Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

The Status Quo and Belief Polarization of Inattentive Agents: Theory and Experiment

Vladimir Novak, Andrei Matveenko, Silvio Ravaioli

Working Paper n. 674

This Version: February, 2021

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

The Status Quo and Belief Polarization of Inattentive Agents: Theory and Experiment

Vladimír Novák*[1], Andrei Matveenko[2], and Silvio Ravaioli[3]

[1]Department of Economics and IGIER, Bocconi University, via Roentgen 1, 20136 Milan
[2]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, 1307 Copenhagen
[3]Department of Economics, Columbia University, 420 West 118th Street, 10027 New York

February, 2021

Abstract

We show that rational but inattentive agents can become polarized, even in expectation. This is driven by agents' choice of not only how much information to acquire, but also what type of information. We present how optimal information acquisition, and subsequent belief formation, depends crucially on the agent-specific status quo valuation. Beliefs can systematically update away from the realized truth and even agents with the same initial beliefs might become polarized. We design a laboratory experiment to test the model's predictions; the results confirm our predictions about the mechanism (rational information acquisition) and its effect on beliefs (systematic polarization).

Keywords: polarization, beliefs updating, rational inattention, status quo, experiment.

JEL codes: C92, D72, D83, D84, D91.

*Corresponding author. E-mail addresses: vladimir.novak@unibocconi.it (V. Novák), anma@econ.ku.dk (A. Matveenko), sr3300@columbia.edu (S. Ravaioli). Many thanks to Filip Matějka for invaluable advice and crucial help in setting up the framework employed in this paper, and to Mark Dean for continuous support on this project. We are grateful for comments from Sandro Ambuehl, Andrew Caplin, Alessandra Casella, Yeon-Koo Che, Alexander Nesterov, Pietro Ortoleva, Jacopo Perego, Debraj Ray, Avner Shaked, Jakub Steiner, Michael Woodford, Jan Zápal, and numerous seminar participants for helpful comments and suggestions. This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreements No. 678081 and No. 740369) and we acknowledge the financial support of grants from the Columbia Experimental Laboratory for Social Sciences. Vladimír Novák was supported by the Charles University in Prague (GA UK projects No. 197216 and No. 616219), and by the H2020-MSCA-RISE project GEMCLIME-2020 GA. no 681228. The authors affirm no financial interest in the outcomes of the experiments detailed herein. All data were collected with the approval of the Columbia University Institutional Review Board (protocol AAAS5801).

1 Introduction

Many decisions, from democratic votes to daily choices,[1] can be characterized as a binary choice between preserving the status quo policy, with a known outcome, and choosing a new policy, the outcome of which is uncertain. Consider the 2016 Brexit referendum or the European climate law to achieve climate-neutrality by 2050 as illustrative examples. The consequences of leaving the EU or of choice adopting the climate law were uncertain at the moment of choice, whereas the opposite – to preserve the status quo – appeared to be more certain in its consequences. Public opinion surveys suggest that such binary policy choices are often associated with an increase in belief polarization in society.[2] More generally, an increase in opinion polarization is a widely discussed phenomenon in both the academic literature[3] and public discourse.

In this paper, we present a new model of belief polarization. The key mechanism is that the relative position of the status quo determines the evolution of opinions through the choice of information structure, which may in turn lead to systematic polarization even between agents with the same prior beliefs. We design a laboratory experiment in which we manipulate the value of the status quo action in order to test the theoretical predictions. The results confirm both predictions about the key mechanism of information acquisition and the qualitative effect on beliefs, although the overall effect is mitigated by behavioral factors. The main mitigating channel is a preference for simple signal structures, a result that generalizes the well-known preference for certainty. Other factors, like risk preference and subjective beliefs, are not sufficient to explain the observed deviations from the predictions.

The key aspect of our approach is that we concentrate our analysis on the role of the value of the status quo in the process of information acquisition and, subsequently, on the possibility of polarization of opinions. Unlike other papers on polarization, we show that polarization can appear even in expectations and without any exogenously imposed biases.

We model the agent to be rationally inattentive, following Sims (1998, 2003), which allows

[1] See Section 2 for an example of such a situation.

[2] The British society a few months after the Brexit referendum was even more polarized than on the referendum day (Smith, 2019). The polarization of the climate change news is documented, for instance, by Chinn, Hart and Soroka (2020); Leiserowitz et al. (2019).

[3] See, for instance Poole and Rosenthal (1984); McCarty, Poole and Rosenthal (2008); Gentzkow, Shapiro and Taddy (2016).

us to account for endogenous information acquisition without imposing any restrictions on the agent's learning process. Since information is plentiful but attention is scarce, the agent chooses to learn the essential pieces of information for her decision problem, whether to adopt the new policy or to preserve the status quo. Thus, the agent endogenously partitions states into groups, separating the states in which payoffs of the new policy are higher or lower than under the status quo, and acquires costly information just to identify which of the two groups contains the realized state. We show under which conditions and in which states of the world the agent systematically updates her belief away from the true realization, even in expectations over the received information. When comparing the choices of agents with different status quo valuations, we demonstrate a possibility of belief polarization, even for the agents with the same prior expected belief about the payoff of the new policy.

We begin our paper with a simple illustrative example. The information acquisition strategies in the example are very limited, but they provide an intuition for our main result – beliefs polarization – and the key mechanism behind it – the state pooling effect (avoiding redundant information about states associated with the same optimal action). In Section 3 we introduce a static decision problem with a rationally inattentive agent, n states and two actions, in the manner of Matějka and McKay (2015). However, in contrast to Matějka and McKay (2015), our focus is on the evolution of beliefs. We study how the mean of beliefs about the payoff of the new policy evolves over the optimal signals. As a result of the state pooling effect, we observe that the agent might update her own belief about the expected payoff of the new policy in the opposite direction from the realized value. We formulate a simple criterion to identify when the agent updates systematically in the wrong direction (Theorem 1). We also show that two agents might become polarized when they either differ in the valuation of the status quo or in their prior expected beliefs about the payoff of the new policy (Theorem 2). Further, we illustrate how behavior resembling over-optimism[4] and over-pessimism can appear as a result of the agent's inattentiveness.

The theoretical results presented rely on several assumptions about a decision maker's preferences and updating process that have been challenged by previous experimental find-

[4]The important consequences of over-optimism can be seen, for instance, in Beaudry and Willems (2018), who show that recessions, fiscal problems, and Balance of Payment-difficulties are more likely to arise in countries where past growth expectations have been overly optimistic.

ings.[5] In Section 4 we introduce an experiment designed to test our theoretical results; in particular the state pooling effect and the presence of belief polarization in expectations. The subjects are presented with a binary choice, and can acquire instrumentally valuable information from advisors (signal structures) before making their decision. In the main task, participants make an information choice followed by an action choice. For the information choice, they are presented with a pair of advisors and can select only one of them. After that, they indicate the chosen action (status quo or new policy) conditional on the observed signal. Our key manipulation consists in varying the value of the status quo, and we expect this to be sufficient to revert the choice of optimal advisor. In two separate tasks we elicit, for each participant, the subjective beliefs about the likelihood of a signal realization, and each state (conditional on the realized signal).

Section 5 presents our main experimental results. Participants do react to the value of the status quo, as predicted by our theoretical model, and display preference for state pooling information structures. Importantly, we document the belief polarization in our laboratory setting. The magnitude of the polarization is mitigated with respect to the predicted magnitude, and we investigate the main behavioral channels that can explain the deviation. We show that the evaluation of information structures is affected by non-instrumental characteristics; in particular, preference for advisors that are characterized by certainty (degenerate posteriors) or simplicity (fewer possible outcomes). In Section 6 we discuss the robustness of our experimental findings and Section 7 concludes the paper with a discussion of the limits of our model and experiment, and directions for future research.

Related literature. This paper contributes to the theoretical and experimental literature on information acquisition, as well as to the broader literature on belief polarization. We highlight the role of rational endogenous information acquisition about the instrumental choice in belief polarization. In particular, modelling the agent as rationally inattentive relieves us of the need to assume exogenously given biases[6] or bimodality of preferences (Dixit

[5]The experimental literature contains systematic evidence of deviation from theoretical predictions in several domains. See, for example Holt and Laury (2002) (risk preferences), Kareev, Arnon and Horwitz-Zeliger (2002) (subjective beliefs), and Ambuehl and Li (2018) (preference over signal structures).

[6]Gerber and Green (1999) review the literature that invokes some biases in learning or perception in order to modify Bayesian updating.

and Weibull, 2007), which are common in the preceding literature. This in turn allows us to move in a new direction, away from findings that the beliefs of Bayesian agents would converge over time and that they will almost surely assign probability 1 to a true state (Savage, 1954; Blackwell and Dubins, 1962). The closest theoretical paper to ours is Nimark and Sundaresan (2019), which studies the question of how inattentiveness can lead to persistent belief polarization. In our model, agents with the same prior beliefs might become polarized in expectation, whereas in Nimark and Sundaresan (2019) two agents with the same prior beliefs always choose the same signal structures and can become polarized only when they receive sufficiently different signal realizations. We consider a multiple states environment with full flexibility in the shape of the received signal, which allows us to discover a state pooling effect that cannot arise in the environment of Nimark and Sundaresan (2019).[7] We contribute to the polarization literature also by providing rational decision theory mechanism that is not driven by the media competition (Perego and Yuksel, 2018; Bernhardt, Krasa and Polborn, 2008) or necessarily partisan conflict (Prior, 2013; Oliveros and Várdy, 2015; Gentzkow, Shapiro and Taddy, 2016) .

Naturally, this paper also adds to the rational inattention literature[8] by studying the evolution of beliefs alongside the manifestation of the crucial implications of incorporating the safe option into the choice set. The main mechanism of our paper - the state pooling effect - is connected with the presence of the status quo policy, a particular form of reference point. However, in contrast with the rich theories of individually determined reference points (see, e.g., Kahneman and Tversky, 1979; Kőszegi and Rabin, 2006; Guney, Richter and Tsur, 2018), the status quo is exogenously given. At the same time, we neither assume the status quo bias (Samuelson and Zeckhauser, 1988) nor aim to study formation of such bias (Ortoleva, 2010; Masatlioglu and Ok, 2014), but our focus is on the evolution of beliefs and not necessarily the chosen actions. Hence, our paper complements the studies documenting how current economic standing, in our setting represented by the status quo policy valuation, influences the action taken by citizens. For instance, consider Fetzer (2019) who claims that

[7]They focus on the persistence of the polarization and thus study a two state environment where the agent might receive binary signals and the information structure of the agent is characterized by error probabilities.

[8]A survey of the literature on rational inattention is provided in Maćkowiak, Matějka and Wiederholt (2020). For a posterior-based approach see Caplin and Dean (2015) and a dynamic discrete choice model is presented in Steiner, Stewart and Matějka (2017).

economic losers were more likely to vote for Brexit, and Dal Bó et al. (2018) who connects economics losers with the rise of Swedish radical right.

The state pooling effect that emerges as a result of the model can be viewed as an endogenously arising categories formation mechanism, a specific decision making heuristic (Gigerenzer and Gaissmaier, 2011), and creates a link to the settings with at most two states that are then often studied in dynamic settings (Che and Mierendorff, 2019). This also differentiates us from the previous literature that assumes exogenously given categories (see, e.g., Suen, 2004; Manzini and Mariotti, 2012; Maltz, 2020).

Finally, the paper adds to the experimental literature. First, in contrast to Charness, Oprea and Yuksel (2018) we focus on variation in the value of the status quo and not on the variation in prior beliefs. Second, our findings give reason for caution for empirical and experimental work inferring preference for information. In particular, this paper provides a disciplined model that suggests how the preference for skewed information might crucially depend on the value of the status quo and thus provides an important channel that is missing in the research on whether people prefer negatively or positively skewed information, e.g. Masatlioglu, Orhun and Raymond (2017). Third, we replicate and extend previous evidence of certainty preference (Ambuehl and Li, 2018) in a three states environment.

2 Example

In this section, we illustrate the logic of our result in a simple example, in which the information acquisition strategy is highly restricted, but still allows us to demonstrate the main mechanism behind the results of the paper: the state pooling effect and its influence on the belief updating process.

Setup. Consider two risk-neutral payoff-maximizing agents, Alice and Bob (A and B). Each agent independently faces a choice between a currently implemented policy (status quo) and a new policy. We interpret these policies broadly. Examples might include staying home (status quo) versus going to a theatre (where a movie has an uncertain quality), voting whether the UK should stay in or leave the EU during the Brexit vote, selecting in a referendum if the country should adopt the climate neutrality laws or not, among many others.

The outcome of the new policy is uncertain and depends on the realized state of the world. There are three possible realizations of the state of the world $s \in \{\text{bad, medium, good}\}$. In the state s, the payoff of the new policy is $v_s \in \mathbb{R}$ and the status quo generates a payoff R. The agents might differ in their valuation of the status quo payoff R_A and R_B, and in their prior probabilities $g_{A,s}$ and $g_{B,s}$ with which each state s is realized for Alice and Bob, respectively. However, in this illustratory example we assume that both agents have a uniform prior belief, that is $g_{A,s} = g_{B,s} = \frac{1}{3}$, $\forall s$, and the assumed payoffs are summarized in Table 1. Thus, the prior expected value of the new policy for both agents is $\mathbb{E}v \approx 0.53$.

	New policy	**Status quo**	
s	v_s	R_A	R_B
bad	0	0.45	0.7
medium	0.6	0.45	0.7
good	1	0.45	0.7

Table 1: Payoff structure assumed in the example.

Signals. Both agents have an opportunity to learn about the realized state of the world, but for the purpose of this example we restrict their learning possibilities. Specifically, assume that they can ask any "yes/no" question of the form "Is the realized state i?" for some strictly positive cost, where $i \in \{\text{bad, medium, good}\}$.

