

# Motivated Skepticism\*

Jeanne Hagenbach<sup>†</sup>      Charlotte Saucet<sup>‡</sup>

November 15, 2022

## Abstract

We experimentally study how individuals read strategically-transmitted information when they have preferences over what they will learn. Subjects play disclosure games in which Receivers should interpret messages skeptically. We vary whether the state that Senders communicate about is ego-relevant or neutral for Receivers, and whether skeptical beliefs are aligned or not with what Receivers prefer believing. Compared to neutral settings, skepticism is significantly lower when it is self-threatening, and not enhanced when it is self-serving. These results shed light on a new channel that individuals can use to protect their beliefs in communication situation: they exercise skepticism in a motivated way, that is, in a way that depends on the desirability of the conclusions that skeptical inferences lead to.

**Keywords:** Disclosure games, hard information, unraveling result, skepticism, motivated beliefs.

**JEL Code:** C72, C91, D82, D91.

---

\*We thank Victor Augias, Kai Barron, Roland Bénabou, Alessandra Casella, Frédéric Koessler, Johannes Leutgeb, Yves Le Yaouanq, Eduardo Perez-Richet, Rémi Suchon, Marie Claire Villeval, Sevgi Yuksel and seminar participants at Sciences Po, PSE, CREST, UCL, WZB Berlin, Universities of Constance, Cologne, Maastricht, Kyoto, Osaka, Sinica, Hong Kong (BU), Taiwan (NTU), the ASFEE conference and the BBU workshop in Amsterdam for helpful comments. We are grateful to Quentin Thevenet for programming the experiment and to Nina Bonge for helping at the WZB-TU Lab. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement n°850996 – MOREV).

<sup>†</sup>CNRS, Sciences Po, CEPR - jeanne.hagenbach@sciencespo.fr

<sup>‡</sup>University of Paris 1 - charlotte.saucet@univ-paris1.fr

# 1 Introduction

In many economic situations, communication is considered as an effective way to reduce information asymmetry. The less informed parties are assumed to be willing to learn the truth and, knowing the objectives of the communicating parties, to be able to make inferences from disclosed and undisclosed information. At the extreme, the *unraveling result* (Milgrom, 1981, Grossman, 1981) establishes that information can be fully learnt by uninformed parties provided that they read incomplete disclosure with skepticism. And indeed, it is often natural to interpret the absence of precise evidence as unfavorable for the communicating party. This paper shows that individuals' ability to exercise skepticism, and therefore the relevance of the unraveling result, importantly depends on whether or not individuals *want* to learn the truth in the first place. Using an online-lab experiment, we investigate how subjects interpret hard information when they have preferences over what they want to be true.

Situations of strategic communication often share the feature that agents are not indifferent about what they learn. Think for example of firms revealing hard information to consumers about products attributes. Consumers wish to know these attributes to best-adapt purchasing decisions, but they may also benefit *per se* from believing that products have particular attributes, such as being environmental-friendly or ethically produced. In advising settings, advisors often communicate with advisees about their abilities or chances of success. Advisees have an interest in learning the truth but they may also be directly affected by the beliefs they hold about their abilities. In these situations, the agents who read information trade-off the need to form accurate beliefs and the wish to hold comforting or pleasant beliefs. Our experiment studies agents who form potentially motivated beliefs using strategically-transmitted information. This way, it brings together the literature on *disclosure games* and the literature on *motivated beliefs*.

The theoretical literature on voluntary disclosure games makes especially sharp predictions about the reading of information in equilibrium. In these games, the information transmitted by the privately-informed Sender to the Receiver is hard in the following sense: a message from the Sender consists of a subset of types that must include the true type. The message is considered precise when it is a singleton set – the type is fully disclosed –, and vague otherwise. In the classical version of these games proposed by Milgrom (1981), the Sender's payoff simply increases with the decision finally taken by the Receiver. The

Receiver’s objective is to take an action that matches the type, or, said differently, to identify the type as accurately as possible. In equilibrium, information is fully revealed by the Sender because different types never pool on sending the same vague message: the highest type has an interest in separating from the lower types and inducing a higher action, which he can always do by fully disclosing; at the next step, it is the second-highest type who should fully disclose, etc. Because of this unraveling mechanism, the Receiver’s equilibrium beliefs after any vague message are *skeptical*, that is, assign probability one to the lowest type disclosed. Skepticism captures the intuitive idea that, when facing vague statements such as “the student is in the top ten” or “at least half of the ingredients are organic”, the rational reading is that the student actually ranked ten and that no more than half of the ingredients are organic.

Our experiment is designed to study the extent to which Receivers form skeptical beliefs as a function of whether or not these beliefs are desirable for them. In the lab, Sender-subjects transmit verifiable information about a state to uninformed Receiver-subjects. Receivers need to guess the state right to maximize their payoffs but, in some treatments, additionally have intrinsic preferences over what they believe about the state. We cross two exogenous variations to change whether the skeptical beliefs are aligned or not with the preferred beliefs of the Receivers, all else equal.

The first variation concerns the type that the Sender communicates about. In *Loaded* treatments, the type corresponds to a noisy measure of the Receiver’s relative performance in a previous IQ test. In *Neutral* treatments, the type is a rank with no particular meaning. The idea behind this variation is that Receivers have preferences over what they believe in *Loaded* treatments but not in *Neutral* treatments.<sup>1</sup> We study how these preferences affect the reading of strategically-disclosed information, while keeping the monetary value of holding accurate beliefs fixed in all treatments.

The second variation concerns the objective of the Sender. In *High* treatments, the Sender’s payoff strictly increases with the Receiver’s guess of the type. As in standard theory, the Sender therefore wants to induce a high guess, so he precisely discloses high types which, in our design, correspond to low ranks. In *Low* treatments, the Sender’s payoff strictly decreases with the Receiver’s guess. Thus, it is now the Sender of relatively high

---

<sup>1</sup>The use of an ego-relevant state is a commonly-used artefact to create preferences over beliefs in the lab (see Schwardmann and Van der Weele (2019) and Zimmermann (2020) for instance).

types who sends vague messages or, said differently, low ranks that are concealed. With the *High/Low* variation, we change whether being skeptical consists in believing the lowest or the highest rank disclosed. To our knowledge, we are the first to consider this variation in an experimental disclosure game. And indeed, it seems redundant in *Neutral* treatments as it makes no difference in theory that the skeptical beliefs correspond to a high or low rank. In *Loaded* treatments however, it affects whether or not the skeptical belief is aligned with the Receiver’s preferred belief, which is the central element of our design. Specifically, in the *High\_Loaded* treatment, the skeptical belief assigns probability one to the highest rank disclosed, a belief we consider *self-serving*. In the *Low\_Loaded* treatment, the skeptical belief assigns probability one to the lowest rank disclosed, a belief we consider *self-threatening*.

We implement the four treatments between subjects. Each subject is given the role of Sender or Receiver, and play ten Sender-Receiver games with random rematch and no feedback. Our data contains 2000 games for which we record the type seen by the Sender, the message sent and the Receiver’s guess. In each game, we measure Receiver’s skepticism by evaluating how close his guess is to the guess he would have made if he held skeptical beliefs given the message seen. The latter guess is called the *skeptical guess*. We also measure skepticism by taking the frequency of skeptical guesses. We test three main pre-registered hypotheses.<sup>2</sup> First, we hypothesize that Receivers’ skepticism is unaffected by the *High/Low* variation in *Neutral* treatments. Second, we hypothesize that Receivers’ skepticism will be at least as high in *High\_Loaded* as in *High\_Neutral*, that is, when skeptical beliefs are self-serving. Third, we hypothesize that Receivers’ skepticism is strictly lower in *Low\_Loaded* than in *Low\_Neutral*, that is, when skeptical beliefs are self-threatening. In short, we expect individuals’ to interpret vague information skeptically when it is good news to them, and less so when it is bad news to them.

We begin the data analysis by checking that Senders and Receivers understood the basics of the game. Senders’ communication strategies have a clear structure that takes into account their reversed objectives: in *High* treatments, Senders of type  $t$  most often disclose that their type is *at least*  $t$  (that is, use the “top ten” kind of messages); in *Low* treatments, Senders of type  $t$  most often disclose that their type is *at most*  $t$ . These strategies are fully-revealing.<sup>3</sup> On their side, Receivers take the evidence into account in the sense that

---

<sup>2</sup>The reference for pre-registration is AEARCTR-0007541.

<sup>3</sup>Additionally, these strategies correspond to the optimal fully-revealing strategy of a Sender facing a Receiver who, with some probability, is not making skeptical inferences but may guess any of the disclosed types.

their guesses are almost always in between the lowest and highest types disclosed. When comparing Receivers' skepticism in the two *Neutral* treatments that serve as benchmarks, the difference is marginal and goes in the direction of a slightly lower level of skepticism in *Low\_Neutral* than in *High\_Neutral*. We attribute this small difference to the Sender's payoff function being slightly more complex in the *Low* than in the *High* condition.

Our main finding is that skepticism is significantly lower when it is self-threatening (as in the *Low\_Loaded* treatment) than when it is not (as in the *Low\_Neutral* treatment). Said differently, individuals read information less skeptically when it implies believing they ranked low in the IQ test, than in neutral environments. In contrast, we find no significant difference in the level of Receivers' skepticism when skepticism is self-serving (as in the *High\_Loaded* treatment) and when it is not (as in the *High\_Neutral* treatment). An explanation for the fact that skepticism is not enhanced when self-serving is that subjects already reach some limit in making skeptical inferences in the *High\_Neutral* treatment.

These two results are important. They first demonstrate that the exercise of skepticism does not depend on the object individuals reason about: subjects are able to make skeptical inferences about their relative IQ in the *High\_Loaded* treatment. Second, they show that individuals exercise skepticism in a way that depends on the conclusions that skeptical inferences lead to. To some extent, this result is in line with previous findings establishing that individuals process good news and bad news asymmetrically to protect their beliefs (see for instance Eil and Rao, 2011 or Möbius et al., 2022). One important difference however is that, in our setting, individuals need to reason about the Sender's strategy to uncover whether a message is good or bad news to them. The uncertainty about the informed party's strategy, which is intrinsic to strategic communication settings, leaves individuals room to avoid reaching disturbing conclusions. And indeed, we report that when a message happens to be bad news, subjects reach skeptical guesses less often. Given that economic agents base most of their decisions on strategically-transmitted information, the motivated reading of such information can have important consequences. In particular, motivated skepticism provides a new, psychological reason for the failure of unraveling which is often observed in the field and contract standard theory.<sup>4</sup>

---

<sup>4</sup>Understanding when unraveling fails is important to decide whether or not to mandate disclosure. Dranove and Jin (2010) participate in this debate by providing a survey of the empirical and theoretical literature on quality disclosure. They report that unraveling is incomplete in many markets and already provide explanations of why this is the case: disclosure may not be costless, the Sender may not be fully informed, etc.

We further develop our main results by evaluating, for every message, the extent to which the skeptical conclusions are desirable or not in *Loaded* treatments. We find that skeptical guesses are significantly less frequent when these guesses are strongly self-threatening than when they are mildly so. On the contrary, skeptical guesses are significantly more frequent when they are strongly self-serving than when they are mildly so. We additionally show that our main results, established at the aggregate level, are confirmed at the individual level. On average, every subject makes a significantly lower fraction of skeptical guesses in the *Low\_Loaded* treatment than in any other treatment. When considering the steps of reasoning that could lead to the guesses we observe, we find that Receivers make fewer steps towards the skeptical guess when this guess is self-threatening than when it is not.

In *Loaded* treatments, we construct ranks using Receivers' performance in an IQ test they complete at the beginning of the experiment. This has two consequences. First, it offers a measure of all Receivers' IQ. We show that Receivers' skepticism is positively correlated to this measure, that is, that subjects who solve more Raven matrices are also better at making skeptical inferences, confirming a finding of Schipper and Li (2020). Motivated skepticism comes in addition to the fact that the exercise of skepticism is, in the first place, limited by agents' cognitive ability to make sophisticated inferences. Second, in *Loaded* treatments, Receivers guess their ranks while having prior beliefs about their relative performance in the IQ test. We elicit these beliefs and show that Receivers' lack of skepticism is not driven by confident Receivers who may be insufficiently skeptical simply because they stick to their positive prior beliefs.

**Related literature.** Our experiment connects the literature on *disclosure games* and the literature on *motivated beliefs*.

