


ORIGINAL PAPER

Audi alteram partem: An experiment on selective exposure to information

Salvatore Nunnari¹  and Giovanni Montanari²

¹Department of Economics, Bocconi University, Milan, ITALY

²Department of Economics, New York University, New York, NY, USA

Corresponding author: Salvatore Nunnari; Email: salvatore.nunnari@unibocconi.it

We are grateful to Gary Charness and audiences at New York University and the 2020 Global Online Meeting of the Economic Science Association for helpful comments. Nunnari gratefully acknowledges financial support from the European Research Council through ERC Grant 852,526 (POPULIZATION).

(Received 13 December 2024; accepted 11 March 2025)

Abstract

We report the results of an experiment on selective exposure to information. A decision maker interested in learning about an uncertain state of the world can acquire information from one of two sources that have opposite biases: when informed on the state, they report it truthfully; when uninformed, they report their favorite state. A Bayesian decision-maker is better off seeking confirmatory information unless the source biased against the prior is sufficiently more reliable. In line with the theory, subjects are more likely to seek confirmatory information when sources are symmetrically reliable. On the other hand, when sources are asymmetrically reliable, subjects are more likely to consult the more reliable source even when prior beliefs are strongly unbalanced and this source is less informative. Our experiment suggests that base rate neglect and simple heuristics (e.g., listen to the most reliable source) are important drivers of the endogenous acquisition of information.

Keywords: base rate neglect; biased information sources; confirmation bias; echo chambers; Information acquisition; laboratory experiment; selective exposure

JEL Codes: C91; D81; D83; D91

1. Introduction

Social scientists have collected ample evidence that people selectively search for and attend to a subset of the available information, ignoring additional evidence (Frey, 1986; Iyengar & Hahn, 2009; Nickerson, 1998). In the words of Berelson and Steiner (1964), ‘people tend to see and hear communications that are favorable or congenial to their predispositions; they are more likely to see and hear congenial communications than neutral or hostile ones.’

This behavior has raised concern: as the availability of media choices has grown, selective exposure to like-minded sources has contributed to a deep partisan divide in news consumption (Barberá et al., 2015; Gentzkow & Shapiro, 2011; Lawrence et al., 2010; Peterson et al., 2021). In turn, this segregation into ‘echo chambers’ has been associated with the observed intensification of partisan sentiment as well as with the recent surge of populist parties in developed democracies (Bakshy et al., 2015; Flaxman et al., 2016; Mann & Ornstein, 2012). Why do we observe this behavior? Recent theoretical

work in economics suggests that individuals might have systematic preferences for information consonant with their beliefs (Mullainathan & Shleifer, 2005) or that, being uncertain about the reliability of an information source, they interpret disconfirming evidence as less credible than confirming evidence and turn their attention towards the source they deem more informative (Gentzkow & Shapiro, 2006). Notably, even when individuals have no uncertainty about sources reliability and regard all media outlets as equally credible, selective exposure to like-minded sources can be a rational choice for an individual who has limited time or attention and can only access or process a subset of the available evidence.

In this paper, we investigate this last mechanism with a laboratory experiment. In particular, we ask the following research questions: How should an attention-constrained but otherwise rational agent optimally acquire information from multiple potential sources with different biases? What is the ability of this normative model to predict the observed demand for (dis)confirmatory information?

In our model, decision-makers have the possibility to acquire a *signal* from one of two *information sources* in order to reduce their uncertainty about a state of the world. Importantly, decision-makers know the conditional distributions of signals for each information source, ruling out any uncertainty about the *reliability* of information sources. We also provide decision-makers with an exogenous prior belief on the state of the world and focus on an abstract decision environment that allows us to minimize the confounding effects of *motivated beliefs*. Once decision-makers observe the signal from the information source of choice, they guess the state of the world they deem more likely and receive a reward only for a correct guess. We manipulate the probability distributions of signals delivered by each information source (in order to control their *relative reliability*) and the prior belief over the state of the world. As a consequence of our manipulations, it is optimal to follow confirmatory information sources in some treatments but not in others. We verify optimal information acquisition in both environments and test for a confirmatory pattern on top of and above what can be explained by rational behavior.

Some predictions of our theory align with observed behavior, while others are not supported by the data. When the two information sources are equally reliable, information acquisition displays a *confirmatory pattern*, as the source supportive of the prior belief is the most consulted one. This is in line with theoretical predictions. On the other hand, when we manipulate the relative reliability of information sources and make the source less supportive of the prior belief the optimal choice, participants display a *disconfirmatory pattern* of information acquisition, regardless of the strength of the prior. This contrasts with the predictions of the model, suggesting decision-makers pay undue attention to the reliability of information sources and underweigh the importance of the ex-ante uncertainty surrounding the phenomenon to learn about. This suggests that the adoption of simple heuristics – e.g., listen to the more reliable source – is an important driver of the endogenous acquisition of information.

