

UC San Diego

UC San Diego Previously Published Works

Title

Why do People not Lie More? Truth is (Wrongly) Believed to be More Persuasive

Permalink

<https://escholarship.org/uc/item/5cs162sm>

Journal

The Economic Journal

ISSN

0013-0133

Authors

Gneezy, Uri

Serra-Garcia, Marta

Publication Date

2026-08-22

DOI

10.1093/ej/ueaf128

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed

Why Don't People Lie More? Truth Is (Wrongly) Believed To Be More Persuasive

Uri Gneezy* and Marta Serra-Garcia†

October, 2025

Abstract

Is truth believed to be more persuasive than falsehood? This paper explores this question using a series of experiments. First, a survey experiment reveals that participants consistently believe the persuasive power of truthful messages is higher than that of lies. Second, two laboratory experiments, in which senders record truthful and false video messages about news events, show that senders *mistakenly* believe their truthful messages will be more believable. Even when incentivized to lie, most senders choose to tell the truth—if tasked with persuading receivers. If not, however, most senders follow the incentives and lie.

JEL Classification: C91, D83, D91, L82.

Keywords: persuasion, lying, experiment.

*Rady School of Management, UC San Diego. Email: ugneezy@ucsd.edu.

†Corresponding author. Rady School of Management, UC San Diego and CESifo. Email: mserragar-cia@ucsd.edu. This research was conducted under IRB 161827. We are grateful to participants at numerous seminars and conferences for their feedback.

1 Introduction

Being persuasive directly influences decision-making, trust, and the ability to achieve mutually beneficial outcomes (e.g., Cialdini 2007; Charness and Dufwenberg 2006; Coffman and Niehaus 2020). Whether negotiating contracts, marketing products, or presenting investment opportunities, effectively convincing others to adopt a particular perspective can lead to successful transactions and long-term relationships. In responding to such acts of persuasion, receivers must work hard to distinguish truth from falsehood.

When senders lie, they manipulate information, potentially manipulating the receiver. When they tell the truth, they provide accurate information that could help the receiver make informed choices, if they believe the message (e.g., Crawford and Sobel 1982). Failing to distinguish between truth and falsehood can have two important repercussions. It can cause the receiver to make Type 1 errors (believing falsehood), or Type 2 errors (disbelieving the truth). Both errors can lead receivers to make poor decisions.

This paper reports on three experiments focused on senders’ beliefs about persuasion, to examine what senders believe about the likelihood that receivers believe their messages, conditional on being true or false.¹ Through these experiments, we show that senders often exhibit miscalibrated beliefs about persuasiveness, overestimating how believable true messages are, relative to false messages.

We start with a survey experiment illustrating that individuals believe the truth is more persuasive than lies. This online experiment, with a quota-representative sample of US individuals ($N = 1,011$), presented individuals with three different scenarios in which a sender was portrayed as telling the truth or lying. The scenarios simulated face-to-face interactions between a sender and a receiver. In all scenarios, the message was kept fixed, and participants were asked whether they thought the receiver would believe the sender when what they said was true and when it was false.

The main findings is that, across all scenarios, individuals believed the message was significantly more likely to be believed when it was true than when it was false, with the difference ranging between 19 and 26 percentage points. This survey experiment provides

¹DellaVigna and Gentzkow (2010) survey empirical evidence on the effects and the drivers of persuasive communication, focusing on how receivers of communication are persuaded.

evidence that in everyday interactions, individuals believe that being truthful is more persuasive than lying.

Is that belief correct? To answer that question and understand how that belief influences communication, our second experiment was based on a sender-receiver game. Our senders were UCSD students in the laboratory who recorded two videos—one in which they told the truth about a news event and another in which they lied about a news event ($N = 200$ videos). For each video, we incentivized senders to convince receivers of the video’s veracity. These videos were then shown to online participants in the role of receivers, who were incentivized to guess whether the sender was being truthful, for each video.

Consistent with the survey experiment, senders expected their true videos to be believed significantly more often than the false ones: They believed true videos would be persuasive 67% of the time, and false videos 51% of the time. These beliefs were miscalibrated and, in line with prior research (e.g., Bond and DePaulo 2006; Belot and Van de Ven 2017; Dwenger and Lohse 2019; Serra-Garcia and Gneezy 2021), receivers were not much better than chance at this task: they believed true videos 56% of the time, and false videos 52% of the time. This experiment revealed a further miscalibration in senders’ beliefs: What they believed about the persuasiveness of true videos was positively correlated with their actual success, but overly optimistic, whereas the corresponding belief about false videos was uncorrelated with their actual persuasiveness.

Together, the findings of these two experiments led us to ask how communication might change when persuasion matters: Are senders’ decisions to be truthful affected by this need to persuade? To answer this question, we conducted our third experiment ($N = 276$ senders), in which we provided senders with a choice to record only one video—either a true or a false one. Across four treatments, we varied the senders’ payoffs: some depended on whether senders told the truth or lied, and some on whether they persuaded a receiver or not.

When we incentivized senders to record a video that persuaded receivers, and the payoff for sending a truthful or a false video was the same, a majority of senders (94%) chose to record a true video. When we increased the incentive to lie, by paying double if the receiver believed a false video, a majority of senders (67%) still chose to record a true video, but the percentage of those who lied rose significantly from 6% to 33%.

