lying. This social-signaling cost comes from the sender not wanting to appear dishonest to the experimenter. An alternative reason to avoid lying is self-

signaling, in which the cost of lying comes from the sender suffering a cost to her identity as a moral person.

The second type of game used to study cheating, namely, the “cheating game,” can help in isolating self-signaling from social-signaling. Fischbacher and Föllmi-Heusi (2013) introduced this game in which a participant observes in private the realization of a random process (e.g., the outcome of a die roll) and then reports the outcome to the experimenter. In this method, the experimenter does not know the actual realization of the die roll, and hence whether a participant is lying. This privacy reduces the social-signaling costs of lying. By observing the distribution of reported outcomes by many participants and comparing it with the statistical one, the experimenter can then infer how truthful the population in the experiment was.

The results in the literature suggest that participants lie more in the deception game than in the cheating game (Gerlach et al., 2019). In addition, participants in the deception game are more sensitive to external factors, such as incentives, than participants in the cheating game (Kajackaite and Gneezy, 2017). One reason for these results could be the weight participants place on self-signaling relative to social signaling, with social signaling having more room in the standard cheating game than in the deception game.

In this paper, we study how the possibilities of social signaling one’s honesty affect lying in cheating games. In the experiment, we use a die-roll game, similar to the one introduced by Fischbacher and Föllmi-Heusi (2013). We ask participants to roll a die and report the number they saw. The higher the number they report, the higher the payoff, which creates an incentive to cheat. In the treatments, we vary (1) *the observability of one’s random draw* and (2) *the observability of one’s report* by the experimenter. We test how the possibilities of socially

signaling one's honesty change the extent of lying on both the extensive (lie or truth) and intensive (size of the lie) margins.

Gneezy et al. (2018) and Abeler et al. (2019) showed that, in cheating games, the observability of the random draw matters (comparison of observed and non-observed games). In particular, they showed that when the experimenter observes the random draw, few participants lie, and when they lie, they mostly lie all the way by reporting the maximum possible outcome. By contrast, when the experimenter does not observe the random draw, more participants lie (extensive margin) and a significant fraction of participants tell partial lies—they lie by reporting a non-maximal outcome (intensive margin). Both papers propose that the main reason for more partial lying in non-observed games is social signaling. Because the experimenter does not know the actual random-draw outcome, reporting a number that is not the maximum makes one appear more honest than reporting the maximum. Such social signaling cannot take place in the observed game, where the experimenter knows the random-draw outcome.

In this paper, similar to Gneezy et al. (2018) and Abeler et al. (2019), we test how the degree of observability of the random draw affects lying. We use a super-observed die-roll game in which the experimenter sees and writes down the die-roll outcome of the participant, an observed die-roll game in which the experimenter sees but does not write down the die-roll outcome of the participant, and a non-observed die-roll game in which the experimenter does not see the die-roll outcome of the participant. We find that the more observed the random outcome is, the less participants lie.

In addition to varying the observability of the random draw, we vary the observability of the report. We introduce a double-blind procedure in which the experimenter observes neither the random outcome nor the report. In such an environment, no social signaling should be taking

place; that is, because the experimenter does not see the participant's report, the participant cannot signal that she is honest. Therefore, we expect to see more lying (extensive margin) and less partial lying (intensive margin) in a double-blind procedure than in the standard procedure.

In contrast to our predictions, the double-blind procedure has no effect on lying. This result is in line with evidence by Fischbacher and Föllmi-Heusi (2013), who also found no effect of a double-blind procedure on lying in a cheating game. There are some possible reasons for why the double-blind procedure has no effect on lying. First, we observe high reports already in the standard non-observed cheating game, which creates a ceiling effect for the double-blind treatments. The extensive margin (the number of liars) cannot increase by much compared to the standard cheating game. This ceiling effect does not, however, explain the persistence of partial lies in the double-blind treatments. As an explanation, we consider that some participants might have internalized a preference for appearing honest and might see themselves as “their own audience” (Dufwenberg and Dufwenberg, 2018). We connect our findings to results of double-blind procedures in other economic games and conclude that this explanation is also in line with previous findings.

Overall, our results provide evidence that whereas the observability of one's random draw leads to less cheating, the observability of one's report has no effect on lying, neither on the intensive nor extensive margin. Our results contribute to a broader literature investigating how changes to the communication protocol can foster information transmission. Papers that study different interventions include Blume et al. (2001), who show that in sender-receiver games, communication with messages that have a priori meanings is more effective than communication with messages that are a priori meaningless; Serra-Garcia et al. (2011) on vague communication;

Khalmetski et al. (2015) on lying by omission; and Blume et al. (2019) on the randomized response technique.

## 2 Experimental design and procedure

We use a between-subjects design with five treatments. All treatments involve the generation of a random number between 1 and 6 using a die roll. After observing the random number, participants report it and are paid the number they report in euros multiplied by two. That is, if they report 1 they get 2 euros, if they report 2 they get 4 euros, and so on. In different treatments, we vary the degree with which the random draw and the report of the random draw are observable by the experimenter. We pre-registered the experimental design and the hypotheses at [aspredicted.com](http://aspredicted.com).<sup>1</sup>

### *Treatments*

Treatment 1: “Super Observed.” The participant rolls a six-sided die in front of the experimenter and is then asked to report the number on a reporting sheet. The experimenter observes the random-draw outcome, says “ok,” and writes the number for herself, and only then moves on to the next participant. This procedure is common knowledge. In this treatment, we can compare the numbers that participants actually saw with their reports, and therewith can see whether and to what extent participants have lied.

---

<sup>1</sup> The links for the pre-registration are at <http://aspredicted.org/blind.php?x=9q7fp6>, and <http://aspredicted.org/blind.php?x=cb2cm2>.

Treatment 2: “Observed.” This treatment has the same structure as the “Super Observed” treatment, but the experimenter does not write the random-draw outcome on a piece of paper before moving on to the next participant. That is, whereas in the “Super Observed” game, participants know the experimenter has written proof of the number rolled, in the “Observed” game, they know the experimenter does not write it down and would need to remember all random-draw outcomes of the session (most sessions consisted of 24 participants) at the time of payment to know whether a participant has lied.

In the “Observed” treatment, we do not know which individuals cheated and by how much, but we can follow the convention in this literature and estimate the level of cheating by comparing the expected theoretical uniform distribution of numbers between 1 and 6 with the actually reported distribution of numbers. Hence, this treatment differs from the standard die roll game (Fischbacher and Föllmi-Heusi, 2013) in that participants are physically observed by the experimenter when rolling the die even though there is no record of the die roll.

Treatment 3: “Basic.” This treatment is similar to the observed one, but participants roll the die in private such that the experimenter cannot observe the outcome. The remaining structure of the game is the same as in the observed treatments. This treatment is similar to the standard die roll game introduced by Fischbacher and Föllmi-Heusi (2013), with some minor variations presented in the section on the Experimental Procedure.

The first three treatments vary observability of the random-draw outcome but keep the observability of the report constant. Both Gneezy et al. (2018) and Abeler et al. (2019) have also varied the observability of the random-draw outcome in their treatments. However, our

procedures differ in some important aspects. Gneezy et al. (2018) and Abeler et al. (2019) implement observability of the random draw by computer recording the random draw outcome. In contrast, in our experiment the observability is increased by a human interaction. The experimenter (and not a computer as in previous studies) observes the participants' random draw directly at the time of the die roll, which makes the observability more salient. In both the Super Observed and Observed treatments, the experimenter stands next to the experimental participants, whilst she or he is rolling the die. In the Super Observed treatment, the observability is especially salient, since the experimenter writes down the random draw outcome, whilst standing next to the participant.

In the first three treatments, at the end of each treatment, participants go to the experimenter to report their number and proceed with the payment. Therefore, the experimenter observes the report of each participant. In treatments 4 and 5, by contrast, we use a double-blind procedure in which the experimenter does not observe the reports of participants.

