



HAL
open science

Motivated Skepticism

Jeanne Hagenbach, Charlotte Saucet

► **To cite this version:**

| Jeanne Hagenbach, Charlotte Saucet. Motivated Skepticism. 2022. hal-03770685v1

HAL Id: hal-03770685

<https://hal.science/hal-03770685v1>

Preprint submitted on 6 Sep 2022 (v1), last revised 18 Mar 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

Motivated Skepticism*

Jeanne Hagenbach[†] Charlotte Saucet[‡]

July 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. Skepticism is lower when skeptical beliefs are self-threatening than in neutral settings. When skeptical beliefs are self-serving, skepticism is not enhanced compared to neutral settings. These results demonstrate that individuals' exercise of skepticism depends on the conclusions of skeptical inferences.

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

JEL Code: C72, C91, D82, D91.

*The paper was initially circulated under the title “Motivated vs. Skeptical Beliefs”. We thank Victor Augias, Kai Barron, Roland Bénabou, Alessandra Casella, Frédéric Koessler, Johannes Leutgeb, Yves Le Yaouanq, Eduardo Perez-Richet, 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 research has received financial support from the European Research Council (grant 850996 – MOREV).

[†]CNRS, Sciences Po, CEPR - jeanne.hagenbach@sciencespo.fr

[‡]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.

Our experiment brings together the literature on *disclosure games* and the literature on *motivated beliefs*. 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. Many situations of strategic communication share this 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 wish to form accurate beliefs but they may form motivated beliefs, that is, form beliefs simply because they are comforting or pleasant.

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. This 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. To do that, we implement two exogenous variations in our experimental Sender-Receiver games. By crossing these two variations, we change whether or not the skeptical beliefs are aligned 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, namely to believe they ranked relatively high, but not in *Neutral* treatments. The use of an ego-relevant state in *Loaded* treatments is an artefact to create preferences over beliefs in the lab, a technique used by Schwardmann and Van der Weele (2019) and Zimmermann (2020) among others. We use it to study how these preferences affect the reading of hard information. While we vary the demand for motivated beliefs across treatments, we keep fixed the monetary value of holding accurate beliefs: both in *Loaded* and *Neutral* treatments, Receiver-subjects earn more money when their guess is closer to the type truly seen by the Sender.

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 types who sends vague messages or, said differently, low ranks that are concealed. With the *High/Low* variation, we change whether the skeptical reading of a vague message consists in believing it comes from the lowest or 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*. In *High_Neutral* and *Low_Neutral* treatments, skeptical beliefs respectively assign probability one to the highest and lowest ranks disclosed but these ranks have no intrinsic value for Receivers.

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.¹ 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 for them, and less so when it is bad news for 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

¹The reference for pre-registration is AEARCTR-0007541.

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. Additionally, they 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. On their side, Receivers take the evidence into account in the sense that their guesses are almost always in between the lowest and highest types disclosed. In our analysis, we compare skepticism in the *Loaded* and *Neutral* treatments while keeping the *High* or *Low* condition fixed. The two *Neutral* treatments thus serve as benchmarks. When comparing Receivers’ skepticism in these 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 reaching an unpleasant conclusion, namely that 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 their 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. Instead, and this is the main message, the exercise of skepticism crucially depends on the conclusions that skeptical inferences lead to. Put differently, individuals exercise skepticism in a motivated way. Motivated skepticism provides a new, psychological reason for the failure of unraveling which is often observed in the field and contradicts standard theory.²

We push our main results further by evaluating, for every message, the extent to which

²Understanding 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 several explanations of why this is the case: disclosure may not be costless, the Sender may not be fully informed, etc.

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 – that is, correspond to a particularly low rank – 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. Again, it is the exact conclusion that Receivers are supposed to reach that affects whether or not they interpret vague message with skepticism. 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 significantly fewer steps towards the skeptical guess when this guess is self-threatening than when it is neutral.

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 examine if they are correlated to Receivers’ guesses. When skepticism is self-serving, subjects with higher priors about themselves are significantly more skeptical than subjects with lower priors. The converse is true when skepticism is self-threatening.

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.³ This result crucially relies on Receivers reading information

³See Okuno-Fujiwara et al. (1990), Seidmann and Winter (1997), Giovannoni and Seidmann (2007) and Hagenbach et al. (2014) for related theoretical works.

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 unraveling in theory and in practice. Jin et al. (2021) and Schwardmann et al. (2021) closely replicate the standard game of Milgrom (1981). In their lab settings, an important fraction of Receivers, but not all, make skeptical inferences when facing absent or vague information. They respectively show that subjects’ ability to make these inferences depends on their experience in playing the disclosure game, and on the language used by the Sender.⁴ Some experiments additionally incorporate realistic perturbations to the standard game, such as costs to disclosure or a probability that the Sender is not informed. 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. 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 may affect unraveling in important ways, a point also made in Loewenstein et al. (2014). Our experiment is the first to consider that the Receivers’ reading of information may depend on the preference they have over beliefs.

The idea that, in some contexts, individuals’ utility is directly impacted by their beliefs is at the center of the recently-growing literature on motivated beliefs (surveyed in Bénabou, 2015).⁵ 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.⁶ In the related experiments, subjects form beliefs about states they intrinsically care about, such as their relative intelligence, beauty, generosity etc. They do so based on

⁴Schipper 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).

⁵Models 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).

⁶These channels include asymmetric information processing, selective recall and motivated information selection. Asymmetric information processing is documented in various experiments such as Eil and Rao (2011), Sharot et al. (2011), Mobius et al. (2011), Charness and Dave (2017) or Drobner and Goerg (2021). 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), Chew et al. (2020), Carlson et al. (2020), Gődker et al. (2020) and Mőller (2021). Motivated information selection, including avoidance and acquisition, 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. A general discussion of the mechanics of motivated reasoning can be found in Epley and Gilovich (2016), and earlier in Kunda (1990).

noisy information provided to them by the experimenter, not by another strategic player. We borrow from this literature the *Neutral / Loaded* variation but adapt it to a context of strategic communication. We then play with the Sender’s objective to affect whether Receivers’ skeptical reading of information leads to good or bad news in *Loaded* treatments. In a related paper, Thaler (2022a) shows that subjects assess the veracity of information sources in directions which correspond to pre-conceived political views. We ask whether individuals make skeptical inferences in a way that allows them to reach favored conclusions.

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 (2022b), Senders adapt their cheap-talk communication to what Receivers want to hear about political issues. In Hagmann and Loewenstein (2017), Senders are paradoxically more efficient in making Receivers change their mind about emotionally-charged topics when appearing less persuasive. Oprea and Yuksel (2021) propose an experiment in which subjects form beliefs about their relative intelligence in a social context. They show that subjects rely more on peers’ beliefs when these beliefs are self-serving. We consider a situation of 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 the design of our experimental benchmarks 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). In the beginning of the game, 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 $u_R(a; t)$ is strictly concave in a and reaches its maximum when a equals t .

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.⁷ As an example, consider $T = \{1, 2, 3, 4, 5\}$. The set of messages available to a Sender of type $t = 1$ is $M(1) = \{\{1\}, \{1, 2\}, \{1, 2, 3\}, \{1, 2, 3, 4\}, \{1, 2, 3, 4, 5\}\}$. Let $M = \cup_{t \in T} M(t)$ be the set of all messages available to the Sender. 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 β_m is *consistent* with m if β_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) β_m is derived from Bayes' rule after any message m sent on the equilibrium path, (iii) for every m , $\sigma_R(m)$ is the action that maximizes the Receiver's payoff given β_m . When the Receiver believes one type with probability one, say type t , his optimal action is $\sigma_R(m) = t$. When the Receiver believes several types with positive probability, the optimal action lies strictly in between the lowest and highest types believed. Point (ii) implies that Receivers' beliefs will be consistent after any message m sent on equilibrium path. We add the requirement, common in disclosure games, that β_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

⁷This 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.

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 β_m is *skeptical* if it assigns probability one to the lowest type in m .

In our experiment, we will observe the Receivers' actions but not directly their beliefs. It is therefore useful to define a *skeptical action* as follows:

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

In the baseline game, the skeptical belief attributes probability one to the lowest type in m , so 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.*

Our experiment will consider neutral settings which closely fit the baseline game, as well as settings in which the Receiver will intrinsically care about the beliefs he holds about t . In the latter settings, we will measure the Receiver's deviations from skeptical beliefs.

