

Experimental Evidence on Misguided Learning*

Lorenz Götte[†] and Marta Kozakiewicz[‡]

September 11, 2024

Abstract

This paper studies how people form beliefs in environments with multiple unknown parameters, some of which are relevant to agents' self-esteem. In particular, we examine how initial bias in beliefs about an ego-relevant characteristic affects learning about the state of the world. Using data from a laboratory experiment, we demonstrate that the learning process of an overconfident agent is *self-defeating*: the agent repeatedly takes suboptimal actions, misinterprets the output, and forms increasingly mistaken beliefs about the state. Therefore, we corroborate the theory of misguided learning formulated by Heidhues et al. (2018). We provide the first empirical evidence that allowing a biased agent to experiment and acquire new information is not only ineffective but in some cases counterproductive. Furthermore, we move beyond the theory as we examine how learning about multiple parameters evolves in ego-relevant and ego-neutral environments.

Keywords: overconfidence, belief formation, learning, experiment

JEL classification: C91, D83

*We thank Ximeng Fang, Botond Kőszegi, Axel Wogroly and Florian Zimmermann for helpful comments. Richard Franck and Zhi Hao Lim provided excellent research assistance. The authors gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (project A01 and B07).

[†]National University of Singapore

[‡]Frankfurt School of Finance and Management; corresponding author: m.kozakiewicz@fs.de

1 Introduction

Many economic decisions require an accurate assessment of the state of the world. Often, more than one decision-relevant aspect is unobservable, and people have to form beliefs *simultaneously* about multiple parameters. Learning in such environments is particularly challenging. Agents need to not only keep track of actions and their consequences but also disentangle the effects of various factors in order to update beliefs about specific parameters. If an agent holds biased beliefs about one parameter, e.g. overconfident beliefs about his own ability, he is likely to make incorrect inferences from the observed data.

Heidhues et al. (2018) show that, in some cases, the learning process goes awry: the agent repeatedly misinterprets the data, takes suboptimal actions, and forms more and more incorrect beliefs about the state of the world.¹ Learning is “misguided” and, since it is the agent who generates the observations that lead him astray, one can describe it as “self-defeating”. Importantly, even if initially the agent has correct beliefs about *all* other aspects of the world, biased perception of ability can start a process that drives beliefs away from the truth. The theory predicts this pattern in a range of applications: it can explain why overconfident individuals exert too little effort, managers do not delegate enough tasks to subordinates, and CEOs engage in unprofitable mergers. It can also explain why additional feedback doesn’t always help to correct one’s actions. Yet, the extent to which people engage in misguided learning has not been examined.

In this paper, we use data from a carefully designed laboratory experiment to provide the first empirical evidence on misguided learning.² We test the comparative statics of the model by Heidhues et al. (2018) and document the learning processes of biased agents. Secondly, we move beyond the model as we investigate how learning about multiple parameters evolves in ego-relevant and ego-neutral environments.

¹An illustrative example considers an overconfident agent who is learning about the state of the world by taking actions and observing output, which also depends on his unknown ability. After observing an unexpectedly low output, the agent does not interpret it as a result of his low ability but concludes that the state must be worse than expected. He adjusts his action to match the new belief about the state. The increased mismatch between the action and the state further lowers the output. This reinforces the agent’s pessimism and leads to an action that, in reality, fits even worse. Over time, the agent takes more and more inadequate actions and becomes increasingly mistaken about the state.

²A laboratory setting is particularly suitable for studying misguided learning. Firstly, it enables us to elicit beliefs multiple times in an incentive-compatible way, providing a precise measure of overconfidence and changes in beliefs (rarely observable in the field). Secondly, it gives us tight control over technology and information available to subjects. In the field setting, technology is usually unobservable, and without the knowledge of the output generating process, one cannot formulate the model’s predictions.

Our experiment integrates all features of the model in a simple way. The main goal was to create an environment in which subjects take actions, observe output and learn about the underlying state of the world. Importantly, the output also depends on an unknown parameter that is relevant to subjects' self-esteem. For this purpose, we used subjects' relative performance in an IQ test taken in the first part of the study.³ Before the main task, we also elicited participants' beliefs about their relative performance.

In the second part of the experiment, participants completed several rounds of a learning exercise. In every round, the task was to estimate an unknown state of the world: a randomly drawn integer between -10 and 10 . Participants had 4 trials to guess the state and were remunerated based on the accuracy of their guesses. After making a guess, each participant received feedback in the form of a real number between 4 and 51 displayed on the individual computer screen. Feedback was determined by the state of the world and one's relative performance in the IQ test. In every trial, the optimal strategy was to enter one's best guess about the state of the world. Therefore, we could directly track participants' belief formation process. After the learning exercise, we again elicited subjects' beliefs about their relative performance.

To help participants correctly interpret the feedback, we provided them with tables to look up which states of the world and relative performances are consistent with the feedback they observed. We did not preclude subjects from considering different performance levels and they were free to choose any combination of the two parameters. Giving subjects the opportunity to reconsider their beliefs about ability allows us to see if misguided learning emerges even in environments where beliefs are not restricted.

We introduced two experimental conditions: treatment and control. In the treatment condition, participants received informative feedback based on their last guess, whereas in the control condition, feedback did not depend on subjects' guesses. Thus, the main mechanism of the model – the interdependence between actions and feedback – was shut down in the control condition. We kept all other features of the experiment unchanged: participants were asked to make four guesses and after each one, a number was displayed

³We decided to use intelligence as an input to the production function for several reasons. Firstly, intelligence is known as a personal characteristic that people deeply care about, so a measure of IQ seems to be a good candidate for a genuinely ego-relevant parameter. Secondly, the literature provides evidence that people have biased assessments of their relative cognitive ability (see, for example, Burks et al., 2013). One would expect misguided learning to arise in this context. Last but not least, cognitive ability as a component of human capital is an actual input to many production functions.

on their computer screens. We use the control condition to exclude alternative explanations, for example, that the effect is an artifact of repeated choice-based elicitation. Rather, we show that it is induced by informative feedback as in Heidhues et al. (2018). The experiment was conducted in November 2017 in BonnEconLab at the University of Bonn. We collected data from 171 subjects, mostly university students.

Furthermore, we designed an ego-neutral condition in which output depends on a parameter that does not affect agents' self-esteem. In this condition, subjects performed the task based on the performance of some other, randomly selected individual who reported similar beliefs (we assume that the performance of another subject is irrelevant to one's ego).⁴ The structure of the experiment was identical to that of the main condition, and it allows us to isolate the effect of ego-relevance of one of the parameters. The data from 155 participants in the ego-neutral condition was collected in November 2018.

Overall, we find strong support for the predictions of the misguided learning model. When overconfident individuals can adjust their actions and learn about the state of the world, repeated feedback leads them to form increasingly mistaken beliefs. Their learning process is self-defeating: overconfident participants tend to attribute an unsatisfactory outcome to the realized state instead of their relative performance, take suboptimal actions, and become pessimistic about the state over time. Importantly, we find a significant difference between the treatment and the control condition, confirming that the effect is driven by the mechanism described in Heidhues et al. (2018).

The effect is more pronounced for participants who are more biased about their ability. We test the model's comparative statics and show that the more overconfident the participant is, the more mistaken about the state he becomes. We also corroborate the qualitative predictions for underconfident and unbiased subjects. The learning process of the underconfident agent is misdirected and its trajectory is different from that of the overconfident agent.⁵ We do not detect any similar pattern in the behavior of unbiased participants. In line with the model, the unbiased subjects immediately learn the true state and take the optimal action in the following trials.

⁴After eliciting beliefs about subjects' relative performance, participants were informed that they will be randomly matched to a person from one of the previous sessions, who took the same IQ test and reported the same beliefs but not necessarily obtained the same IQ test score. Before the main task, we elicited subjects' beliefs about the relative performance of the person matched to them.

⁵We use the term "misguided learning" to describe the learning of biased (over- or underconfident) agents, and we only refer to the overconfident agents' misguided learning as "self-defeating learning".

However, the effects observed in the data are quantitatively lower than the theory predicts. The gap between the theoretical predictions and the observed behavior is caused by some participants updating their beliefs about ability during the experiment. We observe a significant difference in beliefs measured before and after the learning exercise, with 25% of participants revealing unbiased beliefs after the task (compared to 7.6% before the task). Notwithstanding, almost 80% of subjects who were classified as overconfident at the beginning of the experiment remained overconfident, and many of them were engaging in self-defeating learning until the end of the last round.

