Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Audi Alteram Partem:
an Experiment on Selective Exposure to Information

Salvatore Nunnari and Giovanni Montanari

Working Paper n. 650

This Version: June, 2019

Audi Alteram Partem: An Experiment on Selective Exposure to Information

Giovanni Montanari

New York University

Salvatore Nunnari

Bocconi University, IGIER, CEPR

June 5, 2019

Abstract

This paper presents a model of selective exposure to information and an experiment to test its predictions. An agent interested in learning about an uncertain state of the world can acquire information from one of two sources which have opposite biases: when informed on the state, they report it truthfully; when uninformed, they report their favorite state. When sources have the same reliability, a Bayesian agent is better off seeking *confirmatory information*. On the other hand, it is optimal to seek *contradictory information* if and only if the source biased against the prior is sufficiently more reliable. We test these predictions with an online experiment. When sources are symmetrically reliable, subjects are more likely to seek confirmatory information but they *listen to the other side* too frequently. When sources are asymmetrically reliable, subjects are more likely to consult the more reliable source even when prior beliefs are strongly unbalanced and listening to the less reliable source is more informative. Moreover, subjects follow contradictory advice sub-optimally; are too trusting of information in line with a source bias; and too skeptic of information misaligned with a source bias. Our experiment suggests that biases in information processing and simple heuristics — e.g., listen to the more reliable source — are important drivers of the endogenous acquisition of information.

Keywords: Choice under Uncertainty, Information Acquisition, Bayesian Updating, Selective Exposure, Confirmation Bias, Limited Attention, Online Experiment

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. In particular, existing research in sociology, social psychology, political science and economics strongly suggests that people tend to look for information that is consistent with their world view (Sears and Freedman 1967, Frey 1986, Gunther 1992, Klayman 1995, Nickerson 1998, Iyengar and Hahn 2009). In the words of Berelson and Steiner (1968), "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" (pp. 529–530).

This patten of selective exposure to information contrasts with the general wisdom that the other side in a debate should always be heard — a principle that dates back to Ancient Greece[1] and that lies at the heart of the contemporary legal tradition (*audi alteram partem* or *listen to the other side*). Moreover, this behavior has raised concern, both among social scientists and in the public opinion: as the availability of media choices has been growing, selective exposure to like-minded sources has contributed to a deep partisan divide in news consumption (Lawrence et al. 2010, Gentzkow and Shapiro 2011, n.d., Del Vicario et al. 2016, Quattrociocchi et al. 2016, Peterson et al. 2018). This segregation into "echo chambers" has been associated with the observed intensification of partisan sentiment as well as with the recent populist insurgencies in the Western world (Mann and Ornstein 2012, Bakshy et al. 2015, Flaxman et al. 2016).

Why do we observe this behavior? Recent theoretical work in microeconomics suggests that individuals might have systematic preferences for information consonant with their beliefs (Mullainathan and Shleifer 2005) or that, being uncertain about an information source reliability, they interpret disconfirming evidence as less credible than confirming evidence and turn their attention towards the source they deem as more informative (Gentzkow and Shapiro 2006). Notably, even when individuals have no uncertainty about sources reliability

---

[1]Consider Aeschylus, "The Eumenides" 431, 435

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 an online 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?

We introduce a simple model of optimal choice between two different information structures and test experimentally whether our theory can account for the observed patterns of information selection. In our model, decision makers have the possibility to acquire a *signal* from one of two *information structures* in order to reduce their uncertainty about a payoff relevant state of the world. Information structures stochastically map the state of the world to a signal and are biased towards different states. Importantly, decision makers know the conditional distributions of signals for each information structure, ruling out any uncertainty about the *reliability* of information sources. We also provide decision makers with an exogenous prior belief on the true state of the world and focus on an abstract decision environment which allow us to minimize the confounding effects of *desirability bias* or *motivated beliefs*.[2] Once decision makers observe the signal from the information structure of choice, they guess the state of the world they deem more likely and receive a positive payoff only for a correct guess. We manipulate experimentally the probability distributions of signals delivered by each information structure (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 structures in some treatments but not in others. We verify optimal

---

[2]Desirability bias arises when an agent's actions reflect not only a probabilistic belief over possible realizations of the state of the world, but also a desire or preference with respect to such states (Tappin et al. 2017). As Epley and Gilovich (2016) put it, motivated beliefs capture the way "people generally reason their way to conclusions they favor, with their preferences influencing the way evidence is gathered, arguments are processed, and memories of past experience are recalled. Each of these processes can be affected in subtle ways by peoples motivations, leading to biased beliefs that feel objective".

information acquisition in both environments and test for a confirmatory pattern on top 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 more informative, participants display a *dis-confirmatory 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 under-weigh the importance of the ex-ante uncertainty surrounding the phenomenon to learn about.

In order to shed light on the motives underlying information acquisition, we investigate how subjects use the advice received by the information source of choice. 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. Moreover, subjects are insufficiently responsive to information misaligned with a source bias (which, in fact, perfectly reveals the state of the world) and excessively responsive to information aligned with a source bias (in particular, they are excessively responsive to confirmatory advice by the source biased towards the prior when the prior is mildly unbalanced; and excessively responsive to contradictory advice by the source biased against the prior when the prior is strongly unbalanced). This suggests that biases in information processing and simple heuristics — e.g., listen to the more reliable source — are important drivers of the endogenous acquisition of information.

This paper contributes to three strands of literatures. First, our paper contributes to a literature in experimental psychology on how people gather evidence to test hypotheses

(Skov and Sherman 1986, Klayman and Ha 1987, Baron et al. 1988, Slowiaczek et al. 1992).[3] Second, our paper is related to a literature in experimental economics studying learning from new information and documenting deviations from Bayesian inference (Tversky and Kahneman 1971, Grether 1980, Viscusi and O'Connor 1984, Hoffrage et al. 2000, Charness and Levin 2005, Dohmen et al. 2009). Third and more directly, our paper contributes to a recent literature in experimental economics on the choice over sources of information structures with instrumental value (Ambuehl and Li 2018, Duffy et al. 2017, 2018).[4] The most closely related work is Charness et al. (2018). Similarly to some of our treatments (namely, E6 and E8), they consider experimental conditions (labeled *bias by commission*), where decision-makers choose between two information structures which 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 treatments where the two available information structures are biased towards opposite states and might send an incorrect signal with different probabilities (that is, they are asymmetrically reliable).[5]

The remainder of the paper proceeds as follows. In Section 2, we introduce a simple model of choice between information structures and present the testable hypotheses. Section 3 details our experimental design. We describe the experimental results in Section 4. Section 5 concludes and discusses directions for future research.

---

[3]Testing 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 "*not q* implies *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 in not designed to test the ability to construct a logical test but rather the endogenous acquisition of an informative signal.

[4]Less related to this paper, Zimmermann (2014), Falk and Zimmermann (2017), Masatlioglu et al. (2017), Nielsen (2018) investigate the choice over sources of information structures in settings where information has no instrumental value.

[5]Charness et al. (2018) also consider experimental treatments where the two information structures 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 experimental treatments where the two information structures are biased towards opposite states and might fail to send a signal (labeled *bias by omission*). In further contrast with Charness et al. (2018), we use a between-subject rather than within-subject design and run the study online with a sample including both students and non-students rather than in the laboratory with students. See Section 3 for additional details.

2 Task and Theoretical Predictions

Assume there is a binary state of the world, $\theta = \{B, R\}$, and consider a decision maker (DM) who has to make a guess, $a = \{B, R\}$, about the true state. The DM's payoff depends both on this guess and on the realized state of the world:

$$u(a, \theta) = \begin{cases} 1 \text{ if } a = \theta \\ 0 \text{ if } a \neq \theta \end{cases}$$

The DM holds a prior belief, π, about the event $\theta = B$. We consider unbalanced prior beliefs and, without loss of generality, we assume $\pi \in (1/2, 1)$. Before making a guess, the DM acquires a piece of information from an information structure σ. The information structure (or *source*) stochastically maps the state of the world to a *signal* $s = \{b, r\}$, where s signals the state is S (that is, b signals the state is B and r signals the state is R). The DM knows *ex ante* the distribution of signals, ruling out any uncertainty over information structures. We consider two different information structures, **Blue** and **Red**, described by the stochastic matrices in Table 1:

(a) **Blue** Information Structure

	$s = b$	$s = r$
$\theta = B$	1	0
$\theta = R$	λ_B	$1 - \lambda_B$

(b) **Red** Information Structure

	$s = b$	$s = r$
$\theta = B$	$1 - \lambda_R$	λ_R
$\theta = R$	0	1

Table 1: Conditional Distribution of Signals by Information Structure

Each cell of Table 1 displays the probability of observing a signal (column) in a specific state of the world (row). For instance, λ_B represents the probability that **Blue** signals the state is B when it is, in fact, R. The smaller the λ_σ, the more *reliable* the information structure σ: an information structure always reports the true state of the world if $\lambda_\sigma = 0$; it always reports the same signal, regardless of the true state of the world if $\lambda_\sigma = 1$. We assume that both sources are informative, that is, $\lambda_\sigma < 1$, $\sigma = \{R, B\}$. Intuitively, we would say that an information structure is *unbiased* when it always reports the state of the world.

