Value of Superficial Information Characteristics*

Aniol Llorente-Saguer† Santiago Oliveros‡ Ro'i Zultan§

February 20, 2023

Abstract

We study non-instrumental value for information when choosing between symmetric and asymmetric information sources. Although symmetry is a seemingly superficial characteristic, we find a systematic preference for a symmetric source over an asymmetric one. Beliefs consistently overestimate the instrumental value of the symmetric source yet fail to explain the choice pattern. We provide novel evidence for systematic deviations from Bayesian belief updating. While around half of the subjects exhibit base-rate neglect as documented in the literature, we identify a distinct type that neglects new information more than the base rate. We show that these two belief formation processes best organize the choice data.

Keywords: information acquisition, representative heuristic, laboratory experiment

*We thank conference and seminar participants at the Economic Science Association online global meeting, Bar-Ilan University, Ben-Gurion University of the Negev, Causality in Cognition Lab in Stanford University, Tel Aviv University, University of Innsbruck, and University of Warwick for helpful comments and discussion.

†Department of Economics, Queen Mary University of London, Mile End Road, London, E1 4NS, United Kingdom. Email: a.llorente-saguer@qmul.ac.uk.

‡Department of Economics, University of Bristol, Priory Road Complex, Bristol, BS8 1TU, United Kingdom. Email: s.oliveros@bristol.ac.uk.

§Department of Economics, Ben-Gurion University of the Negev, Beer-Sheva, 8410501 Israel. Email: zultan@bgu.ac.il.

1 Introduction

Information derives its instrumental value from its strength to change a subject's mind. A pitted fruit farmer in the process of expanding production only values information conveying future heavy rainfall with high reliability which may make her postpone the expansion. On the other hand, a rice farmer facing the same decision seeks information that can convey with high reliability that the weather is unusually dry in the future. Similarly, recent retirees looking to rebalance their portfolios will keep the same stocks only in the presence of reliable information about steady earnings, while a young person facing a similar rebalancing problem looks for certainty about high earnings at some point in the future even at the expense of initial losses.

This differential need for a particular type of *certainty* is pervasive in economics: the FDA's process of drug approvals is heavily skewed towards avoiding secondary effects for novel drugs without market substitutes at the expense of effectiveness, and skewed towards effectiveness and improvements for drugs that have substitutes in the market; doctors facing potential cancer patients are more worried about failing to detect cancer than misdiagnose it, and request further testing to increase certainty about absence of cancer; voters concerned with budget issues are more prone to get information about the economic plans of candidates than those voters worried about purely ideological issues who are more prone to investigate the character of those candidates. More generally, useful information sometimes may be *biased* towards providing certainty in particular directions.

The certainty that information provides is related to the posteriors' position with respect to the agent's prior. Paradoxically, the most natural presentation of information is by considering the precision of the signals' realization; i.e. the probability that a particular realization correlates with the actual states. In a two state world $\omega \in \{A, B\}$ with a binary signal $S \in \{s_A, s_B\}$, what matters for the decision maker are the bayesian posterior of a state given the realized signal, $\Pr(\omega = A|s_A)$ and $\Pr(\omega = B|s_B)$, but information is presented naturally with the probability of a *signal realization given a state*, $p_A = \Pr(s_A|\omega = A)$ and $p_B = \Pr(s_B|\omega = B)$. Returning to our examples, how good is the current weather technology (ratio p_A/p_B) in truly identifying risks of draught $\Pr(\omega = A|s_A)$ or risk of flood $\Pr(\omega = B|s_B)$ is important for the farmer to decide whether information is useful or not.

There is strong evidence that subjects assessing the value of information fail to construct the right mapping between the priors and the signal to the posteriors; i.e. subjects usually fail to properly construct $\Pr(\omega = A|s_A)$ and $\Pr(\omega = B|s_B)$. Less is known about how the decision problems themselves induce the subjects to behave in non optimal ways. For example, consider a subject making a binary decision with uniform priors over the two states and having to use *symmetric information* $p_A = p_B > 0.5$. Clearly, an agent that mistakes p_A with $\Pr(\omega = A|s_A)$ and p_B with $\Pr(\omega = B|s_B)$ will appear as fully rational and only when asking to choose between different sources, or when facing different priors, it may be revealed that this agent does not necessarily process information properly.

In order to detect the subject's ability to process information the analyst must observe different choices and, in particular, may want to disrupt the equivalence $p_A \neq p_B$. In doing so, the problem changes fundamentally: from a source that appears balanced/symmetric the subject now faces sources that appear unbalanced/asymmetric. Are these superficial characteristics important for the way information is processed? Does these characteristics affect how subject perceives information's value? Do subjects have preferences over these characteristics? Do subjects use information optimally?

In this paper, we focus in understanding the role of a seemingly irrelevant (superficial) characteristic in the demand for information sources: symmetry. In order to isolate the effect of symmetry, our experimental design is minimalistic and presents information in the most intuitive statistical way. We design an experiment in a binary state–binary action–binary signal world, and study whether one source of information providing the same precision in both states (symmetric) is easier to process, or has higher subjective value, than an information source that provides different precision depending on the state (asymmetric).[1]

Subjects face a decision problem and need to collect information from one of the two sources (one symmetric and one asymmetric) before choosing a course of action. We vary their prior beliefs and the precision of the symmetric source, in order to change the value and the relative order of the information sources. We find that subjects in our experiment use information too much in comparison to expected utility maximizers performing proper bayesian updating: information appears to be more valuable than what it actually is. Moreover, symmetric sources appear to be more appealing than asymmetric ones as subjects in our experiment choose the former more often than optimal relatively to the latter.

For agents seeking to use information instrumentally, the need to compare asymmetric sources of information may make statistical inferences more challenging. For a researcher trying to understand demand for information, these potential mistakes present an important challenge. Indeed, what appears to be agents making erratic choices, or choices founded in non instrumental reasons, may actually be agents having problems possessing information. In turns, the researcher inability to distinguish mistakes may hide a perfectly reasonable and rational preference order over information sources.

To account for this challenge we investigate how agents process information. We elicit

[1]Note that this apparent difference is irrelevant as the value of information does depend on the prior belief the agent has and symmetry in precision does not mean symmetry in certainty.

(incentivized) subjective beliefs (posteriors and frequency of signals) so we can recover the behavior of the subjects as if they were purely instrumental bayesian agents. We find that subjective expected utility behavior or deviations form expected utility behavior due to processing mistakes (beliefs formation) cannot account for these apparent anomalies.

Nevertheless, the differences in which agents process information may help understand deviations from purely instrumental value.Using Grether (1980) flexible framework, we perform k-means cluster analysis based on the recovered subjective beliefs. We find systematic deviations from Bayesian updating that can be categorized in two different groups: those that (heavily) overweight information's characteristics relatively to priors and those that (mildly) overweight priors relatively to information's characteristics. Importantly, while both groups neglect priors when forming beliefs, the first group appears to completely ignore them. Following the literature we refer to the first group as *Base rate neglect* (BRN) and to the second group as *Information neglect* (IN).[2]

We then use this categorization to study alternative ways of using information beyond utility maximization. In particular, we find that we can explain heuristically why subjects use information too much and why symmetry is appealing. To fix ideas, consider a subject that completely ignores priors and responds excessively to information characteristics (BRN). This subject's posterior belief about state ω is equivalent to the probability that the information source reports the signal that is more likely in state ω; i.e. $\Pr(\omega = A|s_A) = p_A$ and $\Pr(\omega = B|s_B) = p_B$. For the symmetric source we have $p_A = p_B > 0.5$ so the symmetric source always has value (no matter the prior) and subjects always follow the signal. In other words, for every prior, base rate neglect subjects use information excessively and prefer symmetry over asymmetry inefficiently, two predictions that we confirm in the data.

The rest of the paper is organized at follows. After discussing the relevant literature we present the theoretical model and some basic information results in Section 3 and the experiment and testable implications in Section 4. We discuss our results about choice, use of information, and beliefs formation in Section 5. We relegate a complete description of the experiment's instructions and the visual aid provided to subjects to the Appendices.

[2]Strictly speaking BRN subjects underweight priors and weight information almost appropriately while information neglect subjects underweight both information and priors. The formal discussion for this definition and characterization is relegated to the section 5.5.

