

Mechanisms of Mistrust: A Bayesian Account of Misinformation Learning

Lion Schulz ^{1*}, Eric Schulz ¹, Rahul Bhui ^{2,3}, Peter Dayan ^{1,4}

1 Max Planck Institute for Biological Cybernetics, Tübingen, Germany

2 Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, USA

3 Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA, USA

4 University of Tübingen, Tübingen, Germany

* Corresponding author. lion.schulz@tue.mpg.de

Abstract

From the intimate realm of personal interactions to the sprawling arena of political discourse, discerning the trustworthy from the dubious is crucial. Here, we present a novel behavioral task and accompanying Bayesian models that allow us to study key aspects of this learning process in a tightly controlled setting. In our task, participants are confronted with several different types of (mis-)information sources, ranging from ones that lie to ones with biased reporting, and have to learn these attributes under varying degrees of feedback. We formalize inference in this setting as a doubly Bayesian learning process where agents simultaneously learn about the ground truth as well as the qualities of an information source reporting on this ground truth. Our model and detailed analyses reveal how participants can generally follow Bayesian learning dynamics, highlighting a basic human ability to learn about diverse information sources. This learning is also reflected in explicit trust reports about the sources. We additionally show how participants approached the inference problem with priors that held sources to be helpful. Finally, when outside feedback was noisier, participants still learned along Bayesian lines but struggled to pick up on biases in information. Our work pins down computationally the generally impressive human ability to learn the trustworthiness of information sources while revealing minor fault lines when it comes to noisier environments and news sources with a slant.

Author summary

We are bombarded with information. But how do we learn whom to believe and whom to mistrust? For instance, how do we come to trust one news source's report, while believing that another is biased, produces only useless noise, or might even be lying? And how do we incorporate such possibilities when updating our beliefs? Our work offers a computational and empirical perspective on this learning process. We developed a novel and well-controlled task that allows us to characterize human learning about a host of information sources. We show how people can sometimes be remarkably able to discern lying and helpful sources, even when receiving only uncertain outside feedback. We also show how participants need clear feedback to learn about a news provider's slant.

1 Introduction

We are luckily not alone in a complex and uncertain world, and can rely on others for information (1; 2; 3; 4). Friends guide us through social life. Media guides us through political life. While these sources can be helpful for learning, adaptive creatures also need to be wary of the information they receive from other agents (5). This creates a dual learning problem: We not only need to learn about the issue at hand (“Who is the better candidate for mayor?”) but also about the trustworthiness of an information source (“How reliable is the local newspaper?”).

Formally, this dual learning problem can be understood as Bayesian inference operating on two levels. On the lower level, we need to learn about the state of the world, say the mayor’s qualifications. On the higher level, we need to learn about how the evidence we receive from an information source, say the local TV station, lets us make inferences about the state of the world. This Bayesian formulation of learning whether to trust an information source is present in many fields, including in formal epistemology (5; 6), cognitive science (7; 8), economics (9; 10), and applied trust and safety research (11; 12).

Whether we solve this dual learning problem appropriately is crucial in many walks of life: For example, when we trust the wrong news sources online, we risk being led astray and falling for misinformation (13; 14). At the same time, mistrust in the right experts may leave us uninformed – as in the discourse on climate change (15; 16; 17). At least equally important, psychopathology is tied closely to the way we view what others tell us. High and adaptive “epistemic trust” has been proposed to be one of the most powerful contributors to good mental health (18; 19). Conversely, out-sized credulity can leave us exposed to exploitation (20; 21; 22; 23).

Both personal and political life requires us to learn about information sources of diverse shapes: First, even generally helpful sources are often noisy themselves and may differ in their reliability (7; 24; 25): For example, news websites may generally strive to report the truth but nevertheless differ in their journalistic standards (26). While we should optimally still generally believe their evidence, we should learn to weight more the more reliable source (27).

Another group of sources may be considered unhelpful or untrustworthy: Within these sources, we can draw a crucial distinction: Following a classification by Frankfurt (28), “bullshiting” sources merely produces random noise. In contrast, for a source to “lie”, it needs to possess some underlying knowledge of the world state. Adaptive agents respond differently to these categories: “Bullshit” can safely be ignored, like trash websites online (29). “Lies” in turn can have informational content, and thus license more complex inferences (30; 31; 32). For example, we may interact with a source that constantly tries to confuse and reliably points to a wrong option – letting us infer another option is better, at least in a binary world (8; 33).

Finally, information sources can also have biases or slant: For example, media outlets might preferentially laud one candidate. Indeed, media bias is common and its prevalence and formation has been studied both empirically and theoretically (9; 34; 35; 36; 37; 38). Savvy agents should again take such bias into account: For example, a left-leaning outlet’s endorsement of a progressive candidate carries little information. In turn, a Republican newspaper’s endorsement of a Democratic candidate should carry more weight.

Empirically, considerable research has investigated how well people are able to make such inferences from information sources, often under the mantle of advice-taking (7; 8; 33; 39; 40; 41; 42). Advice-taking is involved in a wide range of tasks where we might be uncertain about our own beliefs, from trivia questions (43) to perceptual judgements (7) and moral decisions (44). The takeaway is that humans are generally able to learn how useful a source is but this ability is also subject to several limitations.

However, little work robustly characterizes the individual trial-by-trial learning computations underlying this dual inference and has investigated unifying learning dynamics underlying the diverse type of sources outlined above.

One crucial determinant for whether we are able to learn about a news source is our level of insight into the ground truth itself: When we have no knowledge about a subject, we are fully at the mercy of a source. In contrast, when we are experts ourselves, or the ground truth is directly observable (like the weather (45)), we can fact-check a source on the go, allowing powerful and quick inferences. Both theoretical Bayesian modeling and empirical investigations support this critical role of having a good feedback signal (5; 6; 7).

Here, we investigate people’s ability to learn about the trustworthiness of information providers (styled as ‘news sources’) in a novel task and capture this behavior using Bayesian models. We confront our participants and models with a diverse set of news sources, including helpful, entirely random (“bullshit”), as well as lying and biased sources. We administer two versions of this task. In one, participants learn with full feedback; in the second, they only receive partial feedback, necessitating semi-supervised learning. The tightly controlled, albeit artificial, set-up of our task lets us control rather precisely the information participants receive on a trial-by-trial basis and therefore lets us pin down their learning. This setup also precludes the influence of prior knowledge or motivated reasoning that might be associated with real-world sources (15; 46; 47). As a result, we can closely track, isolate, and quantify participants doubly Bayesian learning process.

To preview, we show that people are generally able to distinguish between different types of news sources and learn to respond to them adaptively, in partial accordance with Bayesian models. We show how this ability is attenuated when they only receive partial feedback, that participants show a bias towards believing a source is initially helpful, and that they struggle to pick up on biased information sources.

We begin by formalizing the computational problem and describing our paradigm at a high level before offering analyses of the two versions of our task.

2 Methods and Results

2.1 Computational problem and paradigm overview

Imagine you need to make a decision between two political options, such as deciding which policy might be better for your country. We can describe this as a signal detection problem where you need to identify a state s_t (which policy is better) on each trial t . Crucially, instead of having perfect knowledge about the state, you only receive noisy information about this, which we here denote by X_t . This can, for example, represent a news report. From this information, you can form a probabilistic belief about the candidates, denoted by the posterior $p(s_t|X_t)$, using Bayes rule:

$$p(s_t|X_t) \propto P(X_t|s_t)p(s_t) \tag{1}$$

This simple Bayesian update is straightforward when we know how good the news report is, that is how X_t is produced from s_t . As we alluded to, however, we often lack information about exactly this. As a result, we need first to learn how informative the likelihood $P(X_t|s_t)$ is.

We developed a task that captures this learning problem. Participants were told they were citizens of an alien planet that had two types of policies, green and blue ($s_t \in (green, blue)$). On each trial of the paradigm, which policy was actually better was chosen randomly with equal probabilities. Participants saw a piece of information $X_{I,t}$ from an imaginary news station I pertaining to the quality of the policies, and

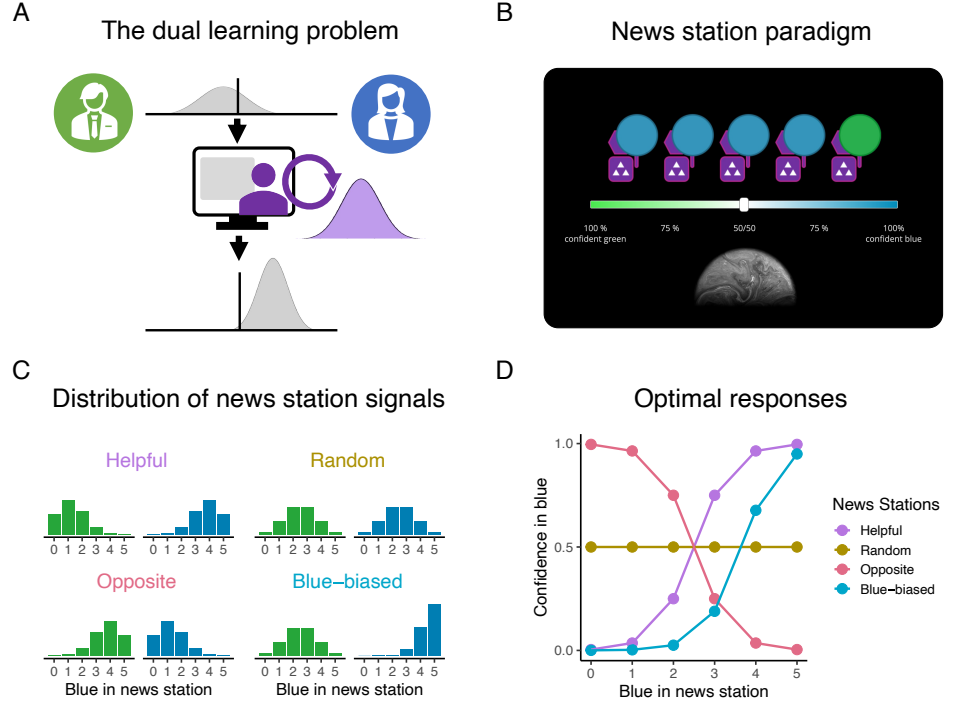


Fig 1. Illustration of paradigm. (A) The dual learning problem underlying how we come to learn about the trustworthiness of a news source. (B) Illustration of the news station task with feedback. Feedback is provided by the planet below the confidence slider lighting up in blue or green after participants have registered their choice. (C) Distributions of blue endorsements in the news stations for the different sources when the ground truth is the green state (in green) and the blue state (in blue). (D) Optimal response curve to the different number of blue endorsements depending on the source.

then had to choose which one was better, and rate their confidence in this choice (see Fig. 1 for an overview).

The news report came in the form of a “News Station Council” which consisted of a panel of $n = 5$ experts that endorsed either blue or green. Here, we let $X_{I,t}$ denote the number of blue endorsements out of the 5 experts so that $X_{I,t} \in \{0, 1, 2, 3, 4, 5\}$. This number of blue endorsements was distributed according to a binomial distribution so that each expert had an independent probability of endorsing blue or green. Crucially, this probability was a function of the news station and the state:

$$X_{I,t} \sim \begin{cases} B(n, b_I) & \text{if } s_t = \textit{blue} \\ B(n, 1 - g_I) & \text{if } s_t = \textit{green} \end{cases} \quad (2)$$

While this is a relatively simple statistical set-up, it allows us to model a host of different information sources. In our experiment, participants were confronted with four distinct sources whose signal distributions we plot in Fig 1 C. We note how these sources generally follow the forms of information providers we outlined in the introduction:

- A **helpful** news source with $b_I = g_I = 0.75$. This source produces more blue endorsements when the blue candidate is better and more green endorsements when the green candidate is better.

- A **random** source with $b_I = g_I = 0.5$. This source is entirely uninformative and does not allow any inferences about the state from its signals.
- An **opposite** source $b_I = g_I = 0.25$. While this source is theoretically as informative as the helpful source, it is inversely so, producing more *green* endorsements when *blue* is in fact the better candidate, and vice-versa.
- A **blue-biased** source with $b_I = 0.9$ and $g_I = 0.5$. In essence, this source is excellent at endorsing a good blue policy but useless at identifying a good green policy. As a result, its overall distribution marginalized over states is skewed towards blue.

How should participants optimally react when interacting with such news stations? Fig. 1 D outlines this, showing the optimal confidence that the blue policy is better, $p(s_t = \textit{blue}|X_t)$, for all possible signals X_t and sources: For the helpful source, the optimal response is straightforward: the more blue endorsements an agent encounters, the more confident it should be in the blue policy. For the random source, news reports should be entirely ignored, that is, the response should always be 50%—recall that there is no information in “bullshit”. For the opposite source, we should optimally invert the response to the helpful source, being more certain that blue is better when we see more green—following Frankfurt, the opposite source represents a simple form of lying. The blue-biased source offers the most complex pattern: While more blue endorsements should still optimally lead to more confidence in the blue candidate, the response curve is shifted: More blue endorsements are necessary to convince an optimal agent of the blue policy than in the helpful case. Indeed, majority endorsements for blue that are only timid should be met with the belief that the green candidate is better: Imagine a US right-wing website that only sheepishly endorsed a Republican candidate over a Democratic candidate. Crucially, participants were only informed about the general fact that news sources might be more or less informative or biased but were not told about these specific sources.

