Irrational Inattention *

Ciril Bosch-Rosa[1], Muhammed Bulutay[2], and Bernhard Kassner[3]

[1]Chair of Macroeconomics, Technische Universität Berlin and Berlin School of Economics
[2]Alfred Weber Institut for Economics, Heidelberg University
[3]Seminar for Economic Policy, Ludwig-Maximilians-Universität München

April 25, 2025

Abstract

This paper proposes the concept of *irrational inattention*, a novel framework integrating overprecision into rational inattention models. Overprecision refers to an overestimation of the accuracy of prior beliefs. In our model, this bias leads to suboptimal updating directly by distorting the perceived value of new information and indirectly by amplifying the effects of attentional costs. We test these predictions using a 2x2 belief-updating experiment, manipulating overprecision through calibration feedback and attentional costs via information complexity. The results show that calibration feedback reduces overprecision and information complexity exacerbates attentional cost. With respect to beliefs, the results indicate that the effect of overprecision on updating depends on information complexity. These findings provide empirical evidence that inattention may arise from biased priors, rather than from rational cost-benefit optimization alone.

Keywords Overconfidence, Rational inattention, Belief updating, Overprecision

JEL Classification C83 · D91 · G41

*Corresponding author: muhammed.bulutay@awi.uni-heidelberg.de. We would like to thank Tilman Fries, Andreas Haufler, Frank Heinemann, Bogdan Klishchuk, Matthias Lang, Dmitry Taubinsky, and Mirko Wiederholt and the participants at the Maastricht Behavioral and Experimental Economics Symposium (M-BEES) 2022, the 12th Workshop on Theoretical and Experimental Macroeconomics, the Experimental Finance Conference 2022 and the 4th Behavioral Macroeconomics Workshop for their insightful comments. We acknowledge the financial support from the Deutsche Forschungsgemeinschaft (DFG) through CRC TRR 190 "Rationality and Competition" (project number 280092119).

1 Introduction

The canonical full information rational expectations (FIRE) framework assumes that agents form beliefs by optimally processing all available information. However, empirical evidence often contradicts this assumption.[1] A prominent explanation for these departures from the Bayesian benchmark is "rational inattention" (Sims, 2003). At the core of this theory is the premise that processing information is costly. Therefore, if the value of processing a unit of information does not exceed its cost, then ignoring that information becomes the "rational" course of action. This theory has emerged as a key framework, not only because it can explain a wide variety of economic phenomena such as financial contagion (Mondria and Quintana-Domeque, 2012) or the slow adjustment of prices to nominal shocks (Woodford, 2001; Sims, 2003), but also because it appeals to a common human behavior—ignoring information.[2]

While the assumption that agents can assess the costs and benefits of acquiring and processing information without bias serves as a useful benchmark, it is unlikely to hold in practice. A large body of literature has shown that individuals deviate significantly from rational or fully Bayesian optimization due to biases such as base rate neglect (Esponda et al., 2024), correlation neglect (Enke and Zimmermann, 2019), environmental complexity (Oprea, 2020), or simply limited cognitive ability (D'Acunto et al., 2021). More recently, the literature has highlighted overprecision as a systematic and pervasive bias shaping various aspects of decision-making and behavior (Kahneman, 2011; Moore et al., 2015; Bosch-Rosa et al., 2024).[3] Overprecision is a specific form of overconfidence in which individuals overestimate the precision of their information (Moore, 2022). As a result, it is closely related to individuals' subjective cost-benefit evaluation of information processing and attention.

This paper examines how overprecision influences (in)attention and shapes individu-

[1]While some studies show underreaction to changes in taxation (Chetty et al., 2009; Taubinsky and Rees-Jones, 2018) or to important monetary policy announcements (Coibion et al., 2021, 2020), others document overreaction to macroeconomic or financial news by households, firms, and experts (Bordalo et al., 2020; Broer and Kohlhas, 2024). See Born et al. (2023) for a recent literature review on the FIRE.

[2]See Maćkowiak et al. (2023) for a recent overview of the rational inattention literature.

[3]Examples include excessive trading in the stock market (Barber and Odean, 2001), the formation of systematic forecasting errors in finance (Deaves et al., 2019), political extremism (Ortoleva and Snowberg, 2015), or the prevalence of "fake news" (Thaler, 2024).

als' use of information. Specifically, we explore how overprecision may cause systematic deviations from optimal attention allocation by distorting the perceived cost-benefit assessments of information processing. Our key insight is that overprecision introduces a two-layered distortion. First, it biases the perceived precision of the agent's prior beliefs, leading them to overestimate the extent of their current knowledge. Second, because agents see fewer gains from processing new signals when they think their baseline knowledge is highly precise, they may exhibit "irrational inattention," allocating less attention than standard rational inattention models would predict.

To formalize these ideas, we develop a model in which an overprecise agent updates its beliefs about an uncertain state after observing a noisy signal. As in the canonical models of rational inattention, the agent can reduce signal noise by paying more attention. This process is costly, and the optimal level of attention is determined by balancing the potential benefits of more precise information against the costs of processing it. However, since the agent is overprecise, it overestimates the accuracy of its priors, distorting the cost-benefit tradeoff that governs optimal attention allocation. This distortion introduces a two-layered bias: overprecision not only skews the perceived value of information but also amplifies the perceived costs of attention. As a result, higher levels of overprecision reduce the marginal value of paying attention, leading agents to allocate less attention when information is costly, further reinforcing their initial bias.

We test our model's predictions using a novel survey task embedded in a pre-registered, incentivized 2x2 online experiment conducted on a large, representative sample of the German population. The survey consists of two blocks. In the first block, participants view a series of pictures depicting groups of people and are asked to estimate the average age of the individuals in each picture at the time it was taken. For each picture, we also elicit participants' expected *absolute* error.[4] Following the "Subjective Error Method" (SEM) introduced in Bosch-Rosa et al. (2024), we measure overprecision as the difference between participants' expected absolute errors and their realized absolute errors across all pictures in the first block. In the second block, participants are shown a new set of pictures and provide estimates of the average age and expected absolute error for each

[4]As an example, imagine we asked about the number of teeth an adult polar bear has. While participants will try to provide their best answer (with an *expected* error of zero), they will also be aware that their answer is most likely incorrect (Yeung and Summerfield, 2012; Bosch-Rosa et al., 2024). Their anticipated deviation from the true value, regardless of direction, the *absolute* error, is what we try to capture.

picture. However, in this block, participants provide these estimates twice; first with no information other than the picture (as in the first block) and then after receiving a "cloud" containing a random subset of all the ages in the picture. This design allows us to study participants' updating process conditional on their overprecision.

To identify the causal effects of overprecision and inattention on belief updating, we implement two experimental treatments designed to independently manipulate these components in the second block of the experiment. For overprecision, we provide feedback to a randomly selected subset of participants, informing them whether their absolute error estimates in the first block of pictures were overprecise or underprecise. The control group receives no feedback. By comparing pre- and post-feedback blocks, we can identify the causal effect of calibration feedback on overprecision. To exogenously increase attention costs, we randomly assign a subset of participants to a treatment that makes information processing harder by including irrelevant "decoy" words—such as "table" or "spaghetti"—into the information cloud, effectively increasing information complexity by adding extraneous bits without altering the signal's precision.

This stylized experimental design includes several unique features. First, although prior research has examined the effects of calibration feedback on overconfidence (e.g., Russo et al., 1992; Moore and Schatz, 2020), these studies measure overprecision using confidence intervals, a method known to be problematic due to respondents' unfamiliarity with them (Teigen and JØrgensen, 2005; Moore et al., 2015). In contrast, we measure overprecision using the SEM, a method that is robust across knowledge domains and heterogeneous populations (Bosch-Rosa et al., 2024).[5] Second, we employ decoy words as a simple yet novel treatment to manipulate attention costs. To our knowledge, the closest example is Ambuehl et al. (2022); Bronchetti et al. (2023), where participants solve two-digit addition problems containing both correct and incorrect answers, and assess the likelihood that at least half are correct. Third, we use a continuous state variable (age) to distinguish between the mean and variance of beliefs. Unlike the "colored dot" paradigm used in previous attention cost experiments (e.g. Dewan and Neligh, 2020), this more naturalistic measure induces variation in priors and, consequently, overprecision.

[5]Another line of literature considers the effect of measurement scales on the amount of measured overprecision (Haran et al., 2010). Our method avoids this issue, as it does not require scales, circumventing the pitfalls associated with this dimension.

Our results confirm the effectiveness of our experimental treatments. Feedback reduced overprecision among participants who were overprecise in the first block of pictures, while increasing it for those who were underprecise. The attention cost treatment also worked as expected: respondents in the "decoy" condition revised their beliefs less, made larger estimation errors, and spent more time studying the information cloud. Regarding belief updating, participants who received feedback indicating they were overprecise placed more weight on the signal in the second block when updating their beliefs. In other words, reducing overprecision led participants to view their prior beliefs as less accurate, thereby shifting the perceived cost-benefit analysis in favor of allocating more attention to new information.

Our pre-registered analysis does not detect the predicted interaction between overprecision and attention costs. However, exploratory analysis suggests that the interaction null result is likely due to feedback influencing both overprecision and perceived task difficulty, introducing a confound that masks the interaction effect in aggregate comparisons. To address this, we introduce a more granular statistical model that isolates the share of belief updating attributable to new information. The results indicate that overprecise participants are systematically less responsive to signals when processing costs are high and make make larger estimation errors in the "decoy" condition. Together, these findings support the model's prediction that *irrational inattention* emerges from the interaction between biased priors and increased information complexity.

Our study contributes to the growing literature on rational inattention by showing how cognitive biases can interact with information-processing costs to generate systematic deviations from optimal attention allocation. As noted by Maćkowiak et al. (2023), an important direction for future research is to extend the rational inattention framework beyond its current application in explaining behavioral anomalies to examine how these anomalies interact with information frictions. In line with this call, we show that overprecision, a well-documented cognitive bias, distorts attention allocation by altering the perceived value of new information. This contributes to a broader shift in the literature, from treating inattention as purely rational to recognizing how miscalibrated priors or distorted mental models can drive it. We also contribute to the literature that studies how non-standard preferences and beliefs shape attention, such as Pagel (2018), who incorporates prospect theory into portfolio choice under inattention, or Bolte and Raymond (2024), who show that payoff-dependent emotions influence attention allocation. In contrast to these preference-driven mechanisms, our results identify a belief-based distortion: overprecision inflates perceived prior accuracy, which in turn reduces responsiveness to costly information. This mechanism is consistent with the framework proposed by Gabaix (2019), in which behavioral phenomena like money illusion and base-rate neglect are driven by internally biased information processing. We extend this perspective by isolating overprecision as a specific cognitive distortion that amplifies attentional frictions and leads to suboptimal belief updating.

We also contribute to the literature on overprecision, a pervasive bias linked to a variety of negative outcomes and described by Bazerman and Moore (2012) as "the mother of all biases." Precisely because of this relevance, trying to reduce overprecision through feedback has been an active area of research (Moore and Healy, 2008; Haran et al., 2010), though such interventions have proven notoriously difficult (Moore et al., 2015; Sanchez and Dunning, 2023). In contrast, we show that if measured using the *subjective error method* (Bosch-Rosa et al., 2024), feedback can reduce overprecision in a simple context in the short run. More importantly, our results show that overprecision can lead to a misallocation of attention, leading people to become "irrationally inattentive."