Using data from the ego-neutral condition, we show that self-defeating learning is more likely to arise and persist when one’s ego is at stake. When the output is based on the IQ test performance of some other, randomly selected individual, misdirected learning of overconfident agents is significantly mitigated. Overconfident participants in the ego-neutral condition are engaging in self-defeating learning to a lesser extent partly because they are willing to revise downwards their beliefs about the ability of their match.⁶

The opposite is true for underconfident agents. Underconfident participants are *more* likely to become unbiased in the ego-relevant condition, that is, when learning about their own ability, compared to similarly underconfident subjects in the ego-neutral condition. The results demonstrate that, when learning involves multiple parameters and some of them are ego-relevant, people are steered to learn along the dimension that brings them higher ego utility. While motivated attribution of ego-relevant outcome is a phenomenon well-known in the psychological literature (see Coutts et al., 2024, for a review of the literature), our paper is the first to demonstrate it in a dynamic setting.

The paper contributes to the growing literature on learning with misspecified models. The burgeoning theoretical literature (see Bohren and Hauser, 2023, for a comprehensive review) has only recently been followed by empirical evidence. One should mention the work on selective attention in learning (Hanna et al., 2014, Graeber, 2023), base-rate neglect (Esponda et al., 2024), redundancy neglect in social learning (Enke and Zimmermann, 2017), overlooking selection problems (Enke, 2020), and misattribution of reference dependence in learning from experience (Gagnon-Bartsch and Bushong, 2022,

⁶Importantly, we are comparing overconfident subjects in the ego-relevant condition with similarly overconfident participants in the ego-neutral control. This allows us to control for any confounding factors and disentangle the effect of agents’ bias from the ego-relevance of the unknown parameter.

Bushong and Gagnon-Bartsch, 2023). Our work adds to this literature by testing the model of misguided learning by Heidhues et al. (2018). A similar problem—learning with biased beliefs about one’s ability—was recently studied by Hestermann and Le Yaouanq (2021). The authors also consider two parameters that are not separately identifiable, but in their model, the agent is learning about both. Unfortunately, their framework is different from our experimental setup, and we cannot directly test the model’s predictions. We contribute to the literature on learning with misspecified models by providing empirical evidence on misguided learning and taking the next step to elucidate how people learn in environments with multiple unknown parameters.

Our work is partially motivated by the behavioral literature on motivated reasoning, which suggests that people might interpret feedback in a self-serving manner (Bénabou and Tirole, 2016, for a review of the literature). A large body of work demonstrates that people use various strategies to manipulate their beliefs to maintain a positive self-view. These strategies include information avoidance (Golman et al., 2017), selective recall (Chew et al., 2020; Huffman et al., 2022; Zimmermann, 2020), and asymmetric updating (see, e.g., Eil and Rao, 2011; Möbius et al., 2022).⁷ Our experiment was not designed to test any of these mechanisms directly, but to examine how motivated reasoning unravels in a more complex, dynamic environment.

We view our paper as complementary to the literature investigating the consequences of holding inaccurate beliefs with some degree of persistence. Overconfidence, a widely-studied phenomenon by both psychologists and economists, is believed to generate great costs for both the individual and the society.⁸ We contribute to the literature concerning the implications of overconfidence as we document its detrimental effect on learning.

The paper is organized as follows. In Section 2, we describe a simplified version of the model and its testable predictions. Section 3 outlines our experimental design and Section 4 presents the empirical results. In Section 5, we discuss the results of the ego-neutral condition. Section 6 concludes.

⁷Perhaps the closest to our work, Coutts et al. (2024) test two different theories of self-attribution bias and show that, although people tend to update more favorably about themselves than about their teammates, they do not attribute the negative outcome to the other player.

⁸Negative consequences of overconfidence include excessive selection into competitive environments (Camerer and Lovo, 1999; Niederle and Vesterlund, 2007), excessive trading (Barber and Odean, 2001), suboptimal investment decisions (Malmendier and Tate, 2005, 2008), and political polarization (Ortoleva and Snowberg, 2015).

2 Theoretical Framework

In this section, we present an example based on a simplified version of the model by Heidhues et al. (2018). Our goal is to provide intuition for the main mechanism, therefore, in the description below, we include only the essential features of the model.⁹ For a complete specification, we refer the reader to the original paper.

2.1 The Model

In each period $t \in \{1, 2, 3, \dots\}$, the agent produces an observable output q_t according to the following production function:

$$q_t = Q(e_t, A, \Phi) = A + \Phi - L(e_t - \Phi), \quad (1)$$

where $e_t \in (\underline{e}, \bar{e})$ denotes the agent's action in period t , $A \in \mathbb{R}$ is the agent's true ability, $\Phi \in (\underline{\phi}, \bar{\phi})$ is the unknown state of the world, and $L(\cdot)$ is a symmetric loss function $L(x) = \frac{1}{3}|x|$. The loss is minimized when the agent matches his action to the state of the world. Let us assume that $\Phi = 0$ and the agent has initially correct beliefs about the state: $\phi_0 = 0$.

In each period, the agent takes an optimal action given his beliefs. To minimize the loss function, he chooses $e^*(\phi) = \phi$. The agent follows a myopic decision rule: the action maximizes the expected output in a given period.^{10,11} In the first period, the optimal action is equal to the agent's initial belief: $e_1^* = \phi_0 = 0$. It produces the following output:

$$q_1 = Q(e_1, A, \Phi) = A - L(0) = A.$$

⁹One way in which we deviated from Heidhues et al. (2018) is that we restricted the number of output realizations per period to one and eliminated the noise component. In the model, the agent observes multiple noisy output realizations in every period and updates his beliefs based on *the average* output in that period (he averages out the random component). We decided not to include this feature, as we were concerned that biases in information aggregation could obscure the results. Moreover, we assume that the agent has degenerate beliefs: he assigns either probability zero or one to the state of the world.

¹⁰The assumption implies that there is no learning motive at play. The agent is neither intentionally experimenting nor gathering data about his environment to make better choices in the future. Misguided learning is a by-product of a sequential, short-sighted optimization.

¹¹We decided not to impose this assumption onto participants in our experiment. However, we expected that the task will induce short-sighted behavior to some extent.

The agent observes the output q_1 and compares it to the output that he expected. The difference between the observed and the expected output depends on the direction and magnitude of the agent's bias.

2.2 Overconfidence and Self-Defeating Learning

An overconfident agent believes that his ability is $\tilde{a} > A$ (it is higher than his actual ability A). After taking an action $e_1^* = \phi_0 = 0$, he expects to observe the output \tilde{q}_1 :

$$\tilde{q}_1 = Q(e_1, \tilde{a}, \phi_0) = \tilde{a}.$$

Importantly, the agent never updates his beliefs about his ability (we discuss this assumption in Section 2.4). He attributes the difference between q_1 and \tilde{q}_1 to the state of the world. Knowing the output and his previous action, the agent derives ϕ from (1). He concludes that the state is *worse* than he thought: $\phi_1 = -\frac{3}{4}(\tilde{a} - A) < \phi_0$.

In Period 2, the agent chooses $e_2^* = \phi_1$. He observes the output $q_2 = A - L(\phi_1) = A - \frac{1}{4}(\tilde{a} - A)$, while he expected to produce $\tilde{q}_2 = \tilde{a} - \frac{3}{4}(\tilde{a} - A)$. Once again, he is surprised by the output and attributes the difference to the state of the world, forming a new belief $\phi_2 = -\frac{9}{8}(\tilde{a} - A)$. As a result, he becomes even more pessimistic about the state:

$$\phi_2 < \phi_1 < \phi_0. \tag{2}$$

With each observation, the agent's beliefs are driven further away from the true state. Hypothesis 1.OC summarizes the learning process of an overconfident agent:

Hypothesis 1.OC (Overconfident Agents)

The learning process of an overconfident agent is self-defeating: in each period, after taking an action and observing the output, the agent forms increasingly pessimistic beliefs about the state.

The change in beliefs in each each period depends on the difference between q_1 and \tilde{q}_1 ,

which is in turn a function of agent's bias ($\tilde{a} - A$). An overconfident agent with a larger bias has higher output expectations relative to a less biased individual. He will be more surprised by the actual output and will attribute this larger difference to the state of the world. As a result, he will form more biased beliefs about the state compared to a less overconfident individual.