This paper contributes to two strands of literatures. First, our paper contributes to a literature in experimental psychology on how people gather evidence to test hypotheses (Baron et al., 1988; Klayman & Young-Won, 1987; Skov & Sherman, 1986; Slowiaczek et al., 1992).¹ Second, our paper contributes to a recent literature in experimental economics on the choice over information sources with instrumental value (Ambuehl, 2021; Ambuehl & Shengwu, 2018; Castagnetti & Schmacker, 2022; Chopra et al., 2024; Duffy et al., 2019; Sharma & Castagnetti, 2023). The most closely related work is Charness et al (2021). Similarly to some of our treatments, they consider experimental conditions (labeled *bias by commission*) where decision-makers choose between two information sources that are biased towards opposite states and might send an incorrect signal with the same probability (that is, they are symmetrically reliable). Contrary to their setting, we investigate experimental

¹Testing an hypothesis means checking whether a statement of the form ' p implies q ' is true. Logically, one can test the same hypothesis by checking whether a statement of the form ' $\text{not } q$ implies $\text{not } p$ ' is true. This means that, in this context, it is difficult to define what it means for information to be confirmatory or contradictory. Our experiment is not designed to test the ability to construct a logical test but rather the endogenous acquisition of an informative signal.

treatments where the two available information sources are biased towards opposite states, and might send an incorrect signal with different probabilities (that is, they are asymmetrically reliable).²

This distinction is critical because, in many real-life contexts, information sources often differ not only in the direction of their bias but also in the magnitude of their bias. For example, political pollsters with opposing partisan leanings might have symmetric reliability, meaning they are equally likely to make errors when the truth misaligns with their bias, such as overestimating support for their preferred party by the same degree. Similarly, financial advisors with different investment philosophies – one favoring high-risk, high-reward strategies and the other favoring conservative, low-risk options – may have symmetric reliability if they are equally prone to misjudging market conditions that do not align with their preferred strategies. In contrast, media outlets often exhibit asymmetric reliability; for instance, a left-leaning outlet might be highly accurate when reporting on topics aligned with progressive values, like climate change, but less reliable when covering issues that challenge those values, while a right-leaning outlet may show the reverse pattern but at different accuracy levels overall. Medical experts provide another compelling example: a cardiologist might be highly reliable in diagnosing heart-related issues but less so for neurological problems, while a neurologist would display the opposite pattern, with differences in the degree of overall reliability depending on the doctor's experience. By focusing on asymmetrically reliable sources, this paper introduces an additional layer of complexity that mirrors real-world challenges and broadens our understanding of selective exposure in these nuanced and practically relevant contexts.

Charness et al. (2021) conclude that 'sub-optimal decision rules [...] emerge here because *it is difficult to correctly reason through information valuation problems*, even in our deliberately simple setting.' Our complementary experimental design allows us to uncover one simple heuristic that individuals rely on in such a complex decision-making environment: when the available sources have different reliability, choose the most trustworthy. If individuals suffer from base-rate neglect (a well documented error in probabilistic reasoning; Bar-Hillel, 1980; Esponda et al., 2024; Kahneman & Tversky, 1973) and are not very responsive to the strength of their prior belief, this simple rule of thumb can also appear normatively appealing.

2. Task and theoretical predictions

Consider a decision-maker (DM) who is uncertain about a state of the world, $\theta = \{B, R\}$, and has to make a guess, $a = \{B, R\}$. The DM earns a reward (normalized to 1) only if this guess matches the state of the world. We denote with π the DM's prior belief that $\theta = B$. We focus on unbalanced priors and, without loss of generality, we assume $\pi \in (1/2, 1)$, that is, the DM's prior is that the state is more likely to be B . Before making a guess, the DM acquires a piece of information from one of two *information sources*, **Blue** and **Red**. Each information source stochastically maps the state of the world to a *signal* $s = \{b, r\}$, as described in Table 1.