A central result in the literature on disclosure games is the *unraveling result* of Milgrom (1981) and Grossman (1981), which establishes that information is fully disclosed by the informed party in equilibrium.<sup>5</sup> This result crucially relies on Receivers reading information in a specific way, unfavorable to the Sender. While the theory is clear, the empirical literature reports that voluntary disclosure is not always complete (see Mathios, 2000, Luca and Smith, 2015 or Bederson et al., 2018 for concrete examples, and Dranove and Jin, 2010 for a survey). Experiments have offered some elements to understand the partial disconnection between

---

<sup>5</sup>See Okuno-Fujiwara et al. (1990), Seidmann and Winter (1997), Giovannoni and Seidmann (2007) and Hagenbach et al. (2014) for related theoretical works.

unraveling in theory and in practice. Jin et al. (2021) and Schwardmann et al. (2021) closely replicate the standard game of Milgrom (1981) in the lab. These works respectively show that subjects’ ability to make skeptical inferences depends on their experience in playing the disclosure game, and on the language used by the Sender.<sup>6</sup> Some experiments additionally incorporate realistic perturbations to the standard game, such as costs to disclosure or a probability that the Sender is not informed, that induce failures of unraveling.<sup>7</sup> Benndorf et al. (2015) observe a low amount of disclosure when Senders are framed into playing the role of workers disclosing their productivity. Their experiment suggests that psychological aspects can affect communication in important ways, a point also made in Loewenstein et al. (2014). Our experiment is the first to consider that the reading of hard information may depend on the preference Receivers have over beliefs. We show that unraveling may fail because individuals do not want to uncover the truth.

The idea that, in some contexts, individuals’ well-being directly depends on their beliefs is at the center of the recently-growing literature on *motivated beliefs* (surveyed in Bénabou, 2015).<sup>8</sup> In the last decade, the experimental literature on that topic has identified various channels that individuals use to reach favored conclusions, sometimes despite contradictory evidence. These channels include asymmetric belief updating, selective recall, or motivated information selection.<sup>9</sup> In the works on asymmetric updating, subjects form beliefs about states they intrinsically care about, such as their relative intelligence or beauty. They do so based on noisy signals provided to them by the experimenter, and can directly assess whether these signals are good or bad news to them. In our work, subjects update their beliefs using information which is *strategically disclosed* by another player. Updating beliefs

---

<sup>6</sup>Schipper and Li (2020) provide additional evidence of an important fraction of unraveling outcomes and link Receivers’ levels of reasoning to their IQ. Hagenbach and Perez-Richet (2018) show that Receivers can exercise skepticism also when the Senders’ payoffs are not monotonic as in Milgrom (1981).

<sup>7</sup>In that vein, King and Wallin (1991) and Dickhaut et al. (2003) find that unraveling fails more when the Sender is less likely to be informed.

<sup>8</sup>Models have been developed in which beliefs directly enter the agent’s utility function through self-image concerns (as in Bénabou and Tirole, 2011), anticipatory emotions or anxiety (as in Kőszegi, 2006, Caplin and Eliaz, 2003 or Schwardmann, 2019) or motivational concerns (Bénabou and Tirole, 2002).

<sup>9</sup>Asymmetric information processing is documented in various experiments such as Eil and Rao (2011), Sharot et al. (2011), Charness and Dave (2017), Drobner and Goerg (2021) or Möbius et al. (2022). Evidence of selective memory, both in the lab and in the field, can be found in Huffman et al. (2019), Zimmermann (2020), Saucet and Villeval (2019), Galeotti et al. (2020), Chew et al. (2020), Carlson et al. (2020), Gödker et al. (2020) and Müller (2021). Motivated information selection is documented in Grossman (2014), Grossman and Van der Weele (2017), Serra-Garcia and Szech (2021), Chen and Heese (2021) and Exley and Kessler (2021). See Golman et al. (2017) for a survey on information avoidance and Gino et al. (2016) for a survey on how subjects reconcile feeling moral while acting egoistically. Thompson and Loewenstein (1992) is a seminal paper on biased search of information. A general discussion of the mechanics of motivated reasoning can be found in Epley and Gilovich (2016), and earlier in Kunda (1990).

therefore requires to reason about the other party’s strategy in the first place. In addition, it is the strategic setting itself (the Sender’s objective, not the message per se) that determines whether a given message is good or bad news to the Receivers. We shed light on a new channel that individuals can use to protect their beliefs: insufficient skepticism.<sup>10</sup>

Only a few papers examine the formation of motivated beliefs in strategic and social settings. In Schwardmann and Van der Weele (2019), Schwardmann et al. (2022), and Solda et al. (2020), individuals convince themselves that a state is true to better persuade others. In Thaler (2022) and Burro and Castagnetti (2022), Senders adapt their cheap talk to what Receivers want to hear about political issues or relative intelligence. Oprea and Yuksel (2021) show that subjects rely more on peers’ beliefs when these beliefs are self-serving. We consider hard information transmission and study motivated deviations from skeptical beliefs.

## 2 Theoretical framework

In this section, we present the well-known theoretical framework that guided our design and give definitions that will help us analyze the data.

**Baseline game.** The baseline game is a version of the classical Sender-Receiver game of Milgrom (1981) and Grossman (1981). The Sender is privately informed of a type  $t$  and sends a costless message  $m$  about it to the Receiver. We assume that  $t$  is initially drawn from a finite set of real numbers  $T$ . Upon receiving  $m$ , the Receiver updates his beliefs about  $t$  and chooses an action  $a \in \mathbb{R}$  which affects both players. The Sender’s payoff  $u_S(a)$  is type-independent and strictly increasing in  $a$ . The Receiver’s payoff is given by  $u_R(a; t) = -(a - t)^2$ , capturing the idea that the Receiver wants his action to match the type.

**Hard information.** When the type is  $t$ , the set of messages available to the Sender is  $M(t)$ . This set contains all the subsets of  $T$  which include  $t$  and are made up of consecutive numbers.<sup>11</sup> As an example, consider  $T = \{1, 2, 3, 4, 5\}$ . The set of messages available to a Sender of type  $t = 1$  is  $M(1) = \{\{1\}, \{1, 2\}, \{1, 2, 3\}, \{1, 2, 3, 4\}, \{1, 2, 3, 4, 5\}\}$ . The set of all

---

<sup>10</sup>In a related work Thaler (2022), shows that subjects assess the veracity of information sources in directions which correspond or not to pre-conceived political views.

<sup>11</sup>This message structure corresponds to the *rich language* considered in Hagenbach and Koessler (2017) or Ali et al. (2021). An alternative language that is commonly considered in theoretical works and experiments is the *simple language*: the Sender of type  $t$  either discloses  $\{t\}$  or  $T$ . The rich language allows for more nuances in disclosure.



messages available to the Sender is  $M = \cup_{t \in T} M(t)$ . With this message structure, a message  $m \in M$  provides *evidence* that the true type is in  $m$ . We talk about a *precise message* when  $m$  is a singleton, and about a *vague message* otherwise. The *size* of message  $m$  is its cardinal.

**Players' strategies and equilibrium concept.** A (pure) strategy of the Sender is a mapping  $\sigma_S(\cdot)$  from  $T$  to  $M$  such that  $\sigma_S(t) \in M(t)$ . The Sender's strategy is *fully revealing* when it is separating:  $\sigma_S(t) \neq \sigma_S(t')$  for every  $t \neq t'$ . A (pure) strategy of the Receiver is a mapping  $\sigma_R(\cdot)$  from  $M$  to the set of actions  $\mathbb{R}$ .  $\beta_m \in \Delta(T)$  is the belief of the Receiver following message  $m$ . We say that a belief  $\beta_m$  is *consistent* with  $m$  if  $\beta_m$  has support in  $m$ .

We solve this game using Perfect Bayesian equilibrium. In equilibrium, (i)  $\sigma_S(\cdot)$  is a best-reply to  $\sigma_R(\cdot)$ , (ii)  $\beta_m$  is derived from Bayes' rule after any message  $m$  sent on the equilibrium path, and (iii) for every  $m$ , the action that maximizes the Receiver's payoff given  $\beta_m$  is  $\sigma_R(m) = E_{\beta_m}(t)$ . Point (ii) implies that Receivers' beliefs will be consistent after any message  $m$  sent on the equilibrium path. We add the requirement, common in disclosure games, that  $\beta_m$  is also consistent after any messages  $m$  sent off path. By extension of consistent beliefs, we define consistent actions:

**Definition 1** An action  $\sigma_R(m)$  is *consistent* with  $m$  if it is optimal for a belief that has support in  $m$ .

In the experimental data, we will examine the consistency of Receivers' actions. It will serve as a check that Receivers have understood that Senders transmit hard evidence.

**Receiver's skepticism.** The notion of skepticism, introduced by Milgrom (1981) and Milgrom and Roberts (1986), plays a central role in disclosure games. In the game described above, the Sender wants the Receiver to take an action which is as high as possible. A Sender of high type should therefore fully disclose his type, which can always be done with a precise message. In contrast, a Sender of low type may have an interest in shrouding information and sending vague messages. It follows that an intuitive way for a Receiver to interpret a vague message is to be skeptical and interpret the message as coming from a Sender of low type. We define skeptical beliefs as follows:

**Definition 2** In the baseline game, the belief  $\beta_m$  is *skeptical* if it assigns probability one to the lowest type in  $m$ .

In our experiment, we will observe the Receivers' actions and not directly their beliefs. We define a *skeptical action* as follows:

**Definition 3** The action  $\sigma_R(m)$  is *skeptical* if it is optimal for the skeptical belief  $\beta_m$ .

In the baseline game, the skeptical action  $\sigma_R(m)$  equals the lowest type in  $m$ .

We now state the unraveling result which establishes that every equilibrium of this game is fully-revealing, and that the Receiver's equilibrium beliefs are always skeptical (Milgrom and Roberts, 1986). The proof is reminded in Appendix A.

**Proposition 1** *Every equilibrium is fully revealing. In equilibrium, the Receiver's beliefs are skeptical after every (on and off-path) message.*

**Alternative game.** In our experiment, subjects will play the baseline disclosure game or an alternative game in which the Sender's payoff is type-independent but strictly decreasing in  $a$ . In that case, Proposition 1 still applies provided that we adapt the definition of skeptical beliefs as follows:

**Definition 4** In the alternative game, the belief  $\beta_m$  is *skeptical* if it assigns probability one to the highest type in  $m$ .

In the alternative game, the skeptical action  $\sigma_R(m)$  equals the highest type in  $m$ . Our experiment is crucially based on the reversal of what it means to be skeptical for a Receiver.

**Selective disclosure.** Proposition 1 establishes that the Sender uses a fully-revealing strategy in every equilibrium. This strategy can be made up of precise or vague messages, all of them being interpreted skeptically by a sophisticated Receiver. Consider now that there is a positive probability that the Receiver is not sophisticated enough to make skeptical inferences and acts as follows: when a type  $t$  is disclosed (that is, contained in the message received), there is always a probability that the Receiver believes this type for sure and takes action  $a = t$ . In this case, a specific fully-revealing strategy is optimal for the Sender: every type  $t$  discloses that the type is *at least*  $t$ , that is, sends the message  $\{t, \dots, t_{sup}\}$  with  $t_{sup}$  the highest type in  $T$ .<sup>12</sup> On the contrary, if the game is the alternative one in which the Sender's payoff is strictly decreasing in the Receiver's action, every type  $t$  discloses that the

---

<sup>12</sup>Milgrom and Roberts (1986) already point to this result which is proved in Hagenbach and Koessler (2017).

type is *at most*  $t$ , that is, sends the message  $\{t_{inf}, \dots, t\}$  with  $t_{inf}$  the lowest type in  $T$ . We refer to these specific communication strategies as *selective disclosure*.

**Receiver’s preference over beliefs.** In our experiment, we will consider neutral settings which closely fit the baseline and alternative games, but also settings in which the Receiver intrinsically cares about the beliefs he holds about  $t$ . We can incorporate this fact into our model by considering that the Receiver’s utility is given by the sum of his material payoff  $u_R(a, t)$  and a psychological payoff that depends directly on his expectation about  $t$ . For a given belief  $\beta_m$ , the Receiver’s utility equals  $-(a - t)^2 + \alpha E_{\beta_m}(t)$  where  $\alpha \geq 0$  parametrizes self-image concerns.<sup>13</sup> With belief  $\beta_m$ , the Receiver’s optimal action is still given by  $E_{\beta_m}(t)$ .

When the Receiver’s utility is  $u_R(a, t)$ , the Receiver’s equilibrium belief is skeptical after every message (Proposition 1). When the Receiver’s utility intrinsically depends on beliefs, he may trade-off the material benefit of holding skeptical beliefs after  $m$  – assigning probability one to  $t_{skep}(m)$  – and the psychological benefit of having a good expectation of  $t$ . Let us assume that, after message  $m$ , the Receiver chooses the belief  $\beta_m$  that maximizes  $-(E_{\beta_m}(t) - t_{skep}(m))^2 + \alpha E_{\beta_m}(t)$  under the constraint that  $\beta_m$  is consistent with  $m$ . The belief  $\beta_m$  is then such that  $E_{\beta_m}(t) = \min\{t_{sup}(m), t_{skep}(m) + \frac{\alpha}{2}\}$ , where  $t_{sup}(m)$  is the highest type in  $m$ . The Receiver inflates his expectation of the type. He does it more when self-image is more important, and up to the highest consistent expectation he can have after  $m$ ,  $t_{sup}(m)$ .

### 3 Experimental design

We present the overall structure of our experiment before describing the two experimental variations that we consider.

#### 3.1 Overall structure

The experiment is made up of two parts. Subjects’ final payoff is the sum of a show-up fee and of the money they made in each part.