Treatment 4: "Double Blind." Participants roll the die in private, report the number on a reporting sheet, and take the money they earned from twelve 1 EUR coins in an envelope that we place in their cubical prior to the experiment (in all treatments, the envelopes containing twelve 1 EUR coins are placed in the cubical).<sup>2</sup> After taking the money, participants place the reporting sheet inside the envelope with the remaining coins, seal the envelope, put the envelope together

---

<sup>2</sup> We place envelopes with money in the cubical in all treatments to keep the endowment effect of money constant over the treatments. The next section describes the physical procedure in detail.



with other participants' envelopes in a box placed at the exit of the laboratory and leave the laboratory.<sup>3</sup> In this treatment, the experimenter cannot attribute the reports to participants.

Treatment 5: "Super Double Blind." In the double-blind treatment, participants might be worried about the sound of the envelope when putting it into the box. If a participant takes all the money, the envelope will not make any sound. To get rid of this concern, in the "Super Double Blind" treatment, we add fake metal coins to the envelope. In addition, to rule out the possibility of participants expecting to sign some receipt at the end (as they usually do in other experiments), we explicitly write in the instructions that they will not sign any receipts at the end of the experiment. Finally, participants receive printed instructions when they enter the laboratory instead of the instructions being put on each participant's desk prior to the session. We incorporate this procedure to make clear to participants that the sheets have no marks to trace them. These three changes to the Double Blind protocol help increase participants' belief that the procedure is indeed completely anonymous. We conducted this treatment to test the robustness of Double Blind's results.

The Double Blind treatment is closely related to the double blind condition implemented by Fischbacher and Föllmi-Heusi (2013). We compare our Double Blind treatment to the Basic treatment, in which participants report their payoff directly to the experimenter. This is different from Fischbacher and Föllmi-Heusi (2013), who implement a control treatment where participants are endowed with a monetary sum and an envelope. Instead of reporting their payoff to the experimenter and handing back the remaining money, participants leave any remaining

---

<sup>3</sup> The experimenter was in the other end of the laboratory, in no close proximity to the box.

money inside their cubicle after the experiment.<sup>4</sup> In this treatment, reports are only observed in the sense that the experimenter can match reports to cubicle numbers after the session. Therefore, our treatment difference between the Double Blind game and control (Basic treatment) is arguably larger, as the Basic treatment removes any ambiguity participants might have about whether their report can be matched to them.

### *Experimental Procedure*

We conducted the experiment in July – September 2019 at the Berlin Experimental Economics Lab. We recruited 696 (44.68% female)<sup>5</sup> participants via ORSEE (Greiner, 2015); none of them participated in more than one session. We conducted 30 experimental sessions, with a session lasting approximately 30 minutes.

After arriving to the lab, participants were randomly seated in cubicles. In the first four treatments, each cubicle was equipped with a die, a pen, an envelope with twelve 1 EUR coins, and a sheet of paper with instructions and a space for reporting the die roll. In the Super Double

---

<sup>4</sup> While Fischbacher and Föllmi-Heusi (2013) also implement a benchmark treatment on which our Basic treatment builds on, they introduce specific control treatments for each treatment manipulation they introduce. They compare results of each of their manipulations to the specific control treatment but not to the benchmark treatment. One reason for the additional controls in Fischbacher and Föllmi-Heusi (2013) is that the participant recruitment criteria do not stay constant across benchmark and treatments, so that any comparison between them might be confounded by selection. Another reason is that, in the benchmark treatment, participants receive money only after they report their number to the experimenter. It is thus not directly comparable to their double-blind procedure, in which participants are endowed with a sum and have to hand money back.

<sup>5</sup> We determined the number of participants by calculating the sample size necessary to detect a minimum treatment effect with 80% power. In particular, in the calculation, we took the reporting behavior in Fischbacher and Föllmi-Heusi (2013) as a benchmark for the Basic treatment and assumed that reporting will change such that the average reported die roll will be 0.6 points lower in Observed and 0.6 points higher in Double Blind. Based on these assumptions, we fixed a sample size and simulated 1000 draws and reports for each observation in the sample. The simulations implied that a sample size of 144 observations per treatment would be sufficient to detect a treatment difference in reporting with a Mann-Whitney test at the 5% level in at least 80% of all cases. Based on our calculation, we aimed at having 144 participants a treatment; however, due to no-shows, the number of participants was slightly lower in some treatments.

Blind treatment, the envelope contained, in addition, ten blank metal coins, and participants received the printed instructions upon entering the laboratory.

After all the participants were seated, they were asked to read the instructions (see Appendix B) and were allowed to ask questions privately. Then, depending on the treatment, participants were asked to roll the die in private or wait for the experimenter to come to their cubicle to observe the die roll. After observing the roll, participants wrote down a number on their reporting sheet and took out the money they earned from the envelope, leaving the remaining money in it.

In the Super Observed, Observed and Basic treatments, each participant then went to the experimenter and gave the experimenter the reporting sheet and the envelope. The experimenter checked whether the remaining money in the envelope matched the report on the reporting sheet. Then, the participant signed a receipt and left the laboratory. Participants knew about this procedure since the beginning of the experiment.

In the Double Blind and Super Double Blind treatments, the experimenter and participants did not interact at the end of the experiment. In these treatments, after taking the money they earned, participants placed the reporting sheets inside the envelopes, put their envelopes in a box placed at the exit, and left the laboratory.

Finally, we were interested to see gender effects in our treatments. We, however, did not want to affect the anonymity perception in the Double Blind procedure. Therefore, we asked the participants at the end of the Double Blind treatments to write down their gender on the envelope. In other treatments, we marked the gender on the reporting sheet when doing the

payouts. Importantly, in the Double Blind treatments, we waited until all participants sealed their envelopes, and only then asked them to write down their gender.<sup>6</sup>

Table 1 presents all of our treatments and the number of participants in each, and Table 2 summarizes the experimental procedure in each treatment.

**Table 1. Summary of Treatments and Number of Participants**

<i>Treatment</i>	<i>Number of participants</i>
<i>Super Observed</i>	<i>135 (37.04% female)</i>
<i>Observed</i>	<i>144 (51.39% female)</i>
<i>Basic</i>	<i>144 (48.61% female)</i>
<i>Double Blind</i>	<i>142 (47.18% female)</i>
<i>Super Double Blind</i>	<i>131 (37.69% female)</i>

<sup>6</sup> The procedure allows us to identify the gender of participants, albeit with some noise. In the Basic and both Observed treatments, the experimenter assesses the gender of participants by visual inspection, which might not coincide with the gender participants self-identify with. After these sessions, we confirmed, however, that in each session, the marked gender composition was the same as recorded on the registration lists in ORSEE. In the Double Blind treatments, the self-reported gender composition matched the registration lists from ORSEE, with one exception in the Super Double Blind treatment, where one participant wrote “diverse” for their gender, differently from the records in ORSEE. We do not include this participant to the analyses of lying differences between men and women.

**Table 2. Experimental procedure in each treatment**

<b>Treatment</b>	<b>Observability of the draw</b>	<b>Observability of the report</b>
<b>Super Observed</b>	Participants roll the die in front of experimenter; experimenter records the die roll.	At the end of the experiment, participants go to the experimenter with the reporting sheet and remaining coins.
<b>Observed</b>	Participants roll the die in front of experimenter.	As in Super Observed.
<b>Basic</b>	Participants roll the die in private.	As in Super Observed.
<b>Double Blind</b>	As in Basic.	At the end of the experiment, participants put reporting sheets and remaining coins in an envelope and place it in a box at the exit of the laboratory.
<b>Super Double Blind</b>	As in Basic.	At the end of the experiment, participants put reporting sheets and remaining coins in an envelope and place it in a box at the exit of the laboratory. Envelopes contain additional metal coins so that the weight and sound of the envelope are not informative about the report.

### *Hypotheses*

Our five treatments vary the degree of the experimenter's observability of the die roll and the report. According to lying models, the degree of observability may affect the individual willingness to lie (Gneezy et al., 2018; Dufwenberg and Dufwenberg, 2018; Abeler et al., 2019; Khalmetski and Sliwka, 2019). Given the theoretical prediction of these models and the experimental evidence presented in Gneezy et al. (2018), we hypothesize that as the degree of observability decreases, participants' willingness to lie increases.