Information acquisition strategy. *Alice's* valuation of the status quo ($R_A = 0.45$) implies that she prefers to choose the new policy when the realized state is medium or good. Therefore, she does not need to distinguish the states "good" and "medium." Thus, her optimal information acquisition strategy in this example is to ask "Is the realized state bad?". Each answer[9] is then directly connected with the action. If the answer is "Yes" it is optimal to stick with the status quo, if the answer is "No" the optimal action is to choose the new policy. For *Bob*, who values the status quo highly, $R_B = 0.7$, the optimal information acquisition strategy is to ask a question "Is the realized state good?". This is the case because only when the realized state is good does he prefer to choose the new policy over the status quo. Note that neither Alice nor Bob would ask more than one question,

[9]We assume that answers are always truthful.

because they would have to pay an additional cost for a second question, but it would not change their chosen action. The fact that the agents do not distinguish some states of the world after optimal acquisition of information, and pool states associated with the same action together, is referred to as a *state pooling effect.*

Belief. We can put ourselves in the shoes of an observer who investigates how agents change their expected payoff of the new policy after they have acquired the desired information. The observer also knows that the realized state is medium. In such a situation both agents have received the answer "Not bad" and "Not good," respectively. Afterwards, Alice knows that the true state is either medium or good, whereas Bob knows the state is bad or medium. Therefore, Bob's posterior expected payoff from the new policy is lower than the prior expected payoff and also lower than the true payoff of the new policy. Alice is facing the symmetrical case, with the posterior expected payoff from the new policy being both higher than the prior one and than the true payoff. As a consequence of the endogenously chosen information, agents update their expected beliefs in opposite directions. In the following sections we show that this result can occur in a more general setup, and in particular when agents can freely choose the form and precision of the signal structure.

Summary of the intuition. As we illustrate by this example, the agents endogenously partition states into two groups, separating the states in which payoffs of the new policy are higher or lower than under the status quo. We denote this partitioning strategy as a state pooling effect. The agent acquires costly information to identify which of the two groups contains the realized state. As a result, the agent might overestimate the probability of the states that belong to the same group as the realized one. One agent can update the expected payoff of the new policy in the opposite direction with respect to its true payoff, and two agents with different status quo payoffs can update their beliefs in opposite directions, so their beliefs become *polarized.*

3 The model

In this section, we describe the general case of the agent's decision problem, introduce a methodology for assessment of beliefs evolution, and present the main theoretical results. The structure of this section is as follows. In the subsections 3.1 and 3.2, we describe the agent's problem, which is a special case of the agent's problem from Matějka and McKay (2015). Subsections 3.3 and 3.4 present our main results about the evolution of the beliefs given the true state of the world. Subsection 3.5 discusses polarization of rationally inattentive agents, that is, polarization of agents that choose any form and precision of the signal structure optimally given their decision problem, with costs based on the Shannon mutual information between prior and posterior beliefs.

3.1 Description of the setup

A single agent faces a problem of discrete choice between two options. The first option, which we refer to as *a new policy,* provides a payoff $v_s \in \mathbb{R}$ that depends on the realized state of the world $s \in S = \{1, \ldots, n\}$, where $n \in \mathbb{N}$, $n \geq 3$. When we say that the state is realized, we mean that the potential outcome of the new policy has become possible to be evaluated. The states are labeled in ascending order $v_1 < v_2 < \ldots < v_n$. The second option, which we refer to as *a status quo policy,* yields a known fixed payoff $R \in \mathbb{R}$. We assume that $v_1 < R < v_n$ in order to exclude trivial cases.[10] That is, there exists a unique $k \in \{1, 2, ..., n-1\}$, such that $v_k \leq R < v_{k+1}$.

The agent is uncertain which state of the world s is going to be realized and we denote her prior belief as a vector of probabilities $\mathbf{g} = [g_1 \;\; g_2 \;\; \ldots \;\; g_n]^T$, where $\mathcal{P}(s = j) = g_j$, $\forall j \in S$; $\sum_{j=1}^{n} g_j = 1$ and $g_j > 0$, $\forall j \in S$. We model the agent to be rationally inattentive in the fashion of Sims (1998, 2003). The agent wishes to select the option with the highest payoff. Prior to making the decision, the agent has a possibility to acquire some information about the actual value of the new policy, which is modeled as receiving a signal $x \in \mathbb{R}$. The distribution of the signals, $f(x, s) \in P(\mathbb{R} \times S)$, where $P(\mathbb{R} \times S)$ is the set of all probability distributions on $\mathbb{R} \times S$, is subject to the agent's choice. Upon receiving a signal, the agent updates her belief using Bayes rule. However, observing a signal is costly and we assume the cost κ to be proportional to the expected reduction in entropy[11] between the agent's prior

[10] If $R \leq v_1$, the safe option is weakly dominated by the risky option, and if $R \geq v_n$ the risky option is weakly dominated by the safe option. In both of these cases the agent does not have incentives to acquire information about the realization of the state of the world.

[11] The entropy $H(Z)$ of a discrete random variable Z with support $\mathcal{Z}$ and probability mass function

and posterior beliefs.

Upon receiving a signal, the agent chooses an action, and her choice rule is modeled as $\sigma(x) : x \to \{\text{new policy}, \text{status quo}\}$. Given the updated belief, the agent chooses the action with the highest expected payoff. The agent's objective is to maximize $\mathbb{E}v_i - \lambda\kappa$, the expected value of the chosen policy less the cost of information. The timing of the decision problem is depicted in Figure 1.

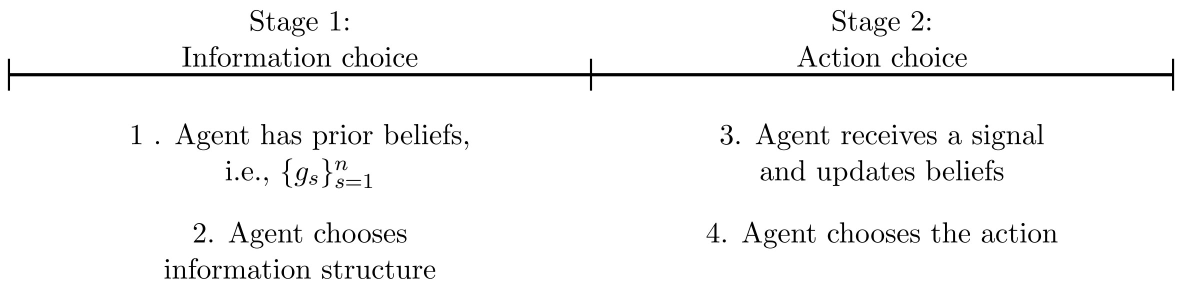

Figure 1: Timing of the events in the problem. The decision problem consists of two stages: an information strategy selection stage and a standard choice under uncertainty stage.

3.2 Agent's problem

The information strategy is characterized by the collection of conditional probabilities of choosing option i in state of the world s : $\mathcal{P} = \{\mathcal{P}(i|s) \big| \; i = 1, 2; \; s \in S\}$, where $i \in \{\text{new policy, status quo}\} = \{1, 2\}$ denotes the option and s is the state. The agent solves:[12]

$$\max_{\{\mathcal{P}(i|s)|i=1,2;\, s\in S\}} \left\{ \sum_{s=1}^{n} \left(v_s \mathcal{P}\left(i = 1|s\right) + R\mathcal{P}(i = 2|s) \right) g_s - \lambda\kappa \right\}, \tag{1}$$

$\mathcal{P}(z) = Pr\{Z = z\}, z \in \mathcal{Z}$ is defined by $H(Z) = -\sum_{z\in\mathcal{Z}} p(z) \log p(z)$. For detailed treatment of entropy, see, for example, Cover and Thomas (2012).

[12] According to Lemma 1 from Matějka and McKay (2015), the choice behavior of the rationally inattentive agent can be found as a solution to a simpler maximization problem that is stated in terms of state-contingent choice probabilities alone.

subject to

$$\forall i: \ \mathcal{P}(i|s) \geq 0 \qquad \forall s \in S \ , \tag{2}$$

$$\sum_{i=1}^{2} \mathcal{P}(i|s) = 1 \qquad \forall s \in S \ , \tag{3}$$

and where

$$\kappa = \underbrace{-\sum_{i=1}^{2} \mathcal{P}(i) \log \mathcal{P}(i)}_{\text{prior uncertainty}} - \sum_{s=1}^{n} \left(\underbrace{-\left(\sum_{i=1}^{2} \mathcal{P}(i|s) \log \mathcal{P}(i|s) \right)}_{\text{posterior uncertainty in state } s} g_s \right) . \tag{4}$$

$\mathcal{P}(i)$ is the unconditional probability that option i will be chosen and is defined as

$$\mathcal{P}(i) = \sum_{s=1}^{n} \mathcal{P}(i|s) g_s, \qquad i = 1, 2.$$

Here κ denotes the expected reduction in entropy between the prior and the posterior beliefs about the choice outcome, $\lambda \geq 0$ is the unit cost of information and thus, $\lambda\kappa$ reflects the cost of generating signals with different precision.

Non-learning areas. An important feature of the solution to the agent's problem is that for the given vector of possible payoffs of the new policy $(v_1, ..., v_n)$, the value of the status quo R and the marginal cost of information λ there exist prior beliefs of the agent for which she decides not to acquire any information. In this case, we say that the agent is in a *non-learning area.* If the agent's prior is such that she decides to acquire at least some information, we say that the agent is in a *learning area.* For such prior beliefs, the unconditional choice probabilities lie in the open interval $(0, 1)$. In order to make the problem meaningful, we require for the rest of the paper that the agent is in the learning area.[13] Consequently, we assume that:

Assumption 1. $0 < \mathcal{P}(i = 1) < 1$.

[13] See Caplin, Dean and Leahy (2019) for the characterization of the necessary and sufficient conditions for solution of the discrete rational inattention problems.

3.3 Description of beliefs evolution

The main aim of this paper is to describe the evolution of the agent's belief, represented by the expected payoff of the new policy, in each state of the world. Let us thus first introduce the main objects used in our analyses.

Prior expected value. The uncertainty in this model is about the realized state of the world and thus about the actual payoff of the new policy. Without the information acquisition stage of the problem the agent would choose the option based on the comparison of the status quo payoff R with the agent's *prior expected value of the new policy* being

$$\mathbb{E}v = \sum_{s=1}^{n} v_s g_s.$$

Posterior conditional expected value. In order to judge how this expected payoff from the new policy changes after the signal is received and the option is chosen, we take the position of an external observer. The observer knows that a realized state of the world is s^* and is interested in the agent's posterior expected belief about the payoff of the new policy v given the realized state s^*. Note that the agent's posterior belief is given by the signal she receives and thus the observer not only wants to know what the expected posterior belief is for a given signal, but is interested in the expected posterior belief about the new policy on average across all possible signals the agent may receive. Since there is a one to one mapping between the selected information structure and consequently chosen action, the posterior expected belief of interest is

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \sum_{i=1}^{2} \left(\sum_{s=1}^{n} v_s \mathcal{P}(s|i) \right) \mathcal{P}(i|s^*),$$

where option $i \in \{1, 2\} = \{\text{new policy}, \text{status quo}\}$. For the optimally behaving rationally inattentive agent, who is solving the problem (1)–(4), the posterior expected belief can be rewritten as:[14]

[14]The derivation of Formula (5) is in Appendix A.

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}. \tag{5}$$

The change in the mean of beliefs about the payoff of the new policy. The primary indicator for the expected belief evolution over the optimal signals that we consider is the change in the mean of beliefs about the payoff of the new policy that can be defined as $\Delta(s^*) = \mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v$. In particular, we are interested in the sign of $\Delta(s^*)$, which informs us whether the posterior expected belief in state s^* is moving from $\mathbb{E}v$ towards v_1 or v_n or stays equal to $\mathbb{E}v$.

Proposition 1. *Given that the realized state of the world is $s^* \in S$, the sign of the change in the mean of beliefs about the payoff of the new policy $\Delta(s^*) = \mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v$ is the same as the sign of $(v_{s^*} - R)$.*

Proof. The proof is presented in Appendix B. □

Proposition 1 significantly simplifies the considerations of the beliefs evolution. At the same time, it also demonstrates the direct link between the value of the status quo R and the updating process.

3.4 Updating in the opposite direction from the realized value

In this section, we show the impact of the value of the status quo on the opinion polarization of inattentive agents.

Definition 1. The agent updates in the *opposite direction from the realized value* v_{s^*} in the state $s^* \in S$, if the condition $(v_{s^*} - \mathbb{E}v) \cdot \Delta(s^*) = (v_{s^*} - \mathbb{E}v)(\mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v) < 0$ is satisfied.

In the following theorem we provide conditions for the presence of the states in which the agent is updating in the opposite direction from the realized value of the new policy.

Theorem 1. *The agent updates in the the opposite direction from the realized value v_{s^*} in the realized state $s^* \in S$, if and only if the inequality $(v_{s^*} - R)(\mathbb{E}v - v_{s^*}) > 0$ holds.*

Proof. We need to show that if $(v_{s^*}-R)(\mathbb{E}v-v_{s^*})>0$ is satisfied, then $(v_{s^*}-\mathbb{E}v)(\mathbb{E}_i[\mathbb{E}(v|i)|s^*]-\mathbb{E}v)<0$ is also satisfied. By Proposition 1, the sign of $v_{s^*}-R$ is the same as the sign of $\mathbb{E}_i[\mathbb{E}(v|i)|s^*]-\mathbb{E}v$, which implies the condition needed. The proof in the opposite direction is analogical. □

Theorem 1 provides a simple criterion for updating in the opposite direction from the realized value of the new policy. The criterion does not require solving the agent's problem and is formulated in the primitives of the model only.

Intuition for the result presented in Theorem 1 is as follows. Due to costly information acquisition, the rationally inattentive agent chooses only the necessary information in order to disentangle whether to select the status quo or the new policy. This leads to the *state pooling effect*, when the agent endogenously divides the states into two categories (the states in which the payoff of the new policy is higher than the payoff of the status quo and the states in which it is lower) and chooses only information that helps to disentangle which of these two categories the realized state s^* is from. Namely, as Proposition 1 states, for all the states s in which $v_s>R$ the expected posterior belief about the value of the new policy is higher than the prior belief; thus all such states are pooled into one category. Similarly, all the states s for which $v_s<R$ are pooled into another category.

It is important to notice that the agent's expected posterior belief is not the same for different realized states, even when such states are from the same category. Notice that so far we have shown that updating in the opposite direction from the realized value of the new policy can occur in some realized state s^*. A natural question arises: How is the difference between the prior and the posterior expected payoff from the new policy $\Delta(s^*)$ influenced by a different realized true state s^*? The answer to this question is provided by the following proposition.

