

Belief-Based Utility and Signal Interpretation*

Marta Kozakiewicz[†]

August 8, 2024[‡]

Abstract

Do people update their beliefs differently after positive versus negative feedback? The existing literature disagrees on the magnitude and direction of the bias. In this paper, I propose a new experiment guided by a simple model of belief choice. The experimental data reveal a strong asymmetry in updating after “good” versus “bad” news. Moreover, I design a control condition that allows a clear identification of belief *manipulation* and provides robust evidence on the underlying mechanism. The results point towards the role of affect (or utility from beliefs shifted by the signal) in asymmetric updating. The proposed method can be applied more broadly to study belief-based utility and its role in belief formation.

Keywords: overconfidence, belief formation, learning, experiment

JEL classification: C91, D83

*The author gratefully acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (project A01). Financial support from the German Research Foundation (DFG project 462020252) is gratefully acknowledged.

[†]Frankfurt School of Finance and Management; email: m.kozakiewicz@fs.de

[‡]The first version of the paper was presented in the academic job market 2021/2022.

1 Introduction

People tend to overestimate their abilities, making costly mistakes as they hold on to their biased beliefs at the expense of accuracy. At the same time, individuals often operate in environments characterized by repeated feedback (examples include academic or workplace settings). If biased beliefs are costly and there is ample opportunity to learn, how can overconfidence persist? One explanation advanced in the literature concerns *asymmetric updating*.¹ It proposes that people update beliefs differently after “good” versus “bad” news, responding more to positive information. Theoretical models involve an agent deriving direct utility from his beliefs so that he has an incentive to adopt overly optimistic beliefs (Brunnermeier and Parker, 2005; Caplin and Leahy, 2019). Although this explanation is consistent with observed overconfidence, the empirical evidence on asymmetric updating is not conclusive. The experimental data brought mixed results, showing optimistic/pessimistic updating or no effect (Benjamin, 2019).² Inconsistent findings can be attributed to methodological limitations, e.g., the absence of a control condition based on the same primitives (Barron, 2021; Drobner and Goerg, 2024), or ignoring the timing of the resolution of uncertainty (Drobner, 2022).

In this paper, I take the next step to fully understand the nature of the bias. Guided by a simple model of belief choice, I design an experiment that exposes the conditions necessary for the asymmetry to arise.³ The experimental data show strong asymmetry in updating after “good” news and elucidate the underlying mechanism. The results provide suggestive evidence of why some work failed to capture the differential response to signals and point toward a more comprehensive theory of belief formation.

The experiment has the following structure. In the treatment condition, participants solve an IQ test and receive a noisy signal about their relative performance. Then, they report a subjective probability that the received signal corresponds to their performance. I designed the task in a way that, according to the model, would enhance the chances of capturing the effect. Furthermore, I introduce a new control condition. Subjects in the control group solve the same IQ test and consider the same signal structure but report their beliefs about *hypothetical* signal realizations. By comparing the two conditions, I separate belief manipulation due to the utility from “good” or “bad” news from distortions caused by other factors. The collected data reveal a substantial asymmetry in the treatment condition: there is a large and highly significant difference in subjects’

¹Other explanations that consider motivated reasoning rather than cognitive processes fall into two categories: information avoidance (see Golman et al., 2017, for a review of the literature) and selective recall (Chew et al., 2020; Huffman et al., 2022; Zimmermann, 2020).

²I review the relevant literature at the end of the introduction.

³The experiment was pre-registered in AEA RCT Registry (Registration No. AEARCTR-0006233).

responses to “good” versus “bad” news. Participants tend to interpret positive signals as 10 pp more likely to be informative (18% increase in relative terms). The data show no asymmetry in the control condition, as predicted by the model. At the same time, there is a significant difference *across* the experimental conditions. It provides evidence of utility-driven belief manipulation that operates in the direction of the preferred state.

The paper makes several contributions. First, the upgraded design of the treatment condition allows me to capture the effect that, while suggested by the theory, was elusive for previous studies. Second, the new control condition offers a benchmark based on the same information over the same unknown state, solving the problem prevalent in the literature on asymmetric updating. Together with the model, the experiment allows a clear identification of belief manipulation and provides robust evidence on the mechanism. It is the first paper that directly tests a model of belief choice and establishes instantaneous utility as the source of asymmetry. It advances the research that, so far, provided only indirect evidence on the underlying process. The proposed method is not limited to the context of IQ but can be applied to study belief formation in other domains.

In Section 2, I describe the design in detail. Subjects first solved an IQ test and then reported a subjective probability of their score falling into the 1st, 2nd, ..., 10th decile of the test score distribution. Thus, I elicited a prior belief distribution over deciles, which I referred to as “ranks”. In the treatment condition, participants received a noisy signal about their performance. They were shown a number between 1 and 10 that could be equal to their rank. The framework was described as follows: *There are two boxes. Box 1 contains ten balls with numbers 1 to 10 written on them (each number occurs exactly once). Box 2 contains ten balls with the same number written on every one of them. That number is equal to your rank.* For example, if a subject’s rank was 4, Box 2 contained ten balls with the number “4” written on them. Subjects were informed that one ball would be randomly drawn from one of the boxes (either box can be selected with equal probability 0.5) and displayed on their computer screen. Their task was to tell us what they thought: Which box did the ball come from? Using an incentive-compatible mechanism, I elicited beliefs that the ball came from Box 2 (with the numbers equal to one’s rank).⁴ The design with two boxes is equivalent to the design used in the literature

⁴The main advantage of eliciting beliefs about the box instead of rank is that it minimizes confounds arising from people’s desire to be consistent (Falk and Zimmermann, 2017). By reframing the question, I avoid asking about the rank multiple times, whereas the composition of Box 2 ensures that I obtain the relevant probability. As a robustness check, I again elicited the entire distribution of beliefs at the end of the study. I confirm that beliefs about the box are consistent with the posterior distribution.

extended to 10 states of the world.⁵ I adopt a richer state and signal space to generate a stronger effect. This is based on the observation that it is more painful to learn that your score was among the worst 10% than to learn that it was below the median.

In the control condition, subjects did not receive a signal but reported their beliefs for every possible signal realization. The procedure, known as the Strategy Method, is commonly used in studies on strategic interactions (Brandts and Charness, 2009). The hypothetical control has an important advantage over alternatives used in the literature: it is based on the same subjective beliefs over the same unknown state. Previous work compared how people update beliefs about an ego-relevant outcome (e.g., one’s performance in an IQ test) and a neutral parameter (e.g., the performance of a robot).⁶ This comparison involves not only learning about different objects but also updating subjective beliefs, possibly multiple priors, and updating objective probabilities given by the experimenter. In contrast, the hypothetical control relies only on the assumption that when a signal is not realized, it does not induce the emotional reaction that distinguishes “good” and “bad” news from neutral information. The present findings are consistent with this assumption.

I took several steps to alleviate concerns about the non-comparability of the two conditions. For example, to minimize differences in the understanding of the task, I also required the treatment group to consider, one by one, every possible number. Participants in the control condition evaluated each number separately, using the same interface as subjects in the treatment condition. The numbers were presented in random order.

In Section 3, I describe a static model of belief choice in the spirit of Brunnermeier and Parker (2005). In the model, an agent forms beliefs about his unknown ability. He starts with a prior and receives a noisy signal with known precision. Then, he chooses a posterior belief facing a trade-off between the utility from the new belief and the costs of belief manipulation. The costs are increasing in the distance between the chosen belief and an “unmanipulated” posterior.⁷ These costs and benefits are all the agent cares for, as the uncertainty is not resolved at the end of the first period (to incorporate

⁵One cannot use the same design as in the literature, because the signal structure becomes too complicated when extended beyond the binary case (see Appendix E). The two-box design introduces 10 states in a way that is easy to explain to participants and allows for a simple elicitation of conditional beliefs.

⁶See, for instance, Coutts (2019), Eil and Rao (2011), and Möbius et al. (2022). A notable exception is a contemporaneous study by Drobner and Goerg (2024) who introduce an exogenous variation in subjects’ perception of the IQ test validity. Their results are consistent with my findings, however, their work differs in terms of methods and the conceptual framework (e.g., belief-based utility is defined over the results of a valid or invalid IQ test, whereas I define utility over beliefs about one’s intelligence). The latter has important implications for the interpretation of experimental manipulation and the results.

⁷This assumption, commonly used in theoretical work, implies that belief formation is partly driven by a rational process. One can view it as a modeling technique, conceptually similar to a dual-self model (Fudenberg and Levine, 2006), describing an internal, subconscious process of coming to a belief.

this feature, participants were informed that their test results will be available to them only one week after the session). Importantly, the utility is derived from beliefs the agent holds *at the moment*. When deciding about a hypothetical signal, the utility stems from the prior that is unaffected by speculations. In this case, there is no incentive for belief manipulation, so the agent reveals the unmanipulated posterior.

In Section 4, I present the results. The main outcome variable is the posterior belief revealed through the decision about the box. I test for asymmetry by comparing beliefs in the treatment condition with 1). the Bayesian benchmark, and 2). the decisions in the control condition. Both serve as a proxy for the unmanipulated posterior. The Bayesian benchmark is calculated based on beliefs elicited in the first part of the study.⁸

In the treatment condition, subjects reported an average probability of 38.5 pp that the signal came from Box 2 (56.1 pp for signals with non-zero prior probability). The average masks substantial heterogeneity in responses to signals. The reports are 10 pp higher after signals “1”, “2”, or “3” (the best signals) compared to the reports after the remaining signals (p-value of one-tailed t-test = 0.005). The result is robust to changes in the definition of a “good” signal and controlling for observables. Importantly, it is not driven by selection—the coefficient is nearly identical if I estimate the effect on a sub-sample of subjects who received a signal from the box with random numbers.

There is no similar relation in the control condition, in line with the model prediction. On average, subjects in the control condition reported a conditional probability of 30 pp. (50.2 pp for signals with non-zero prior probability).⁹ In contrast to the treatment, the reports in the control condition do *not* depend on the signal value. Moreover, the estimated weight placed on the Bayesian benchmark is no different than in the treatment condition. Subjects seem to assess signals in the same way but without the overreaction to positive information. This result provides suggestive evidence that the treatment effect is not due to the different structure of the two conditions (hypothetical vs not) but stems from the belief-based utility. Finally, the “good” news effect is present across the experimental conditions. The difference-in-difference analysis shows that people tend to report 10.1 pp higher probability after receiving “good” news.

⁸I use the prior probability placed on the rank indicated by the signal. As the Bayes’ rule is undefined for $p_0 = 0$, the baseline sample is restricted to subjects who assigned a positive probability to the relevant rank. In a second sample, I include subjects who assessed signals that were close to their prior belief distributions. The decisions after unexpected signals (signals two or more ranks away from one’s prior belief distribution) are analyzed separately and discussed in Section 4.5.1.

⁹One might worry that, with many decisions to make, subjects might gravitate toward the default option of 50 pp. Overall, the share of “default” choices is below 5% in the control condition and even lower (1.3%) in the treatment. The results are robust to excluding these subjects from analysis. Moreover, I show that time effects in the control condition are not driving the results. The results are similar if based on the first five choices made in the control condition or the last five choices instead of all ten.

My work draws on the literature on motivated reasoning. This literature examines various ways in which people distort their beliefs to achieve certain goals (see Bénabou and Tirole, 2016, for a comprehensive review). One strand of this literature studies belief formation when an agent derives utility from his beliefs. The theory predicts that, in this case, the agent will deviate from rationality adopting optimistic beliefs (Brunnermeier and Parker, 2005; Caplin and Leahy, 2019). The empirical evidence is less conclusive—the results from lab experiments differ in the magnitude and the direction of the effect. Some authors document asymmetry in the direction of the preferred state (Drobner, 2022; Drobner and Goerg, 2024; Eil and Rao, 2011; Möbius et al., 2022), others found no asymmetry (Buser et al., 2018; Schwardmann and Van der Weele, 2019; Zimmermann, 2020), or asymmetry in the opposite direction (Coutts, 2019; Ertac, 2011). As the designs differ along several dimensions (see Appendix E), the findings are difficult to reconcile.¹⁰

This paper makes several methodological and conceptual contributions. First, I develop a control condition that enables a causal identification of the effect of “good” news. More importantly, the new control allows testing a model of belief choice, opening the way for studying the utility from beliefs—its functional form and properties in different contexts and across different samples. Such evidence would inform the theory and advance the development of better models of learning with belief-based utility. Moreover, the tools developed in this paper can be applied to other domains. As long as people derive utility from their beliefs, their learning will be guided by similar principles as learning about cognitive ability. Possible applications include the formation of beliefs on politically contentious issues such as climate change or vaccination, learning about others (their trustworthiness or cooperativeness) in social interactions, or updating beliefs about social norms when confronted with changes in socially acceptable behavior.

Second, the results provide a unique insight into the nature of the bias. As asymmetry arises only after the signal, the results point towards the role of affect (the experience of feeling or emotion, formalized as utility from beliefs) in asymmetric updating. This is a noteworthy finding, as it suggests who might be more prone to biased belief formation—special attention should be paid to people experiencing strong emotions and using particular strategies to regulate them.¹¹

¹⁰A recent study by Drobner (2022) shows that the differences in findings can be driven by the differences in the expected timing of the resolution of uncertainty, as people expecting to learn the state sooner might face incentives to form more accurate beliefs. My paper complements this work by providing direct evidence of the role of instantaneous utility from beliefs in asymmetric updating.

¹¹I collected additional data on self-reported emotions experienced before the task and the habitual use of emotion regulation strategies. The data reveal correlations between subjects’ responses to the questionnaires and deviations from the Bayesian benchmark. Since the hypotheses were not pre-registered, the result should be only viewed as a suggestion for future research. I report them in Appendix G.

The final methodological contribution concerns designing a treatment condition that extends the binary state and signal space without introducing a complicated signal structure. It allows the experimenter to generate a quantitatively large shift in beliefs. On account of that, my paper hints at why some of the previous work did not capture the effect. If the increase in utility is not large enough, due to a coarse signal structure or a particular functional form, the manipulation might be hard to detect in the data.¹²

2 Experimental Design

The experiment consisted of two parts. In the first part, subjects completed an IQ test. The second part included the elicitation of prior and posterior beliefs and a stage in which subjects received signals (or considered every possible signal realization in the control condition).¹³ The outline of the experiment is presented in Figure 1. I describe the procedures in detail in the following subsections.

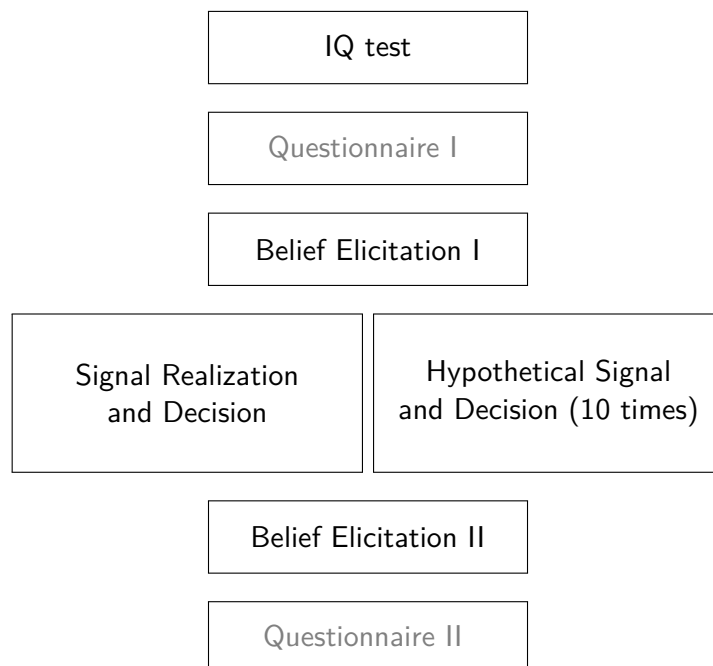


Figure 1: The outline of the experiment.