**Hypothesis 1 (Over-reporting and observability):** The average number reported will decrease with increase in the observability of random-draw outcomes and reports by the experimenter. In particular, we predicted that the average reported number would be:

$$\text{Super Observed} < \text{Observed} < \text{Basic} < \text{Double Blind} < \text{Super Double Blind}.$$

Evidence from cheating games shows that some individuals lie partially; that is, they lie, but they report a number that does not maximize their monetary payoff, typically the second-highest one (Abeler et al., 2019). In Gneezy et al.’s (2018) model, when the experimenter does not observe the random draw, the social image (“social identity”) of reporting a state is decreasing in the fraction of liars who report this state. This provides a rationale for partial lying, because individuals gain a higher social image from reporting a state that is not maximizing their monetary payoff. When the draw is observed, this rationale is shut down because the experimenter can identify partial lies as dishonest actions with certainty.

In our observed treatments, we vary the degree to which the experimenter can detect lying at the time of reporting. With the observed draw, even if the experimenter does not record it, the probability that the experimenter remembers the random-draw outcome is positive. When lie can be detected, social image models predict that participants would prefer full lies over partial lies to at least gain the highest possible monetary payoff. Therefore, we expect more partial lying in the Basic game than in the Observed games. In the Super Observed treatment, the experimenter knows if a participant has lied. In the Observed treatment, the experimenter might not remember individual die rolls at the time of the report, and individuals thus still have a social signaling motive to lie partially. Therefore, the degree of observability in the Observed game presents an intermediate case between the Super Observed and Basic games. Consequently, when the random draw is observed, the incentive to lie partially decreases for Observed games and disappears for Super Observed games. Appendix B formally shows that these predictions follow from a variant of the social image and lying model where individuals are uncertain about whether their reports can be verified. We condense these predictions in Hypothesis 2a.

**Hypothesis 2a (Partial lying by reporting 5 and observability of the random draw):** More partial lying occurs in the Basic treatment than in the Observed treatment, and more partial lying in the Observed treatment than in the Super Observed treatment:

Reported 5s Basic > Reported 5s Observed > Reported 5s Super Observed.

Gneezy et al. (2018) explain the partial lying by the desire of individuals to appear honest. In the Basic, Observed, and Super Observed treatments, the social-image component arises because participants make their report to the experimenter. By contrast, in the Double Blind treatments, social-image concerns should not have a meaningful influence on decisions, because the report is unobserved. Hence, we hypothesize that more partial lying will occur in the Basic treatment than in the Double Blind and Super Double Blind treatments.

**Hypothesis 2b (Partial lying by reporting 5 and observability of the report):** More partial lying occurs in the Basic treatment than in the Double Blind and Super Double Blind treatments:

Reported 5s Basic > Reported 5s Double Blind = Reported 5s Super Double Blind.

If partial lying occurs to preserve social image, the procedure in the Super Observed, Double Blind, and Super Double Blind treatments will result in no partial lies. In the Super Observed treatment, the reason for no partial lying is that participants who decide to lie have no image gains by lying partially, because the experimenter will know they lie with certainty. In both double-blind treatments, the reason for no partial lying is that no outsider is observing the report.

**Hypothesis 2c (Partial lying at the extremes):** Almost no partial lying occurs in the Super Double Blind, Double Blind, and Super Observed treatments.

Hypothesis 2c tests whether partial lies decrease when we deliver complete anonymity or when we observe both the report and the random draw. Combining Hypothesis 1 and Hypothesis 2c implies we expect more participants to lie to the full extent as we decrease observability. In other words, given that the social-image concern decreases as the level of observability decreases, we hypothesize that more participants will report the maximum number as the degree of observability decreases.

**Hypothesis 3 (Full-extent lies and observability):** More full-extent lies will occur when the observability of participants decreases. In particular, we predicted that the fraction of reported 6's would be

$$\text{Super Observed} < \text{Observed} < \text{Basic} < \text{Double Blind} < \text{Super Double Blind}.$$

We also make predictions regarding the interaction between observability and gender. In a meta-study, Abeler et al. (2019) find that female participants report lower numbers than male participants. This finding has been substantiated by an additional meta-study by Gerlach et al. (2019), which documents a robust gender effect across various experimental games that measure dishonesty. Therefore, we expect that, overall, women will report lower numbers than men.

**Hypothesis 4 (Over-reporting and gender):** Across treatments, women will report a lower average number than men.

Moreover, the evidence on gender sensitivity to experimental contexts suggests women are more sensitive to changes in experimental conditions (Croson and Gneezy, 2009; Miller and Ubeda, 2012). Hence, we hypothesize that females' lying behavior will be more affected by both



the observability of random draw and the observability of reports when compared to the lying behavior of men.

**Hypothesis 5a (Gender and observability of the random draw):** The effect of observability of the random draw will be stronger for women than for men.

**Hypothesis 5b (Gender and observability of the report):** The effect of observability of the report will be stronger for women than for men.

### 3 Results


#### *Observability and average reporting*

We use the Wilcoxon signed-rank test in the Super Observed condition to test whether the reported distribution is equal to the distribution of die rolls recorded by the experimenter. In all other treatments, we use the Chi-squared test to determine whether the reported distribution is different from the uniform distribution. In all treatments, we find that the distribution of reported numbers is significantly different from the theoretical distribution of the die rolls ( $p < 0.001$ , two-sided Wilcoxon signed-rank test for the Super Observed condition;  $p < 0.001$ , two-sided Chi-squared test for all other treatments).<sup>7</sup> Figure 1 presents the average reported die-roll outcomes by treatment.

---

<sup>7</sup> In the paper, we use one-sided tests for the comparisons when we had directional hypotheses and two-sided tests otherwise.

**Figure 1: Average reported die rolls**



*Note:* Error bars are 95% confidence intervals. The dashed horizontal line highlights the average report in the Basic condition.

Comparing the results in the Basic treatment with the Super Observed and Observed treatments results, we find that, in line with Hypothesis 1, making the die-roll outcome observed leads to less lying: Lying in the Super Observed game is significantly lower than in the Basic game ( $p < 0.001$ , one-sided Mann-Whitney-U test), and lying in the Observed game is significantly lower than in the Basic game ( $p < 0.001$ , one-sided Mann-Whitney-U test). Comparing reported numbers in the Super Observed and Observed treatments, we find marginally less lying in the Super Observed game than in the Observed game ( $p = 0.068$ , one-sided Mann-Whitney-U test).

**Result 1:** Increasing the observability of the random draw relative to Basic and relative to Observed decreases the average reported die roll.

However, the average and the distribution of reported numbers in Double Blind and Super Double Blind treatments are not significantly different from that of the Basic game ( $p > 0.261$ , one-sided Mann-Whitney U test), rejecting that part of Hypothesis 1.

**Result 2:** Increasing the anonymity of the report relative to the Basic treatment does not affect the average reported die roll.

Taken together, these results provide mixed evidence for Hypothesis 1. Although lying decreases when the observability of the random draw increases, lying is invariant to reducing the observability of the report relative to the Basic treatment.