The choice to remain truthful could be driven by several motives, including senders' intrinsic aversion to lying (e.g. Gneezy 2005; Gneezy et al. 2013; Abeler et al. 2014 and Abeler et al. 2019). In this paper, we propose a potential additional reason for truth-telling, by asking: Is the need to be persuasive also what motivates truth telling? To answer this question, we introduced two treatments in our experiment that used the same payoffs but no longer gave senders any incentive to persuade receivers. Senders were paid solely based on whether they chose to lie or be truthful.

In these two additional treatments, when senders received the same payoff regardless of whether they told the truth or lied, a majority (90%) of them choose to record a true video, as before. However, when we doubled the incentives to lie and removed the persuasion incentive, only 19% of senders choose to record a true video, and a majority (81%) choose to lie. This percentage is significantly larger than the 33% of senders who chose to lie when they had the same incentives but needed to persuade receivers to be paid. It shows that beliefs about how persuasive truths and lies are can strongly determine how honest or dishonest senders are.

Our findings reveal a fundamental bias in how individuals approach communication and persuasion. Related research has shown that individuals can have biased beliefs over the beliefs of others in updating tasks (e.g., Camerer et al. 1989; Gilovich et al. 1998; Biais and Weber 2009; Madarász et al. 2023; Agranov and Detkova 2024). In addition, senders may lie even when incentivized to be truthful—if they anticipate receivers believe in lies (Thaler 2021). In this paper, we focus on a new explanation for why people may choose to tell the truth: They overestimate the power of truth and may therefore choose to tell the truth even when lying might yield a more self-serving outcome. This miscalibrated belief has broader implications for economic interactions, legal environments, and social behaviour. It suggests that senders, who tell the truth because they overestimate how persuasive their statements are, could underinvest in making their communication credible. Additionally, individuals may fail to recognize how others might exploit lies to achieve their favored outcomes and may not be well prepared to defend themselves against deceptive tactics.

2 Experimental Design

We present the design of each of the three experiments we conducted in what follows (further details and instructions for each experiment in Online Appendix B).

2.1 Experiment 1: Survey Experiment

Participants were presented with three scenarios (Eckel and Petrie 2011). The scenarios were described in words and with illustrations (for instructions, see Online Appendix B). The first scenario portrayed a car sale, where the potential buyer asks the seller whether their dog was transported in the car. The seller answers negatively. In the second scenario, two friends are having a meal, and the one who cooked it asks whether their friend liked the food. The friend answers affirmatively. In the third scenario, a student who has written an essay shows it to their friend and asks whether they liked it. The friend answers affirmatively.

In each scenario, participants were asked whether they thought that the answers would be believed in two potential scenarios: if the individual is telling the truth (e.g., they actually liked the meal), and if they are lying (e.g., stating they liked the food, although they did not).

We recruited 1,011 participants on Prolific Academic (Peer et al. 2017), with a recruitment sampling that was quota-representative of the US population in terms of gender, age, ethnicity and political affiliation. By recruiting a diverse pool of participants, we are able to see how extended differences in perceived persuasiveness are across different groups, and whether they are limited to certain groups of the population. For example, with respect to gender (e.g., Dreber and Johannesson 2008).

We hypothesized that participants would believe that truthful individuals are more persuasive, as stated in Hypothesis 1.

Hypothesis 1. *In all three scenarios, participants will believe that individuals are more persuasive when they tell the truth.*

2.2 Experiment 2: Beliefs Experiment

This experiment consisted of two different tasks, depending on participants’ roles—senders or receivers. We describe the design of the experiment for each role in what follows.

Sender. Senders in the experiment were undergraduate students at UC San Diego. They participated in a session that consisted of two parts. In the first part, each sender generated two 30-second videos: one “false” and one “true.” We chose to use videos to allow senders and receivers to use language and facial expressions to persuade and to evaluate the truthfulness of the videos, respectively. This design builds on Serra-Garcia and Gneezy (2021), which studies receivers’ ability to detect lies, with a set of new senders and 200 new videos, to examine senders’ beliefs. This rich design differs from the simpler scenario-based setting of Experiment 1. It allows senders to lie in different ways and both senders and receivers could potentially form beliefs about persuasiveness based on the (wrong) cues from the videos (see, e.g., Serra-Garcia and Gneezy 2025).

Senders sat in front of a computer and saw two screens in random order. When the screen showed a headline and picture from the *New York Times*, we asked them to describe it in the 30-second video. When the screen was blank, they could make up a news event, or report that the screen was blank. If they chose to make up a news event, senders were instructed that they should not describe a real event but rather make one up. We incentivized senders to convince receivers that they were seeing a news event in both cases: for each of the two videos they recorded, senders received \$10 if they convinced a receiver that they were seeing a true news event, and \$0 if they did not.