---

<sup>13</sup>Similar forms of utility can be found in Bénabou et al. (2019), based on Bénabou and Tirole (2006) and Bénabou and Tirole (2011).

**Part 1: IQ test.** Subjects begin by completing a 15-minutes test made up of 15 Raven matrices (Raven, 1936). They earn 0.50 euros per correctly solved matrix. A subject's *performance* is an integer between 0 and 15 that corresponds to the number of correctly solved matrices. When the IQ test is over, we elicit subjects' beliefs about their performance in that test relative to the performance of a benchmark group made up of 99 subjects who did the same IQ test.<sup>14</sup>

**Part 2: Sender-Receiver games.** In the second part of the experiment, subjects play 10 times the same Sender-Receiver game. Before the 10 games start, subjects learn whether they will play the role of Sender or the role of Receiver. Subjects keep their role for the 10 games but are randomly matched in Sender-Receiver pairs at the beginning of every game, which is common knowledge. Each of the 10 games has four steps:

*Step 1* The computer generates a type  $t$  in  $\{1, 2, 3, 4, 5\}$ . In section 3.2 below, we detail how the type is generated, a key element of our experimental manipulation.

*Step 2* The Sender is privately informed about the generated type  $t$ .

*Step 3* The Sender decides which message  $m$  about  $t$  to sent to the Receiver. For any given type  $t$ , the set of messages available to the Sender is restricted to the sets of consecutive types which contain  $t$ .

*Step 4* The Receiver observes the message  $m$  sent in *Step 3* and makes a guess  $a \in [1, 5]$  about the type  $t$ . We allow for guesses with one digit.<sup>15</sup>

When a Sender-Receiver game is over, subjects move to the next game without getting any feedback about the type  $t$  that the Sender had effectively seen or about their realized payoff.

**Payoffs.** Players' payoff functions are common knowledge. The Sender's payoff function only depends on the Receiver's guess  $a$  and not on the type generated in *Step 1* of the game. The Sender's payoff function is part of our experimental manipulation and is detailed in section 3.2 below. The Receiver's payoff function is the same in all treatments and depends both on his guess  $a$  and on the type  $t$ . It is given by the following formula:  $5 - |a - t|$ . When

---

<sup>14</sup>We had previously ran 5 sessions with the only objective to gather the performance of these 99 subjects in the IQ test. The beliefs elicitation procedure and payment is detailed in section 7.2.

<sup>15</sup>See Online Appendix 7 for an example of the screens seen by the Sender and by the Receiver. The complete instructions are given in Online Appendix 6.

the Receiver believes a type  $t'$  for sure, his optimal guess equals  $t'$ . We give subjects the exact formula, and explain them that their payoff is higher when their guess is closer to the true type.<sup>16</sup> In each game, both players' payoffs can range from 1 and to 5 euros. One of the 10 games is randomly selected for payment of Part 2.

## 3.2 Treatments

Our experiment has four treatments that result from the crossing of the two variations described below. The two-by-two design is implemented *between subjects*.

### 3.2.1 Variation 1: *Neutral vs. Loaded type*

We vary exogenously whether the type  $t$ , generated in Step 1 of each Sender-Receiver game, is *Loaded* or *Neutral*.

In the *Loaded* treatments, the type  $t$  corresponds to a measure of the relative performance of the Receiver in the IQ test completed in Part 1. Let us call this loaded type the *IQ-rank* of the Receiver, and explain how it is generated. In *Step 1* of each Sender-Receiver game, the computer randomly selects four subjects from the benchmark group of 99 subjects who did the IQ test previously. For each Receiver, the IQ-rank  $t \in \{1, 2, 3, 4, 5\}$  is then computed by comparing his performance to the performance of these four randomly-selected subjects (with ties broken at random):

- ◇  $t = 1$  when the Receiver has the highest performance in the group of five subjects or, said differently, when the Receiver ranked first,
- ◇  $t = 2$  when the Receiver has the second highest performance in the group of five subjects,
- ◇ etc.

Importantly, a new IQ-rank is computed in every game as four new subjects are randomly selected from the benchmark group in Step 1 of each game.

In the *Neutral* treatments, the type  $t$  also corresponds to a rank but it has no particular meaning. Let us call this type a *neutral rank*, and explain how it is generated. At the very beginning of Part 2, that is, before the 10 Sender-Receiver games are played, an integer

---

<sup>16</sup>The Receiver's payoff function we use is not exactly the same as in the theory section. In a fully-revealing equilibrium, which always exists in our experimental games, the Receiver's action stays skeptical after any message. Proposition 2 in Online Appendix 1 establishes that, with this payoff function, the Receiver's optimal action is equal to the median of his beliefs.

between 0 and 15 is randomly attributed to each Receiver.<sup>17</sup> In Step 1 of each Sender-Receiver game, the computer randomly selects four other integers between 0 and 15. For each Receiver, the neutral rank  $t \in \{1, 2, 3, 4, 5\}$  is then computed by comparing his integer to the four randomly-selected other integers (with ties broken at random):

- ◇  $t = 1$  when the Receiver's integer is the highest in the group of five integers,
- ◇  $t = 2$  when the Receiver's integer is the second highest in the group of five integers,
- ◇ etc.

Importantly, a new neutral rank is computed in every game as four new integers are randomly selected in Step 1 of each game.

We now discuss three important aspects of the first experimental variation.

**Receivers' preferences over beliefs.** In both the *Loaded* and the *Neutral* treatments, the Receiver precisely knows how the type is generated in Step 1 of the games. In particular, he knows the type is a measure of relative intelligence in *Loaded* and a rank with no particular meaning in *Neutral*. The main hypothesis behind the *Loaded* / *Neutral* manipulation is that it affects whether or not the Receiver has intrinsic preferences over what he believes about the type. In the *Loaded* treatment, we assume that the Receiver cares about the type and, everything else equal, has a preference for believing a higher IQ-rank (closer to 1). In the *Neutral* treatment, there is no such intrinsic preference over the rank.

**What do Senders know about the type generation?** In Step 1 of every game, the Sender is fully informed about the type  $t$ . However, in both the *Loaded* and *Neutral* treatments, the type is simply presented to the Sender as a "secret number from the set  $\{1,2,3,4,5\}$ ". In other words, the Sender does not know how  $t$  is generated in the different treatments. The Receiver knows that the Sender does not know how the type is generated. Because Senders do not know that the type has a different meaning for Receivers in the *Loaded* and *Neutral* treatments, their communication strategy are comparable across treatments (provided that the Sender's payoff is the same). This will allow us to focus on the Receivers' reactions facing comparable messages in *Neutral* and *Loaded* treatments. If Senders knew they were transmitting information about Receivers' relative IQ, it could affect their communication

---

<sup>17</sup>This procedure is meant to parallel the fact that, in *Loaded* treatments, Receivers start Part 2 with a fixed IQ performance between 0 and 15.

strategies in potentially complex ways (they could derive pleasure or discomfort disclosing good or bad news about Receivers relative IQ, etc.).

**Comparability between the *Loaded* and *Neutral* treatments.** We designed the *Neutral* treatment to be as possible to the *Loaded* treatment. In the *Loaded* treatment, as the 10 games are played, the Receiver may learn something about his performance in the IQ test, even if the IQ-rank is newly computed in every round. This potential learning, which depends on the information disclosed in every game, may affect the way future messages about the IQ-rank are interpreted. In the *Neutral* treatment, a similar process can occur because the Receiver is initially attributed an integer that is fixed for the 10 games. In both treatments, we do not give Receivers feedback between the games to limit this learning. One difference however remains. In the *Loaded* treatment, the Receiver may have some information about his IQ-rank because he has experienced the IQ test. In the *Neutral* treatment, he only knows his integer has been selected according to a uniform distribution. We examine the impact of Receivers' prior beliefs on their guesses in section 7.2.

### 3.2.2 Variation 2: *High* vs. *Low* guess

We vary exogenously whether the Sender's objective is to induce a *High* or a *Low* guess from the Receiver. This treatment manipulation affects what being skeptical means for the Receiver, as explained in the theory section. We remind that payoffs are common knowledge, so the Receiver knows the Sender's objective.

In the *High* treatments, the Sender's payoff is equal to the Receiver's guess  $a$ . With this payoff, the Sender earns more when the Receiver guesses a higher number or, equivalently, a lower rank. According to the theory, when  $m$  is vague, the Receiver should be skeptical and, as stated in Definition 2, believe the lowest type / highest rank in  $m$ . For instance, if  $m = \{3, 4, 5\}$ , the skeptical belief assigns probability one to  $t = 3$ .

In the *Low* treatments, the Sender's payoff is equal to  $6 - a$ . With this payoff, the Sender earns more when the Receiver guesses a lower number or, equivalently, a higher rank. According to the theory, when  $m$  is vague, the Receiver should be skeptical and, as stated in Definition 4, believe the highest type / lowest rank in  $m$ . For instance, if  $m = \{1, 2, 3\}$ , the skeptical belief assigns probability one to  $t = 3$ .

### 3.2.3 Two-by-two design

By crossing variations 1 and 2, we affect whether, for any given message, the Receiver’s skeptical belief is neutral, self-serving or self-threatening. In the *High\_Loaded* treatment, the skeptical belief is self-serving in the sense that the subject feels better believing the highest IQ-rank disclosed. In the *Low\_Loaded* treatment, the skeptical belief is self-threatening in the sense that the subject feels worse believing the lowest IQ-rank disclosed. Treatments are summarized in Table 1.

Table 1: Summary of Treatments

Treatment	Skeptical belief facing $m$	Intrinsic preferences over beliefs	Skeptical beliefs is
<i>High_Neutral</i>	Highest rank in $m$	None	Neutral
<i>High_Loaded</i>	Highest rank in $m$	Pref. for higher rank	Self-serving
<i>Low_Neutral</i>	Lowest rank in $m$	None	Neutral
<i>Low_Loaded</i>	Lowest rank in $m$	Pref. for higher rank	Self-threatening

### 3.3 Main hypotheses

Our main hypotheses relate to the Receivers’ levels of skepticism in the different treatments. We construct a measure of skepticism that captures the idea that a Receiver, while being consistent, is less skeptical when the distance between his guess and the skeptical guess (defined by Definition 3) is larger. To make the measure comparable across games, we normalize this distance by the maximal distance to the skeptical guess that any consistent guess could have, that is, by the size of the message seen by the Receiver. The measure is denoted  $Sk(a, m)$  and constructed for each game in which the Receiver made a consistent guess  $a$  when faced with a vague message  $m$ .<sup>18</sup> The measure is in  $[0, 1]$  and equals 1 when the Receiver makes the skeptical guess.

**Definition 5** In *High* treatments, for every vague  $m$  and consistent  $a$ , the measure of skepticism is given by:

$$Sk(a, m) = 1 - \frac{a - t_{inf}(m)}{t_{sup}(m) - t_{inf}(m)},$$

where  $t_{inf}(m)$  is the lowest type in  $m$  and  $t_{sup}(m)$  is the highest type in  $m$ .

<sup>18</sup>When considering precise messages, consistent guesses are necessarily skeptical.



In *Low* treatments, for every vague  $m$  and consistent  $a$ , the measure of skepticism is given by:

$$Sk(a, m) = 1 - \frac{t_{sup}(m) - a}{t_{sup}(m) - t_{inf}(m)},$$

where  $t_{sup}(m)$  is the highest type in  $m$  and  $t_{inf}(m)$  is the lowest type in  $m$ .

In the data, skepticism will be evaluated using averages of  $Sk(a, m)$  over games and frequencies of skeptical guesses. We use *Neutral* treatments as benchmarks and formulate our main hypotheses in terms of differences in skepticism between the *Loaded* and *Neutral* treatments. We also formulate an hypothesis regarding Receivers' skepticism in the *High\_Neutral* and *Low\_Neutral*.

**Hypothesis 1:** *Facing vague messages, Receivers are as skeptical in High\_Neutral as in Low\_Neutral.*

From the theoretical point of view, we see no reason a priori to think that skepticism is harder to exercise when it consists in assigning probability one to the lowest type (as in *High\_Neutral*) or to the highest type disclosed (as in *Low\_Neutral*).

**Hypothesis 2:** *Facing vague messages, Receivers are at least as skeptical in High\_Loaded as in High\_Neutral.*

In *High\_Loaded*, the skeptical belief is self-serving for the Receiver. If this has any effect on the level of skepticism, the Receiver should be more skeptical in *High\_Loaded* than in *High\_Neutral*. If Receivers are already fully skeptical in *High\_Neutral*, they will be as skeptical in *High\_Loaded*.

**Hypothesis 3:** *Facing vague messages, Receivers are less skeptical in Low\_Loaded than in Low\_Neutral.*

In *Low\_Loaded*, the skeptical belief is self-threatening for the Receiver. If this has any effect on the level of skepticism, the Receiver should be less skeptical in *Low\_Loaded* than in *Low\_Neutral*.