In each panel of Table 1, each cell displays the probability of observing a signal (column) in a specific state of the world (row). We can interpret λ_σ as a measure of *bias* (or as an inverse measure of *reliability*) of source $\sigma = \{B, R\}$: **Blue** is biased towards B and λ_B represents the probability that it signals the state is B when it is, in fact, R ; **Red** is biased towards R and λ_R represents the probability that it signals the state is R when it is, in fact, B . We assume that both sources are somewhat informative but also somewhat biased – that is, $\lambda_B, \lambda_R \in (0, 1)$. In line with Gentzkow and Shapiro (2006), this simple framework can capture different real-world scenarios: the information source

²Charness et al. (2021) also consider treatments where the two information sources are asymmetrically reliable but biased towards the same state and, thus, there is no trade-off between reliability and direction of the bias (in fact, the two experts can easily be ranked by Blackwell ordering); and treatments where the two information sources are biased towards opposite states and might fail to send a signal (labeled *bias by omission*). In their design, it is the nature of the bias (*commission* versus *omission*) to determine whether it is optimal to consult information sources biased towards or against prior beliefs. In contrast, we achieve this goal by keeping the nature of the bias fixed and varying the sources' relative reliability.

Table 1. Conditional distribution of signals by information sources

(a) Blue Information Sources			(b) Red Information Source		
	$S = b$	$S = r$		$S = b$	$S = r$
$\theta = B$	1	0	$\theta = B$	$1 - \lambda_R$	λ_R
$\theta = R$	λ_B	$1 - \lambda_B$	$\theta = R$	0	1

may be uninformed about the state and report a default signal; it may strategically slant its report when the information it holds is against its favorite state; or its intended signal may inadvertently be distorted.

2.1. Optimal guess for given information source

We characterize the DM’s optimal choice of information source by backward induction. First, we investigate the optimal guess for a given signal received by a given source. Second, we investigate what information source the DM prefers to consult, given the distribution of signals induced by each information source and how the DM will use these signals. In what follows, the notation $a^*(s, \sigma)$ denotes the optimal guess after observing signal s from information source σ . All proofs are in Appendix A.

Proposition 1 (Optimal Guess if Signal from Blue Source) *The DM always follows the signal received from source **Blue**, that is, $a^*(b, \text{Blue}) = B$ and $a^*(r, \text{Blue}) = R$.*

Proposition 2 (Optimal Guess if Signal from Red Source) *The DM always follows a confirmatory signal received from source **Red**, that is, $a^*(b, \text{Red}) = B$. The DM follows a contradictory signal received from source **Red** if and only if the source is sufficiently reliable, that is, $a^*(r, \text{Red}) = R$ if $\lambda_R < \frac{1-\pi}{\pi}$ and $a^*(r, \text{Red}) = B$ otherwise.*

Remember that the DM’s prior belief favors B . When she observes a signal confirming her prior from either source, the DM’s posterior belief that $\theta = B$ is strictly greater than her prior. Thus, in this case, the DM sticks with her prior belief and guesses accordingly. Receiving a signal that disagrees with the source bias – that is, receiving signal b (r) from the **Red** (**Blue**) source – is fully revealing: the DM learns the state with certainty, independently of her prior beliefs and the source reliability. Finally, when she observes signal r from **Red**, the DM’s posterior belief that $\theta = B$ is strictly smaller than her prior. In this case, the optimal guess depends on the model parameters: if **Red** is sufficiently reliable (i.e., λ_R is sufficiently small), it is optimal to follow its signal. Otherwise, the DM is better off ignoring the signal altogether and sticking with the guess induced by her prior belief. The relative size of λ_R must be gauged against the prior belief: the larger the prior in favor of B , the higher the reliability of **Red** required by the DM to follow an r signal from this source.

2.2. Optimal choice of information source

First, consider the expected utility from consulting the source biased in favor of the prior, that is, **Blue**. Since the DM follows any signal received from **Blue**, acquiring information from this source always improves the confidence the DM has in her guess with respect to a decision made without collecting any additional information. Second, consider the expected utility from consulting the source biased against the prior, that is, **Red**. When this source is sufficiently biased — that is, when $\lambda_R \geq \frac{1-\pi}{\pi}$ — the DM guesses B regardless of the signal. In this case, acquiring information from this source does not change the confidence the DM has in her guess with respect to a decision made without collecting any additional information. When, instead, this source is sufficiently reliable — that is, when $\lambda_R < \frac{1-\pi}{\pi}$ — the DM follows any signal received from **Red** and, similarly to **Blue**, acquiring information from this source always improves the confidence the DM has in her guess.

Since consulting the source biased in favor of the prior is always informative while consulting the source biased against the prior is informative only if $\lambda_R < \frac{1-\pi}{\pi}$, the DM is better off consulting **Blue** when $\lambda_R \geq \frac{1-\pi}{\pi}$. When $\lambda_R < \frac{1-\pi}{\pi}$, both sources are informative and the choice involves a trade-off. Intuitively, the DM chooses the source with the smallest probability of misleading signals. If the DM has a perfectly balanced prior, choosing **Red** over **Blue** reduces to $\lambda_R < \lambda_B$. When the prior is unbalanced, the DM has an incentive to choose the information source that is biased towards the prior. She prefers to observe a signal from **Red** only when this information source is sufficiently more reliable than the other. Proposition 3 summarizes this discussion and characterizes this threshold:

Proposition 3 (Optimal Information Source) *The DM acquires information from **Red** if $\lambda_R < \frac{1-\pi}{\pi} \lambda_B$ and acquires information from **Blue** otherwise.*

2.3. Summary of testable hypotheses

Below, we summarize the testable hypotheses that we set out to investigate empirically.