Proposition 2. *The change in the mean of beliefs $\Delta(s^*)$ is an increasing function of s^*.*

Proof. The proof is presented in Appendix C. □

Proposition 2 helps us to identify the set of states $W\subset S$ in which the agent updates in the opposite direction from the realized payoffs v_{s^*} in corresponding states. First, it follows from Proposition 1 that $\Delta(s^* = 1) < 0$ and $\Delta(s^* = n) > 0$. These findings, together with Proposition 2, imply that $\Delta(s^*)$ reaches its minimum in state 1, its maximum in state n and $\Delta(s^*) = 0$ occurs in between. At this point, recall that we have defined state k such that $v_k \leq R < v_{k+1}$. Thus, due to Proposition 1, $\Delta(k) \leq 0$ and $\Delta(k+1) > 0$. The last two inequalities together with Proposition 2 imply that for all $s^* \leq k$ the change of the mean of beliefs $\Delta(s^*) \leq 0$ and that for all $s^* \geq k+1$ holds that $\Delta(s^*) \geq 0$. We know that in states for which the condition $(\mathbb{E}v - v_{s^*}) \cdot \Delta(s^*) > 0$ is satisfied, the agent updates in the opposite direction from the realized payoff of the new policy. Let us assume that the agent's prior expected value of the new policy is $\mathbb{E}v > R$. Then one can see that the agent is updating towards the true payoff of the new policy for all states where $\Delta(s^*)$ is negative. However, updating in the opposite direction from the realized value occurs for all states that have payoffs smaller than $\mathbb{E}v$ and at the same time higher than R (see Figure 2).

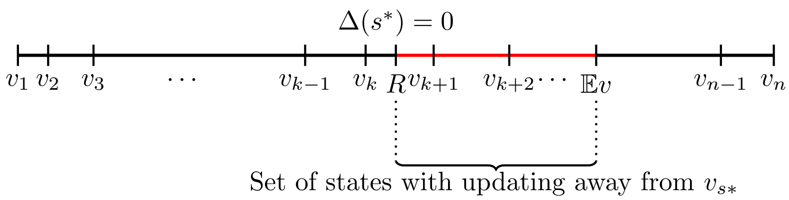

Figure 2: Set of states W (the red part of the line) for $\mathbb{E}v > R$, where the agent updates in the opposite direction from the realized value of the new policy.

If we assume that $\mathbb{E}v < R$ then updating in the opposite direction from the realized value would happen in all states s^* for which it holds that $R > v_s^* > \mathbb{E}v$. We can write that

$$W = \begin{cases} \{s \mid R < v_s < \mathbb{E}v\}, & \text{if } \mathbb{E}v > R, \\ \{s \mid \mathbb{E}v < v_s < R\}, & \text{otherwise.} \end{cases}$$

It is worth noticing that the set of states where the agent updates in the opposite direction from the realized value are those states where payoffs are neither very high nor very low. This has significant implications for predictions when the inattentive agents become polarized. We can also see that the number of states in the set W is determined by the status quo payoff R and by the prior expected value of the new policy.

Symmetry, over-optimism and over-pessimism. From the last observation we can see that even though the value of the status quo R is crucial for the state pooling effect, it is the relation between prior belief $\mathbb{E}v$ and R that determines if the agent would update in the opposite direction from the realized value of the new policy. Importantly, their role is symmetrical. Thus, for a fixed realized state s^* the agent updates in the opposite direction from the realized payoff when $\mathbb{E}v < v_{s^*} < R$ or when $R < v_{s^*} < \mathbb{E}v$. In the former situation the agent has a low prior expected value from the new policy and then updates towards v_1. In the latter situation the prior expected value is quite high and then it is updated upwards towards v_n. Stated differently, in the situation when $\mathbb{E}v < v_{s^*} < R$, the agent is pessimistic about the new policy and consequently becomes even more pessimistic. In the second case, when $R < v_{s^*} < \mathbb{E}v$, the opposite is true. The agent is optimistic, and becomes over-optimistic about the outcome of the new policy.

3.5 Belief polarization

Let us now consider a situation with two agents $j = 1, 2$. The agents are facing the binary choice described in Section 3.2; however, the agents might have (i) different valuations of the same status quo policy R^j and, (ii), different prior beliefs about the value of the same new policy $\mathbb{E}^j v$. The expected posterior belief of agent j about the value of the new policy, conditional on the realized state s^*, is denoted by $\mathbb{E}_i^j[\mathbb{E}(v|i)|s^*]$. The difference between the expected posterior beliefs of agent j in the state s^* and the prior beliefs of agent j is denoted by $\Delta_j(s^*) = \mathbb{E}_i^j[\mathbb{E}(v|i)|s^*] - \mathbb{E}^j v$.

Definition 2. We say that two agents $j = 1, 2$, who are characterized by the pair $(R^j, \mathbb{E}^j v)$ and are choosing between actions $i = \{1, 2\}$, become *polarized in the state* $s^* \in S$ when the following two conditions are satisfied

1. $|\mathbb{E}_i^1[\mathbb{E}(v|i)|s^*] - \mathbb{E}_i^2[\mathbb{E}(v|i)|s^*]| > |\mathbb{E}^1 v - \mathbb{E}^2 v|$.
2. $\Delta_1(s^*) \cdot \Delta_2(s^*) < 0$.

The first condition secures that the expected posterior beliefs in the state s^* of two agents are further apart than the expected prior beliefs are, whereas the second ensures that they update in opposite directions in the state s^*. In the following theorem we provide conditions for the presence of the states of the world in which the agents become polarized.

Theorem 2. *Let us assume that there are two agents $j = 1, 2$, who are characterized by the pair $(R^j, \mathbb{E}^j v)$. If in state of the world $s^* \in S$ the conditions $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) > 0$ and $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) < 0$ hold, then the two agents become polarized in this state of the world.*

Proof. Without loss of generality, let us assume that $\mathbb{E}^1 v > \mathbb{E}^2 v$. For the condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) > 0$ to be satisfied, it is necessary that $v_{s^*} > R^1$. Proposition 1 states that in this case $\Delta_1(s^*) > 0$. Analogously, the second condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) < 0$ holds when $v_{s^*} < R^2$, which further implies that $\Delta_2(s^*) < 0$. That is, two agents $j = 1, 2$ update in different directions and the expected posterior beliefs are farther away from each other than the priors are. Both conditions from Definition 2 are satisfied and the agents, indeed, become polarized in state s^*. □

Even though the focus of the paper is on documenting the belief polarization, a natural question arises of how such belief polarization influences actions taken by the agents. This connection between belief and action should become clear by considering a simple two-period extension of our model.

Example of connection between the belief polarization and actions. Two agents (Alice and Bob) face a choice between different phones. Each of them has a current phone (status quo) with value R_i: they start with different phones with values $R_A = 0.4, R_B = 0.6$. They are considering whether to change phone and buy a model x or y, or keep the current phone. The values of phones x and y are unknown, and the decision makers a) know the prior distribution of possible values, b) can ask a binary question to an expert in order to get partial disclosure of the value for each phone, and c) initially ask only a question about x (myopic, or no initial information about y being available), then only about y.

Suppose phone x has an equal chance of values 0, 0.5, or 1, and a realized value 0.5. Both agents observe the optimal signal I_i according to their own subjective state pooling (based on the status quo), and they reach posteriors $E(v^x|I_A) = 0.75, E(v^x|I_B) = 0.25$. Therefore, Alice would buy the new phone and Bob would keep the old one, and their actions are optimal.

Now phone y appears in the stores, and this phone (with lower uncertainty about the value) has an equal chance of values 0.3, 0.5, or 0.7, and a realized value 0.7. So both agents under perfect information should choose phone y. But Alice and Bob collected different information about x, and this affects their final choice. Bob assigns x an expected value of 0.25, and has a status quo with value 0.6. He learns that $v_y = 0.7$ and buys the new phone. Alice assigns x an expected value of 0.75, so will choose x regardless of the information she can get about y.

3.6 Beliefs' convergence and divergence of beliefs updated in the same direction

In order to draw the whole picture of all possible situations (directions of belief updating), we describe our framework's predictions as to when the beliefs of two agents converge to each other and at the same time move closer to the true value of the new policy.

Definition 3. We say that two agents $j = 1, 2$, who are characterized by the pair $(R^j, \mathbb{E}^j v)$ and are choosing between actions $i = \{1, 2\}$, *converge in their beliefs in the state* $s^* \in S$ when the following two conditions are satisfied

1. $|\mathbb{E}_i^1[\mathbb{E}(v|i)|s^*] - \mathbb{E}_i^2[\mathbb{E}(v|i)|s^*]| < |\mathbb{E}^1 v - \mathbb{E}^2 v|$.

2. $\Delta_1(s^*) \cdot \Delta_2(s^*) > 0$.

Theorem 3. *Let us assume that there are two agents $j = 1, 2$ who are characterized by the pair $(R^j, \mathbb{E}^j v)$. If in state of the world $s^* \in S$ the conditions $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) < 0$ and $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) > 0$ hold, then the two agents converge in their beliefs in this state of the world.*

Proof. Without loss of generality, let us assume that $\mathbb{E}^1 v > \mathbb{E}^2 v$. For the condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) < 0$ to be satisfied, it is necessary that $v_{s^*} < R^1$. Proposition 1 states that in this case $\Delta_1(s^*) < 0$. Analogously, the second condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) > 0$ holds when $v_{s^*} < R^2$, which further implies that $\Delta_2(s^*) < 0$. That is, two agents $j = 1, 2$ update in different directions and the expected posterior beliefs are closer to each other than the priors are. Both conditions from definition 3 are satisfied and the agents indeed converge in their beliefs in the state s^*. □

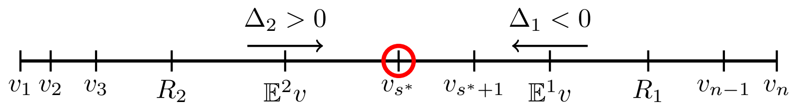

Figure 3: Illustration of the situation when the agents' beliefs converge to each other and to the true value of the new policy.

The situation when two agents converge in their beliefs occurs when agents have different prior expectations of the new policy and when they value the status quo differently. In addition, for both agents, it has to hold that their prior expected value from the new policy $\mathbb{E}^j v$ and valuation of the status quo R^j for the same agent j are close to each other, i.e., both $\mathbb{E}^j v$ and R^j are smaller or larger than v_{s^*}. This situation is illustrated in Figure 3.

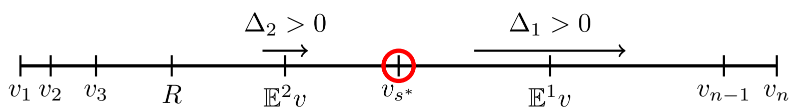

Figure 4: Illustration of the situation when the agents diverge in their belief, updated in the same direction.

Up to this point we have considered only the polarization when two agents are updating in the opposite directions. However, as Figure 4 illustrates, there is also a possibility that posterior expected values of the two agents are further away from each other than their prior expected values, while both agents update in the same direction.

Definition 4. We say that two agents $j = 1, 2$, who are characterized by the pair $(R^j, \mathbb{E}^j v)$ and are choosing between actions $i = \{1, 2\}$, *diverge in their belief updated in the same direction* when in the state $s^* \in S$ the following two conditions are satisfied

1. $|\mathbb{E}_i^1[\mathbb{E}(v|i)|s^*] - \mathbb{E}_i^2[\mathbb{E}(v|i)|s^*]| > |\mathbb{E}^1 v - \mathbb{E}^2 v|$.

2. $\Delta_1(s^*) \cdot \Delta_2(s^*) > 0$.

In Appendix D, we consider an example that demonstrates that such divergence can occur, and also how our results are influenced by the change in the cost of information, but these situations are too complex to be studied analytically.

4 Experimental design

Our theoretical results show that the agent-specific value of the status quo determines the information structure selected. As a consequence, two agents with different values of the status quo might become polarized in expectations. In particular, the rationally inattentive agent chooses to learn whether the outcome of the new policy is better or worse than the status quo, and not to learn the exact state-dependent outcome of the new policy. We have denoted this information strategy as state pooling behavior.

The main results of the model rely on several assumptions about a decision maker's preferences (risk neutrality), ability to estimate probabilities (by correctly updating beliefs), and motivation (information has a purely instrumental value). The experimental literature reports a large amount of evidence that casts doubts on human ability to perform these tasks as accurately as the theory requires, and highlights that belief divergence could be mitigated or enhanced by human biases. We are interested in testing whether belief divergence in expectations, the main result of our model, can occur in a lab setting and whether behavioral components enhance or mitigate its magnitude.

In the following sections of the paper we investigate whether our normative model is also accurate in describing human behavior. We do so by running a lab experiment in which participants are allowed to collect information before making choices under uncertainty. We collect actions and beliefs separately, and combine them to compute a cardinal indicator for beliefs divergence and to compare human behavior and theoretical predictions. Our design allows us to test whether a manipulation of the status quo affects the choice of information source and generates belief polarization in expectations.

4.1 Overview of the experimental design

The experiment comprises four tasks and a final questionnaire. In the first and second tasks, subjects face a binary choice between the opaque box (*risky action*), which contains a single "color ball," the value of which depends on the unknown color, and the transparent box (*safe action*), containing a single ball whose value is known. The color ball is randomly drawn from a box containing three balls (*states*) with different colors (red, yellow, blue) with uniform probability of being selected. The two tasks differ in the way we provide interim information about the color ball. In task 1 four possible advisors (representing degenerate signal structures) are evaluated separately and subjects report their willingness to accept (Becker-DeGroot-Marscha method,[15] BDM thereafter) renunciation of the signal. In each trial for task 2 only two advisors are displayed and the subject makes a binary choice between them. In tasks 3 and 4 we elicit unconditional and conditional beliefs for different advisors, assigned exogenously. We ensure incentive compatibility by paying subjects for a single decision randomly selected from the entire experiment. Subjects never receive feedback about their decisions until the very end of the experiment. Each subject participates in all of the following tasks, in the order listed below.

4.2 Task 1 - Colorblind advisor game

In each round of task 1 subjects (i) choose an action contingent on the advisor and signal received (Figure 5a) and then (ii) indicate for each advisor the willingness to accept renunciation of its signal (Figure 5b). This task is presented at the beginning of the session to familiarize participants with the choice environment and the notion of the advisor.

Subjects play 10 rounds with the same four advisors and different lottery return values. Three of the advisors in this game (named Red, Yellow, and Blue) are described as colorblind to all colors except the subject's own. are able to observe the color ball and report truthfully whether it matches her own name's color.[16] The fourth advisor is named Rainbow and reports every color accurately, without uncertainty. In each round, the subject chooses which box she would pick in each hypothetical advisor/answer scenario (strategy method). Then, the

[15] Becker, DeGroot and Marschak (1964).

[16] For example, the Red advisor returns a signal RED or NOT RED, which is easy to interpret.