We exploit this intuition to introduce the following definition of bias:

Definition 1 *(Bias)*

An information structure σ is biased towards state S if

1. *it signals s when the state is S almost surely;*

2. *it signals s when the state is not S with positive probability.*

Using this definition, the two information structures display a "partisan bias": **Blue** is biased towards B and **Red** is biased towards R. In line with Gentzkow and Shapiro (2006) we argue this model of distortion can accommodate different real-world interpretations: the information source 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. Our simple framework encompasses all these possibilities, since we do not impose any further condition on the information structure. What is relevant to us is the information observed by the DM.

The timing of the events is as follows. Given the prior π, the DM chooses whether to acquire information from **Blue** or **Red**. The DM then observes a signal from the chosen information structure and updates her prior belief. Finally, she submits her guess. The ensuing payoff is 1 if the guess matches the state of the world, and 0 otherwise.

Our object of analysis is the information acquisition behavior of the DM when faced with different information sources, and how this varies with her prior belief and the available sources' reliability. Note that the only goal of the DM in this framework is to maximize her expected payoff, which she can do by picking the most informative source and processing its signal as a Bayesian learner. Indeed, after observing a signal from an information source, the expected payoff from a guess coincides with the posterior belief this guess matches the true state of the world. Hence, posterior beliefs fully characterize the optimal choice of the DM.

We characterize the DM's optimal choice of an 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 structure and how the DM will use these signals. In what follows, the notation $a^\star(s, \sigma)$ denotes the optimal guess after observing signal s from information structure σ. Posterior beliefs are denoted by $Pr(\theta|s, \sigma)$. All proofs are reported in the Appendix.

2.1 Optimal Guess for Given Information Source

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

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

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. This means that, in this case, the DM sticks with her prior belief and guesses accordingly. Receiving a signal which disagrees with the source bias — that is, receiving signal b (r) from the **Red** (**Blue**) information structure — is fully revealing: the DM learns the state with certainty, independently of her prior beliefs and the information 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

We now move a step backward in the DM's problem and consider her decision about what source to acquire information from. First, consider the expected payoff from consulting the source biased in favor of the prior, that is, **Blue**. As discussed above, the DM follows any signal received from this source. Thus, her posterior belief that she is making the correct guess is $Pr(\theta = R | r, \textbf{Blue}) = 1$ following a contradictory signal; and $Pr(\theta = B | b, \textbf{Blue}) = \frac{\pi}{\pi + (1-\pi)\lambda_B}$ following a confirmatory signal. Weighing these posterior beliefs with the unconditional distribution of signals by this source, we get the following expected payoff:

Lemma 3 (Expected Utility from Blue Source) $\mathbb{E}[\textbf{Blue}] = 1 - (1 - \pi)\lambda_B \in [\pi, 1]$ *is increasing in* π *and decreasing in* λ_B.

Accessing this source always improves the confidence the DM has in her guess with respect to the decision she would made without collecting any additional information. The benefit is increasing in the likelihood its bias is in line with the true state and with its reliability.

Consider now the expected payoff from consulting the source biased against the prior, that is, **Red**. As discussed above, when this source is not sufficiently reliable — that is, $\lambda_R \geq \frac{1-\pi}{\pi}$ — the DM guesses B regardless of the signal. Thus, in this case, the expected payoff from consulting this information source is equal to the expected payoff from not consulting any source and deciding on the basis of the prior belief. On the other hand, when $\lambda_R < \frac{1-\pi}{\pi}$, the DM follows any signal received from **Red**. In this case, her posterior belief that she is making the correct guess is $Pr(\theta = R | r, \textbf{Red}) = \frac{1-\pi}{\pi\lambda_R + (1-\pi)}$ following a contradictory signal; and $Pr(\theta = B | b, \textbf{Red}) = 1$ following a confirmatory signal. Weighing these posterior beliefs with the unconditional distribution of signals by this source, we get the following expected payoff:

Lemma 4 (Expected Utility from Red Source) *If $\lambda_R \geq \frac{1-\pi}{\pi}$, $\mathbb{E}[\textbf{Red}] = \pi$. If, instead, $\lambda_R < \frac{1-\pi}{\pi}$, $\mathbb{E}[\textbf{Red}] = 1 - \pi\lambda_R \in [\pi, 1]$, decreasing in π and in λ_R.*

Accessing this source improves the confidence the DM has in her guess only when it is sufficiently reliable. In this case, the benefit is increasing in the likelihood its bias is in line with the true state and with its reliability.

Since consulting the source biased in favor of the prior is always informative, that is, $\mathbb{E}[\textbf{Blue}] > \pi$, 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 information structure with the smallest probability of misleading signals. If the DM has a perfectly balanced prior, that is, $\pi = 1/2$, choosing **Red** over **Blue** reduces to $\lambda_R < \lambda_B$. When the prior is unbalanced, that is, $\pi > 1/2$, the DM has an incentive to choose the information structure which 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 1 summarizes this discussion and characterizes this threshold:

Proposition 1 (Optimal Source) *The DM consults **Red** if and only if $\lambda_R < \frac{1-\pi}{\pi}\lambda_B$ and **Blue** otherwise.*

2.3   Summary of Testable Hypotheses

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

On Information Acquisition

H1 *When information sources are equally reliable, it is optimal to seek confirming information, that is, to acquire information from the source biased towards the prior.*

H2 *When the source biased against the prior is more reliable, it is optimal to acquire information from the more reliable source (even if biased against the prior) if the prior*

is mildly unbalanced; it is optimal to acquire information from the source biased towards the prior (even if less reliable) if the prior is strongly unbalanced.

On Information Processing

H3 *It is always optimal to follow information from the source biased towards the prior.*

H4 *It is always optimal to follow confirming information from the source biased against the prior. Contradictory information from the source biased against the prior should be followed when the prior is mildly unbalanced and ignored when the prior is strongly unbalanced.*

3   Experimental Design

The experiment was conducted online on Prolific, a crowdsourcing platform for academic research. Subjects were recruited from the Prolific database of participants and screened by their characteristics: only American citizens currently residing in the U.S. whose first language was English were eligible for participation. A total of 201 subjects took part in the experiment, and no subject was able to sit for the experiment twice. While not representative of the American population, our sample is more representative than traditional samples composed of undergraduate students at elite universities.[6] The use of web-based experiments is relatively novel in experimental economics. While this methodology requires additional precautions in the design and in the instructions in order to ensure continuous attention[7], research suggests test results are in line with those obtained in more controlled

---

[6]In our sample, age ranged between 19 and 75 years old, with an average age of 33.4 ($N = 192$ out of 201 participants); 51% of participants were women ($N = 198$); 28.4% of participants were students ($N = 197$); 56.3% of participants were full time workers ($N = 192$); 77% of participants had at least some college education ($N = 196$); 75.8% of participants were caucasian; the median personal income was in the \$40k-\$50k bracket ($N = 169$); the median household income was in the \$65k-\$75k bracket ($N = 157$); 50.8% of participants were Democrat (with 19.8% being a Republican and 29.5% being an Independent, $N = 193$). These statistics are based on self-declarations collected by Prolific.

[7]The rest of this Section details how we achieved this.

environments such as laboratories (**?**, Snowberg and Yariv 2018). Instructions and sample screens are reported in Appendix C.[8]

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 first introduced to the problem of guessing the color of a ball (either red or blue) randomly drawn from a jar containing only blue and red balls, for a total of 10 balls. One of our experimental manipulations is participants' prior belief (π) about the state of the world. We provide participants with a prior belief on the color of the extracted ball by displaying information on the distribution of the balls in the jar.

Next, we model the information structures as imperfectly informed "experts". Before making their guess, participants have to consult either the *Blue Expert* or the *Red Expert*, randomly extracted from two population of experts (the Red population and the Blue population). In each population, a certain fraction of experts is informed about the true color of the extracted ball and issues a report revealing such true color. The complementary fraction of experts is uninformed about the color of the extracted ball and always issues the same report. In particular, a random *Blue Expert* is informed with probability $(1 - \lambda_B)$ and uninformed with probability λ_B. When uninformed, a *Blue Expert* always reports the color of the extracted ball is blue. Analogously, picking randomly a *Red Expert* implies he is informed with probability $(1 - \lambda_R)$ and uninformed with probability λ_R. When uninformed, the expert always reports the color of the ball is *red*. Both experts can be consulted for free, but participants are forced to choose only one of them.

Choosing the expert prompts participants to the following screen, where we use the *strategy method* and elicit a subject's guess about the color of the ball conditional on the expert's signal. On the same screen, we elicit a subject's confidence in each of these two conditional guesses, on a scale between 0 and 100.[9] Once these choices have been recorded,

---

[8]The user interface was programmed with oTree (Chen et al. 2016).

[9]As detailed in the Instructions available in Appendix C, we elicit confidence with the following question: "On a scale from 0 to 100, how confident are you about this guess? For example, 0 means that you think it

participants proceed to the final screen where they see the expert's signal, their relevant choice given the signal, the color of the extracted ball, and their payoff.