Hypothesis 2.OC (Overconfident Agents)

An overconfident agent with a larger bias will form more pessimistic beliefs compared to a less overconfident agent, and will end up further away from the true state.

Heidhues et al. (2018) shows that the agent's belief about the state is not decreasing indefinitely but converges to a unique limiting belief ϕ_∞ . This belief induces action and output that exactly matches the agent's expectations, thereby confirming his beliefs. It is stable in the sense that the agent has no incentive to abandon it: his expectations are confirmed, bringing the learning process to an end. It could be found by setting the difference between the actual and the expected outputs to zero: $Q(e^*(\phi_\infty), A, \Phi) - Q(e^*(\phi_\infty), \tilde{a}, \phi_\infty) = 0$. With the loss-function specification, that condition reads:

$$(A - \tilde{a}) + (\Phi - \phi_\infty) - L(\Phi - \phi_\infty) = 0. \tag{3}$$

By rearranging the above equation one can derive a formula for the stable belief ϕ_∞ .¹² It is worth noting that the stable belief is a function of the agent's bias.

2.3 Underconfident and Unbiased Agents

The model also predicts the behavior of underconfident agents. The analysis is analogous, with the only difference that the agent underestimates his true ability, i.e. $\tilde{a} < A$. In Period 1, the agent chooses $e_1^* = \phi_0 = 0$. He observes the output of $q_1 = A$, while he expected to produce \tilde{a} . The agent does not update his beliefs about his ability, but instead he looks for a state that would explain the output – the agent solves (1) for ϕ , given the observed output and his previous action. Since $\tilde{a} < A$, the updated belief ϕ_1 is

¹²In our simple example, the stable belief of an overconfident agent is equal to $\phi_\infty = -\frac{3}{2}(\tilde{a} - A)$.

higher than the underconfident agent's prior: $\phi_1 = -\frac{3}{2}(\tilde{a} - A) > 0$. The agent concludes that the state of the world is *better* than he expected.

In Period 2, the agent chooses $e_2^* = \phi_1$. He observes the output of $q_2 = A + \frac{1}{2}(\tilde{a} - A)$, while he expected to produce $\tilde{q}_2 = \tilde{a} - \frac{3}{2}(\tilde{a} - A)$. The output falls short of his expectations, so he concludes that the state is *worse* than he thought: $\phi_2 = -\frac{3}{4}(\tilde{a} - A)$. The adjustment in the following period goes in the right direction, bringing the agent closer to the true state. In contrast to the overconfident agent, the underconfident agent's misguided learning is *self-correcting*. The model predicts that the underconfident agent's beliefs satisfy:

$$\phi_1 > \phi_0 \quad \wedge \quad \phi_2 < \phi_1. \tag{4}$$

This allows us to formulate the following prediction about the belief-formation process of underconfident agents:

Hypothesis 1.UC (Underconfident Agents)

The learning process of an underconfident agent is self-correcting: after the initial overly optimistic assessment, the agent corrects his beliefs downwards.

In the first period, an underconfident agent with a larger bias is more surprised by the output than an underconfident agent with a smaller bias. Because the agent attributes the entire difference to the state of the world, he becomes more mistaken about the state compared to the less biased individual.

Hypothesis 2.UC (Underconfident Agents)

In the first period, an underconfident agent with a larger bias forms beliefs that are further away from the true state compared to the beliefs of a less underconfident individual.

While we cannot form a hypothesis similar to Hypothesis 2 for underconfident agents in every period, one can use (3) to derive a testable prediction for long-term learning

outcomes of overconfident and underconfident agents.¹³

Hypothesis 3.UC&OC (Stable Belief)

The stable belief of an underconfident (overconfident) agent with a larger bias lies further from the true state than the stable belief of a less underconfident (overconfident) agent.

Last but not least, the model characterizes the behavior of unbiased agents. An unbiased individual correctly evaluates his ability $\tilde{a} = A$. After choosing the optimal action in the first period, $e_1^* = \phi_0 = 0$, he observes exactly the output he expects: $\tilde{a} + \phi = 0 = -L(0)$. The unbiased agent has no reason to update his beliefs any further, implying:

$$\phi_2 = \phi_1 = \phi_0. \tag{5}$$

We summarize it in the following hypothesis:

Hypothesis 1.UB (Unbiased Agents)

The learning process of an unbiased agent is immediate and stable afterwards.

2.4 Quantitative Predictions

It is important to note that the model is based on the assumption that agents never update their beliefs about their ability. Although there is some evidence that people are reluctant to update beliefs about ego-relevant characteristics, especially if prompted to revise them downwards (see, for example, Eil and Rao, 2011), the assumption of no updating is rather extreme. Still, the qualitative predictions of the model will hold on aggregate if beliefs about ability are sufficiently sticky. Even if some agents correctly update their beliefs, as long as the bias is not entirely reduced in the population we will observe misguided learning. In this case, one would expect the effect to be of the same direction, but quantitatively lower than predicted by the model.

¹³While we admit that our setting is more suitable to test the short-term dynamics of the model, we argue that we can provide some evidence on the long-term. We discuss this point in more detail in Section 4.

3 Experimental Procedures

The experiment took place in November 2017 in the Laboratory for Experimental Economics at the University of Bonn. We conducted 8 two-part sessions with 19 to 25 participants each. In sum, we collected data from 171 male participants, mostly students from the university.¹⁴ The first and the second part of the experiment lasted around 45 minutes and 90 minutes, respectively. Participants earned 30 euros on average.

In the first part of the experiment, subjects completed an IQ test and filled out a questionnaire. The second part of the experiment took place one week later, after all subjects had completed the first part, and included the learning exercise as well as the elicitation of both prior and posterior beliefs.¹⁵ Both parts of the experiment were programmed using zTree (Fischbacher, 2007) and completed by subjects on computers in private cubicles. We describe each part in detail below.

3.1 IQ Test and Belief Elicitation

In the first part of the experiment, we evaluated subjects' relative performance in the IQ test, which consisted of 29 standard logic questions. Participants were asked to solve as many of them as possible in 10 minutes. The individual score was calculated based on the number of correctly answered questions minus the number of incorrect answers. To incentivize effort during the test, participants were told that the individual result is important for the next part of the experiment, and their earnings will depend on their scores. After the IQ test, subjects were asked to fill out a questionnaire designed to assess their character traits and individual anxiety levels. At the end of the first session, we reminded participants about the second session one week later, and that they will not be paid unless they show up for the second session.

¹⁴We invited only male subjects as our main research question concerns the consequences of overconfidence, and men are known to be more overconfident than women (Niederle and Vesterlund, 2007). We are not the only study that uses a group of male subjects to investigate overconfidence: see, for example, Burks et al. (2013).

¹⁵To match subjects' data between the sessions without violating anonymity, we followed a special procedure, which included generating private codes that were used to match subjects to cubicles at the beginning of the second session.

Between the sessions, we ranked participants according to their IQ test results. For every subject, we calculated his position in the group. The individual position was defined as a number equal to the percent of participants whose test scores were lower or equal to the score obtained by the subject. We defined 20 equi-length “performance intervals” ranging from 0% to 100% in steps of 5%. Every participant was assigned the performance interval that his position fell under (with 0 – 5% denoting the lowest and 95 – 100% the highest performance interval). We refer to the midpoint of that performance interval as the agent’s *relative performance* (denoted by A).

At the beginning of the second session, we elicited subjects’ prior belief about their relative position (Confidence I) using the crossover method that is known to be incentive-compatible independently of subjects’ risk attitudes (see Schlag et al., 2015). We presented participants with a choice list and asked them to indicate their preferred option in each of the 20 lines. Option A was a lottery with p chance of receiving 5 euros and $1 - p$ chance of receiving 0; the winning probability was increasing from $p = 0.05$ to $p = 1$ in 5% steps. Option B stood for a competition with a randomly selected individual, which granted 5 euros if one’s IQ test score was higher than their partner, and nothing otherwise. A rational individual would choose Option A if and only if p is larger than his perceived relative performance. Therefore, we interpret the switching probability as a measure of confidence in one’s skills. The procedure was explained to subjects in a simple language, with two examples on how to translate one’s beliefs into choices. We followed the same procedure in the second belief elicitation (Confidence II).