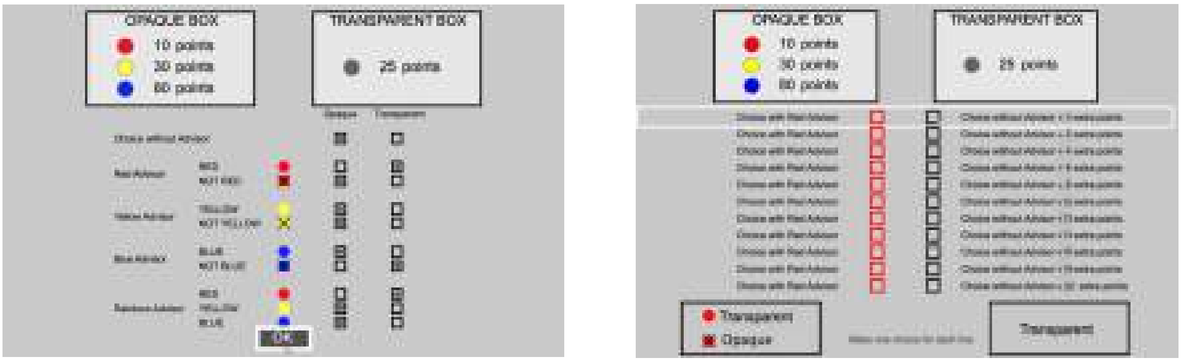

(a) Task 1, Screen 1: Action choice

(b) Task 2, Screen 2: WTA for each advisor

Figure 5: Task 1 - Colorblind advisor game. Left: Subjects choose an action (box) contingent on the advisor and signal received. The possible values of each action are indicated on the top of the screen. Each state (ball color) is equally likely to occur. Right: Subjects indicate for each advisor the willingness to accept renunciation of its signal in a series of binary choices (BDM method). At most one switch is allowed. Action choices selected in the previous stage are reported on the bottom of the screen.

subjects fill out a multiple choice list for each of the four advisors, choosing between pairs of options: "Choice with the X advisor" (X is replaced with the advisor's name) or "Choice without advisor + w extra points," for w between 0 and 20 points, in 2 points intervals.[17]

4.3 Task 2 - Imprecise advisor game

In each round of task 2 subjects (i) choose one advisor between the two options available (Figure 6a) and then (ii) indicate the signal-contingent action for each signal (Figure 6b).

Subjects play 40 rounds with different pairs of advisors and values for the ball in the transparent box.[18] Each round comprises two parts. First, subjects observe a pair of advisors and make a binary choice to select which advisor they want to consult. Subsequently, only the selected advisor is consulted, a signal-contingent binary choice is implemented, and the participant chooses one box based on the signal received. Each advisor is defined as a triplet of state-contingent conditional probabilities of providing a binary signal.

The 40 rounds are designed as a combination of 20 advisor pairs and two values for the ball in the transparent box (safe option). The advisors are selected in order to examine

[17] The value w at which a subject i switches from preferring the former to the latter option reveals her subjective valuation w_I^i.

[18] The ball in the transparent box can take two values (30 and 65 points). The values for the balls in the opaque box are unchanged during the task (10, 50, and 80 points, uniform probability of being drawn).

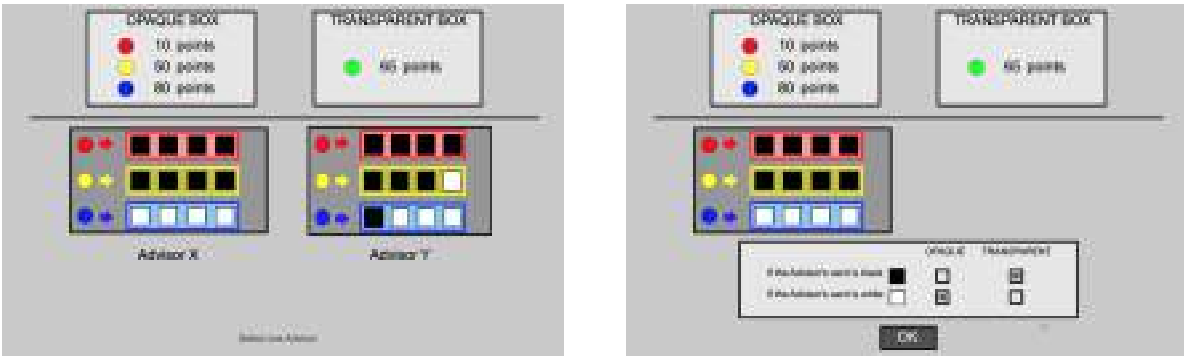

(a) Task 1, Screen 1: Advisor choice

(b) Task 2, Screen 2: Action choice

Figure 6: Task 2 - Imprecise advisor game. Left: Subjects choose one signal structure (advisor) between the two options available. Each advisor is a triplet of state-contingent signal probabilities. Right: Subjects indicate the signal-contingent action for each signal (strategy method).

preference over sources of information and formulate predictions about the effect of the safe option on information collection and posterior beliefs. In particular, 11 out of 20 pairs of advisors are designed such that a Bayesian agent would pick different advisors by changing the safe option.

4.4 Task 3 - Card color prediction game

In each round of task 3 subjects indicate the likelihood of observing each signal for a given advisor (Figure 7a). We elicit the subjects' signal probability beliefs for each of the 20 advisors using a single slider with sensibility to the unitary percentage level. Each round contains a single advisor from those used in task 2 and subjects are asked to report the probability of a black or white card being shown. We incentivize accurate and truthful reporting by using the quadratic loss scoring rule with two states (card colors).

4.5 Task 4 - Ball color prediction game

In each round of task 4 subjects indicate the likelihood of each state given an advisor and signal (Figure 7b). We elicit the subjects' posterior probability beliefs for each of the 20 advisors and for each possible signal realization, using a double slider with sensibility to the unitary percentage level. Each round contains a single advisor from those used in task 2

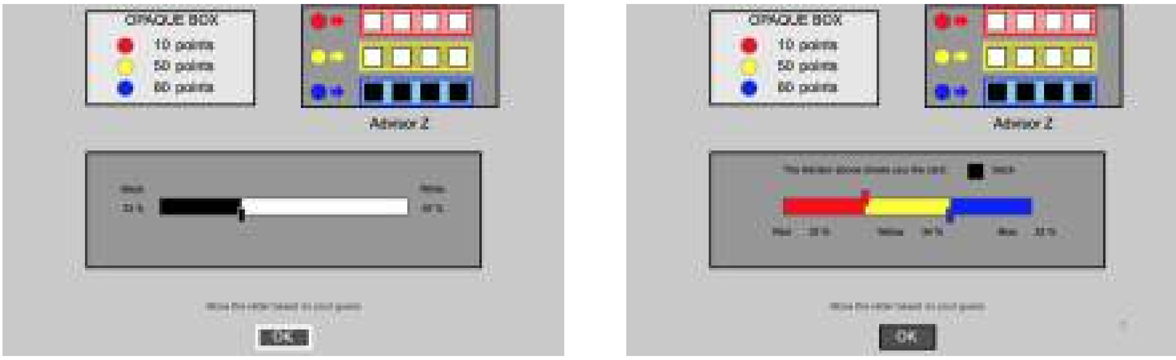

(a) Task 3: Beliefs over signal likelihood

(b) Task 4: Beliefs over state likelihood

Figure 7: Left: Task 3 (Card color prediction game). Subjects indicate the likelihood of observing each signal (card color) for the given advisor. Right: Task 4 (Ball color prediction game). Subjects indicate the likelihood of each state (ball color) given an advisor and signal. In both tasks subjects move the slider(s) and receive a number of points according to the quadratic loss scoring rule described in the instructions.

and one realized signal (black or white card). The subject is asked to report the probability of a red, yellow, or blue ball being in the opaque box after observing the card color. We incentivize accurate and truthful reporting by using the quadratic loss scoring rule with three states (ball colors).

4.6 Questionnaire

The final part of the experiment is a questionnaire designed to collect demographic variables (including field of study and familiarity with Bayes' rule), psychological measures (Life Orientation Test - Revisited,[19] LOT-R hereafter), risk attitude (Holt-Laury risk elicitation method with multiple price list, Holt and Laury 2002, HL hereafter), cognitive ability (five questions from the Raven Progressive Matrices Test, Raven hereafter), as well as questions on the subjects' strategy in the first and second task.[20]

[19] See Scheier, Carver and Bridges (1994).

[20] We did not collect questionnaire information for the first 22 participants to the experiment.

4.7 Procedure

The experiment was run in the CELSS (Columbia Experimental Laboratory for Social Sciences) between August and September 2019.[21] The experiment was coded in MATLAB (Release 2018b) using Psychotoolbox 3 (Psychophysics Toolbox Version 3). Eighty-five volunteers were recruited using the platform ORSEE[22] (Online Recruitment System for Economic Experiments) and were naive to the main purpose of the study. All subjects provided written, informed consent. The whole experiment took on average 85 minutes, including instructions and payment. On completion of the experiment, the subjects received payment in cash according to task performance. Each subject received a \$10 show-up fee, and played for a bonus prize of \$15. In addition, a subject could earn between \$0.10 and \$4 in the risk elicitation task and \$0.50 for each question of the cognitive test, up to \$5. The average payment was \$25. A small number of participants were recruited for each laboratory session (6 on average) in order to facilitate clarification questions during the experiment.

4.8 Hypotheses

Our experiment allows us to test the main predictions of the model, as well as disentangle the possible factors that mitigate or enhance the results with respect to the behavior of an optimal decision maker. The first hypothesis refers to the crucial effect of the status quo on information acquisition.

Hypothesis 1. *A change in the status quo (safe option) generates a reversal in the choice between advisors when such a reversal is optimal.*

We test this hypothesis by collecting choices over information structures under different values of the safe option. The optimal advisor choice would not be sufficient to generate belief polarization. The second hypothesis refers to the result that appears in the title of the paper.

Hypothesis 2. *A change in the status quo (safe option) generates beliefs polarization in expectations.*

[21] The experimental protocol was approved by the Columbia Institutional Review Board, protocol AAAS5801.

[22] See Greiner (2015).

We test this hypothesis by collecting subjective beliefs after receiving a signal, in addition to the advisor choices. There are three main channels that may mitigate or enhance the results with respect to the behavior of an optimal decision maker: non-standard preferences over realized states, biased beliefs, and non-standard preferences for information structures. We summarize these potential confounding factors in three channels.

Channel 1. *Subjects evaluate information structures based on non-standard preferences over the realized state, e.g. risk aversion or risk seeking.*

Channel 2. *Subjects evaluate information structures based on subjective beliefs about the likelihood of signal and state realizations.*

Channel 3. *Subjects evaluate information structures based on non-instrumental characteristics, including ease of processing of the signals (certainty, few possible outcomes).*

Our setup allows us to test the main hypotheses and study the three behavioral channels by collecting detailed data about willingness to pay and binary choices between information structures.

4.9 Experimental investigation

We consider a setting with three possible states and two actions that generate state-contingent payoffs. The actions represent two policies: the current policy (the status quo), whose return $R \in \mathbb{R}$ is known and independent of the state, and a new policy, whose return $v_s \in \mathbb{R}$ is uncertain. The state of the world $s \in S = \{r, y, b\}$ is represented by a color associated with the deterministic return for the uncertain policy: r (red, low return), y (yellow, intermediate return), or b (blue, high return), with $0 < v_r < v_y < v_b < 100$ and $v_r < R < v_b$. An agent with a correct uniform prior belief $\mathcal{P}(s) = \frac{1}{3}$, $\forall s$ observes an informative signal about the state and selects one of the two policies. The return of own choice depends on the selected action and the realized state, and represents the probability, expressed in percents, of receiving a fixed prize k (\$15 in our laboratory experiment).

Information is valuable because it informs the subsequent binary choice between policies. We let $\sigma \in \{0, 1\}$ denote the realization of a stochastic signal that the subject may observe.[23]

[23]Notice that the signal is deterministic in the case of a triplet of degenerate probabilities.

Since we have three states and two possible signal realizations, a signal structure is a triplet of state-dependent probabilities $\mathcal{P}_I(\sigma = 1|s)$. We will refer to such a triplet I as an information source or advisor.[24]

The Bayesian agent represents a natural benchmark to consider the objective value of information in this environment. Let $V(I)$ denote the bonus that renders the agent indifferent between playing the game without additional signals[25] (but receiving additional $V(I)$ "tickets") and playing the game with the signal structure I. The valuation of the information structure I is given by a chosen lottery and by the observed signal

$$V(I) = \sum_{\sigma\in\{0,1\}} \underbrace{\max\left\{\sum_{s\in S} v_s \mathcal{P}_I(s|\sigma), R\right\}}_{=V(\sigma|I)} \mathcal{P}_I(\sigma) - \underbrace{\max\left\{\sum_{s\in S} v_s \mathcal{P}(s), R\right\}}_{=V(\emptyset)}, \tag{6}$$

where $V(\emptyset)$ is the expected value of the action chosen without observing any signal and $V(\sigma|I)$ is the expected value of the action chosen after receiving signal σ from advisor I.

We can generalize the subjective valuation in order to include non-instrumental preference over information. A decision maker i has a subjective valuation $V^i(I)$ of the signal structure I that depends both on the instrumental value $V(I)$ and other characteristics of I; for example the type of "optimistic/pessimistic" information that it provides. We postpone further discussion about possible differences between Bayesian and subjective valuation of information to the results section.

5 Experimental results

This section contains the main results of the experimental investigation. We report aggregate choices between sources of information (Section 5.1) and provide evidence that (i) subjects do react to the value of the status quo as predicted by the theoretical model, (ii) the variation in the status quo leads qualitatively to the belief polarization in our laboratory setting and (iii) we observe preference for state pooling information structures.

[24] Notice that even though the three states are equally likely, the two signals need not be equally likely.

[25] Playing without any additional information is, from a theoretical perspective, equivalent to playing with a purely noisy signal. We prefer to refer to the former case for the sake of clarity.

In section 5.2, we discuss to what extent the main behavioral channels (introduced in section 4.8) cause deviations of the participants' behavior from optimality. We verify that subjects' actions are consistent with the optimal behavior of a risk-neutral agent and that their beliefs about the likelihood of signal and state realizations are close to the optimal ones. Consequently, we identify the evaluation of information structures based on non-instrumental characteristics as a major driver of the observed deviations from optimality.

Finally, we analyze the willingness to pay (WTP) for information structures in the first task (Section 5.3), where we observe that (iv) subjects display compression in their WTP and (v) are willing to pay higher amounts for information about the most desirable state.