#### *Observability and the size of lies*

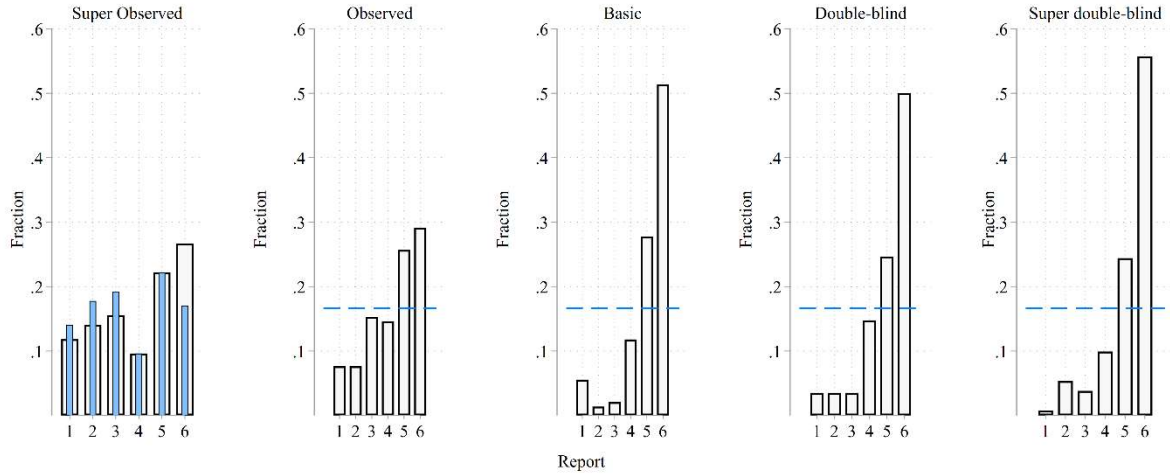
Next, we examine how the distribution of reports changes in the type of lying—full extent and partial—across treatments. We take the reporting frequency of the second-highest state as a measure of partial lying and the reporting frequency of the highest state as a measure of full-extent lying. Figure 2 displays the distributions of reported numbers in each treatment. In the Super Observed treatment, in addition to the reports, we also documented the actual random draws. As can be seen in Figure 2, more than 20% of participants actually rolled a five in the Super Observed game. The conventional methods to measure partial lying (which quantify the fraction of reported fives in excess of  $1/6$ ) would therefore overestimate the extent of partial lies in that treatment.<sup>8</sup> We instead take the fraction of participants who lied by reporting a 5 as a

---

<sup>8</sup> This part of the analysis was not pre-registered.

measure of partial lying in the Super Observed game. We return to the issue of the random-draw distribution in the Super Observed game in the additional results' section.

**Figure 2: Histograms of die-roll reports**



*Note:* The blue vertical bars in the left panel display the observed distribution of random draws in the Super Observed condition. The dashed horizontal lines in the remaining panels display the underlying theoretical distribution of the random draw.

In the Super Observed game, we find that four out of 135 participants over-reported by reporting a 5. We obtain a comparable measure for partial lying in the Observed game by noting that whereas 37 participants reported a 5 in the Observed game, only 24 participants were expected to have drawn a 5, given the sample size of 144. This analysis allows us to establish that *in expectation*, 9.03% of participants in the Observed game lied partially, which is significantly higher than the fraction of participants who lied by reporting 5 in the Super Observed game ( $p = 0.029$ , one-sided Fisher's exact test). We use the same approach to establish that partial lying in the Super Observed game is significantly lower than in the Basic game ( $p = 0.007$ , one-sided Fisher's exact test). However, the frequency of reported 5's is not significantly

different between the Observed and Basic games ( $p = 0.395$ , one-sided Fisher's exact test). Taken together, the results are only partially in line with Hypothesis 2a.

**Result 3:** Observability of the random draw decreases partial lying only if the draw is recorded.

Moving to full-extent lies, we find that fewer participants report a 6 in both observed treatments than in the Basic game ( $p < 0.001$  in both pairwise comparisons; one-sided Fisher's exact tests). The reporting frequencies of 6's are not significantly different when comparing the Observed and Super Observed games ( $p = 0.370$ , one-sided Fisher's exact test). Therefore, Hypothesis 3 is only partially supported by the results. Full-extent lying decreases when making the random draw observable but does not decrease further as the random draw becomes more observable.

**Result 4:** Observability of the random draw decreases full-extent lying.

Result 1 established that lying, as measured by the average reported die roll, decreases in the degree of observability of the random draw. The results on the size of lies further establish that lying decreases when moving from the Basic to Observed game, primarily because of a reduction in full-extent lies as presented in Result 4. This finding is in line with the part of Hypothesis 3 where we predicted that more participants would be encouraged to lie to the full extent as the draw becomes unobservable. Differently, moving from the Observed to Super Observed game mainly affects partial lying. This finding supports Hypothesis 2a —when lies are

detected for sure, participants no longer have an incentive to obscure them by reporting a non-payoff-maximizing state.

As shown in Figure 2, in the double-blind treatments, the reporting frequencies of 5's are higher than  $1/6$ , which is the opposite of what we expected in Hypotheses 2b and 2c. The differences from the theoretical benchmark are significant in both treatments ( $p < 0.015$ , one-sided Binomial test) and imply some participants lied partially in the double-blind treatments. Further, pairwise comparisons between partial lying in the Basic, Double Blind, and Super Double Blind treatments reveal non-significant differences ( $p > 0.310$ , one-sided Fisher's exact test). This result is not in line with Hypothesis 2b.

**Result 5:** Observability of the report does not affect partial lying.

Moreover, full-extent lies, as measured in the frequency of reported 6's, do not significantly differ between the Basic, Double Blind, and Super Double Blind treatments ( $p > 0.274$ , one-sided Fisher's exact test). The result does not support Hypothesis 3.

**Result 6:** Observability of the report does not affect full-extent lying.


### *Gender*

Using the pooled data from all treatments, we find that women report significantly lower numbers than men ( $p = 0.005$ , one-sided Mann-Whitney-U test). This result supports Hypothesis 4.

**Result 7:** Women report lower average numbers than men.

Figure 3 displays the gender differences in reporting for each treatment. It shows that women in almost all treatments report lower numbers than men. These reporting differences become the largest and most significant in the double-blind treatments.

**Figure 3: Average die-roll reports by gender**



*Note:* The p-values are the results of treatment-wise two-sided Mann-Whitney U tests for differences in reported die rolls by gender.

However, the interaction between treatment and gender is small. Table 3 presents regression results that test for heterogeneous treatment effects. No significant interaction effect exists between gender and the Observed or Double Blind treatments. That is, although we find that, on average, over all treatments, women lie less than men, they do not react differently to changes in observability, and therefore Hypothesis 5 has no support.

**Result 8:** Women do not react differently than men to an increase in the observability of the random draw or to a decrease in the observability of the report.

**Table 3. OLS regressions with the report as the dependent variable**

	(1)
Observed treatments	-0.981*** (0.173)
Double blind treatments	0.108 (0.135)
Female	-0.120 (0.071)
Observed treatments X Female	0.057 (0.198)
Double blind treatments X Female	-0.231 (0.152)
Constant	5.149*** (0.109)
Observations	695
$R^2$	0.102

*Note:* Observed treatments is a dummy equal to 1 if the treatment is Observed or Super Observed. Double-blind treatments is a dummy equal to 1 if the treatment is Double Blind or Super Double Blind. Standard errors clustered at the session level are in parentheses. \*\*\* denotes p-value less than or equal to 1%.

*Additional results: Distribution of die rolls in the Super Observed game*

While the die roll distribution in Super Observed is not significantly different from a uniform distribution ( $p=0.173$ , Chi-squared test), it is characterized by a high number of actually rolled fives. To find out how unexpected the distribution is, we simulated 10,000 die rolls for samples with 135 observations (equal to the number of observations in Super Observed). The average die roll in Super Observed (3.59) is in the top third of the averages in the simulation. In addition, we



found that only 6% of the samples have a distribution where the proportion of rolled 5s is higher than 22%, the value in Super Observed. This might raise the question whether our sample size in the other treatments is sufficient to say anything with statistical confidence about treatment differences in lying. In Appendix C, we provide two robustness checks. First, we show that binomial tests that investigate whether a specific number is reported by more than  $1/6$  of all participants do not result in too many false positives even for samples of our size. Second, we provide estimates of the total rate of liars and on the rate of partial liars for different treatments using the Bayesian technique suggested by Hugh-Jones (2019), who provides simulations showing that Bayesian techniques offer more precise measures of the distribution of lying rates than frequentist techniques in small samples. The main results remain similar under the alternative approach.