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 β_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 .⁸ 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 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*.

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

Part 1: IQ test

Subjects begin by completing a test made up of 15 Raven matrices (Raven, 1936). Subjects have 15 minutes to take the test and we remind them that it is frequently used to measure intelligence. 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.⁹

⁸Milgrom and Roberts (1986) already point to this result which is proved in Hagenbach and Koessler (2017).

⁹We 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.

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 send 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 . See Appendix I for an example of the screen seen by the Sender.
- 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. See Appendix I for an example of the screen seen by the Receiver.

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 in this game. When we describe the treatments below, it should become clear why we made this choice of giving no feedback between the games.

Payoffs. The 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 the Receiver believes a type t' for sure, his optimal guess equals t' . We give subjects the exact formula used to compute their payoff, and explain them that their payoff is higher when their guess is closer to the true type.¹⁰ 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.

¹⁰The Receiver's payoff function we use does not have the property of strict concavity in a considered in the theory section. In Appendix A, we discuss our choice and its implication for the theoretical predictions. Proposition 2 in this Appendix establishes that the Receiver's optimal action is equal to the median of his beliefs.

Comprehension questions and final questionnaire. Before playing the ten games, each subject must correctly answer some comprehension questions about (i) how Senders and Receivers are matched, (ii) which messages are available to the Senders, (iii) whether the type is constant across games and (iv) how payoffs are computed. At the end of the experiment, subjects answer a psychological questionnaire (Rosenberg, 2015) and a demographic questionnaire that includes age, gender, educational attainment, etc. Subjects are then informed of their aggregate earnings and leave.

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 between 0 and 15 is randomly attributed to each Receiver.¹¹ 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. As explained in the literature review, this treatment variation relies on previous experiments which study the formation of beliefs by subjects who intrinsically cares about these beliefs.

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 should be the same in these treatments (provided that

¹¹This procedure is meant to parallel the fact that, in *Loaded* treatments, Receivers start Part 2 with a fixed IQ performance between 0 and 15.

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 to Receivers information about their relative IQ, it could affect their communication strategies in potentially complex ways (they could derive pleasure or discomfort disclosing good or bad news about Receivers relative IQ, etc.). For now, we decided to shut down such effects which could, in turn, influence the way Receivers read disclosed information.

Comparability between the *Loaded* and *Neutral* treatments. We designed the *Neutral* treatment in a way that makes it as comparable 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. Despite these efforts to make the *Loaded* and *Neutral* treatments comparable, one difference 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. In section 7.2, we study the correlation between Receivers’ prior beliefs about their performance (elicited in Part 1) and the reading of evidence in *Loaded* treatments.

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* Receiver’s guess. This treatment manipulation affects what being skeptical means for the Receiver, as explained in Remark 1 of 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 a , the Receiver’s guess. 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 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 than holding any other consistent belief. 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 than holding any other consistent belief. Preferred and skeptical beliefs in the four 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 / lowest type in m	None	-
<i>High_Loaded</i>	Highest rank / lowest type in m	Pref. for higher rank	Self-serving
<i>Low_Neutral</i>	Lowest rank / highest type in m	None	-
<i>Low_Loaded</i>	Lowest rank / highest type 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. To formulate these hypotheses, 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. The skeptical guess is equal to the lowest type disclosed in *High* treatments and to the highest type disclosed in *Low* treatments. 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 .¹²

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 .

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 .

This measure is in $[0, 1]$ and equals 1 when the Receiver makes the skeptical guess. In our work, *Neutral* treatments serve as a benchmark, and we formulate our 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*, an experimental comparison that has never been done. We have three main hypotheses. In the data, skepticism will be evaluated using averages of $Sk(a, m)$ over games and frequencies of skeptical guesses.

Hypothesis 1: *Facing vague messages, Receivers are as skeptical in High_Neutral as in Low_Neutral.*

From the theoretical point of view, there is no difference between the two *Neutral* treatments. 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 or to the highest type disclosed.

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*.

¹²When considering precise messages, consistent guesses are necessarily skeptical.

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: 120 in *High_Neutral*, 130 in *High_Loaded*, 118 in *Low_Neutral* and 96 in *Low_Loaded*. The subjects belonged to the subject 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 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. 232 subjects played the role of Senders and 232 the role of Receivers. Each subject was supposed to play 10 Sender-Receiver games but the experiment took place online and a few participants encountered computer bugs. In the data, we drop 16 Receivers for whom at least one game out of 10 could not be played (either because the Sender they were matched with could not send a message, or because 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.¹³ Overall, this leaves 200 Receiver-subjects. We focus our analysis

¹³As shown in Appendix B, our main results are unaffected by the inclusion of these Receivers.

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.¹⁴ Standard errors are clustered at the session level using bootstrapping. In Appendix C, 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 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.¹⁵

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

¹⁴When studying Senders’ communication strategies, the group identifier variable is the Senders’ identity. When studying Receivers’ behavior, the group identifier variable is the Receivers’ identity.

¹⁵The 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.

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	-	67.69	100
2	-	6.11	-	-	-	-	3.93	-	-	0.87	17.90	-	2.18	54.59	14.41	100
3	-	-	6.97	-	-	-	1.49	11.94	-	0.50	4.98	65.17	1.49	2.49	4.98	100
4	-	-	-	16.76	-	-	-	1.16	72.83	-	1.16	5.78	1.16	-	1.16	100
5	-	-	-	-	82.45	-	-	-	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.

sent most often is the same in these two treatments, and the frequency with which this message is sent is never significantly different between *High_Neutral* and *High_Loaded* (except marginally when $t = 4$, $p = 0.076$). Table D.1 in Appendix D separately displays the frequencies of each message conditional on t in *High_Neutral* and *High_Loaded*.

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.

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	72.82	-	-	-	-	10.77	-	-	-	6.15	-	-	3.59	-	6.67	100
2	-	7.10	-	-	-	71.01	1.78	-	-	7.69	0.59	-	2.37	5.33	4.14	100
3	-	-	2.76	-	-	-	5.52	1.10	-	75.69	3.31	2.21	2.76	3.87	2.76	100
4	-	-	-	1.22	-	-	-	3.27	0.41	-	15.51	1.63	62.45	4.90	10.61	100
5	-	-	-	-	0.53	-	-	-	3.68	-	-	6.32	-	14.74	74.74	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 between *Low_Neutral* and *Low_Loaded* (except marginally when the type is 4, $p = 0.089$). Table D.2 in Appendix D separately displays the frequencies of each message conditional on t in *Low_Neutral* and *Low_Loaded*.

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 they type is *at most* t .

4.2 Messages seen by Receivers

We now check that Senders’ strategies generate comparable messages for Receivers in the four treatments. Two remarks are in order when comparing the distribution of messages received across treatments. First, the messages seen in *High* and *Low* treatments are necessarily different because of Senders’ reversed objectives. Second, and even if we take these objectives into account,¹⁶ the distributions of messages cannot be exactly the same in all treatments because they depend on the types seen by Senders and on the noise in their strategies. We report the distribution of types in Table D.3 in Appendix D. Regarding the noise on the Senders’ side, we believe that having human Senders instead of computers sending messages is a strength of our design. Receivers only know the objective of the informed parties and must try to understand their strategies, as in real strategic settings. With computerized Senders, we could have reduced the noise but would have had to fully describe messaging strategies to the Receivers.

To check that Receivers make guesses in comparable situations in the four treatments, we can look at 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* or *High_Neutral* treatments (average sizes are 3.25, 2.98 and 2.90; $p = 0.014$ and $p = 0.024$). When studying Receivers’ skepticism, we

¹⁶By considering that message {5} in *High* treatments is equivalent to message {1} in *Low* treatments, that message {4, 5} in *High* treatment is equivalent to message {1, 2} in *Low* treatments, and so on.

will take into account the size of the messages they saw to control for these variations.¹⁷ 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 equals is not different in *High_Neutral* and *High_Loaded* (2.21 and 2.24, $p = 0.686$). For all vague messages received, the average skeptical guess is not different in *Low_Neutral* and *Low_Loaded* (3.85 and 3.71, $p = 0.362$).