2 Literature Review

There is ample evidence that subjects value information for non-instrumental reasons. For example, Jones and Sugden (2001) and Charness and Dave (2017) find evidence of confirmation bias;[3] Von Gaudecker, Van Soest, and Wengström (2011), Zimmermann (2014), and Falk and Zimmermann (2014) find that there is a time dimension to how agents value information; Eliaz and Schotter (2007) and Eliaz and Schotter (2010) find value for information that comes from confidence seekers subjects.[4] In our paper we also identify a characteristic that induces agents to appear to overvalue information. Unlike in the papers listed above, the characteristic we identify is superficial, irrelevant and is purely statistical. More importantly, we identify the reason why subjects may *like* this characteristic that is not rooted in alternative preferences.

Our setting is minimalistic and, in that sense, our paper directly relates to papers that present information in purely statistical terms and study demand for *biased* information. Masatlioglu, Orhun, and Raymond (2016) study subjects choosing between information sources that have no value. They document evidence that subjects value skewness. Charness, Oprea, and Yuksel (2021) study subjects choosing between asymmetric sources. They document patterns of demand for information for confirmatory, anti-confirmatory, and uncertainty reduction reasons. In their experiment, subjects' mistakes cannot be explained by calculation errors but reflect heuristics highlighting the afortmentioned motives. Montanari and Nunnari (2019) use a similar setup and confirm similar patterns regarding confirmatory and uncertainty reducing motives. They also find that subjects use information suboptimally when information goes against their prior and it comes from the source biased against their own priors.

Unlike these papers we focus on choices between symmetric (equal precision in both states) and asymmetric (different precision in each state) and we vary the prior to account for different value of information. At the same time, we also compare situations when both sources have no value (as in Masatlioglu, Orhun, and Raymond, 2016) and when only one has value (as in some cases in Charness, Oprea, and Yuksel, 2021; Montanari and Nunnari, 2019). We find that subjects rely too much on information and symmetry seems to be a desirable characteristic.

Our paper also relates to the literature that documents systematic deviations from

[3]Rabin and Schrag (1999) study confirmation bias but in understanding information and not seeking information.

[4]There is also evidence that subjects avoid information (Huck, Szech, and Wenner, 2015). Golman, Hagmann, and Loewenstein (2017) discuss theoretical models of information avoidance and empirical evidence of this type of behavior towards information.

bayesian updating.[5] In particular, we study the process of beliefs formation directly as in Grether (1980) and we also find that subjects tend to form posterior beliefs in a predictable but non-bayesian way: people place more weight on the signal relative to the prior (BRN). Our method allows us to identify a more nuance behavior: we find subjects that also neglect information as well as priors.[6]

The heuristics that follow from this belief formation process allow us to explain why in our experiment information appear to have more instrumental value and to explain certain behavior. In that sense, our paper also relates to Wolfe and Fisher (2013), who show that individual differences in BRN are correlated with performance in various tasks.[7]

3 The Model

An agent must guess an unobserved state of nature, ω, which can be either blue or red, i.e., $\omega \in \{B, R\}$. The red state, R, materializes with probability ρ, and the blue one, B, with probability $1 - \rho$. The agent's preferences are represented by the following von Neumann–Morgenstern utility function: $u(g|\omega) = \mathbb{1}_{g=\omega}$.

Before making a guess, the agent can collect information from one of two different sources of information: S or A. Each informational source $k \in \{S, A\}$ provides a binary signal $s \in \{b, r\}$ with distribution contingent on the state. The probability that source k produces signal b in state B is denoted by $p_k^b \equiv \Pr(b|B, k)$ and the probability that the same source produces the signal r in state R is analogously denoted by $p_k^r \equiv \Pr(r|R, k)$. In is useful to define the expected precision of source k for priors ρ

$$\mathbf{p}_k(\rho) = p_k^r \rho + p_k^b(1 - \rho)$$

We assume that under S, the probability that the signal is correct is constant across states $p_S \equiv p_S^b = p_S^r$ which implies that the expected precision is constant for all ρ: $p_S = \mathbf{p}_S(\rho)$. We refer to source S as a *symmetric source*. In contrast, under the source A, the likelihood of receiving the right signal varies across states. We refer to this source A as *asymmetric source*. In particular, we assume that

$$p_A^b < \mathbf{p}_S < p_A^r < 1. \tag{1}$$

[5] See seminal papers by Kahneman and Tversky (1972) and Tversky and Kahneman (1973, 1975).

[6] To our knowledge, this type of behavior has been only partially documented in the psychological literature and in Benjamin, Bodoh-Creed, and Rabin (2019).

[7] Vartanian et al. (2018) used fMRI to study how the cognitive reflection test can explain individual differences in BRN.

That is, the probability of receiving the right signal from source A in state R is higher than from source S but it is lower in state B.

The optimal decision. Without information agents guess the most likely color given the prior: B when $\rho < 0.5$ and R when $\rho > 0.5$. We call these guesses the *default guess (given the prior)*. It follows that the expected utility without information is equal to the expected utility of playing the default guess:

$$\underline{U} = \max\left\{\rho, 1-\rho\right\}.$$

When the priors are extreme, signals from sources with moderate precision cannot induce the agent to choose other color but the default guess. Why is it optimal to discard information? When receiving a signal, the agent updates her belief in the direction of the signal. If the signal goes against the prior and the precision of the source is not strong enough, the new posterior still remains on the wrong side of $\frac{1}{2}$. Hence, for extreme priors and moderate precision sources, both posteriors are on the same side of $\frac{1}{2}$, and the optimal decision is the default guess for both signals yielding expected utility $\underline{U}$.[8]

We say that an information source *has value (for prior ρ)* if the signals induce posteriors that lie on different sides of $\frac{1}{2}$. In the context of our experiment, we say that a source has value if an agent's unique optimal guess after signal r is R and after signal b is B. Given the previous discussion, only when priors are intermediate, $\rho \approx \frac{1}{2}$, sources can have value. In fact, given assumption 1, both sources have value for some intermediate priors.[9]

The following proposition summarizes the discussion above and presents the priors under which each source induces a choice different than the default guess. It describes how a von Neumann–Morgenstern utility maximizer agent *uses information*:

Proposition 1. *Source S has value if and only if* $\max\{1-\rho, \rho\} < p_S$. *Source A has value if and only if* $\max\{1-\rho, \rho\} < \mathbf{p}_A(\rho)$.

If one source does not provide value while the other does, the agent must choose the latter over the former. When both provide value, a von Neumann–Morgenstern utility maximizer agent compares the utility induced by the source S, $U(S) = \mathbf{p}_S$, and the utility provided by the source A, $U(A) = \mathbf{p}_A(\rho)$. The following proposition describes how this agent *chooses information sources* under the relevant parameters of our experiment:

[8]The result follows by exploiting the martingale property of bayesian updating.

[9]Since the signals are informative under both sources, always going against the signal is dominated.

Proposition 2. *If $p_S < \mathbf{p}_A(1-p_S)$, S is never optimal and A is optimal for some ρ; if $p_S > \mathbf{p}_A(p_S)$, A is never optimal and S is optimal for some ρ; otherwise, S is optimal at all $\frac{\rho}{1-\rho} \in \left(\frac{1-p_S}{p_S}, \frac{p_A^r - p_S}{p_S - p_A^b}\right)$ and A is optimal at all $\frac{\rho}{1-\rho} \in \left(\frac{p_A^r - p_S}{p_S - p_A^b}, \frac{p_A^b}{1-p_A^r}\right)$.*

4 The Experiment

To test our theoretical predictions on choice of sources and use of information, we ran a controlled laboratory experiment. The experiment consisted of four parts.

First part. The first and main part of the experiment implemented the decision problem described in Section 3. The parameters describing the asymmetric source were set throughout the experiment to

$$0.5 = p_A^b < p_A^r = 0.95$$

We manipulated the prior $\rho \in \{0.1, 0.45, 0.55, 0.75\}$ within subjects and the precision of the symmetric source $p_S \in \{0.6, 0.75\}$ between subjects. We refer to the treatment with $p_S = 0.6$ as the *low-precision treatment* and to the treatment with $p_S = 0.75$ as the *high-precision treatment*.

This part of the experiment consisted of four blocks, each consisting of ten identical rounds. The task was identical across the 40 rounds other than the prior ρ, which varied from block to block. To differentiate between the different problems implemented in the different blocks, each block had a different pair of colors—blue/red, yellow/brown, green/pink or orange/purple—used as labels for the two possible states of the world. The order of the four blocks, the matching of color pairs to priors, and the order of presentation of the advisors were randomized at the individual level and independently of each other.