Senders were informed about the incentives of receivers, and receivers were informed about the incentives of senders. This information is important, as receivers’ beliefs regarding the motivation of the senders are expected to affect their beliefs and behaviour (e.g., Blume et al. 1998). When considering reaction to news, the source of information is important because it reveals its motivation. For example, when watching a video clip regarding the election on CNN one might have different expectations than when watching it on Fox News, as the two networks may have different motivations. Similarly, a person may believe communication from a car salesperson less than one from a friend (see the discussion in

Gneezy 2005). We did not vary these factors in the current experiment.

Specifically, within the experiment, senders knew that they would be matched with a receiver, who would receive \$1 if they accurately guessed whether the sender’s video is true or false. Senders also knew their videos could be shown to other participants (as we did). At the same time, receivers knew that senders were incentivized to convince the receiver that they were seeing a news event, regardless of what they actually saw.

In the second part of the experiment, we measured senders’ beliefs regarding the persuasiveness of each of their two videos. We told senders that 10 other participants would watch their videos and be incentivized to guess if the video is true or false. They were not given information about how 10 other participants would be recruited. Senders were asked two questions. The first one was: “Out of the 10 participants who will watch your video when you saw a picture and title of a news event, how many do you think will guess that your video was true?” The second one was: “Out of the 10 participants who watch your video when you saw a blank screen, how many do you think will guess that your video was false?” They were allowed to provide any integer number between 0 and 10 as their answer. They received \$2 for each correct guess.

This design allows us to observe the same sender lying and telling the truth, and (unexpectedly) also to measure their beliefs in each situation. By doing so, we decouple behaviour (the choice to lie or tell the truth) from beliefs, complementing previous designs that have elicited first-order and second-order beliefs of senders, after they have chosen whether to tell the truth or lie in abstract sender-receiver games (e.g., Peeters et al. 2015; Grabova et al. 2023).

The sample consists of 100 senders, resulting in a total of 200 videos. Those in which they described a real news story, we describe as truthful; those in which they made up a story, we describe as false. A research assistant checked the video transcripts and determined that when asked, an overwhelming majority of participants did indeed make videos describing made-up news events. To avoid bias, we chose not to exclude any videos based on this criterion.

Senders received payment either for Part 1 or Part 2 of the experiment. Receivers were recruited later as described below. If a sender received compensation for participating

in Part 1, they were matched to a receiver to calculate the sender’s payment. For senders compensated for their participation in Part 2, we randomly selected 10 receivers to determine the correct guess and the sender’s payment. Among the senders, 59% identified as female, 40% as male, and 1% as other.

News Events. We systematically chose pictures and headlines from the front page of the international edition of the *New York Times*. Most news events reported on the front cover deal with international conflicts (e.g., the meeting of South and North Korean leaders), international catastrophes (e.g., an airplane crash), and news from the arts (e.g., a new movie by a renowned actor).

To avoid experimenter selection bias in sampling news events, we decided (before we saw the headlines) to include all front-page headlines with a picture consecutively.

Receivers. Receivers watched 20 videos, randomly drawn from the set of 200 videos recorded. After each video, receivers guessed whether the sender in the video was telling the truth or lying. Receivers knew that half the videos in the study were truthful and half were false, and that 20 videos out of the complete set of 200 would be randomly drawn and shown to them. Receivers were informed that one of the 20 videos they evaluated would be randomly selected. They would be paid \$5 if they guessed correctly whether the sender in the selected video was telling the truth or lying.

To ensure receivers’ attention, the experiment included a question at the beginning that required sound recognition, and participants could not participate if they failed this test. The experiment included understanding questions and a captcha verification. When watching the senders’ videos, receivers had to stay on the video question for at least the duration of the video (25 seconds) but could remain longer if they wished to. The median time they spent on a video was 37.5 seconds—and the 90th percentile spent 58 seconds, less than a minute. As we will see below, their inability to distinguish truth from lies indicates that they did not check outside the experiment to verify the truthfulness of videos.

To test whether receivers believed that truthful videos were generally more likely to be correctly detected than lies, receivers answered two questions about the beliefs of other receivers. They were asked to consider 10 participants in their role and guess how many

were correct when the video was truthful, and when the video was false. They provided their answer on a slider ranging from 0 to 10 and received \$1 for each correct answer. We also elicited receivers’ confidence in their ability to correctly detect true and false videos, in absolute and relative terms.² Following the preregistration we invited 200 participants, and 193 participants passed the attention checks and completed the experiment. They were recruited on Prolific Academic during the Spring of 2023, from the US sample, and their location was verified via an IP check in the survey.

In Serra-Garcia and Gneezy (2021) we find that, in a similar design, receivers’ beliefs about the truthfulness of true and false videos do not differ. In experiments in psychology, receivers have also been found to display a truth default: believing others are truthful most of the time (for a review, see Levine 2014). In this paper, we focus on senders and propose a new hypothesis regarding senders’ beliefs about receivers’ perceptions of their true and false videos: Senders will wrongly believe that true videos are significantly more persuasive than lies.

Hypothesis 2. *Senders believe that truthful videos are more persuasive.*

2.3 Experiment 3: Choice Experiment

Experiment 3 builds on the design of Experiment 2, but instead of being asked to record two videos, senders were asked to choose whether to record just one truthful or false video. There were a total of 276 senders in this experiment, all undergraduate students at UC San Diego.