We can think of the response to a news stations as a psychometric function and can capture this psychometric function via a logistic regression. This regression gives us two parameters that succinctly describe an agent’s or participant’s responses. These parameters should differ between sources: The model’s slope describes how strongly and in what direction an agent integrates a source’s information. It should be positive for the helpful and blue-biased source, meaning the agent increases their confidence in blue being correct the more blue endorsements they see. In turn, the slope should be zero for the random source, meaning that an agent ignores its information, and negative for the opposite source, signifying an inversion of the evidence. The second parameter, the model’s intercept, should be equivalent for all but the blue-biased sources, in whose case it should be negative, describing its shifted response criterion. For convenience, we will subtract 2.5 from the number of blue endorsements so that an unbiased response would have an intercept of 0.

2.2 Learning and response model overview

How would an optimal agent come to give these responses? Generally, we can extend the Bayesian update from Equation 1 to take into account uncertainty about the parameters b_I and g_I of the sources. In more detail, this means that an agent will have a current belief over these parameter, $p(b_{I,t}, g_{I,t}|\mathcal{H}_t)$, that is based on its learning history, $\mathcal{H}_t \in \{X_{t-1}, Y_{t-1}, \mathcal{H}_{t-1}\}$. To update its beliefs about the underlying state given evidence X_t , the agent then marginalizes over this distribution when doing the Bayesian update. In the remainder of this section, we suppress subscript I for readability and

suppose that this learning proceeds equivalently, and independently, for all sources. The state update proceeds as follows:

$$p(s_t|X_t, \mathcal{H}_t) \propto \int_{b_t} \int_{g_t} p(b_t, g_t|\mathcal{H}_t)P(X_t|s_t; b_t, g_t)p(s_t) db_t dg_t \quad (3)$$

Now, to arrive at its belief over the parameters, the second part of the dual learning problem needs to be solved: the higher-level updating. As we mentioned in the introduction, at least some feedback signal or additional knowledge is necessary for this learning to succeed (5; 6; 7). Here, we denote this feedback by Y_t . In this paper, we consider two types of feedback: One where the feedback signal is a full revelation of the ground truth on each trial, $Y_t = s_t$, and another where the signal is merely a noisy read-out of the state, being produced by a distribution ϕ that is a function of the ground truth, so that $Y_t \sim \phi(s_t)$.

Regardless of feedback type, the optimal learning in both cases also proceeds in a Bayesian fashion. Specifically, the agent combines its prior estimate $p(b_t, g_t|\mathcal{H}_t)$ and current information X_t and Y_t to form an updated belief about the two parameters (again dropping I):

$$p(b_{t+1}, g_{t+1}|X_t, Y_t, \mathcal{H}_t) \propto \sum_{s_t} P(X_t|s_t; b_t, g_t)p(b_t, g_t|\mathcal{H}_t)P(Y_t|s_t) \quad (4)$$

This form is equivalent in the full and noisy feedback case but $P(Y_t|s_t)$ takes on different shapes. In the case of full feedback, it merely becomes an indicator $P(Y_t|s_t) \sim \mathbb{1}(s_t = \text{blue})$ so that b_I only gets updated when the state is blue and g_I only when the state is green. In the case of the noisy feedback, $P(Y_t|s_t)$ is then the conditional probability under the model $\phi(s_t)$ and differently weights the updates of b_I and g_I . In brief, in the full feedback case, this update takes the form of a simple and independent closed-form beta-binomial update for b_I and g_I whereas the noisy feedback case requires a more elaborate updating over their joint distribution with no closed-form solution.

To capture participants' expected idiosyncrasies better, we fit this model to their behavior with a total of 9 parameters. These parameters capture a range of phenomena: For example, given the Bayesian set-up of our model, we fit an initial prior to participants, $p(b_{t=0}, g_{t=0})$, letting them be more or less trusting in the early interaction with a news source. We also implement minor biases in the model's learning, for example letting them update quicker towards helpful or unhelpful beliefs, or letting them forget. We note that while these parameters introduce significant degrees of freedom to our model, they nevertheless situate the learning in a Bayesian regime while allowing us to bridge the gap between data and model. We detail both the optimal model computations as well as the parameter fitting in the SI.

2.3 Ethics statement

Participants in both experiments provided typed informed consent in accordance with procedures approved by the Ethics Committee of the Medical Faculty and Medical Clinic at the Eberhard-Karls-University of Tübingen (approval number 734/2019BO1).

2.4 Experiment 1: Full feedback

2.4.1 Experimental details

In the first version of our task, participants received full feedback about the true better policy at the end of each trial. In each trial, participants first saw the news station consisting of five experts and had to rate their confidence. They responded using a

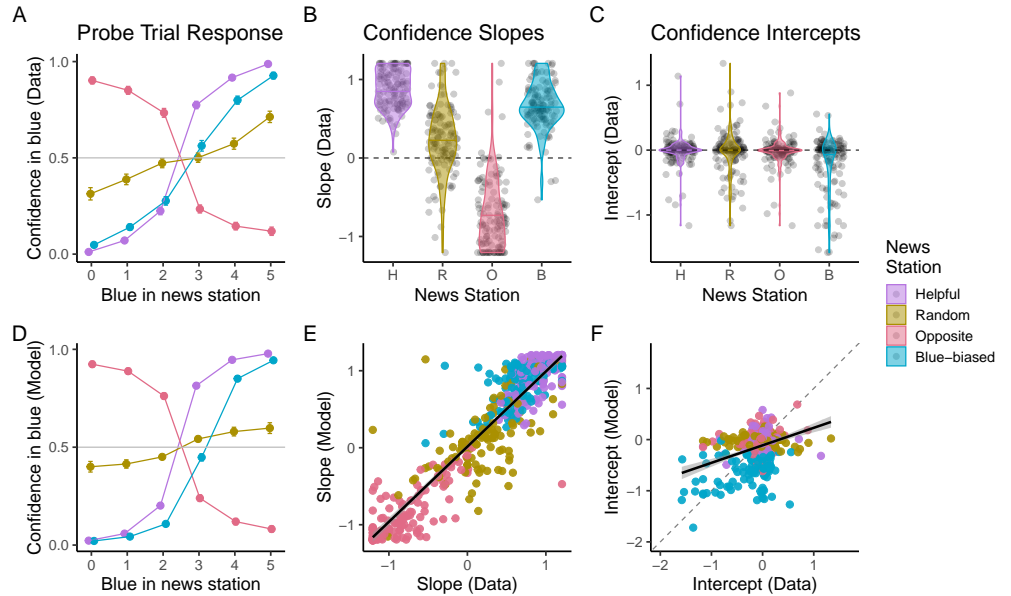


Fig 2. Probe trial results show good learning convergence and fit of Bayesian model in full feedback task. (A) Participants’ psychometric responses to the sources in the probe trials. Dots represent means and error bars represent standard errors of the mean (partially occluded by the dots). (B, C) Distributions of slopes and intercepts fit to individual participants’ probe trial responses. (D) Probe trial predictions of fit model. Dots represent means and error bars represent standard errors of the mean (partially occluded by the dots). (E, F) Correlations of slopes and intercepts fit to individual data and model probe trial responses.

colored slider from “100% sure green” to “100% sure blue” (via “50/50”) (see Fig. 1 B). After they had rated this confidence, participants received feedback about the ground truth. Participants played with each of the sources for a total of 28 trials. We presented the sources in a blocked manner, randomizing the order between participants.

To test participants’ beliefs in the sources in a standardized manner, each block ended with an additional six demarcated probe trials where participants saw all possible constellations of blue and green endorsements $X_{I,t} \in \{0, 1, 2, 3, 4, 5\}$ in random order. Participants did not receive feedback during these trials. At the end of each block, we additionally asked participants how much they thought that the news station improved their decision and how much they trusted the news station. At the end of the experiment, participants filled out a battery of questions relating to trust.

We analyze data from a total of 123 US adult participants collected via Prolific with a broad range of ages and educational backgrounds (see SI for more details).

2.4.2 Results: Probe trials

For a high-level overview of participants’ learning success, we first investigated their responses on the final probe trials. Recall that these showed each participant all possible combinations of source endorsements and so let us investigate their full ‘revealed’ belief in the source.

We plot these responses in Fig. 2 A showing the average responses across participants. This reveals that participants were able to pick up on the source’s

attributes (we provide a statistical analysis below): For the helpful source, on average participants were more confident that blue was better the more blue endorsements they saw. In contrast, for the random source, this integration was diminished, and participants tended to relatively ignore its signals. Participants also inverted the evidence that they received from the opposite source, being more confident in *blue* the more *green* endorsements they saw. Finally, we observed a marked shift in the participants' responses to the blue-biased source: When participants only saw a weak endorsement of blue ($X_I = 3$), their confidence was on average almost indifferent between blue and green.

As we outlined, we can describe this behavior via a logistic regression. We thus fit individual regressions to participant responses (see SI for details). Recall that the regression slope should be positive for the helpful and blue-biased source, zero for the random source, and negative for the opposite source. The second parameter, the model's intercept, should be equivalent for all but the blue-biased sources, in whose case it should be negative. In line with our across-participant analysis, we found that slopes fit to participants' probe trial responses (Fig. 2 B) were positive for helpful and blue-biased sources, lower for the random sources, and negative for the opposite source. This was supported by a significant main effect of source on slope in an ANOVA ($F(3, 488) = 435.40, p < .0001$) and significant differences in Tukey's tests between all sources (all p 's < 0.0001). Participants' slopes for the helpful, blue-biased, and random source were all significantly higher than zero (one-sample t-tests, all p 's < 0.0001), indicating that on average participants still held the random source to be slightly helpful. The opposite source's slopes were significantly lower than 0 (one-sample t-test, $t(122) = -17.10, p < .0001$), indicating that participants on average inverted its evidence, at least slightly.

In turn, there was no significant difference between the intercepts (Fig. 2 C) of the helpful, random, and opposite source (Tukey's tests, all p 's > 0.9) and they were all not significantly different from zero (one-sample t-test, p 's > 0.2). However, the intercept of the blue-biased source differed significantly from all the other sources (ANOVA testing main effect of source, $F(3, 488) = 11.82, p < .001$, and Tukey's tests comparing blue-biased means to other means of other sources, all p 's < 0.0001), and was significantly lower than 0 (one-sample t-test, $t(122) = -5.54, p < .0001$). However, while the intercept distribution was generally shifted negatively for the blue-biased source, we also observed that most participants remained within the ranges of the other three source types.

Participants' responses were well captured by our Bayesian model. This is visible in Fig. 2 D where we plot the Bayesian model's predictions. Our model was also able to capture individual patterns of responses. To check this, we fit the same logistic regressions to the fit model posterior predictions. We observed strong correlations between the psychometric slopes obtained from the data and model ($r = 0.92, p < 0.0001$). These slopes are the main differentiator between the four sources and thus the main indicator for learning success. We found a weaker albeit still significant correlation between the model and data for the intercept ($r = 0.36, p < 0.0001$). This is unsurprising, given that a considerable number of participants did not, as we discussed, manage to realize there was a bias.

2.4.3 Results: Learning dynamics

Our analysis of the probe trials showed that participants were generally able to learn the statistical regularities of these sources - but how did they arrive at these conclusions? To answer this, we next analyzed participants' learning trajectories over the first 28 trials of a block.

First, we again investigated participant averages, tracking their beliefs across the

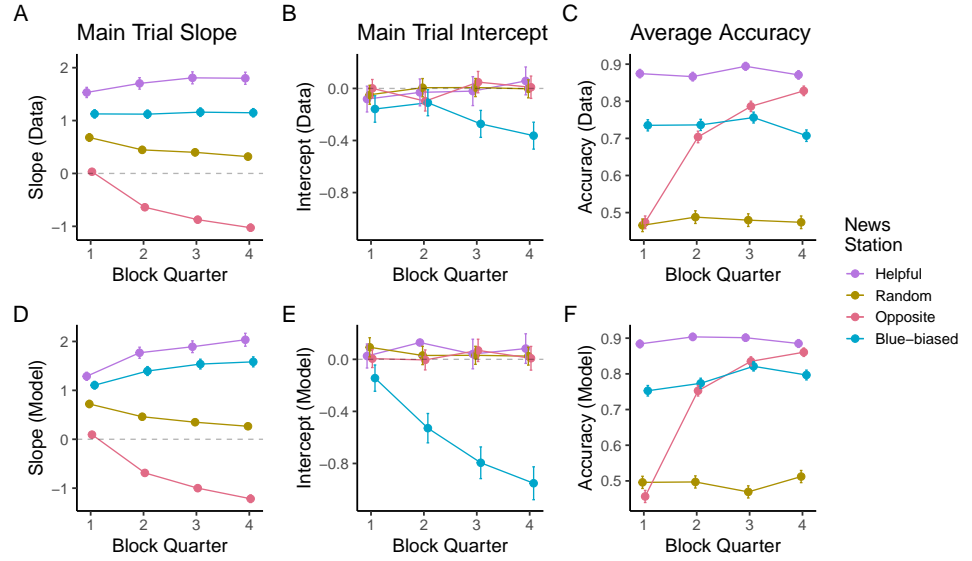


Fig 3. Pooled regression analysis shows signatures of Bayesian-like model in full feedback version. (A, B) Slopes and intercepts fit to pooled participant responses, split by block quarter and source. (C) Accuracy split by block quarter and source. (D - F) The same metrics inferred from fit model. Dots in A, B, D, and E represent the mean regression coefficient estimates, and error bars represent their respective standard errors (partially occluded by the dots). C and D show mean accuracies and the error bars standard errors of the mean (again partially occluded).

interaction with a source. To do so, we fit logistic regression models akin to the ones described above for each source individually, splitting each block into quarters of seven trials and pooling the responses across participants.