¹²For example, the experiments studying updating in the financial domain found little asymmetry (see Barron, 2021). The utility from a “good” signal about a small monetary gain is likely to be lower than the utility from a “good” signal about one’s IQ, which is a known predictor of *all* future earnings.

¹³In my study, subjects received only one signal. The previous work used to introduce more than one signal and report the average effect (e.g., Eil and Rao, 2011). Yet, the way people respond to a sequence of signals is a potential confound that should be minimized when estimating the effect of signal valence.

2.1 IQ Test

In the first part of the experiment, I evaluated subjects' cognitive ability using an IQ test.¹⁴ The test consisted of 29 standard logic questions and participants were asked to solve as many of them as possible in 10 minutes. Individual scores were calculated based on the number of correctly answered questions minus the number of incorrect answers, and subjects were paid 0.75 euro for every point they obtained.

Participants were informed that their earnings from the IQ test will be added to their earnings from the remaining parts of the experiment and paid at the end of the session. They were also informed that, although they will receive the entire sum of money at the end of the study, they will not learn their IQ test score nor how much money they earned in each part. Their test results and the details of their payoffs will be available to them one week after the session. Every participant received a personal link to a website where he could check his (and only his) IQ test result and payment details. This procedure enabled me to minimize the dynamic concerns (e.g., subjects might adopt overly pessimistic beliefs to “prepare” themselves for a disappointing outcome) and focus on the trade-off described in the model.

2.2 Belief Elicitation

At the beginning of the second part, participants were told that they have to complete three tasks, for which they can earn up to 12 euro. They were informed that *one task* will be drawn at random at the end of the session, and they will be paid only for that task. In the first task, I elicited subjects' beliefs about their test scores being in the 1st, 2nd, ..., 9th and 10th deciles of the distribution of the test scores of 300 participants who took the same test in the BonnEconLab in the past. I introduced 10 “ranks”, with Rank 1 denoting the highest rank (assigned to participants whose IQ test scores were higher than or equal to the test scores of 90–100% of former participants), and Rank 10 denoting the lowest rank (defined analogously). The first task was to allocate 100 points among the ten ranks in a way that reflects one's beliefs about their relative performance in the IQ test.

¹⁴Cognitive ability is known to strongly correlate with educational achievement, success in the labor market, and income. As people care deeply about intelligence, a measure of IQ seems to be a good candidate for an ego-relevant parameter. At the same time, there is evidence that people have overconfident beliefs about their cognitive ability, suggesting that belief-based utility might be in play. Lastly, there are established methods to assess it, providing me with a measure that is valid and easy to obtain.

Figure 2: The interface used in the first task (the prior belief elicitation).



A screenshot of the computer interface used by subjects is presented in Figure 2. Participants allocated points by dragging blue arrows on ten scales corresponding to Rank 1 to 10. Subjects were informed they can move the arrows back and forth to correct their choices. The text below each scale informed a participant how many points he allocated to a given rank, and the allocation immediately appeared on the graph to the right. The number above the graph indicated how many points the participant still has to allocate before he can proceed to the next task.

To incentivize truthful reports, I used the binarized scoring rule (Hossain and Okui, 2013) as follows. A random variable X takes one of 10 values: $(1,0,\dots,0,0)$, $(0,1,\dots,0,0)$, \dots , $(0,0,\dots,1,0)$, $(0,0,\dots,0,1)$; the position of 1 indicates the decile into which a subject's IQ test score falls. The agent makes a report $x = (x_1, \dots, x_{10})$, where x_i denotes the share of points allocated to the decile $i \in \{1, \dots, 10\}$. The researcher observes the IQ test score in the k^{th} decile, the agent wins the prize if the QSR for multiple events,

$$s(x, k) = 2x_k - \sum_i x_i^2 + 1,$$

exceeds a uniformly drawn random variable with the support $[0, 2]$.

The procedure was explained to the subjects in a simple way. More importantly, the instructions directly spelled out the main implication of the method: the probability of getting a large prize (12 euro) is maximized when a person allocates the points in a way that reflects her beliefs about her rank. I followed the same procedure during the second belief elicitation at the end of the study. It is worth noting that, during the first belief elicitation, subjects were not aware that they will be asked to state their beliefs one more time.

2.3 The Signal Stage

After the first belief elicitation, participants received instructions for the second task. The task was explained in simple language, using pictures and illustrative examples. It was framed in a neutral way and described as follows. *There are two boxes. Each box contains 10 balls with numbers written on them. Box 1 contains balls with numbers from 1 to 10, and every number appears exactly once. The composition of the second box depends on your rank in the IQ test. Box 2 contains 10 balls that all have one number written on them, and this number is equal to your rank.*

The composition of the boxes of a person assigned Rank 2 is presented in Figure 3. For every participant, a computer program randomly selected one of the two boxes. Next, one ball was drawn from the selected box and displayed on the participant's screen. The participant did not know which box the ball was drawn from, but he knew that either box can be selected with equal probability $\frac{1}{2}$. After seeing the ball, he had to state his beliefs about the box selected by the computer. I used the same incentive-compatible elicitation method as for the prior belief elicitation. Subjects had to allocate 100 points

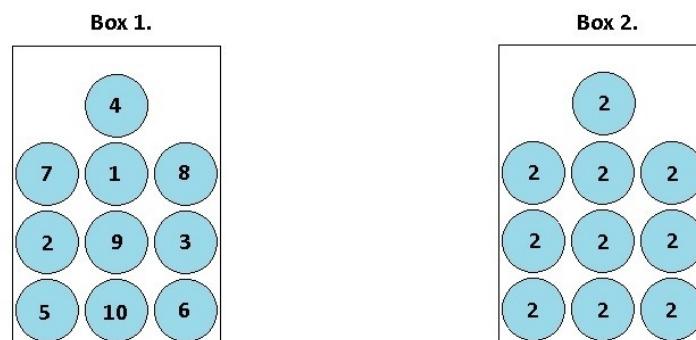


Figure 3: The composition of the boxes of a person whose rank was 2.

between Box 1 and Box 2 in proportions that reflected their beliefs about the source of the signal. The probability of getting a large prize (12 euro) was maximize when a person allocated her points in a way that reflected her true beliefs.

Participants were instructed on how to arrive at the Bayesian posterior given one’s belief.¹⁵ I explained it in two steps with a simple example. First, I demonstrated how a person should allocate her points after different signal realizations if she knew her rank. Then, I showed how a person should allocate her points if she was not sure about her rank, but was assigning a certain probability to it.

Step 1: How should a person ranked 2 allocate her points if she knew for sure that her rank is 2, and saw a ball with the number “2” on it? There are 10 times as many balls with “2” in Box 2 as there are in Box 1, hence it is 10 times as likely that the ball came from the second box. Therefore, the person should allocate 9 points to Box 1, and 10 times as many, 90 points, to Box 2 (the remaining point should be allocated to the box with a higher probability).

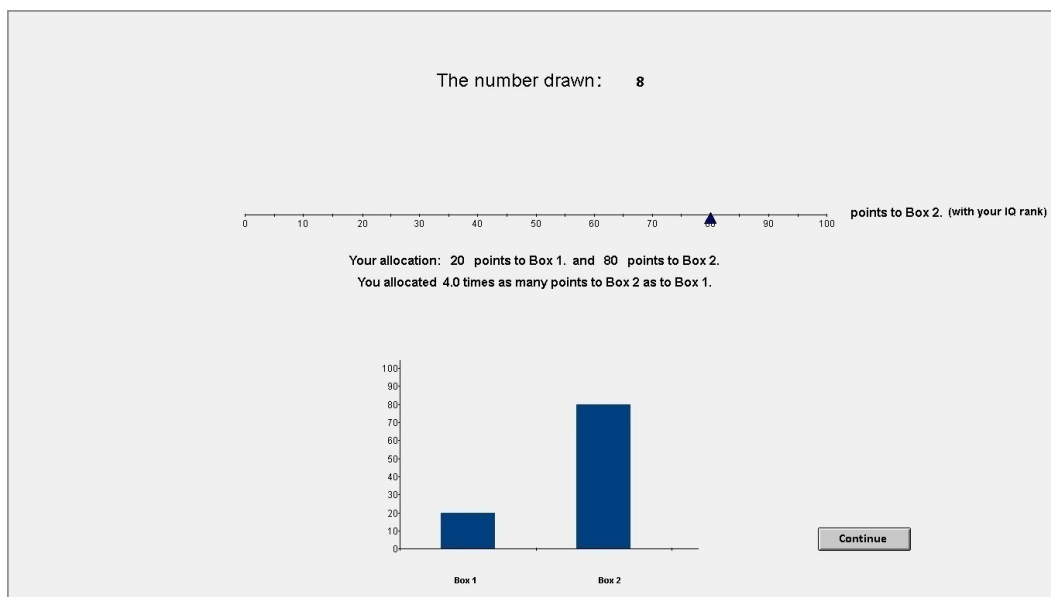
Step 2: What if a person did not know her true rank, but she believed that there is 30% chance that her rank is 2? The same logic applies to this case. One can visualize 30% chance as 3 out of 10 balls in Box 2 having a number “2” on them. In this imaginary case, there are 3 times as many balls with the number “2” on them in Box 2 as in Box 1, implying an allocation of 25 points to Box 1 and 3 times as many (75 points) to Box 2.

A screenshot of the interface used in this task is presented in Figure 4. Importantly, the interface enabled subjects to split the points in the desired proportions without calculating the ratios. This feature was added to minimize computational errors. The text below the scale informed participants about the current allocation and the ratio between the points allocated to the two boxes. By moving the cursor, a subject could choose the number of points corresponding to allocating x times as many points to one of the boxes (with $x \in \{1, 1.1, \dots, 99\}$). The graph below showed the current allocation.

Before proceeding to the signal stage, subjects were required to answer a set of control questions. The questions were designed to check participants’ understanding of the task including the steps necessary to arrive at the Bayesian posterior.

¹⁵Detailed explanations allow to minimize mistakes and ensure subjects’ understanding of the task. They are particularly relevant in the lab environment, which grants a tight control over the belief formation process but is far from natural (in everyday life, it is not common for people to know the precise signal structure and use it to form probabilistic beliefs). Still, there is a risk of framing or creating expectations of what “should” be done in the experiment. Fortunately, I find little evidence in the data of people blindly following the rule.

Figure 4: The interface used in the second task (the signal stage).



2.4 Experimental Conditions

I introduced two experimental conditions: treatment and control. In the control condition, subjects were asked to state their beliefs ex-ante, conditional on a number being drawn. I informed participants that, although they will see every number, their choices are not entirely hypothetical. At the end of the session, one box will be selected by the computer program and one ball will be randomly drawn from the selected box. They will be paid for the decision that corresponds to the number drawn.

This procedure is incentive-compatible as the probability of drawing any number is at least 5%.¹⁶ To alleviate concerns about the non-comparability of the two conditions, I adopted special procedures targeting the issues raised in the literature. One concern is that using the Strategy Method might lead to a better understanding of the game—a consequence of considering the problem from different points of view. In my setup, considering every possible signal may affect belief formation in the control condition. For this reason, I also asked participants in the treatment condition to consider every possible signal realization *before* they saw the actual draw. Subjects were required to go through 10 slides with screenshots of the interface used in the control condition. They were asked to contemplate a hypothetical decision before clicking on the button

¹⁶However, if subjects weigh the cost of cognitive effort against the expected payoff, they may exert less effort in the control condition. In this case, one would expect subjects' decisions to have a higher variance—this hypothesis is not confirmed in the data.

“Continue” which appeared on the screen after 15 seconds. While only the control group was allowed to enter their choices, both groups were required to think about every signal.

Another problem that may arise in the Strategy Method concerns framing subjects with the order of options. I addressed this issue by randomizing the order of numbers displayed in the control condition and the order of slides presented to the subjects in the treatment condition. Moreover, participants in the two conditions used the same interface—the only thing that differed was the headline, which said: “The number drawn” in the treatment condition, and “Consider the number” in the control condition.

2.5 Questionnaires

After each part of the experiment, participants were asked to fill in a questionnaire. The first set of questions, displayed after the IQ test, included a short version of the Big-5 personality test (Gerlitz and Schupp, 2005) and the state-trait anxiety inventory STAI (Spielberger, 1983). The second set of questions, displayed after the three tasks, comprised the Emotion Regulation Questionnaire (ERQ) by Gross and John (2003) and a subset of questions from the Achievement Emotions Questionnaire (AEQ) taken from Pekrun et al. (2011). While Big-5 and STAI are often used in behavioral economics, the last two questionnaires require some explanation.

The ERQ was designed to assess the habitual use of two strategies commonly used to alter emotions. First, one can alleviate the emotional impact of a situation by reinterpreting it in a different way. This emotion regulation strategy, known as *reappraisal*, relies on “applying mental models to the often ambiguous and incomplete information” (Uusberg et al., 2019). The second strategy, *suppression*, involves “inhibiting ongoing emotion-expressive behavior” (Gross and John, 1998, cited in Uusberg et al., 2019). People differ in their use of reappraisal and suppression, and these differences have implications for their experiences of emotions, their behavior in response to those emotions, and general well-being (Gross and John, 2003). To measure the use of these strategies, I administered a 10-item questionnaire developed by Gross and John (2003).¹⁷

The AEQ measures *achievement emotions* (emotions directly linked to achievement activities and outcomes) experienced by students in academic settings (Pekrun et al., 2011). I adopted part of the questionnaire to measure the following test-related emotions: enjoyment, hope, pride, relief, anger, anxiety, shame, and hopelessness.

¹⁷The habitual use of the two strategies is measured by the degree to which subjects agree with particular statements, e.g. “I keep my emotions to myself” or “When I want to feel less negative emotion, I change the way I’m thinking about the situation”.

The questionnaires allow for exploratory investigation of the psychological forces driving the results. The idea presented in this paper is related to research on emotions and decision-making (Lerner et al., 2015). One conclusion from this literature is that emotions may influence decisions via changes in the content of thought, and vice versa.¹⁸ For this reason, emotions and the use of emotion-regulation strategies might play a role in asymmetric updating. Since the hypotheses were not pre-registered, the results are only mentioned in the discussion section. The exploratory data analysis, delegated to Appendix G, is meant to raise a question about the role of emotion regulation in asymmetric updating, which I consider a promising avenue for future research.

3 Theoretical Framework

In this section, I present a one-period model that underlies the experimental design. I formulate testable predictions and describe the empirical strategy used in the analysis.

3.1 The Model

An agent is learning about the state of the world ω (e.g., his cognitive ability) that can be high or low, $\omega \in \{H, L\}$. He has a prior belief about his ability being high p_0 and receives a signal $s \in \{H, L\}$ that induces a posterior belief $p_{1,s}$. The agent derives utility from his beliefs $u(p_t)$, where $t \in \{0, 1\}$ denotes the prior and the posterior belief. The utility function $u(\cdot)$ is increasing in the probability of the high state, concave and twice continuously differentiable. Because of belief-based utility, the agent has an incentive to *manipulate* his beliefs by choosing a different posterior $\tilde{p}_{1,s}$, which enters his utility function $u(\cdot)$. I assume a quadratic cost of belief manipulation that depends on the distance from the unmanipulated posterior $p_{1,s}$. The agent’s utility after a signal s has the following form:

$$U(\tilde{p}_{1,s}) = u(\tilde{p}_{1,s}) - \frac{1}{2\gamma}(p_{1,s} - \tilde{p}_{1,s})^2, \quad (1)$$

where $\gamma > 0$ is the cost parameter. The choice of $\tilde{p}_{1,s}$ that maximizes (1) describes the process of coming to belief about one’s ability.¹⁹ The first-order condition gives us:

$$\gamma u'(\tilde{p}_{1,s}) = \tilde{p}_{1,s} - p_{1,s}. \quad (2)$$

¹⁸A similar hypothesis about anxiety has been recently tested by Engelmann et al. (2019).

¹⁹I assume the process to be outside of the conscious cognition. For this reason, (1) does not include the monetary reward for reporting truthfully—a conscious decision resolved with the elicitation mechanism.