3.2 Learning Task

After the first belief elicitation, participants completed 6 rounds of a learning task. For every participant, we drew one number for each round, with replacement, from the set $\{-10, -9, \dots, 9, 10\}$.¹⁶ We refer to this collection of 6 numbers as an “individual set” and

¹⁶The numbers were drawn from a distribution that put slightly higher weight on numbers in the interval $[-4, 4]$. Participants were not presented the exact distribution but were told that the sum of numbers drawn is equal to zero in every round. We explained that some participants had been assigned the number 0, and among the rest half of the participants had been assigned a positive number, while the other half had the same number with the opposite sign.

to the set containing all feasible numbers “the feasible set”. Participants were reassured that the numbers had been drawn before the experiment started.¹⁷

In each round, participants were guessing one number taken from their individual set without replacement.¹⁸ For each number, they had to make 4 guesses and enter them into the interface one at a time. After each guess, the computer program calculated a payoff according to the formula:

$$\Pi(e, A, \Phi) = 20 + 0.8 \times (28.6 \times A + \Phi - 0.48 |e - \Phi|), \quad (6)$$

where A denotes the agent’s relative performance, Φ is the number drawn, and e refers to his guess. The formula corresponds to the specification of the absolute value loss function. We decided to use this specification because of its simple form and straightforward interpretation. The parameters were chosen such that misguided learning could arise for moderately biased agents. The formula was presented to participants in a descriptive form with an intuitive explanation of the absolute value in terms of distance on the linear scale. We drew subjects’ attention to the fact that the payoff is the higher the closer their guess is to the number drawn (with the highest payoff for the exact match). Participants were informed that, at the end of the experiment, two out of $4 \times 6 = 24$ guesses will be randomly drawn and paid out (with the exchange rate of 0.3).¹⁹

After entering a guess e , every participant received private feedback. The feedback was equal to one’s payoff with an added random component and was displayed on the

¹⁷We informed subjects that the numbers from their individual set had been printed and placed in a sealed envelope next to their seat. They were told not to open the envelopes until the end of the study. As an additional precautionary measure, we placed the envelopes within the sight of the person supervising the session.

¹⁸We framed the task as “guess the number” instead of “guess your ability and the number”, as we aim to test the theory that describes this particular type of situation. We argue that this framing is more suitable to study the implications of overconfidence. In many real-world situations learning about ability is not explicit. For example, when an investor is trading stocks his main task is to generate profits and learn about the market, and not about his ability (even though the profits depend on his analytical skills).

¹⁹One may raise a question whether paying subjects for two elicitation procedures and the learning exercise could induce participants to misreport their beliefs. We admit this possibility, however, we argue that this comment applies only to the second belief elicitation (Confidence II) and does not undermine our main result. Firstly, the instructions for each part were distributed separately, and beforehand participants were not given any information about the remaining tasks. Secondly, in the learning exercise, subjects were not able to influence their payoffs by misreporting their beliefs about themselves. Participants were informed that it is their actual relative performance that enters the payoff function, not their subjective belief.

individual computer screen. The noise was introduced only to ensure that subjects would not be able to infer their ability by matching the feedback to a single identical number in the table.²⁰

Participants were informed that they can infer the actual number they are guessing in a given round from their feedback. Knowing their relative performance A , the last guess e , and the payoff Π , they can calculate the unknown number Φ . However, it requires some arithmetical skills. Considering that computational mistakes could influence the learning behavior and obscure the results, we provided subjects with a tool to help them with the task.

3.2.1 Introducing Tables

Before the learning exercise, every participant was given a set of 21 tables (see Online Appendix A), from which they could obtain the value of Φ using the feedback they received. The tables contained payoffs for every possible combination of e , Φ , and A . The three parameters jointly determine the payoff, and hence the set of two-dimensional tables contains all feasible payoffs. There is one table for each possible guess e (indicated in the title), the rows indicate the relative performance A (performance intervals are listed in the first column), whereas the columns indicate the number Φ (its values are listed in the first row).

We provided participants with detailed instructions on how to correctly read the tables. Firstly, we described how to find the payoff given e , Φ , and A . A user has to look for a table with his guess in the title, and then look for the intersecting cell corresponding to the row with his relative performance and the column with the number. Secondly, we explained that if someone knows the payoff Π , his last guess e , and his relative performance A , he can obtain the value of Φ by reversing the last steps. After finding the right table, the subject should look at the row with his relative performance and search for his payoff in this row. The column in which lies the payoff indicates the number.

²⁰The random component was drawn from the uniform distribution over the interval $[-0.18, 0.18]$ known to the subjects. Importantly, the noise was not big enough to influence the update. Thus, the set-up can be still described using a model without a random component introduced in Section 2.

We presented participants with multiple examples and strongly encouraged them to raise questions when in doubt. Every participant had to answer control questions that not only tested their understanding but also pointed out important aspects of the task. Feedback was only displayed after the first guess and participants were not given any information prior to it. Therefore, the first guess that maximizes the expected payoff was $e = 0$. To avoid misunderstandings, we directly told subjects that it is in their best interest to choose zero as their first guess.

3.2.2 Experimental Conditions and Groups

We introduced two conditions: treatment (we refer to it as “multiple-feedback rounds”) and control (“single-feedback rounds”). The two conditions differed with respect to the information provided to participants after each guess. In the multiple-feedback rounds, participants received feedback calculated according to the formula (6) after each guess.

In the single-feedback rounds, subjects received feedback calculated according to (6) only after their 1st guess. After the 2nd and the 3rd guess computers displayed feedback calculated using the 1st guess in that round. Subjects were notified that no matter what they enter as their 2nd or 3rd guess, the feedback will not reflect their choices. Nevertheless, they were asked to enter their best guess two more times keeping in mind that every guess is equally important for their earnings. By comparing the 3rd and the 4th guess in the multiple-feedback rounds to the corresponding guesses in the single-feedback rounds, one can isolate the effect of informative feedback on misguided learning and prove that the mechanism described in Heidhues et al. (2018) drives the results.

Every participant completed a total of 6 rounds, alternating between the treatment and control conditions. We randomly assigned subjects to two groups (see Table 1), with the first group starting with a single-feedback round and the second group starting with a multiple-feedback round.

Table 1: Experimental Conditions and Groups

Round	Group 1	Group 2
1.	SF	MF
2.	MF	SF
3.	SF	MF
4.	MF	SF
5.	SF	MF
6.	MF	SF

SF – single-feedback round

MF – multiple-feedback round

4 Results

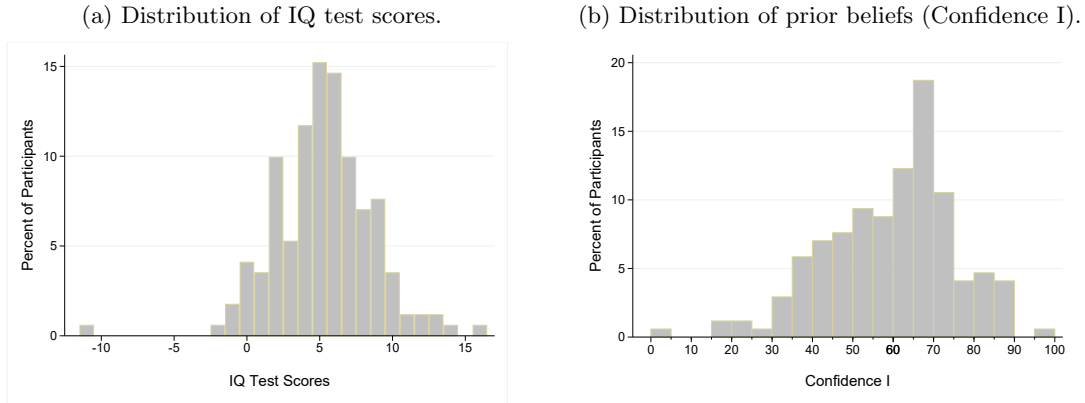
In this section, we present the results of our empirical analysis. Firstly, describe the data on test performance and beliefs, as well as our independent measure of overconfidence. In Section 4.2, we look at learning in the multiple-feedback rounds and test the model’s predictions for overconfident, underconfident, and unbiased subjects. Next, we use the data from the control condition to exclude alternative explanations of the results. Finally, we discuss learning about ability during the experiment in Section 4.3.