Fig. 3 A and B show the results of this analysis: Already in the first quarter, there are markedly different responses to the four sources, with the blue-biased and random showing lower, albeit positive, slopes than the helpful source. The opposite source in turn shows a slope close to zero (panel A). This distinction develops further across the block with the opposite source’s slope becoming negative — that is on average participants began inverting its evidence rather than just ignoring it — and the random source showing a slow decline towards zero — that is participants began to ignore its evidence more. There was a mirrored slight increase in the helpful source’s slope, showing a slightly larger integration of its evidence towards the end of a block. In turn, the average intercepts remained equivalent for all but the blue-biased source, whose intercept declined towards the end of the trial. This lowering intercept reflects learning that the blue-biased source was indeed biased.

Fig. 3 C shows the effects of this learning on participants’ average accuracies, computed by binarizing their confidence ratings. Initially, participants were able to glean no information from the opposite source, as shown by their random performance in the first block quarter. However, as their interaction with the source went on, their judgement also became more accurate, almost reaching the accuracy of the helpful source by the end of the block. Recall that the opposite source in theory conveys just as much information as the helpful source. This general accuracy pattern was supported by a significant source \times trial interaction in a logistic regression predicting accuracy ($\chi^2(3) = 239.43, p < 0.0001$)

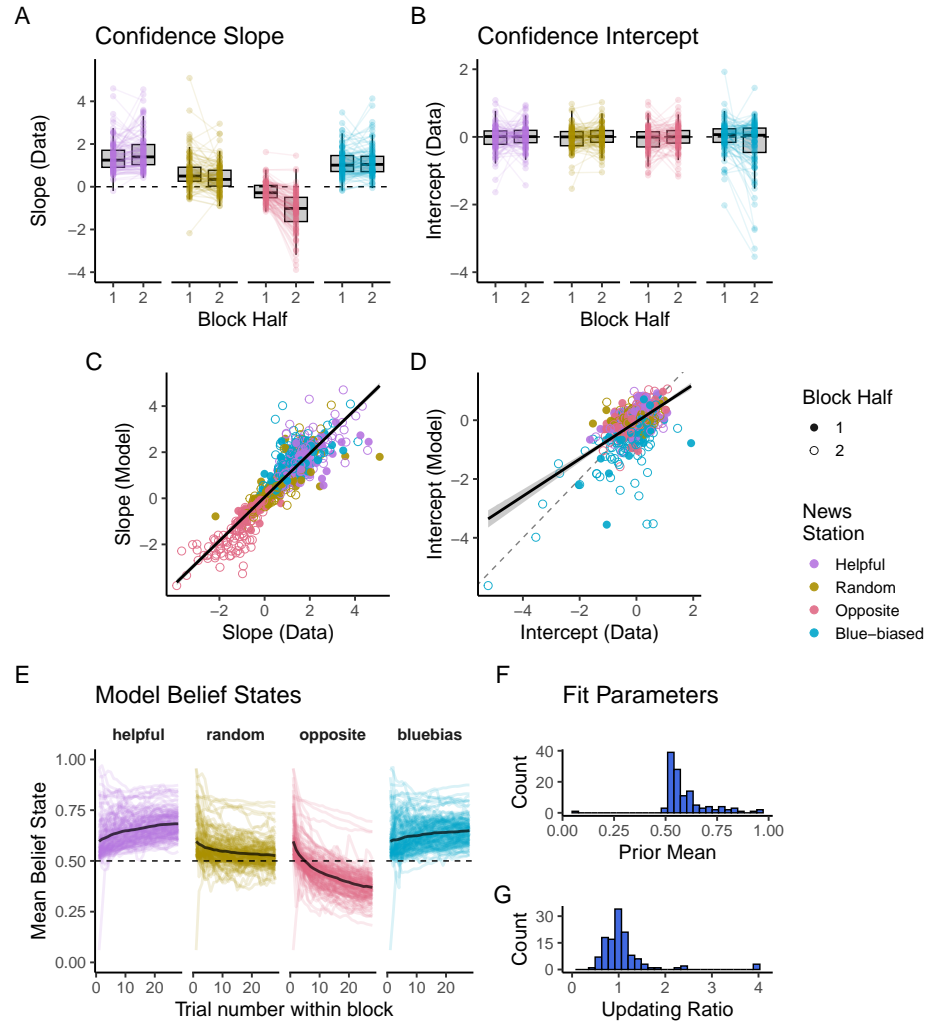


Fig 4. Individual regression analysis highlights learning and model fit in full feedback version. (A, B) Slopes and intercept of psychometric curves fit to first and second half of participant responses. (C, D) Correlation between first and second half slopes and intercepts fit to data and fit model. (E) Trajectory of mean over the two belief states. Coloured lines represent individual participants, and black line mean over participants. (F) Distribution of mean over two prior parameters. (G) Distribution of updating ratio.

Our model was able to capture these dynamics as is evident in Fig. 3 D-F. There, we show the learning dynamics for the posterior predictions of the model. All response patterns are well recovered: Both participants and model are able to develop distinguishable responses to the sources already in the first half, already ignoring the opposite source (Fig. 3 D). In line with our analyses of the probe trials, we also find that the model is able to pick up quicker (and more strongly) on the bias than the participants (Fig. 3 E). Accuracy patterns, however, were comparable between the model and data (Fig. 3 F).

To investigate whether these learning patterns held on an individual level, we next fit logistic regressions individually to the data, although this time split by block half

rather than quarter to ensure a sufficiently large sample per regression. We plot the results of these analyses in Fig. 4 with the slope and intercept values in A and B. We observed a significant two-way source \times block half interaction for both slope and intercept (slope: $F(3, 976) = 27.17, p < .0001$, intercept: $F(3, 976) = 5.99, p < .0001$).

Crucially, participants' slopes in response to the opposite source were significantly lower in the second half of a block compared to the first (post-hoc Tukey's test, $p < 0.0001$). Participant's intercepts showed a decrease when interacting with the blue-biased source ($p = 0.0005$). Despite trends in the expected directions, there was no significant difference between participants' first and second half slopes when interacting with the opposite ($p = 0.30$) and random source ($p = 0.25$). Slopes also did not differ for the blue-biased source ($p > 0.99$). As we would expect given the lack of bias, intercepts for helpful, random, and opposite sources also did not change between the first and second half ($p's > 0.9$).

Fig. 4 C shows that participants' individual patterns were also well captured by our fit model: The slopes we recovered for the two halves again showed a high correlation between model predictions and data ($r = 0.90, p < 0.0001$), showing how participants and models similarly picked up on the correct response patterns. As before, this correlation was still significant for the intercept, although to a lesser degree ($r = 0.57, p < 0.0001$), highlighting the difficulty participants had in learning about the bias (Fig. 4 D).

The model's internal belief states revealed the individual learning process on a more fine-grained level. In Fig. 4 E, we plot these beliefs, that is, what the model assumes to be the binomial probabilities with which the source generates its information. We plot the mean over the two beliefs (formally, $\mathbb{E}[\mathbb{E}[b_{I,t}] + \mathbb{E}[g_{I,t}]]/2$) to highlight the learning process and note key similarities between this plot and the dynamics we displayed in the pooled regression analysis in Fig. 3 A and D.

We highlight two distributions of fit parameters which lets us quantify aspects of these individual learning dynamics: Fig. 4 F shows the fit prior belief (averaging the prior for b_I and g_I). This affirms our model-agnostic analyses using the slopes: Participants' prior tended to be positively skewed ($\mu = 0.60, \sigma = 0.11$). Additionally, the closed-form belief update within our model let us straightforwardly fit a learning bias, asking whether people tended to learn quicker in favor of believing that a source was helpful (updating ratio > 1) versus that a source was unhelpful (updating ratio < 1). This distribution centered around 1 ($\mu = 1.09, \sigma = 0.56$) indicating that participants did not update in a particularly hopeful or unforgiving manner. This was supported by a one-sample t-test that showed, despite outliers, no significant difference from 1 ($t(122) = 1.80, p = 0.074$).

2.5 Experiment 2: Noisy feedback

In the real world, unambiguous feedback is often lacking. In the second experiment, we investigated how well participants were able to cope with noisier feedback when learning how reliable or biased a source was.

2.5.1 Experimental details

Participants again interacted with the four different news sources, presented in a randomized blocked manner. However, instead of fully revealing the better policy at the end of each trial, participants had to rely on a noisier source of information, consisting of a panel of experts that were displayed in the same manner as the news station. Fig. 5 shows the time course of a trial: Participants again first saw the news station consisting out of five green or blue endorsements, and rated an initial confidence (using the same scale as above). Following this, they were shown the Independent Council, also

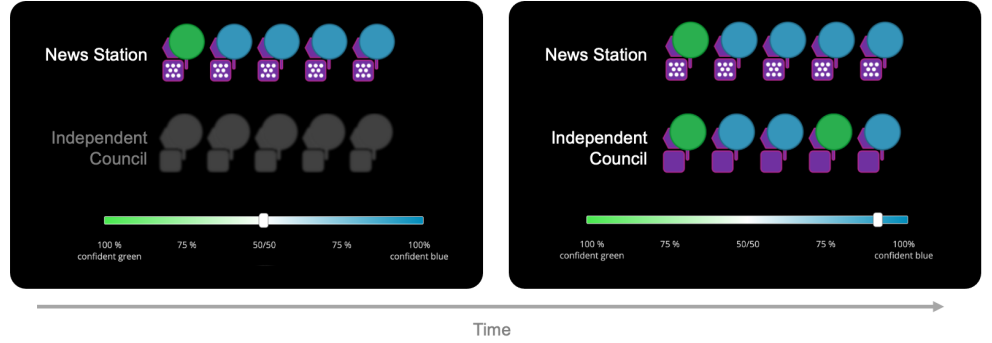


Fig 5. Illustration of trial in noisy feedback condition. Participants first saw the news station experts and rated their confidence based on this. They then saw the independent council and were able to adjust their confidence.

consisting of five green or blue stochastic endorsements. Based on this, they could revise their confidence, using the same slider.

Participants were told that the “Independent Council” was a helpful, albeit sometimes erring, collection of experts. Specifically, the statistical set-up of this expert panel was such that it was equivalent to the helpful source. Here, we denote the number of blue endorsements as Y_I and the parameter producing this evidence with b_Y :

$$Y_{I,t} \sim \begin{cases} B(n, b_Y) & \text{if } s_t = \text{blue} \\ B(n, 1 - b_Y) & \text{if } s_t = \text{green} \end{cases} \quad (5)$$

Here, we set $b_Y = 0.75$. Participants were extensively informed about this quantity. We conducted analyses to ensure that participants used the independent council’s feedback which we report in the supplementary information.

As before, participants interacted with each source for 28 trials before engaging with 6 probe trials. They also again answered the trust and improvement question at the end of each block and a battery of questions at the end of the experiment.

We analyze data from a total of 111 US adult participants, again collected via Prolific and including a range of ages and educational backgrounds (see SI).

2.5.2 Results: Probe trials

We again first analyzed the probe trials at the end of each block (Fig. 6 A-C). This shows how participants displayed individual patterns of responses that were broadly in line with an adaptive response but less so than in the full feedback case: The random source was again associated with a lower slope than the helpful source (Tukey’s test: $p < 0.0001$). However, on average, participants still integrated the random source’s information, with their average slopes being significantly higher than zero (see 6 B, one-sample t-test $t(110) = 11.74, p < .001$). On average, participants managed to invert the opposite source’s evidence ($t(110) = -3.63, p < .001$), but this inversion did not mirror in strength the helpful source. While we observed a minute shift in the response to the biased source the individual intercepts, shown in Fig. 6, were on the whole not significantly different from zero ($t(110) = -0.39, p = 0.696$).