Instructions are displayed in the first screens of the experiment and followed by three multiple-choice questions to verify 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.[10] Each page of the instructions is timed, so that participants must spend no less than a specified amount of time on each page, ranging from 30 to 60 seconds.

Rounds. The discussion above describes one round of the experiment. The experiment consists in a sequence of 5 rounds and it lasts around 10 minutes. In each round, the computer draws the true state of the world and the messages sent by the two experts from the same distributions and independently from any past action or outcome. Playing multiple rounds, subjects familiarize with the structure of the experiment and have room for learning. We opted for a limited number of rounds for two reasons. First, we wanted participants to pay due attention to the instructions and forced them to spend no less than a specified amount of time reading them. Increasing the number of rounds may lead participants to skim quickly through the instructions in order to have more time to formulate each subsequent (remunerated) choice. Second, given the online implementation of the experiment, we wanted to avoid boredom: keeping the number of rounds at a minimum favors attention.

Choices. In each round, we are eliciting five choices from each participant: the expert to consult, the guess on the color of the ball for each possible signal from the chosen expert, and the confidence level surrounding these guesses. The strategy method allows us to record information also for unlikely events. Moreover, we use confidence statements to construct

---

is just as likely that you are right or wrong and 100 means you are sure your guess is correct."

[10]47% of participants (94 out 201) answered all three comprehension questions correctly; 87% (175 out of 201) answered at least two questions correctly; and 99% (199 out of 201) answered at least one question correctly. Results are qualitatively unchanged if we consider only subjects who answered correctly 1, 2 or 3 comprehension questions. Table 6 in Appendix B reports observed behavior in the subsamples determined by the number of questions answered correctly in the comprehension quiz.

a measure of observed posterior beliefs, which we exploit to test for systematic biases in distinct treatments.[11] We do not incentivize confidence statements for two reasons. First, given the experiment was administered online, we deemed as particularly important to keep the tasks simple and the total duration short (below 20 minutes). Explaining carefully a Binary Scoring Rule and guaranteeing a full understanding of the underlying incentives would have required some additional time, possibly more than the duration of the main task, increasing attrition and reducing the quality of decisions throughout the whole experiment. Second, incentivizing the elicitation of posterior beliefs (either with a Binary Scoring Rule or asking for the closest guess to the posterior belief held by a statistician, as in Charness and Dave 2017) could potentially interfere with the fundamental incentive to choose the most informative expert, as it gives an incentive to choose informational sources which are more likely to induce degenerate posterior beliefs. At the same time, in the instructions, we stress the importance of revealing truthful confidence assessments.

Payoffs. On top of earning a fixed amount of \$1 for taking part in the experiment, subjects are remunerated for guessing the color of the ball. In particular, they earn \$1 if their guess matches the color of the drawn ball, \$0 otherwise. At the beginning of the experiment participants are instructed they will play multiple rounds, but only a randomly chosen round will be selected to determine the bonus payment. As discussed above, we do not incentivize the elicitation of a subject's confidence in his/her guess.

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 probability that experts garble their signals, we consider the case where Blue and Red Experts are equally reliable, $(\lambda_R, \lambda_B) = (0.5, 0.5)$, and the case

---

[11]We 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.

(a) $\pi = 0.8$ (b) $\pi = 0.6$

Figure 1: Theoretical Predictions by Treatment

where the Red Expert is more reliable, i.e. $(\lambda_R, \lambda_B) = (0.3, 0.7)$. The combination of these two manipulations lead 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. To understand how our theoretical predictions map to our treatments, consider Figure 1. The red-shaded area represents all combinations of (λ_R, λ_B) such that the Red Expert provides the highest expected payoff, while the blue-shaded area represents all combinations of (λ_R, λ_B) such that it is optimal to consult the Blue Expert. As the slope of the line separating these two areas depends on the prior belief, the right

14

panel is for the mildly unbalanced prior and the left panel for the strongly unbalanced prior. 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. Treatments E8 and E6 allow us to test a crucial prediction of the model, namely that when experts are equally reliable and the prior is unbalanced subjects prefer to get advice from the supportive expert (H1). Treatments S8 and S6 allows us to test the effect of the prior on the incentives to consult the contrarian expert (H2). We can restate the predictions from Section 2.3 in terms of our treatments:

H1 Subjects acquire information from the Blue Expert in both E6 and E8.

H2 Subjects acquire information from the Red Expert in S6, from the Blue Expert in S8.

H3 Subjects always follow signals by the Blue Expert.

H4 Subjects always follow Blue signals by the Red Expert. Subjects follow Red signals by the Red Expert in E6 and S6 but ignore them in E8 and S8.

4   Experimental Results

4.1   Information Acquisition

Table 2 and Figure 2 show 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 66.4% of decisions with mildly unbalanced priors (treatment E6) and in 70.1% 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 more informative in these environments. When the information source biased against the prior (i.e., the Red Expert) is more reliable, the Blue Expert

15

Panel A: Treatment E8 (Equal Reliability, Strongly Unbalanced Prior)

	N	Observed	Theory	p-value vs E6	p-value vs S8
% Chooses Blue Expert	265	70.1	100	0.582	0.000
% Follows Advice if B Says b	186	98.9	100	0.457	0.691
% Follows Advice if B Says r	186	85.5	100	0.962	0.585
% Follows Advice if R Says b	79	98.7	100	0.067	0.358
% Follows Advice if R Says r	79	46.8	0	0.726	0.070

Panel B: Treatment S8 (Skewed Reliability, Strongly Unbalanced Prior)

	N	Observed	Theory	p-value vs E8	p-value vs S6
% Chooses Blue Expert	235	24.3	100	0.000	0.971
% Follows Advice if B Says b	57	98.3	100	0.691	0.124
% Follows Advice if B Says r	57	80.7	100	0.585	0.311
% Follows Advice if R Says b	178	96.6	100	0.358	0.297
% Follows Advice if R Says r	178	68.0	0	0.070	0.212

Panel C: Treatment E6 (Equal Reliability, Midly Unbalanced Prior)

	N	Observed	Theory	p-value vs E8	p-value vs S6
% Chooses Blue Expert	255	66.4	100	0.582	0.000
% Follows Advice if B Says b	169	97.6	100	0.457	0.143
% Follows Advice if B Says r	169	85.8	100	0.962	0.075
% Follows Advice if R Says b	86	91.9	100	0.067	0.021
% Follows Advice if R Says r	86	51.2	100	0.726	0.004

Panel D: Treatment S6 (Skewed Reliability, Mildly Unbalanced Prior)

	N	Observed	Theory	p-value vs E6	p-value vs S8
% Chooses Blue Expert	250	24.0	0	0.000	0.971
% Follows Advice if B Says b	60	91.7	100	0.143	0.124
% Follows Advice if B Says r	60	68.3	100	0.075	0.311
% Follows Advice if R Says b	190	98.4	100	0.021	0.297
% Follows Advice if R Says r	190	78.4	100	0.004	0.212

Table 2: Observed Outcomes by Treatment. Notes: The unit of observation is a decision made by a subject in a round; p-values for comparisons between treatments are based on Probit regressions with standard errors clustered at the subject level.

Figure 2: Information Acquisition by Treatment: Theory vs. Observed Data

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). Moreover, these proportions are statistically indistinguishable from one another, highlighting how relative reliability trumps the importance of the prior in subjects' considerations. This behavior is in contrast with hypothesis H2, since the Red Expert is more informative only when the prior belief is mildly unbalanced. 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 more informative.*

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 source is the more informative or not.*

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 28.8% of decisions 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 75.5% of decisions).

Importantly, these mistake rates come at a cost in terms of accuracy in guessing the state. To show this, 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 the observed guessing accuracy improvement with two benchmarks: the guessing accuracy improvement by hypothetical subjects who choose the same information structure as actual subjects but process the information as Bayesian learners; and the guessing accuracy improvement by hypothetical subjects who choose the optimal information structure and process the information as Bayesian learners. The results show that subjects improve guessing accuracy less than they could in all treatments. Indeed, when experts have asymmetric reliability and the prior is strongly unbalanced, subjects actually make worse guesses than they would simply following their priors. In part, this is due to subjects making sub-optimal use of the information provided by experts (regardless of whether the chosen information structure 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 structure also has a cost in terms of guessing accuracy, especially in treatments S8 and E6.

Finding 3. *In all treatments, subjects frequently acquire information from the less informative source and this leads to sub-optimal learning.*

Treatment	N	Observed	Observed Source & Bayesian Updating	Optimal Source & Bayesian Updating
E8	265	+1.1	+7.2	+8.7
S8	235	-3.4	+1.7	+7.7
E6	255	+6.6	+9.4	+16.8
S6	250	+14.8	+27.2	+28.4

Table 3: Guessing Accuracy Improvement over Prior by Treatment

4.2   Information Processing

To shed light on subjects' choice of information source and understand why they are prone to mistakes, we analyze the use subjects make of the information they obtain from experts. Table 2 and Figure 3 shows the percentage of decisions which follow the advice from the chosen information source, disaggregated by treatment, information source and advice.[12] We define confirmatory advice as a signal which aligns with the prior belief. Pooling together all treatments, subjects follow confirmatory advice from the Blue Expert 97.5% of the time and confirmatory advice from the Red Expert 96.8% of the time. This is in line with our theoretical model: since both information sources are somehow informative, confirmatory advice increases the confidence in the state being the blue one, regardless of the information source it comes from.

