Beyond Value: on the Role of Superficial Characteristics in Demand for Information*

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

October 31, 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. We elicit beliefs and find that participants consistently overestimate the instrumental value of the symmetric source yet fail to explain the choice pattern. The analysis of the belief data reveals two types of systematic deviations from Bayesian belief updating: about half of the subjects exhibit base-rate neglect while the others systematically neglect new information more than the base rate. We leverage this understanding of belief formation processes to provide further evidence of intrinsic preferences for symmetry.

Keywords: information acquisition, representative heuristic, base-rate neglect, laboratory experiment

*We thank Gary Charness, Kirill Pogorelskiy, Mauricio Ribeiro, and conference and seminar participants at the Economic Science Association online global meeting, the Econometric Society European Meeting in Barcelona, 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. This project has received funding by the British Academy (SG171679).

†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@post.bgu.ac.il`.

1 Introduction

Information derives its instrumental value from its strength to change a subject's actions. For instance, consider a pitted fruit farmer who is contemplating a crop expansion. This farmer will only reconsider their plans if they receive trustworthy information indicating an imminent occurrence of heavy rainfall in the future. Conversely, a rice farmer facing a comparable expansion decision would only alter their course of action if they obtains reliable information signaling an uncharacteristically dry weather pattern. Which type of information is more valuable depends on the farmer's preferred action sans information.

This differential need for a particular type of *certainty* informs what type of information people seek.[1] For instance, doctors who encounter potential cancer patients are primarily concerned with avoiding the failure to detect cancer rather than the risk of misdiagnosis. To alleviate their concerns, they often request additional tests to enhance the certainty regarding the absence of the disease. Voters whose primary concern revolves around budgetary issues tend to seek information about candidates' economic plans, while those worried about ideological matters are more inclined to investigate the candidates' character. More generally, valuable information may be prone to provide 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 state. 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 posteriors of a state given the realised signal, $\Pr(\omega = A|s_A)$ and $\Pr(\omega = B|s_B)$. In contrast, it is natural to describe an information source by its precision in each state—the probability of a *signal realization given the state*, formally $p_A = \Pr(s_A|\omega = A)$ 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 drought $\Pr(\omega = A|s_A)$ or risk of flood $\Pr(\omega = B|s_B)$ is essential for the farmer to decide whether the information provided by the weather source is useful or not.

If agents construct the right mapping between the priors and the signal to the posteriors, no other statistical characteristics of the problem should hold relevance. For instance, subjects should not be concerned about whether the information source is symmetric ($p_A = p_B$) or whether a source is fully revealing of one of the states (i.e., $p_A = 1 > p_B > 0$). However, in cases where agents fail to construct this mapping correctly, these seemingly superficial characteristics of the information sources may factor into subjects' decision-making and in-

[1] By certainty, we do not imply that there is perfect knowledge about the state of the world. We rather use it as a graded notion of trust in the prediction, akin to the term "diagnostic certainty" in the clinical literature.

formation processing. This leads to important questions: do these superficial characteristics play a significant role in information processing? Do they influence how subjects perceive the value of information? Do subjects exhibit intrinsic preferences regarding these characteristics?

In this paper, we investigate the role of symmetry—a seemingly irrelevant characteristic—in shaping the demand for information. We present findings from a laboratory experiment where participants obtain a prize if they guess the state of the world correctly. Prior to making their prediction, they choose whether to receive information from a symmetric source or an asymmetric source. To eliminate potential confounds, such as reputation concerns or motivated beliefs, we adopt a minimalist approach and present information to the participants in an abstract setup with binary states and binary signals. In this context, we manipulate participants' prior beliefs and the precision of the symmetric source. We thus vary the instrumental value and the relative order of the two information sources, allowing us to identify systematic preferences for the symmetric source *beyond value*.

Two clear patterns emerge from subjects' behaviour in the experiment. First, regardless of the chosen information source, subjects use information too much. In other words, they often follow the signal realisation even when it would be optimal to disregard it and make their prediction based solely on the prior information. Second, subjects systematically choose the symmetric source over the asymmetric one more frequently than prescribed by the comparison of their instrumental values. These patterns suggest that subjects assign more than the instrumental value to information sources and find symmetry appealing.

To interpret the findings as evidence for intrinsic preferences for information, we must account for potential mistakes in processing the information. To address this issue, we elicit (incentivised) subjective posteriors and expected frequency of signals. We use the elicited beliefs in two distinct ways. First, we construct expected payoff based on subjective beliefs. This *subjective* expected payoff fails to predict choice behaviour better than the *objective* expected payoff. Second, we model the belief updating process using Grether's (1980) flexible framework and perform a k-means cluster analysis based on the model estimates. This analysis reveals that differences in how agents process information help explain the observed patterns in the use of information.

The cluster analysis identifies two groups of subjects: one group (heavily) overweight information relative to priors whereas the other group (mildly) overweight priors relative to information. Notably, while both groups neglect priors when forming beliefs, the first group appears to ignore them completely. Following the literature, we refer to the first group as *Base rate neglect* (BRN) and to the second group as *Information neglect* (IN).[2]

[2]Strictly speaking, BRN subjects underweight priors and weight information almost appropriately while

We then use this categorisation to study alternative ways of using information beyond utility maximisation. In particular, we can explain why some subjects tend to overuse information, while others do not. To fix ideas, consider a pure BRN subject who completely ignores priors and receives information from a symmetric source. First, given the source of information, this subject's posterior belief about state ω after receiving a signal corresponding to that state is simply the source's precision. Second, since the precision of this source is always higher than 0.5, the two possible posteriors of this subject fall on opposite sides of 0.5. Third, since the subject wants to maximise the probability of guessing the state correctly, she always follows the decision recommended by the signal. Therefore, pure BRN agents always perceive symmetric sources as having instrumental value, regardless of the prior. In contrast, the posteriors of our estimated IN types do not exhibit this tendency to overestimate the value of the sources.