Most of these patterns were well-captured by the Bayesian model, as is visible in Fig. 6 D - F. The model’s general response patterns (panel D) matched those of the participants. Again, this was highlighted by the strong correlation between the slopes recovered from fit model predictions and from the data (panel E, $R = 0.80, p < 0.0001$). In line with most participants’ inability to realise the blue-biased sources bias, we found

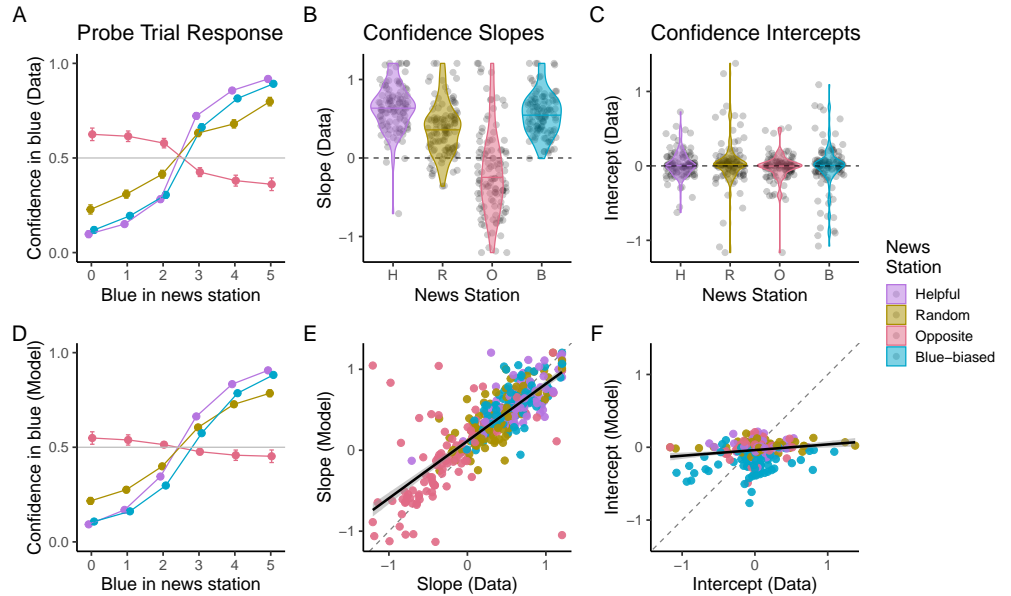


Fig 6. Probe trials in noisy feedback condition highlight diminished learning convergence and model fit. (A) Participants' psychometric responses to the sources in the probe trials. Dots represent means and error bars represent standard errors of the mean (partially occluded by the dots). (B, C) Distributions of slopes and intercepts fit to individual participants' probe trial responses. (D) Probe trial predictions of fit model. Dots represent means and error bars represent standard errors of the mean (partially occluded by the dots). (E, F) Correlations of slopes and intercepts fit to individual data and model probe trial responses.

only a weak correlation between the model predictions and participant data (panel F, $R = 0.17, p = 0.0003$).

2.5.3 Results: Learning dynamics

How did participants reach these beliefs? To investigate this, we again fit regressions to the main block and investigated their dynamics. Fig. 7 A and B show the results of this analysis. In the main block, participants' response slopes showed less distinction between the helpful, random, and blue-biased sources. Indeed, the slopes in response to the random source were only slightly lower than to the helpful source and only developed this difference over the course of the block. While the opposite source's slope is clearly distinguished from the remaining sources, participants' inverting response is not as strong as before, converging at, on average, ignoring the source rather than inverting its evidence. As to be expected from the probe trials, participants' average intercept in response to the blue-biased source barely shifted away from the remaining sources.

The effects of this slower learning are evident in Fig. 7 C where we plot the time course of the participants' average accuracies after having seen the news station (but not having seen the helpful independent council): The positive integration of the evidence coming from the helpful and blue-biased source results in high accuracies throughout the block. Inference from the random source naturally remains at 50 % accuracy, containing no information. In contrast, participants initially fall for the opposite source's misinformation, integrating their information in the incorrect manner:

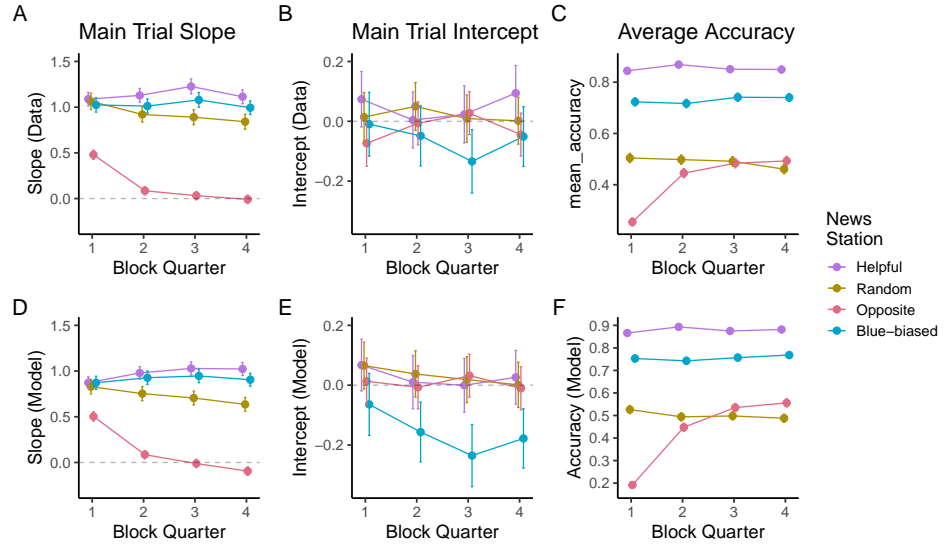


Fig 7. Pooled regression analysis shows signatures of Bayesian-like model in noisy feedback version. (A, B) Slopes and intercepts fit to pooled participant responses, split by block quarter and source. (C) Accuracy split by block quarter and source. (D - F) Same metrics inferred from fit model. Dots in A, B, D and E represent the mean regression coefficient estimates, and error bars represent their respective standard errors. Note how these error bars are partially occluded by the dots. C and D show mean accuracies and the error bars depict standard errors of the mean (again partially occluded).

This leads them to perform below chance accuracy in the first half of the block. Only in the second half of the block do they reach accuracy levels that align with the random source, in line with their ignoring of the opposite source’s information (panel A). This was again supported by a significant within-block trial number \times source effect on accuracy ($\chi^2(3) = 52.52, p < 0.0001$)

Our model was again able to capture these general dynamics as is visible in Fig. 7 D - F: The model’s response slope only slowly and weakly distinguished the helpful, random, and blue-biased source. The model initially weakly integrated the opposite source and then ignored it towards the end of the block (panel D). There was little distinction in the intercepts (panel E). These learning patterns were reflected in the model’s average accuracy (panel F).

Analyses of individual-level responses supported these results, as is shown in Fig. 8. We again found a significant block half \times source type interaction effect in an ANOVA predicting the slopes ($F(3, 880) = 3.43, p = 0.017$, see panel A). Again, participants’ response slope significantly decreased only for the opposite source between the first and second block halves (Tukey’s test, $p = 0.027$). All other sources saw no changes in slope (p ’s > 0.97). Conversely, there was no significant shift in intercepts, supported by a non-significant block half \times source type interaction effect in an ANOVA ($F(3, 880) = 1.56, p = 0.198$).

Our model captured these individual response patterns well, as is visible in panels C and D: Slopes and intercepts fit to the model’s responses correlated with those fit to participants data (slope: $R = 0.88, p < 0.0001$, intercept: $R = 0.59, p < 0.0001$).

These individual learning dynamics were also visible in the fit model’s belief states

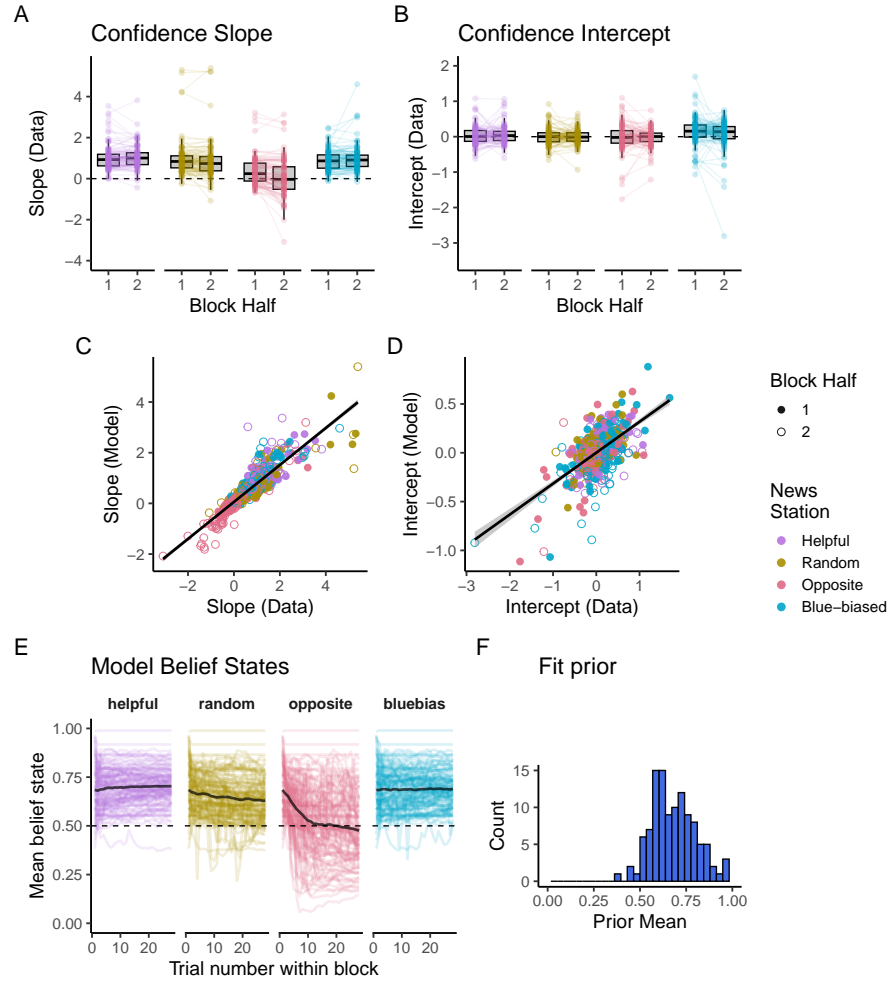


Fig 8. Individual regression analysis highlights learning and model fit in noisy feedback condition. (A, B) Slopes and intercept of psychometric curves fit to first and second half of participant responses. (C, D) Correlation between first and second half slopes and intercepts fit to data and fit model. (E) Trajectory of mean over the two belief states. Coloured lines represent individual participants, black line mean over participants. (F) Distribution of mean over two prior parameters.

which we show in panel E, again pooled across the two beliefs. Learning was mostly evident when interacting with the opposite source, and participants remained largely around their own helpful prior. Indeed, this prior was again skewed positively ($\mu = 0.68, \sigma = 0.11$). We note that due to the non-closed form nature of the belief update, we did not characterize the updating bias as comparable to the full feedback case.

2.6 Summary and experiment comparison

In both experiments, the Bayesian model described participants well but not perfectly. How did these conditions compare?

We first compared participants' average accuracies between the two versions: As to be expected, participants showed higher accuracy in the condition with full feedback,

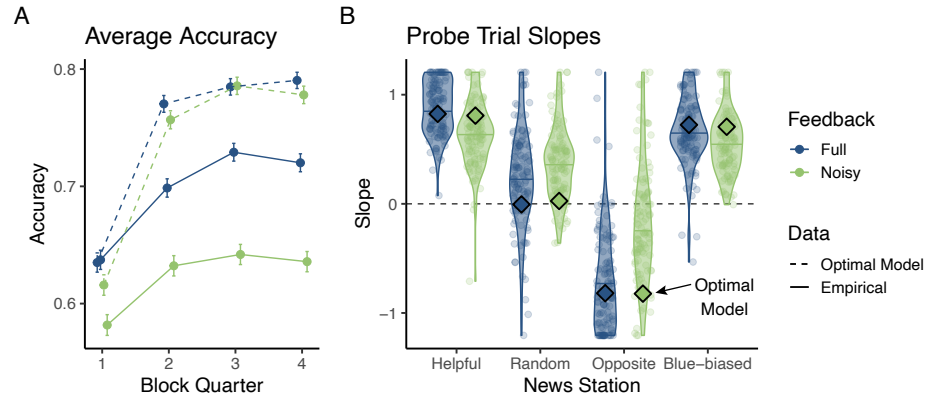


Fig 9. Full feedback version shows better empirical performance. (A) Average accuracy split by block quarter and experiment. We show both the average accuracy from the data (continuous line) and the average accuracy achieved by the optimal parameter-less model (dashed lines) operating with uniform priors. (B) Probe trials slopes by condition and news source. Mean slope values for the optimal model highlighted with framed diamonds.

and this difference increased over the course of the block (Fig. 9 A). This was supported by a significant trial number \times experiment condition effect ($\chi^2(1) = 19.422, p < 0.0001$).

We also checked how far this performance diverged from the optimal model. To do so, we also ran the full normative models without fitting any parameters and operating with flat priors over the initial source probabilities ($b_{t=0,I} = g_{t=0,I} = \text{Beta}(1,1)$). We overlay these performances in Fig. 9, panel A. This highlights that the optimal model significantly outperformed participants in both feedback conditions (significant trial number \times optimal model effect, $\chi^2(1) = 184.17, p < 0.0001$). We also note how the performance differences between the two feedback conditions were small for the optimal model, showing the theoretical inferential power of an optimal semi-supervised model. Indeed, the optimal model licenses powerful inferences in both cases and asymptotes to high performance after the second quarter. In contrast, empirical data showed how participants struggled more with the noisy feedback condition.