*Additional results: Results on the mechanism behind lying behavior in the Observed game*

We investigate whether participants are uncertain whether the experimenter remembers their draw. While we did not measure beliefs directly, we have data on the sequence in which participants made their report to the experimenter. Participants may believe that the experimenter is less likely to remember the original draw if a long time elapses between draw and report. Participants reported by their cubicle number, starting with the lowest number and going to the highest. The reporting procedure took around 15 minutes in total, as each participant individually went to a neighboring room to make the report, sign the receipt, and pick up the payment. This contrasts with a relatively short draw procedure, in which the experimenter observed die rolls by cubicle number, but only took around two minutes altogether. If participants anticipate the time

delay between draw and report, they would be more likely to lie if they are later in the queue. We would therefore expect a positive correlation between cubicle numbers and reports.

Indeed, we find that the reports are positively and marginally significantly correlated with cubicle number in the Observed game (Spearman's  $\rho = 0.156$ ,  $p = 0.063$ ). The correlation is negative and insignificant in both the Basic (Spearman's  $\rho = -0.104$ ,  $p = 0.215$ ) and Super Observed (Spearman's  $\rho = -0.120$ ,  $p = 0.167$ ) games. Interestingly, the average report of early reporters (cubicles 1-12) in the Observed game is 4.06, which almost coincides with 3.96, the average report in the Super Observed game. The average report of late reporters (cubicles 13-25) is 4.56, which is between the average reports of the Super Observed and Basic games (5.09). These results are in line with the notion that individual reports differ depending on their position in the queue. They provide some evidence that beliefs about remembering the actual random-draw outcome influence reporting behavior in the Observed game.

## **4 Discussion and Conclusion**

In this paper, we systematically test how individuals adjust their unethical behavior in response to observability. We differentiate between two types of observability: observability of the random draw and observability of the report. We find that interventions that increase the observability of the random draw can facilitate honesty, whereas interventions that increase the observability of the report have no effect on behavior.

Our results indicate that once perfectly identifying liars is possible (Super Observed treatment), partial lying virtually disappears, which supports social signaling as a motivation for partial lying. The results also illustrate that the effects of interventions to promote honesty critically depend on their implementation. Moving from the Basic to Observed game decreases

the size of lies conditional on lying, whereas moving from the Observed to Super Observed game increases the size of lies conditional on lying. Depending on the relative importance of promoting honesty and making lies easier to detect, different types of interventions may be preferred in different contexts.

In contrast to the insights obtained from the comparison between the Observed and Basic games, we find that changes in the observability of the report do not influence lying behavior. In particular, we find no significant difference between *partial lying* in the Basic game than that in the Double Blind and Super Double Blind games. We were somewhat surprised given that the unobservability of the report takes away the rationale for signaling honesty to the experimenter. A possible explanation for the result is that some participants have internalized a preference for appearing honest and act as if they are observed even when they are not (Dufwenberg and Dufwenberg, 2018); they care about what an observer would infer from their report. Under such preferences, the double-blind treatments do not eliminate all signaling motives.

Furthermore, we do not find more lying on the *extensive margin* in the double-blind games than in the Basic die-roll game. Our results suggest the Basic die-roll game is already capturing a high level of cheating – in this game, the average report made by participants is 5.09. As a consequence of the high level of cheating on the extensive margin in the basic game, we face a ceiling effect, which makes it difficult to observe an even higher level of lying in the double-blind treatments. Future experiments should consider procedural changes that lead to lower baseline lying rates, so that there is room to measure increases in lying in double-blind treatments. Experiments with the general population, for example, might be an attractive option, as studies consistently find that student populations are more likely to lie than the general

population (Abeler et al., 2019). Additional treatments where the draw is observed might also allow for increased lying rates when moving to a double-blind procedure.<sup>9</sup>

Our findings that the double-blind procedure does not affect one's cheating behavior are in line with some of the social-preferences literature. For example, Andreoni and Bernheim (2009) study image concerns in dictator-game experiments. When the selfish option is implemented randomly regardless of the actual dictator choice, dictators become less likely to share their endowment. Because the receiver does not know whether the dictator acted selfishly or whether he was simply unlucky, dictators can hide behind the uncertainty provided by the randomization. Relatedly, Ariely et al. (2009) show participants act more pro-socially if their actions are made public to other participants. At the same time, Barmettler et al. (2012) find robust evidence that going from a standard experimental protocol to double-blind procedures does not reduce pro-social behavior in dictator, ultimatum, and trust games as would be predicted by a purely "social" theory of image concerns.

## References

- Abeler, J., Nosenzo, D., Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4), 1115–1153.
- Andreoni, J., Bernheim, B. D. (2009). Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects. *Econometrica*, 77 (5), 1607–16036.
- Ariely, D., Bracha, A., Meier, S. (2009). Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially. *American Economic Review*, 99 (1), 544–555.

---

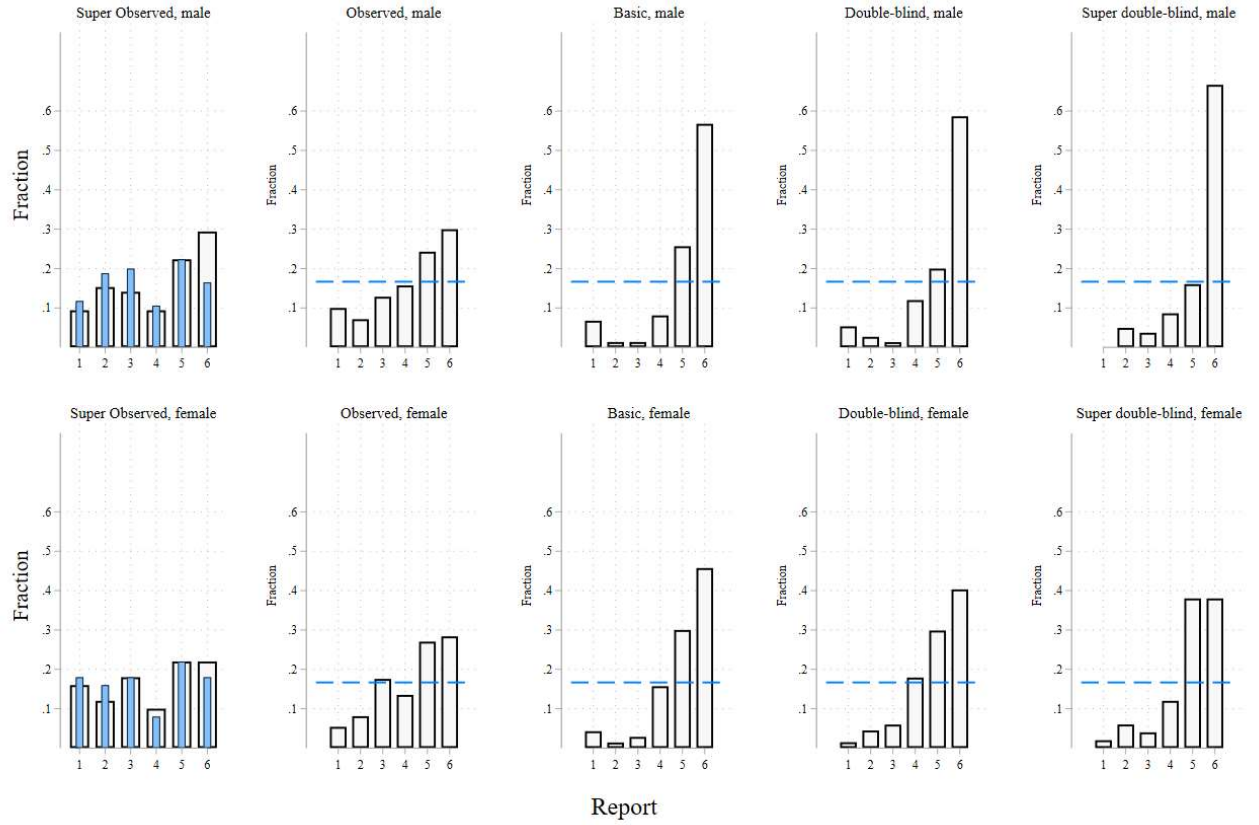
<sup>9</sup> Abeler et al. (2019) implement a treatment with an observed draw and unobserved report and find relatively low lying rates. However, they compare this treatment to a treatment with an unobserved draw and observed report.