5.1 Beliefs polarization in the laboratory experiment

Our model predicts that a change in the status quo creates belief polarization because of the endogenous choice of information structures (advisors). In our experimental design this means that, given the true state, the same decision maker will have different beliefs (ex ante, before the signal realization) based on her status quo value. The experiment contains 11 pairs of trials that can be used to verify whether such polarization occurs. We combine the data collected for the binary advisor choice (task 2) with the subjective beliefs about posterior distribution (task 4) to calculate the magnitude of the observed polarization in the laboratory experiment.

The first hypothesis is that a change in the status quo generates a reversal in the advisor choice when such a reversal is optimal. Figure 8a shows that the hypothesis is confirmed in the trials from Task 2 where we expect to observe the reversal. Advisor I_1, defined here as the best one under $R = 30$, is chosen 66% of the times when it is optimal to do so (11 trials), and only 30% when $R = 65$ (11 trials). The difference between the two treatments is large and highly significant, and confirms our first hypothesis. In Figure 8b we depict the probability of choosing the advisor (out of the pair of advisors labeled I_1 and I_2) as a function of the difference in the instrumental values of the two presented advisors, with all the values computed as in equation 6. The probability of selecting advisor I_1 increases with the difference between the instrumental values of advisor I_1 and advisor I_2. 95% of the trials lie in the first and the third quadrant, that is, whenever the instrumental value of advisor I_1 is higher than advisor I_2 the probability is above or at most equal to 1/2, and vice versa. Probabilities increase almost linearly and not stepwise near the zero, suggesting that subjects do not respond only to the sign of the difference in the instrumental values between the advisors, but to the actual value of the difference.

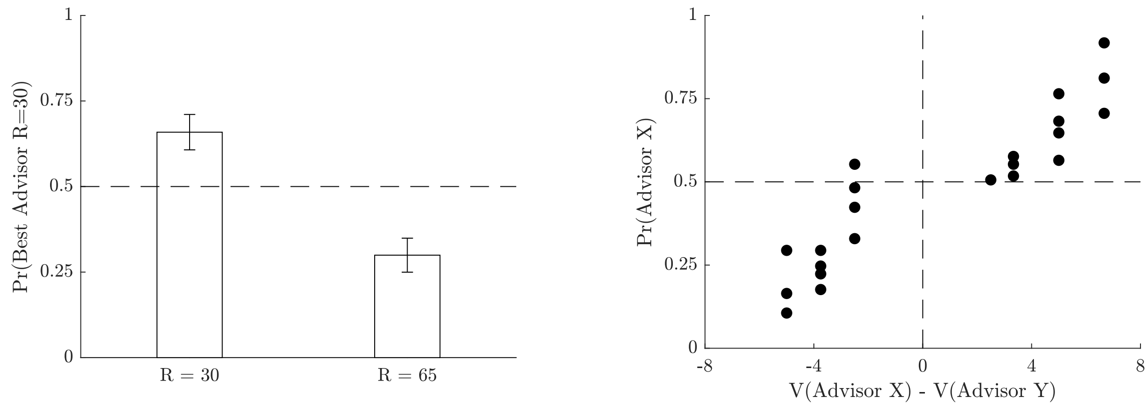

(a)

(b)

Figure 8: Advisor selection probability in Task 2. Left: Probability of selecting I_1, defined here as the best advisor under $R = 30$, for each of the two treatments (11 trials per treatment, 935 observations per treatment). Right: Probability of selecting advisor I_1 based on the difference between instrumental values (22 trials, 85 observations per trial).

Experimental Result 1. *Subjects systematically react to the value of the status quo and choose the optimal advisor (information structure).*

The advisor choice reversal is necessary but not sufficient to generate polarization. If participants are systematically biased in the computation of probabilities, and react less (more) to the new information provided by the signal, this will reduce (increase) the magnitude of the polarization. We elicit subjective beliefs with a strategy-proof mechanism and we verify that the participants hold, on average, very accurate beliefs. In the fourth task of the session we elicit the posterior likelihood of each state, given an advisor and a signal realization. On average, we observe accurate probability estimates, close to the predictions of an optimal Bayesian agent. Results are displayed in Appendix I and show that 1) participants are on average accurate in the estimate of probabilities, 2) we do not observe a systematic difference between estimates involving different states (i.e., we do not have evidence of motivated beliefs, Bénabou, 2015), and 3) the results show mild evidence of conservatism (central tendency of judgement), as vastly reported in experiments with subjective estimates (Hollingworth, 1910; Anobile, Cicchini and Burr, 2012). A linear fit of the average subjective estimates over the unbiased ones confirms the mild conservatism (slope $\beta = 0.825$) and the overall good fit ($R^2 = 0.993$).

By combining advisor choices and posterior beliefs we can finally calculate the magnitude of polarization of beliefs. We consider separately each of the 11 pairs of trials in which the agents should switch advisor because of the status quo manipulation. For every trial and state, we calculate the ex-ante expected value for the risky action (conditional on the state, but not conditional on the signal realization). The *predicted polarization* is the absolute difference between the two ex-ante expected values: one is generated under the low status quo, the other under the high status quo. The *realized polarization* (Figure 9a) is calculated similarly, by replacing the observed human behavior in two stages. First, we replace the posterior beliefs with the average subjective ones (task 4). Second, we replace the advisor choice probabilities with the observed ones (task 2). Participants switch advisors less then predicted, and update their beliefs slightly less than predicted, so both the components reduce the magnitude of the realized polarization. A linear fit of the distribution shows that the realized polarization is, on average, 32% of the predicted one, with little dispersion across pairs of trials, as confirmed by the high $R^2 = 0.892$.

Experimental Result 2. *Variations in the value of the status quo generate ex-ante belief divergence (before the signal realization, and after controlling for the true state) qualitatively analogous to those predicted by the model, but with smaller magnitude.*

State pooling is the key mechanism for our model that determines the advisor switch. We define state pooler advisors as follows.

Definition 5. *An advisor with information structure I is a state pooler under status quo value R when it can provide a signal σ such that the posterior belief for the agent is either $\mathcal{P}(\pi_s \geq R|\sigma) = 0$ or $\mathcal{P}(\pi_s \geq R|\sigma) = 1$.*

When we investigate how the probability of choosing the state pooling advisor depends on the instrumental value of such an advisor (see Figure 9b), we can notice that the probability is greater than 1/2 when it is optimal to select the state pooler, and otherwise it is below 1/2.

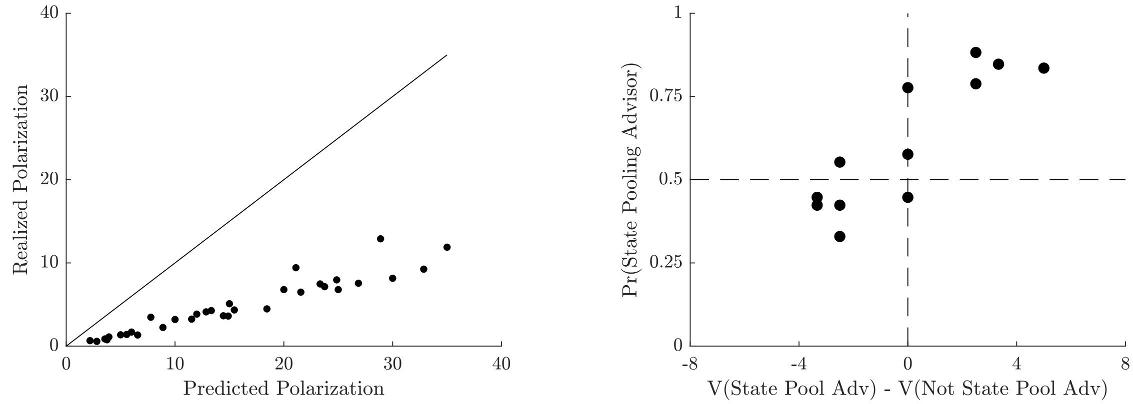

(a) (b)

Figure 9: Left: Realized polarization. Predicted (Bayesian decision maker) and realized polarization (calculated from tasks 2 and 4) in the 11 pairs of trials with predicted advisor switches (3 states per pair of trials, n=85). Right: Probability of choosing the state pooling advisor over the alternative one in the trials with exactly one state pooling advisor. (12 trials, 85 observations per trial).

However, we can also notice that the probability of selecting the state pooler increases with the instrumental value. When the state pooling and non state pooling advisors both have the same instrumental value (0 on x-axis), subjects strictly prefer the state pooling advisor in comparison to a non state pooling advisor, even though it is not more informative.

In the next section we further discuss the quantitative departure from the theoretical prediction and we analyze separately the three main behavioral channels that represent potential confounding factors.

5.2 Behavioral channels of departure from theoretical predictions

Possible deviations from the optimal choice between alternative sources of information in task 2 (imprecise advisor game) can be rationalized by risk preferences, systematically biased beliefs, and non-instrumental preferences over information structures, as discussed in Section 4. We now consider each of these channels separately and discuss whether they provide an explanation for the advisor choices observed.

Risk preferences represent the first behavioral channel. Possible deviations from the optimal choice between alternative sources of information in the imprecise advisor game (task 2) can be consistent with risk aversion, as this would affect the evaluation of each advisor. We design our experiment in order to minimize this concern, by using probability points as prizes (see, e.g., Harrison, Martínez-Correa and Swarthout, 2013) and assigning probabilities between 10% and 90% (to minimize concerns about extreme probability events).

Our analysis confirms that risk aversion does not represent a driver of the participants' departure from the predictions: we summarize below the main results, with more analysis available in Appendix H. In task 2, given an advisor and a signal realization, participants were asked to choose one action corresponding to the risky lottery (opaque box) or the safe option (transparent box). If we assume that the decision maker adopts a CRRA utility function $u_\alpha(x) = \frac{x^{1-\alpha}}{1-\alpha}$, and we use maximum likelihood method to estimate the concavity of the utility function, represented by the risk aversion coefficient α, we obtain $\hat{\alpha} = 0.34$, which indicates a small risk aversion, compared to the null hypothesis $\alpha = 0$ (risk neutral agent). We test the statistical significance of adding the risk aversion coefficient in the action choice model using the likelihood ratio test, and we reject the null hypothesis ($p < 0.001$). The use of risk preferences slightly improves the predictive power, which we measure using the likelihood ratio (Cohen et al., 2013). This measure, representing an R squared statistics for logistic regressions, is $R^2_{risk.neutral} = 0.382$ using risk neutrality and $R^2_{risk.averse} = 0.422$ using risk aversion. Nevertheless, the magnitude of the deviation from risk neutrality is modest and unable to explain a choice reversal for the advisor selection.

Experimental Result 3. *Participants behave similarly to a risk neutral agent, and risk preferences represent a small driver of deviation from optimality in the choice between actions.*

Biased beliefs represent the second behavioral channel. We usually assume that the instrumental value of each advisor is determined using the exact probabilities for signal and state realizations. We relax this assumption, and consider the behavior of an agent who selects the advisors based on the subjective value instead. We compute the subjective advisor value by using the subjective beliefs about signal likelihood and state likelihood (conditional on the observed signal) that we elicit in tasks 3 and 4, respectively. The average responses for these tasks are shown in Appendix G, and both distributions display small conservatism in the estimates.

In Figure 10a we depict the advisor choice probability (from a pair of advisors labeled I_1 and I_2) as a function of the difference in the instrumental values of the two presented advisors, with all the values computed as in equation 6 and using unbiased beliefs. This is similar to Figure 8b, but now we display all 40 trials in the experiment, including those where we do not expect to observe the advisor switch. We compare this result with Figure 10b, where the instrumental value of each advisor is computed using the average subjective beliefs ($V_{\text{subjective}}$). The use of subjective beliefs reduces the predictive power of the model, which we measure using the likelihood ratio (pseudo R squared, as above): $R^2_{unbiased} = 0.563$ using unbiased beliefs and $R^2_{subjective} = 0.225$ using subjective beliefs.

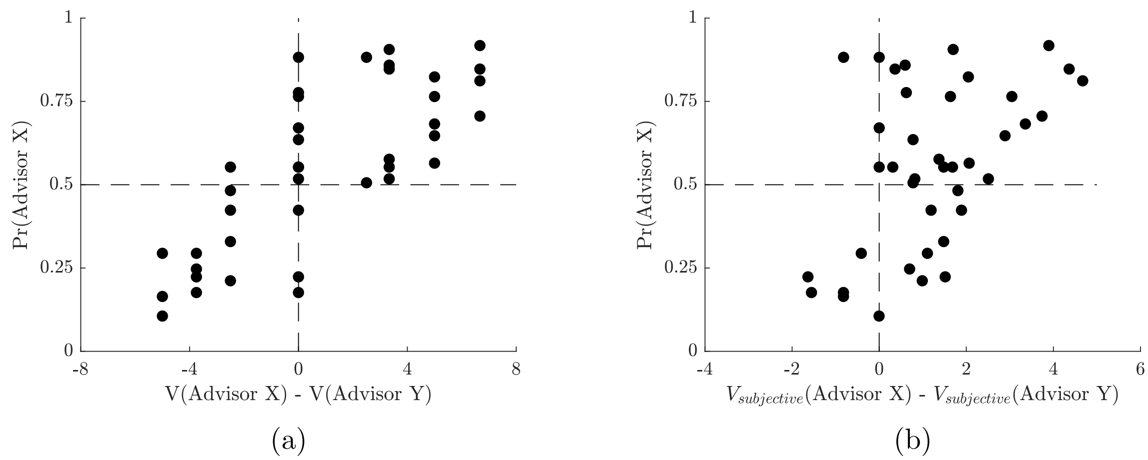

Figure 10: Task 2: Advisor selection probability (all trials). Left: Probability of selecting I_1 based on the difference between instrumental values (40 trials, 85 observations per trial). Right: Instrumental values of the advisors are calculated using the average subjective beliefs elicited in tasks 3 and 4.

Experimental Result 4. *Participants behave similarly to an agent with correct understanding of probabilities. Subjective beliefs about the likelihood of signal and state realizations represent a small driver of deviation from optimality in the choice between advisors.*

Non-instrumental preferences over information structures represent the third behavioral channel. Differently from the previous two approaches, we now assume that the decision makers evaluate some non-instrumental features of the advisor (e.g. simplicity and certainty) and chooses the preferred advisor based on the tradeoff between instrumental and non-instrumental characteristics. This is equivalent to the addition of a cognitive cost for each of the signal structures. Before proceeding, we introduce the definition for a *certain* advisor, which will be useful for the analysis.