Finding 4. *Subjects follow confirmatory advice optimally.*

On the other hand, subjects suffer from biases in interpreting contradictory advice. A Bayesian learner always follows contradictory advice from the Blue Expert, regardless of the prior and the source reliability (indeed, this message perfectly reveals the state). In the experiment, subjects follow a red message by the Blue Expert 85.6% of the time when experts have the same reliability (treatments E6 and E8), and 74.4% of the time when the Blue Expert is less reliable (treatments S6 and S8). The difference between the two pairs of treatments is statistically significant at the 10% level according to a Probit regressions

---

[12]We must note that interpreting these results is complicated, at least in part, by self selection, as subjects choose their information source.
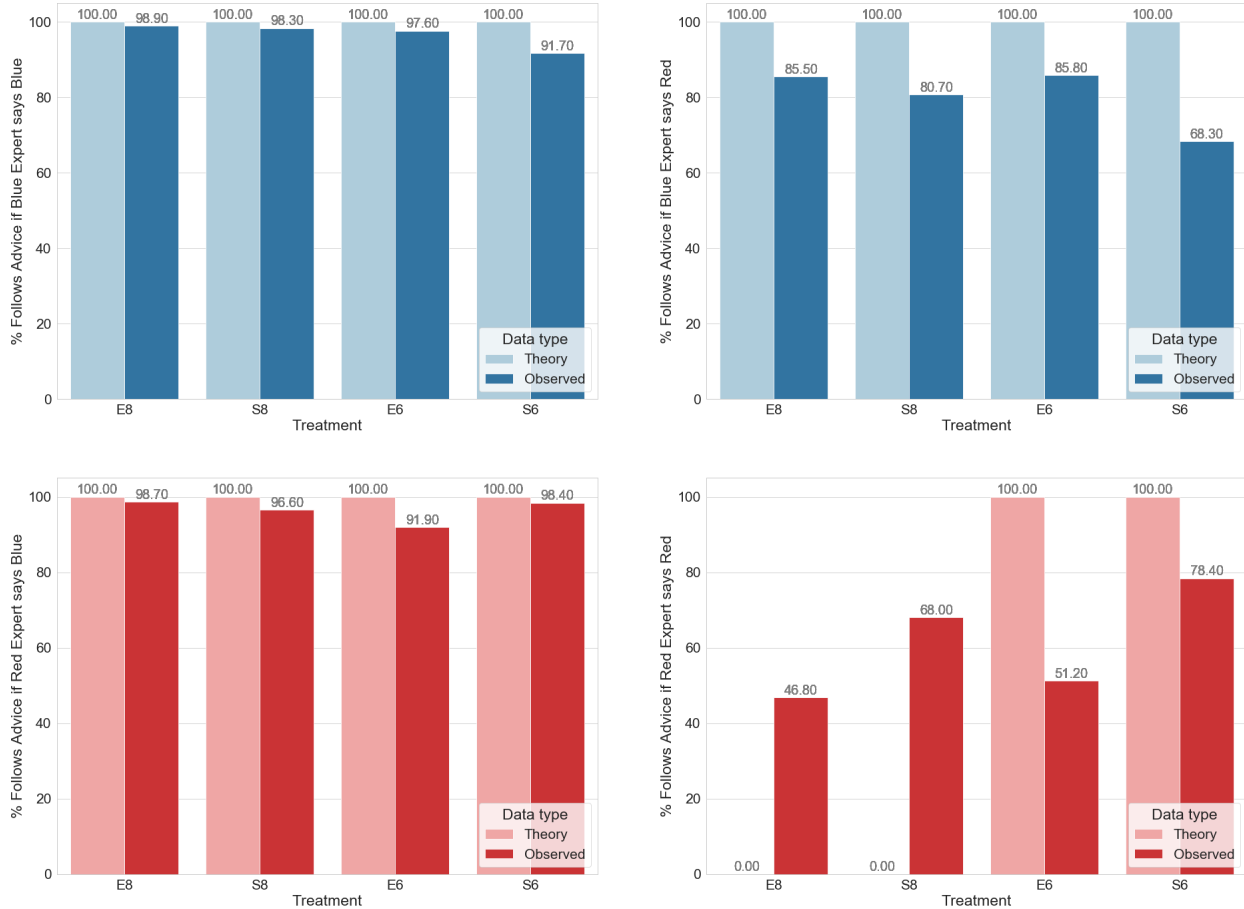
Figure 3: % Following Advice by Treatment and Information Set: Theory vs. Observed

with standard errors clustered at the subject level (p-value = 0.092). A Bayesian learner follows contradictory advice from the Red Expert only when messages from this expert are sufficiently informative. Given our experimental parameters, this is the case only with mildly unbalanced priors. In the experiment, subjects follow a red message by the Red Expert 69.9% of the time when the prior is mildly unbalanced (treatments E6 and S6), and 61.5% of the time when the prior is strongly unbalanced (treatments E8 and S8). The difference between the two pairs of treatments is not statistically significant according to a Probit regressions with standard errors clustered at the subject level (p-value = 0.247). Keeping the prior belief constant, subjects are more likely to follow contradictory advice by the Red Expert when this information source is more reliable: this happens 51.1% of the time in treatment E6 against 78.4% of the time in treatment S6 (p-value = 0.004); and 46.8% of the time in

treatment E8 against 68% of the time in treatment S8 (p-value = 0.070). Theoretically, even if increasing the Red Expert's reliability does increase the sensitivity of posterior beliefs to its advice, contradictory advice by the Red Expert should affect the optimal guess only when the initial belief is not too strong (and for both levels of reliability).

Finding 5. *Subjects follow contradictory advice sub-optimally: they are excessively skeptic of contradictory advice by the expert biased towards the prior and excessively trusting of contradictory advice by the expert biased against the prior.*

To understand why subjects' decision-making after receiving a contradictory signal is different from the Bayesian benchmark, we map the (non-incentivized) confidence levels in the guess to posterior beliefs on the state of the world and analyze them. We follow Charness et al. (2018) and define *responsiveness to information* as follows:

$$\alpha_s = \frac{p_s - p_0}{p_s^{Bay} - p_0}$$

where p_s is the observed posterior belief, p_0 is the prior belief and p^{Bay} is the posterior belief held by a Bayesian learner with the same information. Note that $\alpha_s = 1$ corresponds to Bayesian updating, $\alpha_s < 1$ corresponds to under-responsiveness and $\alpha_s > 1$ corresponds to over-responsiveness. We calculate α_s for each decision and show the average by treatment and information set in Table 4.

Table 4 shows that subjects' posterior beliefs are statistically indistinguishable from those of Bayesian learners when advice is in line with the source bias and the prior is more favorable to this bias: average responsiveness is not statistically different from 1 when the prior is strongly unbalanced and the Blue Expert suggests blue; and when the prior is mildly unbalanced and the Red Expert suggests red. At the same time, subjects are too trusting of advice in line with an expert's bias when the prior is less favorable to this bias: subjects are excessively responsive to a blue message by the Blue Expert when the prior is mildly unbalanced (average responsiveness being 1.5 in E6 and 2.5 in S6); and excessively respon-

Panel A: Treatment E8 (Equal Reliability, Strongly Unbalanced Prior)

	N	Observed	Theory	p-value
Mean Responsiveness if B Says b	186	0.8	1	0.252
Mean Responsiveness if B Says r	186	0.8	1	0.000
Mean Responsiveness if R Says b	79	0.7	1	0.006
Mean Responsiveness if R Says r	79	2.2	1	0.011

Panel B: Treatment S8 (Skewed Reliability, Strongly Unbalanced Prior)

	N	Observed	Theory	p-value
Mean Responsiveness if B Says b	57	1.3	1	0.408
Mean Responsiveness if B Says r	57	0.7	1	0.012
Mean Responsiveness if R Says b	178	0.6	1	0.000
Mean Responsiveness if R Says r	178	1.7	1	0.000

Panel C: Treatment E6 (Equal Reliability, Midly Unbalanced Prior)

Equal Reliability, Prior = 0.6 (E6)	N	Observed	Theory	p-value
Mean Responsiveness if B Says b	169	1.5	1	0.001
Mean Responsiveness if B Says r	169	0.8	1	0.001
Mean Responsiveness if R Says b	86	0.8	1	0.006
Mean Responsiveness if R Says r	86	0.8	1	0.409

Panel D: Treatment S6 (Skewed Reliability, Mildly Unbalanced Prior)

	N	Observed	Theory	p-value
Mean Responsiveness if B Says b	60	2.5	1	0.002
Mean Responsiveness if B Says r	60	0.5	1	0.001
Mean Responsiveness if R Says b	190	0.9	1	0.001
Mean Responsiveness if R Says r	190	1.0	1	0.712

Table 4: Belief Updating (from Confidence Statements) by Treatment. Notes: the unit of observation is a decision made by a subject in a round; p-values for comparison with theory are based on one-sample t-tests with standard errors clustered at the subject level.

sive to a red message by the Red Expert when the prior is strongly unbalanced (average responsiveness being 2.2 in E8 and 1.7 in S8). Moreover, subjects are always too skeptic of advice in conflict with an expert's bias (which, in fact, perfectly reveals the state of the world): the average responsiveness in these cases ranges from 0.5 (in treatment S6 when the Blue Experts says red) to 0.9 (in treatment S6 when the Red Expert says blue) and is statistically different from 1 for all treatments and information sets. Finding 6 summarizes this discussion.