Understanding the beliefs formation process allows us to provide stronger evidence regarding the role of symmetry and its desirability. Given the parameters in the experiment, BRN subjects rank information sources in a clear way independently of the prior. In particular, our design includes situations where BRN subjects should never choose the symmetric source. Nonetheless, they tend to choose the symmetric source even in those situation, revealing intrinsic preferences for this superficial characteristic.

The rest of the paper is organised as 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 regarding 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 Literature Review

There is ample evidence that subjects demand 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 in the demand for information; Eliaz and Schotter (2007, 2010) find a "confidence effect" driving demand for

information neglect subjects underweight both information and priors. See Section 5.3 for the formal discussion for this definition and characterization.

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

information beyond its value.[4] In our paper we identify a reason that induces agents to appear to overvalue information. Unlike in the papers listed above, we identify a purely statistical characteristic that is superficial and irrelevant.[5]

Our paper directly relates to papers that present information in purely statistical terms and study demand for *biased* information. Masatlioglu, Orhun, and Raymond (2023) study subjects choosing between information sources that have no value. They document evidence that subjects demand positively skewed information sources. Charness, Oprea, and Yuksel (2021) study subjects choosing between asymmetric sources. They identify patterns of demand for information for confirmatory, anti-confirmatory, and uncertainty reduction reasons. In their experiment, subjects' mistakes cannot result from calculation errors but reflect heuristics highlighting the above mentioned 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 comes from the source biased against their priors.

Unlike these papers, we focus on choices between symmetric (equal precision in both states) and asymmetric (different precision in each state) sources and vary the value of information by manipulating the prior. This allows us to compare situations when both sources have no value (as in Masatlioglu, Orhun, and Raymond, 2023) and when only one has value (as in some cases in Montanari and Nunnari, 2019; Charness, Oprea, and Yuksel, 2021). We find that subjects rely too much on information, i.e, they use it too much. Symmetry emerges as a desirable characteristic, as subjects choose the symmetric source even at an expected loss.

Ambuehl and Li (2018) also use a minimal framework to study how individuals respond to information. They derived an indirect ranking of information sources by eliciting a lottery equivalence to using each information source. They find intrinsic preferences for certainty revealing information but no difference between symmetric and asymmetric sources. In all of the treatments in Ambuehl and Li (2018), the prior was uniform—eliminating base-rate neglect by design—and all information sources had instrumental value. In contrast, we find that BRN, as well as the comparison between valuable and non-valuable information sources, play a crucial role in identifying the effect of symmetry.

Our paper also relates to the literature documenting systematic deviations from Bayesian

[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 behaviour towards information.

[5]More recently, Guan, Oprea, and Yuksel (2023) found that people prefer information structures that have the same value but with a lower *entropy informativeness*.

updating.[6] In particular, we study the process of belief formation directly as in Grether (1980) to find that subjects tend to form posterior beliefs in a predictable but non-Bayesian way. People neglect both the prior and the information compared to the Bayesian benchmark, but place more weight on the signal relative to the prior (BRN). Our method allows us to identify a more nuanced behaviour: we find that some subjects exhibit BRN, while others neglect information more than they neglect priors.[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, materialises with probability ρ, and the blue one, B, with probability $1-\rho$. The agent obtains one pound if her guess matches the state, i.e., $u(g|\omega) = \mathbb{1}_{g=\omega}$.

Before making a guess, the agent can collect information from one of two different sources of information: the symmetric source S, or the asymmetric source 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

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

Note that given the utility function, $\mathbf{p}_k(\rho)$ is also the expected utility of a subject with prior ρ who follows the signals provided by source k blindly.

We assume that under the symmetric source (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: $p_S = \mathbf{p}_S(\rho)$. In contrast, under the asymmetric 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

$$0.5 \leq p_A^b < \mathbf{p}_S < p_A^r < 1. \tag{1}$$

That is, the probability of receiving the right signal from source A in state R is higher than

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

[7]This type of behaviour has been documented in the psychological literature and in Benjamin, Bodoh-Creed, and Rabin (2019). Ambuehl and Li (2018) use Grether's framework and found heterogeneity in information neglect, assuming away the possibility of base-rate neglect. Wolfe and Fisher (2013) and Vartanian et al. (2018) tested correlations of individual differences in base-rate neglect and various tasks.

from source S but it is lower in state B.

The optimal decision. In this section we characterise the optimal behaviour of a (Bayesian) payoff maximizer. See first that the optimal decision given a posterior, is to choose the most likely colour. Without information agents guess the most likely colour 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:

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

When the priors are extreme, signals from sources with moderate precision cannot induce the agent to choose other colour 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 same side of $\frac{1}{2}$ than the prior. Hence, for extreme priors and moderate precision sources, both potential posteriors are on the same side of $\frac{1}{2}$, and the optimal decision is the default guess for both signals yielding expected utility $\underline{U}$.

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, when priors are not extreme, sources can have instrumental value. In fact, given assumption (1), both sources have value for some intermediate priors.

The following proposition summarises the discussion above and presents the priors under which each source induces a choice different than the default guess. It describes the optimal *use of 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 maximiser 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:

Proposition 2. *A is the optimal choice if* $\frac{\rho}{1-\rho} \in \left(\max\left\{\frac{p_A^r - p_S}{p_S - p_A^b}, \frac{p_A^r}{1-p_A^b}\right\}, \frac{p_A^b}{1-p_A^r}\right)$ *and S is the optimal choice if* $\frac{\rho}{1-\rho} \in \left(\frac{1-p_S}{p_S}, \min\left\{\frac{p_A^r - p_S}{p_S - p_A^b}, \frac{p_S}{1-p_S}\right\}\right)$*; otherwise, following the priors (ignoring the signal) is optimal.*

By assumption (1), both information sources have value for some priors, but under some conditions, one information source always provide a lower expected utility than the other one. For example, if $\frac{p_S}{1-p_S} < \frac{p_A^r}{1-p_A^b}$, whenever S has value A provides a higher expected utility: $p_S < \mathbf{p}_A(1-p_S)$. On the other hand, if $\frac{p_A^b}{1-p_A^r} < \frac{p_S}{1-p_S}$, S is better than A, whenever A has value.