The above condition is the solution to the agent’s problem *after* receiving a signal, that is, the problem faced by participants in the treatment condition. I assume that the unmanipulated belief $p_{1,s}$ is a linear function of the belief formed by an impartial observer—the rational belief $p_{1,s}^B$ (the superscript B denotes the Bayesian update):

$$p_{1,s} = \beta p_{1,s}^B, \tag{3}$$

where $\beta > 0$ captures the degree to which the process of belief formation follows Bayesian updating. I assume this process to be independent of the directional effect arising from belief-based utility.

Asymmetry in Belief Updating

Asymmetric updating describes a situation in which a decision-maker puts higher weight on signals indicating the preferred state compared to the weight he places on the remaining signals (Benjamin, 2019). It is reflected in the manipulation after positive signals being larger than the manipulation after negative signals, which goes in the direction of a less preferred state: $\tilde{p}_{1,L} < p_{1,L}$.²⁰ I define asymmetry as follows:

$$\tilde{p}_{1,H} - p_{1,H} > p_{1,L} - \tilde{p}_{1,L}. \tag{4}$$

In the case of no asymmetry, the left-hand side of (4) is equal to the right-hand side.²¹

Control Condition

What happens when an agent does not receive a signal, but only considers its realization? A hypothetical signal does not change the agent’s belief and does not affect the utility function $u(\cdot)$. The agent keeps deriving utility from his prior belief p_0 . He makes a report $\bar{p}_{1,s}$ about the posterior he would form after the signal s . The agent’s utility is:

$$U(\bar{p}_{1,s}) = u(p_0) - \frac{1}{2\gamma}(p_{1,s} - \bar{p}_{1,s})^2, \tag{5}$$

where the second term denotes the quadratic cost of manipulation of beliefs about the conditional posterior. As previously, the cost depends on the distance from the unmanipulated posterior $p_{1,s}$. The agent’s problem is to maximize (5) by choosing $\bar{p}_{1,s}$. The

²⁰Overweighting negative signals means that the decision-maker manipulates his beliefs by choosing the probability of the low state, $(1 - \tilde{p}_{1,L})$, higher than the unmanipulated probability $(1 - p_{1,L})$. Rearranging gives us the inequality $\tilde{p}_{1,L} < p_{1,L}$, and the extent of manipulation after a negative signal is: $p_{1,L} - \tilde{p}_{1,L}$.

²¹In general, asymmetry can also operate in the direction of a less preferred state. This case would be described by (4) with the reversed inequality sign.

first-order condition gives us:

$$\bar{p}_{1,s} = p_{1,s}. \tag{6}$$

Since the agent has no incentive to manipulate his beliefs, he reports the probability equal to the unmanipulated posterior. Again, I assume that the unmanipulated belief $p_{1,s}$ is a linear function of the rational belief: $\bar{p}_{1,s} = \beta p_{1,s}^B$, with $\beta > 0$.

3.2 Testable Predictions

The overarching question of the paper is whether there is asymmetry in belief formation. I approach this question in two ways. In the first approach, I test the prediction (4) using the Bayesian benchmark as a proxy for the unmanipulated belief $p_{1,s}$. I consider the following regression model:

$$Y_i = \alpha_0 + \alpha_1 Y_i^{Bayes} + \alpha_2 X_i^{signal} + \epsilon_i. \tag{7}$$

The dependent variable Y_i denotes the decision in the main task – how many points a subject allocated to Box 2 after a signal s . This decision reveals the manipulated belief that the state is high after $s = H$ or the state is low after a signal $s = L$. I regress it on an independent variable Y_i^{Bayes} , which denotes the number of points that *should be* allocated according to Bayes’ rule (the Bayesian benchmark), and an indicator variable X_i^{signal} , which takes value 1 if the subject received a “high” signal.²² If people tend to place a higher weight on “good” signals, the difference will be captured by α_2 .

Hypothesis 1.T

Subjects tend to manipulate their beliefs to a larger extent after “good” signals. The coefficient α_2 in (7) is positive.

The equation (2) reveals that the asymmetry stems from the belief-based utility function $u(\cdot)$. Therefore, no asymmetry is expected when deciding about hypothetical signal realizations. The decision of how many points to allocate to Box 2 when considering a signal is only guided by the rational process, as described by (6). I test this prediction by estimating the regression (7) using the data from the control condition.

²²I assume that the scaling parameter β in (3) is the same for “good” and “bad” signals. This assumption is later confirmed in the data. Moreover, I also estimate the regression (7) with a restriction $\alpha_1 = 1$, that is, using the deviations from the Bayesian benchmark $Y_i - Y_i^{Bayes}$ as the dependent variable. All robustness checks can be found in Appendix C.

Hypothesis 1.C

There is no asymmetry in how participants respond to “good” and “bad” signals in the control condition. The coefficient α_2 in (7) is not significantly different from zero.

The causal effect of receiving a “good” signal on belief manipulation can be confirmed in a difference-in-difference analysis. Pooling the data from the treatment and the control condition, I estimate the following regression:

$$Y_i = \beta_0 + \beta_1 Y_i^{Bayes} + \beta_2 Treat_i + \beta_3 X_i^{signal} + \beta_4 Treat_i \times X_i^{signal} + \epsilon_i, \quad (8)$$

where $Treat_i$ is an indicator variable taking value 1 if a subject was assigned to the treatment condition. The coefficient at the interaction term informs us about the effect of receiving “good” news on beliefs in the treatment compared to the control condition.²³

Hypothesis 1.T&C

Subjects tend to manipulate their beliefs after “good” signals to a larger extent in the treatment compared to the control condition. The coefficient β_4 in (8) is positive.

The specifications discussed so far employ the Bayesian benchmark as a proxy for the rational updating process. In the second approach, I propose a different way of modeling unmanipulated beliefs. I use the data from the control condition to predict, for every participant in the treatment group, a counterfactual outcome: what the subject would have decided, had the signal not affected his belief-based utility. To this end, I use the data from the control condition to estimate the following regression:

$$Y_j = \hat{\gamma}_0 + \hat{\gamma}_1 Z_j + \zeta_j, \quad (9)$$

where Z_j is a vector of observables: the subject’s rank, prior beliefs, and the signal under consideration. I use the model to predict the counterfactual outcome \hat{Y}_i for every participant in the treatment group. Next, using the data from the treatment condition,

²³The specification (8) does not take into account that unmanipulated beliefs might have a different effect in the two conditions. As a robustness check, I estimate the same equation adding the interaction $Y_i^{Bayes} \times Treat_i$. The coefficient at the interaction term is not significant, indicating no difference between β in the two conditions.

I estimate the following regression:

$$Y_i - \hat{Y}_i = \gamma_0 + \gamma_1 X_i^{signal} + \zeta_i. \quad (10)$$

The dependent variable is the deviation from the counterfactual outcome, and I regress it on a dummy indicating a “good” signal. If there is asymmetry in belief formation in the direction of the preferred state, the deviations should be greater after positive signals. This effect will be captured by the coefficient at the X^{signal} variable.

Hypothesis 2

Subjects tend to deviate more from the counterfactual outcome after “good” signals. The coefficient γ_1 in (10) is positive.

An alternative specification, allowing for a non-linear relationship between observable characteristics and outcomes, involves matching. As a robustness check, I estimate the nearest neighbor matching model to confirm Hypothesis 2.

4 Results

The experiment was conducted in two waves in summer 2020 and 2023 in BonnEconLab at the University of Bonn. I collected data from 322 participants in the treatment condition and 106 participants in the control condition.²⁴ The experimental sessions lasted around 80 minutes and participants earned 21.4 euro on average. I report the analysis based on the data from 402 participants who made less than three mistakes in five control questions (I excluded 26 participants, that is, 6% of the sample).²⁵

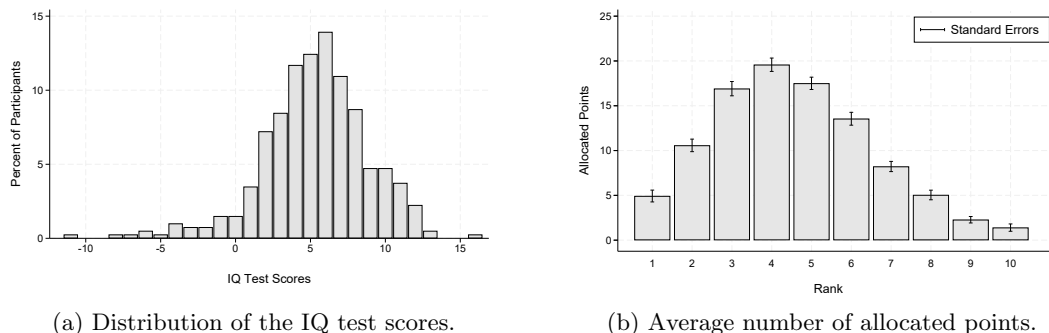
4.1 Raw Data

In this section, I briefly describe the raw data. For a more detailed description, as well as additional figures and tests, see Appendix A. I use observations from both the treatment and the control condition. Panel a) in Figure 5 presents the distribution of the IQ test scores. The distribution has a mean of 5.30 and a standard deviation of 3.58. The

²⁴Women constituted 26% of the sample, with the same share of women in the treatment and in the control condition. The gender differences are gathered in Appendix C.7. While there is little difference in prior beliefs and decisions about signals with non-zero prior probability, women are more likely to believe that an unexpected signal (a signal to which they assigned zero prior probability) is their rank.

²⁵The sample size was estimated based on the number of participants excluded from the analysis in the first wave (13 people, 5.8% of the sample). The results are similar, albeit noisier, if I include observations from mistaken individuals in the analysis.

Figure 5: IQ Test Results and Prior Beliefs.



average rank is equal to 5.54, with a standard deviation of 2.75.²⁶ Importantly, there is no significant difference between the two groups in the average IQ test score or the average rank (see Appendix A.1).

Prior Beliefs about Rank

First, I analyze the aggregate belief distribution. For every rank, I calculate the average number of points allocated by the participants. I present the averages in Panel b) in Figure 5. The distribution is visibly skewed to the right, with the mean of 4.56 and the median of 4. On aggregate, subjects appear to be *overconfident*, as they put a higher probability mass on lower (better) ranks. Second, I examine individual belief distributions. I report the averages of individual measures in Table 1. I look at the average mean belief, the 1st, 2nd, 3rd quartile, and the range. The average mean equals 4.56 and is not different from the average median belief. However, only 45 participants revealed a symmetric belief distribution. Almost half of all subjects (193 participants) revealed a positively skewed belief distribution, and the remaining 164 subjects revealed a nega-

Table 1: Individual belief distributions.

	Mean Belief	Q1	Median	Q3	Range
Mean	4.56	3.82	4.55	5.28	5.02
(Std. Dev.)	(1.69)	(1.70)	(1.73)	(1.78)	(1.58)

²⁶The average rank was not equal to 5.5 because subjects' scores were compared to a different group, as described in Section 2.2.

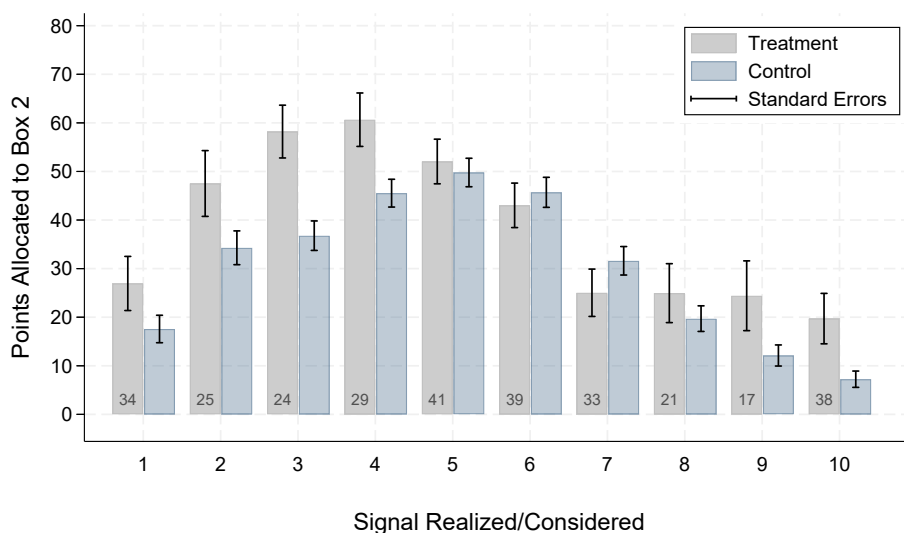
tively skewed distribution. The average difference between the mean and the median in the two groups was 0.24 and -0.23 , respectively. It should be noted that there is a small difference in the average beliefs in the two conditions (0.3 rank, significant at the 5%-level). In the analysis, I always control for prior beliefs or the rational benchmark that is based on subjects' priors.

Decisions in the Main Task

In Figure 6, I present the average number of points allocated to Box 2 in the two conditions. The bars on the left (in light gray color) show the decisions made by subjects in the treatment condition, whereas the bars on the right (in light blue) in the control. The numbers above the x-axis show how many people received a given signal in the treatment condition (in the control, the number is always 101). For example, 24 participants in the treatment condition received a signal “3” and allocated, on average, 58 points to Box 2.

The average number of points allocated in the treatment condition is 38.5 and is 8.5 points higher than the average in the control condition (p-value = 0.001).²⁷ Participants tend to allocate more points after signals “1”, “2”, “3”, or “4” in the treatment compared to the control condition. The average difference equals 13.4, and it is 7.6 points higher

Figure 6: Decisions in the main task in the treatment and control condition.




²⁷When I restrict the sample to participants who assigned a non-zero prior probability to the signal they received, the average increases to 56 points. I note that very few participants allocated exactly 50 points, which was a default option in the main task (the cursor was initially placed in the middle of the scale). The share of decisions equal to the default option was 1.1% in the treatment and 5.2% in the control condition. The results are robust to excluding these observations from the analysis.

than the average difference after the remaining signals.²⁸ A possible explanation for the observed differences is an upward belief manipulation after a positive signal. However, to determine the presence of utility-driven belief manipulation, it is necessary to disentangle the factors driving the effect. To this end, I employ the empirical analysis described in Section 3.2. The results are presented in the same order as the formulated hypotheses. First, I analyze the data from the treatment and the control condition separately, and only afterward, I compare the decisions in the two conditions.

4.2 Data Analysis: Treatment Condition

In this section, I examine subjects' beliefs in the treatment condition. It is important to note that, in the model, belief formation is partly driven by the rational process approximated with the Bayesian update. However, the Bayesian posterior is undefined for a prior equal to zero. That is, Bayes' rule does not specify the posterior of an agent who assigned zero prior probability to the state indicated by the signal. At the same time, due to the signal structure (subjects observed a random number half of the time), this is the case for 40% of participants in the treatment condition. I approach this problem as follows. The main analysis is based on a sample of 173 subjects who received a signal to which they assigned a non-zero prior probability. In the second step, I include in the sample participants who received a signal that was adjacent to the individual prior

Figure 7: Belief manipulation after a signal.



²⁸For signals “1”, “2”, and “3”, which are likely to be considered “good” by the majority of subjects, the difference in points allocated in the two conditions is equal to 12.6 and it is 5.8 points larger than the difference after the remaining signals.

belief distribution.²⁹ The augmented sample includes 212 participants (70% of the treatment group). I analyze subjects’ responses to signals further away from their prior belief distributions separately and discuss them briefly in Section 3.³⁰

Figure 7 shows how belief manipulation, defined by equations (2) and (3), depends on the signal realization. To generate the graph, I use $\beta = \alpha_1$ obtained by estimating (7).³¹ The numbers above the x-axis show how many participants received a given signal. One can notice that after worse signals (indicated by higher numbers), subjects manipulate their beliefs downwards. The pattern is in line with motivated reasoning: participants are *less* convinced that a “bad” signal is their rank.