- H1 *When information sources are equally reliable, it is optimal to acquire information from the source biased towards the prior.*
- H2 *When information sources have different reliability and the prior is mildly unbalanced, it is optimal to acquire information from the more reliable source. Conversely, when the prior is strongly unbalanced, it is optimal to acquire information from the source biased towards the prior, even if it is less reliable.*

3. Experimental design

The experiment was conducted in 2017 on Prolific with 201 U.S. nationals and residents whose first language was English. Instructions are available in Appendix D.³

Setup. The task builds on the classic urn paradigm, which has been extensively used in the experimental literature since Anderson and Holt (1997). Subjects are asked to guess the color of a ball randomly drawn from an urn containing only blue and red balls, for a total of 10 balls. One of our experimental manipulations is participants' prior belief about the state that we control by varying the number of blue and red balls in the urn. We model the information sources as imperfectly informed 'experts.' Before making their guess, participants have to consult either the *Blue Expert* or the *Red Expert*, randomly extracted from two populations of experts. In each population, a certain fraction of experts is informed about the true color of the extracted ball and issues a truthful report revealing such color. The complementary fraction of experts is uninformed about the color of the extracted ball and always issues the same report.⁴ Both experts can be consulted for free, but participants are forced to choose only one of them. We used the *strategy method* to elicit participants' guesses about the color of the ball conditional on the expert's signal. On the same screen, we elicited their confidence in each

³Instructions were followed by three multiple-choice questions to verify that participants understand the details of the experiment. After answering each of these questions, subjects see a commented feedback page with the correct answers and a further explanation of the reasoning leading to the correct answer. Appendix C reports observed behavior in the subsamples determined by the number of questions answered correctly in the comprehension quiz. In addition, participants were required to spend a minimum amount of time on each page of the instructions and could not continue to the following page until a specified amount of time (ranging from 30 to 60 seconds) had elapsed.

⁴While this implementation eases participants' understanding of random variables, participants essentially face two Blue and two Red sources – one accurate and one uninformative. This implementation is equivalent to our theoretical framework for expected utility maximizers, but it may not be neutral for individuals who evaluate uncertain prospects differently. These potential effects are discussed in Section 4.

of these guesses, on a scale between 0 and 100. We used these statements to construct a measure of observed posterior beliefs.⁵

Rounds. The discussion above describes one round of the experiment. The experiment consists in a sequence of five rounds. In each round, the computer draws the state of the world and the messages sent by the two experts from the same distributions and independently from any past action or outcome. At the end of each round, participants learned the expert's signal, their relevant choice given the signal, the color of the extracted ball, and their payoff in that round.

Payoffs. On top of earning a fixed amount of \$1 for taking part in the experiment, subjects are remunerated with \$1 for guessing the color of the ball correctly in a randomly selected round. Since recent research shows that complex incentive schemes systematically bias truthful reporting of beliefs (Danz et al., 2022), we stressed the importance of revealing truthful confidence assessments but did not incentivize these statements.

Treatments. We employ a between-subjects design, where we manipulate the prior belief that the ball drawn from the jar is blue, π , and the relative reliability of the two experts, (λ_R, λ_B) . We consider both a mildly and a strongly unbalanced prior, respectively $\pi = 0.6$ and $\pi = 0.8$. Regarding the sources' bias, we consider the case where Blue and Red Experts are equally reliable, $(\lambda_R, \lambda_B) = (0.5, 0.5)$, and the case where the Red Expert is more reliable, i.e. $(\lambda_R, \lambda_B) = (0.3, 0.7)$. This leads to four experimental treatments:

- **E6:** equal reliability, prior mildly favors ball being blue;
- **E8:** equal reliability, prior strongly favors ball being blue;
- **S6:** skewed reliability (Red is more reliable), prior mildly favors ball being blue;
- **S8:** skewed reliability (Red is more reliable), prior strongly favors ball being blue.

These four treatments have been designed to test the key predictions of the model, as summarized in Section 2: only when the Red Expert is more reliable and when the prior is mildly unbalanced – that is, in treatment S6 – it is optimal to consult the contrarian expert. In all other treatments, it is optimal to consult the supportive expert.