Participants' probe trial responses were also closer to the optimal response in the full feedback condition (Fig. 9 B): Their average slope was higher in response to the helpful source and more negative in response to the opposite source (Tukey's tests comparing the conditions, $p < 0.0001$). The slope did not differ between the conditions for the blue-biased source (Tukey's test $p = 0.40$). We again overlay the optimal model's slopes in panel B which shows how participants' average slopes were close to the optimal model in the full feedback condition. We note how the comparably worse performance in the learning trials (panel A) is due to the much higher variability of participants' slopes (for example some participants positively integrating the random, or ignoring the opposite source).

2.7 Questionnaire analyses

2.7.1 Post-block trust and improvement ratings

How might participants have assessed the news sources' qualities more broadly? To answer this, in both experiments, we asked participants two questions at the end of each block: One relating to how much they trusted the source and one relating to how much they thought the source improved their decisions. Specifically, participants responded




Fig 10. Participants rate trust in sources and how they improve their decisions. (A) Raw responses to post-block trust and improvement questions by source. (B, C) Relationships between second block half psychometric slopes and trust and improvement ratings. Regression lines fit jointly to helpful, random, and blue-biased sources, and separately to the opposite source.

on a continuous scale ranging from “Strongly Agree” to “Strongly Disagree” (via intermediate options) to the following statements: (1) “I trust this news station” and (2) “This news station improves my decisions”.

Analysis of participant responses revealed key distinctions between the sources (see Fig. 10 A): Throughout, participants rated the helpful source highest on both trust and improvement questions (Tukey’s tests against all other sources, $p < 0.0001$). The blue-biased source was similarly rated high, albeit lower than the helpful source in both trust and improvement questions. The random source in turn was rated as neither trustworthy nor as particularly improving decisions and had lower trust and improvement ratings than the helpful and blue-biased sources (Tukey’s tests $p < 0.00001$). A more curious pattern emerged in response to the opposite source: While it had the lowest average trust ratings (Tukey’s tests against all other conditions p ’s < 0.0001), some participants indeed rated it highly for improving their decisions, in line with its general usefulness. Indeed, while all other sources had equivalent trust and improvement ratings (Tukey’s tests p ’s > 0.99), the improvement rating of the opposite source was significantly higher than its trust rating ($p < 0.00001$), and as high as the improvement rating of the random source ($p = 0.9970$)

How did these responses relate to participants’ behavior on the task? To address this, we correlated participants’ psychometric slopes in the second block half to their trust and improvement ratings. We did so jointly for the helpful, random, and blue-biased source, and separately for the opposite source, for which we observed the trust-improvement distinction. The relationship between participants’ trust ratings and their response slopes was best captured by a simple linear regression, highlighting that the more participants integrated the sources’ information, the more they also rated it as trustworthy (Fig. 10 B, regression model comparison in the supplementary material).

This pattern was more complex for the improvement ratings: For the helpful, random, and blue-biased source, the relationship between slopes and ratings was still best captured by a simple linear model. However, for the opposite source, this relationship followed a linear quadratic function: This meant that while some participants who realised they had to invert the source’s information did not reflect this in their improvement ratings (negative slope, low improvement rating), other participants who had a similar realisation (negative slope) did reflect this in high improvement ratings.

2.7.2 Trust and media consumption patterns

How might behavior in our task relate to more real-world behavior and attitudes? To assess this, we administered a number of questionnaires relating to general interpersonal trust as well as trust in media in both experiments. Additionally, we investigated participants' media consumption patterns by asking them about their information sources and the frequency with which they engaged with these sources.

For a more principled approach to these questionnaires, we first performed an exploratory factor analysis on the entire set of 50 questions. This revealed a three factor structure: A first factor captured media consumption and trust, with consumption of (mainstream and social) media and trust in this media loading highly on this factor. A second factor represented both participants' levels of personal mistrust and credulity with respective questionnaire sub-scales displaying corresponding loadings. Interestingly, this factor also had weak negative loadings for the consumption of classical media and weak positive loadings for the consumption of social media. A final factor represented more general levels of interpersonal trust with questions around epistemic trust loading positively and questions around mistrust loading negatively (see supplementary information for more details).

We had no particular set of hypotheses regarding the relationships between these questionnaires and our task. We thus conducted a number of exploratory regression analyses linking the two. These analyses revealed no significant relationship between the factors extracted above and different behavioral components such as game scores, psychometric, or model parameters.

3 Discussion

We investigated people's ability to learn different statistical regularities of information providers. We framed this as a doubly Bayesian learning problem where agents learn about both a ground truth and the trustworthiness of a source, and developed a corresponding model. We showed how people were generally able to distinguish different kinds of information sources and respond to them adaptively. Participants integrated trustworthy information correctly. They also tended to learn to ignore useless information and invert information from a reliably incorrect source. Finally, they—albeit with limitations we discuss below—were able to uncover biases in news sources.

We conducted two experiments, in which participants received either full or noisy feedback. Across these tasks, we showed how participants followed the learning trajectories of a parameterized Bayesian model. Participants showed better performance in the task version with full feedback, in line with the more straightforward nature of the learning process involved. Our models described participants' learning patterns well. This is in line with previous conceptualizations of Bayesian learning about the trustworthiness of information sources, for example theoretically proposed in similar forms by (5) and (6) and empirically shown by (48) and (8).

Two behavioral signatures that differed from a purely optimal learning model emerged in both experiments. First, only few participants were able to learn the bias of the blue-biased source. This source is indeed the hardest to learn - requiring the tracking of two higher-order probabilities instead of one. Whether and how we deal with biased information is an ongoing question. For example, in a field experiment people adaptively took news bias into account for more accurate judgements (36), but an experimental study highlights errors in the way people seek information from biased sources (34).

Second, on average, participants tended to start out believing that our sources were helpful. This led them initially not to ignore the random source, or invert the opposite

source. While such behavior is suboptimal in this specific context, a general bias holding information providers to be useful may be generally adaptive. For example, most (social) media information we encounter tends to be true (49; 50), and people are sensitive to the general distribution of misinformation in an environment (51). Indeed, always starting to learn from scratch is wasteful, and holding sources to be generally useful would be a healthy and adaptive inductive bias in most non-adverse social environments (2; 52). We found a prior leaning more towards a “helpful” source in the more difficult task version with noisy feedback, indicating that participants might have relied more on the prior when learning was harder.

Our fit parameters allowed us to characterize some of these behaviors, along with providing other signatures: Participants used helpful priors, an inference also supported by model-agnostic analyses. In the full feedback condition, our model also allowed us to characterize participants’ learning as being neither particularly hopeful nor unforgiving. The latter is interesting given that other investigations have found optimistic and pessimistic biases in Bayesian and reinforcement learning (53; 54). Particularly when it comes to interpersonal relations, trust is however often purported to be hard to build and easy to lose (55; 56; 57). Similar effects have been shown with regard to news sources in more naturalistic experiments (58). The fact that we do not find such a bias here might hint at this being not a general bias but rather a function of actors in a social environment.

We note that the purpose of our model fitting was to show general alignment between Bayesian-like learning patterns and participant behavior and not to make strong statements about a particular class of model or parameters. Indeed, we rely on a medium to high number of parameters that are partially intertwined. Future modelling will have to disentangle these factors. For example, resource-rational accounts might explain why only some participants learnt the bias (59). To pin down specific effects such as the learning biases further, targeted experiments will also be necessary, for example, ones that might change the quality of a source mid-way through a block (33). This will allow for more robust modelling conclusions (60).

Our post-block questionnaires showed that participants were able to explicitly express their trust in the sources and how strongly they believed that the sources improved their decision-making. We showed how this related to behavior in our task. This points to the validity of our paradigm. The fact that these ratings also dissociate for the opposite source hints that people don’t purely understand trust to be a merely utilitarian construct (61). The ratings also demonstrate how people express divergent trust towards sources we may consider “bullshit” and those that lie (28). It would be interesting to consider how trust and utility trade off: For example, would we rather (pay to) hear from a useful but untrustworthy liar or a less useful but trustworthy source?

Our broader questionnaire measures did not correlate with individual differences in our task. There are several lenses through which to view this finding: First, our tightly controlled task came at the cost of external validity, and the statistical learning task we presented participants with was necessarily abstract. More realistic extensions of our paradigm may present participants with actual news sources like newspapers and with more realistic stimuli like news headlines. This would naturally be accompanied by a drop in control and increase the difficulty of quantifying the individual likelihoods associated with each item.

Second, our pure Bayesian ability to learn the trustworthiness of information sources might only be a marginal factor driver behind our real-world attitudes and media choices. Instead, forces like motivated reasoning (15; 46) might explain more variance. Additionally, we often do not just learn about news sources through direct interaction like we described here but also hear about the reliability of sources from other sources.

This can create triple (or indeed quadruple) learning problems. However, it is generally important to note that we did not set out this study to focus on individual differences, and larger samples are usually necessary to pin down individual differences in this field (62; 63; 64; 65).

The role that sources play in helping people identify misinformation is contested. Research into boosting people to be more robust towards misinformation indeed tries to highlight the source (29; 66). Indeed, people tend to generally have a good assay of the trustworthiness of news providers, at least in a US context (25), and displaying explicit trustworthiness rating may lower their propensity to share fake news (67). However, other research (13) shows that people can pay little regard to the sources of misinformation, making them more vulnerable to being misled.

The inferences our participants and models draw about news providers can be understood as a basic form of theory of mind (68; 69; 70; 71): In essence, they reason and learn about how the sources produce evidence from the ground truth by inverting their model via Bayesian inference. More complex forms of theory of mind can also take into account the beliefs and intentions of other agents (72; 73). The way we deal with those intentions is particularly important, both in media and in psychopathology. When we suspect a source might profit from misleading us, we become more skeptical towards it, and its messages might even backfire (31) or be ignored as cheap talk (74). Computationally, this can be modelled as a recursive theory of mind, where a sender tries to hack the inference process of a receiver and where this receiver might try to infer about such a hacking scheme and defend against it (30; 32). Over- or under-interpreting other agents within this cognitive hierarchy can be understood as both credulity and paranoia, important for different aspects of mental disorders (75; 76). It would be interesting to consider the effects of such recursive modeling in our task.

A defining feature of the modern information sphere is abundant choice. Particularly in the political realm, we can rather freely decide between news sources (77; 78). Future iterations might take into account this information-seeking perspective (79; 80; 81; 82), adding key computational challenges but potentially letting us observe interesting behavioral patterns (see (83) for an empirical investigation of this using a perceptual task). Specifically, agents that have free reign over their sources need to decide when they still need to explore different sources and when they can begin exploiting a source's knowledge (84). Such exploration-exploitation dynamics also open the door for a number of path dependencies that might make participants disengage with potentially mistakenly erroneous news sources quickly, akin to the hot stove effect (85; 86; 87).

In summary, our results have optimistic and pessimistic implications: On the optimistic side, our subjects were generally very successful at their dual Bayesian learning task, showing that the basic components of accurate source assessment and integration have not been dulled. However, that our subjects struggled to learn about the bias of a skewed source even in our rarefied abstract conditions leads to pessimism that we will be able to avoid blandishments exploiting motivated reasoning, and worse. Work on prevention and cure is pressing.

Acknowledgements

We gratefully acknowledge financial support from the Max Planck Society (L.S., E.S., P.D), the Fulbright Program (L.S.), the Humboldt Foundation (P.D.) and the Office of Naval Research (Grant N00014-21-1-2170) and Sloan School of Management (R.B.).

Supplementary Information

Participants

In experiment 1, we analyzed a total of 123 participants. (58 female, 63 male, 2 “other/would rather not say”). The mean age of participants was 38.72 years (SD = 11.07). In experiment 2, we analyzed a total of 111 participants. (54 female, 55 male, 2 “other/would rather not say”). The mean age of participants was 36.77 years (SD = 13.27). We show an overview of these demographics and participants education levels in Fig. 11.

We excluded participants based on two criteria in both experiments: Participants were excluded if they had incorrectly responded to more than two comprehension checks in our instructions, or if they incorrectly responded to a catch-question in the post-task questionnaires.

Participants were remunerated with a base payment and a performance based bonus computed via the quadratic scoring rule based on their confidence judgements during the main block (88) .