Observation 1.¹⁸ 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 guesses are skeptical in subsequent sections, we now make two preliminary checks about Receivers’ guesses.

Consistency of Receivers’ guesses. Applying Definition 1 of the theory section, a guess a is consistent with message m if a is in m . Over all treatments, the percentage of consistent guesses is 98.10%, which is very high.¹⁹ 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 when they received a vague message (99.50%, $p = 0.004$) as mechanically expected. When looking at the very few inconsistent guesses made by Receivers, we find no evidence that inconsistent guesses correspond to higher ranks, that is, that inconsistency could be motivated in *Loaded* treatments.

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

¹⁷The 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.

¹⁸We use the term “result” for findings that are linked to our experimental treatments, and the term “observation” otherwise.

¹⁹This 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.

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 5 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.

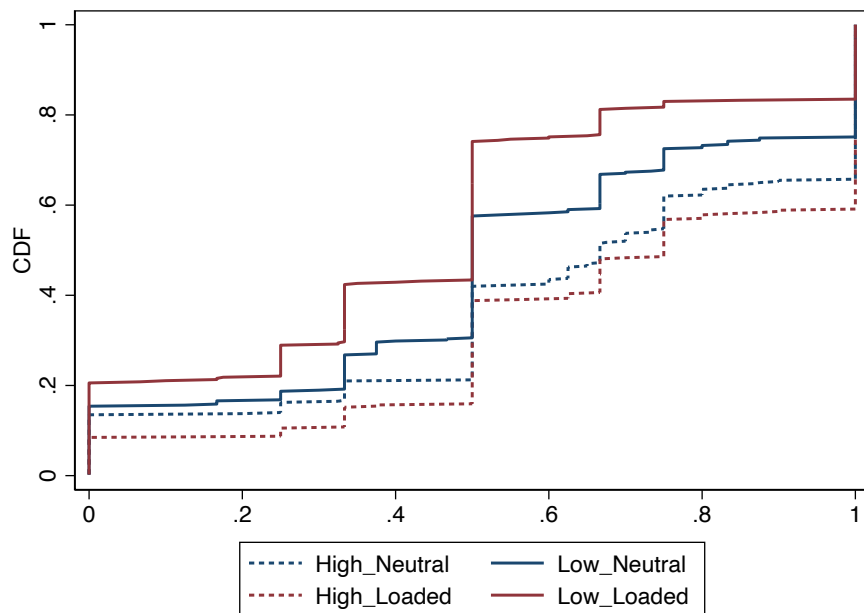
Observation 3. Receivers guess higher ranks in *Loaded* than in *Neutral* treatments.

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 the measure of skepticism. This measure ranges from 0 to 1 and is normalized by the size of the message. 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.

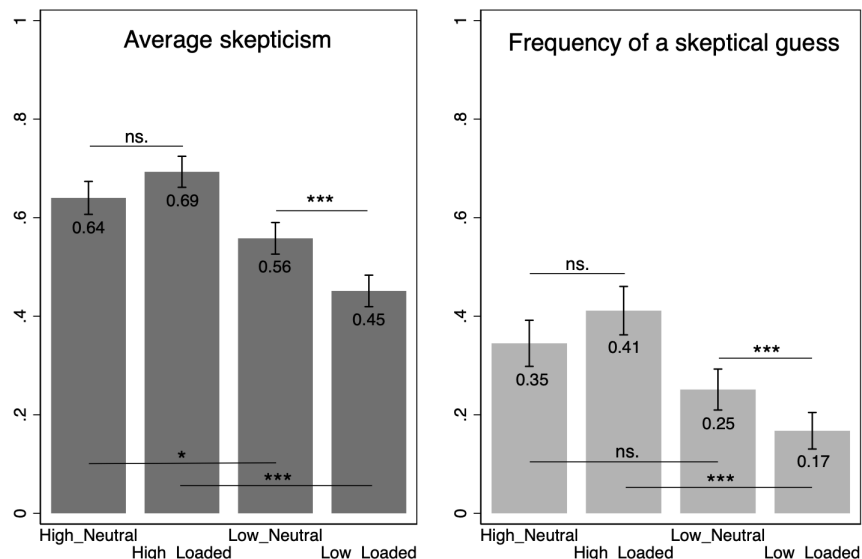


Note: The Figure displays the cumulative distribution function of skepticism, by treatment. The closer the line is to the bottom-right of the box, the more skeptical the Receivers' guesses are. On the contrary, the closer the line is to the top-left of the box, the less skeptical the Receivers' guesses 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 thus 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.²⁰ The above findings are confirmed when looking at the frequencies of skeptical guesses, reported on the right part of Figure 2. The frequency of skeptical guesses however does not differ between the two *Neutral* treatments.

²⁰On 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 E.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, Receivers do not face the same message in every round and do not get any feedback between the rounds. 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 4: Determinants of skepticism

<i>Dep. Var.</i>	Skepticism								
	<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>	0.048 (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 the Appendix, we offer various checks of these results. Table E.1 in Appendix E 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 E.1 additionally shows that the likelihood of a skeptical guess significantly decreases as more types are disclosed.

In Appendix C, we report alternative specifications to test the robustness of the effects presented in Table 4 and Table E.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 Appendix F, we also

propose to look at Receivers’ guesses by classifying them into various categories, ranging from skeptical to not skeptical at all. The analysis of the distributions of guesses in the four treatments confirm the following main results.

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*.

In our experiment, being skeptical allows Receivers to maximize their monetary payoffs when Senders use the selective disclosure strategy, which they mostly do. The lack of skepticism therefore induces monetary losses that we can evaluate simply by taking the distance between the guess made and the skeptical guess given any message. The Receivers’ average loss per game is €0.95 in *High_Neutral*, €0.79 in *High_Loaded*, €1.19 in *Low_Neutral* and €1.35 in *Low_Loaded*. In Appendix G, we additionally show that, given any message size, the loss is always larger in the *Low_Loaded* treatment than in any other treatment.

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 gives the skeptical guess for every vague message in the *High* and *Low* treatments.

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

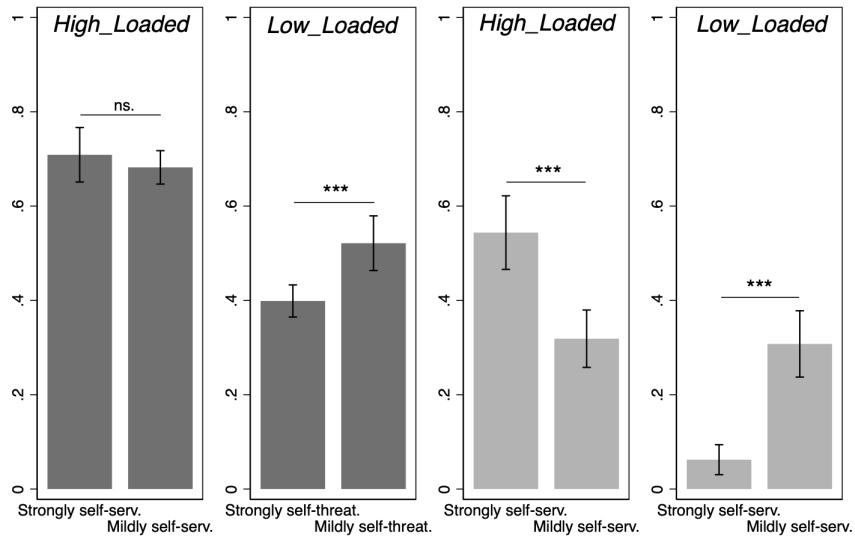
Message	{1, 2}	{2, 3}	{3, 4}	{4, 5}	{1, 2, 3}	{2, 3, 4}	{3, 4, 5}	{1, 2, 3, 4}	{2, 3, 4, 5}	{1, 2, 3, 4, 5}
<i>High</i>	1	2	3	4	1	2	3	1	2	1
<i>Low</i>	2	3	4	5	3	4	5	4	5	5

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.

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 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 3 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, we see that 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 really depends on the conclusion that skepticism leads to.



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 3: Skepticism when it is mildly/strongly self-threatening/self-serving

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 Receivers 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.