4.1 IQ Test Results and Elicited Beliefs

In Figure 1a) we present a histogram of the IQ test results. The scores range from -11 to 16 , with over 90% of participants obtaining between 0 and 10 points. The score distribution is fairly symmetrical, with a mean score of 5.29, and a standard deviation of 3.38. In Figure 1b) we present the distribution of subjects’ beliefs about their relative performance elicited before the main task (Confidence I). The average prior belief equals 59.46% and is significantly higher than the average actual position, 55.25% (p-value = 0.092).²¹ The average participant is overconfident, yet the magnitude of bias in our sample is not very high. The correlation between subjects’ prior beliefs and their actual performance is 0.31 and significant at the 1%-level.

²¹The average actual position is different from 50% as participants with the same IQ test score were, based on our definition, falling together into one performance interval. We decided not to randomly break ties to avoid misattribution of the result to the random component.

Figure 1: IQ test results and subjects’ beliefs about their relative performance.



We classify an agent as *overconfident* (*underconfident*) if he assessed his performance to be higher (lower) than his actual position within the group. An unbiased participant correctly estimated his relative performance. As revealed in Confidence I, there are 79 overconfident, 79 underconfident, and 13 unbiased subjects in our sample. On average, underconfident participants held beliefs 20 percentiles lower than their actual position. The average bias of overconfident subjects was around 30 percentiles, meaning that overconfident participants tend to believe that their relative performance was 30 percentiles higher than it actually was. There is a significant difference in the actual performance of overconfident and underconfident subjects. The low-ranked participants tend to overestimate their relative performance, and the high-ranked subjects tend to underestimate it (we address this issue in the following section). The exact values are presented in Table 2.

Table 2: Performance and beliefs of overconfident, underconfident and unbiased subjects.


	Underconfident	Unbiased	Overconfident
Actual Performance:			
Mean (Std. Dev.)	77.50 (16.41)	62.12 (13.14)	31.87 (19.81)
Prior Beliefs:			
Mean (Std. Dev.)	57.31 (16.71)	62.12 (13.14)	61.17 (15.76)

4.2 Model Predictions Based on Elicited Beliefs

4.2.1 Misguided Learning

First of all, we look at participants' choices in each of the 4 trials. The average guesses of overconfident subjects are presented in Figure 2. The blue line connects subjects' actual guesses, and the black points denote the average number being guessed.²² For every subject, we calculate the 2nd, 3rd and 4th guess predicted by the model, based on the number he was guessing and his bias revealed in Confidence I. The red line connects the average predicted guesses. Although subjects' actual guesses do not coincide with the predicted guesses, the belief path resembles the one predicted by the model. The learning of overconfident agents is self-defeating, with each guess diverging from the true state. We test this formally, by comparing coefficients of a simple regression explaining the difference between a guess and the number with dummy variables, one for each guess (see Table 3). The 2nd guess is significantly lower than the 1st guess (one-tailed test: p-value = 0.000), and the 3rd guess is significantly lower than the 2nd guess (one-tailed test: p-value = 0.019).

Figure 2: Learning of overconfident subjects in multiple-feedback rounds.



²²Although in every round the sum of numbers given to participants was equal zero, we could not predict the way in which they were distributed among the overconfident, the underconfident and the unbiased agents. Thus, the average of the numbers estimated by different groups was not exactly zero.

Although we cannot attest the strict inequality for the 3rd and the 4th guess with similar confidence level, the difference between the 2nd and the 4th guess is highly significant (one-tailed test: p-value = 0.003). Thereby, we confirm the qualitative predictions of the model for overconfident agents. Quantitatively, the effect is around 40% lower than predicted by the model. This may be due to conservatism (under-responsiveness to information known in the literature on asymmetric updating, e.g. Möbius et al., 2022) or subjects learning about their ability over the course of the experiment. We provide evidence for the latter explanation in Section 4.3.

The patterns revealed by the underconfident and unbiased agents also follow the model’s predictions. For the underconfident agents, the 2nd guess is significantly higher than the 1st guess (one-tailed test: p-value = 0.000), and the 3rd guess is significantly lower than the 2st guess (one-tailed test: p-value = 0.000). After receiving the first feedback, underconfident agents tend to overshoot, but in the following guess they correct their predictions downwards – a pattern also visible in Figure 3a). Quantitatively, however, the effect is even more mitigated compared to that of the overconfident agents: it is between 53% to 62% lower than predicted by the model.

Unbiased agents neither overshoot nor become pessimistic about the state over time, as we can see in Figure 3b). Their second guess is indistinguishable from the true state, and there is little change in the following trials. Thereby, we confirm Hypothesis 1 for overconfident, underconfident, and unbiased agents (Hypothesis 1.OC, 1.UC, and 1.UB).

Figure 3: Learning of underconfident and unbiased subjects in multiple-feedback rounds.

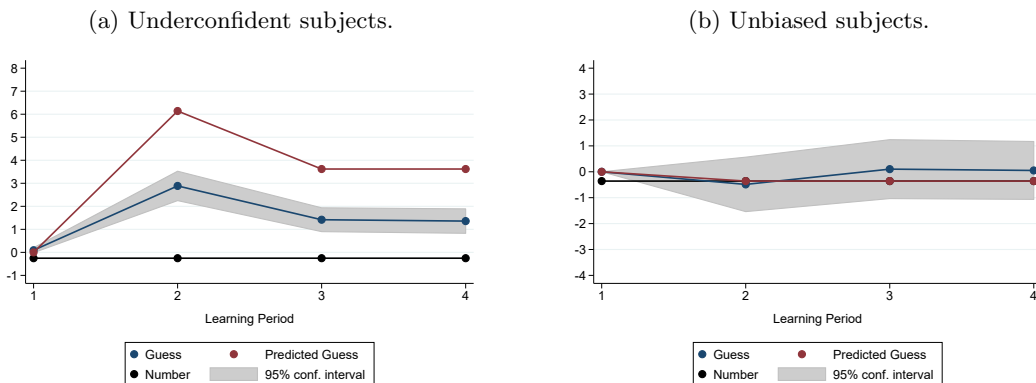


Table 3: Learning in multiple-feedback (MF) rounds.

	Overconfident (1)	Unbiased Agents (2)	Underconfident (3)
Dependent variable: difference between a guess and the number in MF rounds.			
Independent variables: dummy variables for each guess in the MF rounds.			
2 nd guess MF	-3.684*** (0.342)	-0.487 (0.570)	2.793*** (0.381)
3 rd guess MF	-4.245*** (0.391)	0.103 (0.466)	1.325*** (0.291)
4 th guess MF	-4.519*** (0.426)	0.051 (0.829)	1.266*** (0.548)
Const.	-0.004 (0.243)	0.359 (0.574)	0.346 (0.254)
<i>N</i>	948	156	948

Note: The coefficients at the 2nd, 3rd, and 4th guess MF remain unchanged if we control for subjects' relative performance (their actual position in the IQ test score distribution). Standard errors clustered at individual level. Their values in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.2.2 Excluding Alternative Explanations

Before we conclude that participants in our experiment were engaging in misguided learning, we test whether our results were driven by factors outside of the model. For example, the observed patterns might stem from the differences in cognitive ability between underconfident and overconfident subjects, and might not be specific to the environment described in the model. We address this concern with a control condition, in which the main mechanism of the model is switched off.

In the single-feedback rounds, participants received meaningful feedback only after their 1st guess. After the 2nd and the 3rd guess, the number displayed on screen was *independent of the preceding guess*. We kept all other features of the experiment unchanged: as in the multiple-feedback rounds, participants were asked to make four guesses and af-

ter each one, a number was displayed on their computer screens. Subjects were informed that the number displayed after the 2nd and the 3rd guess will be based on their 1st guess. Thus, the feedback after the 2nd and the 3rd guess does not bring any new information. The essential feature of the model – the interdependence between actions, feedback, and beliefs – is no longer present, so misguided learning should not arise. However, if there is a downward trend in beliefs of overconfident agents that is independent of the model mechanism, it should be present in the control condition as well.

Firstly, we show that there is no evidence of self-defeating learning in the single-feedback rounds after the second guess. Again, we compare the coefficients of subsequent guesses in a simple regression (Table 4). The 3rd guess is not significantly lower than the 2nd guess (one-tailed test: p -value = 0.953), and the 4th guess is not significantly

Table 4: Learning in single-feedback (SF) rounds.