Additional results

Noisy feedback condition: responses to independent council

In the noisy feedback condition, participants should use the independent council as a trustworthy, albeit noisy, feedback mechanism. This means that they should respond positively to its information, increasing their confidence in blue when they see more blue endorsements. This shift should be visible regardless of whether participants played with the helpful, the random, opposite or blue-biased source. This was indeed the case, as we show in Fig. 12. There, we plot the average changes of mind, that is the difference between the final confidence after seeing the independent council and the initial confidence after only seeing the news station. Participants shifted their beliefs increasingly towards the blue option the more blue endorsements they saw in the

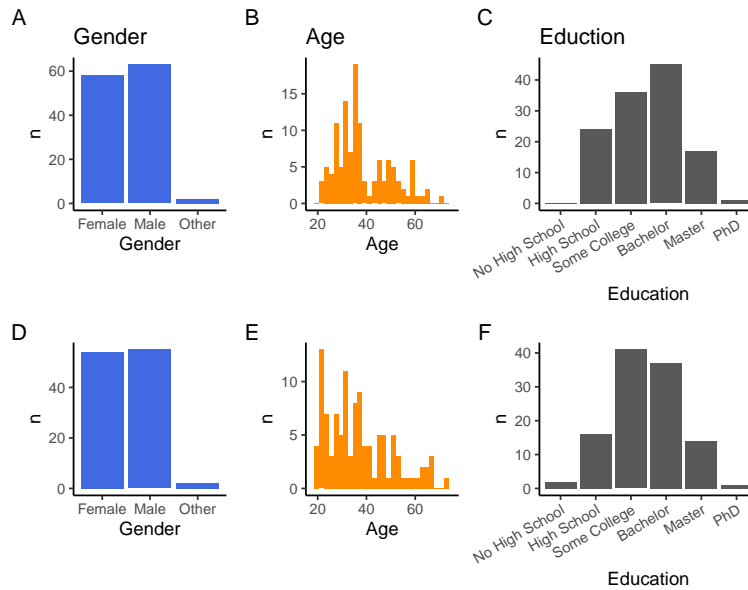


Fig 11. Demographics. (A-C) Experiment 1. (D-F) Experiment 2.

Changes of mind in response to independent council
(Per participant averages)

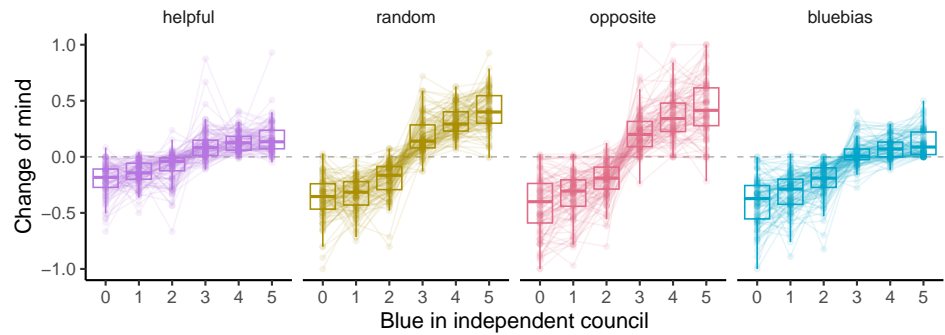


Fig 12. Participants used the independent council as a sensible feedback signal and changed their mind to it accordingly in the noisy feedback version. We plot the changes of mind in response to the independent council as a function of the number of blue endorsement in the independent council. We computed the changes of mind by subtracting the initial confidence (after seeing only the news station) from the final confidence (after additionally seeing the news station). We plot individual averages (lines and dots) and distributions over these (boxplots).

independent council. We note that there are differences in the strength of this shift depending on the news source that participants played with – but that this is due to the joint distributions between the respective news source and independent council, and the accompanying ceiling effects: For example, in the case of the helpful source, participants would have likely already seen a large number of blue endorsements, leaving little room to shift more towards blue when the independent council also favored blue.

Questionnaires and factor analysis

After finishing the main task, participants answered a number of questionnaires related to real-world trust and media consumption patterns:

- The trust question from the World Value Survey. Specifically, participants answered “Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?” on a scale from 0 to 100.
- The epistemic trust questionnaire from Campbell et al. (18). This questionnaire contains three subscales: *Epistemic trust*: A penchant to be open to social learning in “benign social situations”, and to do so adaptively. *Epistemic mistrust*: Tendency to view any information provided by others as unreliable and malign and the resulting hesitancy to rely on such information. *Epistemic credulity*: Absence of “vigilance” towards others, and the resulting tendency to fall prey to deception and misinformation.
- Questions assessing news consumption frequency. Participants were asked “How often did you use the following types of sources to get news in the past week?” with regards to ten different news types, and could respond using four different options (“Never”, “1 - 2 days”, “3 - 4 days”, “5 - 7 days”).¹

¹The news types were: “Cable television news (e.g. CNN, Fox News, MSNBC)”, “National network TV news (e.g. ABC, CBS, NBC)”, “Local television news”, “Social media (e.g. Facebook, Twitter, TikTok, Reddit)”, “Podcasts”, “Blogs”, “Public radio (e.g. NPR)”, “Talk radio”, “News websites or apps”, “Print newspapers”

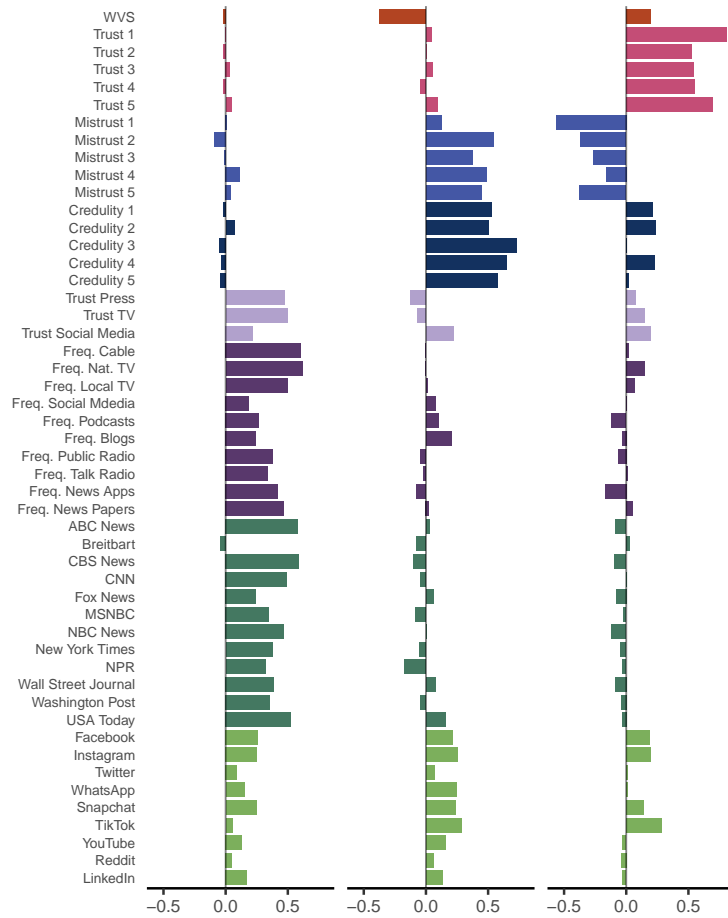


Fig 13. Factor analysis of questionnaire data. Factor loadings for the three factors.

- A questionnaire about the specific news sources participants consumed via a simple yes and no response (“Which of the following media outlets (if any) do you follow to get news?”). ²
- A questionnaire assessing participants’ trust in different types of media (“To what extent do you trust the information that comes from the following?”: “The press (e.g. news apps, newspapers)”, “Television”, “Social media”). Participants could answer using on a five-point scale (“None at all”, “A little”, “A moderate amount”, “A lot”, “A great deal”).

We conducted an exploratory factor analysis on the total set of 50 questions. The number of factors was determined using the Cattell-Nelson-Gorsuch test. We used maximum-likelihood estimation with a an oblique rotation (oblimin). We discuss the results of the factor analysis in the main text and the factor loadings are plotted in Fig 13.

²We assessed the following news organizations: ABC News, Breitbart, CBS News, CNN, Fox News, MSNBC, NBC News, New York Times, NPR, Wall Street Journal, Washington Post, USA Today, Other; and the following social media sources: Facebook, Instagram, Twitter, WhatsApp, Snapchat, TikTok, YouTube, Reddit, LinkedIn, Other

Modelling

Optimal model: Full feedback learning

In the full feedback version of the task, we can model an optimal agent's posterior over the sources' probabilities as two independent beta distributions that both are governed by an α and β parameter. We here drop the subscript I indicating the source but note that this learning process proceeds individually and independently for all sources:

$$b_t | \mathcal{H}_t \sim \text{Beta}(\alpha_{b,t}, \beta_{b,t}) \quad (6)$$

$$g_t | \mathcal{H}_t \sim \text{Beta}(\alpha_{g,t}, \beta_{g,t}) \quad (7)$$

To update its estimate of the two probabilities b_t and g_t , the agent computes the following posterior (here only for b_t but g_t proceeds analogously):

$$p(b_{t+1} | X_t, s_t, \mathcal{H}_t) \propto P(X_t | s; b_t) p(b_t | \mathcal{H}_t) \mathbb{1}(s_t = \text{blue}) \quad (8)$$

Thereby, $P(X_t | s; b_t)$ is again a binomial likelihood, and $p(b_t | \mathcal{H}_t)$ the prior distribution over b_t . The indicator $\mathbb{1}(s_t = \text{blue}) \in \{0, 1\}$ means that we only update our beliefs about b_t when we are in a blue state. Because we are updating a beta-distribution with a binomial likelihood, this update takes the form of a regular beta-update, so that when $s_t = \text{blue}$ (and vice-versa for green):

$$\alpha_{b,t+1} = \alpha_{b,t} + X_t \quad (9)$$

$$\beta_{b,t+1} = \beta_{b,t} + n - X_t \quad (10)$$

The four parameters $\alpha_{b,t}$, $\beta_{b,t}$ and $\alpha_{g,t}$, $\beta_{g,t}$ are thus sufficient statistics for the qualities of the news sources and are the only quantities that need to be carried forward between trials.

Optimal model: Noisy feedback learning

We here outline the computations underlying the agents' update over the sources parameters b_I and g_I in more detail. As we discussed, this is done via an update on the joint distribution of these two parameters. We again drop the source subscript I for convenience:

$$p(b_{t+1}, g_{t+1} | Y_t, X_t, \mathcal{H}_t) = \sum_{s_t} p(b_{t+1}, g_{t+1}, s_t | Y_t, X_t, \mathcal{H}_t) \quad (11)$$

$$\propto \sum_{s_t} p(Y_t, X_t | s_t; b_t, g_t) p(b_t, g_t, s_t | \mathcal{H}_t) \quad (12)$$

$$\propto \sum_{s_t} P(X_t | s_t; b_t, g_t) p(b_t, g_t | \mathcal{H}_t) P(Y_t | s_t) \quad (13)$$

$$\propto [B(X_t; n, b_t) B(Y_t; n, b_Y) + \quad (14)$$

$$B(X_t; n, 1 - g_t) B(Y_t; n, 1 - b_Y)] p(b_t, g_t | \mathcal{H}_t) \quad (15)$$

Essentially, the agent updates its prior belief $p(b_t, g_t | \mathcal{H}_t)$ with the likelihoods of X_t , given this prior belief, but weighted by the likelihood obtained from the noisy feedback sample Y_t , all while marginalizing out the state.

Optimal model: Inferring the ground-truth

After seeing the news station, in both the full feedback and the noisy feedback case, the model updates its beliefs equivalently, marginalizing over its current beliefs about the sources, $p(b_{t,I}, g_{t,I} | \mathcal{H}_t)$. Specifically, it computes the posterior as follows, here for $s_t = \text{blue}$ assuming a flat prior over states $p(s_t) = 0.5$ and dropping the subscript I for convenience.

$$P(s_t = \text{blue} | X_t, \mathcal{H}_t) = \frac{P(X_t | s_t = \text{blue}, \mathcal{H}_t)}{P(X_t | s_t = \text{blue}, \mathcal{H}_t) + P(X_t | s_t = \text{green}, \mathcal{H}_t)} \quad (16)$$

$$= \int_{b_t} \int_{g_t} p(b_t, g_t) \frac{P(X_t | s_t = \text{blue}; b_t, g_t)}{P(X_t | s_t = \text{blue}; b_t, g_t) + P(X_t | s_t = \text{green}; b_t, g_t)} db_t dg_t \quad (17)$$

$$= \frac{\int_{b_t} P(X_t | s_t = \text{blue}; b_t) p(b_t) db_t}{\int_{b_t} P(X_t | s_t = \text{blue}; b_t) p(b_t) db_t + \int_{g_t} P(X_t | s_t = \text{green}; g_t) p(g_t) dg_t} \quad (18)$$

Thereby, $P(X_t | s_t = \text{blue}; b_t)$ is simply the binomial likelihood $B(X_t; n, b_t)$ and in the noisy feedback version $p(b_t)$ is just the marginal distribution $p(b_t) = \int_{g_t} p(b_t, g_t)$ (and vice-versa for $s_t = \text{green}$ and g_t).

In the noisy feedback version, this initial posterior is then integrated with the information from the independent council Y_t to form a final posterior (with no closed form update):

$$P(s_t | X_t, Y_t; b_Y, \mathcal{H}_t) \propto P(Y_t | s_t; b_Y) P(s_t | X_t, \mathcal{H}_t) \quad (19)$$

Model fitting: Parameters

To capture participant idiosyncrasies in learning and confidence reporting, we implemented an eight parameter version of the optimal model, and fit these parameters to participants' behavior using a ninth noise parameter. These parameters are partially shared between the two experiments and models, and partially individual.