4. Experimental results

Figure 1 shows the percentage of decisions where subjects consulted the Blue Expert – that is, the expert biased in favor of the prior – disaggregated by treatment. When information sources are equally reliable, this happens in 66.3% of decisions with mildly unbalanced priors (treatment E6) and in 70.2% of decisions with strongly unbalanced priors (treatment E8). These proportions are statistically different from 50%, according to one-sample tests of proportions (p-values < 0.001). This behavior is in line with hypothesis H1, as the Blue Expert is always the optimal choice in these environments. When the information source biased against the prior (i.e., the Red Expert) is more reliable, the Blue Expert is chosen in 24% of decisions with mildly unbalanced priors (treatment S6) and in 24.3% of decisions with strongly unbalanced priors (treatment S8). These proportions are statistically different from 50%, according to one-sample tests of proportions (p-values < 0.001).

When comparing outcomes across treatments, we use random-effects logistic regressions, which take into account the panel nature of the data (that is, the fact that the same individual contributes more than one observation to the dataset). The estimates from these regressions are presented in Table 2.

⁵We mapped a confidence of 0 – that is, 'I think it is just as likely that I am right or wrong' – to a posterior belief of 0.5 (i.e., indifference between guessing blue and guessing red) and a confidence of 100 – that is, 'I think I am sure my guess is correct' – to a posterior of 1 (i.e., almost certainty in the choice). Intermediate levels of confidence were mapped proportionally to intermediate posteriors between 0.5 and 1.




Fig. 1 Information acquisition by treatment: theory vs. Observed.

Notes: The theoretical probabilities are 100% for E8, S8, E6, and 0% for S6. The observed probabilities are 70.19% for E8 ($N = 265$, $SE = 2.81$), 24.26% for S8 ($N = 235$, $SE = 2.80$), 66.27% for E6 ($N = 255$, $SE = 2.96$), and 24.00% for S6 ($N = 250$, $SE = 2.70$). The black vertical lines represent 95% confidence intervals.

Keeping the sources' relative reliability constant (equal or skewed) and manipulating the prior belief about the state from a mildly unbalanced one (0.6) to a strongly unbalanced one (0.8) does not affect the propensity to consult the Blue Expert (the p-value of E6 vs E8 is 0.595; the p-value for S6 vs S8 is 0.763). On the other hand, keeping the prior belief about the state constant (0.6 or 0.8) and manipulating the relative reliability of the sources from equal to being skewed in favor of Red strongly decreases the chance of consulting the Blue Expert: the difference between E6 and S6 (-42.4%) and the difference between E8 and S8 (-45.8%) are both statistically significant at the 1% level (p-values < 0.0001). This highlights that relative reliability trumps the importance of the prior in subjects' considerations. The regression estimated in the last column of [Table 2](#) confirms that, contrary to the theoretical predictions, subjects are equally sensitive to sources' reliability in treatments with a mildly unbalanced prior and in treatments with a strongly unbalanced prior.⁶ Findings 1 and 2 below summarize this discussion.

Finding 1. *When information sources are equally reliable, subjects are more likely to acquire information from the source biased towards the prior, which is the optimal choice. This behavior is in line with hypothesis H1.*

Finding 2. *When the source biased against the prior is more reliable, subjects are more likely to acquire information from the more reliable source, regardless of the prior and whether this is the optimal choice. This behavior is in contrast with hypothesis H2.*

Even when subjects are more likely to choose the optimal source of information (in treatments E6, E8, and S6), they are prone to mistakes: when information sources are equally reliable, they listen too often to the expert biased against the prior (33.6% of decisions in E6 and 29.8% of decisions

⁶The p-value for the coefficient interacting Skewed Reliability with Strongly Unbalanced Prior is 0.585.

Table 2. Random-effects logistic regressions to estimate ATEs

	(1)	(2)	(3)	(4)	(5)
Treatment E8	0.346 (0.653)				
Treatment S8		−0.254 (0.840)			
Treatment S6			−3.929*** (0.766)		
Treatment E8				5.195*** (0.919)	
Skewed Reliability					−4.203*** (0.780)
Strongly Unbalanced					0.374 (0.720)
Skewed Reliability x Strongly Unbalanced					−0.573 (1.050)
Constant	1.412** (0.479)	−3.058*** (0.637)	1.439** (0.492)	−3.126*** (0.646)	1.537** (0.523)
Baseline	E6	S6	E6	S8	E6
Observations	520	485	505	500	1005
ln(σ²)	2.114*** (0.299)	2.734*** (0.322)	2.176*** (0.308)	2.599*** (0.318)	2.376*** (0.222)

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

in E8); when the Red Expert is more reliable and the uncertainty on the state is sufficiently strong, they listen too often to the expert biased in favor of the prior (24% of decisions in S6). Mistakes are, of course, even more frequent when subjects are more likely to consult the less informative expert (in treatment S8, when this happens in 75.7% of decisions).