	Overconfident (1)	Unbiased Agents (2)	Underconfident (3)
Dependent variable: difference between a guess and the number in the SF rounds. Independent variables: dummy variables for each guess in the SF rounds.			
2 nd guess SF	-3.350*** (0.360)	0.333 (0.786)	3.493*** (0.392)
3 rd guess SF	-2.958*** (0.378)	0.718 (0.779)	3.080*** (0.381)
4 th guess SF	-2.992*** (0.361)	1.051 (0.828)	3.198*** (0.387)
Const.	0.278 (0.258)	-0.513 (0.749)	-0.118 (0.237)
N	948	156	948

Note: The coefficients at the 2nd, 3rd, and 4th guess SF remain unchanged if we control for subjects' relative performance (their actual position in the IQ test score distribution). Standard errors clustered at individual level. Their values in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

lower than the 3rd guess (one-tailed test: p-value = 0.431). The results prove that, for overconfident agents, the belief path in the single-feedback rounds does not exhibit the pattern characteristic of self-defeating learning. While we cannot reject the hypothesis that the learning process of underconfident agents is self-correcting in the single-feedback rounds, the extent of correction is much lower. In the multiple-feedback rounds, participants corrected 55% of the initial overshooting, and the correction in the single-feedback rounds did not exceed 7%.

Secondly, we pool the data from the single- and multiple-feedback rounds and look at the effect of receiving informative feedback on learning. We regress the difference between a subject’s guess and the number on a dummy variable indicating a multiple-feedback round. The results are gathered in Table 2 in Online Appendix B. For overconfident participants, being in a multiple-feedback round increases the negative difference between a guess and the number by -1.57 in the 3rd guess (one-tailed test: p-value = 0.000) and by -1.81 in the 4th guess (one-tailed test: p-value = 0.000).²³ Informative feedback makes overconfident subjects more mistaken both in the 3rd and in the 4th guess.²⁴ As a final test, we regress the difference between the 4th and the 2nd guess on a dummy variable indicating a multiple-feedback round (see Table 4 in Online Appendix B). The coefficient is negative and highly significant: providing overconfident subjects with informative feedback shifts their beliefs downwards by -1.19 , which constitutes 67% of the effect predicted by the model. Moreover, receiving informative feedback affects underconfident but not unbiased subjects, in line with the model predictions. The results for underconfident and unbiased agents are delegated to Online Appendix B.

4.2.3 Individual Heterogeneity

In this section, we analyze how misguided learning depends on subjects’ bias. To this end, we conduct a regression analysis similar to the one presented in the previous section, but we allow for the effect to depend on the bias (the degree of over- and under-confidence measured before the task). The results are gathered in Table 5. The dependent variable

²³While the differences might appear small, they are close to the values predicted by the model (-1.59 in the 3rd guess and -2.05 in the 4th guess).

²⁴The negative sign indicates that overconfident subjects became more pessimistic about the state.

is the difference between the subject’s guess and the number in the multiple-feedback rounds, whereas independent variables include dummy variables indicating consecutive guesses and their interactions with a measure of agents’ bias. “Bias” variable takes values between -1 and 1 , with positive (negative) values characterizing overconfident (underconfident) subjects. We analyze separately the behavior of overconfident and underconfident subjects. However, this time, we include unbiased agents in each group, as they provide a useful benchmark for studying the effect of bias (similar regressions without unbiased subjects could be found in Table 5 in Online Appendix B).

The coefficients at the interaction terms provide evidence for a significant effect of bias on the learning process. For the overconfident subjects, a 10-percentile increase in bias exacerbates mislearning by -0.68 , -0.74 , and -0.76 in the 2nd, 3rd, and 4th guess, respectively. Thus, we confirm Hypothesis 2.OC that more overconfident participants tend to form more pessimistic beliefs and end up further away from the true state compared to less overconfident subjects. Moreover, we observe a similar effect in the group of underconfident subjects. A 10-percentile increase in bias results in additional overestimation of the number by 0.65 in the 2nd guess (underconfident agents’ bias takes negative values, hence the effect goes in the predicted direction). It confirms Hypothesis 2.UC, as more underconfident participants end up further away from the true state after the first feedback.

The results also shed light on Hypothesis 3.UC&OC. While we admit that our setting is more suitable to test the short-term dynamics of the model, we argue that the last guess is informative about the long-term. In our setting, most participants are expected to reach the stable belief within 4 trials.²⁵ The average stable belief is 3.78 for underconfident and -8.42 for overconfident subjects – both values are very close to the average predicted 4th guess of 3.62 and -7.80 , respectively. For this reason, we treat the 4th guess as close enough to the stable belief and test Hypothesis 3.UC&OC. The coefficient at the interaction with the 4th guess in Table 5 informs us about the effect of bias on the end belief. Both for the underconfident and overconfident subjects, more biased individuals end up further away from the true state, in line with the model predictions.

²⁵It is due to the chosen parameters, as well as the discrete action space (if we did not require subjects’ guesses to be integers, the convergence to the stable belief would take longer than 4 periods).

Table 5: The effect of bias on learning in multiple-feedback (MF) rounds.

	Overconfident or Unbiased (1)		Underconfident or Unbiased (2)	
Dependent variable: the difference between a guess and the number in MF rounds.				
Independent variables: dummy variables for each guess and their interactions.				
2 nd guess MF	-1.509***	(0.438)	1.198**	(0.521)
3 rd guess MF	-1.758***	(0.467)	0.225	(0.383)
4 th guess MF	-1.961***	(0.524)	0.179	(0.419)
Bias	-0.914	(1.113)	-1.212	(1.461)
Bias \times 2 nd guess MF	-6.848***	(2.089)	-6.526**	(2.642)
Bias \times 3 rd guess MF	-7.440***	(2.001)	-5.346***	(1.783)
Bias \times 4 th guess MF	-7.600***	(2.227)	-5.278***	(1.914)
Const.	0.277	(0.350)	0.138	(0.372)
N	1104		1104	

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


* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.3 Learning about Ability

As we have already mentioned, there is a substantial gap between subjects' guesses and the decisions predicted by the model. Why is misguided learning less pronounced than the model predicts? While the model is based on the assumption that agents do not change their beliefs about ability, we did not impose this assumption on our subjects.²⁶ As a result, we observe learning about ability over the course of the experiment.

²⁶It was our intention from the beginning to leave participants with an opportunity to revise their beliefs about their cognitive ability. We believe that imposing too many restrictions on subjects' behavior would make the test meaningless, as it would tell us little about how subjects would behave if not restricted. In our view, this design provides a more powerful and interesting test of the theory. Our results show that even without requiring subjects to hold on to their initial assessment, they follow the theoretical predictions as they *choose* to stick to their biased beliefs about their cognitive ability.

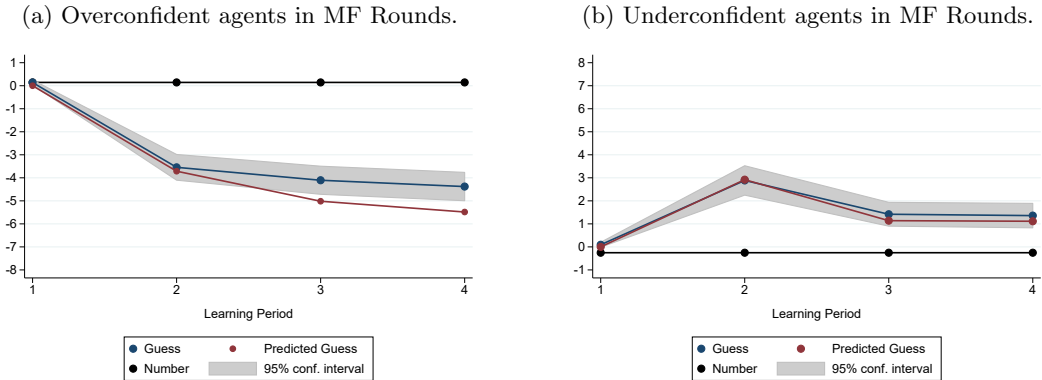
Figure 4: Distribution of subjects' bias before and after the task.