We begin by outlining the **shared parameters**. A first shared parameter by the models relates to the way participants report their confidence in a potentially non-linear way. Specifically, we let the confidence be scaled non-linearly using the parameter κ . Here, we denote the reported confidence as \hat{c}_t and the original confidence $c_t \in \{P(s_t = 0 | X_t; \mathcal{H}_t), P(s_t = 0 | X_t, Y_t; b_Y, b_t, g_t)\}$:

$$\hat{c}_t = c_t^\kappa / (c_t^\kappa + (1 - c_t)^\kappa) \quad (20)$$

Crucially, this transformation only impacts the confidence report but not the probabilistic estimates that get made and carried forward by the model.

Another set of four parameters relate to the priors about the source probabilities, $p(b_{t=0}, g_{t=0})$, an agent might have when approaching this task. We implemented this by instantiating the initial beliefs of the agent as beta distributions with the parameters $\alpha_{b,0}, \beta_{b,0}, \alpha_{g,0}, \beta_{g,0}$, so that agents had the following prior

$$b_{t=0} \sim \text{Beta}(\alpha_{b,0}, \beta_{b,0}) \quad (21)$$

$$g_{t=0} \sim \text{Beta}(\alpha_{g,0}, \beta_{g,0}) \quad (22)$$

In the **full feedback experiment**, we consider three more parameters: Specifically, agents might have a tendency towards updating their confidence in a more or less optimistic manner. We capture this by weighting the beta-binomial update by two

parameters ω_α and ω_β , so that each updates proceeds as following, here for when $s_t = \textit{blue}$:

$$\alpha_{b,t+1} = \alpha_{b,t} + \omega_\alpha X_I \quad (23)$$

$$\beta_{b,t+1} = \beta_{b,t} + \omega_\beta (n - X_I) \quad (24)$$

Furthermore, agents might have a tendency to forget (towards the uniform $Beta(1, 1)$), which we implement via a forgetting rate parameter λ . That is, between each trial noise gets added to the parameter estimates and there is a regression to the mean (for both *green* and *blue* parameters, hence dropping the colour subscript here):

$$\alpha'_t = (1 - \lambda)\alpha_t + \lambda \quad (25)$$

$$\beta'_t = (1 - \lambda)\beta_t + \lambda \quad (26)$$

In the **noisy feedback model**, we consider a similar set of phenomena. However, this is implemented somewhat differently because of the joint belief distribution. Specifically, we also implement a forgetting parameter but instead of merely adding noise and forgetting towards a uniform distribution, we forget towards a specific distribution, thereby implementing both a biased belief update and forgetting. We call this forgetting parameter ν . Specifically, ν weights the current two-dimensional belief state $p(b_t, g_t)$ with a second two-dimensional distribution $p_{\text{forget}}(b, g)$ between trials:

$$p(b_t, g_t)'_t = (1 - \nu) * p(b_t, g_t) + \nu * p_{\text{forget}}(b, g) \quad (27)$$

We parameterize $p_{\text{forget}}(b, g)$ to be a 2-dimensional distribution made out of the outer product of two beta distributions both sharing the two parameters α_{forget} and β_{forget} .

$$p_{\text{forget}}(b, g) = Beta(\alpha_{\text{forget}}, \beta_{\text{forget}})Beta(\alpha_{\text{forget}}, \beta_{\text{forget}}) \quad (28)$$

Model fitting: Parameter estimation

We fit the model to participants' data using maximum likelihood estimation via differential evolution implemented in the DEOptim package in R. Participants were fit individually with one set of parameters for all sources. To obtain a likelihood for our data given a parameterized model, we fit the confidence ratings in the data $c_{t,D}$ as being sampled from a Gaussian distribution centered around the model predictions given parameters θ , $c_{t,M|\theta}$ and a standard deviation σ_c :

$$c_{t,D} \sim \mathcal{N}(c_{t,M|\theta}, \sigma_c) \quad (29)$$

The noise term σ_c was a free parameter in our fitting, resulting in a total of 9 parameters.

In both experiments, we fit the initial confidence reported after participants had seen the news station. In the noisy feedback experiment, we additionally fit the second confidences participant reported after having seen the independent council. In both versions, we also fit the 6 probe trials. The probe trials were fit using the frozen belief state after the final main block trial, without any additional forgetting.

Analysis details

Regression analysis

To obtain slopes and intercepts from our data and model, we fit logistic regressions. To deal with the limited data per regression and participants whose responses showed perfect separation, we used Bayesian generalized linear models implemented in the “arm” package in R using the default Cauchy priors. We fit the following regression analyses:

- For the probe trial analysis, we fit four individual regressions per participant, each time fitting their individual six responses per source
- For the pooled analysis split by block quarter, we pooled all participants’ data, and then split it into bins of 7 trials, fitting one regression per source and block quarter
- For the analysis investigating participants’ individual learning progress, we split their data into the first and last 14 trials, and fit one regression per half and source.

We ran these equivalent regressions on the predictions made by the model fit to participant data. To do so, we simulated behavior from the model using the fit parameters. We used the mean confidence prediction of the fit model with no additional noise. In all regressions, we subtracted 2.5 from the number of of blue endorsements, so that an unbiased response would have a intercept of 0.

Statistical analysis of main task

To further analyze participant behavior, we conducted the following statistical analyses, both for the full and noisy feedback version:

- To investigate whether the slopes and intercepts differed from each other in the probe trials between news sources, we conducted an ANOVA predicting the slope (or intercept, respectively) from the source type. We conducted Tukey’s honest significance test to compare the individual means.
- To investigated whether participant’s slopes and intercept differed from 0, we conducted one-sample t-tests.
- To investigate quality of fit between probe trial slopes and intercepts computed for the data and the model predictions, we computed Pearson’s R.
- To investigate effects of trial and news source on participants’ accuracy, we set up a logistic regression predicting, on a trial-by-trial basis, participants’ accuracy from the trial number (within a block) and the source, as well as their interaction. To test for the significant effects of these predictors, we conducted χ^2 tests.
- To investigate whether slopes and intercepts differed between block halves and news sources, we conducted an ANOVA predicting the slope/intercept from the source type and the block half. We conducted Tukey’s honest significance test to compare the individual means.

We ran the following analyses to compare the two conditions and the optimal model:

- To compare the accuracy achieved between the optimal model and the data, we set up a logistic regression predicting, on a trial-by-trial basis, participants’ accuracy from the trial number (within a block), as well as from the data type (optimal model or data) and the interaction of these two predictors. To test for the significant effect of these predictors, we conducted a χ^2 tests.

Table 1. Comparison of model fit for the post-block trust and improvement ratings: BICs for regression models predicting the post-block trust questionnaires from the betas fit to participants’ responses in the second half of a block.

DV	News Sources	Model	BIC
Trust	All but opposite	Linear	1585.6531
Trust	All but opposite	Quadratic	1589.4885
Trust	All but opposite	Linear and quadratic	1590.7492
Trust	Opposite	Linear	690.0235
Trust	Opposite	Quadratic	696.8924
Trust	Opposite	Linear and quadratic	694.5620
Improve	All but opposite	Linear	1602.8504
Improve	All but opposite	Quadratic	1610.0373
Improve	All but opposite	Linear and quadratic	1608.9850
Improve	Opposite	Linear	720.3958
Improve	Opposite	Quadratic	718.0166
Improve	Opposite	Linear and quadratic	717.8379

- To compare the slopes achieved by the models, we set up an ANOVA predicting the probe trial slopes from the experiment condition and the news source. We conducted Tukey’s honest significance test to compare the individual means.

Trust and improvement ratings

To investigate the different types of responses to the post-block trust and improvement questionnaires, we set up an ANOVA predicting the response from the source type and the response type (trust/improve). We conducted Tukey’s honest significance test to compare the individual means.

To check how participants’ trust and improvement ratings were related to their task behavior, we set up regressions predicting the respective rating from the slopes inferred as above from the participants’ second half responses. We set up three regressions per relationship, (1) one predicting the rating linearly from the slopes, (2) another predicting the ratings from a quadratic function of the slopes, and (3) a final using both a linear and quadratic function. We compared the quality of these fits using the regressions’ respective Bayesian Information Criteria, reported in Table 1.

References

1. FeldmanHall O, Nassar MR. The computational challenge of social learning. *Trends in Cognitive Sciences*. 2021;25(12):1045–1057. doi:10.1016/j.tics.2021.09.002.
2. FeldmanHall O, Shenhav A. Resolving uncertainty in a social world. *Nature Human Behaviour*. 2019;3(5):426–435. doi:10.1038/s41562-019-0590-x.
3. Olsson A, Knapska E, Lindström B. The neural and computational systems of social learning. *Nature Reviews Neuroscience*. 2020; p. 1–16. doi:10.1038/s41583-020-0276-4.
4. Heyes C. Who Knows? Metacognitive Social Learning Strategies. *Trends in Cognitive Sciences*. 2016;20(3):204–213. doi:10.1016/j.tics.2015.12.007.

5. Hahn U, Merdes C, von Sydow M. How Good Is Your Evidence and How Would You Know? *Topics in Cognitive Science*. 2018;10(4):660–678. doi:10.1111/tops.12374.
6. Merdes C, von Sydow M, Hahn U. Formal models of source reliability. *Synthese*. 2021;198(23):5773–5801. doi:10.1007/s11229-020-02595-2.
7. Pescetelli N, Yeung N. The Role of Decision Confidence in Advice-Taking and Trust Formation. *Journal of Experimental Psychology: General*. 2020;doi:10.1037/xge0000960.
8. Vélez N, Gweon H. Integrating Incomplete Information With Imperfect Advice. *Topics in Cognitive Science*. 2019;11(2):299–315. doi:10.1111/tops.12388.
9. Gentzkow M, Shapiro JM, Stone DF. Media Bias in the Marketplace. In: *Handbook of media economics*. North-Holland; 2015. p. 623–645.
10. Liang A, Mu X. Complementary Information and Learning Traps. *The Quarterly Journal of Economics*. 2020;135(1):389–448. doi:10.1093/qje/qjz033.
11. Jøsang A, Ismail R. The Beta Reputation System. In: *Proceedings of the 15th bled electronic commerce conference*. vol. 5; 2002. p. 2502–2511. Available from: <http://databases.si.umich.edu/reputations/>.
12. De Siqueira Braga D, Niemann M, Hellingrath B, De Lima Neto FB. Survey on computational trust and reputation models. *ACM Computing Surveys*. 2018;51(5). doi:10.1145/3236008.
13. Dias N, Pennycook G, Rand DG. Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media. *Harvard Kennedy School Misinformation Review*. 2020;doi:10.37016/mr-2020-001.
14. Pennycook G, Rand DG. The Psychology of Fake News. *Trends in Cognitive Sciences*. 2021;25(5):388–402. doi:10.1016/j.tics.2021.02.007.
15. Druckman JN, McGrath MC. The evidence for motivated reasoning in climate change preference formation. *Nature Climate Change*. 2019;9(2):111–119. doi:10.1038/s41558-018-0360-1.
16. Orchinik R, Dubey R, Gershman SJ, Powell D, Bhui R. Learning from and about climate scientists; 2023. Available from: <https://osf.io/preprints/psyarxiv/ezua5/>.
17. Cook J, Lewandowsky S. Rational Irrationality: Modeling Climate Change Belief Polarization Using Bayesian Networks. *Topics in Cognitive Science*. 2016;8(1):160–179. doi:10.1111/tops.12186.
18. Campbell C, Tanzer M, Saunders R, Booker T, Allison E, Li E, et al. Development and validation of a self-report measure of epistemic trust. *PLoS ONE*. 2021;16(4 April):1–21. doi:10.1371/journal.pone.0250264.
19. Fonagy P, Campbell C, Constantinou M, Higgitt A, Allison E, Luyten P. Culture and psychopathology: An attempt at reconsidering the role of social learning. *Development and Psychopathology*. 2021;(May):1–16. doi:10.1017/S0954579421000092.
20. Teunisse AK, Case TI, Fitness J, Sweller N. I Should Have Known Better: Development of a Self-Report Measure of Gullibility. *Personality and Social Psychology Bulletin*. 2020;46(3):408–423. doi:10.1177/0146167219858641.