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: the main task, two belief elicitation tasks, and survey questions.

First part. The first and main part of the experiment implemented the decision problem described in Section 3. Participants guessed the colour of a triangle (corresponding to the state ω) based on a prior ρ and information from an advisor they chose of two possible advisors (corresponding to sources S and A). The precisions of the asymmetric source were $p_A^b = 0.5$ and $p_A^r = 0.95$ across the experiment. 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 treatments with $p_S = 0.6$ as the *low-precision* treatments and to the treatments with $p_S = 0.75$ as the *high-precision* treatments.

This part of the experiment consisted of forty rounds divided into four blocks. The blocks differed in the prior ρ and in the two possible colours of the triangle—corresponding to the two possible states of the world– which were blue/red, yellow/brown, green/pink, or orange/purple. The order of the priors across blocks, the matching of colour pairs to priors, and the matching of sources (S or A) to advisors (1 or 2) were randomised at the individual level and independently of each other.

The participants' goal in each round was to guess the triangle's colour. Participants knew the prior and chose to receive information either from advisor 1 or from advisor 2. Each advisor was represented by two urns, one for each state, with the colour composition of the twenty balls in an urn reflecting the corresponding source's precisions. For example, consider the task of guessing between blue and red as the triangle's possible colours. In the case of the advisor representing the asymmetric source, the urn corresponding to the case when the triangle was blue contained ten blue balls and ten red balls ($p_A^b = 0.5$), and the urn corresponding to the red triangle contained 19 red balls and one blue ball ($p_A^r = 0.95$).

After choosing the advisor, a signal was generated by drawing a ball from the chosen advisor's urn corresponding to the triangle's true colour. To continue the example above, if the triangle's colour was blue and the subject chose the asymmetric source, the ball was drawn from the urn containing ten blue balls and ten red balls. If the colour of the triangle was red, the ball was drawn from the other urn. Participants did not observe the signal but made two separate guesses, one for each signal.[8] The end-of-round feedback included the triangle colour, the generated signal, the participant's choice given that signal, and their payoff, equal to 100 if they guessed the triangle colour correctly, or zero otherwise.

Second and third parts. The implicit assumption in the predictions of section 3 is that subjects are Bayesian expected-utility maximisers. In these parts, we elicited participants' beliefs to establish whether mistakes in belief updating can explain deviations from the benchmark predictions. In part two of the experiment, we elicited the posteriors for every prior, source, and signal encountered 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.[9] We elicited beliefs with the binarised scoring rule proposed by Hossain and Okui (2013), as implemented in Wilson and Vespa (2016). The order of priors, sources, and their respective colours was the same as in the first part of the experiment.

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

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 organised 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

[8] See Brandts and Charness (2000, 2011) for a discussion on differences between the strategy method and direct-response method in sequential games.

[9] 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.

to pounds 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

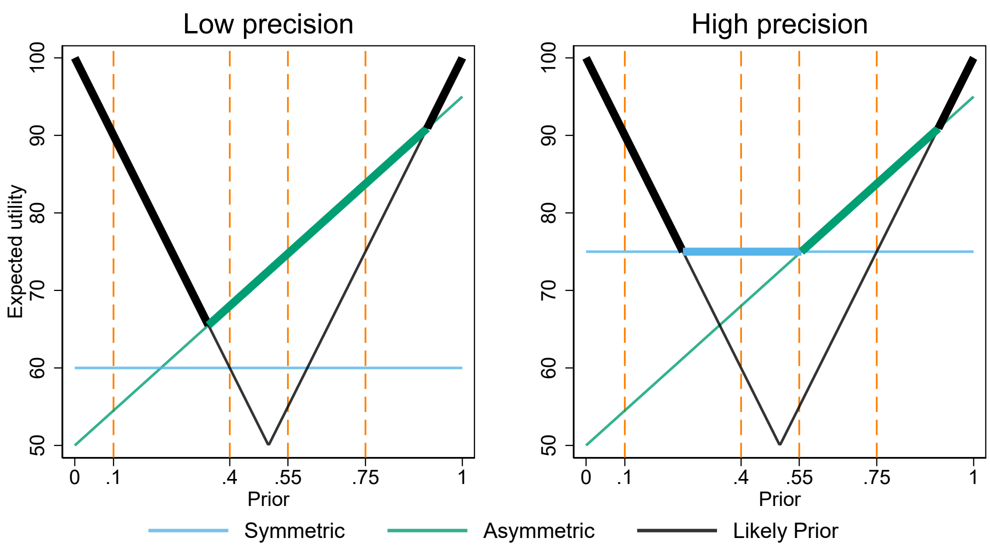

Figure 1: Expected payoff. The vertical dashed lines mark the priors implemented in the experiment. Lines above the "Likely Prior" denote positive information value, while the thicker lines indicate the optimal source choices.

Under the assumption that agents are payoff maximisers, Propositions 1 and 2 provide a series of testable implications regarding use of information and choice of source. Figure 1 depicts the expected utility for different combination of strategies and source of information graphically.[10]

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 $\rho = 0.75$: $\max \rho, 1-\rho$ is higher than the expected payoff of following the source. In these cases, therefore, information should be

[10] Without loss of generality, we set the utility to be equal to the prize; i.e., 100 for winning the prize and zero otherwise.

ignored. Information from source A should be used (source A has value) in both low and high precision treatments at all $\rho \in \{0.4, 0.55, 0.75\}$. Information from source S should be used at $\rho = 0.55$ in both set of treatments, and when $\rho = 0.4$ in the set of high precision treatments.[11]

If agents understand the use of information and follow expected-utility maximization, 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 is enough to determine the choice of information sources.