Irrational inattention is particularly problematic because, unlike standard rational inattention, it does not respond predictably to incentives. While increasing the rewards for accuracy can offset attention costs under rational inattention, our model predicts that overprecision continues to suppress belief updating even when the stakes are high. This prediction is supported by our empirical finding that overprecision reduces responsiveness even when information is easy to process. Beyond validating the model, this distinction has direct practical implications for the design of information policies and behavioral interventions. For example, in economic policy contexts, raising the salience of monetary or fiscal announcements may improve responsiveness among inattentive individuals, but may have limited impact on those exhibiting overprecision. In settings where public information uptake is crucial (e.g., health advisories, financial education, or civic communication) interventions that target biased confidence in prior beliefs may be necessary to promote more accurate belief updating and improved decision-making.

2 Model & Hypotheses

We start with a standard Bayesian updating framework where agent i has a prior belief about the uncertain fundamental θ. These beliefs are normally distributed with mean $\mu_{\theta,i}$ and variance $\sigma^2_{\theta,i}$. The agent uses the perceived variance $\tilde{\sigma}^2_{\theta,i}$ rather than objective variance $\sigma^2_{\theta,i}$ with the difference depending on the agents' overprecision ω_i such that

$$\tilde{\sigma}^2_{\theta,i} = \frac{\sigma^2_{\theta,i}}{\omega_i}$$

with $\omega_i \in (0, \infty)$. Therefore, if an agent is overprecise ($\omega_i > 1$), it perceives its prior belief as *more* precise than it really is. If the agent is underprecise ($\omega_i < 1$), it perceives its prior belief as *less* precise than it really is.

Agents receive an unbiased noisy public signal x about the fundamental:

$$x = \theta + \epsilon,$$

where $\epsilon \sim N(0, \sigma^2_\epsilon)$. The signal noise is orthogonal to the prior belief noise. We follow Fuster et al. (2020) and model inattention by adding an additional individual-specific noise term ψ_i. Therefore, agent i perceives the signal as

$$s_i = x + \psi_i,$$

where ψ_i is assumed to be normally distributed as $N(0, \sigma^2_{\psi,i})$. The variance of this noise symbolizes the level of inattention.

In this setup, the posterior belief of agent i can be written as the weighted average of the signal and the prior mean:

$$E[\theta|s_i] = \beta_i \cdot (\theta + \epsilon + \psi_i) + (1 - \beta_i) \cdot \mu_{\theta,i}, \tag{1}$$

where the weight on the signal (or update rate) can be expressed as the ratio of the variances

$$\beta_i = \frac{\tilde{\sigma}^2_{\theta,i}}{\tilde{\sigma}^2_{\theta,i} + \sigma^2_\epsilon + \sigma^2_{\psi,i}}. \tag{2}$$

Therefore, unlike Fuster et al. (2020), where updates depend solely on signal and attentional costs, in our model, overprecision skews agents' weighting of new information.

2.1 Optimal Updating

The inattention to information, expressed by the variance of the individual specific noise $\sigma^2_{\psi,i}$, is modeled as the agent's choice variable. We assume that the agent's payoff depends on two factors: (i) the expected quadratic belief error and (ii) the cost of attention. Accordingly, agents solve:

$$\max_{\sigma^{-2}_{\psi,i}} -\phi_i \underbrace{E\left[E\left[(\theta - E[\theta|s_i])^2|s_i\right]\right]}_{\sigma^2_{\theta|s_i}} - C(\sigma^{-2}_{\psi,i}) \tag{3}$$

where $\sigma^2_{\theta|s_i}$ is the posterior variance of the belief given the level of the processed signal, ϕ_i is an individual-specific scaling parameter that measures the incentive to hold an accurate posterior, and the function C represents the cost of attending to the signal. Following the literature on rational inattention, we assume that the cost of attention is related to the expected reduction in uncertainty, measured by Shannon entropy:

$$\begin{aligned} C(\sigma^{-2}_{\psi,i}) &= g(H(x) - H(x|s_i)) \\ &= g\left(\frac{1}{2}\ln\left(\frac{\sigma^2_x}{\sigma^2_{x|s}}\right)\right) = g\left(\frac{1}{2}\ln\left(1 + \frac{\tilde{\sigma}^2_{\theta,i} + \sigma^2_\epsilon}{\sigma^2_{\psi,i}}\right)\right). \end{aligned}$$

A common approximation is to assume that the function g is linearly related to the expected reduction in uncertainty such that

$$C(\sigma^{-2}_{\psi,i}) = \frac{\lambda_i}{2}\ln\left(1 + \frac{\tilde{\sigma}^2_{\theta,i} + \sigma^2_\epsilon}{\sigma^2_{\psi,i}}\right) \tag{4}$$

In this case, the parameter λ_i reflects the marginal cost of attention.

By substituting equation (4) into equation (3) and using the expression in (2), we can rewrite the maximization problem as a function of the update rate β_i:

$$-\phi_i(1-\beta_i)\tilde{\sigma}^2_{\theta,i} - \frac{\lambda_i}{2}\ln\left(\frac{1}{1-\beta_i\frac{\tilde{\sigma}^2_{\theta,i}+\sigma^2_\epsilon}{\tilde{\sigma}^2_{\theta,i}}}\right).$$

Solving this problem for β_i yields the optimal solution:

$$\beta_i^* = max\left\{0, \frac{\tilde{\sigma}_{\theta,i}^2}{\tilde{\sigma}_{\theta,i}^2 + \sigma_\epsilon^2} - \frac{\lambda_i}{2\phi_i \tilde{\sigma}_{\theta,i}^2}\right\}. \tag{5}$$

The two-layered effect of overprecision on posterior beliefs is apparent in equation (5): the first term represents the standard Bayesian update rate, adjusted by the prior variance, which is biased by overprecision, while the second term captures the effect of inattention, which is also influenced by overprecision.[6]

2.2 Hypotheses

Assuming equation (5) has an interior solution, we can derive several hypotheses. The first two have to do with changes in overprecision and attention costs:

Hypothesis 1: Increasing overprecision leads to a lower update rate:

$$\frac{\partial \beta_i^*}{\partial \omega_i} = -\frac{\sigma_\epsilon^2/\sigma_{\theta,i}^2}{(1+\omega_i\sigma_\epsilon^2/\sigma_{\theta,i}^2)^2} - \frac{\lambda_i}{2\phi_i\sigma_{\theta,i}^2} < 0$$

Hypothesis 2: Increasing attention cost leads to a lower update rate:

$$\frac{\partial \beta_i^*}{\partial \lambda_i} = -\frac{\omega_i}{2\phi_i \cdot \sigma_{\theta,i}^2} < 0$$

Both hypotheses indicate underreaction to information relative to the FIRE benchmark. However, the mechanisms differ. Hypothesis 1 captures *irrational inattention* as underreaction results from biased beliefs about the accuracy of the prior. In contrast, Hypothesis 2 reflects rational inattention, where underreaction is optimal as agents deliberately limit information processing due to the explicit attention costs. A key implication of both hypotheses is that increasing incentives (ϕ_i) reduces underreaction by raising the value of holding accurate beliefs. However, only underreaction from rational inattention disappears entirely as $\phi_i \to \infty$. Thus, while increasing incentives may fully eliminate underreaction due to attention costs, they are insufficient to offset the effects of overprecision.

The third hypothesis has to do with the interaction between overprecision and attention costs:

[6] Note that in this formulation of equation (5), ϕ_i and λ_i are non-negative to ensure that each term's contribution to the update rate β_i^* is non-positive when considered independently.

Hypothesis 3: Overprecision amplifies the negative effects of attention costs on the update rate:

$$\frac{\partial}{\partial \omega_i}\left(\frac{\partial \beta_i^*}{\partial \lambda_i}\right) = -\frac{1}{2\phi_i \sigma^2_{\theta,i}} < 0$$

This interaction captures the indirect channel through which overprecision influences belief formation via the misallocation of attention. The intuition for this hypothesis lies in the non-linear structure of β, where each increase in the attention cost λ sharply push down β, thereby amplifying the effect of attention costs on belief updating. Figure 1 illustrates this effect by plotting the optimal update rate (vertical axis) as a function of overprecision (horizontal axis) and varying levels of attention cost (depicted by different line types) for the same level of objective prior variance and signal precision. As can be seen, the more overprecise one becomes, the greater the effect of changing the cost of attention, up to extreme levels where the agent does not update at all. This increased sensitivity reflects the twofold influence of overprecision on posterior beliefs: a direct effect on the perceived prior variance, and an indirect effect on the responsiveness to attention costs.

Figure 1: Optimal updating by varying overprecision and cost of attention

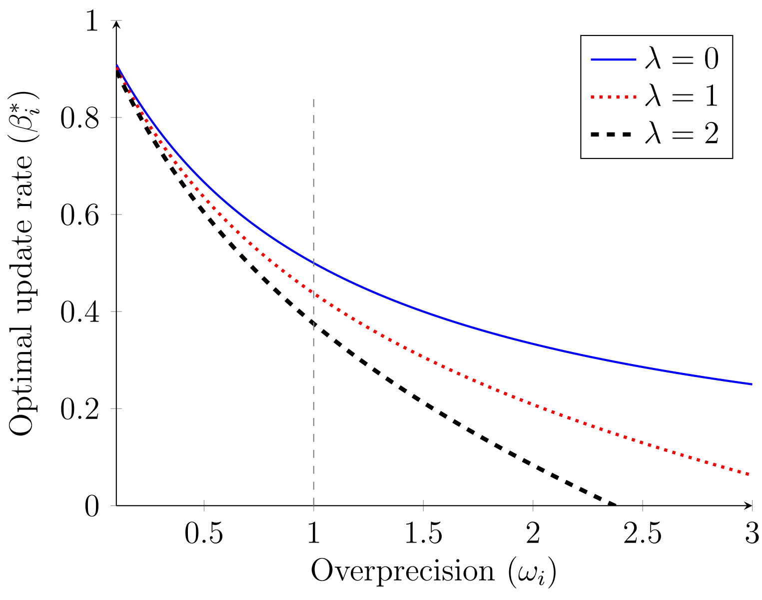

Calibration: $\phi = 2$, $\sigma^2_{\epsilon,j} = 4$, and $\sigma^2_{\theta,i} = 4$

3 Experimental Design

3.1 Design

To test the predictions of our model, we conduct a 2x2 between-subjects online experiment.[7] The experiment consists of two blocks. In the first block, participants view five historical pictures, such as that in Figure 2 (see Appendix D.1 for the full set).[8] For each picture, participants are asked to estimate the average age of the people in the image at the time the picture was taken. After submitting their estimate, participants are asked to assess the absolute distance between their estimate and the true average age in the picture. That is, participants are asked to estimate their "subjective error" (Bosch-Rosa et al., 2024). Both the age estimate and the subjective error are restricted to a range from 0 to 100.

Figure 2: Example of picture used in the online experiment.