Definition 6. *An advisor with information structure I is certain when there exists a state s and a signal σ_s such that $\mathcal{P}(s|\sigma_s) = 1$ and $\mathcal{P}(s|\sigma) = 0 \;\; \forall \sigma \neq \sigma_s$.*

We start with a simple test for the certainty effect, defined as "disproportionate preference for information structures that may perfectly reveal the state of the world" (Ambuehl and Li, 2018). Figure 11a plots the probability of choosing the certain advisor given the difference between the values of the certain and uncertain advisors. It is apparent that the probability does not increase monotonically with the informativeness as in the previous figures. Once it is optimal to choose the certain advisor (the simple advisor is the best one) it jumps to 86%, on average. On the other hand, when it is optimal to choose the uncertain advisor, the probability of selecting the certain advisor is, on average, only 46%.

In order to extend the effect to a larger set of conditions, we introduce here c_I as a discrete measure of complexity (or cost) for the signal structure I:

$$c_I = \sum_{\sigma} \left(\sum_{s} \mathbb{1}(\mathcal{P}(s|\sigma) > 0) - 1 \right). \tag{7}$$

This index counts, for every distinct signal realization, how many states receive positive probability in the posterior beliefs generated by that signal. In our experiment, it takes values between 1 (simple) and 4 (complex). Figure 11b shows that participants tend to prefer advisors with a low index, even after controlling for the advisors' value difference.

We run a series of logit regressions aimed at measuring the relative importance of various features of the advisors. In addition to the instrumental value $V^{\text{Bayes}}(I)$, we use the complexity index c_I and three binary variables to indicate whether the advisor 1) is the best in the pair, 2) provides certainty, and/or 3) provides state pooling. The results from the regression are presented in Table 2, and are consistent with the patterns previously discussed. The instrumental value of the advisor $V^{\text{Bayes}}(I)$ has a significant effect on the probability of the advisor being selected. However, certainty, state pooling, and the general index of complexity are also significant.

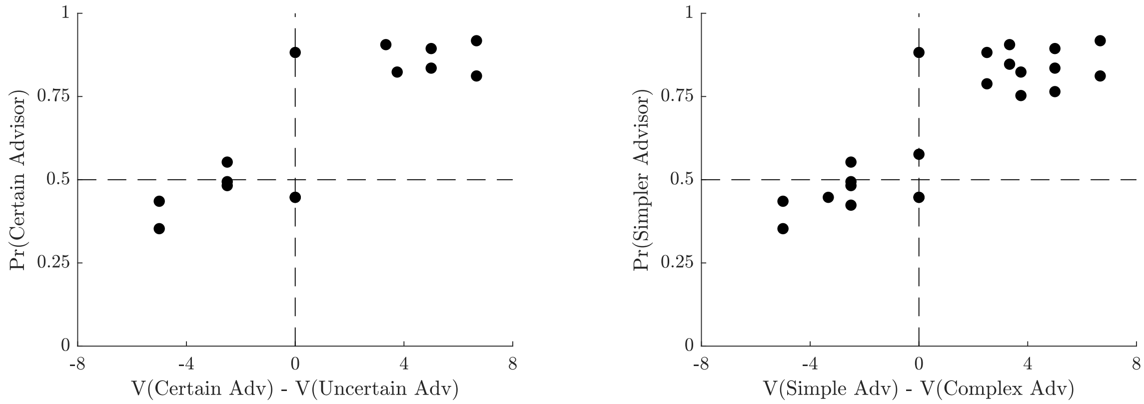

(a) (b)

Figure 11: Preference for certainty and simplicity. Left: Certainty effect. Probability of chosing a certain advisor in the trials containing a certain and an uncertain advisor (14 trials, 85 observations per trial). Right: Simplicity effect. Probability of chosing the simplest advisor (as defined in equation 7) in the trials containing two advisors with different complexity scores (21 trials, 85 observations per trial).

Method: Logit, Dependent variable: Advisor choice

	(1)	(2)	(3)	(4)
Value $V^{\text{Bayes}}(I)$	0.246	0.217	0.235	0.232
	(0.018)	(0.011)	(0.011)	(0.019)
Best Advisor (dummy)	-0.084			-0.007
	(0.096)			(0.102)
Complexity c_I		-0.359		-0.074
		(0.037)		(0.076)
Certainty (dummy)			0.511	0.428
			(0.069)	(0.110)
State Pooling (dummy)			0.404	0.330
			(0.069)	(0.102)
Trials	All	All	All	All
Observations	3,400	3,400	3,400	3,400

Table 2: Advisor choice across all the trials in task 2.

Experimental Result 5. *Participants significantly prefer simple advisors, and advisors providing certainty in particular. This result is robust even after controlling for the instrumental value of the advisors available. This is a major driver of the observed deviation from optimality in the choice between advisors.*

5.3 Willingness to pay for signal structures

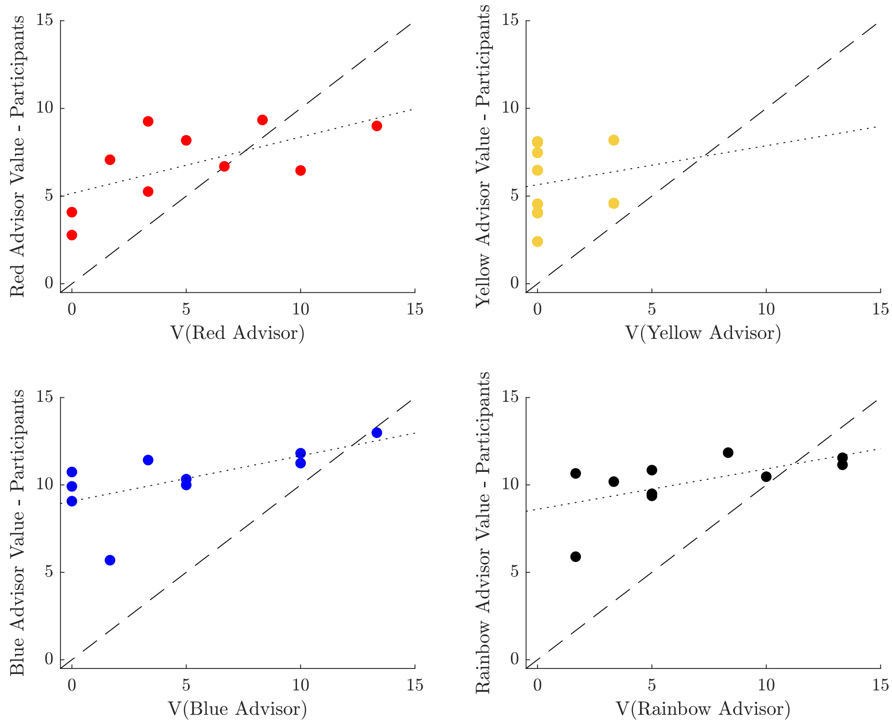

Figure 12: Willingness to pay for information (task 1). Comparison between the average subjective valuations of advisors across participants and the valuations for the optimal decision maker. The color used in each panel indicates the type of advisor: red, yellow, blue, and rainbow advisors (ordered from top-left to bottom-right). Optimal valuation (dashed lines) and linear regression estimates (dotted lines) are shown for comparison. n=85 observations per point in each panel (10 trials, 4 advisors per trial).

The Colorblind advisor game introduced in task 1 provides a different dataset that we can compare with the results from the other tasks. In each of the ten trials we collect signal-contingent actions (risky or safe options) as well as the subjective willingness to pay (WTP) in order to observe a certain signal structure. More precisely, we elicit the willingness to accept (WTA), expressed in probability points of winning the bonus, in exchange for the opportunity of playing the game without the advisor.[26] For each of the four advisors in

[26]Standard theory predicts that the difference between WTA and WTP is negligible when income effects are small. We implemented several measures in order to maintain this difference as negligible: i) we used probability points to minimize risk aversion concerns, ii) we maintain the gain domain during the whole session and iii) we avoid framing the payment in task 1 as a cost/loss. For an exact relation of WTA and WTP see Weber (2003).

the game we elicit subjects' valuation of the advisor using multiple price lists, an incentive-compatible implementation of the BDM mechanism. The red, yellow, and blue advisor provide a binary message, whereas the rainbow advisor fully reveals the true state. Figure 12 shows, for each advisor type, the relation between subjective (averaged across participants) and theoretical advisor value (for an optimal decision maker). We notice that for all the advisors the subjective evaluation tends to exceed the theoretical one (positive intercept) and there is a general positive relation between the two, with subjective values increasing with the theoretical ones, but not as much as the latter (slope coefficients lower than 1). This pattern is known as the compression effect and is well known in experiments with explicit elicitation of WTP for sources of information (e.g. Ambuehl and Li, 2018). We have several cases of advisors whose theoretical value is equal to zero (including for most of the yellow advisors): the observation of a signal from them is not pivotal for the chosen action with respect to the decision without an advisor, yet the subjects invest a significant amount of points to receive this piece of information.

Experimental Result 6. *Participants display a compression effect in their willingness to pay for information structures. They tend to overpay for advisors with low and even zero instrumental value, and their subjective WTP increases with the theoretical values, but with a slope smaller than one.*

The comparison of the plots of different advisors highlights a consistent pattern. The compression effect appears similarly for all four advisors, with a similar slope of the linear regression between observed and theoretical values. At the same time, intercepts are significantly different, with similar values for red and yellow advisors, and much higher levels for blue and rainbow advisors. This difference is aligned with preference for information structure biased in favor of the most desirable (blue) state.

This result is confirmed by running a simple OLS regression of the subjective advisor value using the theoretical value as the regressor. Table 3 shows that the slope is positive but lower than one (compression effect) and the intercept is positive and significantly different from zero (a result analogous to the conservative probability estimates observed in tasks 3 and 4). When we allow the intercept to differ across advisors, we notice that they are not different between the red and yellow advisor, whereas the blue and rainbow advisors receive significantly higher WTPs. This result is consistent with those observed in environments with non-instrumental information (Masatlioglu, Orhun and Raymond, 2017) in which subjects display wishful thinking and desire to observe signals that are more accurate about the positive outcomes. The blue state represents the most desirable outcome in our setting, and it is fully revealed by consulting either the blue or the rainbow advisors. The slope of the curves is not significantly different across advisors (not reported in the table, see Appendix J) confirming that the effect does not arise from a different sensitivity to instrumental value. Instead, it provides evidence in favor of intrinsic (non-instrumental) preference for information structures, similarly to the previously discussed preferences in favor of advisors providing certainty or state pooling in the posterior beliefs.

Experimental Result 7. *Participants are willing to pay significantly higher amounts for advisors that provide evidence in favor of the most desirable state, as well as for advisors that fully reveal the true state.*

Method: OLS, Dependent var: $V^i(I)$

	(1)	(2)	(3)	(4)
Constant	6.66	5.62	3.65	7.58
	(0.166)	(0.231)	(0.299)	(0.339)
Value $V^{\text{Bayes}}(I)$	0.372	0.269	0.430	0.116
	(0.027)	(0.030)	(0.042)	(0.047)
Red advisor		-0.19	-0.31	0.26
		(0.353)	(0.443)	(0.546)
Blue advisor		3.41	3.51	2.59
		(0.349)	(0.496)	(0.486)
Rainbow advisor		2.74	2.70	2.62
		(0.373)	(0.496)	(0.546)
Trials	All	All	$R > v_y$	$R < v_y$
Observations	2520	2520	1260	1260

Table 3: Aggregate valuations of information structures in task 1.

The result is qualitatively robust to the separate analysis of trials with high or low status quo. Columns 3 and 4 contain the regressions run independently, with the two parts of the dataset, divided based on the relative value of the status quo R with respect to the intermediate payoff v_y. Although the signs and significance of the estimates are unchanged, we observe different magnitudes. When the value of the status quo is higher than the intermediate state (column 3) the intercept is lower and the slope steeper (less compression). In this case the decision maker faces a safer choice problem and she reacts more to the incentive represented by the instrumental value of the advisor.

Experimental Result 8. *Participants' WTP for information is more responsive to the instrumental value (less acute compression effect) in safe trials, in which the status quo value is higher than the intermediate payoff.*

6 Heterogeneity across subjects

In this section, we present how the results of the main analysis are robust across individuals, and whether observable characteristics, such as mathematical literacy and risk attitude, can predict the heterogeneity in their behavior. We showed in Figure 9a that at the aggregate level we observe 32% of the predicted polarization. We replicate the analysis for every participant: Figure 13a shows the estimated polarization coefficient $\hat{p}_i$ for every subject, with 0 for no polarization and 1 indicating the magnitude of polarization predicted by our model. We observe substantial heterogeneity, with the full range of possible values, and an average polarization equal to 53% of the prediction. We even encounter a few values above 1; this is possible when subjective beliefs are characterized by base-rate neglect (instead of conservatism, as we observe at the aggregate level).

We previously noted that the polarization coefficient depends on advisor choices (task 2) and posterior beliefs (task 4). We separate the two components: even if subjective beliefs represent, on average, a small deviation from the predictions, the result changes at the individual level, and the average score jumps from 53% to 71%, reducing the missing polarization by over one third.

Another way to recognize the vast heterogeneity in the subjects' behavior is to consider two dimensions of the advisor choice: the choice of the best advisor in terms of the instrumental value, and the choice of the simplest advisor using the complexity index introduced in Section 5.2. It could be the case that most participants suffer from a minor bias in favor of

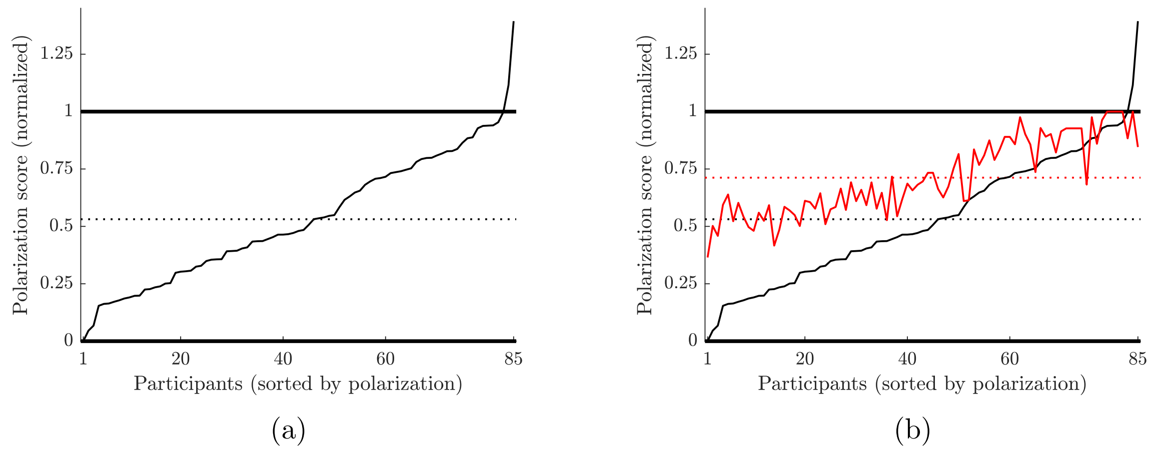

Figure 13: Estimated polarization coefficient $\hat{p}_i$ by subject. Left: Distribution of coefficients, subjects ordered by $\hat{p}_i$. Right: Decomposition of the missing polarization, replacing subjective beliefs with unbiased ones.

simple advisors, or that fewer subjects have strong preferences for them. Figure 14 provides some evidence in favor of the second explanation and suggests that the participants can be categorized into three broad groups based on these two dimensions. A cluster of accurate participants that display little or no bias on the right side, a group of simplicity-driven participants consistently selecting the advisor with lower complexity on the top, and a smaller group of participants whose advisor choices are close to random.