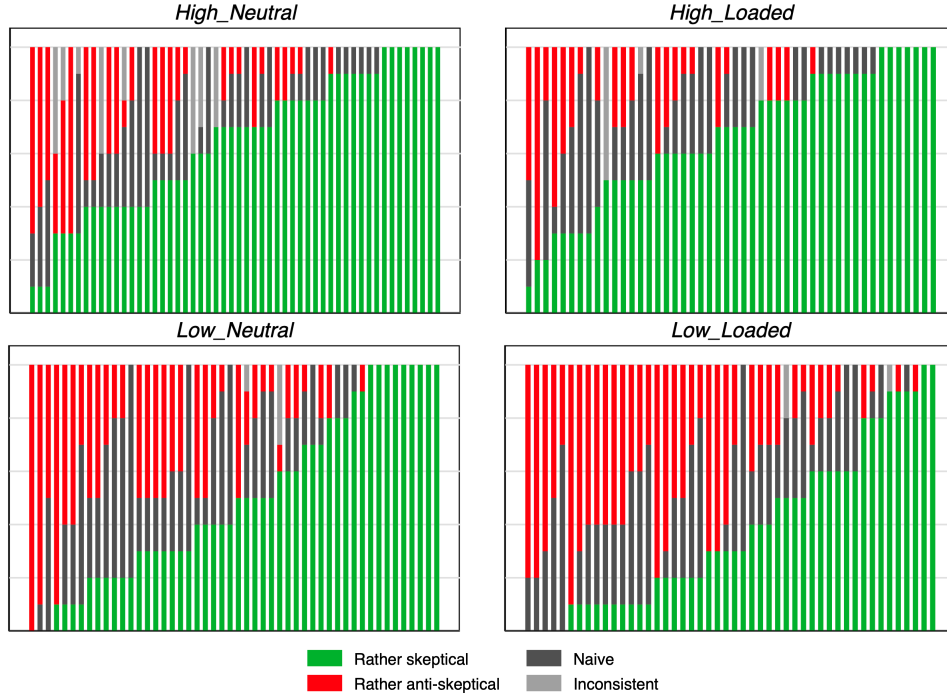
6.1 Individual guesses

On Figure 4, every bar represents the 10 guesses of one individual. There are 200 subjects in total, between 48 and 54 per treatment. 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).²¹

Figure 4 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 two *High* treatments.²² 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

²¹If 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 ($p < 0.001$).

²²The 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).



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 4: Receivers' individual guesses, by category

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).²³ 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 $\{3, 4, 5\}$. 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.4 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 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. Receivers’ depth of reasoning seems affected by the conclusion that this reasoning leads to.

²³These 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.

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><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.

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 IQ and skepticism

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.²⁴ 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 Table E.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 10.14 and 9.14 respectively, $p = 0.026$, ttest).²⁵

We now make a last remark about skepticism in the two IQ groups. Both in the Low-IQ and High-IQ groups, skepticism is lower in *Low_Loaded* than in *Low_Neutral* with the difference being significant only in the Low-IQ group ($p < 0.001$ while $p = 0.153$ for the High-IQ group). We however cannot interpret this finding as showing that Low-IQ subjects are more prone to exercise motivated skepticism than High-IQ subjects. Indeed, because of the way we construct types, the following is true in the *Low_Loaded* treatment: 77.84% of vague messages correspond to strongly self-threatening skeptical guesses in the Low-IQ group

²⁴The message is $\{1, 2, 3, 4, 5\}$ when the IQ-rank is 1, and $\{5\}$ when the IQ-rank is 5.

²⁵Regarding the other comparisons, we report no significant difference between the average level of skepticism in *High_Neutral* and *High_Loaded* and the average performance of Receiver subjects is the same in both treatments (9.18 and 9.93 respectively, $p = 0.166$, ttest). There is no difference in the average performance between the *High_Neutral* and *Low_Neutral* treatments ($p = 0.924$, ttest).

against 38.76% in the High-IQ group. Result 3 establishes that skepticism is lower when strongly self-threatening. If we consider only vague messages that correspond to strongly self-threatening skeptical guesses, the level of skepticism is marginally lower in *Low_Loaded* than in *Low_Neutral* also in the group of High-IQ subjects ($p = 0.101$).

7.2 Beliefs about relative IQ and skepticism

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 that potentially affect the way they interpret Senders' messages. We elicit Receivers' beliefs about their relative performance right after the IQ test. To do so, we insert each subject into a group of 99 subjects who did the same IQ test previously. We then divide 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 ask subjects to give an estimate of the likelihood they belong to each quintile.²⁶ Even if these beliefs about relative performance do not correspond exactly to prior beliefs about IQ-ranks, they give an indication of how well the Receiver thinks he performed in the IQ test.

For simplicity, we summarize Receivers' beliefs about relative IQ by computing each Receiver's *expected quintile*.²⁷ We split the group of subjects into two, depending on whether their evaluation of their relative performance is rather high or low. We say that a subject has *high priors* about his relative performance when his expected quintile is below the median (2.8). We say that a subject has *low priors* when his expected quintile is above the median. We count 104 subjects with high priors and 92 subjects with low priors. We now look at the skepticism of subjects in both groups in the *Low_Loaded* and *High_Loaded* treatments. In the *High_Loaded* treatment, the skeptical guess is self-serving and corresponds to the best IQ-rank disclosed. On average, subjects with high priors make guesses that are significantly more skeptical and reach the skeptical guess significantly more often than subjects with low priors ($p = 0.003$ and $p = 0.002$). The converse is true in the *Low_Loaded* treatment, in

²⁶The sum of the 5 estimates must be equal to 100%. We explain subjects that their payoff is highest if they estimate most accurately their chance to belong to each quintile, and they can see the exact method used to compute their payoff. At the end of the experiment, one quintile is selected at random. The subject's payoff is based on the following formula: $\text{Payoff} = 2 - 2 * [(I - p/100)]^2$, where I is an indicator variable that takes value 1 if the quintile the subject actually belongs to is equal to the selected quintile and 0 otherwise, and p is the subject's estimate in percent.

²⁷To get a subject's expected quintile, we multiply each quintile by the likelihood put on this quintile by the subject, and sum these values over the 5 quintiles. We get the same results if we take an alternative way to summarize a Receiver's beliefs, which is to consider only the quintile that the Receiver estimated as most likely.

which the skeptical guess corresponds to the worst IQ-rank disclosed. On average, subjects with high priors make guesses that are significantly less skeptical and reach the skeptical guess significantly less often than subjects with low priors ($p = 0.017$ and $p = 0.068$).

Observation 5. When skepticism is self-serving, subjects with high priors about relative performance are more skeptical than subjects with lower priors. When skepticism is self-threatening, subjects with low priors about relative performance are more skeptical than subjects with higher priors.

One possible explanation for this observation is that subjects with high priors have a stronger preference than subjects with low priors for thinking highly of themselves, which pushes their guesses to or away from the skeptical guesses. Another possible explanation is that subjects tend to stick to their prior view of their relative performance. A first remark about these explanations is that they are quite hard to disentangle in our experiment because the type that Receivers care about, i.e. the IQ rank, is also a type they have prior information about. We can make two other remarks about the second explanation. First, Receivers' prior beliefs about relative performance in the IQ test probably evolve as they get message about their IQ ranks. To properly look at the effect of prior beliefs on guesses, we would need to track the evolution of beliefs about relative performance over time. Second, if subjects' prior view of their performance affects their guesses, it should depend on how confident subjects are about this prior view. To measure how confident a subject is, we compute the standard deviation of his beliefs about relative performance.²⁸ We see no effect of this standard deviation on the correlation between priors and skepticism.

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

²⁸A subject who attributed 20% chances to belong to each of the five possible quintiles has the largest standard deviation and is considered not very confident. A subject who attributed 100% chances to belong one specific quintile has a null standard deviation and is considered confident.

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.

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. In general, one open question is whether it would be helpful to make Receivers aware of their tendency to read information as preferred. The strategic setting we consider raises another important question about motivated skepticism, namely whether Senders expect it and could therefore 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.²⁹

²⁹Theoretical papers, such as Milgrom and Roberts (1986), Hagenbach and Koessler (2017) or Deversi et al. (2021),

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 evidence.

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. and Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117(3):871–915.
- 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.
- 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.

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.