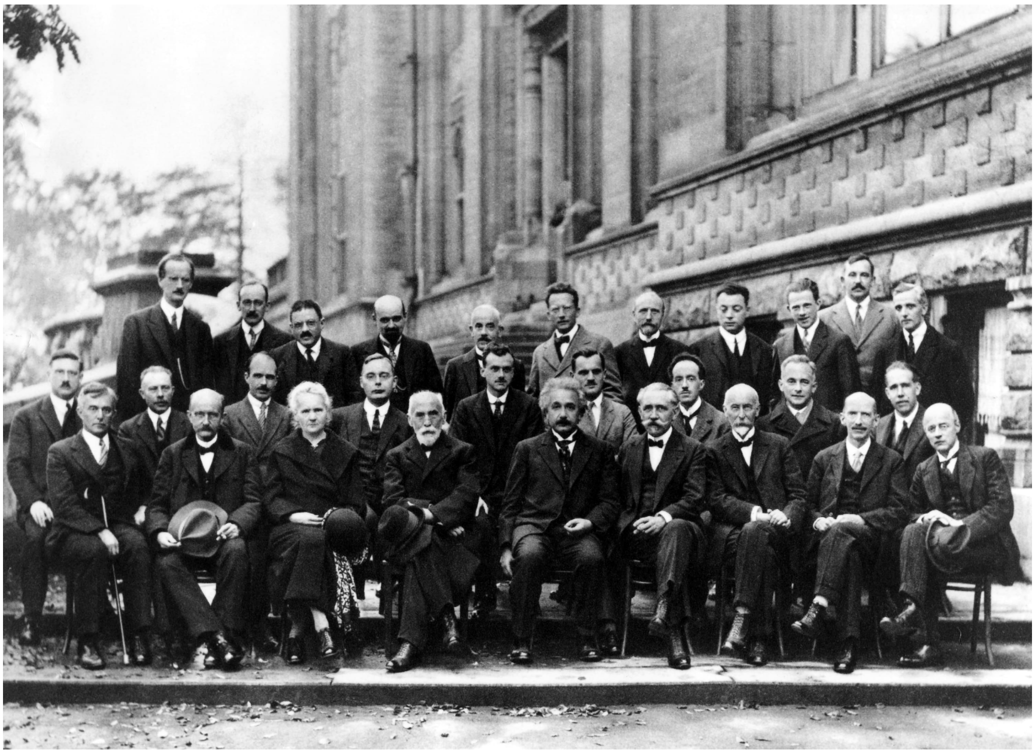

Note: This picture depicts the Fifth Solvay Conference which was held in October 1927. The average age of the people in the picture is 45.83 years. The youngest persons (Paul Dirac and Werner Heisenberg) are 25 years old while the oldest person (Hendrick Lorentz) is 74 years old.

[7]The experimental design and analysis plan are pre-registered at AsPredicted #154320.

[8]The pictures cover a range of topics, including politics, popular culture, and the arts. The order of the pictures for each block was randomized at the participant level.

At the end of the first block, two-thirds of participants are randomly assigned to the *Feedback* treatment, where they receive a feedback about the accuracy of their subjective errors. The remaining third are assigned to the *No-Feedback* treatment and receive no feedback. Specifically, subjects in the *Feedback* treatment receive the following text:

"In more than half of the last five pictures, your answer was [*(a) further away from the correct answer*][*(b) closer to the correct answer*] than you expected"

depending on the performance.[9] We refer to the version of feedback with (a) as *Overprecise* feedback and with (b) as *Underprecise* feedback.

After receiving the feedback, participants begin the second block, during which they once again estimate the average age of individuals in a picture and report their subjective error across five pictures. The difference between Block 1 and Block 2 is that, following the belief and subjective error report, participants are shown an "information cloud" containing fifteen unique ages of people in the picture. Participants are informed that each cloud contains fifteen ages corresponding to individuals in the picture.

Participants are randomly assigned to one of two informational treatments. In the *Easy-Info* treatment, the cloud displays 15 ages corresponding to individuals in the picture. In the *Hard-Info* treatment, the cloud includes 85 additional random words (decoys) interspersed with the ages (see Figure 3 for an example).[10] After viewing the cloud, participants are asked to enter their updated estimate of the average age in the picture.

A graphical summary of the whole experimental design can be found in Figure 4. The 2x2 design of the experiment allows us to observe belief updating by comparing posterior beliefs to initial beliefs as a function of overprecision and attention cost.

[9]We oversample the *Feedback* treatment to ensure sufficient statistical power for each type of feedback.

[10]An alternative approach would have been to allow subjects to endogenously select how many ages to observe for a certain given cost. Instead, we opted to exogenously manipulate attention costs through the use of decoy items in the signal cloud. This strategy offers clean identification of causal effects without introducing strategic behavior or confounding variation in effort or motivation.

Figure 3: Example of a information cloud by *Easy-Info* and *Hard-Info* treatments

(a) *Easy-Info* treatment　　**(b)** *Hard-Info* treatment

Note: Left panel shows the information cloud containing 15 ages of people in the picture. The right panel shows the format in the *Hard-Info* condition, where the same ages are in identical position, along with 85 random words. All participants within each treatment saw the same set of randomly drawn ages and words.

Figure 4: Overview of the experimental design

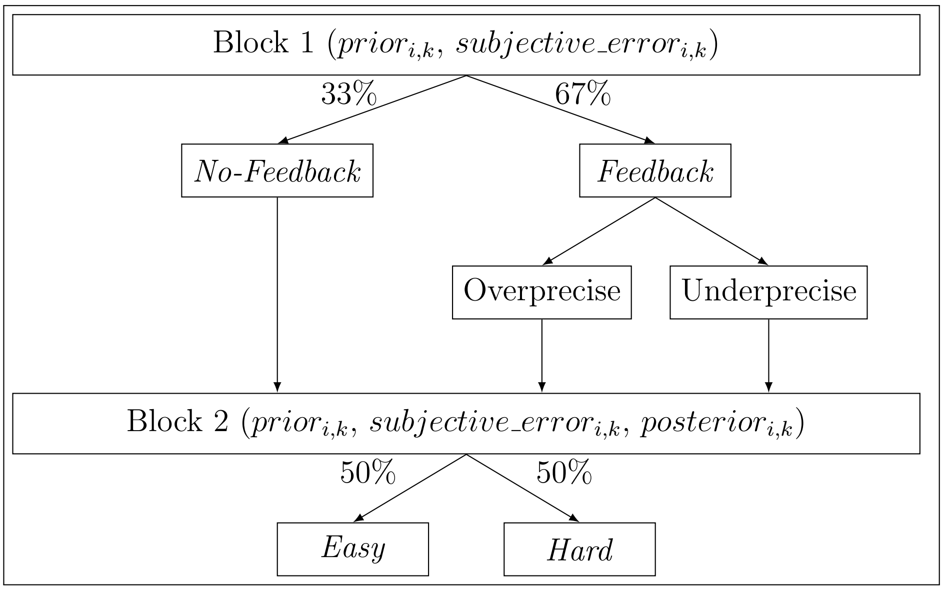

Note: Graphical representation of the experimental design. In both blocks, subjects evaluate five different pictures. The pictures in each block are the same for all subjects, but their order is randomized at the subject level. In Block 1, each participant i stated for each of the five pictures k a belief about the average age in the picture ($prior_{i,k}$) and reported a subjective error ($subjective\_error_{i,k}$). In Block 2, each participant stated a belief about the average age, the subjective error, and a posterior belief ($posterior_{i.k}$) after seeing the signal.

3.2 Incentives & Procedures

The survey was fielded by Bilendi, a professional survey provider, on a representative sample of the German population. Participants were informed that a randomly selected 5% sample would be eligible for a 10 € bonus payment, in addition to the standard fee paid by Bilendi, contingent on the accuracy of their answers. To prevent hedging, participants were told that the bonus payment would depend on the accuracy of a single, randomly selected estimate. Payoffs were determined using the binarized scoring rule (Hossain and Okui, 2013). Following Danz et al. (2022), the formulation of the scoring rule was not directly presented to participants. Instead, they were told that the probability of winning the bonus increases with the accuracy of their answers.[11] Participants had the option to click on a link to read a detailed, formal description of the payoff mechanism.

3.3 Sample & Data Selection

Because Bilendi pays participants based on survey completion rather than accuracy, we screened out participants who failed attention checks or answered fewer than two out of our four instruction comprehension questions correctly. Participants who passed the initial screening and completed the survey were considered "completed" surveys. In total, we collected 1321 completed surveys. Following our pre-registration, we dropped observations where the subjective error exceeded the distance between the prior and $max(prior, 100 - prior)$. If this behavior occurred more than once for the same participant, we removed all of its answers from the dataset. We also excluded any observation where the estimated average age was above 80 or below 20 years; if this occurred more than once for the same participant, we removed all of its answers from the dataset. We also excluded any participant identified as a "speeder", defined as those who completed more than 50% of the prior estimates in less than 25% of the median completion time. From the remaining sample, we retained the first 1,200 participants (our pre-registered sample) and dropped the rest.

The final sample is representative of the adult German population in terms of age (mean of 44.5 years), gender (50% female), education (30% hold a university degree),

[11] The specific wording was: "Je näher Ihre Antwort auf diese zufällig ausgewählte Frage an der richtigen Antwort liegt, desto höher ist die Wahrscheinlichkeit, dass der potenzielle Bonus dann tatsächlich ausgezahlt wird (für weitere Informationen hier klicken). Daher liegt es in Ihrem Interesse, auf jede Frage eine möglichst genaue Antwort zu geben."

and region. Table A.1 in Appendix A summarizes the demographic characteristics of the sample compared to census data. The median participant spent 13 minutes completing the survey.

4 Results

4.1 Overview of Data

Table 1 provides an overview of the data. From left to right, the table reports the name we use for each picture (subjects do not see it), the average age of the people in the picture on the day it was taken, the average prior belief, the average error made by participants in their estimation, the average subjective error, and the average posterior belief, where applicable.[12]

[12] Figure B.2 in Appendix C plots the density of the responses for each picture.

Table 1: Average prior beliefs, errors, subjective errors, and posteriors

Picture description	avg. age	prior	error actual	error subjective	post.
(1) Solvay Conference (1927)	45.83	56.09	10.99	5.23	.
		(8.62)	(7.44)	(3.83)	.
(1) Band Aid (1984)	27.89	29.84	4.86	3.99	.
		(5.91)	(4.11)	(3.07)	.
(1) Obama Cabinet (2009)	54.00	56.72	5.58	4.63	.
		(6.26)	(4.19)	(3.48)	.
(1) Paramount 75th Anniversary (1987)	52.32	54.10	7.45	5.25	.
		(8.76)	(5.23)	(4.01)	.
(1) Sino-German Consultations (2023)	57.04	53.07	5.93	4.53	.
		(6.17)	(4.97)	(3.09)	.
(2) G20 Summit (2017)	60.97	57.76	5.37	5.01	60.96
		(6.15)	(4.71)	(4.73)	(4.49)
(2) The Irascibles (1950)	40.20	46.10	7.20	4.75	42.20
		(6.42)	(5.63)	(3.97)	(4.58)
(2) Australian Labor Party (1901)	40.42	47.32	8.79	5.12	42.50
		(8.48)	(7.23)	(3.99)	(5.26)
(2) Merkel Cabinet (2019)	52.05	51.47	4.81	4.86	51.65
		(5.99)	(4.49)	(4.81)	(4.42)
(2) British Royal Family (2007)	48.65	49.32	7.66	5.91	49.95
		(9.47)	(5.97)	(4.66)	(8.30)

Notes: (1) and (2) refer to the block number. The year in parentheses refers to the year the picture was taken. Pictures within blocks are presented in random order. Standard deviations are shown in parentheses below the means.