The problem was presented to the participants as follows. In the beginning of each round, the color of a triangle was assigned according to the prior probability. The participants' goal was to guess the color of the triangle.

Before making their guess, they had to choose to receive information form source S or A, neutrally framed as advisor 1 and advisor 2. Each source was represented by two urns, one for each state, with the composition of the urns corresponding to the conditional probabilities that define the statistical properties of the source. For example, in the case of the asymmetric source, the urn corresponding to the B state contained ten blue balls and ten red balls ($p_A^b = 0.5$), and the urn corresponding to the red state contained 19 red balls and one blue ball ($p_A^r = 0.95$).

Once the advisor (source) was chosen participants were asked to state their guess for each of the two possible signals. The signal that determined the players' payoff was produced by drawing a ball from the relevant urn and the relevant advisor: if the color of the triangle was blue and the subject chose the source A, the ball was drawn from the the urn containing ten blue balls and ten red balls, and if the color of the triangle was red from the other urn.

At the end of each round, participants received feedback about the color of the triangle, the signal they received (given their chosen advisor), and their payoff, which was equal to 100 if they guessed the color correctly, or zero otherwise.

Second and third parts. The implicit assumption in the predictions of section 3 is that subjects are perfect Bayesians. In order to establish whether deviations can be attributed to preferences or mistakes in updating, we elicited beliefs. In part two of the experiment, we elicited the posteriors for every prior, source and signal experiences in the first part. In part three of the experiment, we elicited the perceived likelihood of receiving each signal for every prior and information source from part one.[10] We elicited beliefs with the binarized scoring rule proposed by Hossain and Okui (2013), as implemented in Wilson and Vespa (2016). The order of priors, sources and their respective colors was the same as in the first part part of the experiment.

Fourth part. The fourth part consisted of a basic questionnaire, a 16-time version of the raven test and a numeracy test. This part was not incentivized.

4.1 Experimental Procedures

The experiment was conducted at the ESSEXLab in December 2019. We ran four sessions with $p_S = 0.6$ and four sessions with $p_S = 0.75$, with a total of 190 participants (93 and 97 respectively). Students interacted through computer terminals, and the experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

All experimental sessions were organized along the same procedure. At the beginning of each part, subjects received detailed written instructions, which an instructor read aloud. Before starting each of the parts, participants had to answer a questionnaire to check their understanding of the experimental design.

To determine payments at the end of the experiment, the computer randomly selected two rounds in each of the four blocks of part one, two rounds from part two and two rounds from part three. Participants earned the total earned in these rounds. Points were converted to

[10] The utility of following source k can be written as $u(k) = \Pr(R|r,k)\Pr(r|k) + \Pr(B|b,k)\Pr(b|k)$, where posteriors are multiplied by the likelihood of each signal.

euros at the rate of 50 points = £1. In total, subjects earned an average of £19.17, including a show-up fee of 5 pounds. Each experimental session lasted approximately two hours.

4.2 Testable implications

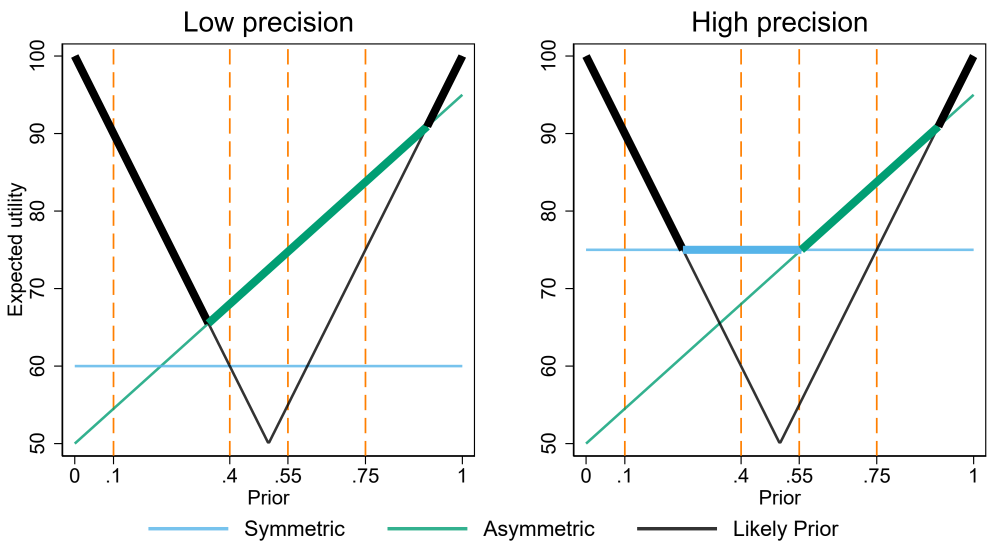

Figure 1: Expected utility. The thick lines mark the optimal payoff by prior. The vertical dashed lines mark the priors implemented in the experiment.

Under the assumption that agents are von Neumann Morgenstern expected utility maximizers Propositions 1 and 2 provide a series of testable implications regarding use of information and choice of source. In Figure 1 we present the the expected utility for different combination of strategies and source of information graphically.

Prediction about use of Information: It is immediate to see that, no source has value when $\rho = 0.1$ and source S has no value when $\rho = 0.75$. Moreover, in these cases, information should be ignored. Information from source A is used (so source A has value) in both treatments at all $\rho \in \{0.4, 0.55, 0.75\}$. Information from source S is used at $\rho = 0.55$ in both treatments, and when $\rho = 0.4$ in the second treatment.[11]

[11]Note that, S source can be ignored at $\rho = 0.4$ for $p_S = 0.6$ or it can be followed as one of the posteriors induced by this signal is equal to 0.5.

If agents understand the use of information, and as expected utility maximizers they do what they are supposed to do, we can infer the choice of source by applying Proposition 2. As stated before, if the prior is 0.1 neither S or A have value and we cannot predict which one will be selected. On the other hand, when $p_S = 0.6$ and $\rho = 0.75$ only A has value and this source is the one that should be chosen. A simple inspection of Figure 1 provides the following results

Prediction about choice of source of information: Source S should be chosen when $\rho = 0.4$ in the high-precision treatment. A should be chosen at all $\rho \neq 0.1$ in the low-precision treatment and when $\rho = 0.75$ in the high-precision treatment. When $\rho = 0.55$ in the high-precision treatment 2 both S and A provide equal value.

The use of comparative statics allow us to present a set of weaker predictions. First, for any $\rho \neq 0.1$, we should see higher demand of source S when $p_S = 0.75$ than the demand of source S when $p_S = 0.6$. Similarly, demand for source A (S) should increase (decrease) with ρ for any $\rho \neq 0.1$. Second, for any $\rho \neq 0.1$ and for a fixed p_S, the probability of choosing R conditional on receiving r should increase with ρ. Similarly, the probability of choosing B conditional on receiving b should decrease with ρ

5 Experimental results

5.1 Testing Predictions

Figure 2 shows the share of times that participants chose each source by treatment and prior.[12] For each source, the figure breaks down choices by whether the signal was followed or ignored. A few observations are readily apparent. First, almost all the direct implications of Propositions 1 and 2 failed to materialize: source S is chosen over source A when A has value. Moreover, when S has no value ($\rho = 0.75$ and $p_S = 0.6$) while A does, more subjects choose and use S over A. Second, information is used when it should not ($\rho = 0.1$) and is ignored when it should be used. That being said, when $\rho = 0.1$ ignoring information is the strategy chosen the most often.

Regarding the weak predictions that follow from comparative statics we have relative success. First, source choices are in line with differences in the expected utility of the source, both between and within treatments. Choice rates for the symmetric source are higher in the high-precision treatment, where it is more informative. Within treatments, participants

[12] All analyses reported here aggregate across the ten decision periods. The only learning apparent in the data is for the low prior $\rho = .1$, where participants ignore the signal in 52.6% of cases in the first period, increasing up to 68.4% in the last period. All of the results are robust to taking only the last five periods in each block.