- 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 Sapiro, 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.
- 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.
- Hagmann, D. and Loewenstein, G. (2017). Persuasion with motivated beliefs.
- 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.
- 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.
- Mobius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. *Management Science*, forthcoming.
- 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*.
- Rosenberg, M. (2015). *Society and the adolescent self-image*. Princeton university press.
- 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. (2022a). The fake news effect: Experimentally identifying motivated reasoning using trust in news. *working paper*.
- Thaler, M. (2022b). The supply of motivated beliefs. *Working paper*.
- Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–61.

Appendices

A Complements to the theory

Proof of Proposition 1 - The unraveling result

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 β_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: β_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.

Proposition 2 - Optimal action for the receivers

Let the Receiver's payoff be $u_R(a, t) = -|a - t|$ and his beliefs be that t is distributed according to the cdf $F(\cdot)$. The action that maximizes the Receiver's expected payoff is a median of F .

Proof. t is a real-valued random variable. First remark that, for any $a \in \mathbb{R}$, we have the following relation:

$$|a - t| = \int_{-\infty}^a \mathbb{1}\{t \leq x\} dx + \int_a^{+\infty} \mathbb{1}\{t \geq x\} dx$$

$F(x) = P(t \leq x)$ is a non decreasing function and is therefore differentiable almost everywhere. Moreover, remark that:

$$E[|a - t|] = \int_{-\infty}^a \underbrace{P(t \leq x)}_{F(x)} dx + \int_a^{+\infty} \underbrace{P(t \geq x)}_{1-F(x)} dx$$

Therefore, the function $u_R(a) = -E[|a - t|]$ is also differentiable almost everywhere and $u'_R(a) = 1 - 2F(a)$ wherever it exists. Finally, remark that the function u is concave: for every $a \in \mathbb{R}$, $u''_R(a) = -2F'(a) \leq 0$. Therefore, the Receiver's optimal action a^* is determined but the FOC, $F(a^*) = \frac{1}{2}$, which implies that a^* is a median of F .

Receivers' payoffs in our experiment

Experiments on disclosure games, including ours, usually use Receiver's payoff functions which have two important properties: the payoff is maximal when the guess matches the true type, and the payoff decreases with the distance between the true type and the guessed type. In most existing experiments on disclosure however, Receiver's payoff functions are more complicated than the function we consider because, instead of presenting the payoff function to the Receiver, the experimenters show the Receivers tables in which they see their (integer) payoffs for every possible pair (a, t) . To present such tables – as is done for instance in Deversi et al. (2021), Jin et al. (2021) and Schipper and Li (2020) –, one needs to restrict the set of guesses available to the Receiver. Instead, we decided to give freedom to the Receivers in making their guesses and choose a relatively simple formula for the payoff: in every game of every treatment, subjects make 5 euros minus the difference (in absolute value) between the guess and the true type. We gave a few examples to Receivers of how a true type and a guess translate into a payoff.

The payoff $u_R(a; t) = 5 - |a - t|$ does not satisfy the property of strict concavity in a assumed in the theory section. As established in Proposition 2, the Receiver's optimal action for a given beliefs corresponds to the median of these beliefs. It follows that the following may be true: when several types are believed by the Receiver with positive probability, the optimal action is not necessarily strictly in between the lowest and highest types believed. Proposition 1 can be modified to apply to the Receivers' payoff we use in the experiment, and would become: The baseline game has a fully-revealing equilibrium. In every fully-revealing equilibrium, the Receiver's belief after every on-path message is skeptical and the Receiver's belief after every off-path message induces a skeptical action. If we consider the perturbation mentioned in Remark 2 of the theory section, namely that there is always a probability that the Receiver is unsophisticated and guesses any of the types disclosed, then we can show that every equilibrium is fully-revealing and that every message sent on or off the equilibrium path is followed by a skeptical action from the sophisticated Receiver.

B Regressions Including Inconsistent Individuals

In this subsection, we present the regressions of Tables 4 and Table E.1 for the data set that includes the 16 individuals that made more inconsistent than consistent guesses.

Table B.1: Determinants of skepticism, including inconsistent players

<i>Dep. Var.</i>	Skepticism								
	<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>	0.046 (0.050)	0.018 (0.044)	0.016 (0.048)	-0.118*** (0.037)	-0.136*** (0.039)	-0.165*** (0.054)	0.046 (0.047)	0.020 (0.046)	0.017 (0.049)
1 if <i>Low</i>							-0.055 (0.038)	-0.059* (0.032)	-0.044 (0.037)
1 if <i>Low_Loaded</i>							-0.165*** (0.062)	-0.158*** (0.059)	-0.166*** (0.063)
IQ perf.		0.026*** (0.004)	0.026*** (0.007)		0.022*** (0.007)	0.020** (0.008)		0.024*** (0.004)	0.0224*** (0.005)
Rounds dummies		✓	✓		✓	✓		✓	✓
<i>Demo.</i>			✓			✓			✓
Cons.	0.642*** (0.021)	0.348*** (0.056)	0.458*** (0.150)	0.586*** (0.030)	0.365*** (0.058)	0.124 (0.228)	0.642*** (0.023)	0.385*** (0.049)	0.426*** (0.141)
<i>N</i>	803	803	803	832	832	832	1635	1635	1635

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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.2: Determinants of a skeptical guess, including inconsistent players

<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>	0.057 (0.055)	0.032 (0.054)	0.022 (0.057)	-0.132*** (0.049)	-0.151*** (0.036)	-0.176** (0.073)	0.056 (0.053)	0.035 (0.048)	0.027 (0.047)
1 if <i>Low</i>							-0.037 (0.044)	-0.039 (0.037)	-0.027 (0.044)
1 if <i>Low_Loaded</i>							-0.194*** (0.075)	-0.188*** (0.065)	-0.189*** (0.071)
Mess. size	-0.158*** (0.019)	-0.159*** (0.018)	-0.161*** (0.017)	-0.092*** (0.016)	-0.093*** (0.013)	-0.0925*** (0.016)	-0.126*** (0.013)	-0.127*** (0.013)	-0.127*** (0.013)
IQ performance		0.023** (0.011)	0.023** (0.011)		0.025*** (0.007)	0.022** (0.010)		0.021*** (0.006)	0.019*** (0.006)
Rounds dummies		✓	✓		✓	✓		✓	✓
<i>Demo.</i>			✓			✓			✓
Cons.	0.892*** (0.065)	0.567*** (0.106)	0.668*** (0.178)	0.630*** (0.088)	0.345*** (0.093)	0.0138 (0.325)	0.784*** (0.047)	0.509*** (0.069)	0.561*** (0.185)
<i>N</i>	803	803	803	832	832	832	1635	1635	1635

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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C 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 C 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 C.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

D Additional Tables and Figures

Table D.1: Sender’s communication strategy in *High* treatments

Treatment <i>High_Neutral</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	4.88	-	-	-	-	1.63	-	-	-	9.76	-	-	16.26	-	67.48	100
2	-	3.42	-	-	-	-	4.27	-	-	0.85	16.24	-	3.42	53.85	17.95	100
3	-	-	7.87	-	-	-	2.25	15.73	-	-	-	61.80	2.25	4.49	5.62	100
4	-	-	-	15.79	-	-	-	-	81.05	-	-	2.11	1.05	-	-	100
5	-	-	-	-	87.93	-	-	-	5.17	-	-	2.59	-	2.59	1.72	100
Total	1.11	0.74	1.30	2.78	18.89	0.37	1.30	2.59	15.37	2.41	3.52	11.11	5.00	12.96	20.56	100

Treatment <i>High_Loaded</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	5.66	-	-	-	-	3.77	-	-	-	8.49	-	-	14.15	-	67.92	100
2	-	8.93	-	-	-	-	3.57	-	-	0.89	19.64	-	0.89	55.36	10.71	100
3	-	-	6.25	-	-	-	0.89	8.93	-	0.89	8.93	67.86	0.89	0.89	4.46	100
4	-	-	-	17.95	-	-	-	2.56	62.82	-	2.56	10.26	1.28	-	2.56	100
5	-	-	-	-	73.61	-	-	-	15.28	-	-	6.94	-	1.39	2.78	100
Total	1.25	2.08	1.46	2.92	11.04	0.83	1.04	2.50	12.50	2.29	7.08	18.54	3.75	13.33	19.38	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 $t = 5$ in the *High_Neutral* treatment send the precise message $m = \{5\}$ 87.93% of the time. It is 73.61% in *High_Loaded*.