The table shows that participants are engaged with the task, as the average age estimate is quite accurate.[13] The data also indicates overprecision, as the average absolute

[13]No participant provided the exact correct answer (to two decimal places) for any picture, nor did

error exceeds the average subjective error for all pictures in the first block. We also observe that the signal influences belief updating as the average posterior belief is closer to the true average age than the corresponding prior in all but one instance. Finally, the gap between subjective and actual errors narrows in the second block, showing that feedback messages help participants calibrate their subjective error estimates.

4.2 Manipulation Checks

The goal of our experimental design is to exogenously manipulate overprecision and attention costs. In this subsection, we analyze whether our interventions had the desired effect.

Overprecision

Following Bosch-Rosa et al. (2024), we define overprecision for subject i and picture k as:

$$overprecision_{i,k} = |error_{i,k}| - subjective\_error_{i,k},$$

where $error_{i,k}$ is the actual error made by subject i in estimating the average age of people in picture k, and $subjective\_error_{i,k}$ is their estimated subjective error for that picture. We then employ a Difference-in-Differences approach to assess the effect of *Overprecise* and *Underpercise* feedback. Specifically, we run an OLS regression where $overprecision_{i,k}$ is regressed on a dummy variable that takes value one if the participant is in the feedback treatment, the order of the picture, and the interaction between these two variables.[14] This regression is estimated separately for overprecise and underprecise participants, allowing us to compare the feedback-induced overprecision shifts across the two groups.

Figure 5 plots the estimated effect of feedback on overprecision across rounds, along with 90% confidence intervals. It is clear that participants who are initially overprecise become less overprecise after the *Overprecise* feedback, and vice versa for those with *Underprecise* feedback.

anyone consistently round to the correct answer for all 10 pictures, suggesting that help from search engines or large language models is unlikely.

[14]Formally, the regression model is: $overprecision_{i,k} = \alpha + \beta_1 \times \text{feedback}_i + \beta_2 \times \text{picture\_order}_k + \beta_3 \times (\text{feedback}_i \times \text{picture\_order}_k) + \epsilon_{i,k}$

Figure 5: Overprecision (relative to control) across pictures and feedback treatments

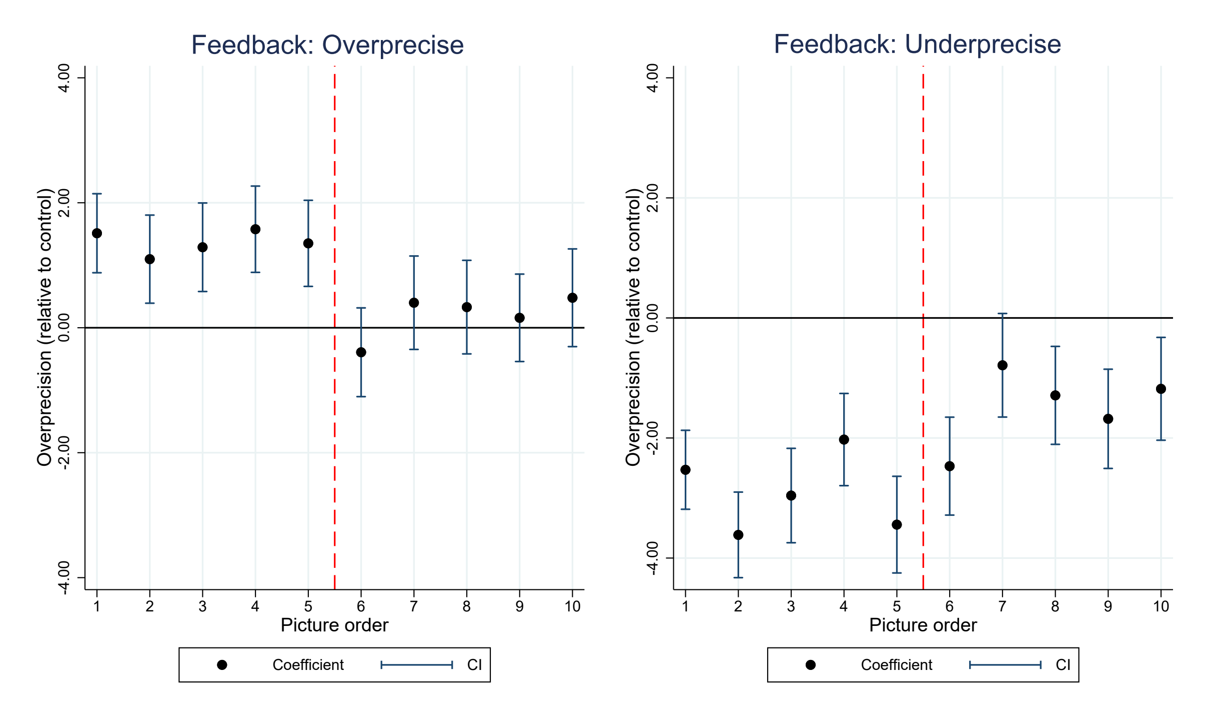

Notes: Estimated overprecision by picture and feedback treatment indicators. The dependent variable in the regressions is overprecision, and the regressors are indicators for pictures (by order of appearance), feedback treatments, and the interaction between the two. The red dashed line indicates the start of the second block. Bounds are 90% confidence intervals. Standard errors are clustered at the individual level.

As a robustness check, examine the effect of feedback on average overprecision within block, defined as:

$$diff\_overprecision_i = \frac{1}{5}\left(\sum_{k=6}^{10} overprecision_{i,k} - \sum_{k=1}^{5} overprecision_{i,k}\right).$$

We find that this difference is negative for those who received *Overprecise* feedback (mean $= -1.65$, p < 0.001, $N = 475$) and positive for those who received *Underprecise* feedback (mean $= 0.99$, p < 0.001, $N = 321$). Overall, it is clear that feedback effectively shifts participants' over and underprecision.

Attention Costs

To assess whether decoys in information clouds increase attention costs, we compare the *Easy-Info* and *Hard-Info* treatments using three proxies for information processing costs. The first proxy measures belief adjustment after exposure to the information cloud

Figure 6: CDF's of the adjustment of the prior and error of the posterior across treatments

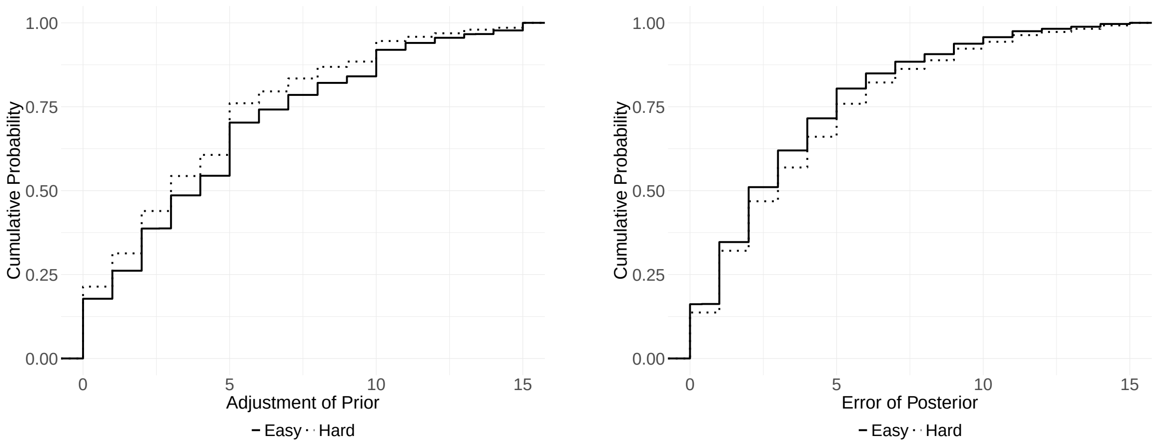

Note: The left panel plots the CDF of the adjustment made to the prior after seeing the signal across the *Easy-Info* (continuous line) and *Hard-Info* treatments (dotted line). The right panel plots the CDF of the error made by participants in their posterior belief in the *Easy-Info* (continuous line) and *Hard-Info* treatments (dotted line). For readability, we limit the horizontal axis of both figures to a maximum of 15 adjustments and error, Figure B.6 in Appendix C plots the figure with the whole support.

($adjustment_{i,k} = |posterior_{i,k} - prior_{i,k}|$). If attention costs are higher in *Hard-Info*, participants should adjust their posterior beliefs less than in *Easy-Info*. The second proxy, measures posterior accuracy ($error_{i,k} = |correct_{i,k} - posterior_{i,k}|$). If attention costs are higher in *Hard-Info*, posterior estimates should be less accurate than in *Easy-Info*. Finally, if attention costs are higher in *Hard-Info*, participants should spend more time processing the information cloud (i.e., $time\_cloud_{i,k}$) in *Hard-Info* than in *Easy-Info*.

Figure 6 plots the cumulative density functions (CDFs) of $adjustment_{i,k}$ (left panel) and $error_{i,k}$ (right panel), separated by treatment. Participants in the *Easy-Info* treatment adjust their beliefs more, as indicated by the weak stochastic dominance of the *Easy-Info* CDF over the *Hard-Info* CDF for $adjustment_{i,k}$. Similarly, the posterior beliefs are more accurate in *Easy-Info*, as the *Hard-Info* CDF weakly stochastically dominates that of *Easy-Info* CDF for $error_{i,k}$. We confirm these visual results by regressing $adjustment_{i,k}$ and $error_{i,k}$ on a treatment indicator with clustered errors at the individual level and picture fixed effects. The estimated coefficient for *Hard-Info* is significant and negative for $adjustment_{i,k}$ ($\widehat{\alpha} = -0.734$, $p < 0.01$) and significant and positive for $error_{i,k}$ ($\widehat{\alpha} = 0.497$, $p < 0.01$), consistent with higher attention costs under *Hard-Info*..

The interpretation of treatment effects on $time\_cloud_{i,k}$ is more nuanced, as higher information processing costs may lead some participants to reduce their attention. Nevertheless, Table 2 shows that the *Hard-Info* treatment is associated with significantly longer viewing times when controlling for outliers with median quantile or Huber robust regressions. Taken together, the evidence from all three time proxies supports the interpretation that the *Hard-Info* treatment increases attention costs.