- 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.
- Blume, A., DeJong, D. V., Kim, Y.-G., & Sprinkle, G. B. (2001). Evolution of Communication with Partial Common Interest. *Games and Economic Behavior*, 37(1), 79–120.
- Blume, A., Lai, E. K., & Lim, W. (2019). Eliciting Private Information with Noise: The Case of Randomized Response. *Games and Economic Behavior*, 113(1), 356–380.
- Cappelen, A. W., Sørensen, E. T., & Tungodden, B. (2013). When do we lie? *Journal of Economic Behavior and Organization*, 93, 258–265.
- Croson, R., Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, 47(2), 448–474.
- Dreber, A., Johannesson, M. (2008). Gender differences in deception. *Economics Letters*, 99, 197–199.
- Dufwenberg, M., Dufwenberg, M. A. (2018). Lies in disguise—A theoretical analysis of cheating. *Journal of Economic Theory*, 175, 248–264.
- Fischbacher, U., Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, 11(3), 525–547.
- Fries, T., & Parra, D. (2021). Because I (don't) deserve it: Entitlement and lying behavior. *Journal of Economic Behavior and Organization*, 185, 495–512.
- Garbarino, E., Slonim, R., & Villeval, M. C. (2018). A method to estimate mean lying rates and their full distribution. *Journal of the Economic Science Association*, 4(2), 136–150.
- Gerlach, P., Teodorescu, K., Hertwig, R. (2019). The truth about lies: A meta-analysis on dishonest behavior. *Psychological Bulletin*, 145(1), 1–44.
- Gneezy, U., Kajackaite, A., Sobel, J. (2018). Lying Aversion and the Size of the Lie. *American Economic Review*, 108(2), 419–453.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114–125.

- Hugh-Jones, D. (2019). True lies: Comment on Garbarino, Slonim and Villeval (2018). *Journal of the Economic Science Association*, (2018).
- Kajackaite, A., Gneezy, U. (2017). Incentives and cheating. *Games and Economic Behavior*, 102, 433–444.
- Khalmetski, K., Rockenbach, B., & Werner, P. (2017). Evasive lying in strategic communication. *Journal of Public Economics*, 156, 59–72.
- Khalmetski, K., Sliwka, D. (2019). Disguising Lies—Image Concerns and Partial Lying in Cheating Games. *American Economic Journal: Microeconomics*, 11 (4), 79-110.
- Miller, L., Ubeda, P. (2012). Are women more sensitive to the decision-making context? *Journal of Economic Behavior and Organization*, 83(1), 98–104.
- Rode, J. (2010). Truth and trust in communication: Experiments on the effect of a competitive context. *Games and Economic Behavior*, 68, 325–338.
- Serra-Garcia, M., van Damme, E., & Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, 73(1), 244–261.
- Sutter, M. (2009). Deception through telling the truth?! Experimental evidence from individuals and teams. *Economic Journal*, 119, 47–60.
- Wang, J. T. Y., Spezio, M., & Camerer, C. F. (2010). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, 100(3), 984–1007.

## Appendix A. Additional descriptive statistics

Figure A1: Histograms of die-roll reports by gender




*Note:* The blue vertical bars in the left panels display the observed distribution of random draws in the Super Observed game. The dashed horizontal lines in the remaining panels display the underlying theoretical distribution of the random draw.

Table A1. Distribution of observed and reported numbers in the Super Observed treatment

Report	Observed						Total
	1	2	3	4	5	6	
1	16	0	0	0	0	0	16
2	0	19	0	0	0	0	19
3	0	0	21	0	0	0	21
4	0	1	0	12	0	0	13
5	1	0	2	0	27	0	30
6	2	4	3	1	3	23	36
Total	19	24	26	13	30	23	135

Figure A2. Scatter plot of observed numbers and reports in the Super Observed treatment





## Appendix B. Theoretical analysis: effect of increasing the observability of the draw

Khalmetski and Sliwka (2019) (K&S) present a model with a population of agents and an observer. Each agent privately draws a state  $y \in \mathcal{Y} = \{0, 1, \dots, K\}$  and makes a report  $x \in \mathcal{Y}$  to the observer. The state is distributed uniformly. Based on the report, an agent receives a direct payoff equal to  $x$ . In addition, the agent experiences intrinsic and extrinsic lying costs. Intrinsic lying costs are fixed costs  $l > 0$  that agents experience whenever their report is not equal to the draw. The fixed cost parameter is distributed among the population according to a distribution  $F(l)$  with  $F(l) = 0$  for  $l \leq 0$  and  $F(l) > 0$  for  $l > 0$ . Extrinsic costs depend on the observer's belief that an agent has lied conditional on the report,  $P(y \neq x|x)$ . We will refer to this term as the agents' reputation. Agents have a utility function:

$$u(l, x, y) = x - 1_{x \neq y}l - \eta P(y \neq x|x),$$

where  $\eta > 0$  is a number that denotes the agents' image concern, i.e., the weight they put on their reputation. A final assumption of the model is that the maximum lying cost is not too small;  $F(K + \eta) < 1$ .

The setup of the game directly maps into the present experiment, where participants are in the place of the agents and the experimenter takes on the role of the observer.<sup>10</sup> K&S show that the model generates a Perfect Bayesian Equilibrium that is broadly in line with the empirical predictions on die-roll games. In equilibrium, the highest state  $K$  is always reported by some liars

---

<sup>10</sup> If  $K = 5$ . The only meaningful difference is that in our experiment, payoffs range from (2-12) in steps of two while they range from (0- $K$ ) in steps of one in the model. This is a simple re-normalization of the payoffs which does not affect any qualitative predictions of the model.

with positive probability. Agents lie only if they draw a number smaller than some threshold  $x_L$  and report a number larger or equal than  $x_L$ . As a result, the report distribution is increasing in the payoff each state provides, and there is partial lying if the image concern ( $\eta$ ) is not too small.

### *Introducing degrees of observability*

We consider the impact of our observed treatments on the model's predictions. Specifically, we consider the case where agents hold a belief  $\gamma \in [0,1]$  about the probability that the observer knows their draw at the time of making a report. This case is briefly discussed in section IV of K&S. When the observer knows the draw, an agent's reputation is either one or zero, depending on whether they lie or not. The utility function becomes:

$$\begin{aligned} u(l, x, y; \gamma) &= x - 1_{x \neq y} l - \eta[\gamma + (1 - \gamma)P(y \neq x|x)] \\ &= x - 1_{x \neq y}(l + \eta\gamma) - \eta(1 - \gamma)P(y \neq x|x). \end{aligned}$$

The introduction of  $\gamma$  adds an additional fixed cost of lying,  $\eta\gamma$ , as the agent will suffer an extrinsic lying cost of  $-\eta$  whenever the observer knows the draw. It further reduces the social image coefficient by a factor  $(1 - \gamma)$ .

After an increase in  $\gamma$ , lying decreases on the aggregate because fixed costs increase. At the same time, as the social image weight decreases, agents have a smaller motive to lie partially.

Therefore, fewer states are reported by liars. The following proposition makes this point formally by providing comparative statics of the effects of increasing  $\gamma$  on lying at the intensive and extensive margins.

**Proposition 1.** *If the agents' belief about the observer knowing their draw,  $\gamma$ , increases, then:*

1.  $x_L$  weakly increases.
2. The likelihood that an agent lies strictly decreases.