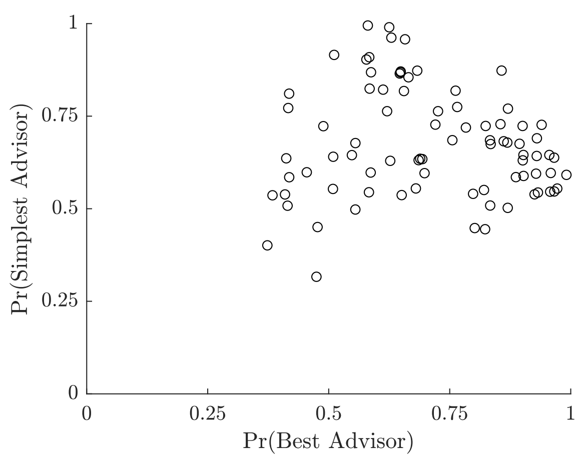

Figure 14: Distribution of participants' advisor choices: probability of choosing the best advisor (based on instrumental value) and simplest advisor (based on the complexity score).

Finally, we use the observable characteristics of the participants to predict the heterogene-

Method: OLS, Dependent variable: Polarization score

	Baseline (1)	Full (2)	Baseline (3)	Full (4)
Risk attitude (Holt and Laury)	−0.52	−0.50	−0.27	−0.26
	(0.16)	(0.16)	(0.24)	(0.25)
Fluid intelligence (Raven test)	0.13	0.10	0.20	0.07
	(0.11)	(0.14)	(0.12)	(0.15)
Familiar with Bayes rule	0.03	0.02	0.10	0.12
	(0.10)	(0.10)	(0.11)	(0.09)
Analytical studies	0.09	0.10	0.06	0.07
	(0.09)	(0.10)	(0.10)	(0.11)
LOT-R scale		−0.03		−0.06
		(0.04)		(0.05)
SUPERSTITION scale		−0.03		−0.01
		(0.04)		(0.05)
RISK scale		−0.02		−0.07
		(0.04)		(0.04)
Observations	63	63	63	63
Demographic Controls			✓	✓

Table 4: Polarization score.

ity in the polarization score presented in Figure 13b . We combine demographic information with two additional tasks (Holt-Laury test of risk attitude and Raven matrices as a measure of cognitive ability) and a final questionnaire with questions about the field of study, mathematical literacy, and other tests (Revised Life Orientation Test, LOT-R, as a measure of optimism, questions on superstition, and questions on risk attitude). Table 4 shows that neither of our measures of mathematical aptitude and cognitive style is significantly associated with our measure of within-subject polarization. Risk attitude, measured by the Holt and Laury test, shows a negative coefficient (a high risk seeking score is associated with a low polarization score), but the effect disappears once we introduce the demographic controls.

7 Discussion and Conclusion

Opinions about proposed policies and pertinent issues often become polarized. The literature provides several explanations of the phenomenon, including (among others) preference for information which confirms existing beliefs, imperfect memory, and interpretation of ambiguous evidence as confirming existing beliefs. In this paper we explore a new source of belief polarization that arises as a consequence of the state-pooling effect when information is costly to acquire.

We find that the valuation of the status quo plays an important role in determining the direction of belief updating, as it directly affects the information acquisition strategy. In our interpretation, the agent partitions the states of the world into categories instead of distinct states. This partition into categories is determined exactly by the valuation of the status quo. If the two agents have different valuations of the status quo, their information acquisition is such that they might diverge in their opinions ex ante (before the realization of the private signal).

The large number of assumptions required by the model may cast doubts on whether belief divergence can emerge from human behavior. We introduce an experiment in which we manipulate the value of the status quo, and we observe that this exogenous variation is able to generate belief polarization. We qualitatively replicate the model's prediction, and observe that the magnitude of the polarization is lower than predicted. We explore the possible drivers of this difference and conclude that intrinsic (non-instrumental) preferences for information comprise the leading factor.

Our paper sheds new light on the problem of opinion polarization in society. Although our analysis focuses on beliefs, the implications of our results easily extend to other variables of interest, including actions and ability to infer agents' preferences from search behavior. In terms of inference of the agents' type from search behavior, the reader can find more details in Appendix L. We consider a platform, similar to Facebook, that can access both actions (likes) and information collection (search) of its users, and we show that the search behavior can be a powerful predictor of the agent's private type.

We acknowledge the limitations of our model and experiment and encourage further exploration of this important phenomenon in several directions. Our model considers individual decision making, and there is space for extensions in different strategic environments, including strategic voting and team coordination. Another limit of our analysis lies in the restriction to the binary action space, and we encourage exploration of the problem with larger action and state spaces: this feature would allow the creation of several endogenous categories and provide a connection with models of categorical thinking. On the experimental side, our design is restricted to binary information decisions (pick one out of two advisors, or buy-no-buy choices), despite the model being much more flexible. We encourage extension of the paradigm to different settings in which a larger signal space is associated with a cost for the signal accuracy (either explicit, in experimental currency, or implicit, for example time required to process the information available). Finally, on the empirical side, we encourage future research to test the implications of our model on referendum data, either with direct intervention on the sources of information available, or using identification strategies that capture different ease of access to media.

References

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

Anobile, Giovanni, Guido Marco Cicchini, and David C. Burr. 2012. "Linear Mapping of Numbers onto Space Requires Attention." *Cognition*, 122(3): 454–459.

Beaudry, Paul, and Tim Willems. 2018. "On the macroeconomic consequences of overoptimism." *National Bureau of Economic Research*, No. w24685.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak. 1964. "Measuring utility by a single-response sequential method." *Behavioral science*, 9(3): 226–232.

Bénabou, Roland. 2015. "The economics of motivated beliefs." *Revue d'économie politique*, 125(5): 665–685.

Bernhardt, Dan, Stefan Krasa, and Mattias Polborn. 2008. "Political polarization and the electoral effects of media bias." *Journal of Public Economics*, 92(5-6): 1092–1104.

Blackwell, David, and Lester Dubins. 1962. "Merging of Opinions with Increasing Information." *The Annals of Mathematical Statistics*, 33(3): 882–886.

Caplin, Andrew, and Mark Dean. 2015. "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review*, 105(7): 2183–2203.

Caplin, Andrew, Mark Dean, and John Leahy. 2019. "Rational inattention, optimal consideration sets and stochastic choice." *The Review of Economic Studies*, 86(3): 1061–1094.

Charness, Gary, Ryan Oprea, and Sevgi Yuksel. 2018. "How do people choose between biased information sources? Evidence from a laboratory experiment." *Working Paper.*

Che, Yeon-Koo, and Konrad Mierendorff. 2019. "Optimal dynamic allocation of attention." *American Economic Review*, 109(8): 2993–3029.

Chinn, Sedona, P. Sol Hart, and Stuart Soroka. 2020. "Politicization and Polarization in Climate Change News Content, 1985-2017." *Science Communication*, 42(1): 112–129.

Cohen, Jacob, Patricia Cohen, Stephen G. West, and Leona S. Aiken. 2013. *Applied multiple regression/correlation analysis for the behavioral sciences.* Routledge.

Cover, Thomas, and Joy Thomas. 2012. *Elements of information theory.* John Wiley & Sons.

Dal Bó, Ernesto, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne. 2018. "Economic losers and political winners: Sweden's radical right." *Unpublished manuscript, Department of Political Science, UC Berkeley.*

Dixit, Avinash K., and Jörgen W Weibull. 2007. "Political polarization." *Proceedings of the National Academy of Sciences*, 104(18): 7351–7356.

Fetzer, Thiemo. 2019. "Did austerity cause Brexit?" *American Economic Review*, 109(11): 3849–86.

Gentzkow, Matthew, Jesse Shapiro, and Matt Taddy. 2016. "Measuring Polarization in High-Dimensional Data: Method and Application to Congressional Speech." NBER Working Paper Series.

Gerber, Alan, and Donald Green. 1999. "Misperceptions about perceptual bias." *Annual Review of Political Science*, 2(1): 189–210.

Gigerenzer, Gerd, and Wolfgang Gaissmaier. 2011. "Heuristic decision making." *Annual review of psychology*, 62: 451–482.

Greiner, Ben. 2015. "Subject pool recruitment procedures: organizing experiments with ORSEE." *Journal of the Economic Science Association*, 1(1): 114–125.

Guney, B., M. Richter, and M. Tsur. 2018. "Aspiration-based choice." *Journal of Economic Theory*, 176: 935–956.

Harrison, Glenn W., Jimmy Martínez-Correa, and J. Todd Swarthout. 2013. "Inducing risk neutral preferences with binary lotteries: A reconsideration." *Journal of Economic Behavior & Organization*, 94: 145–159.

Hollingworth, Harry Levi. 1910. "The Central Tendency of Judgment." *The Journal of Philosophy, Psychology and Scientific Methods*, 7(17): 461–469.

Holt, Charles A., and Susan K. Laury. 2002. "Risk aversion and incentive effects." *American economic review*, 92(5): 1644–1655.

Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47: 263–291.

Kareev, Yaakov, Sharon Arnon, and Reut Horwitz-Zeliger. 2002. "On the misperception of variability." *Journal of Experimental Psychology: General*, 131(2): 287.

Khaw, Mel Win, Ziang Li, and Michael Woodford. 2019. “Cognitive Imprecision and Small-Stakes Risk Aversion.” *NBER Working Paper*, no. 24978.

Kőszegi, B., and M. Rabin. 2006. “A model of reference-dependent preferences.” *The Quarterly Journal of Economics*, 121(4): 1133–1165.

Leiserowitz, Anthony, Edward Maibach, Seth Rosenthal, John Kotcher, Matthew Ballew, Matthew Goldberg, Abel Gustafson, and Parrish Bergquist. 2019. “Politics & Global Warming.” *Yale University and George Mason University. New Haven, CT: Yale Program on Climate Change Communication.*

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2020. “Rational Inattention: A Review.” *CEPR Discussion Papers (No. 15408).*

Maltz, A. 2020. “Exogenous Endowment—Endogenous Reference Point.” *The Economic Journal*, 130(625): 160–182.

Manzini, P., and M. Mariotti. 2012. “Categorize then choose: Boundedly rational choice and welfare.” *Journal of the European Economic Association*, 10(5): 1141–1165.

Masatlioglu, Yusufcan, and Efe A. Ok. 2014. “A Canonical Model of Choice with Initial Endowments.” *The Review of Economic Studies*, 81(2): 851–883.

Masatlioglu, Yusufcan, A. Yesim Orhun, and Collin Raymond. 2017. “Intrinsic information preferences and skewness.” *Ross School of Business Paper.*

Matějka, F., and A. McKay. 2015. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” *American Economic Review*, 105(1): 272–298.

Matějka, Filip, and Alisdair McKay. 2015. “Rational inattention to discrete choices: A new foundation for the multinomial logit model.” *American Economic Review*, 105(1): 272–298.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2008. *Polarized America: The dance of ideology and unequal riches.* MIT Press Books.

Mosteller, Frederick, and Philip Nogee. 1951. “An Experimental Measurement of Utility.” *Journal of Political Economy*, 59: 371–404.

Nimark, Kristoffer P., and Savitar Sundaresan. 2019. “Inattention and Belief Polarization.” *Journal of Economic Theory*, 180: 203–228.

Oliveros, Santiago, and Felix Várdy. 2015. “Demand for slant: How abstention shapes voters’ choice of news media.” *The Economic Journal*, 125(587): 1327–1368.

Ortoleva, Pietro. 2010. “Status quo bias, multiple priors and uncertainty aversion.” *Games and Economic Behavior*, 69(2): 411–424.

Perego, Jacopo, and Sevgi Yuksel. 2018. "Media competition and social disagreement." *Working Paper.*

Poole, Keith T., and Howard Rosenthal. 1984. "The polarization of American politics." *The Journal of Politics*, 46(4): 1061–1079.

Prior, Markus. 2013. "Media and political polarization." *Annual Review of Political Science*, 16: 101–127.

Samuelson, William, and Richard Zeckhauser. 1988. "Status Quo Bias in Decision Making." *Journal of Risk and Uncertainty*, 1(1): 7–59.

Savage, Leonard J. 1954. *The foundations of statistics.* Dover Books.

Scheier, Michael F., Charles S. Carver, and Michael W. Bridges. 1994. "Distinguishing optimism from neuroticism (and trait anxiety, self-mastery, and self-esteem): a reevaluation of the Life Orientation Test." *Journal of personality and social psychology*, 67(7): 1063.

Sims, Christopher A. 1998. "Stickeness." *Carnegie-Rochester Conference Series on Public Policy*, 49: 317–356.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3): 665–690.

Smith, Matthew. 2019. "Most voters have only become more sure about their EU referendum vote." *https://yougov.co.uk/topics/politics/articles-reports/2019/04/16/most-voters-have-only-become-more-sure-about-their.*

Steiner, Jakub, Colin Stewart, and Filip Matějka. 2017. "Rational Inattention Dynamics: Inertia and Delay in Decision-Making." *Econometrica*, 85(2): 521–553.

Suen, Wing. 2004. "The self-perpetuation of biased beliefs." *The Economic Journal*, 114(495): 377–396.

Weber, Thomas A. 2003. "An exact relation between willingness to pay and willingness to accept." *Economics Letters*, 80(3): 311–315.