We can summarize the three hypotheses as follows: Skepticism in *Low\_Loaded* < Skepticism in *Low\_Neutral* = Skepticism in *High\_Neutral* ≤ Skepticism in *High\_Loaded*.

### 3.4 Implementation

A total of 464 subjects participated in the experiment. Subjects belonged to the pool of the WZB-TU Lab in Berlin, which is mostly made up of students from the Technical University of Berlin. They were invited to virtual experimental sessions, of about 20 subjects each, on Zoom. Once checked in, each subject received a link to start and run the experiment on his own computer (while staying on Zoom in the presence of the experimenter). The experiment was programmed with z-Tree (Fischbacher, 2007) and run using z-Tree unleashed (Duch et al., 2020). Experimental sessions took around one hour and subjects earned on average 19.51 euros (s.d.=2.13). The experiment had been pre-registered (AEARCTR-0007541).

### 3.5 Data analysis

**Sample restrictions.** In the data, we drop 16 Receivers for whom at least one game out of 10 could not be played because of computer bugs (either the Sender they were matched with could not send a message, or the Receiver himself could not make a guess). We also drop 16 Receivers who, in more than half of the games they played, made a guess which was not consistent with the evidence contained in the message received. We believe that these subjects did not understand well that Senders were disclosing hard evidence about the type.<sup>19</sup> Overall, this leaves 200 Receiver subjects. We focus our analysis on the 2000 games they played: 540 games in *High\_Neutral*, 480 in *High\_Loaded*, 500 in *Low\_Neutral* and 480 in *Low\_Loaded*. In the data set, each observation is a game that consists of a true type  $t$ , a message  $m$  sent by the Sender and a Receiver’s guess  $a$ .

**Statistical tests.** Unless noted otherwise, for all statistical tests we report p-values obtained from random-effects linear regressions on panel data with the Senders’ or Receivers’ identifiers as the group variable and the rounds as the time variable.<sup>20</sup> Standard errors are clustered at the session level using bootstrapping. In Appendix B, we provide robustness checks of the tests reported in the main text by exploring alternative specifications. They include (i) accounting for the bounded (sometimes binary) nature of the dependent variable

---

<sup>19</sup>As shown in Online Appendix 3, our main results are unaffected by the inclusion of these Receivers. We also explored the alternative possibility, in the vein of Jin et al. (2021), that Receivers are inconsistent because they care about the Senders’ payoffs. We find no support for this hypothesis.

<sup>20</sup>When studying Senders’ communication strategies, the group identifier variable is the Senders’ identity. When studying Receivers’ behavior, the group identifier variable is the Receivers’ identity.

by using Probit or Tobit models when appropriate, (ii) using linear regressions without considering panel data structure, and (iii) clustering at the individual level rather than at the session level.

## 4 Experimental results: first steps

### 4.1 Senders' communication strategies

We first describe Senders' strategies and, in particular, check that they account for the Senders' reversed objective in *High* and *Low* treatments.

**Senders' strategy in *High* treatments.** Table 2 reports the frequency with which each message is sent conditionally on the Sender observing each type  $t$ . In the *High* treatments, in which Senders want Receivers to make a high guess (closer to 5), Senders of type  $t = 5$  most often send the precise message  $m = \{5\}$  while Senders of type  $t = 1$  most often send the vaguest message  $m = \{1, 2, 3, 4, 5\}$ . In fact, even if there are variations across types, Senders of type  $t$  most often disclose that the type is *at least*  $t$ : they do so 67.84% of the time over all types. This disclosure strategy corresponds to selective disclosure described in the theory section (baseline games).<sup>21</sup>

Table 2: Senders' communication strategy in *High* treatments

Treatments <i>High</i>																
Type	Message															Total
	{1}	{2}	{3}	{4}	{5}	{1,2}	{2,3}	{3,4}	{4,5}	{1,2,3}	{2,3,4}	{3,4,5}	{1,2,3,4}	{2,3,4,5}	{1,2,3,4,5}	
1	5.24	-	-	-	-	2.62	-	-	-	9.17	-	-	15.28	-	<b>67.69</b>	100
2	-	6.11	-	-	-	-	3.93	-	-	0.87	17.90	-	2.18	<b>54.59</b>	14.41	100
3	-	-	6.97	-	-	-	1.49	11.94	-	0.50	4.98	<b>65.17</b>	1.49	2.49	4.98	100
4	-	-	-	16.76	-	-	-	1.16	<b>72.83</b>	-	1.16	5.78	1.16	-	1.16	100
5	-	-	-	-	<b>82.45</b>	-	-	-	9.04	-	-	4.26	-	2.13	2.13	100
Total	1.18	1.37	1.37	2.84	15.20	0.59	1.18	2.55	14.02	2.35	5.20	14.61	4.41	13.14	1.99	100

*Note:* The Table reports the frequency with which each message is sent conditionally on the Sender observing each type  $t$ , in the *High* treatments. Numbers in red highlight the most frequently sent message for each type. For instance, Senders of type  $t = 5$  send the precise message  $m = \{5\}$  82.45% of the time.

Regarding the comparison between *High\_Neutral* and *High\_Loaded*, we remind that Senders see the same instructions in these two treatments. For each type, the message sent

<sup>21</sup>The experiment of Deversi et al. (2021) in which the Senders can use a rich/flexible language corresponds to our *High\_Neutral* treatment with 6 instead of 5 possible types. The authors also report that the messages sent most often by the Senders almost perfectly coincide with the prediction of selective disclosure.

most often is the same in these two treatments, and the frequency with which this message is sent is never significantly different (except marginally when  $t = 4$ ,  $p = 0.076$ ).<sup>22</sup>

**Senders' strategy in *Low* treatments.** Table 3 gives the frequency with which each message is sent conditionally on the Sender observing each type  $t$ . In the *Low* treatments, in which Senders want Receivers to make a low guess (closer to 1), Senders of type  $t = 1$  most often send the precise message  $m = \{1\}$  while Senders of type  $t = 5$  most often send the vaguest message  $m = \{1, 2, 3, 4, 5\}$ . Senders of type  $t$  most often disclose that the type is *at most*  $t$ : they do so 70.82% of the time over all types. This disclosure strategy also corresponds to selective disclosure described in the theory section (alternative games).

Table 3: Senders' communication strategy in *Low* treatments

Treatments <i>Low</i>																
Type	Message															Total
	{1}	{2}	{3}	{4}	{5}	{1,2}	{2,3}	{3,4}	{4,5}	{1,2,3}	{2,3,4}	{3,4,5}	{1,2,3,4}	{2,3,4,5}	{1,2,3,4,5}	
1	<b>72.82</b>	-	-	-	-	10.77	-	-	-	6.15	-	-	3.59	-	6.67	100
2	-	7.10	-	-	-	<b>71.01</b>	1.78	-	-	7.69	0.59	-	2.37	5.33	4.14	100
3	-	-	2.76	-	-	-	5.52	1.10	-	<b>75.69</b>	3.31	2.21	2.76	3.87	2.76	100
4	-	-	-	1.22	-	-	-	3.27	0.41	-	15.51	1.63	<b>62.45</b>	4.90	10.61	100
5	-	-	-	-	0.53	-	-	-	3.68	-	-	6.32	-	14.74	<b>74.74</b>	100
Total	14.49	1.22	0.51	0.31	0.10	14.39	1.33	1.02	0.82	16.53	4.59	2.04	17.24	5.71	19.69	100

*Note:* The Table reports the frequency with which each message is sent conditionally on the Sender observing each type  $t$ , in the *Low* treatments. Numbers in red highlight the most frequently sent message for each type. For instance, Senders of type  $t = 1$  send the precise message  $m = \{1\}$  72.82% of the time.

Regarding the comparison between *Low\_Neutral* and *Low\_Loaded*, we remind again that Senders see the same instructions in these two treatments. For each type, the message sent most often is the same in these two treatments, and the frequency with which this message is sent is never significantly different (except marginally when the type is 4,  $p = 0.089$ ).<sup>23</sup>

**Result 1.** Senders' strategies are well-structured: in *High* treatments, Senders of type  $t$  most often disclose that the type is *at least*  $t$ ; in *Low* treatments, Senders of type  $t$  most often disclose that the type is *at most*  $t$ .

## 4.2 Messages seen by Receivers

The messages seen by Receiver in *High* and *Low* treatments are obviously different because of Senders' reversed objectives. We however can check that Receivers face comparable mes-

<sup>22</sup>Online Appendix 2 gives the frequencies of each  $m$  conditional on  $t$  in *High\_Neutral* and *High\_Loaded* separately.

<sup>23</sup>Online Appendix 2 gives the frequencies of each  $m$  conditional on  $t$  in *Low\_Neutral* and *Low\_Loaded* separately.

sages in the four treatments by considering two objects. The first is the size of the messages Receivers see in the various treatments, since it may be easier to make skeptical inferences when a smaller set of types is disclosed. The average size of the messages seen is not different between the *High\_Neutral* and *High\_Loaded* treatments (2.90 and 3.01,  $p = 0.463$ ). Messages are slightly vaguer in the *Low\_Neutral* treatment than in the *Low\_Loaded* (average sizes are 3.25 and 2.98;  $p = 0.014$ ). When studying Receivers’ skepticism, we will take into account the size of the messages they saw to control for these variations.<sup>24</sup> The second object is the skeptical guesses that correspond to the vague messages that Receivers see. For all vague messages received, the average skeptical guess is neither different between *High\_Neutral* and *High\_Loaded* (2.21 and 2.24,  $p = 0.686$ ), nor different between *Low\_Neutral* and *Low\_Loaded* (3.85 and 3.71,  $p = 0.362$ ).

**Observation 1.**<sup>25</sup> When making their guesses, Receivers are in comparable situations in all treatments. In the two *High* treatments, the messages seen by Receivers are similar in terms of size and of skeptical guesses that these messages induce. In the two *Low* treatments, the messages seen by Receivers are similar in terms of size and skeptical guesses that these messages induce, with slightly vaguer messages in *Low\_Neutral*.

### 4.3 Preliminary checks about Receivers’ guesses

Before studying whether Receivers’ guesses are skeptical, we make two preliminary checks about these guesses.

**Consistency of Receivers’ guesses.** Applying Definition 1 of the theory section, a guess is consistent with message  $m$  if it is in  $m$ . Over all treatments, the percentage of consistent guesses is 98.10%, which is very high.<sup>26</sup> It ranges from 95.74% in *High\_Neutral* to 99.37% in *Low\_Loaded*, with no significant differences between the treatments. The rate of consistency is lower when Receivers received a precise message (92.25%) than a vague message (99.50%,  $p = 0.004$ ) as is mechanically expected. When looking at the very few inconsistent guesses

---

<sup>24</sup>The measure of skepticism  $Sk(a, m)$  is normalized by the size of the message. When considering the frequency of skeptical guesses, we will use message size as a control variable.

<sup>25</sup>We use the term “result” for findings that are linked to our experimental treatments, and the term “observation” otherwise.

<sup>26</sup>This percentage is in line with the levels of consistency reported in Hagenbach and Perez-Richet (2018) or Schipper and Li (2020) who also allow for vague messages.

made by Receivers, we find no evidence that inconsistent guesses correspond to higher ranks.

**Observation 2.** Receivers' guesses are consistent with the evidence provided by Senders.

**Receivers' guesses in *Neutral* and *Loaded* treatments.** The Receivers' monetary payoff function is the same in all treatments. In *Loaded* treatments, we assume that Receivers additionally prefer to hold higher beliefs about their IQ-rank, so they should guess higher ranks in these treatments. Pooling all messages, Figure D.1 in Appendix D reports the distribution of Receivers' guesses in the *Neutral* and *Loaded* treatments. It shows that the frequency with which Receivers guess they ranked 2 or 3 is indeed higher in the *Loaded* than in the *Neutral* treatments ( $p = 0.024$  and  $p = 0.022$ , resp.). On the contrary, the frequency with which Receivers guess they ranked 5 is lower in the *Loaded* than in the *Neutral* treatment ( $p = 0.093$ ). We take this as an indication that our first experimental variation was effective.

**Observation 3.** Receivers guess higher ranks in *Loaded* than in *Neutral* treatments.


## 5 Receivers' skepticism

In this section, we study Receivers' skepticism by testing the three hypotheses presented in section 3.3. For every vague message  $m$  and consistent guess  $a$ , skepticism is measured using the formula  $Sk(a, m)$  from Definition 5. We remind that the closer the Receiver's guess is to the skeptical guess, the higher this measure. Overall, we look at the 1605 games in which the message is vague and the guess consistent, which corresponds to 80.25% of our data set.

### 5.1 Main results

Figure 1 gives an overview of our main findings. It displays the cumulative distribution function of the measure of skepticism for each treatment. The closer the line is to the bottom-right of the box, the more skeptical the Receivers' guesses are. On that figure, we see that the level of skepticism is rather high in the *High* treatments, without a big difference between the *High\_Neutral* and the *High\_Loaded* treatments. The level of skepticism seems lower in the *Low* treatments, with a bigger difference between the *Low\_Neutral* and the *Low\_Loaded*

treatments. Figure 1 suggests that the levels of skepticism in the four treatments is indeed ordered as hypothesized at the end of section 3.3.