Prediction about choice of source of information. Source S should be chosen when $\rho = 0.4$ in the high-precision treatments. Information source A should be chosen at all $\rho \neq 0.1$ in the low-precision treatments and when $\rho = 0.75$ in the high-precision treatments. When $\rho = 0.55$ in the high-precision treatments, 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

5.1.1 Choice and Use of Information

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 materialise: source S is chosen over source A even when

[11] Note that there is no prediction regarding use of the S source at $\rho = 0.4$ for $p_S = 0.6$ as one of the signals induces a posterior of 0.5.

[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.

source A has higher value. Moreover, in some instances when S has no value while A does ($\rho = 0.4$ and $p_S = 0.6$, and $\rho = p_S = 0.75$), more subjects choose and use S over those that choose and use 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 (which is not the case with any $\rho \neq 0.1$).

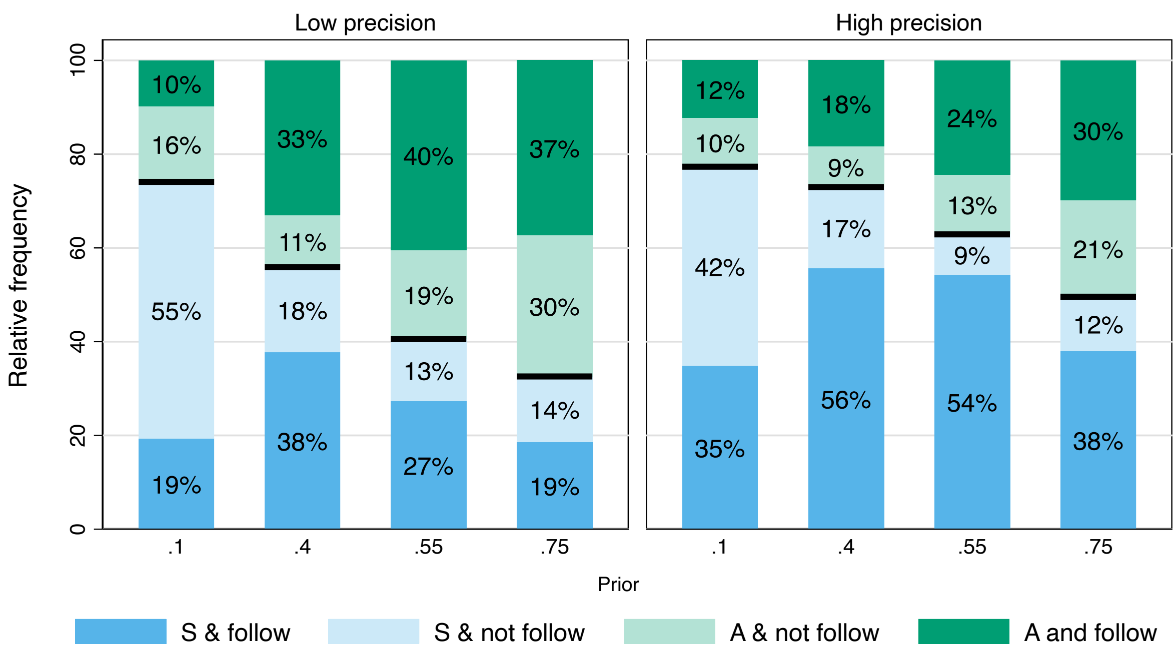

Figure 2: Source choice and decisions.

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 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 are in line with the statistical properties of the information. Notwithstanding, subjects systematically use information*

[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 ($p < .001$).

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 treatments than under the low-precision treatments. Second, when $\rho = 0.75$, only source A should be chosen, but source S is often chosen over it. In particular, source S and source A are chosen with the same frequency in the high-precision treatments. Third, source S is followed more often than 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 even when the sources have no value 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 characteristics of information sources that explain choice and use of information beyond the instrumental value of information.*

5.1.2 Subjective beliefs

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 any gain in expected payoff from following the signal over following the prior, where the expected payoff is calculated based on the subjective probabilities elicited in Parts 2 and 3 of the experiment.

Table 1 provides some evidence that subjective value may rationalise 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.

Thus, the *belief* data shows systematic deviations from Bayesian updating that may rationalise deviations from maximising expected utility in their *choices*. As a first test of this explanation, we construct expected payoff based on the elicited subjective beliefs and look at the rates of payoff-maximising strategies.[14] Across the whole experiment, participants

[14] There are a total of eight possible strategies that result from combining two factors: the choice of source (S or A) and a colour guess for each possible signal realisation.

Table 1: 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 Parts 2 and 3 of the experiment. Percentages indicate the share of cases in which the subjective beliefs indicate no value or positive value conditional on objective value.

chose payoff-maximising strategies in 46.0% of the rounds. If distortions in belief updating are behind this low rate, then the share of optimal strategies should increase when calculated based on the *subjective* beliefs. In contrast, we observe a decrease to 42.1%, indicating that distortions in belief updating fail to rationalise choices.

5.2 An hedonic approach

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 over information characteristics and the observed behaviour 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

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

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 for each possible signal it can generate. We estimate the behavioural 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 that choose the symmetric source, 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]

5.2.1 The role of expected utility

A payoff maximiser is characterised in equation (2) by $\delta = \lambda = \xi = 0$ and by $\pi(s_k)$ equal to the objective value. Column (1) in Table 2 presents the results of the restricted version and Column (2) the results of the unrestricted version of Equation (2), with the parameters scaled by a factor of 100.[16] The results support the qualitative observations in Result 2: the AIC and BIC values indicate that the unrestricted model has a better fit than that of the baseline model.

Table 2: 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. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

[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.

[16] We comment on the other columns in the following sections.