Table 2: Decoy signals and the information time

Dep. var.:	$time\_cloud_{i,k}$		
	(1)	(2)	(3)
	OLS	Median	Robust
Hard-Info	0.716	1.426***	1.286***
	(1.144)	(0.260)	(0.192)
constant	18.078***	10.151***	10.451***
	(1.116)	(0.318)	(0.235)
Picture fixed effects	Yes	Yes	Yes
N	6000	6000	6000

Notes: The table reports results from regressions in which time spent looking at information is regressed on a treatment indicator and picture fixed effects. Columns (1) to (3) report OLS estimates, median quantile estimates, and robust regression (with Huber weights) estimates. Standard errors are clustered at the individual level for (1). Asterisks indicate results of two-tailed t-tests: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.3 Treatment Effects on Belief Updating

Having shown that both treatments work as intended, we study thier effects on the updating behavior of participants following the pre-registered regression:

$$\begin{aligned}
posterior_{i,k} - prior_{i,k} = \alpha + (signal_k - prior_{i,k}) \times \\
(\beta_0 + \beta_1 \times \text{Overprecise} + \beta_2 \times \text{Underprecise} + \beta_3 \times \text{Hard-Info} + \\
\beta_4 \times \text{Overprecise} \times \text{Hard-Info} + \beta_5 \times \text{Underprecise} \times \text{Hard-Info}) + \\
\theta_j \times picture_k + \varepsilon_{i,k}. \quad (6)
\end{aligned}$$

In this regression, participants' belief revision (the dependent variable) is modeled as a function of the difference between the signal (i.e., the average age in the information cloud) and the prior, the treatment combination, and picture fixed effects.

Equation (6) is motivated by our theoretical model, in which a Bayesian agent updates beliefs as a weighted average of the prior and the signal:

$$posterior_{i,k} = \beta_{i,k} \times signal_k + (1 - \beta_{i,k}) \times prior_{i,k}.$$

Rearranging the terms in equation (6) allows us to interpret the sum of $\widehat{\beta_T}$ as the update rate for each given treatment combination $T \in [0, 5]$.[15]

Table 3 reports the regression results based on equation (6).[16] The main results are in line with those reported in Table 1. First, participants in the benchmark group (*No-Feedback* and *Easy-Info* treatments) place more weight on the signal than on their prior belief, as indicated by $\widehat{\beta_0} > 0.5$. Second, participants who receive *Overprecise* feedback ($\widehat{\beta_1}$) update their beliefs significantly more than those without feedback. However, while those who receive *Underprecise* feedback ($\widehat{\beta_2}$) has a negative coefficient, as predicted by our model, it is not statistically different from zero. One possible explanation could be the smaller sample size of underprecise subjects (27%) or respondents react more strongly to being told they are overprecise than to being told they are underprecise. Third, increasing the cost of attention ($\widehat{\beta_3}$) reduces the amount of belief updating. Taken together, these results form the basis of the first set of results:

> **Result 1**: Reducing overprecision increases the update rate, consistent with Hypothesis 1.

> **Result 2**: Increasing the cost of attention decreases the update rate, consistent with Hypothesis 2.

---

[15]For example, the update rate for the combination of *Hard-Info* and *Overprecise* feedback treatments is given by $\widehat{\beta} = \widehat{\beta_0} + \widehat{\beta_1} + \widehat{\beta_3} + \widehat{\beta_4}$.

[16]Since we pre-registered directional hypotheses, the reported t-test results are one-tailed.

Table 3: Belief updating by treatments

	(1) *Revision*	(2) *Revision*
$\beta_0 : signal - prior$	0.595***	0.798***
	(0.028)	(0.245)
β_1 : Overprecise $\times$ $(signal - prior)$	0.065**	0.075**
	(0.037)	(0.033)
β_2 : Underprecise $\times$ $(signal - prior)$	-0.012	-0.006
	(0.041)	(0.036)
β_3 : Hard-Info $\times$ $(signal - prior)$	-0.138***	-0.130***
	(0.041)	(0.034)
β_4 : Overprecise $\times$ Hard-Info $\times$ $(signal - prior)$	-0.011	-0.018
	(0.056)	(0.046)
β_5 : Underprecise $\times$ Hard-Info $\times$ $(signal - prior)$	0.080	0.079
	(0.061)	(0.051)
$\alpha : constant$	1.389***	2.014***
	(0.108)	(1.079)
Picture fixed effects	Yes	Yes
Covariates	No	Yes
N	5946	5862
R^2	0.57	0.58

Notes: The table reports results from OLS regressions in which belief revision is regressed on the difference between the signal value and the prior, the treatment indicators, and their interaction. The second column adds demographic covariates for the regression in the first column. Robust standard errors clustered at the individual level are reported. Asterisks indicate the results of one-tailed t-tests: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. $\hat{\beta}_5$ is in the opposite direction of the pre-registered hypothesis, so the one-tailed test is not significant.

Our third hypothesis (Hypothesis 3) predicts that overprecision amplifies the negative impact of attention costs on belief updating. Specifically, we expect an interaction between interventions that change overprecision (ω) and the marginal cost of information processing (λ). However, we do not find evidence consistent with this prediction, as neither of the interaction estimates is statistically different from zero.[17] Therefore, we conclude that:

> **Result 3**: We do not find evidence that overprecision amplifies the effect of attention costs on update rates when the cost of attention changes.

Finally, our data shows an asymmetry between *Underprecise* and *Overprecise* feedback. As shown in Table 3 and the manipulation checks in Section 4.2, participants are more responsive to *Overprecise* feedback than to *Underprecise* feedback. One likely reason is that *Overprecise* feedback is perceived as more critical or surprising as it tells participants they were more confident than they should have been, whereas *Underprecise* feedback is more affirming and may feel less consequential. This difference in perceived salience could explain the stronger response to *Overprecise* feedback. Moreover, this pattern is consistent with Haaland et al. (2023), who emphasize that information interventions tend to have stronger effects when individuals' prior beliefs are more misaligned with the content of the message.[18] Similarly, our results show that feedback significantly shifts beliefs for participants who were initially overprecise, while it has little impact on those who were underprecise. These two mechanisms—greater saliency of *Overprecise* feedback and treatment effect heterogeneity—likely interact, suggesting that our intervention primarily acts by correcting biased priors, with a particularly strong effect when feedback is both surprising and negative.

> **Result 4**: We find evidence that feedback interventions produce an asymmetric impact. *Overprecise* feedback on participants' calibration reduces their overprecision, whereas *Underprecise* feedback has less influence on increasing it.

5 Discussion

Our findings indicate that reducing overprecision through feedback increases the update rate, while raising attention costs via the *Hard-Info* treatment decreases it. These results

[17] Although a one-tailed test in the positive direction would yield significance at the 10% level for $\widehat{\beta_5}$, this was not the hypothesized effect.

[18] In our setting, Overprecise feedback challenges participants,Äô belief in the accuracy of their initial estimates and confidence judgments, while Underprecise feedback largely confirms that participants were cautious or uncertain.

support our hypotheses that overprecision and attention costs independently influence how participants revise their beliefs in response to new information. However, contrary to our third hypothesis, the pre-registered regression model does not detect a statistically significant interaction between overprecision and attention costs. One potential explanation for the lack of an interaction effect is a violation of the *ceteris paribus* assumption. Our feedback treatment, while intended to influence overprecision, may have unintentionally affected participants' perceptions of task complexity. In particular, underprecise participants may have been nudged to perceive the task as "simpler" than anticipated (e.g., "This isn't that hard after all") by the Underprecise feedback, while overprecise participants may have been nudged into perceiving the task as "more complex" than anticipated (e.g., "This is trickier than I thought"). These shifts introduce a confounding channel through which feedback affects not only overprecision but also participants' perception of the environment's difficulty, thereby altering their mental effort or attention allocation. As a result, the hypothesized interaction between overprecision and attention costs—by which overprecise individuals should respond more strongly to changes in information costs—may be masked or dampened as we are measuring a conflation of the direct effect of shifting overprecision *and* a secondary effect stemming from changes in perceived task complexity (see Appendix B for a graphical example).

To empirically evaluate this explanation, we look at the post-experimental survey, where we included a question on the perceived task difficulty. Specifically, we asked participants:

> "How difficult was it for you to view (visually) the pictures and word clouds with the information we showed you in the survey?",

where the responses ($difficulty_i$) are on a Likert scale (1: Not at all difficult - 5: Very difficult).

Table 4 presents the results of regressing $difficulty_i$ on treatment indicators, their interactions, and demographic controls. The estimates show that participants who received Overprecise feedback perceived the task as significantly more complex, suggesting a shift in perceived complexity. Moreover, the absence of interaction effects between the difficulty and Overprecise feedback treatment suggests that feedback alone had a uniform effect across all participants. This broad shift in perceived difficulty may prompt individuals to see everything as more complex, thereby altering their priors, effort, or level of caution in ways that the original difference-based specification of eq. (6) cannot disentangle.

Table 4: Self-reported difficulty of task by treatments

	(1) difficulty	(2) difficulty
Hard-info	0.189*	0.191*
	(0.114)	(0.114)
Overprecise	0.344***	0.346***
	(0.113)	(0.114)
Underprecise	-0.023	0.001
	(0.106)	(0.107)
Hard-info × *Overprecise*	-0.201	-0.206
	(0.160)	(0.161)
Hard-info × *Underprecise*	-0.180	-0.204
	(0.159)	(0.160)
Constant	2.059***	2.573***
	(0.077)	(0.271)
N	1195	1183
R^2	0.02	0.03
Covariates	No	Yes

Notes: The table reports results from a OLS regression in which self-reported difficulty of the task is regressed on treatment indicators and their interactions. The second column adds demographic covariates for the regression in the first column. Robust standard errors clustered at the individual level are reported. Asterisks indicate the results of two-tailed t-tests: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

To address this problem, rather than relying on the coarse treatment-control comparison, we adopt a more granular "update ratio" ($U_{i,k}$) as:

$$U_{i,k} = \frac{posterior_{i,k} - prior_{i,k}}{signal_k - prior_{i,k}}. \tag{7}$$

This ratio measure is conceptually closer to the theoretical parameter β in eq. 1, as it can be interpreted directly as the Bayesian weight on the signal.[19] Focusing on $U_{i,k}$ allows us to isolate responsiveness to new information by capturing the fraction of the gap between the prior and the signal that is closed by the posterior.

Moreover, $U_{i,k}$ serves as a self-normalizing metric as it compares belief revision not to a fixed prior, but to the prior that participants hold after receiving feedback. Regardless of whether feedback changes participants' perception of the task difficulty, the denominator $signal_k - prior_{i,k}$ adapts accordingly, allowing us to measure their responsiveness to the signal. For example, if Overprecise feedback causes a participant to anchor their priors closer to the midpoint of 50 years (perhaps because now the task seems harder), $U_{i,k}$ still captures how far their posterior moves relative to the prior, preserving interpretability. In this way, $U_{i,k}$ remains a valid proxy for the theoretical weight β_i even when attitudes toward the task shift. While both the update ratio and the average overprecision are measured after the feedback intervention, and therefore cannot be used to identify causal effects *per se*, their interaction provides insights into the *mechanism* itself. Specifically, we interpret the observed relationship between post-feedback overprecision and updating behavior as evidence that overprecision modulates responsiveness to complex information, consistent with the model's predictions. In this sense, $U_{i,k}$ allows us to recover the theoretically predicted interaction that remained undetected in the coarser, treatment-level difference-in-differences approach.