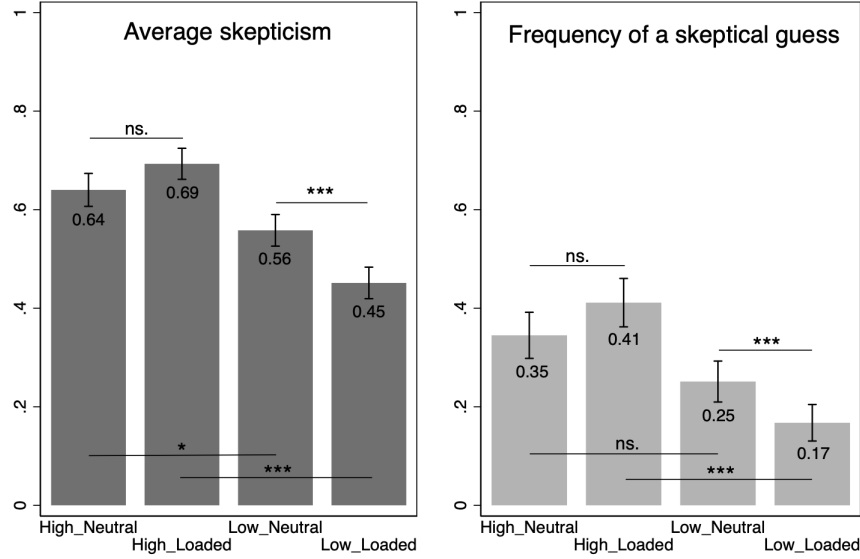


*Note:* The Figure displays the cumulative distribution function of skepticism, by treatment. The closer the line is to the bottom-right of the box, the more skeptical the Receivers' guesses are. On the contrary, the closer the line is to the top-left of the box, the less skeptical the Receivers' guesses are.

Figure 1: Cumulative distribution of skepticism, by treatment

The left part of Figure 2 displays the average value of  $Sk(a, m)$  by treatment and gives another representation of our findings. First, the average level of skepticism is higher in *High\_Loaded* than in *High\_Neutral* but the difference is not significant ( $p = 0.341$ ). This indicates that Receivers' skepticism is not significantly enhanced when the skeptical belief is self-serving. Second, the average level of skepticism in *Low\_Loaded* is strongly and significantly lower than in *Low\_Neutral* ( $p = 0.005$ ). On average, Receivers are less skeptical when skepticism involves believing a low IQ-rank than when it involves believing a low neutral rank. Third, the average levels of skepticism in *Low\_Neutral* and *High\_Neutral* are marginally different from each other ( $p = 0.076$ ). This small difference could be explained by the fact that it is easier for Receivers to understand the objective of a Sender whose payoff equals their guess than the one of a Sender whose payoff equals 6 minus their guess.<sup>27</sup> The above findings are confirmed when looking at the frequencies of skeptical guesses, reported on the right part of Figure 2.

<sup>27</sup>On the Sender side, an observation from previous Tables 2 and 3 points to the same direction: there is a higher rate of full disclosure of type 5 in *High* treatments than of type 1 in *Low* treatments ( $p = 0.056$ ).



*Note:* The left part of the Figure displays the average level of skepticism, by treatment. The right part of the Figure displays the average frequency of a skeptical guess, by treatment. Black segments are 95% confidence intervals. P-values are from random-effects linear regressions on panel data with the Receivers' id as the group variable and the rounds as the time variable. Standard errors are clustered at the session level using bootstrapping. See columns (1), (4) and (7) in Table 4 (left part of the Figure) and columns (1), (4) and (7) in Table C.1 (right part of the Figure). \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

Figure 2: Receivers' average skepticism, by treatment

Table 4 shows the determinants of the Receivers' level of skepticism. Columns (1) to (3) focus on the *High* condition. The coefficient of the *Loaded* dummy is small and insignificant in all three specifications, which confirms that the average level of skepticism is not significantly different between the *High\_Neutral* and *High\_Loaded* treatments. Columns (4) to (6) focus on the *Low* condition. The estimated negative coefficient of the *Loaded* dummy confirms that the level of skepticism is substantially and significantly lower in the *Low\_Loaded* than in the *Low\_Neutral* treatment. Columns (7) to (9) report coefficients of the full difference-in-differences specification. The coefficients of the interaction term is negative and significant, confirming the findings from columns (1)-(6). The estimated negative coefficient of the *Low* dummy shows the marginal difference between the *Neutral* treatments.

Columns (2), (5) and (8) include two additional control variables. First, they include the Receivers' performance in the IQ test, which we happen to measure for all Receivers in Part 1 of the experiment. We see that better cognitive abilities are significantly and positively correlated to Receivers' skepticism. We further explore the role of performance in the IQ test in section 7.1. Second, they include dummy variables for each round of play, which show no clear overall learning trend. In fact, learning may be hindered by the fact that there are



relatively few rounds, that Receivers do not face the same message in every round and do not get any feedback between the rounds (see Online Appendix 4 for more details). Columns (3), (6) and (9) additionally control for various demographic variables which include age, gender, educational attainment, socio-professional category, social class and experience in participating in experiments.

Table C.1 in Appendix C replicates Table 4 considering the measure of whether or not the guess is skeptical. When looking at the Receivers' likelihood to make a skeptical guess rather than at the distance between the Receiver's guess and the skeptical guess, the results of Table 4 are confirmed.

Table 4: Determinants of skepticism

<i>Dep. Var.</i>	<i>High treatments</i>			<i>Skepticism</i>			<i>Difference-in-difference</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1 if <i>Loaded</i>	0.048 (0.050)	0.017 (0.046)	0.016 (0.050)	-0.109*** (0.036)	-0.124*** (0.037)	-0.156*** (0.059)	0.048 (0.052)	0.022 (0.049)	0.019 (0.050)
1 if <i>Low</i>							-0.062* (0.035)	-0.065* (0.034)	-0.051 (0.039)
1 if <i>Low_Loaded</i>							-0.157** (0.063)	-0.149** (0.061)	-0.157** (0.065)
IQ performance		0.026*** (0.006)	0.026*** (0.009)		0.019*** (0.007)	0.017* (0.009)		0.023*** (0.005)	0.021*** (0.006)
Rounds dummies		✓	✓		✓	✓		✓	✓
<i>Demo.</i>			✓			✓			✓
Cons.	0.639*** (0.020)	0.348*** (0.064)	0.458** (0.187)	0.577*** (0.028)	0.391*** (0.061)	0.164 (0.261)	0.639*** (0.020)	0.399*** (0.056)	0.464*** (0.154)
<i>N</i>	789	789	789	816	816	816	1605	1605	1605

*Note:* The Table reports random-effects linear regressions on panel data with the Receivers' id as the group variable and the rounds as the time variable. Standard errors (in parentheses) are clustered at the session level using bootstrapping. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Overall, our results validate Hypothesis 2 (no difference between *High\_Neutral* and *High\_Loaded*) and Hypothesis 3 (lower skepticism in *Low\_Loaded* than in *Low\_Neutral*). Regarding the two *Neutral* treatments, we observe small differences in skepticism, so it is less clear that Hypothesis 1 is validated. In Appendix B, we report alternative specifications to test the robustness of the effects presented in Table 4 and Table C.1. Our validation of Hypotheses 2 and 3 are robust to all specifications considered. Robustness checks of Hypothesis 1 provide more mitigated results, the level of skepticism between the two *Neutral* treatments being either not significantly or marginally significantly different depending on the specification. In Online Appendix 5, we additionally compute the monetary losses that are induced by Receivers' lack of skepticism, and show that the loss is always larger in the

*Low\_Loaded* treatment than in any other treatment. Our results about Receivers' skepticism can be summarized as follows:

**Result 2.** (a) Receivers' skepticism is not significantly higher when it is self-serving than when it is not. (b) Receivers' skepticism is significantly lower when it is self-threatening than when it is not. (c) The average level of skepticism is marginally lower in *Low\_Neutral* than in *High\_Neutral*.

To some extent, Result 2 is in line with the main result from previous works on asymmetric updating (mentioned in the literature review), namely that subjects' inferences conform more closely to Bayes' rule in response to good than to bad news. However, this result is established in situations which drastically differ from the situations we consider. First, in the experimental literature on asymmetric updating, subjects usually can directly assess whether a signal is good or bad news. For instance, subjects receive signals about whether their performance in a test puts them in the lower or upper half of a group of subjects. In our work, Receivers cannot assess whether a message is good or bad from its face value; it is the strategic setting – *High* or *Low* treatment – and skeptical inferences in this setting, that allows them to assess how good a message is for them. Second, our subjects update their prior beliefs using a message which is strategically transmitted to them, and not a noisy signal whose generating process is known. It follows that, in all treatments alike, Receivers need to reason about the Sender's strategy before using it to update their beliefs. The uncertainty about the Sender's strategy may leave room for Receivers to form comfortable beliefs.

## 5.2 More or less self-serving or self-threatening skeptical guesses

We now dig further into previous findings by evaluating the extent to which a skeptical guess is self-threatening or self-serving in *Loaded* treatments. Table 5 presents the skeptical guess for every vague message in the *High* and *Low* treatments.

In the *High\_Loaded* treatment, we will say that skeptical guesses that correspond to ranks 3 or 4 are *strongly self-serving*, while skeptical guesses that correspond to 1 or 2 are *mildly self-serving*. This may seem a bit counter-intuitive but the extent to which a skeptical guess is self-serving is evaluated in comparison to other consistent guesses. In fact, when messages are  $\{3, 4\}$ ,  $\{4, 5\}$  or  $\{3, 4, 5\}$ , making a consistent guess which is not skeptical means

Table 5: Skeptical guess for every possible vague message in the *High* and *Low* treatments

Message	{1, 2}	{2, 3}	{3, 4}	{4, 5}	{1, 2, 3}	{2, 3, 4}	{3, 4, 5}	{1, 2, 3, 4}	{2, 3, 4, 5}	{1, 2, 3, 4, 5}
<i>High</i>	1	2	<b>3</b>	<b>4</b>	1	2	<b>3</b>	1	2	1
<i>Low</i>	2	3	<b>4</b>	<b>5</b>	3	<b>4</b>	<b>5</b>	<b>4</b>	<b>5</b>	<b>5</b>

*Note:* The Table reports the skeptical guess for every possible vague message in the *High* and *Low* treatments. For instance, in the *High* treatments, the skeptical guess corresponding to message  $m = \{1, 2\}$  is  $a = 1$ . It is  $a = 2$  in the *Low* treatments. Strongly self-serving and strongly self-threatening skeptical guesses appear in bold in the *High* and *Low* lines respectively.

guessing a particularly low rank (4 or 5). In that sense, the skeptical guess is strongly self-serving after these messages. For other messages, making a guess that is not skeptical means guessing a relatively high rank such as 2 or 3. In that sense, for messages such as  $\{1, 2\}$  or  $\{1, 2, 3\}$ , the skeptical guess is mildly self-serving. Similarly, in the *Low\_Loaded* treatment, we will say that skeptical guesses that correspond to ranks 2 or 3 are *mildly self-threatening*, while skeptical guesses that correspond to 4 and 5 are *strongly self-threatening*. Strongly self-serving and strongly self-threatening skeptical guesses appear in bold in Table 5.

Figure D.2 in Appendix D gives the average skepticism (in dark grey) and the frequency of skeptical guesses (in light grey), for mildly/strongly self-serving/self-threatening skeptical guesses. In the *High\_Loaded* treatment, Receivers' make the skeptical guess significantly more frequently when it is strongly self-serving than when it is mildly self-serving ( $p < 0.001$ ). In the *Low\_Loaded* treatment, the average level of skepticism and the frequency of a skeptical guess are significantly lower when the skeptical guess is strongly self-threatening than when it is mildly self-threatening ( $p < 0.001$  and  $p < 0.001$ ). This findings demonstrate that Receivers' exercise of skepticism depends on the precise conclusion that skepticism leads to.


**Result 3.** (a) In the *High\_Loaded* treatment, Receivers make skeptical guesses significantly more often when skepticism is strongly self-serving than when it is mildly so. (b) In the *Low\_Loaded* treatment, the average level of skepticism and the frequency of skeptical guesses is lower when skepticism is strongly self-threatening than when it is mildly so.

## 6 Receivers' skepticism at the individual level

In this section, we study guesses at the individual level and consider the steps of reasoning that could lead to each possible guess. To do that, we look at the 10 guesses of each individual, including guesses that are inconsistent or follow precise messages.

## 6.1 Individual guesses

On Figure 3, every bar represents the 10 guesses of one individual. We classify each guess into the 3 following categories: when  $Sk(a, m) > 0.5$ , the guess  $a$  is *rather skeptical* (in green); when  $Sk(a, m) < 0.5$ ,  $a$  is *rather anti-skeptical* (in red); when  $Sk(a, m) = 0.5$ ,  $a$  is naive (in dark grey).<sup>28</sup>



*Note:* The Figure displays the distribution of guesses in our five categories, by individual. Each bar corresponds to the 10 guesses of one individual. Each color corresponds to a category of guess.