Finding 6. *Subjects are insufficiently responsive to information misaligned with a source bias and excessively responsive to information aligned with a source bias.*

5 Conclusions

This paper formalized a model of selective exposure based on Bayesian updating, and tested its predictions through an online experiment. We ask two research questions: when is it rational to seek (dis)confirmatory information? Do experimental subjects behave according to rationality or do we need to impose additional structures? We modeled the problem of selective exposure to information as a choice between experts with different reliability. 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 source when experts are equally reliable; at the same time, these trends switch to dis-confirmatory attitudes as soon as the expert biased against the prior becomes more informative, 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 experts, or must pay a (possibly heterogeneous) price to receive messages from an information structure.

References

Ambuehl, Sandro and Shengwu Li, "Belief Updating and the Demand for Information," *Games and Economic Behavior*, 2018, *109*, 21–39.

Anderson, Lisa R and Charles A Holt, "Information Cascades in the Laboratory," *American Economic Review*, 1997, *87* (5), 847–862.

Bakshy, Eytan, Solomon Messing, and Lada A Adamic, "Exposure to ideologically diverse news and opinion on Facebook," *Science*, 2015, *348* (6239), 1130–1132.

Baron, Jonathan, Jane Beattie, and John C Hershey, "Heuristics and Biases in Diagnostic Reasoning II: Congruence, Information, and Certainty," *Organizational Behavior and Human Decision Processes*, 1988, *42* (1), 88–110.

Berelson, Bernard and Gary A Steiner, *Human Behavior: An Inventory of Scientific Findings*, Springer, 1968.

Charness, Gary and Chetan Dave, "Confirmation Bias with Motivated Beliefs," *Games and Economic Behavior*, 2017, *104*, 1–23.

_ **and Dan Levin**, "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect," *American Economic Review*, 2005, *95* (4), 1300–1309.

_ **, Ryan Oprea, and Sevgi Yuksel**, "How Do People Choose between Biased Information Sources? Evidence from a Laboratory Experiment," *Unpublished Manuscript*, 2018.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTree: An Open-Source Platform for Laboratory, Online, and Field Experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Dohmen, Thomas J, Armin Falk, David Huffman, Felix Marklein, and Uwe Sunde, "The Non-Use of Bayes Rule: Representative Evidence on Bounded Rationality," *Unpublished Manuscript*, 2009.

Duffy, John, Ed Hopkins, and Tatiana Kornienko, "Lone Wolf or Herd Animal? Information Choice and Social Learning," *Unpublished Manuscript*, 2017.

_ **, _ , _ , and Mingye Ma**, "Information Choice in a Social Learning Experiment," *Unpublished Manuscipt*, 2018.

Epley, Nicholas and Thomas Gilovich, "The Mechanics of Motivated Reasoning," *The Journal of Economic Perspectives*, 2016, *30* (3), 133–140.

Falk, Armin and Florian Zimmermann, "Beliefs and Utility: Experimental Evidence on Preferences for Information," *Unpublished Manuscript*, 2017.

Flaxman, Seth, Sharad Goel, and Justin M Rao, "Filter Bubbles, Echo Chambers, and Online News Consumption," *Public Opinion Quarterly*, 2016, *80* (S1), 298–320.

Frey, Dieter, "Recent Research on Selective Exposure to Information," *Advances in Experimental Social Psychology*, 1986, *19*, 41–80.

Gentzkow, Matthew and Jesse M Shapiro, "Media Bias and Reputation," *Journal of Political Economy*, 2006, *114* (2), 280–316.

_ **and** _ , "Ideological Segregation Online and Offline," *Quarterly Journal of Economics*, 2011, *126* (4), 1799–1839.

Grether, David M, "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 1980, *95* (3), 537–557.

Gunther, Albert C, "Biased press or biased public? Attitudes toward media coverage of social groups," *Public Opinion Quarterly*, 1992, *56* (2), 147–167.

Hoffrage, Ulrich, Samuel Lindsey, Ralph Hertwig, and Gerd Gigerenzer, "Communicating Statistical Information," *Science*, 2000, *290* (5500), 2261–2262.

Iyengar, Shanto and Kyu S Hahn, "Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use," *Journal of Communication*, 2009, *59* (1), 19–39.

Klayman, Joshua, "Varieties of Confirmation Bias," *Psychology of Learning and Motivation*, 1995, *32*, 385–418.

_ **and Young-Won Ha**, "Confirmation, Disconfirmation, and Information in Hypothesis Testing.," *Psychological Review*, 1987, *94* (2), 211.

Lawrence, Eric, John Sides, and Henry Farrell, "Self-Segregation or Deliberation? Blog Readership, Participation, and Polarization in American Politics," *Perspectives on Politics*, 2010, *8* (1), 141–157.

Mann, Thomas E. and Norman J. Ornstein, *It's Even Worse Than It Looks: How the American Constitutional System Collided with the New Politics of Extremism*, Basic Books, New York, NY, 2012.

Masatlioglu, Yusufcan, A Yesim Orhun, and Collin Raymond, "Intrinsic Information Preferences and Skewness," *Unpublished Manuscript*, 2017.

Mullainathan, Sendhil and Andrei Shleifer, "The Market for News," *American Economic Review*, 2005, *95* (4), 1031–1053.

Nickerson, Raymond S, "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises," *Review of General Psychology*, 1998, *2* (2), 175.

Nielsen, Kirby, "Preferences for the Resolution of Uncertainty and the Timing of Information," *Unpublished Manuscript*, 2018.

Peterson, Erik, Sharad Goel, and Shanto Iyengar, "Echo Chambers and Partisan Polarization: Evidence from the 2016 Presidential Campaign," *Unpublished Manuscript*, 2018.

Quattrociocchi, Walter, Antonio Scala, and Cass R Sunstein, "Echo Chambers on Facebook," *Discussion Paper No. 877, Harvard Law School*, 2016.

Sears, David O and Jonathan L Freedman, "Selective Exposure to Information: A critical Review," *Public Opinion Quarterly*, 1967, *31* (2), 194–213.

Skov, Richard B and Steven J Sherman, "Information-Gathering Processes: Diagnosticity, Hypothesis-Confirmatory Strategies, and Perceived Hypothesis Confirmation," *Journal of Experimental Social Psychology*, 1986, *22* (2), 93–121.

Slowiaczek, Louisa M, Joshua Klayman, Steven J Sherman, and Richard B Skov, "Information Selection and Use in Hypothesis Testing: What is a Good Question, and What is a Good Answer?," *Memory & Cognition*, 1992, *20* (4), 392–405.

Snowberg, Erik and Leeat Yariv, "Testing the Waters: Behavior across Participant Pools," *Unpublished Manuscript*, 2018.

Tappin, Ben M, Leslie van der Leer, and Ryan T McKay, "The Heart Trumps the Head: Desirability Bias in Political Belief Revision," *Journal of Experimental Psychology: General*, 2017, *146* (8), 1143.

Tversky, Amos and Daniel Kahneman, "Belief in the Law of Small Numbers.," *Psychological Bulletin*, 1971, *76* (2), 105.

Vicario, Michela Del, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi, "The Spreading of Misinformation Online," *Proceedings of the National Academy of Sciences*, 2016, *113* (3), 554–559.

Viscusi, W Kip and Charles J O'Connor, "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?," *American Economic Review*, 1984, *74* (5), 942–956.

Zimmermann, Florian, "Clumped or Piecewise? Evidence on Preferences for Information," *Management Science*, 2014, *61* (4), 740–753.

Appendix A: Proofs

Proof of Lemma 1 Assume the DM observes an r signal from **Blue**. The posterior belief $Pr(R|r, \textbf{Blue})$ is computed via Bayes Rule:

$$Pr(R|r, \textbf{Blue}) = \frac{Pr(R) \cdot Pr(r, \textbf{Blue}|R)}{Pr(R) \cdot Pr(r, \textbf{Blue}) + Pr(B) \cdot Pr(r, \textbf{Blue})} = \frac{(1-\pi)(1-\lambda_B)}{(1-\pi)(1-\lambda_B)} = 1$$

hence the signal is fully revealing and implies the highest possible expected payoff. It follows $a^*(r, \textbf{Blue}) = R$. On the other hand, after observing b from **Blue**, action B yields a higher payoff than R provided that $Pr(B|b, \textbf{Blue}) > Pr(R|b, \textbf{Blue})$, that is provided

$$\frac{\pi}{\lambda_B + \pi(1-\lambda_B)} > \frac{\lambda_B(1-\pi)}{\lambda_B + \pi(1-\lambda_B)}$$

which can be restated as $\lambda_B < \frac{\pi}{1-\pi}$ and the inequality always holds for $\pi > 0.5$.