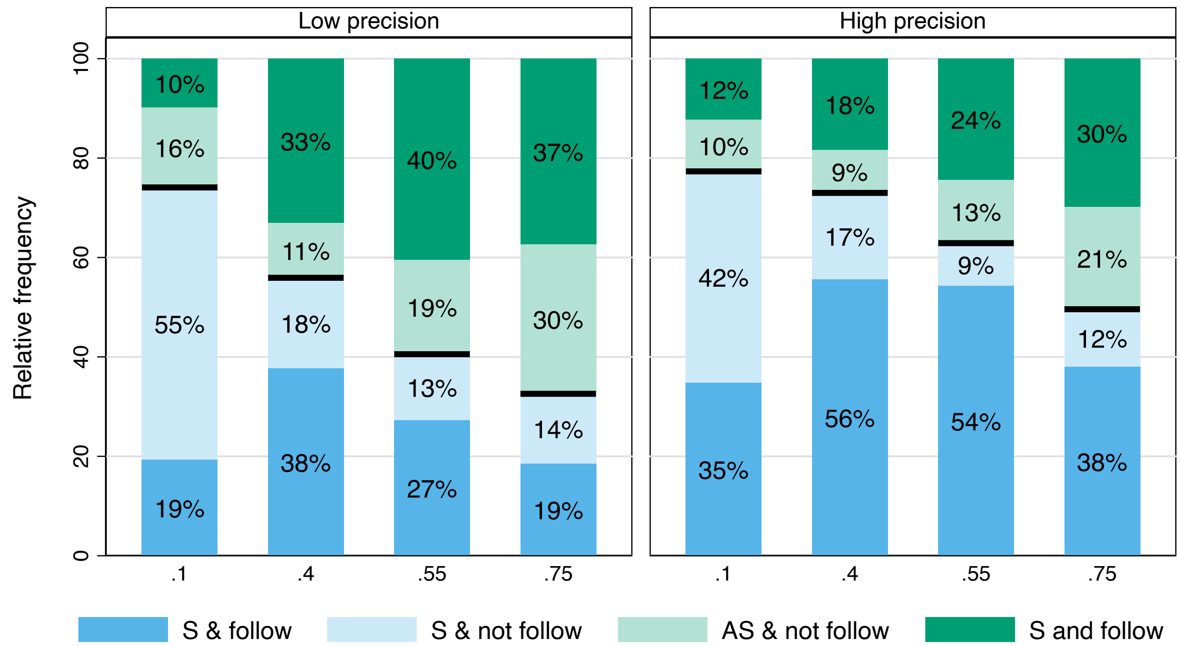

Figure 2: Source choice and decisions.

choose the symmetric source less often as the priors increase.[13] Second, actual use of the signal broadly follows the informativeness of the sources, as information is used more often for intermediate priors.

Result 1. *Qualitatively, changes in choice and use of information is in line with the statistical properties of the information. Notwithstanding, subjects systematically use information when they should not, ignore information when they should use it, and choose source S when they should not.*

The results of the experiment also provide a series of interesting and somehow unexpected results. First, although no source provides value when $\rho = 0.1$, the source that provides the highest expected utility (source S) is chosen the most. Moreover, it is used more often under the high-precision treatment than under the low-precision treatment. Second, when $\rho = 0.75$, only source A should be chosen, but source S is chosen in both treatments and at the same rate in the high-precision treatment. Third, source S is followed more often than

[13] At the individual level, 69 of the 190 (36.3%) participants are monotonic in the sense that they never choose the symmetric source more often for higher priors. Additional 59 participants for a total of 128 (67.4%) exhibit only one deviation from monotonicity. A permutation test reveals that these two numbers are significantly higher than expected by chance, given the choice distribution for each treatment and prior ($p < .001$).

source A even when the former has no value and the latter does, as in the high-precision treatment when $\rho = 0.75$. This suggests that expected utility may play a role in selecting sources and that symmetry may be important for the subjects when comparing sources. At the same time, choosing a source may induce subjects to use the information even though they should ignore it.

Result 2. *Expected utility and symmetry may be characteristic of sources that explain behavior and use of information beyond the instrumental value of information.*

5.2 The role of symmetry

To formally test and quantify the deviations from the model predictions, we fit a binary choice model under the assumption that participants have fixed preferences and the observed behavior results from noisy implementation of these preferences. The probability of implementation error is assumed to be a decreasing value of the cost of error. We use maximum likelihood estimation using the logit regression

$$Prob(s_k) = \frac{e^{\gamma \cdot u(s_k)}}{\sum_{s \in S} e^{\gamma \cdot u(s)}},$$

where γ is a precision parameter and s_k is a strategy in the set S of the eight possible strategies resulting from crossing the choice of source and guess.[14] We estimate the behavioral patterns in the data by allowing $u(s_k)$ to include terms for choosing the symmetric source, for following the signal, and for following the symmetric source specifically. That is,

$$u(s_k) = \pi(s_k) + \delta I^S + \lambda I^F + \xi I^{FS}, \tag{2}$$

where $\pi(s_k)$ is the expected payoff from following strategy s_k. I^S is an indicator for the four strategies where the symmetric source is chosen, I^F is an indicator for the two strategies that follow the signal, and I^{FS} is an indicator for the strategy of following the symmetric signal.[15] Thus, the estimated parameters reflect the cost in terms of probability of winning the prize that individuals are willing to pay for receiving information from the symmetric source or for following the signals.

Column (2) in Table 1 presents the results of this estimation, with the parameters scaled

[14] For each source, there are the four options of guessing color 1, guessing color 2, following the signal, or going against the signal.

[15] In this setting we want to assess preferences over characteristics; as such the concept of value is ill-defined as we are assuming that information is demand for other reasons beyond pure ability of changing actions.

Table 1: Regression estimation for choice and use of information.

	(1) Baseline	(2) Preferences	(3) Subjective beliefs	(4) Base-rate neglect	(5) Information neglect
δ	–	8.217***	12.38***	8.280*	7.479***
		(4.59)	(5.52)	(2.48)	(3.76)
λ	–	13.44***	18.26***	28.04***	3.258
		(5.39)	(5.65)	(6.28)	(1.20)
ξ	–	2.946	0.989	3.156	3.241
		(1.29)	(0.36)	(0.82)	(1.18)
γ	0.0608***	0.0535***	0.0461***	0.0485***	0.0593***
	(27.49)	(22.21)	(17.19)	(13.33)	(16.67)
N	7600	7600	7600	3480	4120
log lik.	$-11,894$	$-11,116$	$-11,490$	$-4,634$	$-6,243$
AIC	$23,791$	$22,239$	$22,989$	$9,277$	$12,494$
BIC	$23,798$	$22,267$	$23,017$	$9,302$	$12,520$

Notes: t-statistics in parentheses based on robust standard errors clustered on individuals. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

by a factor of 100.[16] The results support the observations made above. The model fit is better than the baseline model that includes only the probability of winning the prize, presented in Column (1). Participants tend to choose the symmetric source even if by doing so they lose an estimated 8.2 percentage points of winning the prize.

After receiving a signal, they tend to follow it at a cost of up to 13.4 percentage points of winning the prize. We do not find, however, significant evidence that participants are more likely to follow the signal with the symmetric source than the asymmetric one.[17]

Result 3. *Participants tend to choose the symmetric source and to follow any signal at a non-negligible cost.*

5.3 The role of expected utility

The payoffs $\pi(s_k)$ used in Table 2 were chosen in line with the (objective) instrumental value for information. That is, we assumed that subjects were proper Bayesian expected utility maximizers agents. The difference between instrumental value for information and expected

[16]we refer to the other columns in the following sections.

[17]This finding is robust to conditioning on the precision of the symmetric source.

utility is only relevant when information is useless; i.e. when posteriors beliefs do not change the ex-ante optimal action.

We conclude that it is not the relative expected utility what drives the choices of subjects when this differs from the relative value.[18] In order to argue that there are other heuristics determining choices or that source S is fundamentally different, it remains to show that the mistakes in forming posteriors beliefs cannot explain the data we observe. That is, to consider the subjecetive value, with imperfectly formed posteriors, as the driver behind the choices.

5.4 Subjective value

Table 2: Objective and subjective value of sources.

	Objective value			
	Symmetric source		Asymmetric source	
	No value	Positive value	No value	Positive value
Subjective value				
No value	49.7%	25.8%	70.5%	59.5%
Positive value	50.3%	74.2%	29.5%	40.5%
N	473	287	190	570

Notes: Subjective value determined based on subjective beliefs elicited in Phases 2 and 3 of the experiment.

The observed choice patterns fail to support the basic predictions in the model. It is possible, however, that distortions in probability perceptions and belief updating drive these patterns. As an initial step to explore this possibility, we calculated the subjective value as the expected payoff based on the subjective probabilities elicited in Phases 2 and 3 of the experiment.