Figure 3: Receivers' individual guesses, by category

Figure 3 clearly shows that Receivers behave in various ways, from the ones making 10 rather skeptical guesses (13%) to the ones making no rather skeptical guess at all (4%). Over all treatments, almost half of the subjects (44%) make the three kinds of guesses. On average, over 10 guesses, a Receiver makes 5.54 rather skeptical guesses, 2.05 naive guesses and 2.21 rather anti-skeptical guesses. It is interesting to look at how these numbers vary across treatments. A first observation is that they do not vary significantly between the

<sup>28</sup>If a Receiver is lost about the inference he is supposed to make about the disclosed ranks, he may “naively” pick the one in the middle simply because it is focal. This may be especially true for messages of odd size. We indeed observe a significantly higher fraction of naive guesses after messages of size 3 and 5 than after messages of size 2 and 4 ( $p < 0.001$ ).

two *High* treatments.<sup>29</sup> There are however differences between the two *Low* treatments: on average, the number of rather skeptical guesses per person is higher in *Low\_Neutral* than in *Low\_Loaded* (5.06 and 3.85,  $p = 0.064$ , ttest); the number of rather anti-skeptical guesses per person is lower in *Low\_Neutral* than in *Low\_Loaded* (2.56 and 3.54,  $p = 0.075$ , ttest); there is no difference in the average number of naive guess per person between the *Low* treatments ( $p = 0.564$ , ttest). In short, when moving from the *Low\_Neutral* to the *Low\_Loaded* treatment, there is on average one guess per person that switches from the category rather skeptical to the category rather anti-skeptical. Finally, if we compare the two *Neutral* treatments, we see that the average number of rather skeptical guesses per person is significantly higher in *High\_Neutral* than in *Low\_Neutral* (6.46 and 5.06,  $p = 0.017$ ).

**Result 4.** (a) Receivers behave similarly when skepticism is self-serving and when it is not. (b) Receivers make significantly fewer rather skeptical guesses and significantly more rather anti-skeptical guesses when skepticism is self-threatening than when it is not. (c) Receivers make significantly more rather skeptical guesses in *High\_Neutral* than in *Low\_Neutral*.

## 6.2 Steps of reasoning

In this subsection, we take an alternative look at Receivers' behavior by considering the steps of reasoning that could lead to each guess they made. The heuristic for strategic reasoning in disclosure games is borrowed from the works of Hagenbach and Perez-Richet (2018) and Schipper and Li (2020).<sup>30</sup> We explain the procedure by considering, as an example, message  $\{3, 4, 5\}$  seen in a *High* treatment. At the first step of reasoning, the Receiver understands that the Sender's messages are constrained by the truth, so he makes consistent guesses. At the second step, the Sender draws conclusions from this consistency. In particular, for a Sender of type 5, the message  $\{5\}$  leads to a weakly higher payoff than the vague message  $\{3, 4, 5\}$ . At the third step, the Receiver updates his possible interpretation of the message  $\{3, 4, 5\}$  and eliminates type 5 as a possible Sender, so he makes a guess in  $[3, 4]$ . At the fourth step, the Sender of type 4 is then weakly better-off fully disclosing than sending  $\{3, 4, 5\}$ .

<sup>29</sup>The average number of rather skeptical guesses is 6.46 in *High\_Neutral* and 6.69 in *High\_Loaded* ( $p = 0.663$ , ttest). The average numbers of naive guesses are 1.56 and 1.88 respectively, ( $p = 0.363$ , ttest). The average numbers of rather anti-skeptical guesses are 1.56 and 1.27 respectively ( $p = 0.467$ , ttest).

<sup>30</sup>These works propose procedures of elimination of strategies that are based on different concepts, respectively iterated elimination of weakly dominated strategies and iterated elimination of obviously dominated strategies, but make the same predictions in the monotonic games we study.

At the fifth step, the Receiver eliminates 4 as a possible type and makes the skeptical guess 3. For every message, the procedure terminates in a finite number of steps and the only Receiver’s guess that survives the procedure is the skeptical guess. Note that reaching the skeptical guess requires a number of steps that increases with the size of the message. Table D.1 in Appendix D gives the number of steps required to make each guess conditional on each possible message.

We now can count the steps of reasoning that Receivers make in the various treatments. We consider that a Receiver makes  $k$  steps of reasoning if at least one of his 10 guesses required  $k$  steps of reasoning. Note that we cannot see if a Receiver makes many steps if he only sees messages of small size. However, as argued in subsection 4.2, Receivers see messages of similar sizes in all treatments and we make between-treatments comparisons. In addition, even if Receivers do not all see messages of every possible size, 84.00% of subjects see messages of size 4 or 5. In every column of Table 6, we give the fraction of Receivers who make  $k$  steps.

Table 6: Fraction of Receivers (in %) reaching each reasoning step, by treatment

	1 step	3 steps	5 steps	7 steps	9 steps
<i>High_Neutral</i>	100	92.59	72.22	44.44	22.22
<i>High_Loaded</i>	100	93.75	95.42	50.00	14.58
<i>p-value</i>	-	<i>0.817</i>	<i>0.106</i>	<i>0.575</i>	<i>0.323</i>
<i>Low_Neutral</i>	100	96.00	86.00	54.00	24.00
<i>Low_Loaded</i>	100	93.75	68.75	18.75	4.17
<i>p-value</i>	-	<i>0.613</i>	<i>0.041</i>	<i>&lt;0.001</i>	<i>0.005</i>
<i>p-value (Neutral)</i>	-	<i>0.456</i>	<i>0.086</i>	<i>0.330</i>	<i>0.830</i>

*Note:* The Table reports the fraction of Receivers who make  $k$  steps of reasoning. P-values come from proportion tests. One observation by individual.

Table 6 shows that all Receivers make the first reasoning step, and more than 90% of Receivers perform 3 steps. Receivers’ make similar numbers of reasoning steps in the two *High* treatments. If we look at *Low* treatments however, a significantly larger fraction of Receivers perform 5, 7 or 9 steps in *Low\_Neutral* than in *Low\_Loaded*. This observation suggests that Receivers make fewer steps of reasoning towards the skeptical guess when this guess is self-threatening than when it is not. Hence, Receivers’ depth of reasoning seems affected by the conclusion that this reasoning leads to.

**Result 5.** (a) Receivers make similar numbers of reasoning steps towards the skeptical guess when skepticism is self-serving and when it is not. (b) Receivers make fewer steps of reasoning towards the skeptical guess when skepticism is self-threatening than when it is not. (c) Receivers make similar numbers of reasoning steps in *High\_Neutral* and *Low\_Neutral*.

## 7 Receivers' IQ and beliefs about relative IQ

In Part 1 of our experiment, we measure all Receivers' performance in an IQ test. We also elicit Receivers' beliefs about their relative performance in this test. In this final section, we examine how Receivers' performance and beliefs correlate with their guesses.

### 7.1 Skepticism and IQ

We begin by checking if there is a link between Receivers' performance in the IQ test and their ability to make skeptical inferences. Table 4 shows that, over all treatments, there is a small but significant, positive correlation between performance in the IQ test and average skepticism. To look at this correlation at the individual level, we split the group of Receivers into two based on the median performance in the IQ test: 93 subjects have a relatively *Low IQ* (they answered correctly up to 9 questions in the IQ test) and 107 subjects have a relatively *High IQ* (they answered correctly 10 to 15 questions in the IQ test). Pooling all treatments and vague messages, the average level of skepticism is significantly higher in the high IQ group than in the low IQ group (0.62 and 0.54 respectively,  $p = 0.004$ ). The result that a higher IQ is associated to more skepticism is in line with Schipper and Li (2020) who document that subjects' performance in a Raven IQ test are positively correlated to levels of reasoning in disclosure games.

**Observation 4.** Receivers' skepticism is positively correlated to their performance in the IQ test.

In our experiment, we use Receiver's performance in the IQ test to construct type (IQ-ranks) in *Loaded* treatments. This has a subtle effect: a Receiver's performance affects the messages he sees in the *Loaded* treatments. To understand why this is the case, consider the Sender's communication strategy that consists in disclosing that the type is at least  $t$

in *High\_Loaded*. According to this strategy, vaguer messages are sent when the Receivers' IQ-rank is higher.<sup>31</sup> The converse is true in *Low\_Loaded*: vaguer messages are sent when the Receivers' IQ-rank is lower. Even if the correlation between Receivers' performance and IQ-rank is not perfect, it follows that subjects facing vague messages have, on average, a lower performance in the IQ test in *Low\_Loaded* than in *High\_Loaded*. One may then wonder whether, in the light of the correlation between IQ and skepticism, it could be that skepticism is lower in *Low\_Loaded* simply because subjects have lower performance in the IQ test. A first argument against this possibility is that we control for the performance of Receivers when studying skepticism in our regressions (Tables 4 and C.1). Second, and more importantly, we compare skepticism between the *Low* treatments in which the performance of Receivers is actually higher in *Low\_Loaded* than in *Low\_Neutral* (the average performance of Receivers is 10.14 and 9.14 respectively,  $p = 0.026$ , ttest).

## 7.2 Skepticism and prior beliefs about IQ performance

In *Loaded* treatments, Receivers form beliefs about their IQ-ranks while having experienced the IQ test in Part 1 of the experiment. Receivers may therefore have some beliefs about their performance in the test that potentially affect the way they interpret Senders' messages. In particular, if individuals are very confident about their performance and tend to stick to this prior view, they may mechanically appear not very skeptical in the *Low\_Loaded* treatment (where being skeptical means guessing a low rank), and rather skeptical in the *High\_Loaded* treatment (where being skeptical means guessing a high rank).

To investigate the role of prior beliefs, we elicited, in all treatments, Receivers' beliefs about their relative performance right after the IQ test. Precisely, we inserted each subject into a group of 99 subjects who had done the same IQ test previously. We then divided the group of 100 individuals into five quintiles – from the 20% of individuals with the best performance to the 20% of individuals with the worst performance – and asked subjects to give an estimate of the likelihood they belong to each quintile. Even if these beliefs about relative performance do not correspond exactly to prior beliefs about the IQ-ranks generated in the Sender-Receiver games, they give an indication of how well the Receiver thinks he performed. On average, Receivers consider that they are more likely to belong to the second

---

<sup>31</sup>The message is  $\{1, 2, 3, 4, 5\}$  when the IQ-rank is 1, and  $\{5\}$  when the IQ-rank is 5.



and third quintiles than to the fourth and fifth quintiles.<sup>32</sup> This shows that individuals are rather confident regarding their performance in the IQ test, which is consistent with the large literature on overconfidence (e.g. Moore and Healy, 2008; Merkle and Weber, 2011).

To investigate whether the lower skepticism in *Low\_Loaded* could be driven by the average tendency to confidence, we check that our results are robust to the exclusion of the most confident subjects.<sup>33</sup> There are several ways to identify (and exclude) these subjects. A first, simple way to evaluate confidence is to consider the quintile that each Receiver estimated most likely. We consider that the most confident Receivers are either the ones who estimate that they are most likely to belong to the first quintile, or the ones who estimate that they are most likely to belong to one of the first two quintiles. Excluding these subjects, Figure 4 replicates Figure 2, the figure presenting the main results in Section 5. We see that both the average level of skepticism and the frequency of a skeptical guess stays lower in *Low\_Loaded* than in *Low\_Neutral*.

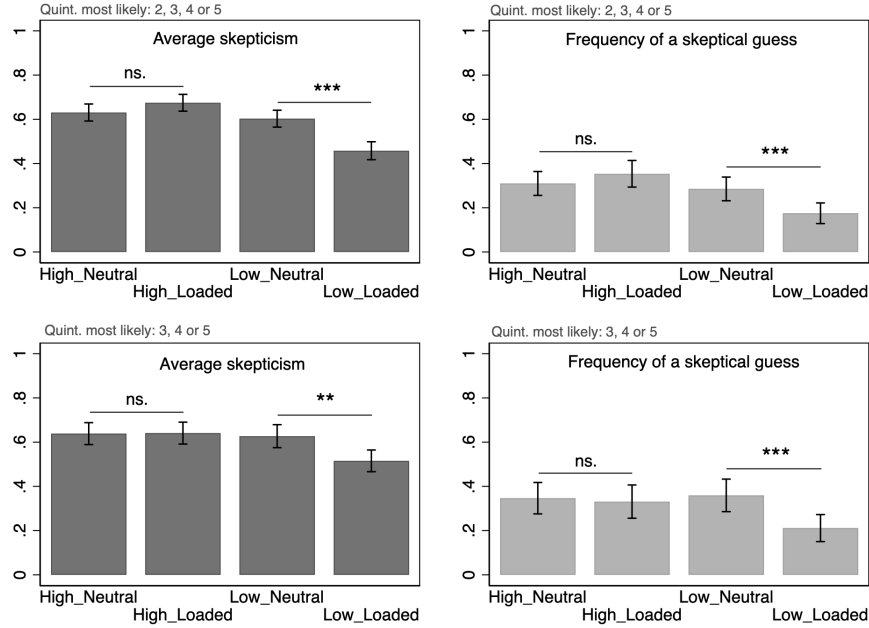
A second way to evaluate a Receiver's confidence is to examine whether his belief about performance, i.e., the distribution of the estimated likelihood to belong to each quintile, is skewed to the left. We consider that the most confident subjects are the ones whose belief differs significantly from the uniform belief at the 1% level (Chi2 test) and is skewed to the left (considers the first and second quintiles more likely than the fourth and fifth). Figure 5 replicates Figure 2 excluding these confident subjects. It shows that, even among the less confident subjects, the average level of skepticism and the frequency of a skeptical guess is lower in *Low\_Loaded* than in *Low\_Neutral* ( $p = 0.039$  and  $p = 0.012$ , resp.).