The experiment consisted of two parts. In Part 1, senders were told that they would be making a video and that they had two options: they chose to see a picture and description of a news event from the *New York Times*, in which case they would be recording a truthful video, or they chose to see a blank screen and make up a news event, in which case they

²Receivers watched the 20 videos in blocks of five, presented in random order. After each block, we asked receivers to state their confidence in their absolute ability to detect lies: “How many of the 5 video guesses you just made do you believe are correct?” Receivers earned \$1 each time their belief was correct. After watching all 20 videos, we elicited receivers’ relative confidence in their ability, using the question, “Compared with previous participants in this experiment, how well do you think you did?” They were asked to choose a quartile and earned \$1 if their answer was correct.

would be choosing to lie.

In Part 2, senders answered two unannounced questions, in each case reporting their beliefs about the persuasiveness of a true and a false video made by other senders who had participated in the experiment before them (in Experiment 2). These questions were based on those asked in Experiment 2, but now asked about the beliefs regarding the persuasiveness of other senders, rather than the participants themselves.

We told all senders that we would randomly draw Part 1 or Part 2 to determine their payment. If Part 1 was selected, senders went on to record the video they chose in a different room. If Part 2 was selected, they were paid according to the accuracy of their beliefs. Part 1 was drawn 10% of the time, which meant that only 22 videos were recorded for this experiment. Part 2 was drawn 90% of the time.

We used a 2-by-2 between-subjects design that varied whether the sender’s payoff was determined by the receiver’s belief, and the amount of the sender’s payoff when they chose to lie. In “10-10 Receiver,” senders received \$10 if the receiver believed the sender was truthful in the video—regardless of whether the video was truthful or false. In “20-10 Receiver,” again senders were paid if the receiver believed the sender was truthful, but their payment depended on which video they recorded: If receivers believed them, senders received \$20 if they recorded a false video and \$10 if they chose to record a truthful video.

The other two treatments involved the same payment amounts, but payment did not depend on whether receivers believed the videos. In 10-10 Fixed and 20-10 Fixed, senders were paid regardless of whether they were believed by a receiver.

We hypothesized that individuals would respond to incentives. Specifically, we hypothesized that they would be truthful almost all the time in the 10-10 treatments and would lie in some cases in the 20-10 treatments (Hypothesis 3). We also hypothesized that individuals would be more likely to lie if they did not have to persuade a receiver than if they had to persuade them (Hypothesis 4).

Hypothesis 3. *In the 20-10 Receiver and Fixed treatments there is more lying than in the 10-10 Receiver and Fixed treatments.*

Hypothesis 4. *In the 20-10 Fixed treatment there is more lying than in the 20-10 Receiver*

treatment.

Additionally, Experiment 3 allows us to understand the relative relevance of beliefs about persuasion for the decision to tell the truth. For example, how important is it compared with lying aversion. If truth-telling is entirely by senders’ lying aversion, the impact of incentives to lie should be similar in the Receiver as in the Fixed treatments, as “turning off” the motive to persuade should not affect truth-telling significantly. By contrast, if individuals in the 20-10 Receiver treatment choose to tell the truth partly because of their beliefs about the persuasiveness of true videos, turning off the motive to persuade (in the 20-10 Fixed treatment) should reduce truth-telling. Therefore, this treatment identifies whether individuals put some weight on their beliefs about persuasiveness when choosing whether to tell the truth.

All experiments were pre-registered on Aspredicted.org and are shown in Online Appendix C. Further details of the implementation together with the experimental instructions are presented in Online Appendix B.

3 Results


In the following sections we report the results of the three experiments, starting with the survey experiment (Experiment 1).

3.1 Survey Experiment

As shown in Figure 1, when individuals consider the interaction between two individuals regarding a car sale, they are significantly more likely to believe the truth will be persuasive (46%) compared to a lie (27%). The same pattern emerges if individuals consider a scenario in which the persuasiveness of a friend’s writing is evaluated by the same criteria (87% vs. 68%). In the case of a dinner invitation, individuals believe that telling the host the dinner was good will almost always be believed when it is true (98%) but only 71% of the time when it is a lie. In all cases the difference is statistically significant (t -tests, p -value < 0.01), in line with Hypothesis 1. Detailed regression results are presented in Online Appendix Table

A.1.

Figure 1: Believed Persuasiveness in Survey Experiment




Notes: Believed persuasiveness in the three scenarios considered in the survey experiment. Scenario A. is a car sale and regards a question about the owner's dog, scenarios B and C regard interactions in which an individual asks their friend whether they liked their writing or their dinner, respectively. Respondents indicated whether the (same) statements would be believed when they were true and when they were false. Number of respondents is 1,011, with 6 observations per respondent. The gray bars indicate 95% confidence intervals.

Figure 2 shows that this result is found across different subgroups of the population. Controlling for other individual characteristics and scenario-fixed effects (Table D.1), we do not observe significant gender differences in beliefs about persuasion. With regards to age, a comparison of the youngest age group (18-24) and those who are older (25-34) shows that those who are older are 9 percentage points (p.p.) significantly more likely to believe a message that is false (t -stat=2.89, p -value< 0.01) and 13 p.p less likely to believe a message that is true (t -stat= 4.00, p -value< 0.01). Across ethnicities, we find that Black participants believe significantly less in the persuasiveness of true messages (8 p.p., t -stat= 2.95, p -value< 0.01) when controlling for other individual characteristics.