We then estimate the model:

$$U_{i,k} = \alpha + \gamma_1 \cdot avg\_op_i + \gamma_2 \cdot (avg\_op_i \times Hard\text{-}Info_i) + picture_k + \varepsilon_{i,k}, \tag{8}$$

where avg_op_i is the average overprecision of participant i across the five pictures of block 2, $Hard\text{-}Info_i$ is a dummy variable equal to one if the participant was assigned to the Hard-Info condition, and $picture_k$ a dummy identifying each picture k. The coefficients γ_1 and γ_2 capture, respectively, the direct effect of overprecision and its interaction with processing costs. Since $U_{i,k}$ directly proxies for $\beta_{i,k}$, any reduction in the update ratio under high-cost conditions for participants with higher overprecision supports the theoretical amplification effect described in Section 2. Another advantage of this model is

[19]From eq. 1 we have that $posterior_{i,k} = \beta_{i,k} \cdot signal_k + (1 - \beta_{i,k}) \cdot prior_{i,k}$, which can be rearranged as the update ratio $\frac{posterior_{i,k} - prior_{i,k}}{signal_k - prior_{i,k}} = \beta_{i,k} = U_{i,k}$.

that it provides us with a direct test of whether the "update weight" ($U_{i,k}$) responds to overprecision and its interaction with processing costs.

Table 5 reports the results using OLS, a quantile regression, and a Huber-robust regression. Across all specifications, we find a robust and statistically significant negative interaction between overprevision and the *Hard-info* treatment, confirming Hypothesis 3. While, the main effect of average overprecision is small and statistically insignificant, this might be because when information is easy to process, even overprecise respondents remain responsive to the signal. It is only when processing costs are high that overprecision meaningfully reduces responsiveness. This result provides direct evidence for the mechanism of irrational inattention, as overprecision amplifies the effect of information complexity on belief updating.

It is important to understand the difference in findings between the pre-registered model and the granular update-ratio approach. The former estimates average treatment effects across experimental groups and detects direct effects of feedback and information complexity on belief updating. However, it does not account for variation in overprecision within treatment groups. As a result, it may attribute part of the effect of overprecision to the treatment itself, capturing it as a direct effect. By contrast, the update ratio model normalizes how far individuals move relative to the prior-signal gap and explicitly isolates individual overprecision and its interaction with information complexity. In doing so, it reveals that overprecision alone does not consistently reduce updating; its effect emerges only when processing costs are high.

Table 5: Updating Treatment Interaction

	(1) update	(2) update	(3) update
avg_op	0.029	0.005	0.002
	(0.025)	(0.004)	(0.003)
hard_info	-0.117	-0.099***	-0.075***
	(0.178)	(0.026)	(0.019)
hard_info × *avg_op*	-0.089**	-0.018***	-0.011**
	(0.042)	(0.006)	(0.005)
constant	0.759***	0.758***	0.686***
	(0.114)	(0.030)	(0.023)
N	6000	6000	6000

Notes: OLS, quantile regression, and Huber robust regressions of the "update ratio" on the average overprecision of participants ($overprecision_i$), the information treatment ($hard_i info_i$), and picture fixed effects (not reported). Errors clustered at the individual level in the OLS. Asterisks indicate the results of two-tailed t-tests: * $p < 0.2$, ** $p < 0.1$, *** $p < 0.05$.

Finally, as a complementary test of our model, we examine whether the interaction between overprecision and the information costs affect the accuracy of participants' responses. After all, if reduced updating reflects irrational inattention, it should result not only in lower responsiveness to the signal (as shown in Table 5), but also in less accurate final beliefs. Table 6 reports the results of regressing the absolute posterior error on overprecision, the information treatment dummy, and its interaction. We find a positive and statistically significant interaction between overprecision in the robust specifications (columns (2) and (3)). That is, once we control for extreme values, participants who are both overprecise and assigned to the Hard-Info treatment make systematically larger mistakes, providing further evidence that overprecision leads to information loss when signal processing is costly, consistent with the lower update ratio shown in Table 5 and with Hypothesis 3.

Table 6: Posterior Absolute Error Interaction

	(1) error	(2) error	(3) error
avg_op	0.149**	0.061**	0.069***
	(0.065)	(0.025)	(0.017)
$hard\_info$	0.386	0.140	0.020
	(0.374)	(0.132)	(0.090)
$hard\_info \times avg\_op$	0.049	0.183***	0.148***
	(0.171)	(0.035)	(0.024)
$constant$	2.668***	1.933***	2.474***
	(0.161)	(0.146)	(0.100)
N	5962	5962	5962

Notes: OLS, quantile regression, and Huber robust regressions of the absolute error in the posterior on the average overprecision of participants ($overprecision_i$), the information treatment ($hard_info_i$), and picture fixed effects (not reported). Errors clustered at the individual level in the OLS. Asterisks indicate the results of two-tailed t-tests: * $p < 0.2$, ** $p < 0.1$, *** $p < 0.02$.

6 Conclusion

This paper identifies overprecision as a source of *irrational inattention.* In our model, overprecise agents overestimate the precision of their information. This bias distorts the cost-benefit analysis of paying attention to new information. As a result, overprecision leads to suboptimal levels of inattention.

Our model's novelty lies in integrating overprecision within the rational inattention framework. Traditional models of rational inattention assume agents are unbiased in evaluating both prior information and processing costs. By relaxing this assumption, our model introduces a two-layered bias: first, because agents are overprecise, they overestimate the accuracy of their prior beliefs. Second, this distortion in prior beliefs leads agents to underweight the value of new signals, resulting in a misallocation of attention resources. This dual distortion not only reduces responsiveness to new information among overprecise agents but also amplifies the effects of information complexity. An important implication of our model is that, while standard inattention can be mitigated by increasing incentives for accuracy, overprecision distorts belief updating even when stakes are high. This distinction highlights the unique challenge of *irrational inattention* and the negative impact that overprecision can have on individual's daily life (Bosch-Rosa et al., 2024).

We test our model's predictions using a novel 2x2 pre-registered experiment that independently manipulates overprecision and attention costs. By shocking subjects' overprecision through feedback and manipulating informational complexity, we provide causal evidence of the impact that overprecision has on belief updating. Specifically, we show that reducing overprecision via feedback increases responsiveness to information while increasing attentional costs diminishes it. However, we do not detect a significant interaction between overprecision and attention costs. In an exploratory analysis, we provide evidence that the null result is likely due to the unintended effect of feedback on participants' task complexity: informing participants that they were overprecise not only has an impact on their overprecision, but also shifts the perceived complexity of the task, which confounds our pre-registered model. To address this, we move beyond binary treatment comparisons and estimate a more granular model that isolates belief updating by directly linking the extent of belief updating toward the signal to participants' level of overprecision. Using this approach, we find that overprecise participants become especially inattentive when information is costly to process, confirming the core mechanism of *irrational inattention*.

Overall, our paper contributes both to the theoretical and empirical literature on rational inattention, offering a novel approach to how cognitive biases influence decision-making under information constraints. On the theoretical side, we present a structured framework that integrates overprecision, a well-documented behavioral basis, into a rational inattention model. Our model suggests that what has traditionally been considered "rational" inattention may, in fact, be irrational when agents hold biased priors, as this leads them to underweight new information even when it is valuable and accessible. An important implication of our model is that while standard inattention can be mitigated by increasing incentives, overprecision resists such interventions, highlighting the importance for policy makers to understand the channel that drives individuals to be inattentive.

Empirically, we provide strong evidence that overprecision can be reduced through feedback. While previous studies have documented the presence of overprecision in various domains (e.g., Moore, 2022; Bosch-Rosa et al., 2024), few have explored whether and how it can be corrected, particularly in contexts involving real-time decision-making under uncertainty. Our results show that targeted feedback effectively recalibrates participants' priors, reducing overprecision and, consequently, increasing responsiveness to new information. Moreover, the ability to experimentally induce changes in overprecision allows us to causally identify how biased priors affect attention allocation—confirming our theoretical results and demonstrating the presence of *irrational inattention*, a type of inattention driven by biased beliefs rather than rational cost-benefit considerations alone.

References

Ambuehl, S., A. Ockenfels, and C. Stewart (2022): “Who opts in? Composition effects and disappointment from participation payments,” *Review of Economics and Statistics*, 1–45. Cited on page 4.

Barber, B. M. and T. Odean (2001): “Boys will be Boys: Gender, Overconfidence, and Common Stock Investment,” *The Quarterly Journal of Economics*, 116, 261–292. Cited on page 2.

Bazerman, M. H. and D. A. Moore (2012): *Judgment in Managerial Decision Making*, Free Press. Cited on page 6.

Bolte, L. and C. Raymond (2024): “Emotional Inattention,” *Available at SSRN 4776766*. Cited on page 5.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020): “Overreaction in Macroeconomic Expectations,” *American Economic Review*, 110, 2748–2782. Cited on page 2.

Born, B., Z. Enders, and G. J. M$\surd$*ller*(2023) : “*OnFIRE*, *News*, *andExpectations*,″ *Tech.Rep.*18259, *CEPRDiscussionPapers*, *cEPRPress*, *Par*

Bosch-Rosa, C., B. Kassner, and S. Ahrens (2024): “Overconfidence and the Political and Financial Behavior of a Representative Sample,” Berlin School of Economics Discussion Papers 0051, Berlin School of Economics. Cited on pages 2, 3, 4, 6, 11, 17, 30, and 31.

Broer, T. and A. N. Kohlhas (2024): “Forecaster (mis-) behavior,” *Review of Economics and Statistics*, 106, 1334–1351. Cited on page 2.

Bronchetti, E. T., J. B. Kessler, E. B. Magenheim, D. Taubinsky, and E. Zwick (2023): “Is Attention Produced Optimally? Theory and Evidence From Experiments With Bandwidth Enhancements,” *Econometrica*, 91, 669–707. Cited on page 4.

Chetty, R., A. Looney, and K. Kroft (2009): "Salience and Taxation: Theory and Evidence," *The American Economic Review*, 99, 1145–1177. Cited on page 2.

Coibion, O., Y. Gorodnichenko, E. S. Knotek II, and R. Schoenle (2021): "Average Inflation Targeting and Household Expectations," Working Paper 27836, National Bureau of Economic Research. Cited on page 2.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte (2020): "Inflation expectations as a policy tool?" *Journal of International Economics*, 124, 103297. Cited on page 2.

Danz, D., L. Vesterlund, and A. J. Wilson (2022): "Belief elicitation and behavioral incentive compatibility," *American Economic Review*, 112, 2851–2883. Cited on page 14.

Deaves, R., J. Lei, and M. Schröder (2019): "Forecaster Overconfidence and Market Survey Performance," *Journal of Behavioral Finance*, 20, 173–194. Cited on page 2.

Dewan, A. and N. Neligh (2020): "Estimating information cost functions in models of rational inattention," *Journal of Economic Theory*, 187, 105011. Cited on page 4.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2021): "Human Frictions in the Transmission of Economic Policies," Working Paper 29279, National Bureau of Economic Research. Cited on page 2.

Enke, B. and F. Zimmermann (2019): "Correlation neglect in belief formation," *The review of economic studies*, 86, 313–332. Cited on page 2.