5.2.2 The role of symmetry

The estimated parameters in (2) reflect the cost in terms of the probability of winning the prize that individuals are willing to pay for receiving information from the symmetric source or for following the signals. The results show that participants tend to choose the symmetric source even if by doing so they lose an estimated 8.2 percentage points of winning the prize. Furthermore, after receiving a signal, participants 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 a signal from the symmetric source than a signal from the asymmetric source.[17]

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

5.2.3 The role of subjective value

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 Equation (2) with the subjective expected payoff $\pi^s(s_k)$ calculated by substituting the objective probabilities with the elicited beliefs. If distortions in subjective probabilities explain the patterns in Result 3, the estimates for δ and λ should go to zero once the estimation accounts for these distortions.[18] The results, presented in Column (3) of Table 2, are quite the opposite, as the fit of the model decreases compared to the main model.

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

5.3 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.

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

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

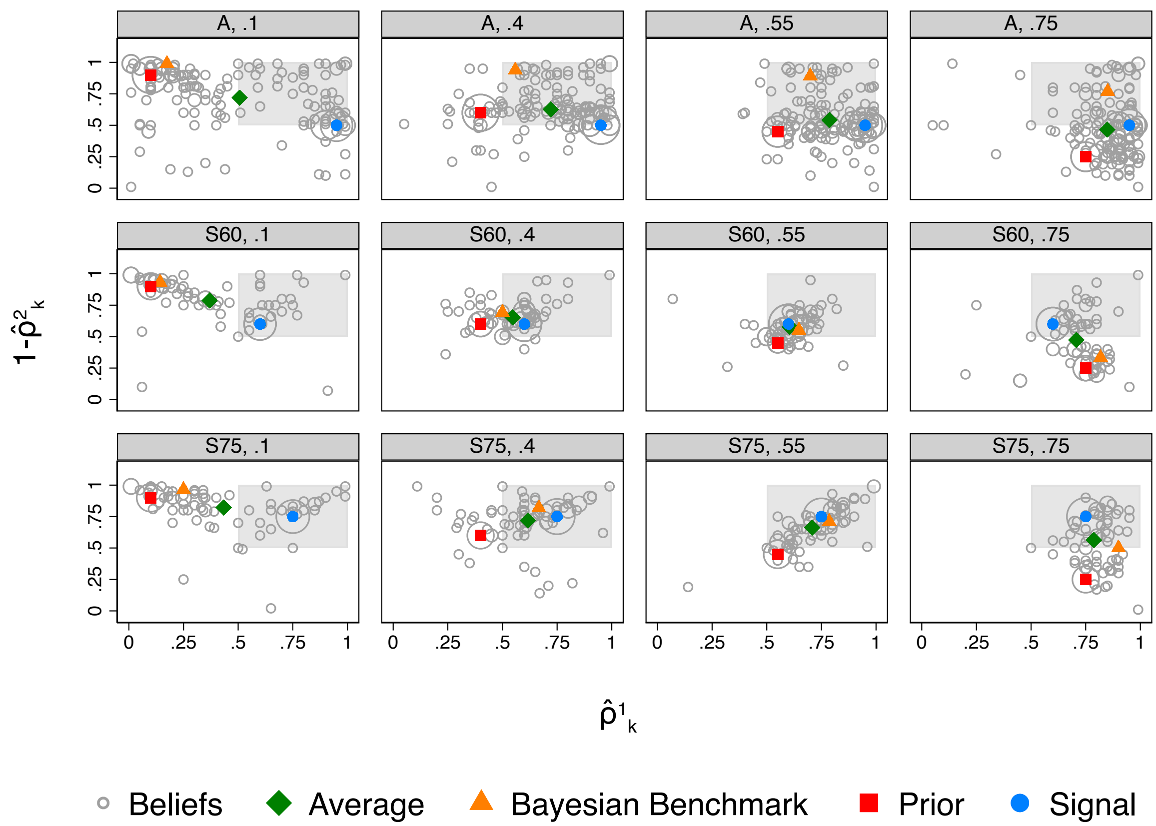

Figure 3: Elicited posteriors.

Many of the subjective posteriors are around the signal odds (circles), 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., Kahneman and Tversky, 1973; Bar-Hillel, 1980; Grether, 1980; Kahneman and Tversky, 1996). Note also that the opposite behaviour is almost as prevalent in the data: many of the subjective posteriors are centered around the priors. We refer to such behaviour below as *Information Neglect* (IN) to contrast with BRN.

We estimate the following generalised Bayes' Rule model based on Grether (1980, 1992) using the subjective posterior data from Part 2 of the experiment:

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

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.[19] The parameter α indicates assigning a high probability to the first colour (the colour 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. Bayes' Rule is a special case of the model, where $\beta_1 = \beta_2 = 1$ and $\alpha = 0$. 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 behaviour, respectively.[20]

Model 1 in Table 3 presents the estimation results based on OLS regression with standard errors clustered on subjects.[21] The estimate for β_1 is higher than 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. As expected, α is only significantly positive for the asymmetric source. Nevertheless, 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

[19] Experimental papers using different versions of this model include Grether (1992), Holt and Smith (2009), Ambuehl and Li (2018), and Coutts (2019).

[20] The literature has referred to generated posteriors that are less extreme than the Bayesian posteriors as "conservatism" (Edwards, 1968; Clippel and Zhang, 2022; Möbius et al., 2022). In Grether's (1980) framework, if $\alpha = 0$ and the priors are uniform, for example, then $\beta_1 < 1$ induces conservatism.

[21] The results are essentially identical for alternative specifications, including fixed- and random-effects for subjects.

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.

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 in the aggregate. The BRN behaviour 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.

5.3.1 Individual heterogeneity

We next estimated the model independently for each participant. Figure 4 plots the individual estimations for (β_1, β_2).[22] 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 underweighting of information, and observations below the $\beta_2 = 1$ indicate underweighting of priors in belief formation.

[22] Only three out of the 190 participants fall out of the region depicted in the figure.