Table 2 provides some evidence that subjective value may rationalize the observed choices. Even when the symmetric source has no value, the subjective value is positive for over half of the observations. On the other hand, the subjective value assigned to the asymmetric source is zero 60% of the time when the source has a positive value. Thus, the subjective probabilities overestimate the value of the symmetric source and underestimate the value of the asymmetric source.

[18] In part this may be due to the difference when no source of information has value may be captured by *symmetry*.

These systematic deviations from Bayesian updating apparent in the *belief* data may explain why our participants deviate from maximizing expected utility in their *choices*. Across the whole experiment, participants chose payoff-maximizing strategies in 46.0% of the rounds. As a first step of testing whether subjective beliefs can explain this low correspondence with the predictions of standard theory, we tested whether the share of optimal strategies increases when calculated based on the *subjective* beliefs. The share, in fact, decreases to 42.1%, indicating that distortions in belief updating is unable to rationalize choices.

We test the role of subjective beliefs in determining choices more formally by estimating a new model, replacing the expected payoff $\pi(s_k)$ in (2) with the subjective expected payoff $\pi^s(s_k)$ calculated by substituting the objective probabilities with elicited beliefs. If distortions in subjective probabilities explain the deviations in choice, the estimates for δ and λ should go to zero once these distortions are accounted for.[19] The results, presented in Column (3) of Table 1, are quite the opposite. The fit of the model decreases compared to the main model, and the estimates for both parameters do not decrease—and even increase by a factor of one-third to one-half.

Result 4. *Distortions in subjective probability perception are unable to explain away the choice patterns.*

5.5 Belief updating

We turn now to a deeper analysis of the belief updating process. Figure 3 plots the elicited posteriors across the experimental environments. Several observations are immediate. The average subjective posteriors (diamonds) generally do not correspond to the Bayesian posterior (triangles). In most cases, whenever the Bayesian posterior indicates positive value, so does the average subjective posterior. Nonetheless, there is a large variance and, as we saw earlier, subjective posteriors often indicate value when the source has no value objectively and vice versa.

In fact, many of the subjective posteriors are around the signal odds (circles). This is consistent with the recurring observation in the literature that people neglect the priors when forming beliefs, a phenomenon known as the representative heuristic or *Base-Rate Neglect* (BRN, see, e.g., Bar-Hillel, 1980; Grether, 1980; Kahneman and Tversky, 1973, 1996). Almost as prevalent in the Figure is the opposite observation, namely posteriors that are centered around the priors. We will return to them below but we anticipate here that we use the term *Information Neglect* (IN) to refer to that behavior in contraposition to BRN.

[19] Given that belief elicitation is noisy, it is possible that the estimates diminish yet remain significant.

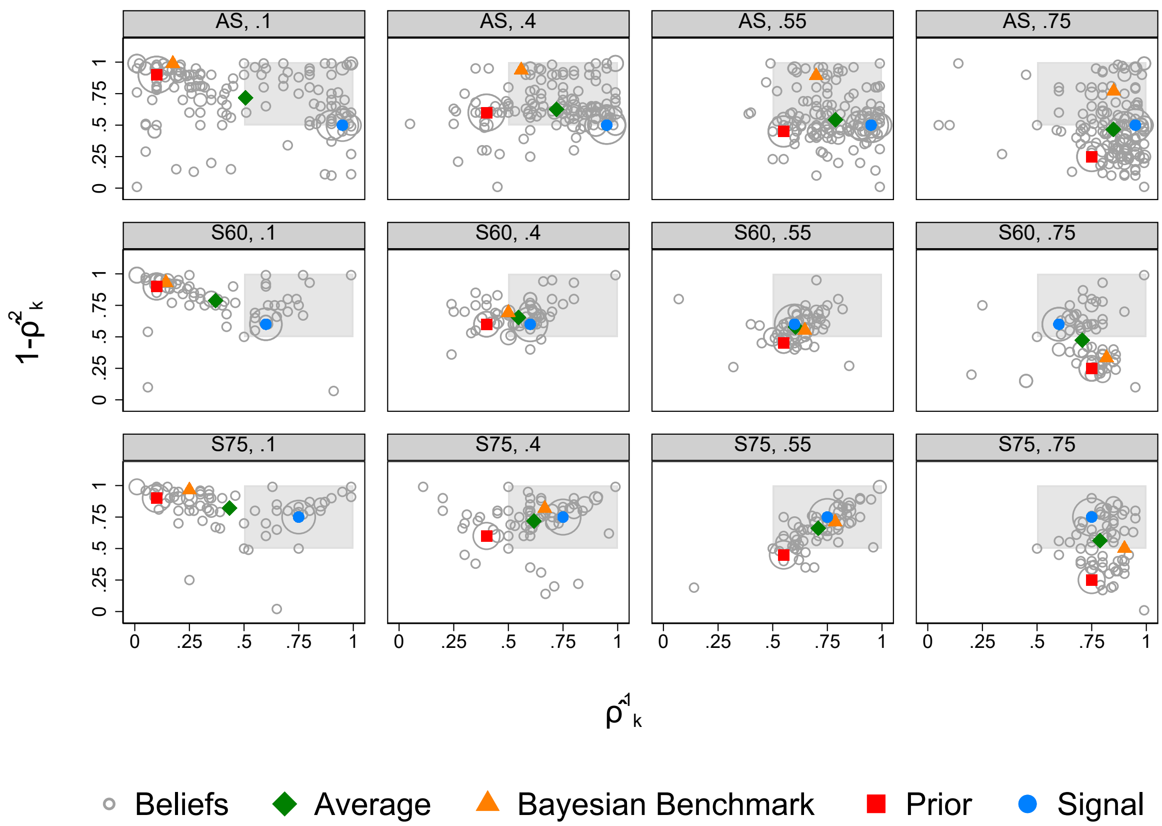

Figure 3: Elicited posteriors.

Table 3: Belief updating.

	α	β_1	β_2
Model 1			
	0.489	0.529	0.517
	(0.032)	(0.029)	(0.036)
Model 2			
Symmetric	−0.028	0.787	0.519
	(0.016)	(0.051)	(0.036)
Asymmetric	1.102	0.644	0.515
	(0.065)	(0.035)	(0.040)
Model 3			
BRN, Symmetric	−0.006	1.064	0.138
	(0.023)	(0.078)	(0.034)
BRN, Asymmetric	1.545	0.974	0.097
	(0.102)	(0.050)	(0.035)
IN, Symmetric	−0.056	0.545	0.828
	(0.023)	(0.052)	(0.039)
IN, Asymmetric	0.744	0.377	0.853
	(0.064)	(0.030)	(0.045)

Notes: Robust standard errors clustered on subjects in parentheses.

Following Grether (1980, 1992), we take the subjective posterior data from Phase 2 of the experiment to estimate the following model of beliefs formation for each signal realization s, which generalizes Bayes' Rule.

$$\ln OR\,(\hat{\rho}^s_{it}) = \alpha_i + \beta_1 \ln LR(p^s_k) + \beta_2 \ln OR\,(\rho) + \varepsilon_{it}, \tag{3}$$

where $OR(x) = \frac{x}{1-x}$ is the odds ratio of x, $\hat{\rho}^s_{it}$ is the reported posterior of subject i in environment t, $\mathrm{LR}(p^s_k)$ is the likelihood ratio of source k, and ε_{it} is a random noise variable with mean zero. The parameter α indicates assigning a high probability to the first color (the color that the asymmetric source is more likely to indicate). Parameters β_1 and β_2 reflect the weight placed on the signal and on the prior, respectively. Fixing $\alpha = 0$, $\beta_1 = 1$, $\beta_2 = 1$ yields the Bayesian posterior. Whenever $\beta_1 < 1$, the subject underweight the signal's contribution to form posteriors and, whenever $\beta_2 < 1$ underweight the prior's contribution to form posteriors, reflecting IN and BRN behavior, respectively.