Figure 2 presents descriptive statistics for participants’ guesses about the state of the world by treatment and information set. Regardless of the treatment, the vast majority of participants (that is, between 91.7% and 100%) uses the available information optimally and guesses Blue when either expert says Blue. Participants are more reluctant to guess Blue when either expert says Red, but this is only partially due to Bayesian thinking: when the optimal guess is indeed Red (that is, in all treatments when consulting the Blue expert and in treatments with a mildly unbalanced prior when consulting the Red expert), they guess Blue between 14.5% (E8 and Blue expert) and 48.8% of the time (E6 and Red expert). When the optimal guess is Blue, they do so only between 32% and 53.2% of the time.

To quantify the cost of these mistakes (at both the information acquisition and information processing stages), Table 3 reports the average guessing accuracy improvement over the prior – that is, the change in the probability of correctly guessing the state relative to simply following the prior – disaggregated by treatment. We compare this with two benchmarks: the guessing accuracy improvement by hypothetical subjects who choose the same information source as actual subjects but process the information as Bayesian learners; and the guessing accuracy improvement by hypothetical subjects who choose the optimal information source and process the information as Bayesian learners.⁷

⁷The theoretical accuracy improvements over the prior when choosing the optimal source and updating beliefs as a Bayesian learner (which coincide with the empirical ones only in the limit as the sample size grows larger) are + 10% for E8, + 6% for S8, + 20% for E6 and + 22% for S6.

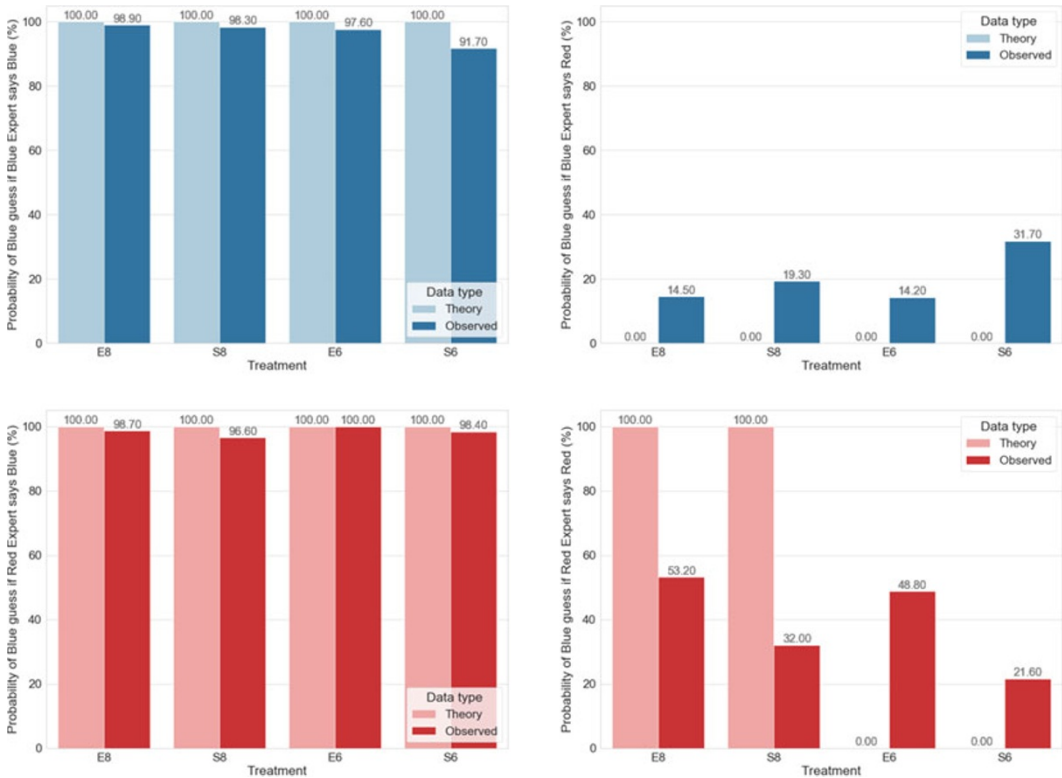


Fig. 2 Guess on the state by treatment and information set: theory vs. Observed data.