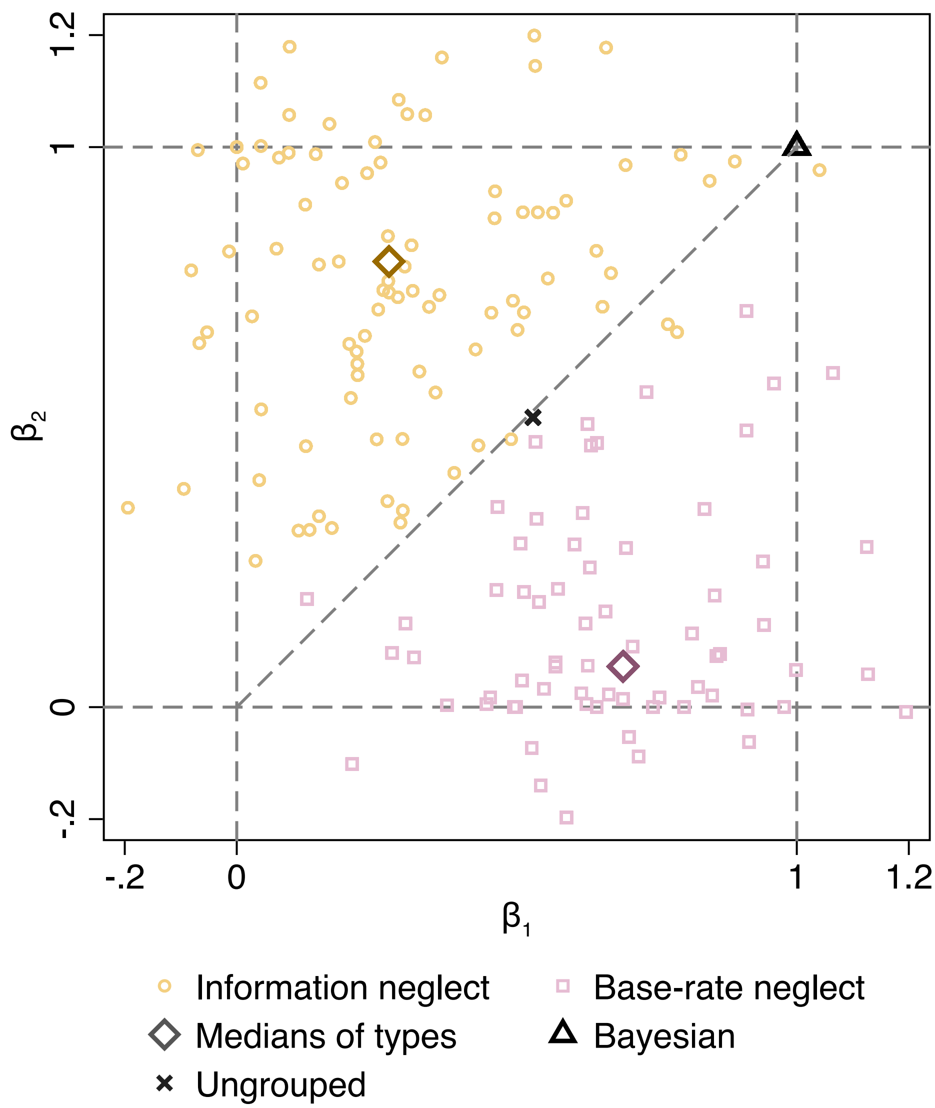

Figure 4: Individual parameters.

Given the substantial heterogeneity in the parameters, we ran a series of k-means cluster analyses to explore whether participants can be categorised based on their β_1 and β_2. The analyses support the existence of two distinct types.[23] Figure 4 plots the two types as squares and circles, respectively. The large diamonds mark the median parameters for the two types.

To characterise the way in which the two types process information differently, we estimated Model 3 in Table 3. This model allows the parameters to vary by source and type. While both types underweight priors, one type overweights information over the prior—reflected in a high ratio β_1/β_2—whereas the other type overweights the prior over information—reflected in a low ratio β_1/β_2. Accordingly, we refer to the two types as *Base-Rate Neglect* (BRN) and Information Neglect (IN), respectively. The BRN type, on average, respond almost exclusively to the signal and not to the prior. In contrast, the IN type underweights both the prior and the signal, but assigns a higher weight to the prior.

Result 6. *Although all subjects underweight priors, there are two distinct types. Approximately half of participants heavily overweight information over priors with the rest 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 the participants' belief updating process correlates with their intrinsic preferences for symmetry and overuse of information. Columns (4) and (5) in Table 2 present separate estimations for the two types. The parameter δ is fairly similar across types and statistically significant for both. The value and significance of the parameter λ, however, vary dramatically between types. Compared to the estimate of 13.44 for the full sample, the estimate of λ for the IN subjects drops to a non-significant 3.26. In contrast, the estimate for the BRN subjects is 28.04. Thus, the people who ignore the prior when reporting posteriors tend to neglect the fact that following the prior and ignoring the signal may be optimal when making decisions.[24]

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.*

Although this result may be striking at first sight, it reflects certain consistency between two independent tasks: beliefs formation (second and third part of our experiment) and use of information (first part of our experiment). If subjects follow their posteriors when using

[23]The resulting types are robust to variations in the initial means and to including the α estimate. In contrast, categorisation 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 $2,101$.

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

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.

5.3.2 The role of symmetry revisited

Using the parameter estimates from Table 3, we computed the implied posterior beliefs for each type and each source to provide stronger evidence on the role of symmetry. Table 4 presents the results of this exercise. While the implied beliefs of the average IN type vary across the different treatments, the implied beliefs of the average BRN type are fairly constant. Two observations are readily apparent with regard to the BRN type. First, the implied posteriors are such that all information sources have value for all priors. Second, these posteriors also mean that the asymmetric source always has higher value than the low-precision symmetric source.[25]

Figure 5 shows the share of times that participants chose each source by treatment and prior conditional on type. BRN subjects use information more than IN subjects for all sources and priors. Furthermore, they choose the symmetric information with low precision over the asymmetric source too often.