Model 1 in Table 3 presents the estimation results.[20] The estimate for β_1 is higher than

[20] Based on OLS regression with standard errors clustered on subjects. The results are essentially identical for alternative specifications, including fixed- and random-effects for subjects.

the estimate for β_2, qualitatively replicating the result in Grether (1980). In contrast to Grether (1980), we find that both parameters are smaller than 1, indicating that subjective posteriors are closer to 0.5 than either the prior or the signal prescribe. The estimate for α is positive and significant. This result is, naturally, only meaningful for the asymmetric source, which defines the directionality of α. Model 2 includes additional coefficients for the source and its interactions with the model parameters. Indeed, we see that α is only significantly positive for the asymmetric source. This result is somewhat surprising, as it indicates the mere expectation of receiving a certain signal leads participants to assign a higher probability to the state of the world that the anticipated signal point to—regardless of the actual realization.

Result 5. *Participants deviate from Bayes' Law in a systematic way: participants exhibit base-rate neglect and information neglect on average. The BRN behavior is consistent with previous findings.*

As anticipated above, Figure 3 not only shows that many of the subjective posteriors are around the signal odds (circles) but also subjects whose posteriors are centered around the priors. We now turn to studying heterogeneity in formation of beliefs.

5.5.1 Individual heterogeneity

We next estimated the model independently for each participant. Figure 4 plots the individual estimations for (β_1, β_2).[21] Recall that perfect Bayesian updating appear as $\beta_1 = \beta_2 = 1$ and $\alpha = 0$. Observations to the left of the $\beta_1 = 1$ line reflect underweigthing of information (IN), and observations below the $\beta_2 = 1$ indicate underweighting of priors (BRN) in belief formation. We ran a 2-means cluster analysis to categorize the participants based on the two parameters.[22] To characterize the two resulting types, we estimated Model 3 in Table 3, which allows the parameters to vary by source and type.

While both types underweight priors, one type is characterized by a high ratio β_1/β_2, reflecting the fact that this type overweights information over priors. The other type is characterized by a low ratio β_1/β_2 reflecting the fact that this type underweights information over priors. Although, strictly speaking, all subjects show base-rate neglect behavior, because the first type equates the posteriors almost perfectly with the signal odds we reserve the name *Base-Rate Neglect* (BRN) to them.

[21] Three of the 190 participants fall out of the region depicted in the figure.

[22] The resulting types are robust to variations in the initial means and to including the α estimate. In contrast, categorization into three types (or more) is not robust. The Caliński-Harabasz pseudo-F stopping value for two clusters is $2,599$. The values for three to five cluster range from $2,098$ to $2,101$.

In contraposition to the (pure) BRN behavior we refer to those with a low ratio β_1/β_2 as *Information neglect* (IN): the IN type underweights both the prior and the signal, but assigns a higher weight to the prior. The base-rate neglect individuals appear as squares in Figure 4 and the information neglect individuals as circles. The large diamonds mark the mean parameters for the two types.

Result 6. *Although all subjects show some level of BRN behavior, there are two types, with around half of participants heavily overweighting information over priors and half of the of participants mildly overweighting priors over information in forming beliefs.*

Having identified the types based on subjective beliefs, we return to the choice data to test whether and how the belief updating style of the participants correlates with their strategies. Columns (4) and (5) in Table 1 present separate estimations for the two types. Although types do not differ with regard to their preference for the symmetric source, they differ on the tendency to follow the signal. This is captured by the parameter λ, which is only significant for base-rate neglect types. That is, the same people who ignore the prior when reporting posteriors also neglect the fact that following the prior and ignoring the signal may be optimal.[23]

Result 7. *The elicited beliefs are meaningful. They identify two types differing not only in their style of belief updating, but also in the way they use information.*

This result may be striking at first sight but, actually, reflects certain consistency between beliefs formation and use of information. If subjects follow their posteriors when using information, and posteriors are constructed relying heavily on information and neglecting the priors, the priors should not play a role when using information. Alternatively, the strategies used by BRN are consistent with their inability to form beliefs properly.

5.5.2 Contrarian updating

The analysis suggests two heuristics used prominently by different individuals. One of the most striking predictions of the use of heuristics to form posteriors is that they may *reverse* the meaning of a signal. That is, a signal in one direction pushes posteriors away from the priors in the opposite direction.We refer to this phenomenon as *contrarian updating.*[24]

For base-rate neglect, contrarian updating results directly from interpreting the precision of the signal as the posterior.[25] For example, consider a perfect BRN participant facing a

[23]Recall, however, that the subjective posteriors are not able to explain the tendency to follow the signal on an individual basis.

[24]Benjamin, Bodoh-Creed, and Rabin (2019) use the term *extreme moderation effect* and discuss supporting evidence (Bar-Hillel, 1980; Griffin and Tversky, 1992).

[25]Formally, $\hat{p}_k^s = p_k^s$

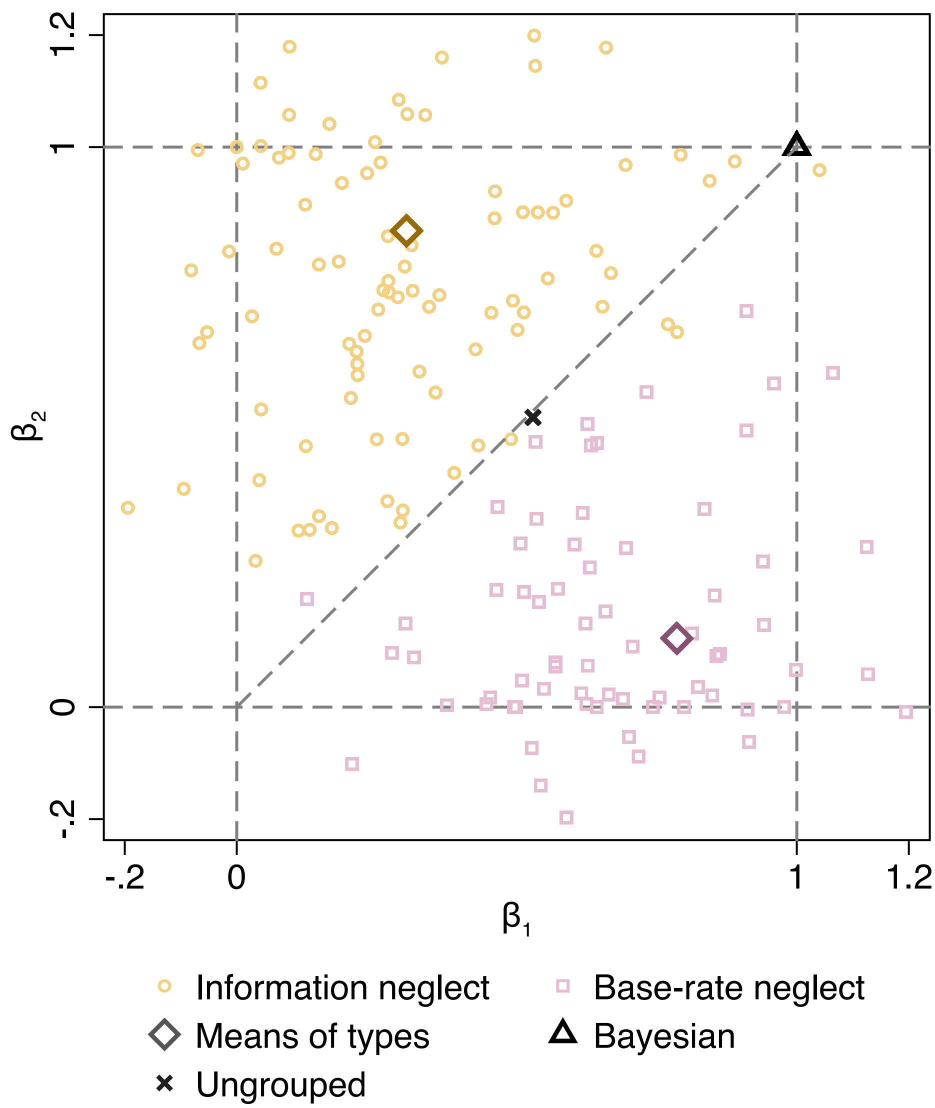

Figure 4: Individual parameters.

symmetric signal with precision $p = 0.75$ (as in the high-precision treatment). The probability that this subject assigns to the state being blue following a b signal is 0.75 (and 0.25 following a red one). If the prior is 0.1 (for red), then a b signal would lead the participant's subjective posterior for the blue state to (incorrectly) *decrease* from 0.9 to 0.75. In general, contrarian updating due to base-rate neglect occurs whenever the prior is more extreme than the precision.