Table 3. Average guessing accuracy improvement over prior by treatment

Treatment	N	Observed Source & Observed Guess	Observed Source & Optimal Guess	Optimal Source & Optimal Guess
E8	265	+1.13 (3.30)	+7.17 (3.05)	+8.68 (2.97)
S8	235	-3.40 (3.88)	+1.70 (3.72)	+7.66 (3.48)
E6	255	+6.67 (4.07)	+9.41 (4.00)	+16.86 (3.78)
S6	250	+14.80 (4.24)	+27.20 (3.90)	+28.40 (3.85)

Notes: Since in all treatments the prior is that θ is more likely to be B , the counterfactual probability of correctly guessing θ without any additional information is given by the empirical frequency of $\theta = B$. Thus, in columns 3–5, we compute the average guessing accuracy improvement as the difference between the empirical frequency of a correct guess (in three different scenarios) and the empirical frequency of $\theta = B$. The experimental software generates a state of the world and a signal for each source (independently for each participant and each round) before the participant chooses the source. This allows us to construct the counterfactual in the last column (where, if the participant chose the sub-optimal source, we use the signal from the other source, unobserved by the participant but available in our dataset). We report standard errors in parentheses.

Finding 3. *Subjects improve guessing accuracy less than they could in all treatments. Indeed, when experts have asymmetric reliability and the prior is strongly unbalanced, subjects make worse guesses than they would simply following their priors.*

This result is due, in part, to subjects making sub-optimal use of the information provided by experts (regardless of whether the chosen information source was optimal or not): the improvement in average accuracy that could be obtained without changing information source but adopting Bayesian inference ranges between 2.8% (in treatment E6) to 12.4% (in treatment S6). At the same

time, choosing a suboptimal information source also has a cost in terms of guessing accuracy, especially in treatments S8 and E6.

In order to shed light on the motives underlying information acquisition, Appendix B investigates how subjects use the advice received by the information source of choice.⁸ We find that, as predicted, subjects are deferential to confirmatory advice. On the other hand, subjects follow contradictory advice sub-optimally: they are excessively skeptic of contradictory advice by the source biased towards the prior and excessively trusting of contradictory advice by the source biased against the prior. This pattern is particularly pronounced in treatment S8, where the most substantial mistakes are observed.

To understand this deviation from the Bayesian benchmark, in the same Appendix, we analyze posterior beliefs about the state of the world. With the exception of the case where learning is easiest (that is, when the prior is strongly unbalanced and the Blue expert says *b*), observed posterior beliefs are different from those of a Bayesian learner: the change with respect to the prior is excessive when advice is in line with the source bias and insufficient when advice is against the source bias. For example, participants often fail to act on the fully revealing *r* signal from the Blue source. We conclude that subjects are excessively responsive to information aligned with a source bias and insufficiently responsive to information misaligned with a source bias (which, in fact, perfectly reveals the state of the world).

Overall, these findings suggest that participants' behavior can be explained by three key drivers: base rate neglect, a reliability heuristic, and a certainty seeking heuristic. First, participants appear to underweight the prior probability of the state of the world when deciding which source to consult. For example, even in treatments where the prior is strongly unbalanced, participants often choose the more reliable source, ignoring the fact that the prior heavily favors one state. This behavior is consistent with base rate neglect, a well documented cognitive bias in probabilistic reasoning, where individuals fail to adequately incorporate prior probabilities into their judgments.

Second, participants' preference for the more reliable source in treatments with asymmetric reliability highlights the use of a simple heuristic: "choose the most reliable source."

While this heuristic simplifies decision-making, it leads to suboptimal information acquisition in scenarios where the less reliable source is optimal given the prior. Third, ambiguity aversion provides another possible explanation for participants' behavior, particularly in their preference for more reliable sources. Ambiguity-averse individuals might seek to maximize their chances of receiving a signal that identifies the states with certainty, and a more reliable source is indeed more likely to give advice misaligned with its bias, an unambiguous signal that completely removes any uncertainty about the state of the world. The chance that the Red experts says blue is increasing in both the Red experts' reliability and in the prior belief that the state is B (while the chance that the Blue expert says red is decreasing in this belief), making the Red expert particularly appealing for certainty-seeking individuals in treatment S8, the treatment where we observe the largest incidence of mistakes. At the same time, ambiguity aversion alone cannot account for all observed patterns: participants' underreaction to fully revealing information, as revealed both by their guesses and their posterior beliefs about the state of the world, contradicts the certainty-seeking heuristic ('choose the information structure most likely to give unambiguous signals') associated with ambiguity aversion.

These behavioral tendencies – base rate neglect and reliance on heuristics – have important implications for understanding how individuals navigate complex environments with conflicting and biased information sources. They suggest that decision-makers prioritize cues like reliability and clarity over statistical optimality, especially when faced with asymmetry in source trustworthiness. Such patterns mirror real-world challenges where individuals often rely on heuristics or biases to simplify decisions, sometimes at the expense of optimal outcomes.