Result 8. *BRN subjects choose the symmetric source over the asymmetric inefficiently, revealing an intrinsic preference for this characteristic.*

5.3.3 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 realisation in one direction pushes posteriors away from the priors in the opposite direction. We refer to this phenomenon as *contrarian updating*.[26]

For a base-rate neglect agent facing information from a symmetric source, contrarian updating is easy to conceptualise. The posterior of a *pure* BRN agent facing a symmetric source is equal to the signal precision. In particular, for the high-precision symmetric source, 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 leads the participant's subjective posterior for the blue state to (incorrectly) *decrease* from 0.9 to 0.75. In

[25]The implied posteriors do not rank the asymmetric source and the high-precision symmetric source without further assumptions on the way in which subjects aggregate the posteriors.

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

Table 4: Implied Posterior Beliefs for average types.

	Information Neglect					
	S Low Precision		S High Precision		Asymmetric	
Prior	b signal	r signal	b signal	r signal	b signal	r signal
0.10	**0.109**	0.161	0.078	0.218	**0.119**	0.291
0.40	0.351	0.457	0.271	0.552	0.385	0.655
0.55	0.472	0.582	0.380	0.670	0.512	0.761
0.75	0.653	**0.745**	0.563	0.810	0.693	0.872
	Base Rate Neglect					
	S Low Precision		S High Precision		Asymmetric	
Prior	b signal	r signal	b signal	r signal	b signal	r signal
0.10	**0.323**	0.531	**0.186**	0.703	**0.287**	0.876
0.40	0.379	0.591	0.226	0.752	0.324	0.894
0.55	0.399	0.611	0.241	0.767	0.337	0.899
0.75	0.429	**0.640**	0.264	0.788	0.356	0.907

Notes: This table presents the implied posterior probability that the state is R for the two different types. We calculated the posteriors for each type based on the estimated models in Table 3. A source has (perceived) instrumental value if the posterior after a b signal is smaller than 0.5 and when the posterior after a r signal is larger than 0.5.

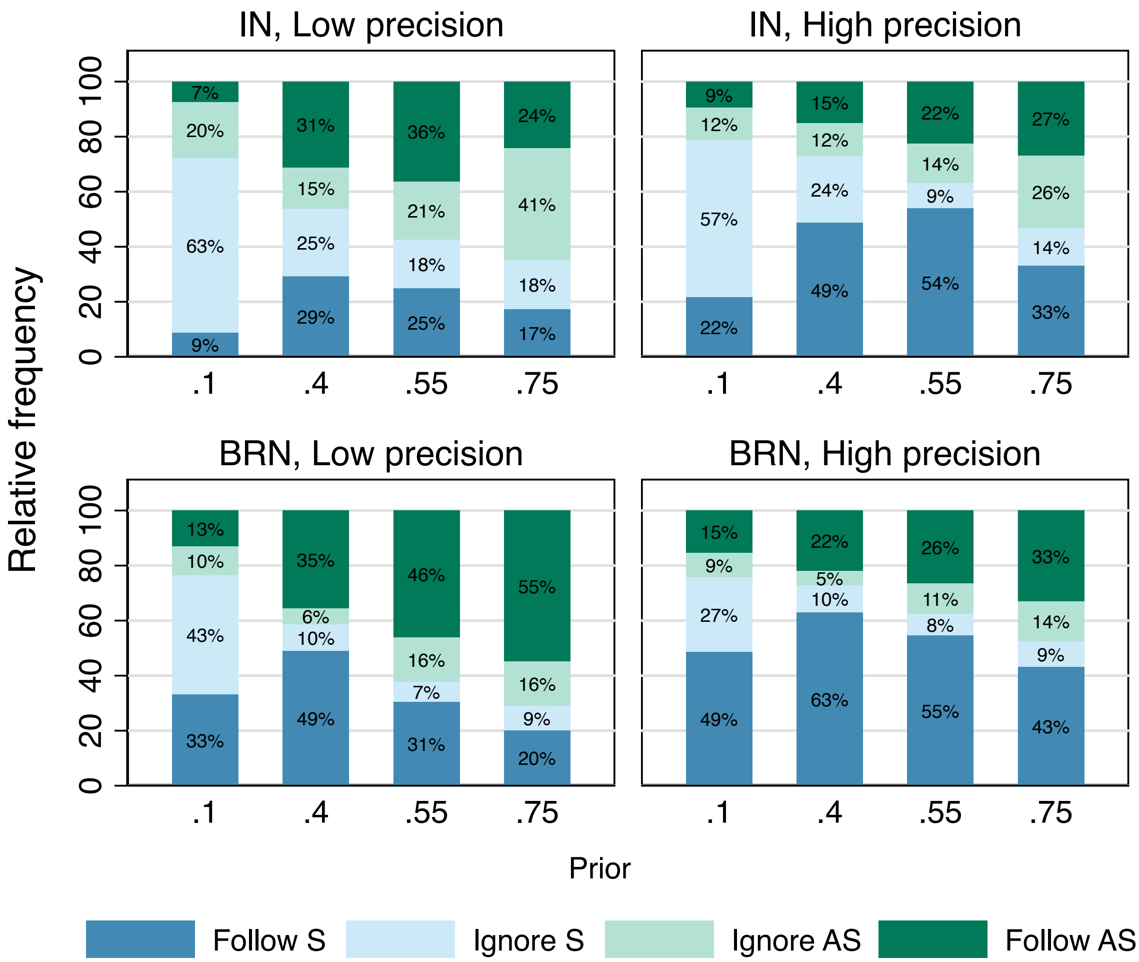

Figure 5: Source choice and decisions by type.

Table 5: 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: Proportion of reported posteriors exhibiting contrarian updating. Numbers in bold indicate situations where the posterior "moves" in the wrong direction.

general, contrarian updating due to base-rate neglect is more likely to occur whenever the prior is more extreme than the precision.

Both BRN and IN representative types show instances of contrarian updating in our experiment. There are seven situations in our experiment in which the implied posteriors of the average types exhibit contrarian updating. These appear in bold in Table 4 as posteriors that are higher than the prior after a b signal and posteriors that are lower than the prior after an r signal. We present the rates of contrarian updating for the two types across the experiment in Table 5. In relative terms, it becomes evident that the treatments identified theoretically as contrarian updating instances in BRN types exhibit the highest frequency in the data.