Table D.2: Sender’s communication strategy in *Low* treatments

Treatment <i>Low_Neutral</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	76.60	-	-	-	-	12.77	-	-	-	4.26	-	-	1.06	-	5.32	100
2	-	1.28	-	-	-	71.79	1.28	-	-	10.26	-	-	2.56	7.69	5.13	100
3	-	-	2.56	-	-	-	7.69	-	-	76.92	5.13	1.28	5.13	1.28	-	100
4	-	-	-	0.74	-	-	-	1.47	-	-	11.76	2.21	71.32	1.47	11.03	100
5	-	-	-	-	0.88	-	-	-	1.75	-	-	5.26	-	11.40	80.70	100
Total	14.40	0.20	0.40	0.20	0.20	13.60	1.40	0.40	0.40	14.40	4.00	2.00	20.80	4.40	23.20	100

Treatment <i>Low_Loaded</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	69.31	-	-	-	-	8.91	-	-	-	7.92	-	-	5.94	-	7.92	100
2	-	12.09	-	-	-	70.33	2.20	-	-	5.49	1.10	-	2.20	3.30	3.30	100
3	-	-	2.91	-	-	-	3.88	1.94	-	74.76	1.94	2.91	0.97	5.83	4.85	100
4	-	-	-	1.83	-	-	-	5.50	0.92	-	20.18	0.92	51.38	9.17	10.09	100
5	-	-	-	-	-	-	-	-	6.58	-	-	7.89	-	19.74	65.79	100
Total	14.58	2.29	0.62	0.42	0.00	15.21	1.25	1.67	1.25	18.75	5.21	2.08	13.54	7.08	16.04	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 $t = 1$ in the *Low_Neutral* treatment send the precise message $m = \{1\}$ 76.60% of the time. It is 69.31% in *Low_Loaded*.

Table D.3: Distribution of type in each treatment (in %)

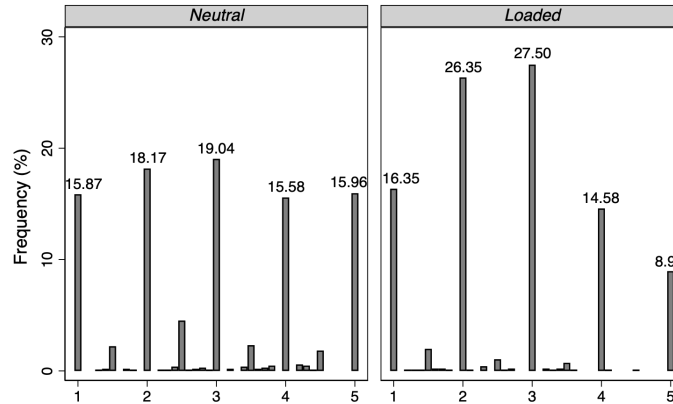
Type	<i>High_Neutral</i>	<i>High_Loaded</i>	<i>Low_Neutral</i>	<i>Low_Loaded</i>
1	22.78	22.08	18.80	21.04
2	21.67	23.33	15.60	18.96
3	16.48	23.33	15.60	21.46
4	17.59	16.25	27.20	22.71
5	21.48	15.00	22.80	15.83
Total	100	100	100	100

Note: The Table reports the distribution of type in each treatment. For instance, in the *High_Neutral* treatment, 28.78% of Senders' are of type $t = 1$.

Table D.4: Reasoning steps

Message	<i>High treatments</i>					Message	<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]	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 5: Distribution of Receivers' Guesses

E Alternative measure of skepticism

Table E.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.

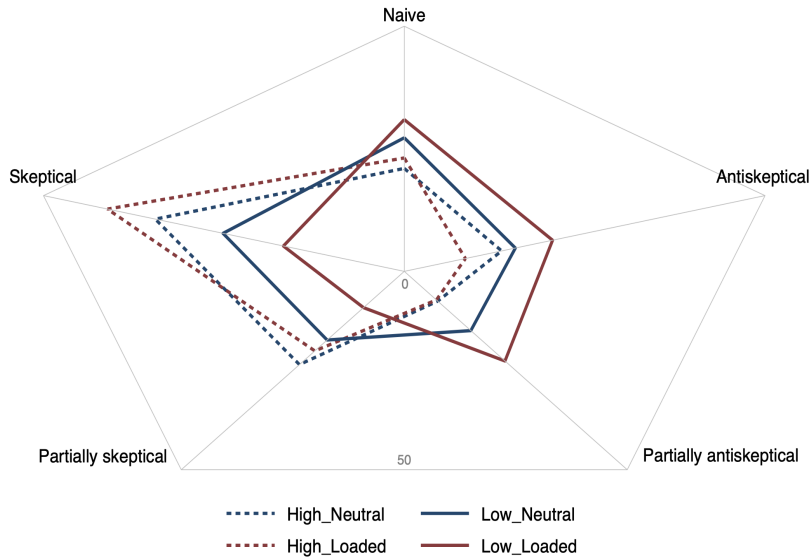
Table E.1: Determinants of a skeptical guess

<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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F Classifying guesses

We can classify Receivers' guesses into five categories: when $Sk(a, m) = 1$, the guess a after m is *skeptical*; when $Sk(a, m) \in]0.5, 1[$, a is *partially skeptical*; when $Sk(a, m) = 0.5$, a is *naive*; when $Sk(a, m) \in]0, 0.5[$, a is *partially anti-skeptical*; when $Sk(a, m) = 0$, a is *anti-skeptical*. Figure 6 represents, for each treatment, the distribution of the guesses into these 5 categories. Every axis corresponds to a category and each point on an axis is the frequency (in %) of guesses in this category (see Table F.1 below for the exact frequencies).



Note: The Figure displays, for each treatment, the distribution of the guesses based on our five categories: a guess is either Skeptical ($Sk(a, m) = 1$), Partially skeptical ($Sk(a, m) \in]0.5, 1[$), Naive ($Sk(a, m) = 0.5$), Partially anti-skeptical ($Sk(a, m) \in]0, 0.5[$) or Anti-skeptical ($Sk(a, m) = 0$).

Figure 6: Distribution of Receivers' guesses, by category and treatment

The graph confirms Result 2. First, the pentagons for *High_Neutral* and *High_Loaded* look very similar. In fact, the percentage of guesses in each category is not different between these two treatments (see Table F.1 below). In contrast, the pentagon for *Low_Loaded* is more shifted to the right than the pentagon for *Low_Neutral*, meaning more anti-skeptical and partially anti-skeptical guesses for the former treatment. Next, we see that a substantial fraction – between 21.00% and 30.96% – of guesses corresponds exactly to the rank in the middle of the disclosed set. We observe that the percentage of naive guesses is marginally higher in the *Low* than in the *High* treatments ($p = 0.083$). This suggests again that it may be slightly more complex for subjects to play *Low* than *High* treatments.

Table F.1: Percentage of guess by category and treatment

	<i>Skeptical</i>	<i>Partially Skeptical</i>	<i>Naive</i>	<i>Partially Anti-skeptical</i>	<i>Anti-Skeptical</i>	<i>Total</i>
<i>High_Neutral</i>	34.50	23.50	21.00	7.50	13.50	100
<i>High_Loaded</i>	41.13	20.05	23.14	7.20	8.48	100
<i>p-value</i>	<i>0.300</i>	<i>0.421</i>	<i>0.713</i>	<i>0.977</i>	<i>0.210</i>	
<i>Low_Neutral</i>	25.12	17.3	27.25	14.93	15.40	100
<i>Low_Loaded</i>	16.75	9.14	30.96	22.59	20.56	100
<i>p-value</i>	<i>0.043</i>	<i>0.131</i>	<i>0.438</i>	<i>0.103</i>	<i>0.282</i>	
<i>p-value (Neutral)</i>	<i>0.125</i>	<i>0.192</i>	<i>0.347</i>	<i>0.001</i>	<i>0.855</i>	
<i>Total</i>	29.28	17.51	25.61	13.08	14.52	100

Note: The Table reports the percentage of guess by category and treatment. For instance, 34.50% of guesses in *High_Neutral* are classified as *skeptical*. 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.