Esponda, I., E. Vespa, and S. Yuksel (2024): "Mental models and learning: The case of base-rate neglect," *American Economic Review*, 114, 752–782. Cited on page 2.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2020): "Expectations with Endogenous Information Acquisition: An Experimental

Investigation," Working Paper 24767, National Bureau of Economic Research. Cited on page 7.

Gabaix, X. (2019): "Behavioral inattention," in *Handbook of behavioral economics: Applications and foundations 1*, Elsevier, vol. 2, 261–343. Cited on page 6.

Haaland, I., C. Roth, and J. Wohlfart (2023): "Designing information provision experiments," *Journal of Economic Literature*, 61, 457–502. Cited on page 23.

Haran, U., D. A. Moore, and C. K. Morewedge (2010): "A simple remedy for overprecision in judgment," *Judgment and Decision Making*, 5, 467–476. Cited on pages 4 and 6.

Hossain, T. and R. Okui (2013): "The binarized scoring rule," *Review of Economic Studies*, 80, 984–1001. Cited on page 14.

Kahneman, D. (2011): *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux, 1st ed. Cited on page 2.

Maćkowiak, B., F. Matějka, and M. Wiederholt (2023): "Rational inattention: A review," *Journal of Economic Literature*, 61, 226–273. Cited on pages 2 and 5.

Mondria, J. and C. Quintana-Domeque (2012): "Financial Contagion and Attention Allocation," *The Economic Journal*, 123, 429–454. Cited on page 2.

Moore, D. A. (2022): "Overprecision is a property of thinking systems." *Psychological Review*. Cited on pages 2 and 31.

Moore, D. A. and P. J. Healy (2008): "The trouble with overconfidence," *Psychological Review*, 115, 502–517. Cited on page 6.

Moore, D. A. and D. Schatz (2020): "Overprecision increases subsequent surprise," *Plos one*, 15, e0227084. Cited on page 4.

Moore, D. A., E. R. Tenney, and U. Haran (2015): "Overprecision in Judgment," in *The Wiley Blackwell Handbook of Judgment and Decision Making*, John Wiley & Sons, Ltd, 182–209. Cited on pages 2, 4, and 6.

Oprea, R. (2020): "What makes a rule complex?" *American economic review*, 110, 3913–3951. Cited on page 2.

Ortoleva, P. and E. Snowberg (2015): "Overconfidence in Political Behavior," *American Economic Review*, 105, 504–35. Cited on page 2.

Pagel, M. (2018): "A news-utility theory for inattention and delegation in portfolio choice," *Econometrica*, 86, 491–522. Cited on page 5.

Russo, J. E., P. J. Schoemaker, et al. (1992): "Managing overconfidence," *Sloan management review*, 33, 7–17. Cited on page 4.

Sanchez, C. and D. Dunning (2023): "Are experts overconfident?: An interdisciplinary review," *Research in Organizational Behavior*, 100195. Cited on page 6.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50, 665–690. Cited on page 2.

Taubinsky, D. and A. Rees-Jones (2018): "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *The Review of Economic Studies*, 85, 2462–2496. Cited on page 2.

Teigen, K. H. and M. JØrgensen (2005): "When 90% confidence intervals are 50% certain: On the credibility of credible intervals," *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 19, 455–475. Cited on page 4.

Thaler, M. (2024): "The fake news effect: Experimentally identifying motivated reasoning using trust in news," *American Economic Journal: Microeconomics*, 16, 1–38. Cited on page 2.

Woodford, M. (2001): "Imperfect Common Knowledge and the Effects of Monetary Policy," Working Paper 8673, National Bureau of Economic Research. Cited on page 2.

Yeung, N. and C. Summerfield (2012): "Metacognition in human decision-making: confidence and error monitoring," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367, 1310–1321. Cited on page 3.

A Additional Tables

Table A.1: Demographic characteristics of the sample

	Sample	Census†
age	44.49	44.6
	(14.25)	
female	0.504	0.507
	(0.500)	
German	0.938	0.854
	(0.242)	
eastern states	0.196	0.194
	(0.397)	
university degree (or more)	0.312	0.335
	(0.463)	
	(0.496)	
N	1200	

Notes: East refers to people living in Berlin, Brandenburg, Mecklenburg-Vorpommern, Sachsen, Sachsen-Anhalt, or Thüringen. †: Census data are from Destatis and are as of 2022, except for education statistics, which are as of 2019, and for the adult population (age$\geq$ 15). Standard deviations are shown below the means.

B Ceteris Paribus

Figure B.1: Optimal updating if feedback treatment had secondary effects on cost of attention

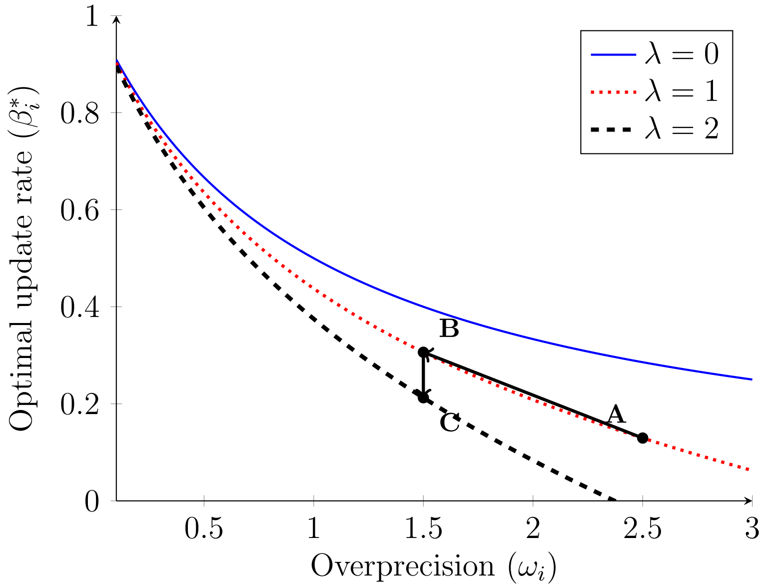

Calibration: $\phi = 2$, $\sigma^2_{\epsilon,j} = 4$, $\sigma^2_{\theta,i} = 4$.

Figure B.1 illustrates the case where feedback treatment has secondary effects on cost of attention. Consider an overprecise participant positioned on the dotted line (point A) who receives feedback in the *Overprecise* treatment. We predict that this participant would decrease their overprecision and increase their update rate (point B). However, the "task-is-harder-than-I-thought" effect might lead the participant to update less due to cost of attention (point C). If true, this pattern might explain the lack of an interaction effect.

C Additional Figures

Figure B.2: Density of the prior beliefs for each picture

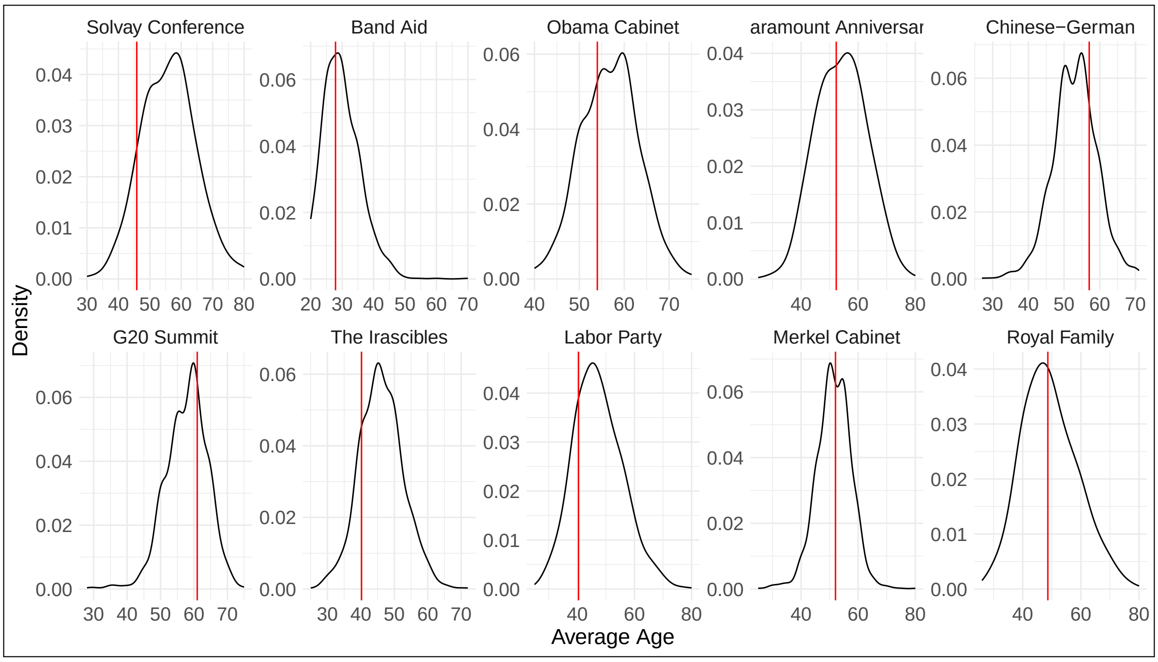

Notes: The vertical red line marks the correct answer. Note that the vertical axis differs for each question.

Figure B.3: Change in overprecision across blocks by feedback type

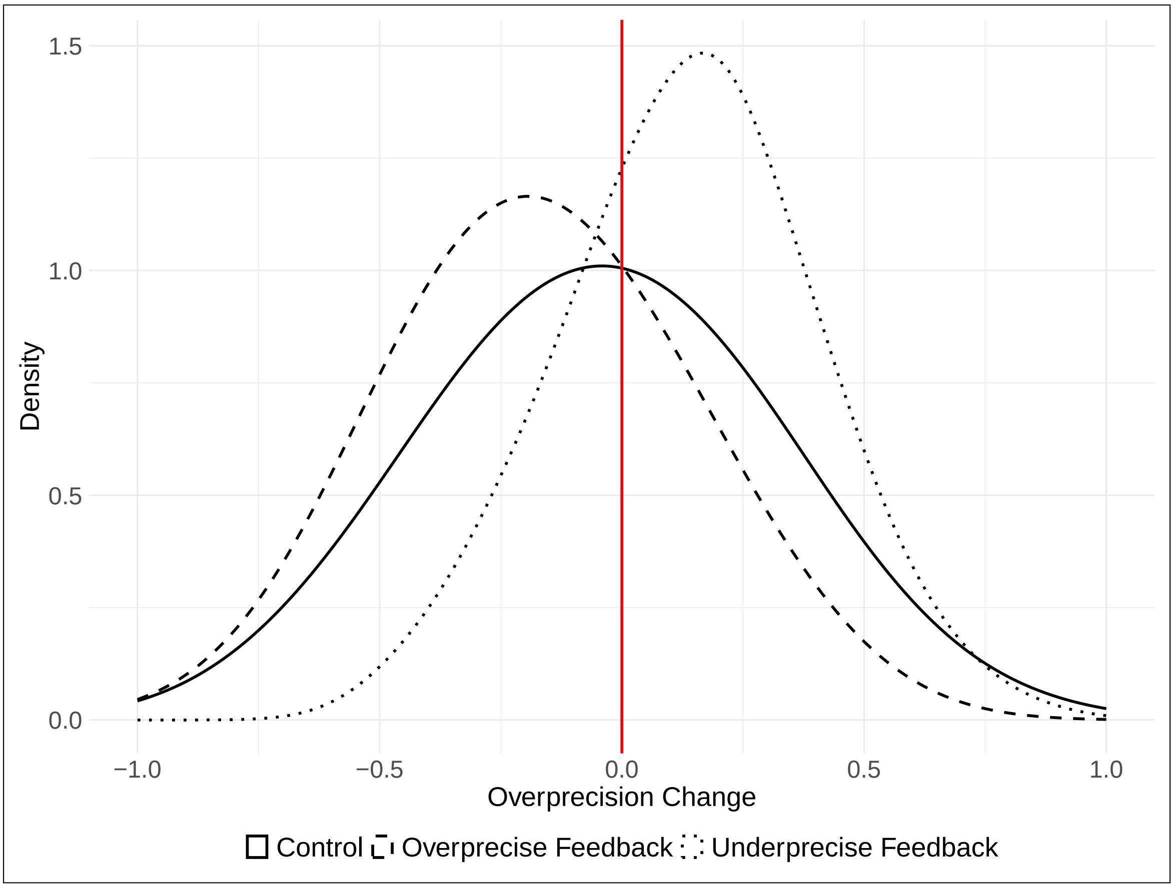

Notes: The figure shows the density of the difference in average overprecision (binary) between the pictures in the second and first blocks.