Proof of Lemma 2 When the DM observes a b signal from **Red**, it updates $Pr(B|b, \textbf{Red}) = \frac{(1-\lambda_R)\pi}{(1-\lambda_R)\pi} = 1$. It immediately follows $a^*(b, \textbf{Red}) = B$. When the DM observes r from **Red**, it is optimal to guess R whenever

$$Pr(R|r, \textbf{Red}) > Pr(B|r, \textbf{Red}) \iff \frac{1-\pi}{1 - (1-\lambda_R)\pi} > \frac{\pi\lambda_R}{1 - (1-\lambda_R)\pi}$$

which can be re-arranged as $\lambda_R < \frac{1-\pi}{\pi}$, as we wanted to show.

Proof of Lemma 3 The expected payoff from accessing **Blue** is computed as:

$$\begin{aligned}
\mathbb{E}[\textbf{Blue}] &= [\pi + (1-\pi)\lambda_B] \frac{\pi}{\lambda_B + \pi(1-\lambda_B)} + (1-\pi)(1-\lambda_B) \\
&= \frac{\pi^2 + (1-\pi)\pi\lambda_B + (1-\pi)(1-\lambda_B)[\lambda_B + \pi(1-\lambda_B)]}{\lambda_B + \pi(1-\lambda_B)} \\
&= \frac{\pi^2 + (1-\pi)[\pi\lambda_B + (1-\lambda_B)\lambda_B + \pi(1-\lambda_B)^2]}{\lambda_B + \pi(1-\lambda_B)} \\
&= \frac{\pi^2 + (1-\pi)\pi\lambda_B + (1-\pi)(1-\lambda_B)[\lambda_B + \pi(1-\lambda_B)]}{\lambda_B + \pi(1-\lambda_B)} \\
&= \frac{\pi[\lambda_B + \pi(1-\lambda_B)] + (1-\pi)(1-\lambda_B)[\lambda_B + \pi(1-\lambda_B)]}{\lambda_B + \pi(1-\lambda_B)} \\
&= \pi + (1-\pi)(1-\lambda_B) \\
&= 1 - (1-\pi)\lambda_B \in [\pi, 1]
\end{aligned}$$

and is clearly increasing in π and decreasing in λ_B.

Proof of Lemma 4 When $\lambda_R \geq \frac{1-\pi}{\pi}$, signals from **Red** are uninformative and therefore

ignored. It follows:

$$\mathbb{E}[\mathbf{Red}] = \pi(1 - \lambda_R) + [(1 - \pi) + \pi\lambda_R] \left[\frac{\pi\lambda_R}{1 - (1 - \lambda_R)\pi} \right]$$

$$= \mathbb{E}_{Pr(\theta|b,\mathbf{Red})}[u(\theta|b)] + \mathbb{E}_{Pr(\theta|r,\mathbf{Red})}[u(\theta|b)]$$

$$= \mathbb{E}_{Pr(\theta|s,\mathbf{Red})}[\mathbb{E}[u(\theta|b)|s, \mathbf{Red}]]$$

$$= \mathbb{E}_\pi[u(\theta|b)] \text{ by LIE}$$

$$= \pi$$

On the other hand, when $\lambda_R < \frac{1-\pi}{\pi}$ the DM follows contradictory signals from **Red**. The expected payoff becomes:

$$\mathbb{E}[\mathbf{Red}] = \pi(1 - \lambda_R) + [(1 - \pi) + \pi\lambda_R] \left(\frac{1 - \pi}{1 - (1 - \lambda_R)\pi} \right)$$

$$= \frac{\pi(1 - \lambda_R) - \pi^2(1 - \lambda_R)^2 + (1 - \pi)^2 + \pi(1 - \pi)\lambda_R}{1 - (1 - \lambda_R)\pi}$$

$$= \frac{\pi(1 - \lambda_R)[1 - \pi(1 - \lambda_R)] + (1 - \pi)[1 - \pi + \pi\lambda_R]}{1 - (1 - \lambda_R)\pi}$$

$$= \pi(1 - \lambda_R) + (1 - \pi)$$

$$= 1 - \pi\lambda_R \in [1 - \pi, 1]$$

which is decreasing in π and in λ_R.

Proof of Proposition 1 We distinguish two cases, namely $\lambda_R < \frac{1-\pi}{\pi}$ and $\lambda_R \geq \frac{1-\pi}{\pi}$.

First consider the case where $\lambda_R \geq \frac{1-\pi}{\pi}$, i.e. signals from **Red** are uninformative. Using the previous results, the DM prefers **Blue** over **Red** as long as

$$\mathbb{E}[\mathbf{Blue}] \geq \mathbb{E}[\mathbf{Red}] \iff 1 - (1 - \pi)\lambda_B \geq \pi \iff 1 - \pi \geq (1 - \pi)\lambda_B \iff 1 \geq \lambda_B$$

which always holds by construction. Hence if $\lambda_R \geq \frac{1-\pi}{\pi}$, the DM chooses to access source **Blue**.

Consider next the case $\lambda_R < \frac{1-\pi}{\pi}$. In this range signals from **Red** are informative and always followed. The DM prefers **Blue** over **Red** whenever

$$\mathbb{E}[\mathbf{Blue}] \geq \mathbb{E}[\mathbf{Red}] \iff 1 - (1 - \pi)\lambda_B \geq 1 - \pi\lambda_R \iff \lambda_R \geq \frac{1 - \pi}{\pi}\lambda_B$$

In particular, the DM accesses source **Red** when $\lambda_R < \frac{1-\pi}{\pi}\lambda_B$ and Blue otherwise. Putting together the two cases, the result in the proposition follows.

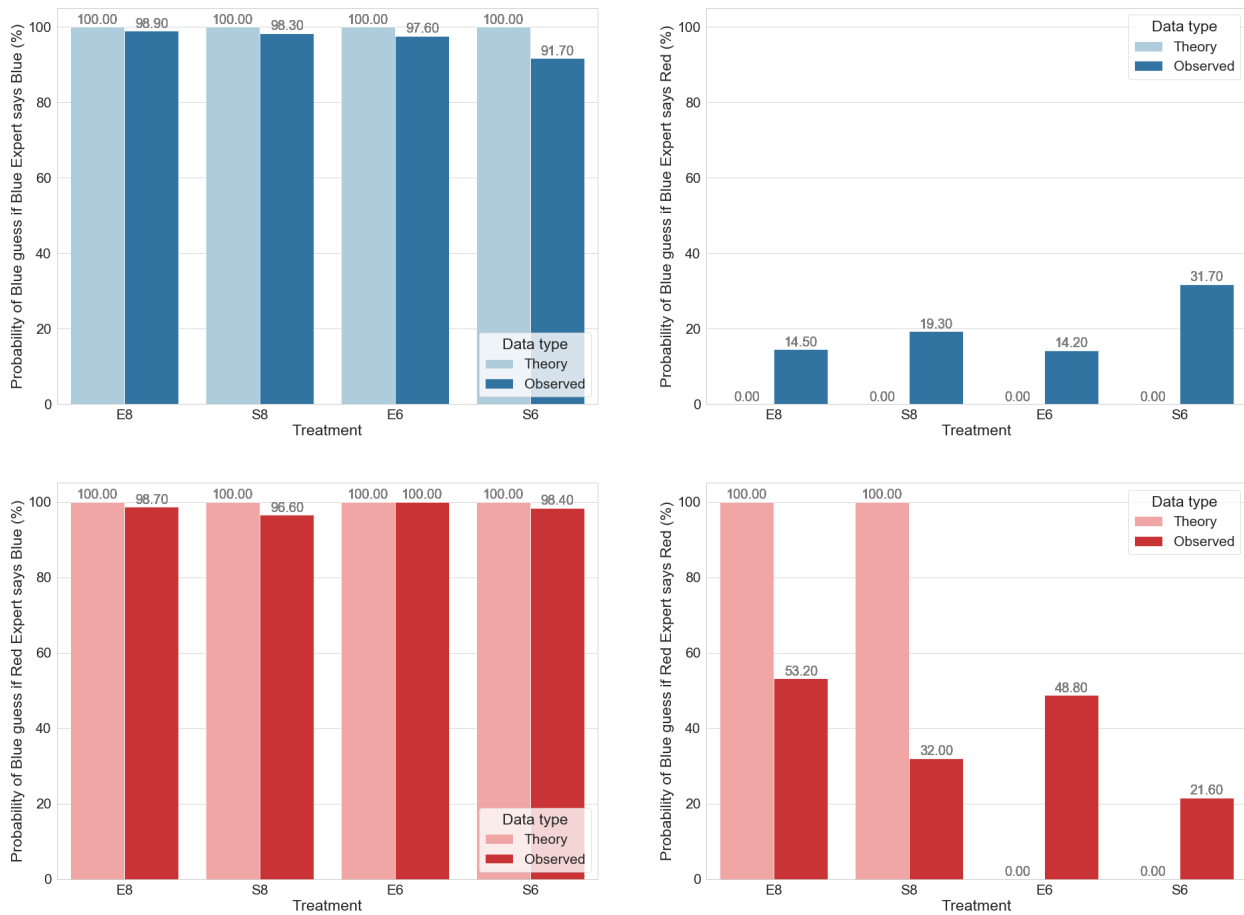# Appendix B: Additional Figures and Tables