Result 9. *Subjective posteriors exhibit substantial contrarian updating, which is inconsistent with value-based models but consistent with the predicted behaviour by BRN representative type.*

6 Conclusion

In this paper, we study how people choose between information sources and use the acquired information. We find systematic deviations from the optimal strategies. Participants systematically, (i) choose a symmetric information source over an asymmetric one, and (ii) overuse information from both sources. Subjective beliefs cannot explain 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 (Wolfe and Fisher, 2013; Vartanian et al., 2018), 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.

The posterior beliefs of BRN individuals are fairly stable across priors. In particular, BRN participants receiving information from a symmetric source 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 BRN heuristic leads to a systematic order between sources leading to stronger evidence of intrinsic preferences for symmetry: the low precision symmetric source is always dominated by the asymmetric source for BRN subjects who, nevertheless, choose and use the low precision symmetric source.

Although our experiment is presented in a purely statistical fashion, our results shed light on the current debate on biased news sources. While symmetry is a superficial characteristic in statistical terms, it is sometimes praised as a feature of information as it appears in the social world. In the debate on public discourse, for example, the virtues and flaws of *balancedness* are discussed in terms of substance and context (Tuchman, 1972) even when it leads to distorting the facts. M. T. Boykoff and J. M. Boykoff (2004), in the context of the climate change debate, argues that the discrepancy between the scientific consensus and the public perception can be attributed to news outlets giving "both sides" equal weight while reporting on the issue. Clarke (2008), in the context of the link between autism and vaccines, similarly blames the desire for balance, pointing out how it allowed the discredited study Wakefield et al. (1998) to survive.[27] We provide evidence for a preference based demand for

[27]See also Dearing (1995) for other instances of discrepancies between scientific consensus and news media reporting.

symmetry—even in a context-free environment—and its effects on decision making. We confirm that subjects value information sources that present states and actions in a superficially similar way, independently of the actual informational content and show that this demand for "balancedness" stems from an inherent preference for this characteristic.

References

Ambuehl, Sandro and Shengwu Li (2018). Belief updating and the demand for information. *Games and Economic Behavior* 109, pp. 21–39.

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.

Boykoff, Maxwell T and Jules M Boykoff (2004). Balance as bias: global warming and the us prestige press. *Global environmental change* 14(2), pp. 125–136.

Brandts, Jordi and Gary Charness (2000). Hot vs. cold: sequential responses and preference stability in experimental games. *Experimental Economics* 2, pp. 227–238.

— (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics* 14, pp. 375–398.

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.

Clarke, Christopher E (2008). A question of balance: the autism-vaccine controversy in the british and american elite press. *Science communication* 30(1), pp. 77–107.

Clippel, Geoffroy de and Xu Zhang (2022). Non-bayesian persuasion. *Journal of Political Economy* 130(10), pp. 2594–2642.

Coutts, Alexander (2019). Good news and bad news are still news: experimental evidence on belief updating. *Experimental Economics* 22(2), pp. 369–395.

Dearing, James W (1995). Newspaper coverage of maverick science: creating controversy through balancing. *Public Understanding of Science* 4(4), p. 341.

Edwards, Ward (1968). Conservatism in human information processing. *Formal representation of human judgment*.

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.

Guan, Menglong, Ryan Oprea, and Sevgi Yuksel (2023). Too much information.

Holt, Charles A and Angela M Smith (2009). An update on bayesian updating. *Journal of Economic Behavior & Organization* 69(2), pp. 125–134.

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.

— (1996). On the reality of cognitive illusions. *Psychological review* 103(3), pp. 582–596.

Masatlioglu, Yusufcan, A. Yeşim Orhun, and Collin Raymond (2023). Intrinsic information preferences and skewness. *American Economic Review* 113(10), pp. 2615–2644.

Möbius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat (2022). Managing self-confidence: theory and experimental evidence. *Management Science* 68(11), pp. 7793–7817.

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.

Tuchman, Gaye (1972). Objectivity as strategic ritual: an examination of newsmen's notions of objectivity. *American Journal of sociology* 77(4), pp. 660–679.

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

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

Vartanian, Oshin, Erin L Beatty, Ingrid Smith, Kristen Blackler, Quan Lam, Sarah Forbes, and Wim De Neys (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.

Wakefield, Andrew J., Simon H Murch, Andrew Anthony, John Linnell, David M Casson, Mohsin Malik, Mark Berelowitz, Amar P Dhillon, Michael A Thomson, Peter Harvey, et al. (1998). Retracted: ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children. *The Lancet* 351(9103), pp. 637–641.

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 triangle is Blue (Prob 40%) — The triangle is Red (Prob 60%)

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.

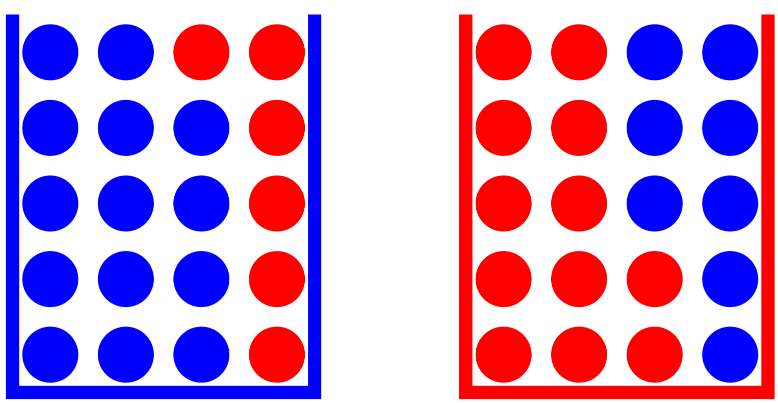

Blue urn — Red urn

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.