G Monetary losses due to lack of skepticism

The lack of skepticism therefore induces monetary losses that we now try to evaluate. To do that, we can simply look at the difference, in euros, between the payoff the Receiver would have got when guessing right and the payoff he effectively got. This corresponds simply to the distance between the Receiver’s guess and the true type. If we compute the average of this loss per treatment, restricting ourselves to vague messages, we obtain the following: 0.94 euro in *High_Neutral*, 0.73 euro in *High_Loaded*, 1.20 euro in *Low_Neutral* and 1.17 euro in *Low_Loaded*. Losses are higher in *Low* than in *High* treatments ($p < 0.001$), and not significantly different between the *Neutral* and *Loaded* treatments ($p = 0.202$). However, the limit of this measure is that losses may be due both to Receivers’ being not skeptical enough and to Senders’ mistakes in the communication strategy.

To focus on Receivers, we construct another measure of monetary loss that considers the distance between the Receiver’s guess and the skeptical guess he should have made *given* the message sent by the Sender. If we compute the average of this loss per treatment, restricting ourselves to vague messages and consistent guesses, we obtain the losses presented in the Table G, both at the aggregate level and conditionally on the size of the message received. Given any message size, the loss is always larger in the *Low_Loaded* treatment than in any other treatment ($p < 0.001$ for *High_Loaded*, $p = 0.002$ for *High_Neutral* and $p = 0.126$ for *Low_Neutral*). When the message is of size 4 and 5, the differences between the average loss in the *Low_Neutral* and *Low_Loaded* is strongly significant ($p < 0.001$). When the message is of size 3, it is significant at the 10% level ($p = 0.076$).

Table G.1: Receivers’ Payoff Loss

	Maximal loss if consistent guess	Loss			
		<i>High_Neu</i>	<i>High_Loa</i>	<i>Low_Neu</i>	<i>Low_Loa</i>
All		0.95€	0.79€	1.19€	1.35€
Mess. size 2	1€	0.33€	0.25€	0.39€	0.52€
Mess. size 3	2€	0.69€	0.63€	0.87€	1.04€
Mess. size 4	3€	1.04€	0.97€	1.48€	1.78€
Mess. size 5	4€	1.66€	1.33€	1.72€	2.29€

Note: The Table reports the average Receivers’ losses, in Euros, by treatment and message size. For instance, in the *High_Neutral* treatment, Receivers’ lose on average 0.95 by lack of skepticism. It is 0.79 in *High_Loaded*.

H Instructions – Online Appendix

Subjects do not see what appears in italic. Instructions are in English and given along the way on the computer screens. The instructions for Part 1 are common to all treatments. The instructions for Part 2 vary depending on the treatment and the role of the subject. Overall, there are two sets of instructions for Senders – High and Low – and four sets for Receivers, one for each treatment. For the sake of brevity, we only report below the instructions corresponding to the High_Loaded treatment.

Welcome!

If you have any question during the experiment, please use the zoom chat.

The experiment has 2 parts.

- ◇ Part 1: IQ-test,
- ◇ Part 2: 10 rounds of a game

Instructions are given along the way. It is in your interest to read them carefully! At the end of the experiment, you will receive the earnings made in the 2 parts plus a 5 euros show-up fee. You will additionally receive a participant's fee of 5 euros for today's online experiment. You will be paid on your PayPal account within 48h after the end of the experiment.

Part 1

Part 1 consists of a Raven IQ-test, a test frequently used to measure intelligence. It measures the ability to reason clearly and grasp complexity. Performance in the test is often associated with educational success and high future income.

The test has 15 questions. For every question, you will see a pattern with a missing piece. Your task is to complete the pattern by choosing one of the pieces that are proposed to you. You will have 15 minutes to answer all the questions.

Payoff: you will earn 0.50 cents for each correct answer. In this part, you can earn up to 7.5 euros ($15 * 0.50$).

Subjects then see an example of a Raven matrix and are given the correct answer. Then they move on to the 15 questions, one per screen.

(Future receivers only) The same IQ test was also done by a large number of participants who previously came at WZB-TU lab. We have divided the group into five quintiles, from the 20% of participants who had the best performance in the IQ-test, to the 20% of participants who had the worst performance in the IQ-test. What do you think is the likelihood that you belong to each quintile? On the next screen, you must state an estimate for each of the 5 quintiles. The sum of the 5 estimates must equal 100%.

Payoff: You will be paid for this short estimation task. Your payment will be highest if you estimate your chances to belong to each quintile as accurately as possible. The maximal payment for the estimation task is 2 euros, negative payment is not possible. If you are interested, here is the detailed information about the payoff: One quintile will be picked at

random by the computer. You will be paid according to the following formula: Your payoff (euros) = $2 - 2 * (I - p/100)^2$ where I is an indicator variable that takes value 1 if the quintile you actually belong to is equal to the quintile drawn at random and 0 otherwise, and p is your estimate in percent.

Receiver-subjects fill in a table with their 5 estimates.

Part 2

Part 2 consists of 10 rounds of a 2-player game. The computer has randomly assigned you the role of Sender or the role of Receiver. You will learn your role on the next screen and keep it for the 10 rounds. In each round, if you are a Sender, you will be randomly matched with a Receiver, and vice-versa. You will never know the identity of the other player and this player can be new in each round.

Subjects then see either “Today, you are a Receiver” or “Today, you are a Sender”.

**** Instructions for SENDERS ****

Description of the game: Each round of the game has 4 steps.

- ◇ **Step 1:** In each round, the computer program will generate a secret number, which is 1, 2, 3, 4, or 5.
- ◇ **Step 2:** You will be informed of the secret number. The Receiver will not know this secret number, but his/her task is to guess it.
- ◇ **Step 3:** Before the Receiver guesses the secret number, you can send him/her information about it. This information will take the form of a set of numbers, with the only constraint that the secret number must be part of the set. For instance, if the secret number is 3, then you can send any of the sets of numbers given in the table below:

Available sets of numbers	
{1, 2, 3, 4, 5}	<input type="checkbox"/>
{1, 2, 3, 4}	<input type="checkbox"/>
{2, 3, 4, 5}	<input type="checkbox"/>
{1, 2, 3}	<input type="checkbox"/>
{2, 3, 4}	<input type="checkbox"/>
{3, 4, 5}	<input type="checkbox"/>
{2, 3}	<input type="checkbox"/>
{3, 4}	<input type="checkbox"/>
{3}	<input type="checkbox"/>

- ◇ **Step 4:** The information you will give to the Receiver will be displayed on his/her screen, and the Receiver will finally make his/her guess. His/her guess can be any number between 1 and 5. We allow for guesses with one digit with the idea that a Receiver who, for instance, thinks 1 and 2 are equally likely can guess 1.5, or that a Receiver who, for instance, thinks it is either 4 or 5 but most likely 5 can make a guess between 4 and 5 but closer to 5. Once the Receiver made his/her guess, the round is over and a new round starts.

In each round, the payoffs are as follows. The Receiver knows these payoffs too.

Your Payoff: Your payoff corresponds exactly to the guess of the Receiver in this round.

$$\text{Your payoff} = \text{guess of the Receiver.}$$

Simply put, you earn more when the Receiver guesses a higher secret number.

Receiver's payoff: The Receiver's payoff depends on how close is his/her guess to the secret number.

$$\text{Receiver's payoff} = 5 - |\text{guess} - \text{secret number}|$$

where $|\text{guess} - \text{secret number}|$ is the distance between the guess and the secret number in the round. For instance, if the Receiver correctly guesses the secret number, he/she gets 5 euros. If the Receiver guesses 4 while the secret number is 3, he/she gets 4 euros ($5 - 1$). Simply put, the Receiver earns more when his/her guess is closer to the secret number.

Summary: To sum up, each round of the game goes as follows:

1. The computer generates a secret number.
2. You are informed about this secret number.
3. You can give information to the Receiver about it.
4. The Receiver receives this information and guesses the secret number.

You earn more when the Receiver guesses a higher number. The Receiver earns more when his/her guess is closer to the true secret number. Part 2 of the experiment ends after 10 rounds of the game. One of the 10 rounds will be randomly selected at the end of the experiment for effective payment in this part.

After some comprehension questions, the game starts. In every round, the Sender is shown the secret number and has to choose the message about this number that he/she wants to send to the receiver.

**** Instructions for RECEIVERS ****

Description of the game: Each round of the game has 4 steps.