**Observation 5.** Skepticism is significantly lower when it is self-threatening than when it is not, even when excluding confident Receivers who may lack skepticism simply because they stick to their positive prior beliefs.

---


<sup>32</sup>Pooling Receivers from all treatments, Figure D.3 in Appendix D displays the averages of the estimated likelihood to belong to each quintile.

<sup>33</sup>We will apply our exclusion criterion to all treatments alike. Even if confidence does not play a role in *Neutral* treatments (as Receivers guess neutral ranks), it will ensure that we compare Receivers with similar IQ levels, types and messages across treatments.



Note: The top-left part of the Figure displays the average level of skepticism, by treatment, excluding subjects who estimate most likely that they belong to the first quintile. The top-right part of the Figure displays the average frequency of a skeptical guess, by treatment, excluding subjects who estimate most likely that they belong to the first quintile. The bottom-left and bottom-right parts of the Figure exclude subjects who estimate most likely that they belong to the first or to the second quintile. Black segments are 95% confidence intervals. P-values are from random-effects linear regressions on panel data with the Receivers' id as the group variable and the rounds as the time variable. Standard errors are clustered at the session level using bootstrapping. \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

Figure 4: Receivers' average skepticism by treatment, without confident subjects (quintile estimated as most likely)



Note: The left part of the Figure displays the average level of skepticism, by treatment, excluding confident subjects. The right part of the Figure displays the average frequency of a skeptical guess, by treatment, excluding confident subjects. Black segments are 95% confidence intervals. P-values are from random-effects linear regressions on panel data with the Receivers' id as the group variable and the rounds as the time variable. Standard errors are clustered at the session level using bootstrapping. \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

Figure 5: Receivers' average skepticism by treatment, without confident subjects (distribution of priors)

## 8 Conclusion

We designed an experiment to study how individuals read strategically-transmitted information when they have preferences over what they want to learn. In our setting, the reading of information precisely consists in making skeptical inferences when faced with vague statements. We vary whether Receivers make inferences about loaded or neutral types and, more importantly, whether these inferences lead to attractive or unattractive conclusions. When skeptical inferences lead to a view that Receivers dislike to hold, we find that skepticism is low. This is true while Receivers are able to exercise skepticism in neutral settings or when skeptical inferences lead to good news for them. At the aggregate level, skepticism is measured using the distance to the skeptical guess and the frequency of skeptical guesses. We also measure it at the individual level showing, in particular, that Receivers' levels of reasoning towards skeptical guesses are lower when skepticism is self-threatening than in any other case.

A central question in the literature on disclosure is whether or not disclosure should be mandated to help individuals take as well-informed decisions as possible. Our work suggests that it may be particularly important to mandate disclosure of information when it conveys truths that are hard to face. We can think for instance about companies disclosing the extent to which their production processes respect the environment through the use of various labels. The absence of green label should be interpreted as bad news. Consumers have various reasons not to interpret the absence of labels as bad news, ranging from the mere fact that reasoning is costly or complex to the fact that absent labels do not necessarily attract attention and trigger reasoning. We shed light on a psychological reason that may reinforce consumers' tendency to avoid reasoning when faced with unclear information, namely that it is sometimes more comfortable not to know.

Adding to the existing literature on motivated beliefs, we report motivated reading of information in a strategic situation. In our experiment, Receivers are well aware of the Senders' objective and probably pushed by the experimental setting to think about what Senders want them to guess. Our results suggest that the will to read through another player's strategy does not push Receivers to read information in a particularly accurate way. The strategic setting we consider also raises the important question of whether Senders expect motivated skepticism and could exploit it. Going back to the example of green labels

mentioned above, firms expecting the lack of skepticism from consumers can stay relatively inactive regarding the protection of the environment. In our experiment, we do not study the equilibrium effect of the Receivers' lack of skepticism as, to focus on Receivers, we do not make Senders' aware of the meaning of the state they disclose for the Receivers.<sup>34</sup>

Our results leave open another interesting question, namely how exactly Receivers reach the favored conclusion, sometimes by going against a relatively simple game-theoretical logic that they otherwise understand. One possibility is that Receivers understand well that vague messages correspond to types that have no interest in revealing themselves more precisely, but they exploit the wiggle room left by the noise in Senders' strategies. Said differently, Receivers may convince themselves, when skepticism is bad news, that the Senders probably did a mistake, disclosing more low ranks than what is optimal for instance. Studying what happens if this room were reduced could advance our understanding of the motivated reading of information in strategic settings.

## References

- Ali, S. N., Lewis, G., and Vasserman, S. (2021). Voluntary disclosure and personalized pricing. *Review of Economic Studies*, forthcoming.
- Bederson, B. B., Jin, G. Z., Luca, Leslie, P., Quinn, A. J. M., and Zou, B. (2018). Incomplete disclosure: Evidence of signaling and countersignaling. *American Economic Journal: Microeconomics*, 10:41–66.
- Bénabou, R. (2015). The economics of motivated beliefs. *Revue d'économie politique*, 125(5):665–685.
- Bénabou, R., Falk, A., and Tirole, J. (2019). Narratives, imperatives, and moral reasoning. *mimeo*.
- Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117(3):871–915.
- Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American economic review*, 96(5):1652–1678.
- Bénabou, R. and Tirole, J. (2011). Identity, morals and taboos: Beliefs as assets. *Quarterly Journal of Economics*, 126(2):805–855.
- Benndorf, V., Kübler, D., and Normann, H.-T. (2015). Privacy concerns, voluntary disclosure of information, and unraveling: an experiment. *European Economic Review*, 75:43–59.

---

<sup>34</sup>Theoretical papers, such as Milgrom and Roberts (1986), Hagenbach and Koessler (2017) or Deversi et al. (2021), have already studied the possibility that the Receiver is not fully skeptical in disclosure games. It induces failures of the unraveling mechanism when the language available to the Sender is simple, that is, when he can disclose his type fully or not at all. When the language is rich as in our experiment, the lack of skepticism pushes the Sender to use selective disclosure, which our Senders mostly already do.

- Burro, G. and Castagnetti, A. (2022). Will i tell you that you are smart (dumb)? deceiving others about their iq or about a random draw. *Journal of Behavioral and Experimental Economics*, 100:101920.
- Caplin, A. and Eliaz, K. (2003). Aids policy and psychology: A mechanism-design approach. *The RAND Journal of Economics*, 34:631–646.
- Carlson, R. W., Maréchal, M. A., Oud, B., Fehr, E., and Crockett, M. J. (2020). Motivated misremembering of selfish decisions. *Nature communications*, 11(1):1–11.
- Charness, G. and Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104:1–23.
- Chen, S. and Heese, C. (2021). Fishing for good news: Motivated information acquisition. *Working paper*.
- Chew, S. H., Huang, W., and Zhao, X. (2020). Motivated false memory. *Journal of Political Economy*, 128(10):3913–3939.
- Deversi, M., Ispano, A., and Schwardmann, P. (2021). Spin doctors: An experiment on vague disclosure. *European Economic Review*, 139:103872.
- Dickhaut, J. W., Ledyard, M., Mukherji, A., and Sapra, H. (2003). Information management and valuation: An experimental investigation. *Games and Economic Behavior*, 44(1):25–53.
- Dranove, D. and Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48:935–963.
- Drobner, C. and Goerg, S. J. (2021). Motivated belief updating and rationalization of information. *Working paper*.
- Duch, M. L., Grossmann, M. R., and Lauer, T. (2020). z-tree unleashed: A novel client-integrating architecture for conducting z-tree experiments over the internet. *Journal of Behavioral and Experimental Finance*, 28:100400.
- Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–38.
- Epley, N. and Gilovich, T. (2016). The mechanics of motivated reasoning. *Journal of Economic perspectives*, 30(3):133–40.
- Exley, C. L. and Kessler, J. B. (2021). Information avoidance and image concerns. *Economic Journal*, Revise & Resubmit.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178.
- Galeotti, F., Saucet, C., and Villeval, M. C. (2020). Unethical amnesia responds more to instrumental than to hedonic motives. *Proceedings of the National Academy of Sciences*, 117(41):25423–25428.
- Gino, F., Norton, M. I., and Weber, R. A. (2016). Motivated bayesians: Feeling moral while acting egoistically. *Journal of Economic Perspectives*, 30(3):189–212.
- Giovannoni, F. and Seidmann, D. J. (2007). Secrecy, two-sided bias and the value of evidence’. *Games and Economic Behavior*, 59(2):296–315.

- Gödker, K., Jiao, P., and Smeets, P. (2020). Investor memory. *Working paper, available at SSRN 3348315*.
- Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24:461–483.
- Grossman, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science*, 60(11):2659–2665.
- Grossman, Z. and Van der Weele, J. J. (2017). Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, 15(1):173–217.
- Hagenbach, J. and Koessler, F. (2017). Simple versus rich language in disclosure games. *Review of Economic Design*, 21.
- Hagenbach, J., Koessler, F., and Perez-Richet, E. (2014). Certifiable pre-play communication: Full disclosure. *Econometrica*, 82:1093–1131.
- Hagenbach, J. and Perez-Richet, E. (2018). Communication with evidence in the lab. *Games and Economic Behavior*, 112:139.
- Huffman, D., Raymond, C., and Shvets, J. (2019). Persistent overconfidence and biased memory: Evidence from managers. *American Economic Review*, conditional acceptance.
- Jin, G. Z., Luca, M., and Martin, D. (2021). Is no news (perceived as) bad news? an experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, 12:141–173.
- King, R. R. and Wallin, D. E. (1991). Market induced information disclosure: An experimental markets investigation. *Contemporary Accounting Research*, 8:170–197.
- Kőszegi, B. (2006). Emotional agency. *The Quarterly Journal of Economics*, 121(1):121–155.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological bulletin*, 108(3):480.
- Loewenstein, G., Sunstein, C. R., and Golman, R. (2014). Disclosure: Psychology changes everything. *Annual Review of Economics*, 6:391–419.
- Luca, M. and Smith, J. (2015). Strategic disclosure: The case of business school rankings. *Journal of Economic Behavior & Organization*, 112:17–25.
- Mathios, A. D. (2000). The impact of mandatory disclosure on product choices: An analysis of the salad dressing market. *Journal of Law and Economics*, 43(2):651–77.
- Merkle, C. and Weber, M. (2011). True overconfidence: The inability of rational information processing to account for apparent overconfidence. *Organizational Behavior and Human Decision Processes*, 116(2):262–271.
- Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12:380–391.
- Milgrom, P. and Roberts, J. (1986). Relying on the information of interested parties. *Rand Journal of Economics*, 17(1):18–32.
- Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science*.

- Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological review*, 115(2):502.
- Müller, M. (2021). Experimental evidence on selective memory of big life decisions using 10 years of panel survey data. *Working paper*.
- Okuno-Fujiwara, A., Postlewaite, M., and Suzumura, K. (1990). Strategic information revelation. *Review of Economic Studies*, 57:25–47.
- Oprea, R. and Yuksel, S. (2021). Social exchange of motivated beliefs. *Journal of the European Economic Association*.
- Raven, J. C. (1936). Mental tests used in genetic studies: The performance of related individuals on tests mainly educative and mainly reproductive. *MSc Thesis, University of London*.
- Saucet, C. and Villeval, M. C. (2019). Motivated memory in dictator games. *Games and Economic Behavior*, 117:250–275.
- Schipper, B. C. and Li, Y. X. (2020). Strategic reasoning in persuasion games: an experiment. *Games and Economic Behavior*, 121:329–367.
- Schwardmann, P. (2019). Motivated health risk denial and preventative health care investments. *Journal of Health Economics*, 65:78–92.
- Schwardmann, P., Deversi, M., and Ispano, A. (2021). Spin doctors: An experiment on vague disclosure. *European Economic Review*, 193.
- Schwardmann, P., Tripodi, E., and Van der Weele, J. J. (2022). Self-persuasion: Evidence from field experiments at two international debating competitions. *American Economic Review*, 112(4):1118–46.
- Schwardmann, P. and Van der Weele, J. (2019). Deception and self-deception. *Nature human behaviour*, 3(10):1055–1061.
- Seidmann, D. J. and Winter, E. (1997). Strategic information transmission with verifiable messages. *Econometrica*, 65(1):163–169.
- Serra-Garcia, M. and Szech, N. (2021). The (in) elasticity of moral ignorance. *Management Science*.
- Sharot, T., Korn, C. W., and Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nature neuroscience*, 14(11):1475–1479.
- Solda, A., Ke, C., Page, L., and Von Hippel, W. (2020). Strategically delusional. *Experimental Economics*, 23(3):604–631.
- Thaler, M. (2022). The fake news effect: Experimentally identifying motivated reasoning using trust in news. *working paper*.
- Thompson, L. and Loewenstein, G. (1992). Egocentric interpretations of fairness and interpersonal conflict. *Organizational Behavior and Human Decision Processes*, 51(2):176–197.
- Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–61.