Figure 4: Guess on the State by Treatment and Information Set: Theory vs. Observed Data

Figure 5: Posterior Beliefs by Treatment and Information Set: Theory vs. Observed Averages

E8	Round 1	Round 2	Round 3	Round 4	Round 5
% Chooses Blue Expert	69.8	66.0	66.0	73.6	75.5
% Guesses Blue Ball if B Says b	97.3	100.0	97.1	100.0	100.0
% Guesses Blue Ball if B Says r	10.8	8.6	20.0	17.9	15.0
% Guesses Blue Ball if R Says b	100.0	94.4	100.0	100.0	100.0
% Guesses Blue Ball if R Says r	43.8	44.4	50.0	57.1	76.9
Mean Posterior if B Says b	87.1	88.7	85.5	87.4	86.5
Mean Posterior if B Says r	12.4	11.6	18.5	19.1	18.9
Mean Posterior if R Says b	92.3	90.7	97.2	96.8	93.1
Mean Posterior if R Says r	43.2	48.1	49.3	53.9	63.5
S8	Round 1	Round 2	Round 3	Round 4	Round 5
% Chooses Blue Expert	25.5	25.5	17.0	25.5	27.7
% Guesses Blue Ball if B Says b	100.0	100.0	100.0	100.0	92.3
% Guesses Blue Ball if B Says r	8.3	25.0	12.5	16.7	30.8
% Guesses Blue Ball if R Says b	94.3	100.0	97.4	97.1	94.1
% Guesses Blue Ball if R Says r	25.7	31.4	33.3	31.4	38.2
Mean Posterior if B Says b	90.0	84.0	89.7	88.8	83.1
Mean Posterior if B Says r	16.9	24.2	9.4	20.4	30.5
Mean Posterior if R Says b	90.3	94.2	92.2	92.0	88.5
Mean Posterior if R Says r	31.7	34.0	36.1	37.5	42.5
E6	Round 1	Round 2	Round 3	Round 4	Round 5
% Chooses Blue Expert	72.5	58.8	62.7	66.7	70.6
% Guesses Blue Ball if B Says b	100.0	100.0	96.9	94.1	97.2
% Guesses Blue Ball if B Says r	13.5	16.7	9.4	14.7	16.7
% Guesses Blue Ball if R Says b	100.0	90.5	84.2	88.2	100.0
% Guesses Blue Ball if R Says r	50.0	47.6	42.1	47.1	60.0
Mean Posterior if B Says b	84.4	83.5	81.6	79.3	80.6
Mean Posterior if B Says r	14.1	15.5	10.2	13.8	17.9
Mean Posterior if R Says b	98.2	89.4	81.3	86.8	98.0
Mean Posterior if R Says r	48.2	46.8	41.9	46.3	51.5
S6	Round 1	Round 2	Round 3	Round 4	Round 5
% Chooses Blue Expert	30.0	16.0	30.0	22.0	22.0
% Guesses Blue Ball if B Says b	86.7	100.0	86.7	90.9	100.0
% Guesses Blue Ball if B Says r	33.3	25.0	13.3	45.5	45.5
% Guesses Blue Ball if R Says b	100.0	92.9	100.0	100.0	100.0
% Guesses Blue Ball if R Says r	17.1	28.6	22.9	23.1	15.4
Mean Posterior if B Says b	76.9	91.6	78.8	75.1	83.8
Mean Posterior if B Says r	30.7	27.5	18.7	41.4	38.9
Mean Posterior if R Says b	96.0	90.8	95.4	96.4	96.2
Mean Posterior if R Says r	29.1	32.0	28.9	29.7	29.2

Table 5: Observed Outcomes by Treatment and Round, All Subjects

31

E8	1+ Correct Answers (N = 265)	2+ Correct Answers (N = 245)	3 Correct Answers (N = 145)
% Chooses Blue Expert	70.7	70.2	73.8
% Guesses Blue Ball if B Says b	97.0	98.8	100.0
% Guesses Blue Ball if B Says r	0.0	15.1	11.2
% Guesses Blue Ball if R Says b	95.1	98.6	97.4
% Guesses Blue Ball if R Says r	48.8	54.8	65.8
Mean Posterior if B Says b	82.6	87.1	87.0
Mean Posterior if B Says r	0.9	14.6	12.6
Mean Posterior if R Says b	94.5	95.1	93.2
Mean Posterior if R Says r	49.0	51.7	60.9

S8	1+ Correct Answers (N = 235)	2+ Correct Answers (N = 205)	3 Correct Answers (N = 75)
% Chooses Blue Expert	24.3	22.4	18.7
% Guesses Blue Ball if B Says b	98.2	97.8	100.0
% Guesses Blue Ball if B Says r	19.3	15.2	7.1
% Guesses Blue Ball if R Says b	96.6	96.9	98.4
% Guesses Blue Ball if R Says r	32.0	34.6	31.1
Mean Posterior if B Says b	86.9	86.9	83.0
Mean Posterior if B Says r	21.2	19.1	9.6
Mean Posterior if R Says b	91.5	91.7	90.7
Mean Posterior if R Says r	36.3	37.7	35.0

E6	1+ Correct Answers (N = 255)	2+ Correct Answers (N = 210)	3 Correct Answers (N = 140)
% Chooses Blue Expert	66.3	65.2	70.7
% Guesses Blue Ball if B Says b	97.6	97.8	97.0
% Guesses Blue Ball if B Says r	14.2	6.6	0.0
% Guesses Blue Ball if R Says b	91.9	93.2	95.1
% Guesses Blue Ball if R Says r	48.8	52.1	48.8
Mean Posterior if B Says b	81.9	82.6	82.6
Mean Posterior if B Says r	14.3	7.6	0.9
Mean Posterior if R Says b	90.0	92.6	94.5
Mean Posterior if R Says r	46.6	48.3	49.0

S6	1+ Correct Answers (N = 240)	2+ Correct Answers (N = 205)	3 Correct Answers (N = 110)
% Chooses Blue Expert	22.1	20.0	11.8
% Guesses Blue Ball if B Says b	90.6	90.2	84.6
% Guesses Blue Ball if B Says r	28.3	34.1	7.7
% Guesses Blue Ball if R Says b	98.4	99.4	100.0
% Guesses Blue Ball if R Says r	21.4	20.1	14.4
Mean Posterior if B Says b	80.0	82.3	74.2
Mean Posterior if B Says r	28.1	29.3	5.8
Mean Posterior if R Says b	95.1	97.2	98.7
Mean Posterior if R Says r	29.8	29.0	26.9

Table 6: Observed Outcomes by Performance in Comprehension Quiz, All Rounds

Equal Reliability, Prior = 0.8 (E8)	N	Observed	Theory
Mean Posterior if B Says b	186	87.1	88.9
Mean Posterior if B Says r	186	16.2	0
Mean Posterior if R Says b	79	94.0	100
Mean Posterior if R Says r	79	50.9	66.7

Skewed Reliability, Prior = 0.8 (S8)	N	Observed	Theory
Mean Posterior if B Says b	57	86.9	85.1
Mean Posterior if B Says r	57	21.2	0
Mean Posterior if R Says b	178	91.5	100
Mean Posterior if R Says r	178	36.3	54.5

Equal Reliability, Prior = 0.6 (E6)	N	Observed	Theory
Mean Posterior if B Says b	169	81.9	75.0
Mean Posterior if B Says r	169	14.3	0
Mean Posterior if R Says b	86	90.0	100
Mean Posterior if R Says r	86	46.6	42.9

Skewed Reliability, Prior = 0.6 (S6)	N	Observed	Theory
Mean Posterior if B Says b	60	80.3	68.2
Mean Posterior if B Says r	60	30.7	0
Mean Posterior if R Says b	190	94.9	100
Mean Posterior if R Says r	190	29.9	31.0

Table 7: Posterior Beliefs (from Confidence Statements) by Treatment.

Appendix C: Sample Instructions (Treatment E8)

Experimental instructions were delivered in the initial screens of the experiment. We report here the complete text and figures of these screens, including the comprehension quiz and the practice round. Page titles, as they appeared on the participants' screen, are in bold.

WELCOME
Welcome!
Thank you for agreeing to participate in this experiment!
This is an experiment designed to study how people make decisions.
The whole experiment will last around 10 minutes.
In addition to your participation fee, you will be able to earn a bonus payment.
Your bonus payment will depend on your choices so, please, read the instructions carefully.
We will use only one decision to determine your bonus payment but all decisions are equally likely to be selected so all choices matter. The instructions describe how your choices affects your earnings. They are composed of three pages and include a comprehension question at the end of each page.
Please, devote at least 5 minutes to the instructions and the comprehension questions.
Once you start the experiment, we require your complete and undistracted attention.
When you are ready to start, please click the button below:

NEXT

INSTRUCTIONS/1: YOUR TASK
In each round, there will be a jar, like the one you see below, containing 8 **BLUE** balls and 2 **RED** balls.