The estimation results are gathered in Table 2. The results in the first two columns are based on the baseline sample and in the last two columns are based on the augmented sample. For participants who assigned a zero prior probability to the signal, I replaced the prior with the smallest value feasible in the experiment (1%). In every regression in Table 2, the dependent variable is the number of points allocated to Box 2 (the box with numbers equal to one’s rank). In the first specification, I regress it on the number

Table 2: The effect of a “good” signal on beliefs in the treatment condition.

Good Signal	9.928*** (3.604)		8.916*** (3.251)	
Signal Value		-1.975*** (0.728)		-1.710*** (0.627)
Bayes	0.904*** (0.091)	0.860*** (0.091)	0.807*** (0.057)	0.776*** (0.058)
N	173	173	212	212

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best three signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first two columns are based on a sample of subjects who assigned a non-zero prior probability to the signal. In the last two columns, I also include participants who received a signal adjacent to their prior belief distribution.

²⁹A signal is considered adjacent if it is one rank lower (higher) than the first (the last) rank assigned a non-zero prior probability. For example, if a subject with prior beliefs $p = [0, 0, 0.3, 0.5, 0.2, 0, 0, 0, 0, 0]$, where i -th element denotes the probability placed on Rank i , received a signal 3, 4, or 5, he would be included in the main sample. If the signal was 2 or 6, he would be included in the augmented sample.

³⁰In Appendix B, I also present the results based on 1) a sample of subjects who received signals up to two ranks away from their prior belief distributions, and 2) a sample of all participants.

³¹A similar graph with $\beta = 1$, that is, the simple deviations from the Bayesian benchmark, is presented in Appendix C.4.

of points that should be allocated according to the Bayes' rule and a dummy indicating a "good" signal. The "Good Signal" variable takes the value 1 if a subject received one of the best three signals (the result is robust to changes in the definition, i.e., using two, four, or five best signals, see Appendix C.3).

The coefficient at the "Good Signal" variable equals 9.928 and is significant at the 1%-level (p-value of two-tailed t-test = 0.007). Receiving a signal "1", "2", or "3" has a positive effect on the number of points allocated to Box 2. The estimate does not change much when I include subjects who received a signal that was outside of their prior distribution but close enough to be deemed probable (the third column in Table 2).³²

Result 1

Subjects tend to manipulate their beliefs to a larger extent after receiving a "good" signal. The coefficient at the "Good Signal" variable is positive.

In the second specification, I regress the dependent variable on a discrete variable "Signal Value", which takes values from 1 to 10 and denotes the signal realization. The coefficient is negative and significant: observing a higher (worse) signal makes participants deviate more in the negative direction. One can notice in Figure 7 that the negative relation breaks down after the worst signal: 7 participants who received a signal "10" allocated as many points to Box 2 as the subjects who received the best signal "1". If I control for those participants, the coefficient at the "Signal Value" increases to -2.830 and becomes significant at the 1%-level (p-value of two-tailed t-test = 0.001).

Lastly, I show that the effect is not due to selection bias. In a setting with informative signals, selection bias might be a problem because high-ability (low-ability) subjects are more likely to get better (worse) signals. To exclude the possibility of ability-related factors driving the results, I ran the same regressions on a sample of subjects who received a signal from Box 1. Those participants observed a random number—a signal unrelated to their ability. The results can be found in Appendix C.6. The estimates are very similar, supporting the claim that the effect is due to the differential treatment of "good" versus "bad" news and does not stem from the differences in updating between more and less cognitively able individuals.

³²The results are robust to controlling for subjects' gender, rank, and measures of the individual belief distribution. Moreover, I obtain the same results if I use the deviation from the Bayesian benchmark as a dependent variable (this specification is equivalent to restricting $\alpha_1 = 1$). All robustness checks can be found in Appendix C.

4.3 Data Analysis: Control Condition

The analysis of hypothetical choices is based on the data from 101 participants in the control condition. Since every subject in the control group made ten decisions, it leaves me with 1010 data points. In Appendix B.6, I present the average belief manipulation depending on the signal under consideration (an equivalent of Figure 7 for the control condition). The average deviation is slightly below zero and, in contrast to the treatment condition, there is no downward pattern.

In order to test Hypothesis 1.C, I analyze the data in the same way as the data from the treatment condition. Table 3 gathers the results of the same regressions as in Table 2 estimated on the data from the control condition. As before, I only include decisions about signals to which subjects assigned non-zero prior probability. The regression estimates reveal that there is no significant relation between a “good” signal (regardless of the definition) and the conditional choice.

Table 3: The effect of a “good” signal on beliefs in the control condition.

Good Signal	-0.291 (3.269)		-1.925 (2.725)	
Signal Value		-0.943 (0.800)		-0.494 (0.586)
Bayes	0.869*** (0.073)	0.867*** (0.072)	0.719*** (0.040)	0.716*** (0.040)
N	483	483	652	652

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

Standard errors clustered at the individual level. Their values in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best three signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first two columns are based a sample of subjects who assigned non-zero prior probability to the signal. In the last two columns, I also include participants who received a signal adjacent to their prior belief distribution.

Result 2

There is no asymmetry after “good” nor “bad” signals in the control condition. The coefficient at the “Good Signal” variable is not significantly different from zero.

Additionally, I examine whether the weight placed on the Bayesian benchmark is the same in the two conditions. To this end, I compare the coefficient at the “Bayes” variable in every regression in Table 2 to the corresponding coefficient in Table 3. In every case, I cannot reject the hypothesis that the two coefficients are equal. This result substantiates the assumption that β in (3) is the same in the two conditions.

Result 3


The relationship between the Bayesian benchmark and the manipulated belief, captured by the parameter β , is not significantly different in the two conditions.

This result provides suggestive evidence that the estimated β captures the underlying rational process (which should be the same in the two conditions), and the difference between the two conditions is *not* due to their structure (hypothetical vs not) but stems from the utility from beliefs.

4.4 Data Analysis: Treatment Effect

In this section, I compare the decisions in the treatment to the control condition. In Figure 8, I show how the treatment effect depends on the signal realization. Each bar represents the difference between the average number of points allocated in the treatment and the control condition. I interpret this difference as the belief manipulation following the signal. In contrast to the previous sections, the comparison does not involve the Bayesian benchmark. One can notice that subjects tend to manipulate their beliefs

Figure 8: Belief manipulation after a signal.



upwards after better signals: the difference is positive and significant for signals “2”, “3”, and “4”. This observation is confirmed in the regression analysis presented in Table 4. In the first three columns, the estimates are based on observations regarding signals to which subjects assigned non-zero prior probability. In the last three columns, I include decisions about signals that were adjacent to the prior belief distributions. The dependent variable is the number of points allocated to Box 2. In the first specification, I regress it on the Bayesian benchmark, an indicator variable “Treatment”, the “Good Signal” variable (defined as previously), and its interaction with the treatment dummy.

The coefficient at the interaction term is interpreted as the effect of receiving a “good” signal compared to the neutral benchmark—the decisions regarding hypothetical signals in the control condition. The effect is positive and significant at the 5%-level (p-value of two-tailed t-test = 0.029). Its size is substantial: 10.105 constitutes 18% of the average decision made in the treatment condition. The result is robust to changes in the definition of a “good” signal, controlling for rank and measures of belief distribution, and using the deviations from Bayes as the dependent variable. Moreover, the effect is not driven by selection—the estimates do not change if I restrict the sample to subjects

Table 4: The effect of signal valence in the treatment condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Treatment	2.112 (2.475)			0.834 (2.229)		
Good Signal	-0.269 (3.249)	11.485** (4.957)	11.420* (6.628)	-1.838 (2.704)	13.924*** (3.949)	12.779** (6.377)
Treat × Good	10.105** (4.604)			10.440** (4.029)		
Bayes	0.878*** (0.059)			0.738*** (0.034)		
N	656	656	656	864	864	864

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the individual level.

Note: The dependent variable in Specification (1) is the number of points allocated to Box 2. “Good Signal” is an indicator variable that takes the value 1 if a subject received one of the best three signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs.

Estimates in the first three columns are based on a sample of participants who assigned a non-zero prior probability to the signal received or considered. In the last three columns, the sample also includes subjects who received a signal adjacent to their prior belief distribution.

The dependent variable in Specifications (2) and (3) is the deviation from the counterfactual outcome. The counterfactual was based on values predicted in Specification (2) and with the nearest neighbor matching in Specification (3). In both approaches, to predict the outcome, I used the following variables: the signal received, the prior assigned to it, the subject’s rank, and the median belief. The reported errors are bootstrap standard errors based on 500 replications.

who observed a random number. Lastly, the results are not driven by time effects or learning in the control condition. The estimates are very similar when based on either 1) the first five decisions made in the control condition or 2) the last five decisions. All robustness checks can be found in Appendix C.

Result 4

Subjects manipulate their beliefs after “good” signals to a larger extent in the treatment than in the control condition. The coefficient at the interaction term is positive.

In the remaining columns in Table 4, I report the results of estimating the regression (10). The dependent variable is the deviation from the counterfactual outcome constructed using the data from the control condition. The results in columns two and five are based on the outcomes predicted with the regression model (9) and (10). The results in columns three and six are based on the counterfactual constructed using the nearest neighbor matching with the number of neighbors set to three. In both specifications, I use the following observable characteristics: the signal received, the prior assigned to it, the subject’s rank, and the median belief. I report the results based on the baseline sample of participants with non-zero prior beliefs (columns three and four) and the augmented sample (columns five and six). The deviations from the counterfactual are regressed on a variable indicating a “good” signal (defined as previously). As we can see in Table 4, the coefficients at the “Good Signal” variable are positive and significant—after “good” news, participants tend to deviate in the direction of the preferred state.

Result 5

Subjects tend to deviate more from the counterfactual outcome after “good” signals. The coefficient at the “Good Signal” variable is positive.

The results provide further support for the model and reveal the mechanism behind asymmetric updating. When receiving a “good” signal, people tend to interpret it as more informative compared to what they would say about the same signal ex-ante. Because the thought of being smart feels good, they would not discard the signal but rather persuade themselves that it is accurate. The effect operates in the direction of overconfidence, implying that this well-known bias might be a consequence of asymmetric updating.

4.5 Discussion

There are several points that should be noted to complete the analysis. In Section 4.5.1, I briefly describe subjects' responses to signals that were far from their prior belief distributions. The consistency between subjects' beliefs about the box and the posterior belief distributions is discussed in Section 4.5.2. In Section 4.5.3, I comment on the questionnaire data, providing additional evidence for the mechanism behind the results.

4.5.1 Prior probability of zero

How do people update after a signal indicating a state they assigned a prior probability of zero? In Appendix B.3, I analyze the data on “unexpected” signals. The sample includes 89 observations in the treatment condition (29% of participants in the treatment) and 358 observations in the control condition (35% of observations in the control). Figure 14 in Appendix B.3 shows subjects' decisions in the two conditions. One can notice that beliefs depend on the signal value in a different way than in the main analysis. Participants in the treatment condition do not perceive lower (better) signals as more likely to be informative.³³ At the same time, there is a significant effect of a “good” signal in the control condition (see Table 9 in Appendix B.3). Subjects tend to allocate 8 points more to Box 2 when considering a signal “1”, “2” or “3” (the effect is significant at the 5%-level).³⁴ The difference-in-difference analysis reveals a significant negative effect of a “good” signal in the treatment condition. Participants tend to be more skeptical of positive unexpected signals when they receive them.

The results suggest another force (in addition to the costs of belief manipulation) limiting the extent of motivated reasoning. News “too good to be true” is discounted more than the unexpected “bad” news. This could explain why overconfidence is not a universal trait. The results also show the importance of distinguishing possible (or expected) versus unexpected information. Exploring this feature of individual feedback is one possible direction for future research.

4.5.2 Consistency

The data from the additional belief elicitation is described in Appendix D. On average, subjects in the baseline sample allocated 32.86 points to the rank corresponding to the signal they received. The difference between the prior and the posterior belief amounts to 11.36 percentage points and is highly significant (p-value one-tailed t-test =0.000). The

³³The effect of a “good” signal is not significant unless I exclude the outlier visible on Figure 14, in which case it becomes significant at the 10%-level.

³⁴The result does not change much if I define a signal “4” to be a “good” signal.

change in beliefs varies with the signal value as expected: it is 80% higher after signals 1 to 4 than after signals 5 to 10 (p-value of one-tailed t-test = 0.0097). A regression analysis in Appendix D reveals two insights. First, there is a strong correlation between the decisions about the box and the posterior about the rank, which remains highly significant even when I control for the Bayesian benchmark. This result proves that participants correctly understood the main task and based the decisions on their beliefs about the rank. Second, when I account for the decisions in the main task, signal valence has no effect on the posterior belief. In other words, there is *no additional* asymmetry in how subjects translate the decision about the box into the beliefs about the rank. This result confirms that the two-box design can be used to study the extent of asymmetric updating. Still, the additional belief elicitation should be interpreted with caution. I describe the caveats of eliciting beliefs multiple times at the end of Appendix D.

4.5.3 Questionnaires

In Appendix G, I present the analysis of subjects' responses in questionnaires. I regress participants' decisions in the main task on the Bayesian benchmark, achievement emotions, emotion-regulation strategies, and personality traits. It is worth noting that none of the BIG-5 personality traits correlates with the decisions in the main task. The only variables that correlate with subjects' beliefs are the negative achievement emotions and the habitual use of reappraisal. The results are different for signals close to one's prior belief distribution and the unexpected signals (this further supports the claim that the two cases involve different processes). Reappraisal has a positive effect when assessing signals close to one's priors: subjects who are above the median in their use of reappraisal allocate 7.8 points more to Box 2 (p-value of two-tailed t-test = 0.01). In consequence, they deviate significantly less from the Bayesian benchmark.³⁵ When evaluating a signal far from one's prior belief distribution, the effect of reappraisal diminishes and the decisions become driven by the achievement emotions. People who report experiencing more negative emotions (as compared to the median report) tend to allocate 11.3 points more to Box 2 (p-value two-tailed t-test = 0.01). Similarly, those who are less anxious (based on the responses in the STAI questionnaire) tend to allocate 9.6 fewer points to Box 2 (p-value two-tailed t-test = 0.03). As a result, more upset or anxious participants end up further away from the Bayesian benchmark.³⁶

³⁵One possible explanation is that people using reappraisal are able to re-interpret the situation to align their beliefs with the evidence at hand.

³⁶I obtain the same results when I look at the deviations from the counterfactual, as in (10), instead of the Bayesian benchmark.

The results suggest that belief formation is affected by 1). the current emotional state, and 2). the ability to handle one’s emotions with the use of reappraisal. This points toward the role of beliefs in managing emotional states, as suggested by the psychological literature (Lerner et al., 2015). More work is needed to understand the dynamics of emotion regulation and its implication for belief formation.

5 Conclusions

There is mounting evidence that people derive utility not only from physical outcomes but also from their beliefs about the current or future state. The belief-based utility is likely to be the driving force behind overconfidence, the demand for (and the avoidance of) information, and belief polarization. Yet, the way it influences people’s actions and beliefs is not fully understood. My study takes the next step toward explaining its role by revealing how belief-based utility shapes the way we interpret new information.

References

- Barron, Kai (2021). “Belief updating: does the ‘good-news, bad-news’ asymmetry extend to purely financial domains?” In: *Experimental Economics* 24.1, pp. 31–58.
- Bénabou, Roland and Jean Tirole (2016). “Mindful Economics: The Production, Consumption, and Value of Beliefs”. In: *Journal of Economic Perspectives* 30.3, pp. 141–164.
- Benjamin, Daniel J (2019). “Errors in probabilistic reasoning and judgment biases”. In: *Handbook of Behavioral Economics: Applications and Foundations 1 2*, pp. 69–186.
- Brandts, Jordi and Gary Charness (2009). “The strategy method: A survey of experimental evidence”. Working Paper.
- Brunnermeier, Markus K and Jonathan A Parker (2005). “Optimal expectations”. In: *American Economic Review* 95.4, pp. 1092–1118.
- Buser, Thomas, Leonie Gerhards, and Joël Van Der Weele (2018). “Responsiveness to feedback as a personal trait”. In: *Journal of Risk and Uncertainty* 56.2, pp. 165–192.
- Caplin, Andrew and John V Leahy (2019). “Wishful Thinking”. Working Paper.
- Chew, Soo Hong, Wei Huang, and Xiaojian Zhao (2020). “Motivated false memory”. In: *Journal of Political Economy* 128.10, pp. 3913–3939.
- Coutts, Alexander (2019). “Good news and bad news are still news: Experimental evidence on belief updating”. In: *Experimental Economics* 22.2, pp. 369–395.
- Drobner, Christoph (2022). “Motivated beliefs and anticipation of uncertainty resolution”. In: *American Economic Review: Insights* 4.1, pp. 89–105.