# Appendices

## A The unraveling result

**Proof of Proposition 1.** First, consider by contradiction an equilibrium in which at least two types of the Sender, say  $t$  and  $t'$  with  $t > t'$ , send the same (necessarily vague) message  $m$ . Facing  $m$ , the Receiver's beliefs  $\beta_m$  assign a positive probability to both types  $t$  and  $t'$ . The higher type  $t$  then has a strict interest in deviating and sending the precise message  $m' = \{t\}$  because  $\sigma_R(m') = t > \sigma_R(m)$ . This demonstrates that the Sender uses a fully-revealing strategy in every equilibrium. Second, assume that, when facing the vague message  $m$ , the Receiver's beliefs are not skeptical:  $\beta_m$  assigns some positive probability to a type in  $m$  which is not the lowest. Then, the lowest type in  $m$ , say  $t$ , has an interest in deviating from sending the message prescribed by any fully-revealing strategy and sending message  $m$  (to get action  $\sigma_R(m) > t$ ). This is in contradiction with the former point which establishes that the Sender's strategy is fully-revealing in every equilibrium.

## B Statistical Tests

The p-values reported in the main text come from random-effects linear regressions on panel data with the Senders' or Receivers' id as the group variable and the rounds as the time variable. Standard errors are clustered at the session level using bootstrapping. While this specification has the advantage of being portable (we can use the same throughout the paper), it does not directly account for the sometimes limited nature of our dependent variables (for instance, *skepticism* is bounded between 0 and 1, *is\_skeptical* is a dummy, etc.).

Table B provides robustness checks of the tests reported in the main text by exploring alternative specifications. These include (i) directly accounting for the bounded nature of the dependent variable by using Probit or Tobit models when appropriate, (ii) using linear regressions without considering panel data structure, and (iii) clustering at the individual rather than at the session level.

Overall, our main results are fairly robust: Hypotheses 2 and 3 are rejected in none of the 24 specifications considered. Evidence are more mixed regarding Hypothesis 1: 11 specifications reject Hypothesis 1 while 13 fail to reject it. This is consistent with our main result 2(c) which concludes that the aggregate level of skepticism is lower in *Low\_Neutral* than in *High\_Neutral*, but only marginally.



Table B.1: P-values of statistical tests

Model Subject Cluster Panel	Linear RE Session ✓	Linear RE Id ✓	Linear RE Session	Linear RE Id	Tobit RE Session	Tobit RE Id	Probit RE Session	Probit RE Id
<i>Without controls</i>								
$Skept_{High\_Neu} = Skept_{Low\_Neu}$	0.075	0.159	0.031	0.066	0.022	0.086		
$Skept_{High\_Loa} \geq Skept_{High\_Neu}$	0.335	0.278	0.334	0.227	0.285	0.203		
$Skept_{Low\_Loa} < Skept_{Low\_Neu}$	0.004	0.017	0.006	0.018	0.002	0.030		
$Skept_{Low\_Loa} < Skept_{High\_Loa}$	<0.001	<0.001	<0.001	<0.001	0.001	<0.001		
$IsSkept_{High\_Neu} = IsSkept_{Low\_Neu}$	0.181	0.401	0.057	0.163			0.033	0.174
$IsSkept_{High\_Loa} \geq IsSkept_{High\_Neu}$	0.271	0.299	0.370	0.347			0.347	0.341
$IsSkept_{Low\_Loa} < IsSkept_{Low\_Neu}$	0.005	0.033	0.013	0.036			0.006	0.028
$IsSkept_{Low\_Loa} < IsSkept_{High\_Loa}$	<0.001	<0.001	0.002	<0.001			<0.001	<0.001
<i>With controls</i>								
$Skept_{High\_Neu} = Skept_{Low\_Neu}$	0.126	0.239	0.046	0.126	0.050	0.168		
$Skept_{High\_Loa} \geq Skept_{High\_Neu}$	0.736	0.729	0.918	0.909	0.808	0.793		
$Skept_{Low\_Loa} < Skept_{Low\_Neu}$	0.002	0.003	0.008	0.006	0.002	0.010		
$Skept_{Low\_Loa} < Skept_{High\_Loa}$	<0.001	<0.001	0.001	<0.001	<0.001	<0.001		
$IsSkept_{High\_Neu} = IsSkept_{Low\_Neu}$	0.272	0.489	0.080	0.205			0.047	0.219
$IsSkept_{High\_Loa} \geq IsSkept_{High\_Neu}$	0.587	0.642	0.982	0.984			0.957	0.961
$IsSkept_{Low\_Loa} < IsSkept_{Low\_Neu}$	0.002	0.010	0.013	0.020			0.005	0.012
$IsSkept_{Low\_Loa} < IsSkept_{High\_Loa}$	<0.001	0.001	0.004	0.001			<0.001	<0.001

## C Alternative measure of skepticism

Table C.1 replicates Table 4 considering a dependant variable which equals 1 if the guess is skeptical and 0 if not. The coefficient of the treatment dummy in columns (1) to (3) is small and insignificant, which shows that the Receivers' likelihood to make the skeptical guess is not significantly different in *High\_Neutral* and *High\_Loaded*. We also see no significant difference in the likelihood to make a skeptical guess between the *Neutral* treatments. On the contrary, the estimated negative coefficient of the treatment dummy in columns (4) to (6) reveals that the Receivers' likelihood to make the skeptical guess is substantially and significantly lower in *Low\_Loaded* than in *Low\_Neutral*. All specifications control for the size of the message received by the Receiver. Whatever the size, the skeptical guess always corresponds to guessing one specific rank among the ones disclosed. When the number of disclosed ranks gets larger, it may become mechanically less likely or cognitively harder to make the skeptical guess. The coefficient of the message size is negative and significant indicating that the likelihood to make a skeptical guess indeed decreases as more ranks are disclosed.

Table C.1: Determinants of a skeptical guess

<i>Dep. Var.</i>	= 1 if the guess is skeptical, 0 if not								
	<i>High treatments</i>			<i>Low treatments</i>			<i>Difference-in-difference</i>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1 if <i>Loaded</i>	0.065 (0.059)	0.038 (0.053)	0.028 (0.063)	-0.118*** (0.044)	-0.135*** (0.032)	-0.173*** (0.059)	0.065 (0.059)	0.042 (0.058)	0.032 (0.060)
1 if <i>Low</i>							-0.043 (0.034)	-0.045 (0.031)	-0.033 (0.040)
1 if <i>Low_Loaded</i>							-0.191*** (0.068)	-0.183*** (0.068)	-0.187** (0.081)
Mess. size	-0.158*** (0.018)	-0.159*** (0.019)	-0.161*** (0.016)	-0.089*** (0.014)	-0.090*** (0.013)	-0.090*** (0.014)	-0.124*** (0.013)	-0.125*** (0.013)	-0.126*** (0.013)
IQ performance		0.022* (0.0134)	0.023* (0.014)		0.022*** (0.007)	0.019* (0.011)		0.019*** (0.007)	0.018** (0.008)
Rounds dummies		✓	✓		✓	✓		✓	✓
<i>Demo.</i>			✓			✓			✓
Cons.	0.889*** (0.060)	0.578*** (0.113)	0.593*** (0.198)	0.604*** (0.073)	0.357*** (0.096)	0.0482 (0.337)	0.773*** (0.046)	0.523*** (0.067)	0.588*** (0.214)
<i>N</i>	789	789	789	816	816	816	1605	1605	1605

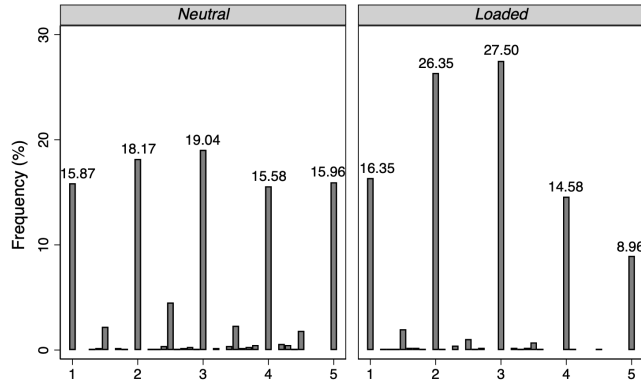
*Note:* The Table reports random-effects linear regressions on panel data with the Receivers' id as the group variable and the rounds as the time variable. Standard errors (in parentheses) are clustered at the session level using bootstrapping. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## D Additional Tables and Figures

Table D.1: Reasoning steps


Message	<i>High treatments</i>					<i>Low treatments</i>				
	1 step	3 steps	5 steps	7 steps	9 steps	1 step	3 steps	5 steps	7 steps	9 steps
{1}	1	-	-	-	-	{1}	1	-	-	-
{2}	2	-	-	-	-	{2}	2	-	-	-
{3}	3	-	-	-	-	{3}	3	-	-	-
{4}	4	-	-	-	-	{4}	4	-	-	-
{5}	5	-	-	-	-	{5}	5	-	-	-
{1, 2}	[1,2]	1	-	-	-	{1, 2}	[1,2]	2	-	-
{2, 3}	[2,3]	2	-	-	-	{2, 3}	[2,3]	3	-	-
{3, 4}	[3,4]	3	-	-	-	{3, 4}	[3,4]	4	-	-
{4, 5}	[4,5]	4	-	-	-	{4, 5}	[4,5]	5	-	-
{1, 2, 3}	[1,3]	[1,2]	1	-	-	{1, 2, 3}	[1,3]	[2,3]	3	-
{2, 3, 4}	[2,4]	[2,3]	2	-	-	{2, 3, 4}	[2,4]	[3,4]	4	-
{3, 4, 5}	[3,5]	[3,4]	3	-	-	{3, 4, 5}	[3,5]	[4,5]	5	-
{1, 2, 3, 4}	[1,4]	[1,3]	[1,2]	1	-	{1, 2, 3, 4}	[1,4]	[2,4]	[3,4]	4
{2, 3, 4, 5}	[2,5]	[2,4]	[2,3]	2	-	{2, 3, 4, 5}	[2,5]	[3,5]	[4,5]	5
{1, 2, 3, 4, 5}	[1,5]	[1,4]	[1,3]	[1,2]	1	{1, 2, 3, 4, 5}	[1,5]	[2,5]	[3,5]	[4,5]

*Note:* The Table gives the number of steps required to make each guess conditional on each possible message. In each cell, we report the guess that corresponds, for a given message (in row), to a given number of steps (in column). For example, in the *High treatments* (left table), making a guess equal to 4 conditional on message {4, 5} requires 3 steps of reasoning.



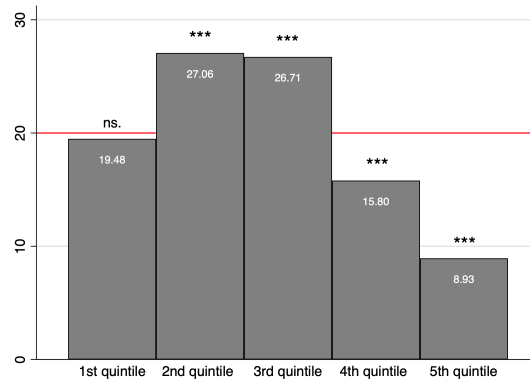
Note: The Figure displays the frequency of Receivers' guesses in the *Neutral* and *Loaded* treatments. For instance, 19.04% of the guesses in the *Neutral* treatments were  $a = 3$ . It is 27.50% in the *Loaded* treatments.

Figure D.1: Distribution of Receivers' Guesses



Note: The Figure displays the average level of skepticism (dark grey) and the frequency of skeptical guesses (light grey) in the *High\_Loaded* treatment (resp. *Low\_Loaded*), when the skeptical guess is mildly or strongly self-serving (resp. self-threatening). Black segments are 95% confidence intervals. P-values are from random-effects linear regressions on panel data with the Receivers' id as the group variable and the rounds as the time variable. Standard errors are clustered at the session level using bootstrapping. \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

Figure D.2: Skepticism when it is mildly/strongly self-threatening/self-serving



Note: The Figure displays the averages of the estimated likelihood to belong to each quintile. For instance, individuals estimate that they have on average 19.48% chances to belong to the first quintile. Significance levels are from two-sided t-tests that the estimated percentage is different to 20%. \*\*\*  $p < 0.01$ .

Figure D.3: Receivers' average likelihood to belong to each quintile