In Figure 4, we present the distributions of participants' bias before and after the task (based on Confidence I and Confidence II). The average bias of underconfident subjects decreased from 20.2 to 6 percentiles after the learning exercise, and the average bias of the overconfident subject decreased from 29.4 to 16.5 percentiles. The changes in mean beliefs are statistically significant both for the overconfident and the underconfident subjects. After the learning exercise, 18% of overconfident and 30% of underconfident subjects became unbiased. Nevertheless, Confidence II reveals that a significant portion of the sample held incorrect beliefs even after the learning exercise, and many of them were engaging in misguided learning till the very last round.

The data from Confidence I and II tells us little about the changes in subjects' beliefs about ability *during* the learning exercise. Fortunately, the experimental design enables us to divulge the beliefs about one's relative performance with few additional assumptions (see Online Appendix C.1). We assume that participants update their beliefs about ability at the beginning of each round, and use their initial guesses to obtain a measure of those updated beliefs (we use the 2nd guess, as in the 1st guess subjects were instructed to enter 0). The round-to-round changes in subjects' beliefs about ability are described in detail in Online Appendix C.2. Here, let us only point out their implications for the model's performance. We use beliefs revealed from the 2nd

Figure 5: Model’s predictions based on revealed beliefs.



guess to calculate the model’s predictions for the 3rd and the 4th guess. The results for the underconfident and overconfident agents are presented in Figure 5.

The average predicted 3rd and 4th guess (the red line) is now much closer to the average actual guess (the blue line). The better fit is reflected in the estimates of how well the model fits the data. The model based on revealed beliefs explains 73.5% variation in the choice data (the 3rd and the 4th guess), compared to 52.3% if we use its predictions based on elicited beliefs (see Online Appendix C.3). We conclude that the difference between our initial theoretical predictions and the actual guesses is due to participants learning their ability during the task. If we control for changes in beliefs from round to round, the model closely tracts subjects’ behavior.

5 Comparison with Ego-neutral Environment

We hypothesize that our results are driven in part by participants’ tendency to interpret feedback in a self-serving manner. We designed an additional control condition to test whether motivated reasoning is driving our results. In this condition, participants were learning about two parameters that were *both* ego-neutral. We used the same experimental design, with the only difference being that subjects performed the main task based on the performance parameter of another subject.²⁷ We assume that the performance

²⁷The almost identical experimental design enables to control for possible confounds such as, for example, the way subjects’ attention was directed during the experiment.

of another individual is irrelevant to one’s ego. Participants were informed that each of them will be randomly matched to another subject who completed the same IQ test and revealed similar beliefs through the same elicitation procedure. Before the main task, we elicited subjects’ beliefs about the relative performance of the participant matched to them and distinguished overconfident, underconfident, and unbiased agents (with respect to their partner’s performance). We again elicited beliefs about the performance of the matched partner after the task.

We collected data from 151 male participants, mostly students from the University of Bonn. There is no significant difference between the two groups in relative performance nor initial bias about own performance (see Table 16 in Online Appendix D.2). In the control group, 73 subjects were classified as overconfident about their partners’ performance, 73 as underconfident, and 9 as unbiased. In what follows, we refer to a control subject as overconfident (underconfident) if he overestimated (underestimated) his match’s performance. Note that participants’ were matched based on their beliefs and *not* on their bias. One consequence is that the average initial bias of overconfident subjects in the ego-neutral condition is lower than that of the overconfident subjects in the ego-relevant condition (see Table 6 and Online Appendix D.2 for a detailed description). We address this problem by controlling for the performance of the decision-maker, the initial bias defined with respect to the relevant variable, or both.

Table 6: Initial bias of overconfident and underconfident subjects in two conditions.

	Ego-neutral	Ego-relevant		Diff < 0	Diff ≠ 0	Diff > 0
Underconfident	-0.218 (0.014)	-0.202 (0.014)	p-value:	0.212	0.424	0.788
N	73	79				
Overconfident	0.247 (0.018)	0.293 (0.020)	p-value:	0.043	0.086	0.957
N	73	79				

Note: Initial bias is defined with respect to own performance in the ego-relevant condition and with respect to the match’s performance in the ego-neutral condition.

5.1 Learning about the Number

In this section, we show how the distance between the agent’s guess and the estimated number differs in the two conditions. As a measure of distance we use the absolute difference between a guess and the number. Table 7 presents the results of a regression analysis based on decisions of overconfident subjects. The distance between the agent’s last guess and the number is higher in the ego-relevant condition. The effect persists when we control for the initial bias (the second column in Table 7), the relative performance (the third column), or both (the fourth column).²⁸ Overconfident participants in the ego-relevant condition end up *more* mistaken about the state compared to similar subjects in the ego-neutral condition. In the last three columns in Table 7, we test for the treatment effect using the nearest neighbor matching estimator. We match participants based on the initial bias (Specification 5), on the relative performance (Specification 6), or both (Specification 7). They yield similar results.²⁹

Table 7: The effect of ego-relevance on learning of overconfident subjects.

<i>Dependent variable: the absolute difference between the 4th guess and the number.</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ego-relevant	1.143** (0.539)	1.168** (0.507)	1.698*** (0.522)	1.630*** (0.492)	1.203*** (0.364)	1.570*** (0.382)	1.229*** (0.375)
Controls 1	No	Yes	No	Yes			
Controls 2	No	No	Yes	Yes			
Adjustment Type	R	R	R	R	M	M	M
Observations	456	456	456	456	456	456	456

Note: The dependent variable is the absolute difference between the 4th guess and the number. The sample includes only overconfident participants. “Ego-relevant” indicates assignment to the ego-relevant condition (learning about own ability). Controls 1 include the initial bias of the decision-maker. Controls 2 include the relative performance of the decision-maker. The adjustment type in the first specifications is regression (R), and in the last ones matching (M). In the matching estimator, observations are matched to the nearest neighbor based on the initial bias (Specification 5), the relative performance of the decision maker (Specification 6), and both (Specification 7). In Specification 1-4, standard errors clustered at the individual level. In Specification 5-7, consistent standard errors as in Abadie and Imbens (2006). Their values in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

²⁸Similar regressions for the 2nd and the 3rd guess are presented in Online Appendix D.3.

²⁹As a final test, we add to specifications 1-4 a control for the model’s predictions (decisions implied by the model). The coefficients of interest are very similar (see Online Appendix D.3).

Table 8: The effect of ego-relevance on learning of underconfident subjects.

<i>Dependent variable: the absolute difference between the 2nd guess and the number.</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ego-relevant	-0.849** (0.403)	-0.740** (0.374)	-1.099*** (0.393)	-0.976*** (0.367)	-0.758** (0.305)	-1.056*** (0.315)	-0.816*** (0.312)
Controls 1	No	Yes	No	Yes			
Controls 2	No	No	Yes	Yes			
Adjustment Type	R	R	R	R	M	M	M
Observations	456	456	456	456	456	456	456

Note: The dependent variable is the absolute difference between the 2nd guess and the number. The sample includes only underconfident participants. “Ego-relevant” indicates assignment to the ego-relevant condition (learning about own ability). Controls 1 include the initial bias of the decision-maker. Controls 2 include the relative performance of the decision-maker. The adjustment type in the first specifications is regression (R), and in the last ones matching (M). In the matching estimator, observations are matched to the nearest neighbor based on the initial bias (Specification 5), the relative performance of the decision maker (Specification 6), and both (Specification 7). In Specification 1-4, standard errors clustered at the individual level. In Specification 5-7, consistent standard errors as in Abadie and Imbens (2006). Their values in parentheses.

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

What makes participants in the ego-relevant condition more mistaken about the state of the world? In the ego-relevant condition, overconfident agents might be reluctant to abandon their model of the world, as it would require them to admit to lower performance. Consequently, they will be less willing to correct their guesses compared to overconfident agents in the ego-neutral condition.³⁰

Our hypothesis finds support in the data from underconfident agents. In the ego-relevant condition, underconfident participants tend to overshoot significantly less compared to similarly underconfident subjects in the ego-neutral condition. The effect is highly significant even after controlling for the initial bias and relative performance (see Table 8). The sign of the effect is opposite to that of the overconfident agents – the ego-relevance of the task makes underconfident agents *less* misguided. However, the direction is consistent with motivated reasoning: in the ego-relevant condition, underconfident agents are more willing to abandon their previously held beliefs, as it allows

³⁰Nonetheless, misguided learning is not entirely eliminated in the control condition, pointing towards the role of biased beliefs as its main source (the analysis of the control data analogous to Section 4.2 could be found in Online Appendix D.1). Our results suggest that misguided learning can emerge in ego-neutral settings, although it is not as pronounced if agents are more willing to update their beliefs.

them to admit that they performed better than expected. This interpretation is also supported by the data on learning about ability presented in the next section.