The computer will randomly draw ONE ball out of this jar. All balls are equally likely to be drawn.
In each round, your task will be to guess whether the ball drawn by the computer is **BLUE** or **RED**.
Before proceeding to the next page, please answer the comprehension question below:
Without any additional information, what do you know about the ball drawn by the computer?

- It is more likely that it is **BLUE**

- It is more likely that it is **RED**

- It is just as likely that it is **BLUE** as that it is **RED**

Please spend at least 30 seconds on this page. Read the instructions carefully! :-)

NEXT

FEEDBACK/1
Correct!
The urn contains 10 balls in total: 8 **BLUE** balls and 2 **RED** balls.
The computer draws one ball completely at random: each of the 10 balls is equally likely to be drawn.
This means that there are 8 chances out of 10 that the computer draws a **BLUE** ball and 2 chances out of 10 that the computer draws a **RED** ball.
Thus, without any additional information, you know that the ball is more likely to be **BLUE**.

NEXT

INSTRUCTIONS/2: GETTING ADVICE
Before you make your assessment, you can consult an expert.
The expert you consult might be informed about the ball drawn by the computer. If he knows the color, he will report it to you. If he does not know the color, he will simply report to you his preferred color.
There are 10 **BLUE** experts and 10 **RED** experts. You choose whether you want to hear from a BLUE expert or a RED expert. If you choose a BLUE expert, the computer randomly picks one BLUE expert to advise you. If you choose to hear from a RED expert, the computer randomly picks one RED expert.

If you get advice from a **BLUE** expert:

- 5 out of 10 **BLUE** experts are informed about the ball
- If the ball is BLUE:
    - An informed BLUE expert says "The ball is BLUE"?
    - An uninformed BLUE expert says "The ball is BLUE"?
- If the ball is RED:
    - An informed BLUE expert says "The ball is RED"?
    - An uninformed BLUE expert says "The ball is BLUE"?

35

If you get advice from a **RED** expert:

- 5 out of 10 **RED** experts are informed about the ball

- If the ball is BLUE:

  - An informed RED expert says "The ball is BLUE"?
  - An uninformed RED expert says "The ball is RED"?

  If the ball is RED:

  - An informed RED expert says "The ball is RED"?
  - An uninformed RED expert says "The ball is RED"?

  Before proceeding to the next page, please answer the comprehension question below:

If a **BLUE** expert says "The ball is **RED**", which of the following is true?

- You know for sure that the ball is **BLUE**

- You know for sure that the ball is **RED**

- The ball is more likely to be **RED** but you ot know this for sure.

- The ball is more likely to be **BLUE** but you do not know this for sure.

NEXT

FEEDBACK/2

Correct!

A BLUE expert says "The ball is RED"? only if he is informed and the ball is, in fact, RED. In all other cases, he says "The ball is BLUE".

This means that, if you get advise from a BLUE expert, and he says "The ball is RED", then you know for sure that the ball is RED. Remember that not all BLUE experts are informed (only 5 out of 10).

Similarly, a RED expert says "The ball is BLUE" only if he is informed and the ball is, in fact, BLUE. In all other cases, he says "The ball is RED".

This means that, if you get advise from a RED expert, and he says "The ball is BLUE", then you know for sure that the ball is BLUE. Remember that not all RED experts are informed

(only 5 out of 10).

INSTRUCTIONS / 3: GUESS THE COLOR AND EARN MONEY!

After you choose what expert to consult, but before you are revealed his message, you will be asked to make your best guess about the color of the ball, depending on what you will hear from the expert.
Since you can receive two different messages, you will be asked two questions:

What is your guess about the color of the ball, if the expert says "The ball is **BLUE**"?
What is your guess about the color of the ball, if the expert says "The ball is **RED**"?

After you submit your answers, the computer will report you the expert's message and will use as your guess for this round the answer to the corresponding question. For example, if the expert you consulted says "The ball is **BLUE**", the computer will use as your guess the answer you gave to the first question above. If, instead, the expert says "The ball is **RED**", the computer will use as your guess the answer you gave to the second question above.

Your guess will determine your bonus payment in the following way:

You will earn $1 if your guess matches the true color of the ball.
You will earn $0 if your guess does not match the true color of the ball.

In addition, you will be asked how confident you are of each of your guesses, on a scale between 0 and 100. For example, 0 indicates that you think it is just as likely that you are right or wrong (that is, you think that it is just as likely that the ball is **BLUE** or **RED**), while 100 indicates that you are sure you picked the right color (that is, you think you know for sure whether the ball is **BLUE** or **RED**).

These assessments do not affect your bonus payment but it is very important to us that you make your choice carefully and that you report to us what you really believe.

Before proceeding to the next page, please answer the comprehension question below:

Consider this example. Your guesses are that the ball is BLUE if the expert says BLUE; and that the ball is RED if the expert says RED. The expert says "The ball is BLUE"? and the true color of the ball is BLUE. What is your bonus payment in this round?

- $1 because your guess is BLUE and it coincides with the actual color of the ball.

- $0.50 because only one of your two guesses coincides with the actual color of the ball.

- $0 because your guess is RED and it does not coincides with the actual color of the ball.

Please spend at least 60 seconds on this page. Read the instructions carefully! :-)

NEXT

FEEDBACK/3

Correct!

Only one guess matters for your bonus payment.
The guess that matters depends on the message you receive from the expert.
Since you do not know what message you will receive, make both guesses carefully.

If the expert says "The ball is **BLUE**", the guess that matters for your bonus payment is the answer to the question: What is your guess about the color of the ball, if the expert says "The ball is **BLUE**"?
If the expert says "The ball is **RED**", the guess that matters for your bonus payment is the answer to the question: What is your guess about the color of the ball, if the expert says "The ball is **RED**"?

In this example, the expert said BLUE; your guess, conditional on the expert saying BLUE, was BLUE and, thus, your guess for this round was: BLUE.
The ball randomly drawn by the computer was BLUE too. This means that your guess coincided with the ball drawn by the computer and, thus, you earned \$1. You earn \$0 if your guess does not match the color of the ball.

NEXT

GET READY FOR THE GAME!

You will play 5 rounds of this game.
The computer will randomly pick one round to determine your bonus payment but all rounds are equally likely to be selected so all choices matter.
In each round, there are a new jar with 10 balls, 10 new BLUE Experts, and 10 new RED Experts. The chance the computer draws a RED ball or a BLUE ball from the jar, as well as the chance that the expert you consult is informed or uninformed are not affected in any way by what happened in the previous rounds.

When you are ready to start with Round 1, please click the button below.
Please spend at least 30 seconds on this page. Read the instructions carefully! :-)

NEXT

PRACTICE ROUND - WHOSE ADVICE DO YOU WANT?

There is a jar containing 8 **BLUE** balls and 2 **RED** balls.
The computer has randomly drawn **ONE** ball out of this jar.

Your task is to guess whether the ball drawn by the computer is BLUE or RED.

Before you make your guess, you can get advice from a **BLUE** or a **RED** expert.
If you get advice from a **BLUE** expert:

- If the ball is BLUE:

  - An informed BLUE expert says "The ball is BLUE"?
  - An uninformed BLUE expert says "The ball is BLUE"?

- If the ball is RED:

  - An informed BLUE expert says "The ball is RED"?
  - An uninformed BLUE expert says "The ball is BLUE"?

If you get advice from a **RED** expert:

- If the ball is BLUE:

  - An informed RED expert says "The ball is BLUE"?
  - An uninformed RED expert says "The ball is RED"?

  If the ball is RED:

  - An informed RED expert says "The ball is RED"?
  - An uninformed RED expert says "The ball is RED"?

Remember that 5 out of 10 BLUE experts are informed and 5 out of 10 RED experts are informed.

Which expert do you want to hear from?

(a) Blue Expert   (b) Red Expert

NEXT

PRACTICE ROUND - GUESS THE COLOR!

You decided to consult a BLUE Expert.
What is your guess about the color of the ball, if the expert says "The ball is BLUE"?

BLUE

RED

On a scale from 0 to 100, how confident are you about this guess? For example, 0 means that you think it is just as likely that you are right or wrong and 100 means you are sure your guess is correct.

CONFIDENCE

What is your guess about the color of the ball, if the expert says "The ball is RED"?

BLUE

RED

On a scale from 0 to 100, how confident are you about this guess? For example, 0 means that you think it is just as likely that you are right or wrong and 100 means you are sure your guess is correct.

CONFIDENCE

NEXT

PRACTICE ROUND - RESULTS

You decided to consult a **BLUE Expert**.
This expert reported "The ball is **BLUE**".
Your guess, given the expert's report, was: **BLUE**.
The ball randomly drawn by the computer in this round was **BLUE**.
Your earnings in this round are $1.00.
When you are ready to start with the next round, please click the button below.

NEXT