21. Shao J, Du W, Lin T, Li X, Li J, Lei H. Credulity rather than general trust may increase vulnerability to fraud in older adults: a moderated mediation model. *Journal of Elder Abuse & Neglect*. 2019;31(2):146–162. doi:10.1080/08946566.2018.1564105.
22. Fischer P, Lea SEG, Evans KM. Why do individuals respond to fraudulent scam communications and lose money? The psychological determinants of scam compliance. *Journal of Applied Social Psychology*. 2013;43(10):2060–2072. doi:10.1111/jasp.12158.
23. Barnby JM, Dean RJ, Burgess H, Kim J, Teunisse AK, Mackenzie L, et al. Increased persuadability and credulity in people with corpus callosum dysgenesis. *Cortex*. 2022;155:251–263. doi:10.1016/j.cortex.2022.07.009.
24. Pornpitakpan C. The Persuasiveness of Source Credibility: A Critical Review of Five Decades' Evidence. *Journal of Applied Social Psychology*. 2004;34(2):243–281. doi:10.1111/j.1559-1816.2004.tb02547.x.
25. Pennycook G, Rand DG. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*. 2019;116(7):2521–2526. doi:10.1073/pnas.1806781116.
26. Lin H, Lasser J, Lewandowsky S, Cole R, Gully A, Rand DG, et al. High level of correspondence across different news domain quality rating sets. *PNAS Nexus*. 2023;2(9):pgad286. doi:10.1093/pnasnexus/pgad286.
27. Olsen K, Roepstorff A, Bang D. Knowing whom to learn from: individual differences in metacognition and weighting of social information. *PsyArXiv*. 2019; p. 20.
28. Frankfurt HG. *On bullshit*. Princeton, NJ: Princeton University Press; 2005.
29. Hertwig R, Engel C. *Deliberate Ignorance: Choosing Not to Know*. MIT Press; 2021.
30. Alon N, Schulz L, Rosenschein JS, Dayan P. A (Dis-)information Theory of Revealed and Unrevealed Preferences: Emerging Deception and Skepticism via Theory of Mind. *Open Mind*. 2023;7:608–624. doi:10.1162/opmi_a00097.
31. Bhui R, Gershman SJ. Paradoxical effects of persuasive messages. *Decision*. 2020;7(4):239–258. doi:10.1037/dec0000123.
32. Oey LA, Schachner A, Vul E. Designing and Detecting Lies by Reasoning About Other Agents. *Journal of Experimental Psychology: General*. 2022;doi:10.1037/xge0001277.
33. Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS. Learning the value of information in an uncertain world. *Nature Neuroscience*. 2007;10(9):1214–1221. doi:10.1038/nn1954.
34. Charness G, Oprea R, Yuksel S. How do People Choose Between Biased Information Sources? Evidence from a Laboratory Experiment. *Journal of the European Economic Association*. 2021;19(3):1656–1691. doi:10.1093/jeea/jvaa051.
35. Gentzkow M, Shapiro J. Media Bias and Reputation. *Journal of Political Economy*. 2006;114(2):280–316. doi:10.1086/499414.

36. Jo D. Better the Devil You Know: An Online Field Experiment on News Consumption. 2020;.
37. Hu L, Kovach M, Li A. Learning News Bias: Misspecifications and Consequences; 2023. Available from: <http://arxiv.org/abs/2309.08740>.
38. Groseclose T, Milyo J. A Measure of Media Bias*. *The Quarterly Journal of Economics*. 2005;120(4):1191–1237. doi:10.1162/003355305775097542.
39. van der Plas E, David AS, Fleming SM. Advice-taking as a bridge between decision neuroscience and mental capacity. *International Journal of Law and Psychiatry*. 2019;67:101504. doi:10.1016/j.ijlp.2019.101504.
40. Diaconescu AO, Mathys C, Weber LAE, Daunizeau J, Kasper L, Lomakina EI, et al. Inferring on the Intentions of Others by Hierarchical Bayesian Learning. *PLoS Computational Biology*. 2014;10(9). doi:10.1371/journal.pcbi.1003810.
41. Hertz U, Palminteri S, Brunetti S, Olesen C, Frith CD, Bahrami B. Neural computations underpinning the strategic management of influence in advice giving. *Nature Communications*. 2017;8(1):2191. doi:10.1038/s41467-017-02314-5.
42. Hertz U, Bell V, Raihani N. Trusting and learning from others: immediate and long-term effects of learning from observation and advice. *Proceedings of the Royal Society B: Biological Sciences*. 2021;288(1961):20211414. doi:10.1098/rspb.2021.1414.
43. Fiedler K, Hütter M, Schott M, Kutzner F. Metacognitive myopia and the overutilization of misleading advice. *Journal of Behavioral Decision Making*. 2019;32(3):317–333. doi:10.1002/bdm.2109.
44. Siegel JZ, Mathys C, Rutledge RB, Crockett MJ. Beliefs about bad people are volatile. *Nature Human Behaviour*. 2018;2(10):750–756. doi:10.1038/s41562-018-0425-1.
45. Kaufmann RK, Mann ML, Gopal S, Liederman JA, Howe PD, Pretis F, et al. Spatial heterogeneity of climate change as an experiential basis for skepticism. *Proceedings of the National Academy of Sciences*. 2017;114(1):67–71. doi:10.1073/pnas.1607032113.
46. Bayes R, Druckman JN. Motivated reasoning and climate change. *Current Opinion in Behavioral Sciences*. 2021;42:27–35. doi:10.1016/j.cobeha.2021.02.009.
47. Tappin BM, Van Der Leer L, McKay RT. The heart trumps the head: Desirability bias in political belief revision. *Journal of Experimental Psychology: General*. 2017;146(8):1143–1149. doi:10.1037/xge0000298.
48. Pescetelli N, Hauperich AK, Yeung N. Confidence, advice seeking and changes of mind in decision making. *Cognition*. 2021;215(October 2020). doi:10.1016/j.cognition.2021.104810.
49. Allen J, Howland B, Mobius M, Rothschild D, Watts DJ. Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*. 2020;6(14):eaay3539. doi:10.1126/sciadv.aay3539.
50. Marchal N, Kollanyi B, Neudert LM, Howard PN. Junk News During the EU Parliamentary Elections: Lessons from a Seven-Language Study of Twitter and Facebook. University of Oxford; 2019.

51. Altay S, Lyons B, Modirrousta-Galian A. Exposure to Higher Rates of False News Erodes Media Trust and Fuels Skepticism in News Judgment; 2023. Available from: <https://psyarxiv.com/t9r43/>.
52. Claessens S, Fischer K, Chaudhuri A, Sibley CG, Atkinson QD. The dual evolutionary foundations of political ideology. *Nature Human Behaviour*. 2020;4(4):336–345. doi:10.1038/s41562-020-0850-9.
53. Sharot T, Garrett N. Forming Beliefs: Why Valence Matters. *Trends in Cognitive Sciences*. 2016;20(1):25–33. doi:10.1016/j.tics.2015.11.002.
54. Dorfman HM, Bhui R, Hughes BL, Gershman SJ. Causal Inference About Good and Bad Outcomes. *Psychological Science*. 2019;30(4):516–525. doi:10.1177/0956797619828724.
55. Fischer I, Harvey N. Combining forecasts: What information do judges need to outperform the simple average? *International Journal of Forecasting*. 1999;15(3):227–246. doi:10.1016/S0169-2070(98)00073-9.
56. Slovic P. Perceived Risk, Trust, and Democracy. *Risk Analysis*. 1993;13(6):675–682. doi:10.1111/j.1539-6924.1993.tb01329.x.
57. Rothbart M, Park B. On the confirmability and disconfirmability of trait concepts. *Journal of personality and social psychology*. 1986;50(1):131.
58. Altay S, Hacquin AS, Mercier H. Why do so few people share fake news? It hurts their reputation. *New Media & Society*. 2022;24(6):1303–1324. doi:10.1177/1461444820969893.
59. Bhui R, Lai L, Gershman SJ. Resource-rational decision making. *Current Opinion in Behavioral Sciences*. 2021;41:15–21. doi:10.1016/j.cobeha.2021.02.015.
60. Wilson RC, Collins AGE. Ten simple rules for the computational modeling of behavioral data. *eLife*. 2019;8:1–35. doi:10.7554/eLife.49547.
61. Montague RP, Lohrenz T, Dayan P. The three R's of trust. *Current Opinion in Behavioral Sciences*. 2015;3:102–106. doi:10.1016/j.cobeha.2015.02.009.
62. Schulz L, Rollwage M, Dolan RJ, Fleming SM. Dogmatism Manifests in Lowered Information Search under Uncertainty. *Proceedings of the National Academy of Sciences*. 2020;117(49):31527–31534.
63. Rollwage M, Dolan RJ, Fleming SM, Rollwage M, Dolan RJ, Fleming SM. Metacognitive Failure as a Feature of Those Holding Radical Beliefs. *Current Biology*. 2018;28(24):4014–4021.e8. doi:10.1016/j.cub.2018.10.053.
64. Zmigrod L, Eisenberg I, Bissett P, Robbins TW, Poldrack R. A Data-Driven Analysis of the Cognitive and Perceptual Attributes of Ideological Attitudes. 2020;(April 2020). doi:10.31234/osf.io/dgaxr.
65. Zmigrod L, Rentfrow PJ, Robbins TW. Cognitive underpinnings of nationalistic ideology in the context of Brexit. *Proceedings of the National Academy of Sciences*. 2018;115(19):E4532 – E4540. doi:10.1073/pnas.1708960115.
66. Breakstone J, Smith M, Connors P, Ortega T, Kerr D, Wineburg S. Lateral reading: College students learn to critically evaluate internet sources in an online course. *Harvard Kennedy School Misinformation Review*. 2021;doi:10.37016/mr-2020-56.

67. Celadin T, Capraro V, Pennycook G, Rand DG. Displaying News Source Trustworthiness Ratings Reduces Sharing Intentions for False News Posts. *Journal of Online Trust and Safety*. 2023;1(5). doi:10.54501/jots.v1i5.100.
68. Jara-Ettinger J. Theory of mind as inverse reinforcement learning. *Current Opinion in Behavioral Sciences*. 2019;29:105–110. doi:10.1016/j.cobeha.2019.04.010.
69. Premack D, Woodruff G. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*. 1978;1(4):515–526. doi:10.1017/S0140525X00076512.
70. Barnby JM, Bellucci G, Alon N, Schilbach L, Bell V, Frith C, et al.. Beyond Theory of Mind: A formal framework for social inference and representation; 2023. Available from: <https://osf.io/preprints/psyarxiv/cmgu7/>.
71. Rusch T, Steixner-Kumar S, Doshi P, Spezio M, Gläscher J. Theory of mind and decision science: Towards a typology of tasks and computational models. *Neuropsychologia*. 2020;146:107488. doi:10.1016/j.neuropsychologia.2020.107488.
72. Jara-Ettinger J, Gweon H, Schulz LE, Tenenbaum JB. The Naïve Utility Calculus: Computational Principles Underlying Commonsense Psychology. *Trends in Cognitive Sciences*. 2016;20(8):589–604. doi:10.1016/j.tics.2016.05.011.
73. Gmytrasiewicz PJ, Doshi PJ. A framework for optimal sequential planning in multiagent settings. *Journal of Artificial Intelligence Research*. 2005;24:49 – 79.
74. Farrell J, Matthew R. Cheap Talk. *Journal of Economic Perspectives*. 1996;10(3):103–118.
75. Alon N, Schulz L, Dayan P, Barnby JM. Between prudence and paranoia: Theory of Mind gone right, and wrong. *ICML 2023 Workshop on Theory of Mind in Communicating Agents*. 2023;.
76. Barnby JM, Dayan P, Bell V. Formalising social representation to explain psychiatric symptoms. *Trends in Cognitive Sciences*. 2023;27(3):317–332. doi:10.1016/j.tics.2022.12.004.
77. Hills TT. The Dark Side of Information Proliferation. *Perspectives on Psychological Science*. 2019;14(3):323–330. doi:10.1177/1745691618803647.
78. Hills TT, Todd PM, Lazer D, Redish AD, Couzin ID, Bateson M, et al. Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*. 2015;19(1):46–54. doi:10.1016/j.tics.2014.10.004.
79. Sharot T, Sunstein CR. How people decide what they want to know. *Nature Human Behaviour*. 2020;4(1):14–19. doi:10.1038/s41562-019-0793-1.
80. Schulz L, Fleming SM, Dayan P. Metacognitive Computations for Information Search: Confidence in Control. *Psychological Review*. 2023;doi:10.1037/rev0000401.
81. Gottlieb J, Oudeyer PY. Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*. 2018;19(12):758–770. doi:10.1038/s41583-018-0078-0.
82. Gottlieb J, Oudeyer PY, Lopes M, Baranes A. Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*. 2013;17(11):585–593. doi:10.1016/j.tics.2013.09.001.

83. Carlebach N, Yeung N. Flexible use of confidence to guide advice requests. *Cognition*. 2023;230:105264. doi:10.1016/j.cognition.2022.105264.
84. Schulz E, Gershman SJ. The algorithmic architecture of exploration in the human brain. *Current Opinion in Neurobiology*. 2019;55:7–14. doi:10.1016/J.CONB.2018.11.003.
85. Rich AS, Gureckis TM. The Limits of Learning: Exploration, Generalization, and the Development of Learning Traps. *Journal of Experimental Psychology: General*. 2018;147(11):1553–1570.
86. Dayan P, Roiser JP, Viding E. The first steps on long marches: The costs of active observation. In: Savulescu J, Roache R, Davies W, Loebel JP, Davies W, Savulescu J, et al., editors. *Psychiatry Reborn: Biopsychosocial psychiatry in modern medicine*. Oxford University Press; 2020. p. 0. Available from: <https://doi.org/10.1093/med/9780198789697.003.0014>.
87. Shook NJ, Fazio RH. Political ideology, exploration of novel stimuli, and attitude formation. *Journal of Experimental Social Psychology*. 2009;45(4):995–998. doi:10.1016/j.jesp.2009.04.003.
88. Staël von Holstein CAS. Measurement of subjective probability. *Acta Psychologica*. 1970;34:146–159. doi:10.1016/0001-6918(70)90013-2.