5.2 Learning about Own versus Other’s Ability

The data from the second belief elicitation (Confidence II) reveals that, in the ego-neutral condition, 33 participants became unbiased about the ability of their match (compared to 38 participants in the ego-relevant condition). While the fraction of subjects who became unbiased is almost the same in the two conditions, the composition of types differs. In the ego-relevant condition, 30% of underconfident and 18% of overconfident participants revealed unbiased beliefs about their ability after the task. In the ego-neutral condition, these proportions are reversed: 18% of underconfident and 27% of overconfident subjects were classified as unbiased after the task.

The results in Table 9 provide evidence that overconfident participants in the ego-relevant condition are less likely to become unbiased compared to similarly overconfident participants in the ego-neutral control. At the same time, underconfident participants are more likely to become unbiased when learning about their own ability. Importantly, the effect is present if we control for the relative performance of the decision-maker and

Table 9: The effect of ego-relevance on becoming unbiased after the task.

<i>Dependent variable: dummy taking value 1 if a subject became unbiased after the task.</i>								
	<i>Overconfident</i>				<i>Underconfident</i>			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Ego-relevant	-0.097 (0.068)	-0.077 (0.071)	-0.168** (0.070)	-0.147* (0.075)	0.126* (0.069)	0.129* (0.071)	0.159** (0.072)	0.165** (0.075)
Controls 1	No	Yes	No	Yes	No	Yes	No	Yes
Controls 2	No	No	Yes	Yes	No	No	Yes	Yes
Observations	152	152	152	152	152	152	152	152

Note: The dependent variable is a binary variable indicating whether subject became unbiased after the task, as revealed in Confidence II. “Ego-relevant” indicates assignment to the ego-relevant condition. Controls 1 include the initial bias with respect to the relevant performance. Controls 2 include the relative performance of the decision-maker. Standard errors clustered at the individual level. Their values in parentheses.

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

the initial bias. The sign of the effect is indicative of motivated reasoning: overconfident subjects are less inclined to learn that they performed worse than expected, and underconfident subjects are more inclined to learn that they did better.

5.3 Discussion: Learning about Multiple Parameters

Our results show that self-defeating learning is more likely to arise and persist when one's ego is at stake. Overconfident participants, reluctant to revise their beliefs about ability downwards, are bound to become mistaken about the state of the world. On the other hand, underconfident agents are *more* willing to correct their beliefs about their own ability, making them *less* susceptible for mislearning in ego-relevant settings. The results also indicate that, when learning involves multiple parameters and some of them are ego-relevant, people will be steered to learn along the dimension that brings them higher ego utility. In the case of overconfident agents, this means holding onto their inflated beliefs, while for underconfident agents – revising them upwards. Still, as we have seen, neither underconfident nor overconfident subjects go the full length in updating or holding on to their biased beliefs about ability. More research is needed to understand how people make the trade-off between ego utility and the expected benefit of learning the state.

6 Conclusions

Successful decision-making often requires forming beliefs about various characteristics of the environment. However, learning about multiple parameters is rarely independent: the way an agent updates his beliefs about one aspect might influence his reasoning about other parameters. In particular, if the agent overestimates his ability, he may repeatedly misinterpret the data and fail to take the optimal action time after time, falling into a vicious circle of misguided learning. In this paper, we experimentally test subjects' propensity to engage in this kind of behavior. The results corroborate the theory formulated by Heidhues et al. (2018) and demonstrate that misguided learning is a real-world phenomenon that is likely to afflict biased agents. As long as people

hold on to their overconfident beliefs, they will continue to misread the data and form erroneous beliefs about their environment. The problem is aggravated when agents hold overconfident beliefs about characteristics they care about: their reluctance to revise their beliefs downwards exacerbates the tendency to mislearn. Allowing agents to experiment and acquire new information is, in these cases, counterproductive.

References

- Abadie, Alberto and Guido W Imbens (2006). “Large sample properties of matching estimators for average treatment effects”. In: *Econometrica* 74.1, pp. 235–267.
- Barber, Brad M and Terrance Odean (2001). “Boys will be boys: Gender, overconfidence, and common stock investment”. In: *The Quarterly Journal of Economics* 116.1, pp. 261–292.
- 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.
- Bohren, J Aislinn and Daniel N Hauser (2023). “The Behavioral Foundations of Model Misspecification: A Decomposition”. Working Paper.
- Burks, Stephen V, Jeffrey P Carpenter, Lorenz Goette, and Aldo Rustichini (2013). “Overconfidence and social signalling”. In: *The Review of Economic Studies* 80.3, pp. 949–983.
- Bushong, Benjamin and Tristan Gagnon-Bartsch (2023). “Reference dependence and attribution bias: evidence from real-effort experiments”. In: *American Economic Journal: Microeconomics* 15.2, pp. 271–308.
- Camerer, Colin and Dan Lovallo (1999). “Overconfidence and Excess Entry: An Experimental Approach”. In: *American Economic Review* 89.1, pp. 306–318.
- Chew, Soo Hong, Wei Huang, and Xiaojian Zhao (2020). “Motivated false memory”. In: *Journal of Political Economy* 128.10, pp. 3913–3939.
- Coutts, Alexander, Leonie Gerhards, and Zahra Murad (2024). “What to blame? Self-serving attribution bias with multi-dimensional uncertainty”. In: *The Economic Journal* 134.661, pp. 1835–1874.
- 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.
- Enke, Benjamin (2020). “What you see is all there is”. In: *The Quarterly Journal of Economics* 135.3, pp. 1363–1398.

- Enke, Benjamin and Florian Zimmermann (2017). “Correlation neglect in belief formation”. In: *The Review of Economic Studies* 86.1, pp. 313–332.
- Esponda, Ignacio, Emanuel Vespa, and Sevgi Yuksel (2024). “Mental models and learning: The case of base-rate neglect”. In: *American Economic Review* 114.3, pp. 752–782.
- Fischbacher, Urs (2007). “z-Tree: Zurich toolbox for ready-made economic experiments”. In: *Experimental Economics* 10.2, pp. 171–178.
- Gagnon-Bartsch, Tristan and Benjamin Bushong (2022). “Learning with misattribution of reference dependence.” In: *Journal of Economic Theory* 203, p. 105473.
- Golman, Russell, David Hagmann, and George Loewenstein (2017). “Information avoidance”. In: *Journal of Economic Literature* 55.1, pp. 96–135.
- Graeber, Thomas (2023). “Inattentive inference”. In: *Journal of the European Economic Association* 21.2, pp. 560–592.
- Hanna, Rema, Joshua Schwartzstein, and Sendhil Mullainathan (2014). “Learning Through Noticing: Theory and Evidence From a Field Experiment”. In: *The Quarterly Journal of Economics*, pp. 1311–1353.
- Heidhues, Paul, Botond Köszegi, and Philipp Strack (2018). “Unrealistic expectations and misguided learning”. In: *Econometrica* 86.4, pp. 1159–1214.
- Hestermann, Nina and Yves Le Yaouanq (2021). “Experimentation with self-serving attribution biases”. In: *American Economic Journal: Microeconomics* 13.3, pp. 198–237.
- 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.
- Malmendier, Ulrike and Geoffrey Tate (2005). “Does overconfidence affect corporate investment? CEO overconfidence measures revisited”. In: *European Financial Management* 11.5, pp. 649–659.
- (2008). “Who makes acquisitions? CEO overconfidence and the market’s reaction”. In: *Journal of Financial Economics* 89.1, pp. 20–43.
- 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.
- Niederle, Muriel and Lise Vesterlund (2007). “Do women shy away from competition? Do men compete too much?” In: *The Quarterly Journal of Economics* 122.3, pp. 1067–1101.
- Ortoleva, Pietro and Erik Snowberg (2015). “Overconfidence in political behavior”. In: *American Economic Review* 105.2, pp. 504–35.

Schlag, Karl H, James Tremewan, and Joël J Van der Weele (2015). “A penny for your thoughts: a survey of methods for eliciting beliefs”. In: *Experimental Economics* 18.3, pp. 457–490.

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