Figure B.4: Density of the answers for each picture. The vertical red line marks the correct answer. Note that the vertical axis is the same for all graphs.

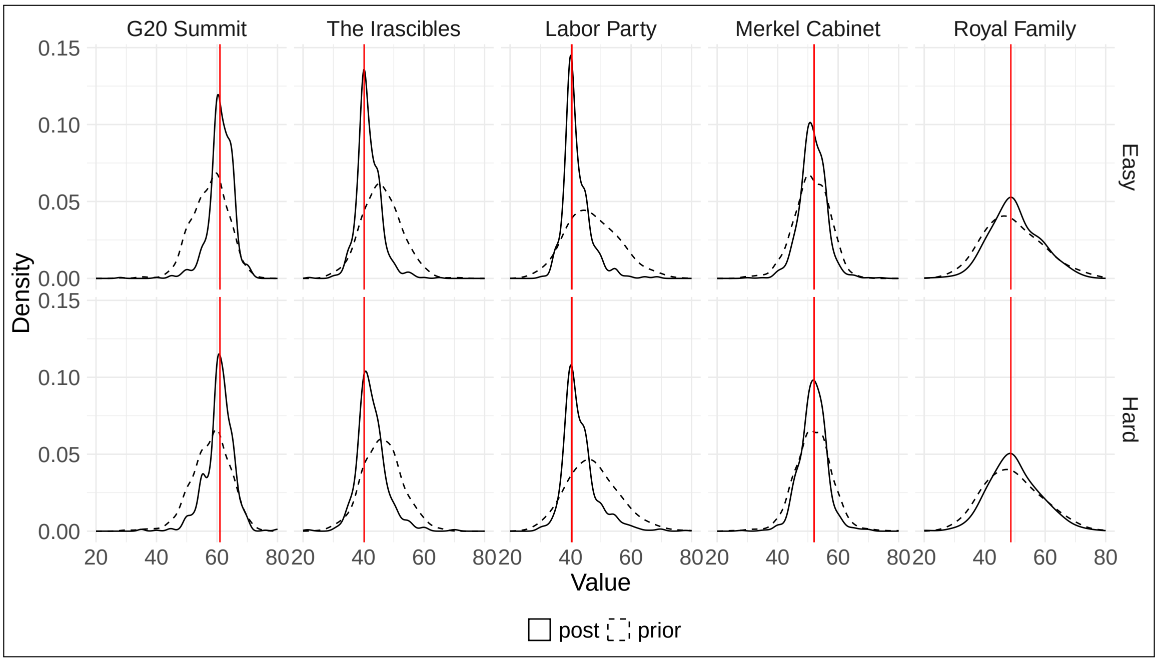

Figure B.5: Density of the answers for each picture. The vertical red line marks the correct answer. Note that the vertical axis differs for each question.

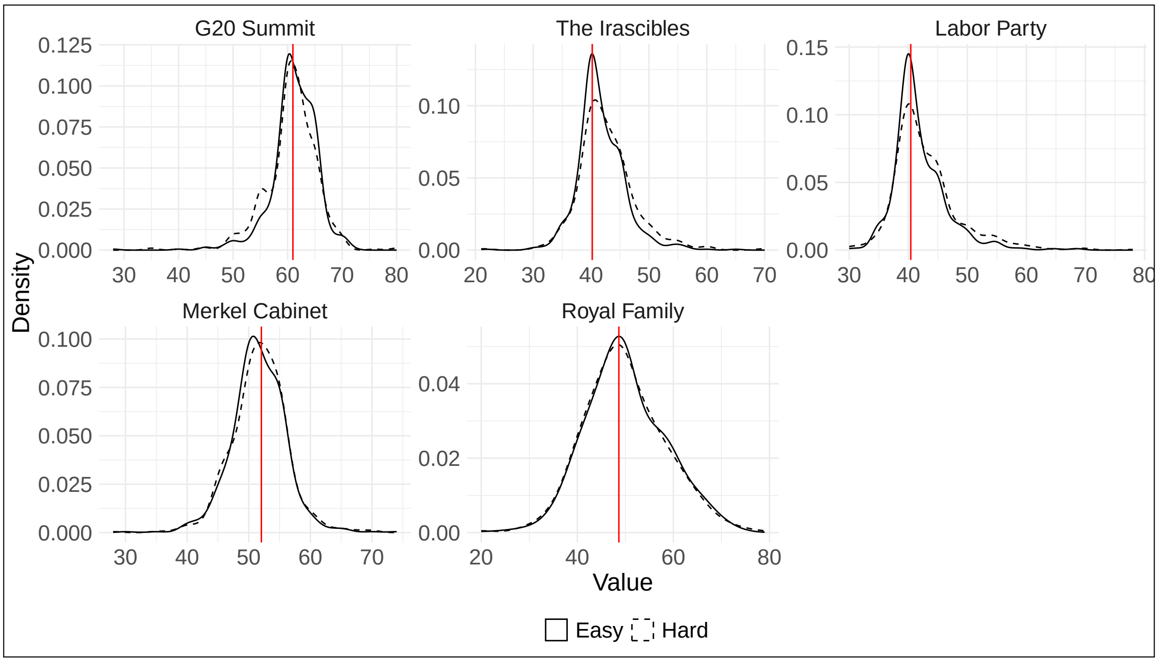

Figure B.6: CDF's of the adjustment of the prior and error of the posterior across treatments

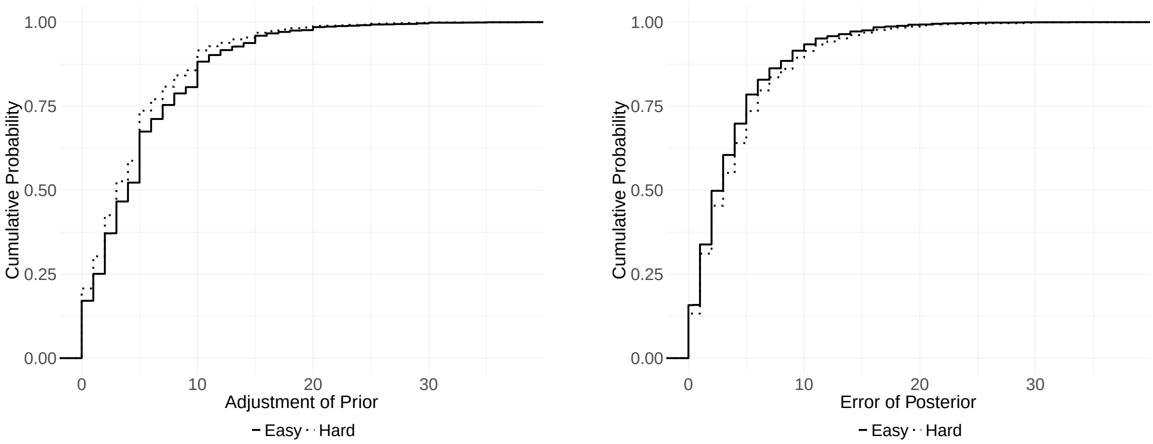

Note: The left panel plots the CDF of the adjustment made to the prior after seeing the signal across the *Easy-Info* (continuous line) and *Hard-Info* treatments (dotted line). The right panel plots the CDF of the error made by participants in their posterior belief in the *Easy-Info* (continuous line) and *Hard-Info* treatments (dotted line).

D Experimental Material

D.1 Pictures

D.1.1 First Block

Figure C.1: The Solvay Conference (1927)

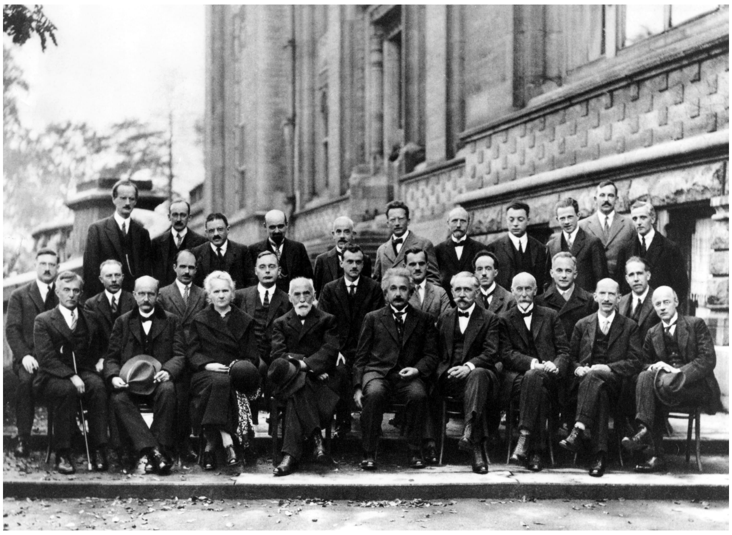

Figure C.2: The Band Aid (1984)

Figure C.3: The Cabinet of the US President B. Obama (2009)

Figure C.4: 75th Anniversary of Paramount (1987)

Figure C.5: The Chinese-German Consultations (2023)

D.1.2 Second Block

Figure C.6: G20 Hamburg Summit (2017)

Figure C.7: The Irascibles (1951)

Nina LeenTime Life Pictures/Getty Images

Life magazine's portrait of the Abstract Expressionist artists known as 'The Irascibles,' 1951. Front row: Theodore Stamos, Jimmy Ernst, Barnett Newman, James Brooks, and Mark Rothko; middle ow: Richard Pousette-Dart, William Baziotes, Jackson Pollock, Clyfford Still, Robert Motherwell, and Bradley Walker Tomlin; back row: Willem de Kooning, Adolph Gottlieb, Ad Reinhardt, and Hedda Sterne

Figure C.8: The Australian Labor Party (1901)

Figure C.9: The Cabinet of German Chancelier A. Merkel (2019)

Figure C.10: The British Royal Family (2007)