⁸We must note that interpreting these results is complicated, at least in part, by self selection, as subjects choose their information source.

5. Conclusion

This paper formalized a model of selective exposure based on Bayesian updating, and tested its predictions through a laboratory experiment. We ask two research questions: when is it rational to seek (dis)confirmatory information? Do people behave according to rationality or do we need to impose additional structures? Overall, our experiment suggests that explaining selective exposure to information sources with Bayesian inference has some limitations: in line with Bayesian learning, we do observe confirmatory patterns in the selection of information when sources are equally reliable; at the same time, these trends switch to disconfirmatory attitudes as soon as the source biased against the prior becomes more reliable, with no role for the strength of prior beliefs. We see many possible directions for future research: while we study the simplest possible setup to investigate selective exposure to information sources, it would be interesting to investigate more complex environments where decision-makers have the opportunity to collect multiple pieces of information from sources, or must pay a (possibly heterogeneous) price to receive messages from a source.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/esa.2025.8>.

Statements and Declarations. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

References

- Ambuehl, S. (2021). Can incentives cause harm? Tests of undue inducement. *Unpublished Manuscript*.
- Ambuehl, S., & Shengwu, L. (2018). Belief updating and the demand for information. *Games and Economic Behavior*, 109, 21–39.
- Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *American Economic Review*, 87(5), 847–862.
- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239), 1130–1132.
- Barber'a, P., John, T. J., Jonathan, N., Joshua, A. T., & Richard, B. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, 26(10), 1531–1542.
- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44(3), 211–233.
- Baron, J., Beattie, J., & Hershey, J. C. (1988). Heuristics and biases in diagnostic reasoning II: Congruence, information, and certainty. *Organizational Behavior and Human Decision Processes*, 42(1), 88–110.
- Berelson, B., & Steiner, G. A. (1964). *Human behavior: An inventory of scientific findings*. New York: Harcourt, Brace & World.
- Castagnetti, A., & Schmacker, R. (2022). Protecting the ego: Motivated information selection and updating. *European Economic Review*, 142, 104007.
- Charness, G., Oprea, R., & Yuksel, S. (2021). How do people choose between biased information sources? Evidence from a laboratory experiment. *Journal of the European Economic Association*, 19(3), 1656–1691.
- Chopra, F., Haaland, I., & Roth, C. (2024). The demand for news: accuracy concerns versus belief confirmation motives. *The Economic Journal*, 134(661), 1806–1834.
- Danz, D., Vesterlund, L., & Wilson, A. J. (2022). Belief elicitation and behavioral incentive compatibility. *American Economic Review*, 112(9), 2851–2883.
- Duffy, J., Hopkins, E., Kornienko, T., & Mingye, M. (2019). Information choice in a social learning experiment. *Games and Economic Behavior*, 118, 295–315.
- Esponda, I., Vespa, E., & Yuksel, S. (2024). Mental models and learning: The case of base-rate neglect. *American Economic Review*, 114(3), 752–782.
- Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1), 298–320.
- Frey, D. (1986). Recent research on selective exposure to information. *Advances in Experimental Social Psychology*, 19, 41–80.
- Gentzkow, M., & Shapiro, J. M. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2), 280–316.
- Gentzkow, M., & Shapiro, J. M. (2011). Ideological segregation online and offline. *Quarterly Journal of Economics*, 126(4), 1799–1839.
- Iyengar, S., & Hahn, K. S. (2009). Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1), 19–39.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.
- Klayman, J., & Young-Won, H. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2), 211.

- Lawrence, E., Sides, J., & Farrell, H. (2010). Self-segregation or deliberation? Blog readership, participation, and polarization in American politics. *Perspectives on Politics*, 8(1), 141–157.
- Mann, T. E., & Ornstein, N. J. (2012). *It's even worse than it looks: How the American constitutional system collided with the new politics of extremism*. New York, NY: Basic Books.
- Mullainathan, S., & Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4), 1031–1053.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175.
- Peterson, E., Goel, S., & Iyengar, S. (2021). Partisan selective exposure in online news consumption: Evidence from the 2016 presidential campaign. *Political Science Research and Methods*, 9(2), 242–258.
- Sharma, K., & Castagnetti, A. (2023). Demand for information by gender: An experimental study. *Journal of Economic Behavior & Organization*, 207, 172–202.
- Skov, R. B., & Sherman, S. J. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, 22(2), 93–121.
- Slowiaczek, L. M., Klayman, J., Sherman, S. J., & Skov, R. B. (1992). Information selection and use in hypothesis testing: What is a good question, and what is a good answer? *Memory & Cognition*, 20(4), 392–405.