The documented differences in believed persuasiveness could stem from actual differences in persuasiveness, if receivers of truths and lies can tell them apart. The existing literature on lie detection in psychology (e.g., Bond and DePaulo 2006) and economics (e.g., Belot and Van de Ven 2017; Serra-Garcia and Gneezy 2021) suggests that receivers rather exhibit significant difficulties in distinguishing truths and lies and hence that the documented

Figure 2: Believed Persuasiveness in Survey Experiment, by Socio-demographic Group




Notes: Believed persuasiveness across different sociodemographic groups (gender, age, political affiliation and ethnicity) based on estimated averages controlling for other individual characteristics and scenario fixed effects. Number of respondents is 1,011, with 6 observations per respondent. The gray bars indicate 95% confidence intervals.

differences in believed persuasiveness are incorrect. However, there is substantial heterogeneity in the existing studies of lie detection and, in the above scenarios, we cannot compare believed persuasiveness to actual persuasiveness. We thus analyze beliefs in Experiment 2, where such a comparison is possible.

3.2 Beliefs Experiment

After recording truthful and false videos, senders reported their beliefs about the persuasiveness of their videos. Senders believe that 67% of receivers will believe them when they are telling the truth. By contrast, they believe that 51% of receivers will believe them when they are lying. These beliefs overestimate the persuasiveness of telling the truth, as shown in Figure 3. When senders tell the truth, they persuade receivers in 56% of the cases, which is significantly less often than the 67% predicted by senders (t -stat= 4.39, p -value<0.01, Table A.3).

Figure 3: Sender beliefs about persuasiveness and actual persuasiveness



Notes: Senders' believed persuasiveness of false (red) and true videos (blue) is shown in solid bars. The actual persuasiveness for receivers of false and true videos is shown in white bars. The gray bars indicate 95% confidence intervals.


When senders lie, they persuade receivers that they are telling the truth in 52% of the cases, which is not statistically different from the senders' prediction of 51% (t -stat = 0.25, p -value = 0.81, Table A.3). These results provide evidence that senders overestimate the persuasiveness of true videos, in line with Hypothesis 2. On average, their believed persuasiveness for false videos coincides with their actual persuasiveness.

Average beliefs about persuasiveness, however, mask significant heterogeneity. To further explore the relationship between believed and actual persuasiveness we examine their correlation at the video level. A positive correlation indicates that senders who are more persuasive also correctly anticipate being more persuasive. A correlation of zero (or close to 0) indicates that senders cannot correctly differentiate when lies will be more persuasive for receivers.

Figure 4 shows the correlation between senders' believed persuasiveness (x -axis) and how frequently receivers were actually persuaded (y -axis). Figure 4(a) presents the results for true videos, while Figure 4(b) presents the results for false videos. For true videos, senders' beliefs about persuasiveness are positively correlated with actual persuasion. Every 10-percentage-point increase in the senders' belief is associated with a 2.4 percentage point

increase in receivers' likelihood of believing the video is true (t -stat= 2.67, p -value< 0.01, see Table A.2). This positive correlation shows that senders can differentiate between true videos that are highly persuasive and highly unpersuasive. But since most observations lie below the 45-degree line, senders overestimate persuasion when they tell the truth.

Figure 4: Senders' beliefs regarding persuasiveness and actual persuasiveness



Notes: The correlation between senders' beliefs about persuasiveness and receivers' actual persuasion (belief that the video is true) when the video is true (panel (a)) and when the video is false (panel (b)).

In contrast, for false videos, senders' predictions are uncorrelated with receivers' beliefs. As shown in Figure 4(b), senders who believe that their lie is more persuasive are not more likely to actually persuade receivers (t -stat= -0.35, p -value= 0.68, see Table A.2). For lies, senders exhibit a limited ability to predict persuasiveness.

At the individual level, a sender can believe that when compared to false videos, true videos are more persuasive, equally persuasive, or less persuasive. The data show that 62% of senders believe that their truthful video will be more persuasive. The believed advantage of true videos ranges from 10 percentage points to 90 percentage points and is 34 percentage points on average. Additionally, 16% of senders believe that they will be equally persuasive as lies. The remaining 22% believe that their truthful video will be less persuasive, with a difference of 24.5 percentage points (the distribution of relative persuasiveness is

shown in Online Appendix A). Hence, most senders believe that truthful videos will be more persuasive.

The results of Experiment 2 show that senders are miscalibrated in their beliefs regarding persuasiveness, believing that truthful videos are significantly more persuasive than they are. This misperception is shown by a majority of senders.