In our experiment, base-rate neglect leads to contrarian updating in five situations, in which the prior is more extreme than the signal's precision: prior 0.1 and a b signal from any of the three sources (S low precision, S high precision, and AS); prior 0.4 and a b signal from the asymmetric source; and prior 0.75 and an r signal from the low-precision symmetric source.[26]

Consider also the implications of the information neglect heuristic for contrarian updating. Information-neglect types form posteriors that are close or even equal to the priors. Therefore, they may show contrarian updating due to even small amounts of noise in the reported posteriors. We use these two observations to provide conclusive evidence for heuristic updating by comparing contrarian updating across the two types. We predict more contrarian updating from BRN types in the five situations in which the signal's precision is less extreme than the prior. In seventeen combinations of source and signal, the precision is more extreme than the prior. In these cases we expect noise to lead to more contrarian updating by the IN types.

Table 4 presents the rates of contrarian updating for the two types across the experiment. As the heuristic predicts, BRN types exhibit significantly more contrarian updating in all five predicted cases. IN types exhibit more contrarian updating in fifteen of the seventeen relevant cases.[27]

Result 8. *Subjective posteriors exhibit substantial contrarian updating, which is inconsistent with value-based models. These patterns are consistent with heuristic-based reasoning and the categorization into two types based on different heuristics.*

6 Conclusion

In this paper, we study how people choose between information sources and how they use the information from these sources. We find systematic deviations from the optimal strategies.

[26]There are also two situations in which the signal precision equals the prior (prior 0.4 and a b signal from the low-precision symmetric source and prior 0.75 and an r signal from the high-precision source). In these cases, BRN implies no updating and any amount of noise may lead to contrarian updating.

[27]Out of these fifteen cases, the difference is significantly different than zero in ten of them. See table 5 in the appendix for the proportion tests of these comparisons.

Table 4: Contrarian updating.

	Information Neglect					
	S Low Precision		S High Precision		Asymmetric	
Prior	b signal	r signal	b signal	r signal	b signal	r signal
0.10	**0.64**	*0.15*	**0.64**	*0.10*	**0.76**	*0.14*
0.40	0.43	*0.17*	*0.34*	*0.18*	**0.6**	*0.09*
0.55	*0.26*	*0.17*	*0.24*	*0.04*	*0.51*	*0.07*
0.75	*0.19*	**0.34**	*0.10*	0.14	*0.15*	*0.1*
	Base Rate Neglect					
	S Low Precision		S High Precision		Asymmetric	
Prior	b signal	r signal	b signal	r signal	b signal	r signal
0.1	**0.88**	*0.00*	**0.87**	*0.02*	**0.82**	*0.02*
0.40	0.45	*0.00*	*0.11*	*0.02*	**0.61**	*0.05*
0.55	*0.00*	*0.08*	*0.06*	*0.02*	*0.09*	*0.02*
0.75	*0.08*	**0.7**	*0.00*	0.11	*0.03*	*0.13*

Notes: Numbers in bold indicate situations where the prior is more extreme than the signal precision and vice versa for numbers in italics.

Participants systematically choose a symmetric information source over an asymmetric one and overuse information. Subjective beliefs cannot account for these deviations, suggesting that participants have non-instrumental value for information.

We focus on the process of belief formation to understand these deviations. The beliefs elicited in the experiment reveal how participants perceive and process information. We identify two different heuristics to form beliefs: those that heavily weight priors over information and those that heavily weight information over priors. The latter type exhibits base-rate neglect, a much-studied phenomenon. While some studies have considered heterogeneity in BRN tendencies (Vartanian et al., 2018; Wolfe and Fisher, 2013), to the best of our knowledge we are the first to document a dichotomous categorization into types and relate these types to differences in the use of information.

Base-rate neglect, sometimes called the representative heuristic, can be conceptualized as conflating posteriors with precisions. In our setting, BRN participants understand the probability of the state being Red given an r signal to be the same as the probability of receiving an r signal given the state being Red (cf. Benjamin, Bodoh-Creed, and Rabin, 2019). The use of the heuristic leads to systematic overestimation of the value of the symmetric source and underestimation of the asymmetric source. Note, however, that this is a result of the specific setup, and in particular our choice of setting the precision of the asymmetric source to be 0.5 in the Blue state. This design choice allows us to identify heuristic thinking by generating clear predictions based on the heuristic with regard to symmetry and contrarian updating. The conclusion is quite dramatic. Economists often consider and analyze information acquisition and use in terms of value. In contrast, our results suggest that the decision process, which relies on heuristics, drives information-related decisions. The same process may lead choice patterns that are consistent with different formulations in terms of value *depending on the problem* at hand.

References

Bar-Hillel, Maya (1980). The base-rate fallacy in probability judgments. *Acta Psychologica* 44(3), pp. 211–233.

Benjamin, Dan, Aaron Bodoh-Creed, and Matthew Rabin (2019). Base-rate neglect: foundations and implications. mimeo.

Charness, Gary and Chetan Dave (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior* 104, pp. 1–23.

Charness, Gary, Ryan Oprea, and Sevgi Yuksel (2021). How do people choose between biased information sources? evidence from a laboratory experiment. *Journal of the European Economic Association* 19(3), pp. 1656–1691.

Eliaz, Kfir and Andrew Schotter (2007). Experimental testing of intrinsic preferences for nonlnstrumental information. *The American economic review* 97(2), pp. 166–169.

— (2010). Paying for confidence: an experimental study of the demand for non-instrumental information. *Games and Economic Behavior* 70(2), pp. 304–324.

Falk, Armin and Florian Zimmermann (2014). Beliefs and utility: experimental evidence on preferences for information. working paper.

Fischbacher, U. (2007). Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), pp. 171–178.

Golman, Russell, David Hagmann, and George Loewenstein (2017). Information avoidance. *Journal of Economic Literature* 55(1), pp. 96–135.

Grether, David M. (1980). Bayes rule as a descriptive model: the representativeness heuristic. *Quarterly Journal of Economics* 95(3), pp. 537–557.

— (1992). Testing Bayes rule and the representativeness heuristic: some experimental evidence. *Journal of Economic Behavior & Organization* 17, pp. 31–57.

Griffin, Dale and Amos Tversky (1992). The weighing of evidence and the determinants of confidence. *Cognitive psychology* 24(3), pp. 411–435.

Hossain, Tanjim and Ryo Okui (2013). The binarized scoring rule. *Review of Economic Studies* 80(3), pp. 984–1001.

Huck, Steffen, Nora Szech, and Lukas M. Wenner (2015). More effort with less pay: on information avoidance, belief design and performance. Tech. rep. Working Paper, UCL.

Jones, Martin and Robert Sugden (2001). Positive confirmation bias in the acquisition of information. *Theory and Decision* 50(1), pp. 59–99.

Kahneman, Daniel and Amos Tversky (1972). Subjective probability: a judgment of representativeness. *Cognitive psychology* 3(3), pp. 430–454.

— (1973). On the psychology of prediction. *Psychological Review* 80(4), pp. 237–251.

Kahneman, Daniel and Amos Tversky (1996). On the reality of cognitive illusions. *Psychological review* 103(3), pp. 582–596.

Masatlioglu, Yusufcan, A. Yesim Orhun, and Collin Raymond (2016). Preferences for non-instrumental information and skewness. working paper.

Montanari, Giovanni and Salvatore Nunnari (2019). Audi alteram partem: an experiment on selective exposure to information. Tech. rep. Working Paper.

Rabin, Matthew and Joel L. Schrag (1999). First impressions matter: a model of confirmatory bias. *The Quarterly Journal of Economics* 114(1), pp. 37–82.

Tversky, Amos and Daniel Kahneman (1973). Availability: a heuristic for judging frequency and probability. *Cognitive psychology* 5(2), pp. 207–232.

— (1975). Judgment under uncertainty: heuristics and biases. In: *Utility, probability, and human decision making*. Springer, pp. 141–162.

Vartanian, Oshin et al. (2018). The reflective mind: examining individual differences in susceptibility to base rate neglect with fmri. *Journal of cognitive neuroscience* 30(7), pp. 1011–1022.

Von Gaudecker, Hans-Martin, Arthur Van Soest, and Erik Wengström (2011). Heterogeneity in risky choice behavior in a broad population. *The American Economic Review* 101(2), pp. 664–694.

Wilson, A. and Emanuel Vespa (2016). Paired-uniform scoring: implementing a binarized scoring rule with non-mathematical language. working paper.

Wolfe, Christopher R. and Christopher R. Fisher (2013). Individual differences in base rate neglect: a fuzzy processing preference index. *Learning and Individual Differences* 25, pp. 1–11.