- Drobner, Christoph and Sebastian J Goerg (2024). “Motivated belief updating and rationalization of information”. In: *Management Science*.
- Eil, David and Justin M. Rao (2011). “The good news-bad news effect: Asymmetric processing of objective information about yourself”. In: *American Economic Journal: Microeconomics* 3.2, pp. 114–138.
- Engelmann, Jan, Maël Lebreton, Peter Schwardmann, Joel J van der Weele, and Li-Ang Chang (2019). “Anticipatory anxiety and wishful thinking”. Working Paper.
- Ertac, Seda (2011). “Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback”. In: *Journal of Economic Behavior & Organization* 80.3, pp. 532–545.
- Falk, Armin and Florian Zimmermann (2017). “Consistency as a signal of skills”. In: *Management Science* 63.7, pp. 2197–2210.
- Fudenberg, Drew and David K Levine (2006). “A dual-self model of impulse control”. In: *American Economic Review* 96.5, pp. 1449–1476.
- Gerlitz, Jean-Yves and Jürgen Schupp (2005). “Zur Erhebung der Big-Five-basierten persönlichkeitsmerkmale im SOEP”. In: *DIW Research Notes* 4.
- Golman, Russell, David Hagmann, and George Loewenstein (2017). “Information avoidance”. In: *Journal of Economic Literature* 55.1, pp. 96–135.
- Gross, James J and Oliver P John (2003). “Individual differences in two emotion regulation processes: implications for affect, relationships, and well-being.” In: *Journal of Personality and Social Psychology* 85.2, p. 348.
- Hossain, Tanjim and Ryo Okui (2013). “The binarized scoring rule”. In: *Review of Economic Studies* 80.3, pp. 984–1001.
- Huffman, David, Collin Raymond, and Julia Shvets (2022). “Persistent overconfidence and biased memory: Evidence from managers”. In: *American Economic Review* 112.10, pp. 3141–3175.
- Lerner, Jennifer S, Ye Li, Piercarlo Valdesolo, and Karim S Kassam (2015). “Emotion and decision making”. In: *The Annual Review of Psychology* 66.
- Möbius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat (2022). “Managing self-confidence: Theory and experimental evidence”. In: *Management Science* 68.11, pp. 7793–7817.
- Pekrun, Reinhard, Thomas Goetz, Anne C Frenzel, Petra Barchfeld, and Raymond P Perry (2011). “Measuring emotions in students’ learning and performance: The Achievement Emotions Questionnaire (AEQ)”. In: *Contemporary Educational Psychology* 36.1, pp. 36–48.
- Schwardmann, Peter and Joel Van der Weele (2019). “Deception and self-deception”. In: *Nature Human Behaviour* 3.10, pp. 1055–1061.
- Spielberger, Charles D (1983). “State-Trait Anxiety Inventory for Adults”. In: *PsycheTESTS Dataset*.

Uusberg, Andero, Jamie L Taxer, Jennifer Yih, Helen Uusberg, and James J Gross (2019). “Reappraising reappraisal”. In: *Emotion Review* 11.4, pp. 267–282.

Zimmermann, Florian (2020). “The dynamics of motivated beliefs”. In: *American Economic Review* 110.2, pp. 337–61.

A Descriptive Statistics

A.1 Differences between participants in the two conditions

Table 5: Differences between participants in the two conditions.

	Treatment	Control		Diff < 0	Diff \neq 0	Diff > 0
IQ test score	5.349	5.168	<i>p-value:</i>	0.669	0.661	0.331
Rank	5.442	5.832	<i>p-value:</i>	0.109	0.218	0.891
<i>Measures of Belief Distribution:</i>						
Mean Belief	4.463	4.817	<i>p-value:</i>	0.038	0.076	0.962
1 st Quartile	3.738	4.069	<i>p-value:</i>	0.045	0.090	0.955
Median Belief	4.481	4.812	<i>p-value:</i>	0.044	0.088	0.956
3 st Quartile	5.193	5.530	<i>p-value:</i>	0.050	0.100	0.950
Range	5.096	4.782	<i>p-value:</i>	0.041	0.083	0.958
N	301	101				

Table 6: Differences between participants in the two conditions.

	Treatment	Control		Diff < 0	Diff \neq 0	Diff > 0
BIG-5: Extr	14.10	13.90	<i>p-value:</i>	0.663	0.674	0.337
BIG-5: Cons	14.10	14.06	<i>p-value:</i>	0.544	0.913	0.456
BIG-5: Open	14.67	13.97	<i>p-value:</i>	0.956	0.087	0.044
BIG-5: Neur	12.65	13.35	<i>p-value:</i>	0.070	0.141	0.930
BIG-5: Agree	15.29	15.31	<i>p-value:</i>	0.479	0.958	0.521
STAI: State	58.98	58.10	<i>p-value:</i>	0.777	0.446	0.223
STAI: Trait	58.34	57.47	<i>p-value:</i>	0.763	0.473	0.237
N	301	101				

A.2 Decisions in the two conditions

Table 7: Beliefs and decisions about signals with non-zero prior probability.

	Treatment	Control		Diff < 0	Diff \neq 0	Diff > 0
Prior Beliefs*	21.502 (1.093)	20.910 (0.641)	<i>p-value:</i>	0.669	0.661	0.331
*Probability placed in Belief Elicitation I on the signal.						
Main Task**	56.121 (2.043)	50.242 (1.284)	<i>p-value:</i>	0.991	0.018	0.009
**Number of points allocated to Box 2 in the main task.						
Bayes***	61.291 (1.362)	60.325 (0.832)	<i>p-value:</i>	0.725	0.549	0.275
***Points that should allocated according to Bayes' rule.						
N	173	483				

Note: Standard errors in parentheses.

Table 8: Beliefs and decisions about signals with zero prior probability.

	Treatment	Control		Diff < 0	Diff \neq 0	Diff > 0
Main Task*	14.664 (2.157)	11.512 (0.930)	<i>p-value:</i>	0.927	0.146	0.073
*Number of points allocated to Box 2 in the main task.						
N	128	527				


Note: Standard errors in parentheses.

B Additional Results

B.1 Results based on the baseline sample (within priors)

In Figure 9, I present the raw data on participants' decisions on signals indicating a rank assigned a non-zero prior probability. The remaining graphs and tables can be found in the main text. The difference between the average decisions in the two conditions concerning the best three signals is equal to 12.43, is significant at the 1% level (p-value one-tailed t-test =0.006), and is 9 points higher than the difference between the two conditions after the remaining signals. If I include the number “4” in the definition of a “good” signal, the difference between the average decision in the two conditions is equal to 12.86 (significant at the 1% level, p-value one-tailed t-test =0.0004) and 12.5 points higher than the difference after the remaining signals.

Figure 9: Decisions in the main task (within priors).



B.2 Results based on the augmented sample (within priors ± 1 rank)

Figure 10: Decisions in the main task (within priors ± 1 rank).




Figure 11: Belief manipulation after a signal (within priors ± 1 rank).




Figure 12: Belief manipulation after a signal (Treatment minus Control).





Figure 13: Changes in the number of points allocated to the rank = signal.



B.3 Unexpected signals (higher/lower than within prior ± 1 rank)

Figure 14: Decisions in the main task (higher/lower than within prior ± 1 rank).

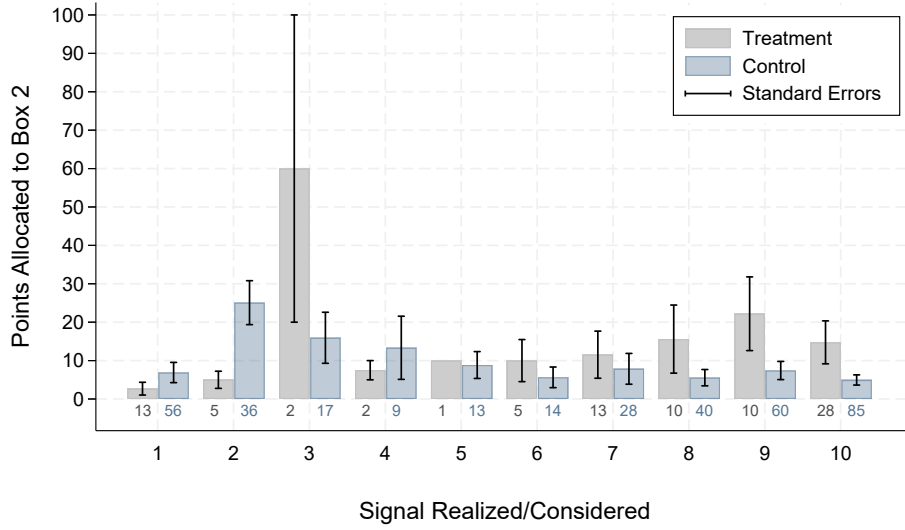


Figure 15: Belief manipulation after a signal (higher/lower than within prior ± 1 rank).




Figure 16: Belief manipulation after a signal (Treatment minus Control).




Figure 17: Changes in the number of points allocated to the rank = signal.




Table 9: The effect of a “good” signal far from prior belief distribution.

	<i>Treatment group</i>		<i>Control group</i>		<i>Both groups</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment					8.211** (3.38)	8.682** (3.46)
Good Signal 1-3	-5.725 (6.44)		7.798** (3.46)		7.798** (3.45)	
Good Signal 1-4		-6.077 (6.23)		7.979** (3.60)		7.979** (3.60)
Treat × Good 1-3					-13.523** (6.78)	
Treat × Good 1-4						-14.056** (6.58)
N	89	89	358	358	447	447

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” is an indicator variable taking value 1 if a subject received one of the best 3 signals, that is, a signal “1”, “2” or “3”. “Good Signal 1-4” is defined analogously, with the best 4 signals. “Treatment” is an indicator variable taking value 1 if a subject was assigned to the treatment condition. The first two columns are based on participants in the treatment condition, whereas the results in column 3 and 4 are based on observations from the control. All results are based on participants who received or considered a signal that was more than 1 rank away from their prior belief distribution.

Table 10: The effect of a “good” signal far from prior belief distribution.

	<i>Treatment group</i>		<i>Control group</i>		<i>Both groups</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment					6.957** (3.18)	7.393** (3.26)
Good Signal 1-3	-9.260 (5.62)		5.373* (3.03)		5.373* (3.03)	
Good Signal 1-4		-9.128* (5.42)		5.742* (3.26)		5.742* (3.26)
Treat × Good 1-3					-14.633*** (4.47)	
Treat × Good 1-4						-14.869*** (4.65)
N	87	87	355	355	442	442

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” is an indicator variable taking value 1 if a subject received one of the best 3 signals, that is, a signal “1”, “2” or “3”. “Good Signal 1-4” is defined analogously, with the best 4 signals. “Treatment” is an indicator variable taking value 1 if a subject was assigned to the treatment condition. The first two columns are based on participants in the treatment condition, whereas the results in column 3 and 4 are based on observations from the control. All results are based on participants who received or considered a signal that was more than 1 rank away from their prior belief distribution, **excluding subjects with extreme beliefs** ($p_1 = 1$).

B.4 Results based on the second augmented sample (within prior ± 2 ranks)

Figure 18: Decisions in the main task (within prior ± 2 rank).




Figure 19: Belief manipulation after a signal (within prior ± 2 rank).

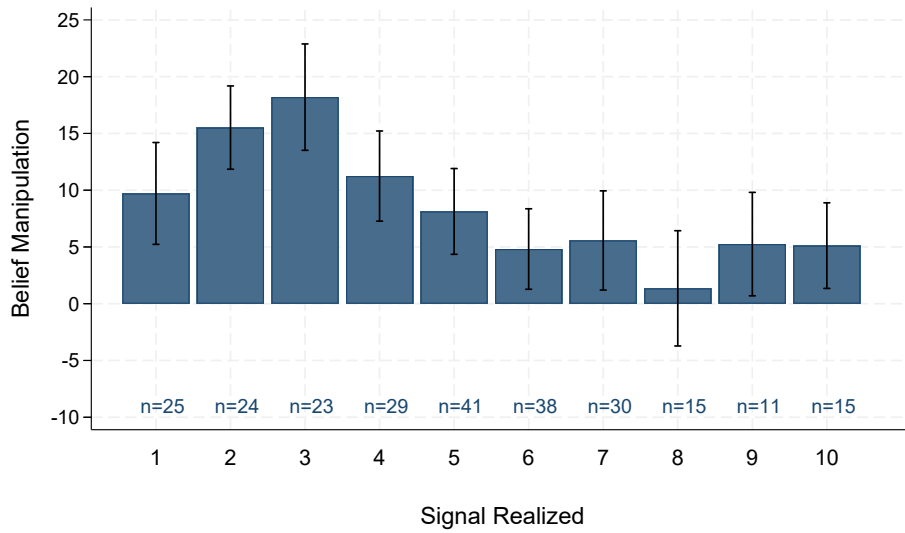


Figure 20: Belief manipulation after a signal (Treatment minus Control).




Figure 21: Changes in the number of points allocated to the rank = signal.




Table 11: The effect of a “good” signal in the treatment condition.

	<i>Treatment group</i>		<i>Control group</i>		<i>Both groups</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment					1.199 (1.96)	8.068** (4.00)
Good Signal	7.960*** (2.99)		-0.494 (2.47)		-0.411 (2.45)	
Display		-1.333** (0.55)		-0.549 (0.47)		-0.530 (0.47)
Treat × Good					8.255** (3.84)	
Treat × Display						-0.891 (0.68)
Bayes	0.818*** (0.05)	0.791*** (0.05)	0.745*** (0.04)	0.739*** (0.04)	0.762*** (0.03)	0.751*** (0.03)
Observations	251	251	793	793	1044	1044

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. Results are based on an augmented sample of participants whose priors were not far from the signal (no more than 2 ranks away from the prior belief distribution). Columns 1 and 2 shows the effect of signal valence in the treatment condition. Columns 3 and 4 present the effect of signal valence in the control condition. The last two columns contain the results of the difference-in-difference analysis based on subjects from both conditions.

B.5 Results based on the entire sample (all participants)

Figures 22, 23, and 24 present the equivalent of Figures 7, 8, and 29 based on the data from the entire sample (irrespective of subjects' prior beliefs).

Figure 22: Belief manipulation after a signal (all participants).




Figure 23: Belief manipulation after a signal (Treatment minus Control).





Figure 24: Changes in the number of points allocated to the rank = signal.



B.5.1 Regression analysis (all participants)

Table 12: The effect of a good signal on beliefs in the treatment condition.

Good Signal	5.179*	5.391*	9.928***
	(2.948)	(2.957)	(3.834)
Outside Prior		5.343	8.879
		(5.752)	(6.040)
Good × Out			-11.068*
			(5.989)
Bayes	0.813***	0.893***	0.904***
	(0.045)	(0.098)	(0.097)
N	301	301	301

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

Standard errors in parentheses.

Table 13: The effect of a good signal on beliefs in the control condition.

Good Signal	1.697 (2.361)	1.761 (2.366)	-0.291 (3.268)
Outside Prior		5.896 (3.838)	4.597 (4.232)
Good \times Out			3.817 (3.405)
Bayes	0.780*** (0.038)	0.872*** (0.072)	0.869*** (0.073)
N	1010	1010	1010

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

Standard errors in parentheses.

Table 14: The effect of signal valence in the treatment condition.