Can senders’ misperceptions be explained by existing mechanisms of belief biases? One potential mechanism could be “truth bias,” by which senders anticipate receivers to believe they are truthful most of the time (Levine 2014). But this mechanism does not explain the asymmetry between senders’ beliefs when they tell the truth and when they lie. Similarly, an alternative mechanism could be “projective thinking” (Madarász 2012; Madarász 2023; and Madarász et al. 2023): Senders may project their information onto receivers’ beliefs. Such an explanation is consistent with senders’ beliefs when they tell the truth, but inconsistent with their beliefs when they lie, because they do not anticipate receivers being systematically less likely to believe their lies than they actually are. Hence, existing mechanisms do not provide a coherent explanation, suggesting that the misperception we document is tightly linked to truth telling: When senders tell the truth, they overestimate how often they will be believed, but not when they lie.

This asymmetry between instances of lying and truth-telling is not shown by receivers. When receivers are asked about the beliefs of other receivers, at the end of their experiment, they believe that other receivers will correctly detect true videos 59% of the time, while their accuracy is 56% (t -test, p -value= 0.09). Similarly, they believe that other receivers will correctly detect false videos 53% of the time, while their accuracy is 49% (t -test, p -value= 0.01). Hence, receivers overestimate the ability of other receivers to detect the veracity of videos, to a similar extent if they are true and if they are false.


In what follows we investigate the importance of believed persuasiveness in determining the *sender’s* choice to be truthful or lie.

3.3 Choice Experiment

3.3.1 Senders' Choice to Tell the Truth

When given the choice whether to lie or tell the truth, if the reward is the same, senders tell the truth, independent of the need to persuade a receiver. Figure 5 shows that in both 10-10 treatments, 10-10 Receiver and 10-10 Fixed, 94% and 90% of senders tell the truth, respectively (t -stat= -0.61 , p -value= 0.54 , see Table A.4).

Figure 5: Average rates of truth-telling by treatment



Notes: Senders' choice to tell the truth in each treatment is shown in solid bars. The gray bars indicate 95% confidence intervals. $N = 276$.

When there is an incentive to lie, in the 20-10 treatments, senders strongly respond to the need to persuade a receiver. If the sender needs to persuade the receiver, in 20-10 Receiver, 67% of senders still tell the truth despite the incentive to lie. Hence, though senders react to the incentive to lie (t -stat= -4.32 , p -value <0.01), a large majority of senders still choose to tell the truth.

This high frequency of truth-telling despite incentives could be explained by lying aversion—despite the expected payoff gain, senders dislike lying.

But, when the dependence on the receiver is removed, a minority of senders tell the

truth. In 20-10 Fixed, only 18% of senders choose to tell the truth. This result reveals a much higher elasticity to incentives, since the drop in truth-telling for the Fixed treatment is 45 percentage points larger than the drop for the Receiver treatment (t -stat= -5.10 , p -value < 0.01). In line with Hypotheses 3 and 4, we find that incentives to lie lead to a decrease in truth-telling and that the decrease is significantly stronger when senders no longer need to persuade the receiver.

Importantly, we observe that, when the need to persuade receivers is removed, the incentive to lie leads to a large decrease in truth-telling and a large majority of senders lying (in the 20-10 Fixed treatment). This drop suggests that pure lying aversion is not a main driver of senders' choices to tell the truth in this setting. It points to an additional mechanism for why individuals may choose to tell the truth: they believe it is more persuasive.

Since previous research on gender differences in truth-telling (e.g., Dreber and Johannesson 2008; Abeler et al. 2019) has found that female senders are more truthful, we explore whether there are gender differences in truth-telling across treatments. As shown in Table A.4 in Online Appendix A, female senders are 8.8 percentage points more likely to tell the truth (t -stat= 1.87 , p -value= 0.06). Females also react differently to incentives and the dependence on the receiver. In the Receiver treatments, male senders are 50 p.p. less likely to tell the truth, while female senders are 17 p.p. less likely, revealing that when female senders need to persuade a receiver, they exhibit a significantly weaker decrease in truth-telling than male senders (χ^2 test, p -value= 0.01). By contrast, in the Fixed treatments, female senders exhibit a decrease in lying (55 p.p. compared to 25 p.p.) that is stronger than men, and marginally significant at a statistical level (χ^2 test, p -value= 0.09).

3.3.2 Sender's Beliefs about Others' Persuasiveness

At the end of Experiment 3, we also elicited senders' beliefs about persuasiveness. Specifically, we elicited their beliefs about the persuasiveness of other senders (those in Experiment 2). We explore their beliefs about persuasiveness in each treatment, with the caveat that their beliefs were elicited after their choices (our main outcome variable) and hence should be interpreted with care. In all treatments, senders directionally believe that true videos are more likely to persuade receivers (see Figure A.3).

The difference in senders’ beliefs is 15 p.p. on average, and it is economically and statistically significant in all treatments except 20-10 Fixed treatment (see Table A.5). In that treatment, where a majority of senders lie, true videos are believed to be significantly less persuasive than in other treatments (8.5 p.p., t -stat= 2.11, p -value= 0.04).

4 Interpreting Decisions to Lie

We have seen that senders overestimate the persuasiveness of truths, and that the believed persuasiveness of truths appears to be an important driver of the decision to lie: When senders do not depend on receivers’ beliefs to make their decision, and when they have sufficient incentives to lie (20-10 Fixed), very few senders tell the truth (18%).