**Proof.** If more than one state is reported by liars in equilibrium, agents must be indifferent between reporting any of those states, because lying costs are fixed. Otherwise, if some agent strictly prefers to lie and report, say,  $K$ , then every other agent who lies should also strictly prefer to report  $K$ . In equilibrium, the gross payoff from lying (without fixed lying costs) is

$$\rho = x - \eta(1 - \gamma)\varphi_x \text{ for all } x \geq x_L(\rho^*(\gamma)),$$

where  $\varphi_x$  is the observer's belief that an agent reporting  $x$  lied. The fraction of liars who draw a state  $y < x_L$  thus is

$$F(\rho - \eta\gamma - y).$$

Summing up over states, we find that the proportion of agents that lie is

$$P(\text{lie}) = \frac{1}{K+1} \sum_{y=0}^{x_L(\rho)-1} F(\rho - \eta\gamma - y). \quad (1)$$

A different way to derive the proportion of liars is by noticing that

$$P(\text{lie}) = \sum_{x=x_L(\rho)}^K P(\text{report } x)P(\text{lie}|\text{report } x).$$

Replacing  $P(\text{report } x) = \frac{1}{K+1} \times \frac{1}{1-\varphi_x}$  and  $P(\text{lie}|\text{report } x) = \varphi_x$ , we arrive at

$$P(\text{lie}) = \frac{1}{K+1} \sum_{x=x_L(\rho)}^K \frac{\varphi_x}{1-\varphi_x}. \quad (2)$$

Finally, we can replace  $\varphi_x$  in equation (2) by  $\varphi_x = (x - \rho)/\eta(1 - \gamma)$  to get to

$$P(\text{lie}) = \frac{1}{K+1} \sum_{x=x_L(\rho)}^K \left( \frac{x - \rho}{\eta(1 - \gamma) - x + \rho} \right). \quad (3)$$

In equilibrium, equations (1) and (3) must coincide. We can define a function

$$\theta(\rho, \gamma) = \sum_{y=0}^{x_L(\rho)-1} F(\rho - \eta\gamma - y) - \sum_{x=x_L(\rho)}^K \left( \frac{x - \rho}{\eta(1 - \gamma) - x + \rho} \right) = 0. \quad (4)$$

This implicitly defines a function  $\rho^*(\gamma)$  such that  $\theta(\rho^*(\gamma), \gamma) = 0$ . K&S show that a solution to this problem always exists and that  $\rho^*$  is unique.<sup>11</sup> Consider some  $0 \leq \gamma' < \gamma'' \leq 1$ . We fix  $\rho^*(\gamma')$  and ask how  $\theta(\rho^*(\gamma'), \gamma)$  changes in response to an increase in  $\gamma$ . The first term in equation (4) decreases in  $\gamma$  while the second term increases. Therefore,  $\theta(\rho^*(\gamma'), \gamma)$  decreases in  $\gamma$ . This implies that

$$\theta(\rho^*(\gamma'), \gamma'') < \theta(\rho^*(\gamma'), \gamma') = 0 = \theta(\rho^*(\gamma''), \gamma'').$$

Since  $\theta(\rho, \gamma)$  increases in  $\rho$ , it follows that  $\rho^*(\gamma'') > \rho^*(\gamma')$ . This holds for any  $0 \leq \gamma' < \gamma'' \leq 1$  and we conclude that  $\rho^*(\gamma)$  increases in  $\gamma$ . As  $x_L$  weakly increases in  $\rho^*$ , increasing  $\gamma$  weakly increases  $x_L$ .

---

<sup>11</sup> See K&S, especially proposition 1 and theorem 1 for a detailed discussion and also for a proof that  $x_L(\rho)$  weakly increases in  $\rho$ . To see that there is always a unique  $\rho^*$ , note that the first term in (4) increases in  $\rho^*$  while the second term decreases. Therefore  $\theta'(\rho) > 0$ , which implies a unique solution. Results by K&S imply that both terms are continuous functions and that  $\theta(\min\{0, K - \eta\}) < 0$  and  $\theta(K) > 0$ , which ensures existence.

(ii). We will argue by contradiction. Assume lying increases after an increase in  $\gamma$  and consider equation (1). A necessary condition for lying to increase in  $\gamma$  is that

$$\rho'(\gamma) - \eta > 0. \quad (5)$$

We have

$$\rho(\gamma) = x - \eta(1 - \gamma)\varphi_x(\gamma) \text{ for all } x \geq x_L(\rho^*(\gamma)),$$

Which has a derivative

$$\rho'(\gamma) = \eta\varphi_x(\gamma) - \eta(1 - \gamma)\varphi_x'(\gamma).$$

plugging into equation(5), the inequality becomes

$$\eta(\varphi_x(\gamma) - 1) - \eta(1 - \gamma)\varphi_x'(\gamma) > 0.$$


This condition can only hold if  $\varphi_x'(\gamma) < 0$ , as  $\varphi_x^*(\gamma) < 1$ . However, equation (2) shows that  $P(\text{lie})$  decreases as  $\varphi_x$  decreases, a contradiction.

## Appendix C. Simulations

In the die roll game, participants draw a number from a uniform distribution without replacement. We only have information about the actual distribution in the Super Observed game. In the other treatments, we rely on the underlying theoretical distribution to make inferences. Hence, all claims regarding lying rates, treatment differences, etc., only hold in expectation. However, we can assess whether our sample sizes are large enough to use standard frequentist inference, which depends on the convergence speed of die rolls. In a related exercise, Fries and Parra (2021) show that frequentist methods produce accurate confidence intervals for the average die roll in the die roll game when the sample is about 100 observations or more.

For measuring partial lying, we compare the frequency of reported 5s with the expected probability of success of 16.66% using a binomial test. Given the random nature of the distributions, the realization of the die rolls might generate a distribution with a lot of 5s. In this case, we could infer that there is partial lying when participants report the 5s truthfully. We simulate die rolls to assess how our sample size could be problematic to draw inferences about partial lying. Figure C.2 plots the proportion of significant binomial tests in 10,000 simulations for different sample sizes at various significance levels using two-sided p-values; we corroborate that the probability of a type I error is below the significance level for each sample size.

Figure C2. Proportion of significant binomial tests in a random sample by the significance level



Note: Simulations are based on 10,000 repetitions for each sample size. The dashed line presents simulated binomial tests using a significance level of 10% and the solid blue line uses a 5% significance level.


### *Lying rates using Hugh-Jones (2019) method*

As a final robustness check, we use the Hugh-Jones (2019) method, which uses a Bayesian approach instead of a frequentist approach. Hugh-Jones presents a method to estimate lying rates in binary lying games. To estimate lying rates, we binarize our games by splitting the die rolls into “high” (numbers 5 and 6) and “low” (numbers 1-4), as suggested by Garbarino et al. (2018).<sup>12</sup> The technique estimates the excess number of high reports to what would be expected under full honesty, where we would expect 1/3 of the sample to report “high”. In the Super Observed game, we can directly estimate the lying rates by classifying lies as “high,” true reports as “low,” and where we would expect zero “high” reports under full honesty. We also drop

<sup>12</sup> The lying models by Gneezy et al. (2018) and Khlametski and Sliwka (2019) predict that lying decisions are essentially binary, even if the game has more than two states, i.e. that there is a set of “high” states that is lied at and that there is a set of “low” states where individuals might lie from. These models predict that all states that are over reported are high states while other states are low.

observations that rolled a six in Super Observed since they did not have an incentive to lie. In all estimations we impose a uniform prior. The lying rates of each treatment using this method are presented in Figure C.3. If one doubts that a sample size is too small, the Bayesian method should be better suited to conduct statistical analyses. As we can see in Figure C.3, the resulting patterns are similar to what we presented in the paper.

Figure C3. Lying rates using Hugh-Jones (2019) method by treatment



*Note:* Error bars are 95% confidence intervals (highest density regions).

Finally, we use the method to estimate the number of participants lying to 5 in the Super Observed and Observed games. In the Super Observed game, the estimated number of participants from states other than 5 lying to 5 is 3.85% (95% CI: 0.007%-7.00%) and it is 11.26% (95% CI: 3.00%-19.68%) in the Observed game.



## **Appendix D. Instructions**

### **/Treatment: Basic/ Instructions**

Welcome to our experiment!

Please read the instructions carefully. If you have a question, please raise your hand. We will then come over to you and answer your question in private.

During the experiment you are not allowed to use electronic devices or to communicate with other participants. If you violate these rules, we must exclude you from the experiment and the payout.

You have already received 5 Euros for showing-up to the experiment. During the experiment you will be able to earn additional money, as we explain on the next page (please turn the page).

On your desk you have a six-sided die and an envelope containing 12 Euros.

We will soon ask you to roll the die in private (you will be told at what time you are supposed to roll the die), and report the number that came up.