Treatment	3.215* (1.88)	3.234* (1.88)	3.964* (2.37)	2.112 (2.47)
Good Signal	1.721 (2.35)	1.782 (2.35)	1.440 (3.07)	-0.269 (3.25)
Treat \times Good	3.455 (3.62)	3.626 (3.61)	3.669 (3.68)	10.105** (4.60)
Bayes	0.787*** (0.03)	0.878*** (0.06)	0.877*** (0.06)	0.878*** (0.06)
Outside Prior		5.822* (3.18)	5.965* (3.54)	5.056 (3.62)
Treat \times Out			-1.641 (3.40)	2.444 (3.86)
Good \times Out			0.609 (2.95)	3.795 (3.38)
Treat \times Good \times Out				-14.771** (6.53)
N	1311	1311	1311	1311

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

B.6 Belief manipulation in the control condition

In this section, I look at belief manipulation in the control condition. Figure 25 and 26 show how the extent of belief manipulation depends on a signal (belief manipulation is defined by equations (2) and (3)). To generate each graph, I use $\beta = \alpha_1$ obtained by estimating (7) in a given sample (the baseline or all subjects in the control condition). In the case of signals indicating a rank assigned a zero prior probability, I replace the prior with the lowest value feasible in the experiment (1%). I present the results on the same scale (from -40 to 20) to facilitate comparisons.

One can note that there is little manipulation after signals to which subjects assigned non-zero prior probability (Figure 25). The estimated β is equal to 0.87, and for most signals, participants only slightly deviate from the rational benchmark. The average deviation from rationality is equal to -2.24 and, while it is significantly different from zero, it is mostly driven by subjects' responses to the worst signals. There is little belief manipulation for signals 1 to 7.

If I augment the sample to include decisions regarding signals to which participants assigned zero prior probability, the estimated β drops to 0.78, as the added participants significantly deviate from the Bayesian benchmark. They deviate in the positive direction, hence the belief manipulation in Figure 26 takes values above zero. The average manipulation increases to 3.83.

Figure 25: Belief manipulation in the control condition (within priors).





Figure 26: Belief manipulation in the control condition (all signals).



C Robustness checks

C.1 Excluding extreme beliefs

Table 15: The effect of a good signal on beliefs in the treatment condition.

Good Signal	10.075*** (3.61)		9.053*** (3.25)	
Signal Value		-2.185*** (0.74)		-1.864*** (0.63)
Bayes	0.893*** (0.09)	0.842*** (0.09)	0.802*** (0.06)	0.766*** (0.06)
N	172	172	211	211

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best three signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first two columns are based a sample of subjects who assigned non-zero prior probability to the signal. In the last two columns, I also include participants who received a signal adjacent to their prior belief distribution. I exclude 1 subject with extreme beliefs ($p_1 = 1$).

Table 16: The effect of a good signal on beliefs in the control condition.

Good Signal	-0.291 (3.27)		-2.350 (2.64)	
Signal Value		-0.943 (0.80)		-0.393 (0.56)
Bayes	0.869*** (0.07)	0.867*** (0.07)	0.724*** (0.04)	0.723*** (0.04)
N	483	483	651	651

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

Standard errors clustered at the individual level. Their values in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best three signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first two columns are based a sample of subjects who assigned non-zero prior probability to the signal. In the last two columns, I also include participants who received a signal adjacent to their prior belief distribution. I exclude subjects with extreme beliefs ($p_1 = 1$).

Table 17: The effect of signal valence in the treatment condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Treatment	1.926 (2.48)			0.639 (2.23)		
Good Signal	-0.276 (3.25)	11.627** (4.78)	11.255* (6.76)	-2.278 (2.61)	14.423*** (3.73)	12.726* (6.67)
Treat \times Good	10.292** (4.61)			11.075*** (3.96)		
Bayes	0.875*** (0.06)			0.742*** (0.03)		
Observations	655	655	655	862	862	862

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable in Specification (1) is the number of points allocated to Box 2. “Good Signal” is an indicator variable that takes the value 1 if a subject received one of the best three signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. In both conditions, I exclude subjects with extreme beliefs ($p_1 = 1$).

Estimates in the first three columns are based on a sample of participants who assigned a non-zero prior probability to the signal received or considered. In the last three columns, the sample also includes subjects who received a signal adjacent to their prior belief distribution.

The dependent variable in Specifications (2) and (3) is the deviation from the counterfactual outcome. The counterfactual was based on values predicted in Specification (2) and with the nearest neighbor matching in Specification (3). In both approaches, to predict the outcome, I used the following variables: the signal received, the prior assigned to it, the subject’s rank, and the median belief. The reported errors are bootstrap standard errors based on 500 replications.

C.2 Controlling for rank and median belief

Table 18: The effect of a “good” signal in the treatment condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Good Signal	9.206** (4.583)	9.855** (4.742)	10.069* (5.791)	9.155** (3.832)	10.904*** (4.069)	11.579** (4.576)
Bayes	0.869*** (0.094)	0.945*** (0.097)	0.900*** (0.100)	0.787*** (0.059)	0.809*** (0.060)	0.785*** (0.062)
N	173	173	173	212	212	212

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Specification (1) includes controls for subjects’ median beliefs. Specification (2) includes controls for subjects’ rank. In Specification (3), I control for both. Results in the first three columns are based on decisions about signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: decisions about signals adjacent to one’s prior belief distribution.

Table 19: The effect of a “good” signal in the control condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Good Signal	0.008 (3.761)	-0.605 (3.242)	-0.493 (3.809)	-1.991 (2.889)	-1.935 (2.683)	-2.119 (2.877)
Bayes	0.880*** (0.075)	0.862*** (0.075)	0.871*** (0.079)	0.718*** (0.041)	0.716*** (0.041)	0.715*** (0.041)
N	483	483	483	652	652	652

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Specification (1) includes controls for subjects’ median beliefs. Specification (2) includes controls for subjects’ rank. In Specification (3), I control for both. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals adjacent to one’s prior belief distribution. Standard errors clustered at the participant level.

Table 20: The effect of a “good” signal in the treatment vs control condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Treatment	1.743 (2.469)	2.290 (2.495)	2.038 (2.578)	0.985 (2.220)	0.872 (2.229)	1.139 (2.293)
Good Signal	-0.237 (2.677)	-0.448 (2.320)	-0.504 (2.684)	-1.693 (2.114)	-1.792 (1.970)	-1.801 (2.115)
Treat × Good	10.459** (4.538)	11.481** (4.707)	11.776** (4.775)	9.886** (4.015)	12.280*** (4.122)	11.705*** (4.168)
Bayes	0.883*** (0.050)	0.880*** (0.049)	0.882*** (0.051)	0.737*** (0.029)	0.740*** (0.029)	0.738*** (0.029)
Observations	656	656	656	864	864	864

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Specification (1) includes controls for subjects’ median beliefs. Specification (2) includes controls for subjects’ rank. In Specification (3), I control for both. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals adjacent to one’s prior belief distribution.

C.3 Alternative definition of a “good” signal

Table 21: The effect of a “good” signal in the treatment condition.

Good Signal 1-2	11.657*** (4.269)			8.502** (3.810)		
Good Signal 1-4		9.586*** (3.280)			8.797*** (3.003)	
Good Signal 1-5			8.199** (3.470)			8.269*** (3.105)
Bayes	0.929*** (0.093)	0.860*** (0.091)	0.841*** (0.093)	0.813*** (0.058)	0.779*** (0.057)	0.765*** (0.058)
N	173	173	173	212	212	212

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-n” is an indicator variable taking value 1 if a subject received one of the best 2, 4, or 5 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal.

Table 22: The effect of a “good” signal in the control condition.

Good Signal 1-2	2.590 (3.935)			-1.540 (3.167)		
Good Signal 1-4		-0.083 (2.952)			-0.293 (2.578)	
Good Signal 1-5			2.306 (2.963)			2.027 (2.619)
Bayes	0.874*** (0.074)	0.869*** (0.071)	0.866*** (0.071)	0.718*** (0.041)	0.720*** (0.040)	0.715*** (0.040)
N	483	483	483	652	652	652

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-n” is an indicator variable taking value 1 if a subject received one of the best 2, 4, or 5 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal. Standard errors clustered at the participant level.

Table 23: The effect of a “good” signal in the treatment vs control condition.

Treatment	3.338 (2.232)	0.802 (2.703)	1.290 (3.351)	2.185 (2.017)	-0.000 (2.393)	-0.209 (2.876)
Good Signal 1-2	2.656 (2.894)			-1.381 (2.408)		
Treat × Good 1-2	8.644 (5.325)			9.207** (4.653)		
Good Signal 1-4		-0.084 (2.092)			-0.294 (1.806)	
Treat × Good 1-4		9.650** (4.078)			9.317** (3.659)	
Good Signal 1-5			2.320 (2.138)			1.942 (1.826)
Treat × Good 1-5			5.754 (4.211)			6.728* (3.701)
Bayes	0.888*** (0.049)	0.867*** (0.049)	0.859*** (0.049)	0.740*** (0.029)	0.734*** (0.029)	0.726*** (0.029)
N	656	656	656	864	864	864

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-2” is an indicator variable taking value 1 if a subject received one of the best 2 signals, that is, a signal “1” or “2”. “Good Signal 1-4” and “Good Signal 1-5” are defined analogously. “Treatment” is an indicator variable taking value 1 if a subject was assigned to the treatment condition. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution.

C.4 Results with a restriction $\alpha_1 = 1$ (deviations from Bayes, unweighted)

Figure 27: Deviations from Bayes in the treatment condition.




Figure 28: Deviations from Bayes in the control condition.

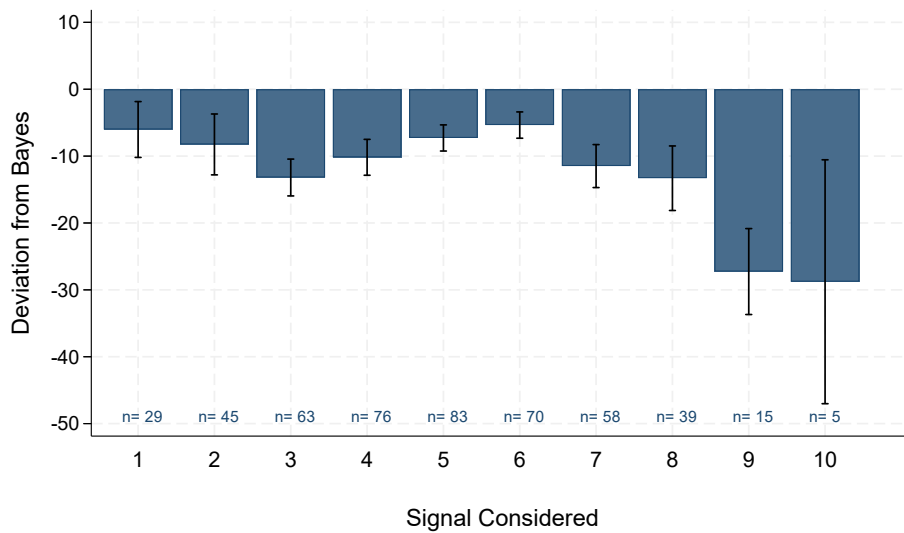


Table 24: The effect of a “good” signal in the treatment condition.

Good Signal 1-3	10.270*** (3.591)			9.783*** (3.319)		
Good Signal 1-4		9.188*** (3.283)			7.712** (3.086)	
Signal Value			-1.879** (0.728)			-1.409** (0.643)
Observations	173	173	173	212	212	212

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the deviation from the Bayesian benchmark (the number of points allocated to Box 2 minus the number of points that should be allocated given one’s prior beliefs). “Good Signal 1-3” is an indicator variable taking value 1 if a subject received one of the best 3 signals, that is, a signal “1”, “2” or “3”. “Good Signal 1-4” and “Good Signal 1-5” are defined analogously. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal.

Table 25: The effect of a “good” signal in the control condition.

Good Signal 1-3	0.037 (3.237)			-0.706 (2.738)		
Good Signal 1-4		-0.033 (2.982)			-0.315 (2.661)	
Signal Value			-0.918 (0.815)			-0.135 (0.602)
N	483	483	483	652	652	652

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the deviation from the Bayesian benchmark (the number of points allocated to Box 2 minus the number of points that should be allocated given one’s prior beliefs). “Good Signal 1-3” is an indicator variable taking value 1 if a subject received one of the best 3 signals, that is, a signal “1”, “2” or “3”. “Good Signal 1-4” and “Good Signal 1-5” are defined analogously. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal. Standard errors are clustered at the participant level.

Table 26: The effect of a “good” signal in the treatment vs control condition.

Treatment	1.956 (2.408)	0.863 (2.716)	9.364* (4.823)	-0.395 (2.258)	-0.690 (2.506)	8.889** (4.340)
Good Signal 1-3	0.037 (2.315)			-0.706 (2.051)		
Treat × Good 1-3	10.233** (4.490)			10.488** (4.142)		
Good Signal 1-4		-0.033 (2.102)			-0.315 (1.892)	
Treat × Good 1-4		9.221** (4.095)			8.027** (3.831)	
Signal Value			-0.918* (0.487)			-0.135 (0.394)
Treat × Signal			-0.961 (0.916)			-1.274 (0.796)
N	656	656	656	864	864	864

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the deviation from the Bayesian benchmark (the number of points allocated to Box 2 minus the number of points that should be allocated given one’s prior beliefs). “Good Signal 1-3” is an indicator variable taking value 1 if a subject received one of the best 3 signals, that is, a signal “1”, “2” or “3”. “Good Signal 1-4” and “Good Signal 1-5” are defined analogously. “Treatment” is an indicator variable taking value 1 if a subject was assigned to the treatment condition. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution.

C.5 First vs last decisions in the control

Table 27: The effect of a good signal on beliefs in the control condition.

<i>Decisions:</i>	<i>1-5</i>	<i>6-10</i>	<i>1-5</i>	<i>6-10</i>
Good Signal	0.625 (3.862)	-0.954 (3.746)	-3.157 (3.218)	-0.700 (3.330)
Bayes	0.940*** (0.103)	0.803*** (0.103)	0.709*** (0.051)	0.729*** (0.057)
N	253	230	351	301

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

Standard errors clustered at the individual level. Their values in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best three signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first two columns are based a sample of subjects who assigned non-zero prior probability to the signal. In the last two columns, I also include participants who received a signal adjacent to their prior belief distribution. Results in the columns 1 and 3 are based on the first 5 decisions in the control condition. Results in the columns 2 and 4 are based on the last 5 decisions in the control condition.

Table 28: The effect of signal valence in the treatment condition.

<i>Decisions:</i>	<i>1-5</i>	<i>6-10</i>	<i>1-5</i>	<i>6-10</i>
Treatment	2.668 (2.851)	1.444 (2.667)	0.530 (2.589)	1.137 (2.365)
Good Signal	0.547 (3.852)	-0.954 (3.746)	-3.038 (3.194)	-0.522 (3.311)
Treat \times Good	9.456* (5.102)	10.667** (4.969)	11.665*** (4.379)	9.219** (4.476)
Bayes	0.925*** (0.070)	0.843*** (0.072)	0.744*** (0.038)	0.759*** (0.041)
Observations	426	403	563	513

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable that takes the value 1 if a subject received one of the best three signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Results in the first two columns are based on a sample of subjects who assigned non-zero prior probability to the signal. In the last two columns, I also include participants who received a signal adjacent to their prior belief distribution. Results in columns one and three are based on the first 5 decisions in the control condition (and all observations from the treatment). Results in columns two and four are based on the last 5 decisions in the control condition (and all observations from the treatment).

C.6 Guessing a random number

Table 29: The effect of a “good” signal in the treatment condition.

Good Signal 1-3	10.723*			9.268*		
	(6.388)			(5.561)		
Good Signal 1-4		9.794*			8.880*	
		(5.856)			(5.201)	
Signal Value			-2.372**			-2.165**
			(1.162)			(0.990)
Observations	68	68	68	91	91	91

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” (“Good Signal 1-4”) is an indicator variable taking value 1 if a subject received one of the best 3 (4) signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. In every specification, I control for the Bayesian benchmark. The sample includes only participants in the treatment condition who saw a random number. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal.