The data from both experiments can be combined to allow us to learn about what may explain the decisions that senders make. We ask whether senders’ beliefs about the persuasiveness of true and false videos, combined with lying costs and risk aversion or risk neutrality, can provide a coherent potential explanation for the data patterns observed. We use the sender’s beliefs about persuasiveness from Experiment 2 to compute their believed expected payoffs from choosing to record a true or a false video. To consider their expected utility, we consider two possible degrees of risk aversion: when the sender is risk neutral, and when the sender is risk averse, with CRRA of 0.7, over the experimental payments (Rabin 2000). These are illustrative calculations to descriptively examine what kinds of preferences of the marginal sender could explain the data. For simplicity, we employ a CRRA of 0.7, which fits behaviour in the experiment. The details are provided in Online Appendix A.

In the 10-10 Receiver treatment of Experiment 3 very few senders lie. This pattern is consistent with the beliefs observed in senders in Experiment 2 and a marginal sender who is risk averse and has lying costs of 1.50 (or 3 if risk neutral). In the 20-10 Receiver treatment, 33% of senders in Experiment 3 chose to lie. The 33rd sender (out of 100), with a CRRA of 0.7, has an expected “cost” (disutility) from telling the truth of 1.69, not considering lying costs. Hence, lying costs of 1.69 would make this sender indifferent between telling the truth and lying in the 20-10 condition. The estimated strength of lying costs (1.50 and 1.69) suggests that a similar strength of moral costs from lying can rationalize behaviour in both

Receiver treatments.

We next consider whether these values are in line with behaviour in the Fixed treatments. When senders are not compensated for lying and their payoffs are independent of receivers (10-10 Fixed), we would expect no senders to lie and we see a low rate of lying (10%). When senders receive an additional \$10 for lying (20-10 Fixed), although senders believe the truth is more persuasive, the very high payoff for lying combined with mild lying costs would entice senders to lie, consistent with a high rate of lying (82%). Overall, behaviour in Experiment 3 is consistent with moderate lying costs and risk aversion combined with beliefs as in Experiment 2.

5 Conclusion

Persuasion is central to many economic interactions. In negotiations, work environments, and social interactions, people have the choice between being truthful or lying to achieve their goals. A key factor is how persuasive they think they will be when lying and telling the truth, i.e., their beliefs about how the receiver will react to true and false messages.

In this paper we study these beliefs in three experiments and consistently find that individuals overestimate the relative persuasiveness of truths. In a large-scale survey experiment, observers of lies and truths believe that senders will believe them significantly more when they tell the truth. In the laboratory, senders who are incentivized to persuade also believe their truthful messages will be more persuasive than their lies. But these beliefs are wrong, as truthful and false messages are similarly persuasive to receivers.

These beliefs are important because the need to convince the receivers affects the decision to lie. By varying the incentives and need to persuade senders, we show that a majority of senders choose to be truthful when their payoffs depend on convincing the receiver, despite high monetary incentives to lie. By contrast, when the need to persuade is removed, a majority of senders choose to lie, attracted by the high payoff it provides.

If senders were sophisticated and understood that receivers believe lies approximately the same as truth, then senders who do not have a high intrinsic cost of lying would maximize their profits by lying more often. The mistaken beliefs can sustain a high level of truth

telling despite the ease with which receivers can be deceived. As such, our results highlight an important aspect of why people do not lie more.

As always, it is risky to generalize from one set of experiments. While we aimed to design experiments that objectively test our hypotheses, there might be aspects that affect the results. For example, in Experiment 2 and Experiment 3, it might be hard to come up with a story on the spot, and so telling the truth is cognitively easier for the sender than telling a lie, leading them to believe that telling the truth is more convincing.

Lying is a complex and intriguing behaviour that has attracted attention from various disciplines. Research in psychology estimates that individuals lie, on average, between one and two times per day, revealing a certain level of lying in everyday interactions (DePaulo et al. 1996; Serota et al. 2010). This frequency prompts inquiries from a moral standpoint into strategies that can reduce the prevalence of lying. In contrast, the classic economic model, which assumes no cost associated with lying and posits that individuals will lie if it maximizes their material payoffs (e.g., Akerlof 1978; Crawford and Sobel 1982), raises the opposite question: why don't people lie more frequently? Despite abundant opportunities, individuals consistently choose not to engage in lying behaviour, challenging the assumption that people always act to maximize their material gains.

One explanation for the limited prevalence of lying is the concern for one's reputation. The classic economic model, incorporating standard assumptions, can explain why individuals opt for honesty in repeated interactions, driven by a desire to maintain a trustworthy image (see the discussion in Mailath and Samuelson 2006).

A second explanation is based on the concept of "lying aversion," proposing that individuals prefer achieving outcomes without resorting to lying. This perspective aligns with moral and ethical considerations, which introduce psychological costs to lying, deviating from standard economic assumptions. Empirical evidence supports the lying aversion explanation, demonstrating that participants in experiments often choose not to lie, even when telling the truth is costly and the likelihood of detection is low (Gneezy 2005; Fischbacher and Föllmi-Heusi 2013; see Abeler et al. 2019, for a review).