Zimmermann, Florian (2014). Clumped or piecewise? evidence on preferences for information. *Management Science* 61(4), pp. 740–753.

A Appendix: Experimental Instructions

Welcome and thank you for taking part in this experiment. Please remain quiet and switch off your mobile phone. It is important that you do not talk to other participants during the entire experiment. Please read these instructions very carefully; the better you understand the instructions the more money you will be able to earn. If you have further questions after reading the instructions, please give us a sign by raising your hand out of your cubicle. We will then approach you in order to answer your questions personally. Please do not ask aloud.

During the experiment all sums of money are listed in ECU (for Experimental Currency Unit). Your earnings during the experiment will be converted to pounds at the end and paid to you in cash. The exchange rate is 50 ECU = £1. The earnings will be added to a participation payment of £5.

In this experiment you will not interact with other participants. Your earnings will depend on your decisions and on chance, as will be explained in the following. The experiment is divided into three parts. Your final payoff will be the sum of your earnings in the three parts.

After the experiment, we will ask you to complete a short questionnaire, which we need for the statistical analysis of the experimental data. The data of the questionnaire, as well as all your decisions during the experiments will be anonymous.

Please, stay in your seat until all other subjects have finished the experiment. The lab administrator will let you know when to stand up.

INSTRUCTIONS PART I

This part consists of **four** blocks, and each block consists of **ten** rounds. The instructions are **identical** for each round. Your task in each round will be to **guess the colour of a triangle**. As explained below, you will know **how the computer chooses** the colour of the triangle. You will also be able to choose a computerised advisor to receive **advice** from regarding the colour of the triangle.

The Triangle Colour

The **colour of the triangle** will be **randomly** chosen at the **beginning of each round** to be one of two colours. For example, the triangle can be either blue ▲ or red ▲. The **two possible colours** will change from block to block, and will be announced at the beginning of each block. In these instructions, we will use blue and red as in the example above.

The **probability** that the triangle is blue ▲ or red ▲ will also change from block to block and will be announced at the beginning of each block. For example, the triangle is blue ▲ with 40% probability, and red ▲ with 60% probability. This will be presented on the computer screen as:

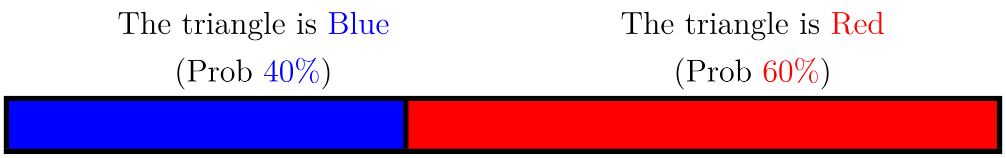

The advisors

You will not know the colour of the triangle, and **your objective will be to guess the colour of the triangle**. For that, you will receive an advice from an advisor that will be simulated by the computer. You will be able to choose **one of two advisors**.

Each advisor **has two urns**, a **blue urn** and a **red urn** (corresponding to each of the colours). Each urn is filled with 20 balls, some of which are blue and the others are red. **There are always more blue balls in the blue urn than in the red urn** (and, therefore, more red balls in the red urn than in the blue urn).

Which advice an advisor gives you depends on a **random ball** drawn from the **urn that corresponds to the true colour of the triangle**. That is,

- if the colour of the triangle is blue ▲, a ball will be drawn from the blue urn
- if the colour of the triangle is red ▲, a ball will be drawn from the red urn.

For example, the blue urn may contain 14 blue balls and 6 red balls, and the red urn may contain 8 blue balls and 12 red balls, as in the figure below. As you can see, there are more blue balls in the blue urn than in the red urn.

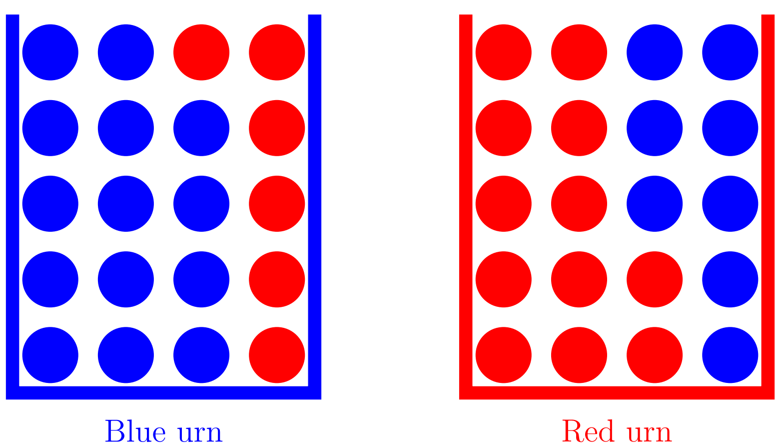

Choosing advisors

In each round you will be asked to choose between **two advisors**. The computer will present you with the **two urns** of each of the **two advisors**. You should consider these urns carefully, and decide **which advisor you prefer to receive advice from**. The two advisors will be the same for all rounds within a block.

Guessing the colour of the triangle

After choosing the advisor, you will be asked to guess the **colour of the triangle**. Your choice can **depend on the advice** you receive, in the following way. The computer will ask you to guess the colour of the triangle **twice**. Once for the case that the advisor you chose gives you a blue advice, and again for the case that the advisor you chose gives you a red advice. After you have made the choice, and without observing the advice, the computer will determine the advice by randomly selecting a ball from the urn of the advisor you chose that corresponds to the colour of the triangle. That is, if the colour of the triangle is blue ▲, the computer will choose a ball from the blue urn. If the colour of the triangle is red ▲, the computer will choose a ball from the red urn. The colour of the **drawn ball** will determine the **advice you receive**, in the following way: if the ball drawn is blue, the computer will implement the guess that you chose after a blue advice, and if the colour of the drawn ball is red, the computer will implement the guess that you chose after a red advice. Your **payoff** will be determined by the **guess corresponding to the advice** and the colour of the triangle.

Summary of the round

1. The computer chooses the colour of a triangle according to a known rule (probability).
2. You choose one of two advisors.
3. You choose which colour to guess for each advice.
4. The computer determines the advice and follows your guess.

Your Payoff

Your payoff will depend on **whether you guessed the colour of the triangle correctly**. If your guess is correct, you will receive 100 ECU. Otherwise you will receive 0 ECU.

Information at the end of each Round

At the end of each round, you will receive the following information about the round: the **colour of the triangle**, which **advisor** you chose, what **advice** you received from the advisor, what was your **guess**, and your **payoff** for the round.

Final Earnings

At the end of the experiment, the computer will randomly select **two rounds** in each of the **four blocks**, which makes a total of **eight** rounds. You will receive the payoffs that you had earned in **each** of these selected rounds. Each of the 40 rounds has the **same chance** of being selected.

Control Questions

Before starting the experiment, you will have to answer some control questions in the computer terminal. Once you and all the other participants have answered all the control questions, Round 1 will begin.

B Appendix: additional analyses

Table 5: Proportion tests for contrarian updating.

		Low precision		High Precision		Asymmetric	
		r signal	b signal	r signal	b signal	r signal	b signal
.1	IN	0.64	0.15	0.64	0.10	0.76	0.14
	BRN	0.88	0.00	0.87	0.02	0.82	0.02
	z-value	-2.55**	2.57**	-2.65***	1.61	-0.98	2.79***
	p-value	0.011	0.010	0.008	0.108	0.326	0.005
.4	IN	0.43	0.17	0.34	0.18	0.60	0.09
	BRN	0.45	0.00	0.11	0.02	0.61	0.05
	z-value	-0.15	2.74***	2.75***	2.57**	-0.10	1.13
	p-value	0.877	0.006	0.006	0.010	0.919	0.260
.55	IN	0.26	0.17	0.24	0.04	0.51	0.07
	BRN	0.00	0.07	0.06	0.02	0.09	0.02
	z-value	3.53***	1.35	2.40**	0.53	6.22***	1.45
	p-value	0.000	0.177	0.016	0.594	0.000	0.146
.75	IN	0.19	0.34	0.10	0.14	0.15	0.10
	BRN	0.07	0.70	0.00	0.11	0.03	0.13
	z-value	1.57	-3.44***	2.23**	0.50	2.61***	-0.64
	p-value	0.118	0.001	0.026	0.615	0.009	0.520