Table 30: The effect of a “good” signal in the control condition.

Good Signal 1-3	-0.525			-1.762		
	(3.246)			(2.695)		
Good Signal 1-4		-0.040			0.044	
		(2.901)			(2.559)	
Signal Value			-0.891			-0.543
			(0.792)			(0.595)
Observations	431	431	431	584	584	584

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” (“Good Signal 1-4”) is an indicator variable taking value 1 if a subject received one of the best 3 (4) signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. In every specification, I control for the Bayesian benchmark. The sample includes only participants in the control condition who were guessing a number different than their rank. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal.

Table 31: The effect of a “good” signal in the treatment vs control condition.

Treatment	1.815 (3.527)	1.297 (3.865)	11.126 (7.756)	1.989 (3.254)	2.071 (3.530)	13.196** (6.136)
Good Signal 1-3	-0.550 (3.240)			-1.794 (2.686)		
Treat × Good 1-3	10.921 (6.767)			11.131* (5.668)		
Good Signal 1-4		-0.040 (2.898)			0.040 (2.556)	
Treat × Good 1-4		8.876 (6.281)			8.373 (5.495)	
Signal Value			-0.897 (0.790)			-0.554 (0.593)
Treat × Signal			-1.175 (1.390)			-1.496 (1.077)
N	499	499	499	675	675	675

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” (“Good Signal 1-4”) is an indicator variable taking value 1 if a subject received one of the best 3 (4) signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. “Treatment” is an indicator variable taking value 1 if a subject was assigned to the treatment condition.

In every specification, I control for the Bayesian benchmark. The sample includes only participants in the control condition who were guessing a number different than their rank. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution.

C.7 Gender Effects

Table 32: Differences between men and women.

	Men	Women		Diff < 0	Diff ≠ 0	Diff > 0
IQ test score	5.323	5.250	<i>p-value:</i>	0.572	0.856	0.428
Rank	5.517	5.602	<i>p-value:</i>	0.392	0.784	0.607
<i>Measures of Belief Distribution:</i>						
Mean Belief	4.508	4.716	<i>p-value:</i>	0.137	0.274	0.863
1 st Quartile	3.752	4.009	<i>p-value:</i>	0.090	0.179	0.910
Median Belief	4.497	4.704	<i>p-value:</i>	0.144	0.289	0.856
3 st Quartile	5.233	5.398	<i>p-value:</i>	0.205	0.410	0.795
Range	4.986	5.102	<i>p-value:</i>	0.258	0.516	0.742
N	294	108				

Table 33: Decisions about signals with non-zero prior probability.

	Men	Women		Diff < 0	Diff ≠ 0	Diff > 0
Decision Treatment	56.146 (2.491)	56.060 (3.563)	<i>p-value:</i>	0.508	0.984	0.492
N	123	50				
Decision Control	49.905 (1.501)	51.198 (2.486)	<i>p-value:</i>	0.329	0.659	0.671
N	357	126				

Note: Decision Treatment (Control) denotes the number of points allocated to Box 2 after observing a signal (considering a signal) in the treatment (control) condition. Standard errors in parentheses.

Table 34: Decisions about signals with zero prior probability.

	Men	Women		Diff < 0	Diff ≠ 0	Diff > 0
Decision Treatment	12.740 (2.190)	20.438 (5.538)	<i>p-value:</i>	0.061	0.123	0.939
N	96	32				
Decision Control	11.214 (1.057)	12.388 (1.949)	<i>p-value:</i>	0.292	0.583	0.709
N	393	134				

Note: Decision Treatment (Control) denotes the number of points allocated to Box 2 after observing a signal (considering a signal) in the treatment (control) condition. Standard errors in parentheses.

Table 35: The effect of a “good” signal in the treatment condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Good Signal	9.801*** (3.605)	9.381** (4.336)	9.947* (5.466)	8.777*** (3.252)	7.684* (3.902)	9.709** (4.763)
Female	3.866 (3.617)	3.446 (4.347)	4.501 (4.466)	3.698 (3.279)	2.567 (3.966)	1.919 (4.031)
Female × Good		1.369 (7.811)	-0.015 (7.934)		3.593 (7.059)	3.635 (7.256)
Bayes	0.915*** (0.092)	0.915*** (0.092)	0.961*** (0.098)	0.812*** (0.057)	0.812*** (0.058)	0.812*** (0.060)
N	173	173	173	212	212	212

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Specification (1) includes dummy for gender (1 if female). Specification (2) includes the gender dummy and its interaction with a good signal. In Specification (3), I add a control for a participant’s rank. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution.

Table 36: The effect of a “good” signal in the control condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Good Signal	-0.260 (3.259)	-1.478 (3.797)	-1.829 (3.791)	-1.924 (2.726)	-1.950 (3.250)	-1.974 (3.220)
Female	1.695 (3.035)	0.369 (3.555)	-0.512 (3.729)	0.281 (2.551)	0.251 (3.091)	-0.445 (3.333)
Female \times Good		4.834 (7.053)	4.959 (7.311)		0.101 (5.687)	0.136 (5.755)
Bayes	0.869*** (0.073)	0.871*** (0.072)	0.865*** (0.074)	0.718*** (0.040)	0.718*** (0.040)	0.716*** (0.041)
N	483	483	483	652	652	652

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Specification (1) includes dummy for gender (1 if female). Specification (2) includes the gender dummy and its interaction with a good signal. In Specification (3), I add a control for a participant’s rank. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution. Standard errors clustered at the participant level.

Table 37: The effect of a “good” signal in the treatment vs control condition.

	(1)	(2)	(3)	(1)	(2)	(3)
Treatment	2.086 (2.399)	1.578 (2.701)	1.292 (2.821)	0.818 (2.164)	-0.034 (2.429)	0.254 (2.540)
Good Signal	-0.224 (2.309)	-1.204 (2.572)	-1.461 (2.673)	-1.832 (1.966)	-2.118 (2.204)	-1.887 (2.283)
Treat × Good	9.969** (4.473)	9.711** (4.485)	10.732** (5.331)	10.389*** (3.962)	10.244** (3.971)	9.248* (4.722)
Female	2.253 (2.016)	0.642 (2.669)	0.364 (2.784)	1.098 (1.762)	-0.038 (2.339)	0.227 (2.437)
Female × Good		3.859 (4.477)	4.874 (5.315)		1.104 (3.847)	0.206 (4.486)
Female × Treat		1.887 (4.517)	2.931 (5.393)		3.138 (4.038)	2.083 (4.862)
Female × Good × Treat			-3.508 (9.880)			3.406 (8.736)
Bayes	0.880*** (0.049)	0.882*** (0.049)	0.883*** (0.049)	0.739*** (0.029)	0.739*** (0.029)	0.739*** (0.029)
N	656	656	656	864	864	864

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal” is an indicator variable taking value 1 if a subject received one of the best 3 signals. “Bayes” is the number of points that should be allocated given one’s prior beliefs. Specification (1) includes dummy for gender (1 if female). Specification (2) includes the gender dummy and its interaction with a good signal and the treatment dummy. In Specification (3), I add a triple interaction of gender, treatment, and good signal. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution.

D Consistency

In this section, I examine the data from the second belief elicitation. First, I look at changes in participants' beliefs and show that beliefs about the box did translate to beliefs about the rank reported in the final belief elicitation. I consider the number of points allocated to the rank corresponding to the received signal. On average, subjects allocated 32.86 points to the relevant rank in Belief Elicitation II. This value is 11.36 points higher than the number of points allocated to the same rank in Belief Elicitation I.

Figure 29 shows how the difference in the number of allocated points depends on the signal realization (a similar graph for the augmented sample is presented in Appendix B.4). For example, 16 subjects who received a signal “2” allocated, on average, 15 points more to Rank 2 in the second belief elicitation. Therefore, they revealed a 15 pp higher probability that “2” is their rank. One can notice that the change in beliefs depends on the signal value. The average change in beliefs is 80% higher after signals 1 to 4, compared to signals 5 to 10 (p-value of one-tailed t-test = 0.0097). The difference remains positive and significant if I control for prior beliefs or the Bayesian benchmark.

I examine the relation between the decisions about the box and the posterior about the rank using regression analysis. The dependent variable is the number of points allocated to the relevant rank in Belief Elicitation II. I regress this value on two independent

Figure 29: Changes in the number of points allocated to the rank = signal.




Table 38: The effect of the decision about the box on the posterior belief.

	(1)	(2)
Points Bayes	0.473*** (0.087)	0.480*** (0.089)
Points Box 2	0.391*** (0.064)	0.385*** (0.065)
Good Signal		1.280 (3.127)
Constant	-20.963*** (4.927)	-21.501*** (5.110)
N	173	173

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to the rank corresponding to the received signal. “Points Bayes” is the number of points that should be allocated based on the prior belief and the signal. “Points Box 2” is the number of points allocated to Box 2 (indicative of one’s rank) in the main task. The sample is restricted to subjects who assigned non-zero prior probability to the signal.

variables. “Points Bayes” denotes the number of points that should be allocated based on the prior belief and the signal. It does not involve subjects’ actual decisions about the boxes. In contrast, the independent variable “Points Box 2” is based solely on the decision about the boxes. It denotes the number of points allocated to Box 2. The results are gathered in Table 38. The estimates in the first column show that both variables have a positive and significant effect on the final belief about the rank. In the second specification, I add an independent variable “Good Signal” defined as in the previous section. The coefficient is not significant, meaning that there is no additional effect of a “good” signal beyond the effect it had on the decision about the box.

The results validate the assumption that beliefs about Box 2 reveal subjects’ beliefs about the rank (which might seem obvious due to the signal structure—Box 2 contains only numbers equal to one’s rank). One can also view the second belief elicitation as an additional check that subjects understood the main task.

Additional Comments

Several points should be kept in mind when interpreting the data. In experiments that measure beliefs multiple times, one common problem is people’s desire to be seen as

consistent decision-makers (Falk and Zimmermann, 2017). Despite my best efforts to ensure anonymity and instruct subjects to treat each part of the experiment independently, these motives might influence beliefs revealed in Belief Elicitation II.³⁷ In this case, the results would provide a lower bound on the effect.

Second, while I explained in intuitive terms how one can translate prior beliefs into a posterior about the box, I provided no such guidance on how to translate the prior distribution and the signal into the posterior distribution about the rank, nor did I explain how to arrive at the posterior distribution given one's beliefs about the box. One implication is that I no longer have control over what people think to be the right course of action. For this reason, one should not expect the results to perfectly align with the Bayesian posterior or the decision about the box. Nevertheless, if the decision about the box revealed the subjects' actual beliefs, it should be positively correlated with the posterior about the rank. This is exactly what I observe in the data.

³⁷This concern is alleviated in the main analysis, for two reasons. First, I elicited posterior beliefs in a different way: by asking about the box. Second, the analysis is based on a comparison between the treatment and the control condition, and there is no reason to suspect that consistency motives operate differently in the two conditions.

E Literature: Design Comparison

The experiment design developed in this paper differs from the designs used in the literature. My main goal was to develop an updating task that induces a strong emotional reaction to a signal. I compare it to the experiments conducted in the past in Table 39. In this review, I focus on papers that study updating about ego-relevant characteristics and do so by asking subjects to update their beliefs about their *relative* performance.³⁸ The papers gathered in the first column in Table 39 are categorized based on one of the relevant design features. In the second column, I describe the design used in my experiment. The last column presents the rationale behind choosing this particular feature for my work.

One design feature that requires an additional comment is the information structure. In almost all of the work reviewed in this section, the information structure follows the scheme presented in Figure 30.³⁹ There are two states of the world H and L indicating whether one's score was in the upper or the lower half of the test score distribution. Subjects receive a signal that is informative about the state with known precision, e.g., 75% in Möbius et al. (2022). However, this signal structure becomes more complicated if extended to a larger signal and state space (see Figure 31) and I am not aware of any experimental work that implements it. Papers that used 10 states of the world, Eil and Rao (2011) and Zimmermann (2020), use binary signals shown in Figure 32. A signal informs a subject whether or not he ranked higher than another participant who was randomly drawn from a group of 10 (I denote these binary signals with H and L). The precision depends on the state and, for the first signal, takes one of the following values: 55.6%, 66.7%, 77.8%, 88.9% or 100% (for the second signal it is 50%, 62.5%, 75%, 87.5% or 100%, as comparisons are made without replacement).

³⁸For a review of the literature on learning about absolute performance as well as updating about non-ego-relevant parameters, I refer the reader to Barron (2021) and Coutts (2019). An even broader review of the literature on errors in probabilistic reasoning can be found in Benjamin (2019).

³⁹See Table 39 for the references.

Figure 30: Design used in the literature (2 states).

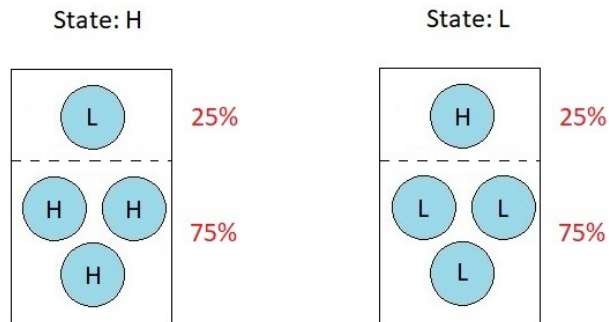


Figure 31: Design used in the literature extended to 10 states.

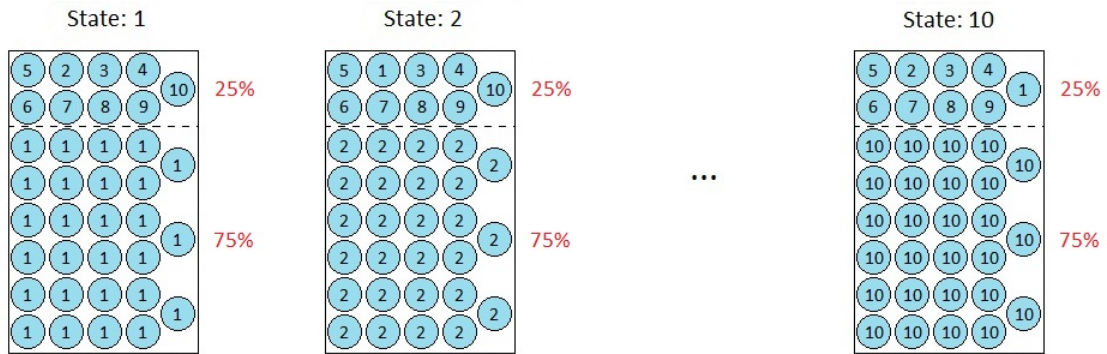



Figure 32: Design used in Eil and Rao (2011) and Zimmermann (2020).



The design commonly used in the literature (Figure 30) extended to 10 states can be simplified by distinguishing two urns: one with balls indicating the state (the “IQ” urn), and one with every possible number (the “Random” urn).⁴⁰ Figures 33 and 34 present the simplified design for 2 and 10 states of the world. The information structure in Figure 33 is equivalent to the one depicted on Figure 30, assuming either urn can be selected with equal probability $\frac{1}{2}$. If the state is H , a ball indicating H is drawn with probability $0.5 \cdot 0.5 + 0.5 \cdot 1 = 0.75$, the same as in Figure 30. Similarly, Figure 34 is equivalent to the information structure in Figure 31 with the signal precision of 55%.

Figure 33: Design developed in this paper (2 states).

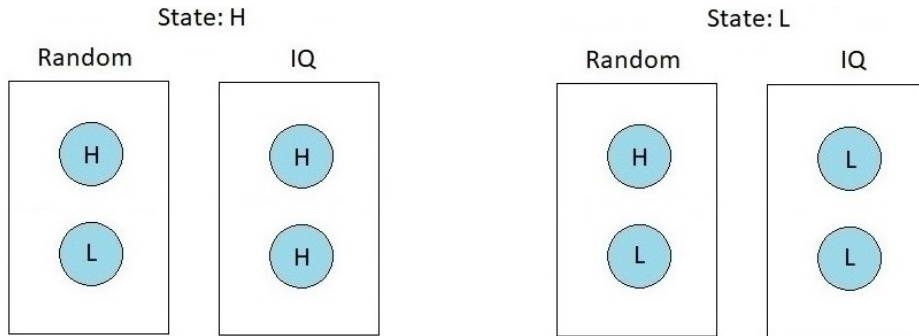



Figure 34: Design developed in this paper (10 states).