- ◇ **Step 1:** As you know, the same IQ-test that you did earlier was also done by a large number of previous participants. In each round of the game, the computer will randomly select 4 previous participants. Together with these 4 participants, you will form a group of 5 participants. Within this group, the computer program will compare the performances in the IQ-test. It will then compute your IQ-rank for the round as follows:
 - ◇ If you have the highest perf. in the group of 5, your IQ-rank will be 1.
 - ◇ If you have the second highest perf. in the group of 5, your IQ-rank will be 2.

- ◊ If you have the third highest perf. in the group of 5, your IQ-rank will be 3.
- ◊ If you have the fourth highest perf. in the group of 5, your IQ-rank will be 4.
- ◊ If you have the lowest perf. in the group of 5, your IQ-rank will be 5.
- ◊ If you have the same perf. as other participants in the group, the computer program randomly decides the ranking between these participants and yourself.

In each round, your IQ-rank will be 1, 2, 3, 4 or 5. The higher your IQ-rank, the better you performed in the IQ-test relative to the 4 randomly selected participants. A lower IQ-rank also means that you made more mistakes in the IQ-test than these 4 participants.

Note: In each round, the computer randomly selects 4 new previous participants whose performance is compared to yours, so your IQ-rank can change across rounds.

- ◊ **Step 2:** The Sender will be informed of a number, which is 1, 2, 3, 4 or 5. This number corresponds to your IQ-rank, but the Sender does not know that this is the case. For him/her, this number has no particular meaning. You will not be informed of this IQ-rank, but your task is to guess it.
- ◊ **Step 3:** Before you guess your IQ-rank, the Sender will give you information about it. This information will take the form of a set of numbers, with the only constraint that your IQ-rank must be part of the set. Said differently, your IQ-rank is always one of the numbers sent by the Sender. For instance, if your IQ-rank is 3, then the Sender can send you any of the sets of numbers given in the table below:

Available sets of numbers	
{1, 2, 3, 4, 5}	<input type="checkbox"/>
{1, 2, 3, 4}	<input type="checkbox"/>
{2, 3, 4, 5}	<input type="checkbox"/>
{1, 2, 3}	<input type="checkbox"/>
{2, 3, 4}	<input type="checkbox"/>
{3, 4, 5}	<input type="checkbox"/>
{2, 3}	<input type="checkbox"/>
{3, 4}	<input type="checkbox"/>
{3}	<input type="checkbox"/>

- ◊ **Step 4:** The information given by the Sender will be displayed on your screen, and you will finally make your guess. Your guess can be any number between 1 and 5. We allow for guesses with one digit with the idea that a Receiver who, for instance, thinks 1 and 2 are equally likely can guess 1.5, or that a Receiver who, for instance, thinks it is either 4 or 5 but most likely 5 can make a guess between 4 and 5 but closer to 5. Once you made your guess, the round is over and a new round starts.

In each round, the payoffs are as follows. The Sender knows these payoffs too.

Sender's Payoff: The Sender's payoff corresponds exactly to your guess in this round.

$$\text{Sender's payoff} = \text{your guess.}$$

For instance, if you guess 5 (the lowest IQ-rank), the Sender gets 5 euros. If you guess 1.4, the Sender gets 1.40 euros. Simply put, the Sender earns more when you guess a lower IQ-rank (closer to 5).

Your Payoff: Your payoff depends on how close is your guess to your IQ-rank.

$$\text{Your payoff} = 5 - |\text{your guess} - \text{your IQ-rank}|$$

where $|\text{your guess} - \text{your IQ-rank}|$ is the distance between your guess and your IQ-rank in the round. For instance, if you correctly guess your IQ rank, you get 5 euros. If you guess 4 while your IQ-rank is 3, you get 4 euros ($5 - 1$). Simply put, you earn more when your guess is closer to your IQ-rank.

Summary: To sum up, each round of the game goes as follows:

1. Your IQ-rank is computed.
2. The Sender is informed about a number that corresponds to your IQ-rank. For him/her, this number has no meaning.
3. The Sender gives you information about this number.
4. You receive this information and guess your IQ-rank.

You earn more when your guess is closer to your IQ-rank. The Sender earns more when you guess a lower IQ-rank (closer to 5). Part 2 of the experiment ends after 10 rounds of the game. One of the 10 rounds will be randomly selected at the end of the experiment for effective payment in this part.

No feedback: In the experiment, you will never receive more information about your IQ-ranks than the information given by the Senders.

After some comprehension questions, the 10 repetitions of the game start. In every round, the Receiver is told that his/her IQ-rank has been computed, shown the information given by the Sender and asked to guess his/her IQ-rank.

At the end of the 10 rounds, all subjects go through a self-esteem questionnaire, and answer some questions about themselves (gender, age, education etc.). They learn their aggregate payoff and the experiment ends.

I Screens – Online Appendix

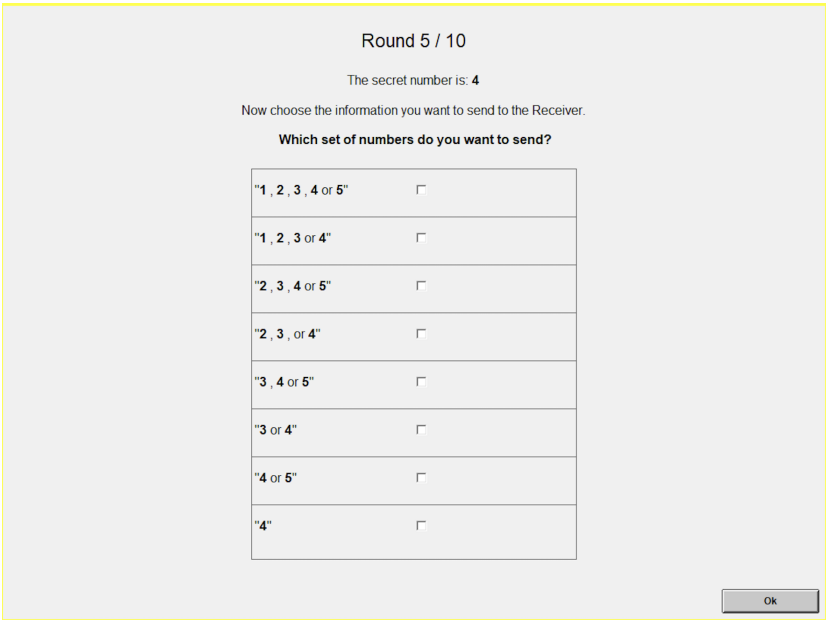


Figure 7: Example of a screen seen by a Sender of type 4

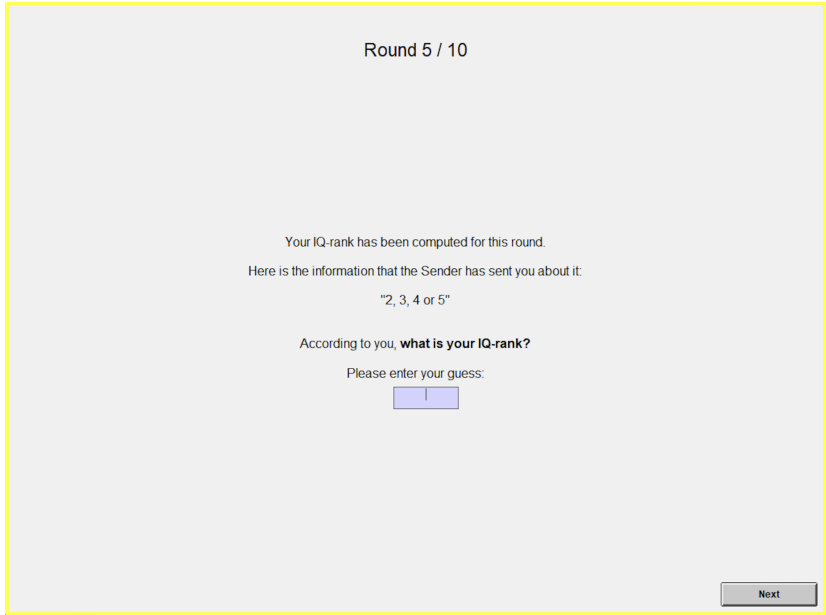


Figure 8: Example of a screen seen by a Receiver who received message {2, 3, 4, 5}