The findings from this paper provide valuable insights into a third possible explanation: senders' (incorrect) beliefs that true statements are more likely to be believed. By examining

the role of individuals’ beliefs in others’ trust, we offer a perspective on why people do not engage in lying more frequently that is missing in the literature. Understanding these factors is essential for comprehending lying behaviour and has important implications for fostering trust, reducing lying, and promoting ethical decision-making in various domains.

There are many reasons why participants might think that the truth is more persuasive. One reason is that they might misestimate the signals sent by body language, the message content, presentation confidence, etc. In a recent paper (Serra-Garcia and Gneezy 2025), we find similar results regarding the limited ability to detect true and false statements in the context of TV show (Golden Balls). However, in Experiment 1, the scenarios based on a simple conversation with the same content, for truth and lies, revealed the same qualitative finding: the truth is believed to be more persuasive.

Our findings are related to “projective thinking” (Madarász 2012; Madarász et al. 2023) – senders believing that their true messages may seem obviously true to receivers –, but they do not explain the difference between true and false messages. When it comes to lies, senders’ beliefs do not exhibit the same projection. Their beliefs are consistent with doubt about what makes lies persuasive. To better understand the perceived asymmetry between truth and lies, future research could use incentivized “controlled” tasks, such as guessing the veracity of a report by a sender in a cheap talk game and testing the accuracy of beliefs in such designs.

References

- Abeler, J., Becker, A., and Falk, A. (2014). Representative evidence on lying costs. *Journal of Public Economics*, 113:96–104.
- Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4):1115–1153.
- Agranov, M. and Detkova, P. (2024). Beliefs of others: an experiment. Technical report, National Bureau of Economic Research.
- Akerlof, G. A. (1978). The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in economics*, pages 235–251. Elsevier.
- Belot, M. and Van de Ven, J. (2017). How private is private information? the ability to spot deception in an economic game. *Experimental Economics*, 20:19–43.
- Biais, B. and Weber, M. (2009). Hindsight bias, risk perception, and investment performance. *Management Science*, 55(6):1018–1029.
- Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *The American Economic Review*, 88(5):1323–1340.
- Bond, C. F. and DePaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review*, 10(3):214–234.
- Camerer, C., Loewenstein, G., and Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy*, 97(5):1232–1254.
- Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6):1579–1601.
- Cialdini, R. B. (2007). *Influence: The psychology of persuasion*. Collins New York.
- Coffman, L. and Niehaus, P. (2020). Pathways of persuasion. *Games and Economic Behavior*, 124:239–253.

- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50:1431–1451.
- DellaVigna, S. and Gentzkow, M. (2010). Persuasion: empirical evidence. *Annual Review of Economics*, 2(1):643–669.
- DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., and Epstein, J. A. (1996). Lying in everyday life. *Journal of personality and social psychology*, 70(5):979.
- Dreber, A. and Johannesson, M. (2008). Gender differences in deception. *Economics Letters*, 99(1):197–199.
- Dwenger, N. and Lohse, T. (2019). Do individuals successfully cover up their lies? evidence from a compliance experiment. *Journal of Economic Psychology*, 71:74–87.
- Eckel, C. C. and Petrie, R. (2011). Face value. *American Economic Review*, 101(4):1497–1513.
- Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, 11(3):525–547.
- Gilovich, T., Savitsky, K., and Medvec, V. H. (1998). The illusion of transparency: biased assessments of others’ ability to read one’s emotional states. *Journal of personality and social psychology*, 75(2):332.
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1):384–394.
- Gneezy, U., Rockenbach, B., and Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior & Organization*, 93:293–300.
- Grabova, I., Nielsen, H., and Weizsäcker, G. (2023). Attempting to detect a lie: Do we think it through? Technical report, Discussion Paper.
- Levine, T. R. (2014). Truth-default theory (tdt) a theory of human deception and deception detection. *Journal of Language and Social Psychology*, 33(4):378–392.

- Madarász, K. (2012). Information projection: Model and applications. *The Review of Economic Studies*, 79(3):961–985.
- Madarász, K. (2023). Limited perspective taking in strategic communication. Technical report, mimeo, LSE.
- Madarász, K., Danz, D., and Wang, S. (2023). Projective thinking: Model, evidence, and applications. *Evidence, and Applications (May 31, 2023)*.
- Mailath, G. J. and Samuelson, L. (2006). *Repeated games and reputations: long-run relationships*. Oxford university press.
- Peer, E., Brandimarte, L., Samat, S., and Acquisti, A. (2017). Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163.
- Peeters, R., Vorsatz, M., and Walzl, M. (2015). Beliefs and truth-telling: A laboratory experiment. *Journal of Economic Behavior & Organization*, 113:1–12.
- Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292.
- Serota, K. B., Levine, T. R., and Boster, F. J. (2010). The prevalence of lying in america: Three studies of self-reported lies. *Human Communication Research*, 36(1):2–25.
- Serra-Garcia, M. and Gneezy, U. (2021). Mistakes, overconfidence, and the effect of sharing on detecting lies. *American Economic Review*, 111(10):3160–3183.
- Serra-Garcia, M. and Gneezy, U. (2025). Improving human deception detection using algorithmic feedback. *Management Science*.
- Thaler, M. (2021). The supply of motivated beliefs. *arXiv preprint arXiv:2111.06062*, 385.