⁴⁰One could also distinguish the two urns along the dashed line in Figure 31, with the Random urn containing all numbers except the one that indicates the state. This design, however, lacks the intuitive interpretation of “a random urn” from which *any number* can be drawn with *the same* probability, hence it might be more difficult to explain to the participants.

Table 39: Literature review: design comparison.

Other Work	This Paper	Goal
1. Number of signals:		
<ul style="list-style-type: none"> – more than 1 signal <p>Buser et al., 2018; Coutts, 2019; Drobner and Goerg, 2024; Eil and Rao, 2011; Möbius et al., 2022; Zimmermann, 2020.</p>	<ul style="list-style-type: none"> – 1 signal 	<ul style="list-style-type: none"> – separating reaction to signals from information aggregation.
<ul style="list-style-type: none"> – 1 signal <p>Drobner, 2022; Ertac, 2011; Schwardmann and Van der Weele, 2019.</p>		
2. State space, signal space, signal precision:		
<ul style="list-style-type: none"> – 2 states (above or below 50%; above or below 85% in Coutts, 2019), – 2 signal values, – signal precision: 67% <p>Coutts, 2019; Drobner, 2022; Drobner and Goerg, 2024.</p>	<ul style="list-style-type: none"> – 10 states (deciles of the distribution) – 10 signal values 	<ul style="list-style-type: none"> – richer state space and signal space to induce a stronger emotional reaction to a signal (based on the observation that it is more painful for subjects to be in the bottom 10% than in the bottom 50%).
<ul style="list-style-type: none"> – 2 states (above or below 50%) – 2 signal values – signal precision: 70% <p>Buser et al., 2018.</p>		
<ul style="list-style-type: none"> – 2 states (above or below 50%) – 2 signal values – signal precision: 75% <p>Möbius et al., 2022; Schwardmann and Van der Weele, 2019.</p>		
<ul style="list-style-type: none"> – 3 states (lower 20%, middle 60%, or upper 20%) – 2 signal values – perfectly informative but coarse signals <p>Ertac, 2011.</p>		
<ul style="list-style-type: none"> – 10 states (deciles of the distribution) – 2 signal values – signal precision depends on the state: 56%, 67%, 78%, 89% or 100%. <p>Eil and Rao, 2011; Zimmermann, 2020.</p>		

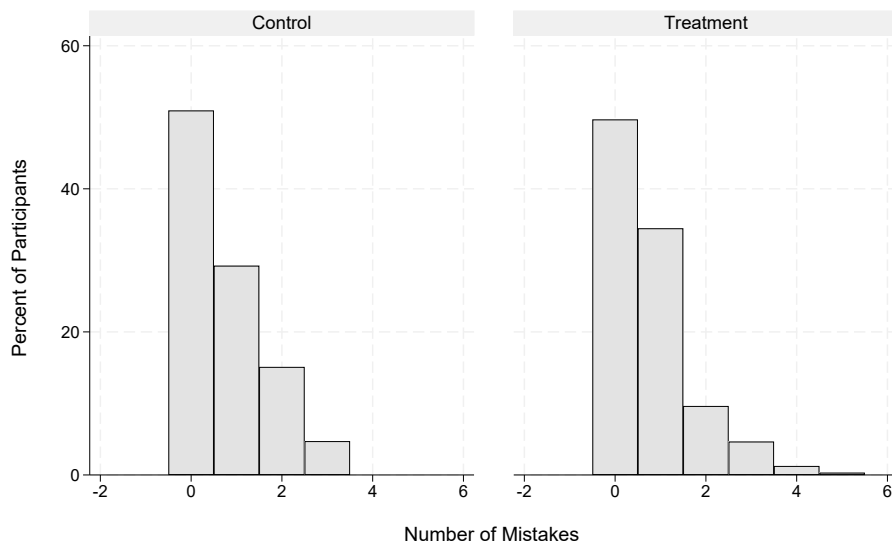
Other Work	This Paper	Goal
3. Information structure and implementation:		
<ul style="list-style-type: none"> – information structure as in Figure 30 – a signal is true or false with precision known to the subjects <p>Buser et al., 2018; Coutts, 2019; Drobner and Goerg, 2024; Möbius et al., 2022; Schwardmann and Van der Weele, 2019. Drobner, 2022, uses the same information structure (Figure 30), but the signal is a comparison with another subject.</p>	<ul style="list-style-type: none"> – information structure shown in Figure 34 – it is equivalent to Figure 31 with a signal precision of 55%. 	<ul style="list-style-type: none"> – it would not be possible to introduce richer state and signal space using any other information structure from the literature.
<ul style="list-style-type: none"> – information structure as in Figure 32 – a signal is a pairwise comparison with another subject <p>Eil and Rao, 2011; Zimmermann, 2020.</p>		
<ul style="list-style-type: none"> – a signal is always true, but only reveals whether the subject is in the top or the bottom half of the distribution, and not precisely the state <p>Ertac, 2011.</p>		
4. Comparison group:		
<ul style="list-style-type: none"> – a group of 4 <p>Drobner, 2022; Schwardmann and Van der Weele, 2019.</p>	<ul style="list-style-type: none"> – 300 subjects 	<ul style="list-style-type: none"> – a larger comparison group makes it more difficult to use reappraisal to lessen the impact of the negative signal (e.g., in the case of a group of four, one can easily attribute a negative signal to being assigned to a particularly strong pair of subjects). When there is another way of “explaining” a bad signal, there may be no need for (costly) belief distortion.
<ul style="list-style-type: none"> – a group of 8 <p>Buser et al., 2018.</p>		
<ul style="list-style-type: none"> – a group of 10 <p>Eil and Rao, 2011; Ertac, 2011; Zimmermann, 2020.</p>		
<ul style="list-style-type: none"> – a group larger than 10 <p>Coutts, 2019; Drobner and Goerg, 2024; Möbius et al., 2022.</p>		
5. Timing of revealing information:		
<ul style="list-style-type: none"> – In most of the papers mentioned above it is unclear whether and when subjects expected the resolution of uncertainty (see Drobner, 2022, for a comprehensive literature review). This problem was noticed and tested in the recent work by Drobner (2022). 	<ul style="list-style-type: none"> – online access one week after the session 	<ul style="list-style-type: none"> – to describe the behavior with a one-period model without the dynamic concerns – to bring the design closer to the real-world situations: grades are rarely immediate, need to be checked etc.

F Including Confused Individuals

F.1 Mistakes in control questions

The experimental tasks required a good understanding of the instructions. For this reason, before the main task, participants had to solve a set of control questions. While I allowed confused participants to finish the experiment, I collected the data on the number of mistakes and removed the most mistaken subjects from the main analysis. In the end, I excluded 25 participants who made three or more mistakes in five control questions.⁴¹ Additionally, one participant reported to the assistant after the session that he mixed up the two boxes. His cubicle number was noted and this observation was removed from the analysis (the participant with a number 340).

Figure 35: Mistakes in the five control questions in the two conditions.



⁴¹One of the control questions (Question 2) was slightly different in the two conditions, and participants in the control condition made more mistakes in their version. I did not include this question in the measure of subjects' mistakes.

F.2 Results based on a sample including confused participants

Table 40: The effect of a “good” signal in the treatment condition.

Good Signal 1-3	9.625*** (3.575)			8.598*** (3.186)		
Good Signal 1-4		8.201** (3.268)			7.532** (2.952)	
Signal Value			-1.760** (0.712)			-1.527** (0.607)
N	183	183	183	224	224	224

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” (“Good Signal 1-4”) is an indicator variable taking value 1 if a subject received one of the best 3 (4) signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. In every specification, I control for the Bayesian benchmark. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal.

Table 41: The effect of a “good” signal in the control condition.

Good Signal 1-3	-0.140 (3.139)			-1.586 (2.626)		
Good Signal 1-4		-0.599 (2.863)			-0.567 (2.488)	
Signal Value			-0.853 (0.780)			-0.474 (0.569)
N	502	502	502	680	680	680

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” (“Good Signal 1-4”) is an indicator variable taking value 1 if a subject received one of the best 3 (4) signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. In every specification, I control for the Bayesian benchmark. Results in the first three columns are based on subjects who assigned non-zero prior to the signal. Results in the last three columns are based on an augmented sample of participants whose priors were not far from the signal.

Table 42: The effect of a “good” signal in the treatment vs control condition.

Treatment	2.077 (2.434)	0.966 (2.704)	9.115* (5.520)	0.695 (2.191)	0.005 (2.412)	9.119** (4.401)
Good Signal 1-3	-0.126 (3.120)			-1.483 (2.605)		
Treat × Good 1-3	9.715** (4.484)			9.731** (3.891)		
Good Signal 1-4		-0.599 (2.856)			-0.570 (2.482)	
Treat × Good 1-4		8.758** (4.346)			8.355** (3.839)	
Signal Value			-0.854 (0.779)			-0.454 (0.567)
Treat × Signal			-0.896 (1.045)			-1.137 (0.793)
N	685	685	685	904	904	904

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors in parentheses.

The dependent variable is the number of points allocated to Box 2. “Good Signal 1-3” (“Good Signal 1-4”) is an indicator variable taking value 1 if a subject received one of the best 3 (4) signals. “Signal Value” refers to the received signal. It takes values from 1 to 10, with higher values indicating worse signals. “Treatment” is an indicator variable taking value 1 if a subject was assigned to the treatment condition.

In every specification, I control for the Bayesian benchmark. Results in the first three columns are based on observations regarding signals to which participants assigned non-zero prior probability. Results in the last three columns are based on an augmented sample: observations regarding signals that were not far from the prior belief distribution.

G Questionnaires

Table 43: The effect of other variables on decisions in the treatment condition.

	<i>Non-zero prior</i>			<i>Prior equal to zero</i>		
	(1)	(2)	(3)	(1)	(2)	(3)
Extroversion	-0.404 (0.47)	-0.405 (0.48)	-0.587 (0.48)	-0.095 (0.51)	0.109 (0.51)	0.352 (0.55)
Conscientiousness	-0.247 (0.58)	-0.049 (0.60)	-0.139 (0.60)	0.348 (0.65)	0.658 (0.66)	0.879 (0.68)
Openness	-0.225 (0.49)	-0.335 (0.50)	-0.521 (0.50)	1.022 (0.64)	0.889 (0.64)	0.885 (0.64)
Neuroticism	-0.361 (0.42)	-0.482 (0.61)	-0.478 (0.62)	0.755 (0.56)	0.002 (0.72)	-0.001 (0.73)
Agreeableness	0.292 (0.63)	0.244 (0.65)	0.123 (0.66)	-0.141 (0.72)	-0.016 (0.71)	0.108 (0.72)
Anxiety State		0.318 (0.24)	0.338 (0.24)		-0.636** (0.31)	-0.605* (0.31)
Anxiety Trait		-0.349 (0.27)	-0.422 (0.27)		-0.088 (0.34)	-0.060 (0.35)
Reappraisal			3.672** (1.81)			-2.493 (2.24)
Suppression			-3.110 (2.03)			2.590 (2.54)
N	166	166	166	124	124	124

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. The sample includes only participants in the treatment condition. Estimates in the first three columns are based on a sample of participants who assigned a non-zero prior probability to the signal received. The number of observations (166) is different from our baseline sample (173), because in one session, the data from the final questionnaire was lost (due to a human error). In the last three columns, the regressions are based on the remaining observations—participants who assigned a zero prior to the rank corresponding to the received signal. In the first three columns, I control for the Bayesian benchmark. In the last three, this control is omitted as the Bayesian benchmark is undefined.

Table 44: The effect of signal valence in the treatment condition.

	<i>Non-zero prior</i>		<i>Prior equal to zero</i>	
	(1)	(2)	(1)	(2)
Emotion: Enjoyment	-1.295 (1.29)	-1.677 (1.30)	-0.325 (1.44)	-0.430 (1.44)
Emotion: Hope	1.051 (1.43)	0.917 (1.42)	-1.328 (1.98)	-1.197 (2.01)
Emotion: Pride	1.933 (1.60)	2.030 (1.59)	0.190 (2.28)	0.180 (2.31)
Emotion: Relief	0.440 (1.72)	0.320 (1.72)	0.229 (2.24)	0.170 (2.25)
Emotion: Anger	-0.633 (1.26)	-0.844 (1.26)	-2.222 (1.50)	-2.407 (1.51)
Emotion: Anxiety	1.798 (2.02)	2.208 (2.04)	5.393* (2.82)	5.530* (2.83)
Emotion: Shame	0.741 (1.36)	0.618 (1.35)	-2.082 (1.83)	-2.291 (1.88)
Emotion: Hopelessness	-0.709 (2.06)	-0.925 (2.05)	5.697** (2.62)	5.707** (2.62)
Reappraisal		3.162* (1.79)		-2.703 (2.03)
Suppression		-2.415 (1.93)		1.363 (2.31)
N	166	166	124	124

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. The sample includes only participants in the treatment condition. Estimates in the first two columns are based on a sample of participants who assigned a non-zero prior probability to the signal received. The number of observations (166) is different from the baseline sample (173), because in one session, the data from the final questionnaire was lost (due to a human error). In the last three columns, the regressions are based on the remaining observations—participants who assigned a zero prior to the rank corresponding to the received signal. In the first three columns, I control for the Bayesian benchmark. In the last three, this control is omitted as the Bayesian benchmark is undefined.

G.1 Negative Emotions and Reappraisal

Table 45: The effect of emotional state and reappraisal (indicator variables).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Reappraisal	6.055*		5.944*	9.429**		5.353	5.237
	(3.36)		(3.38)	(4.70)		(3.39)	(4.63)
Negative Emotions		-1.776	-1.280	1.621			
		(3.33)	(3.32)	(4.29)			
Negative × Reapp				-7.214			
				(6.76)			
Anxiety State					5.638*	4.858	4.752
					(3.38)	(3.40)	(4.47)
Anxiety State × Reapp							0.252
							(6.83)
N	166	166	166	166	166	166	166

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Table 46: The effect of emotional state and reappraisal (indicator variables).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Reappraisal	7.820**		7.798**	11.924***		7.305**	5.918
	(3.03)		(3.04)	(4.35)		(3.05)	(4.28)
Negative Emotions		-0.884	-0.596	2.662			
		(3.04)	(3.00)	(3.87)			
Negative × Reapp				-8.091			
				(6.09)			
Anxiety State					4.831	3.894	2.726
					(3.01)	(3.01)	(3.92)
Anxiety State × Reapp							2.843
							(6.12)
N	205	205	205	205	205	205	205

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. The sample includes only participants in the treatment condition. Estimates in the first two columns are based on a sample of participants who assigned a non-zero prior probability to the signal received. The number of observations in Table 45 (166) is different from the baseline sample (173), because in one session, the data from the final questionnaire was lost (due to a human error). In Table 46, the regressions are based on an augmented sample.

Table 47: The effect of emotional state and reappraisal (indicator variables).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Reappraisal	1.668 (4.59)		1.768 (4.48)	3.542 (6.45)		3.076 (4.56)	1.044 (7.32)
Negative Emotions		11.326** (4.34)	11.341** (4.35)	12.649** (5.54)			
Negative \times Reapp				-3.450 (9.00)			
Anxiety State					-9.636** (4.42)	-10.050** (4.47)	-11.231** (5.58)
Anxiety State \times Reapp							3.332 (9.38)
N	124	124	124	124	124	124	124

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The dependent variable is the number of points allocated to Box 2. The sample includes only participants in the treatment condition. The regressions are based on a sample of participants who assigned a zero prior probability to the rank corresponding to the received signal.