The number you report will determine how much money you will be paid. You will be paid the equivalent in Euros to the number you report, multiplied by 2. In other words, if you report 1, you receive 2€, if you report 2, you receive 4€, if you report 3, you receive 6€ and so on.

The following table illustrates all possible payoffs:

Reported number	1	2	3	4	5	6
Payoff	2 €	4 €	6 €	8 €	10 €	12 €

There will be no further rounds in this experiment. That is, you participate in the task described above only once and this task is the only task you will have today in the session.

When every participant is done, we will ask you to go to the other room and sign a receipt. Please bring this sheet, the envelope with the remaining money and the seat number with you. Please leave the pen on your desk.

If you have any questions please raise your hand and we will come to you.

Now please wait until you are told to roll the die.

{participants are asked to roll the die}

---

**Please report the number that came up in the die roll.**

**The number that I saw is: \_\_\_\_\_**

Now please take the amount of Euros equivalent to the number you reported multiplied by two out of the envelope and put it in your wallet/pocket.

Thank you for participating!

**/Treatment: Observed/  
Instructions**

Welcome to our experiment!

Please read the instructions carefully. If you have a question, please raise your hand. We will then come over to you and answer your question in private.

During the experiment you are not allowed to use electronic devices or to communicate with other participants. If you violate these rules, we must exclude you from the experiment and the payout.

You have already received 5 Euros for showing-up to the experiment. During the experiment you will be able to earn additional money, as we explain on the next page (please turn the page).

On your desk you have a six-sided die and an envelope containing 12 Euros.

We will soon ask you to roll the die in front of the experimenter (the experimenter will come over to you at the time at which you are supposed to roll the die), and report the number that came up.

The number you report will determine how much money you will be paid. You will be paid the equivalent in Euros to the number you report, multiplied by 2. In other words, if you report 1, you receive 2€, if you report 2, you receive 4€, if you report 3, you receive 6€ and so on.

The following table illustrates all possible payoffs:

Reported number	1	2	3	4	5	6
Payoff	2 €	4 €	6 €	8 €	10 €	12 €

There will be no further rounds in this experiment. That is, you participate in the task described above only once and this task is the only task you will have today in the session.

When every participant is done, we will ask you to go to the other room and sign a receipt. Please bring this sheet, the envelope with the remaining money and the seat number with you. Please leave the pen on your desk.

If you have any questions please raise your hand and we will come to you.

Now please wait until the experimenter will come to you for the die roll.

{participants are asked to roll the die}

---

**Please report the number that came up in the die roll.**

**The number that I saw is: \_\_\_\_\_**

Now please take the amount of Euros equivalent to the number you reported multiplied by two out of the envelope and put it in your wallet/pocket.

Thank you for participating!

**/Treatment: Double blind/  
Instructions**

Welcome to our experiment!

Please read the instructions carefully. If you have a question, please raise your hand. We will then come over to you and answer your question in private.

During the experiment you are not allowed to use electronic devices or to communicate with other participants. If you violate these rules, we must exclude you from the experiment and the payout.

You have already received 5 Euros for showing-up to the experiment. During the experiment you will be able to earn additional money, as we explain on the next page (please turn the page).

On your desk you have a six-sided die and an envelope containing 12 Euros.

We will soon ask you to roll the die in private (you will be told at what time you are supposed to roll the die), and report the number that came up.

The number you report will determine how much money you will be paid. You will be paid the equivalent in Euros to the number you report, multiplied by 2. In other words, if you report 1, you receive 2€, if you report 2, you receive 4€, if you report 3, you receive 6€ and so on.

The following table illustrates all possible payoffs:

Reported number	1	2	3	4	5	6
Payoff	2 €	4 €	6 €	8 €	10 €	12 €

There will be no further rounds in this experiment. That is, you participate in the task described above only once and this task is the only task you will have today in the session.

When all participants will be done, we will ask you to put your envelopes in the brown box at the entrance to the room. You can leave the lab immediately afterwards. Please leave everything else on your desk.

If you have any questions please raise your hand and we will come to you.

Now please wait until you are told to roll the die.

{participants are asked to roll the die}

---

**Please report the number that came up in the die roll.**

**The number that I saw is: \_\_\_\_\_**

Now please take the amount of Euros equivalent to the number you reported multiplied by two out of the envelope and put it in your wallet/pocket.

Please place this reporting sheet inside the envelope with the remaining money and seal the envelope.

Thank you for participating!

**/Treatment: Super observed/  
Instructions**

Welcome to our experiment!

Please read the instructions carefully. If you have a question, please raise your hand. We will then come over to you and answer your question in private.

During the experiment you are not allowed to use electronic devices or to communicate with other participants. If you violate these rules, we must exclude you from the experiment and the payout.

You have already received 5 Euros for showing-up to the experiment. During the experiment you will be able to earn additional money, as we explain on the next page (please turn the page).

On your desk you have a six-sided die and an envelope containing 12 Euros.

We will soon ask you to roll the die in front of the experimenter (the experimenter will come over to you at the time at which you are supposed to roll the die), and report the number that came up. The experimenter will watch you roll the die and note the number you rolled on her/his sheet.

The number you report will determine how much money you will be paid. You will be paid the equivalent in Euros to the number you report, multiplied by 2. In other words, if you report 1, you receive 2€, if you report 2, you receive 4€, if you report 3, you receive 6€ and so on.

The following table illustrates all possible payoffs:

Reported number	1	2	3	4	5	6
Payoff	2 €	4 €	6 €	8 €	10 €	12 €

There will be no further rounds in this experiment. That is, you participate in the task described above only once and this task is the only task you will have today in the session.

When every participant is done, we will ask you to go to the other room and sign a receipt. Please bring this sheet, the envelope with the remaining money and the seat number with you. Please leave the pen on your desk.

If you have any questions please raise your hand and we will come to you.

Now please wait until the experimenter will come to you for the die roll.

{participants are asked to roll the die}

---

**Please report the number that came up in the die roll.**

**The number that I saw is: \_\_\_\_\_**

Now please take the amount of Euros equivalent to the number you reported multiplied by two out of the envelope and put it in your wallet/pocket.

Thank you for participating!



**/Treatment: Super double blind/  
Instructions**

Welcome to our experiment!

Please read the instructions carefully. If you have a question, please raise your hand. We will then come over to you and answer your question in private.

During the experiment you are not allowed to use electronic devices or to communicate with other participants. If you violate these rules, we must exclude you from the experiment and the payout.

You have already received 5 Euros for showing-up to the experiment. During the experiment you will be able to earn additional money, as we explain on the next page (please turn the page).

On your desk you have a six-sided die and an envelope containing 12 Euros (in 1€ coins) and 10 blank metal coins.

We will soon ask you to roll the die in private (you will be told at what time you are supposed to roll the die), and report the number that came up.

The number you report will determine how much money you will be paid. You will be paid the equivalent in Euros to the number you report, multiplied by 2. In other words, if you report 1, you receive 2€, if you report 2, you receive 4€, if you report 3, you receive 6€ and so on.

The following table illustrates all possible payoffs:

Reported number	1	2	3	4	5	6
Payoff	2 €	4 €	6 €	8 €	10 €	12 €

There will be no further rounds in this experiment. That is, you participate in the task described above only once and this task is the only task you will have today in the session.

After you will report the number, we will ask you to take your payout out of the envelope, place this sheet inside the envelope with the remaining money and blank metal coins, and then seal the envelope.

When all participants will be done, we will ask you to put your envelopes in the brown box at the entrance to the room. You can leave the lab immediately afterwards. Please leave everything else on your desk. **Differently from other experiments, no receipts will be signed at the end of this session.**

If you have any questions please raise your hand and we will come to you.

Now please wait until you are told to roll the die.

{participants are asked to roll the die}

---

**Please report the number that came up in the die roll.**

**The number that I saw is:** \_\_\_\_\_

Now please take the amount of Euros equivalent to the number you reported multiplied by two out of the envelope and put it in your wallet/pocket. Note: The envelope contains additional blank metal coins so that one cannot indicate from the sound of the envelope how much money a participant earned.

Please place this reporting sheet inside the envelope with the remaining money and seal the envelope.

Thank you for participating!