A Appendix: Derivation of formula 5

The agent's posterior belief about the payoff of the new policy v, given the fixed state s^* for option $i \in \{\text{status quo, new policy}\} = \{1, 2\}$ is

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \mathcal{P}(i = 1|s^*)\mathbb{E}(v|i = 1) + \mathcal{P}(i = 2|s^*)\mathbb{E}(v|i = 2).$$

We use the following lemma in order to substitute for the conditional probabilities $\mathcal{P}(i|s^*)$ $\forall i$.

Lemma 1. *Conditional on the realized state of the world $s \in S$, the probability of choosing a new policy for $\lambda > 0$ is implicitly defined by:*

$$\mathcal{P}(i=1|s) = \frac{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}},$$

the probability of choosing the status quo is

$$\mathcal{P}(i=2|s) = \frac{(1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}},$$

where $\mathcal{P}(i=1)$ is the unconditional probability of choosing a new policy.

When $\lambda = 0$, the agent chooses from two available options: the new policy or the status quo, the option with the highest value with probability one.

Proof. Lemma 1 is a direct consequence of Theorem 1 from Matějka and McKay (2015). □

Thus, after the substitution for the conditional probabilities and applying Bayes rule, we obtain

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \frac{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \cdot \sum_{s=1}^{n} v_s g_s \frac{e^{\frac{v_s}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} +$$
$$+ \frac{(1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \cdot \sum_{s=1}^{n} v_s g_s \frac{e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

Lemma 1 shows that

$$\mathcal{P}(i=1|s^*) = \frac{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

Thus,

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

B Appendix: Proof of proposition 1

First we prove the following lemma, which we further use to prove Proposition 1.

Lemma 2. *Relations* $\mathcal{P}(i=1|s^*) \gtreqless P(i=1)$ *for* $0 < \mathcal{P}(i=1) < 1$ *are equivalent to* $v_{s^*} \gtreqless R$.

Proof. After substitution for the conditional probabilities, the conditions $\mathcal{P}(i=1|s^*) \gtreqless P(i=1)$ can be rewritten as

$$\frac{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \gtreqless \mathcal{P}(i=1),$$

which are equivalent to

$$\left(\mathcal{P}(i=1) - \mathcal{P}^2(i=1)\right)\left(e^{\frac{v_{s^*}}{\lambda}} - e^{\frac{R}{\lambda}}\right) \gtreqless 0.$$

For $0 < \mathcal{P}(i=1) < 1$ the term in the first parenthesis is always positive. Therefore, the left hand side of the inequality is positive when $v_{s^*} > R$ and negative for $v_{s^*} < R$. □

Now we can continue with the proof of Proposition 1.

Proof. In order to solve the agent's problem given by equations 1 - 4 we need to find $\mathcal{P}(i=1)$ and $\mathcal{P}(i=2)$ defined as $\mathcal{P}(i=2) = 1 - \mathcal{P}(i=1)$. These probabilities have to be internally consistent, i.e., $\mathcal{P}(i) = \sum_{s=1}^{n} \mathcal{P}(i|s)g_s$. After dividing both sides of these conditions by $P(i)$ we obtain the following conditions

$$1 = \sum_{s=1}^{n} \frac{e^{\frac{v_s}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} g_s, \quad \text{if } \mathcal{P}(i=1) > 0,$$

$$1 = \sum_{s=1}^{n} \frac{e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} g_s, \quad \text{if } \mathcal{P}(i=2) > 0.$$

The difference in these two equations is

$$\sum_{s=1}^{n} \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} g_s = 0.$$

For k, which holds that $v_k \leq R \leq v_{k+1}$ we can further write the equation above as

$$\frac{e^{\frac{v_k}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_k}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} v_k g_k = -\sum_{s \neq k} \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} v_k g_s. \tag{8}$$

We will use the last equation for determining the sign of $\Delta(s^*)$, which can be written as

$$\Delta(s^*) = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \sum_{i=1}^{n} v_s g_s,$$

$$\Delta(s^*) = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \sum_{i=1}^{n} v_s g_s \frac{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}},$$

$$\Delta(s^*) = \sum_{i=1}^{n} v_s g_s \frac{(\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1))(e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}})}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}},$$

$$\Delta(s^*) = (\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1)) \cdot \sum_{s=1}^{n} v_s g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

Substituting equation (8) into the sum in the last equation we obtain

$$\Delta(s^*) = (\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1)) \left[\sum_{s \neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right].$$

The expression in the square brackets is positive, because for the above-defined k the sign of $(v_s - v_k)$ and the sign of $e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}$ are the same. Hence $\Delta(s^*)$ has the same sign as $(\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1))$ that further, by Lemma 2, has the same sign as $(v_{s^*} - R)$. □

C Appendix: Proof of proposition 2

Proof. We are interested in the monotonicity of $\Delta(s^*)$ when the true state of the world s^* is changing. In appendix B we derived that

$$\Delta(s^*) = (\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1)) \left[\sum_{s \neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right].$$

Let us consider two states of the world s_1^* and s_2^*, such that $s_1^* > s_2^*$. Demonstrating that $\Delta(s_1^*) - \Delta(s_2^*) \geq 0$ would prove the monotonicity of $\Delta(s^*)$.

$$\begin{aligned}\Delta(s_1^*) - \Delta(s_2^*) = & \left[\sum_{s \neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right] \cdot \\ & \cdot (\mathcal{P}(i=1|s_1^*) - \mathcal{P}(i=1) - \mathcal{P}(i=1|s_2^*) + \mathcal{P}(i=1)) = \\ & = \left[\sum_{s \neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right] \cdot \\ & \cdot (\mathcal{P}(i=1|s_1^*) - \mathcal{P}(i=1|s_2^*)) = \\ & = \left[\sum_{s \neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right] \cdot \\ & \cdot \left(\frac{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \frac{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right).\end{aligned}$$

The term in the square brackets is positive, so the sign of $\Delta(s_1^*) - \Delta(s_2^*)$ is determined by the sign of the term in the round brackets.

Let us show that

$$\frac{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \frac{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} > 0.$$

The last inequality is equivalent to

$$\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}}\left(\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}}+(1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}\right)-$$
$$-\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}}\left(\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}}+(1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}\right)>0,$$

$$(1-\mathcal{P}(i=1))\,e^{\frac{v_{s_1^*}}{\lambda}}e^{\frac{R}{\lambda}}-(1-\mathcal{P}(i=1))\,e^{\frac{v_{s_2^*}}{\lambda}}e^{\frac{R}{\lambda}}>0,$$

which, in turn, is equivalent to

$$e^{\frac{v_{s_1^*}}{\lambda}}>e^{\frac{v_{s_2^*}}{\lambda}}.$$

The last inequality holds, so $\Delta(s^*)$ is an increasing function of s^*.

$\square$

D Appendix: Comparative statics

In this section we explore the influence of the model parameters (marginal cost of information, value of status quo) on the magnitude of the change in the mean of beliefs. Specifically, we are interested in the following questions. How much does the expected posterior belief about the value of the new policy differ from the prior expected value? What is the role of the cost of information? Does the model predict that the agents become more polarized in a situation with a higher marginal cost of information? Does the actual valuation of the status quo have an influence on the value of $\Delta(s^*)$ or does it have an influence only on whether the agent is updating towards or away from the realized value of the new policy? All these questions are too complex to be answered analytically for the general case. Therefore, we take advantage of an example with three states and two actions. This problem is a simple benchmark and its solution exhibits the basic features of solutions to the problems with n states and 2 actions. The solution we analyze in this section is symbolic.

In the scenarios that we consider we use several different values of the status quo R and of marginal costs of information λ. All the parameter values are summarized in Table 5.

v_1	v_2	v_3	g_1	g_2	g_3	R_1	R_2	λ_1	λ_2	λ_3	λ_4
0	$\frac{1}{2}$	1	$g \in (0, \frac{2}{3})$	$\frac{1}{3}$	$\frac{2}{3} - g$	$\frac{3}{8}$	$\frac{5}{8}$	$\frac{1}{8}$	$\frac{1}{4}$	$\frac{3}{8}$	1

Table 5: Parameters used in this section

Note that keeping prior probability of state 2, g_2, fixed, we can vary prior probability of state 1, g only between $(0, \frac{2}{3})$. Also, $\mathbb{E}v$ can vary only from $\frac{1}{6}$ to $\frac{5}{6}$. To solve the problem (1)-(4) it is necessary to find the unconditional probabilities $\mathcal{P}(i = 1)$ and $\mathcal{P}(i = 0)$, which we then use for finding the conditional probabilities. First, for a given set of parameters, the unconditional probability $\mathcal{P}(i = 1)$ as a function of $\mathbb{E}v$ for different values of λ is shown in Figure 15.

For $\mathbb{E}v$ close to $\frac{1}{6}$ and $\frac{5}{6}$, the agent does not process any information and chooses with certainty the status quo and the new policy, respectively. With increasing marginal cost of information, the area in which she chooses with certainty grows. In the middle region, the agent acquires information, and the unconditional probability of selecting the new policy is

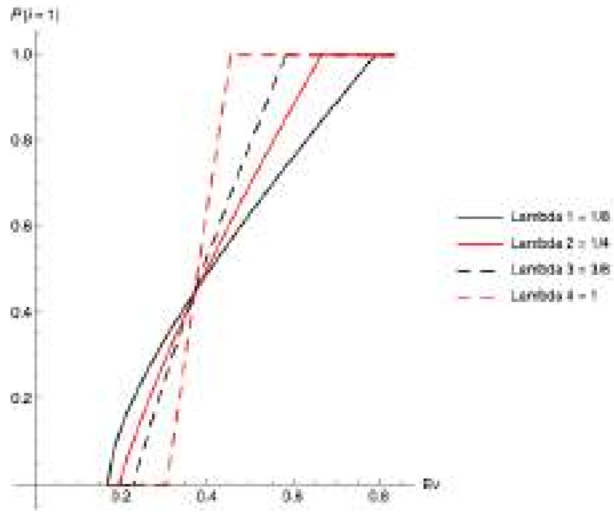

Figure 15: $\mathcal{P}(i = 1)$ as function of $\mathbb{E}v$ for different λ, R_1

an increasing function of the prior expected value of the new policy. With an increase in marginal cost of information λ, the small changes in $\mathbb{E}v$ translate into larger changes in the $\mathcal{P}(i = 1)$.

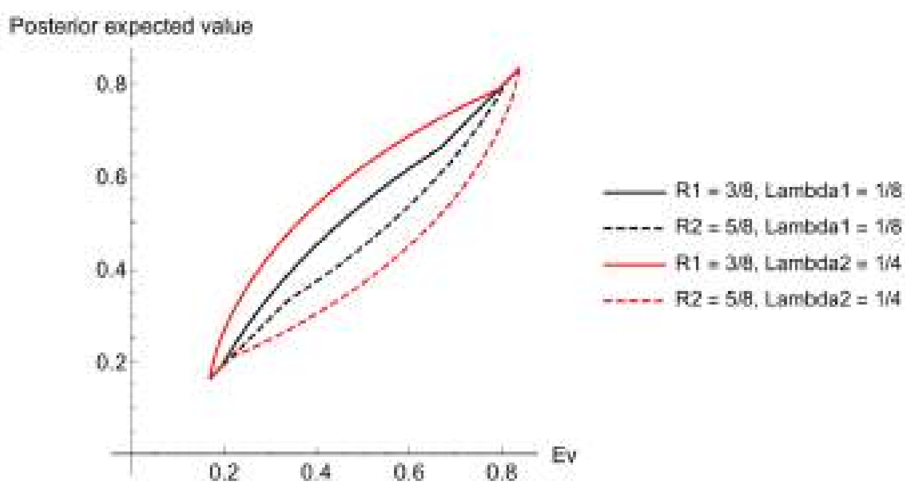

Figure 16: $\mathbb{E}_i[\mathbb{E}(v|i)|s^* = 2]$ as a function of $\mathbb{E}v$ for different levels of R and λ. The solid lines are the case with R_1 and dashed with R_2. Black corresponds to cases with λ_1 and red is used for λ_2.

In order to observe how the change in the mean of beliefs is influenced by the parameters, see Figure 16, which depicts $\mathbb{E}_i[\mathbb{E}(v|i)|s^* = 2]$ as a function of $\mathbb{E}v$ for different levels of R and λ. In line with Proposition 1, different R change the direction of updating. Moreover, for this example, the role of the marginal cost of information is clear from this figure. The lower the marginal cost of information ($\lambda_2 < \lambda_1$), the further away the prior expected values are from the posterior expected value of the new policy. This is also manifested by the fact that the agent is learning even for the prior beliefs, where she was not acquiring information for λ_1. Therefore, in our example, when the cost of information is lower, the polarization of agents is more severe.

Figure 17 provides another perspective on how the change in the mean of beliefs is influenced by the parameters, and directly depicts $\Delta(s^* = 2)$ as a function of $\mathbb{E}v$ for $R_1 = 3/8$ and $\lambda_2 = 1/4$. The figure corresponds to the situation when the agent is updating to the right, towards v_3. The red region indicates the region where the agent is updating away from the realized value of the new policy v_2. An interesting insight is that since the maximal value of $\Delta(s^* = 2)$ is achieved for prior beliefs, which are close to the payoff associated with true state $v_{s^*=2} = 1/2$, it is possible that someone who is updating towards the realized value of the new policy can move her belief from something lower than v_2 to something higher than v_2. Moreover, we observe that the more the agent is optimistic about the new policy, she updates less when she updates towards v_3 (see the decreasing part of Figure 17). This is not surprising in this example, due to keeping the g_2 fixed and $\sum_{s=1}^{3} g_s = 1$.

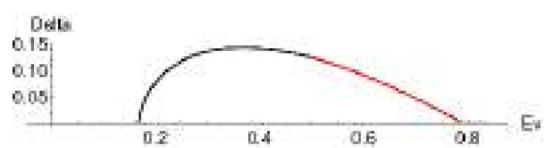

Figure 17: $\Delta(s^* = 2)$ as a function of $\mathbb{E}v$ for $R_1 = 3/8$ and $\lambda_2 = 1/4$. The red area depicts the region of updating in the opposite direction from the realized value.

This figure also provides an example in which two agents diverge in the beliefs while they update in the same direction. We remind the reader of the illustrative figure of the situation in Figure 4, where we assume that the agents' valuations of the current policy are the same. Thus, since two agents differ only in their prior expectations about the new policy, it is sufficient to look at how a single agent's change in the mean of beliefs $\Delta(s^* = 2)$ depends on $\mathbb{E}v$. We are interested in finding two prior expected beliefs for which there is divergence of posterior beliefs. To do so, we need to find two points such that Δ for the left point is lower than Δ for the right point. In our example the red part of the plot is a decreasing function.

This means that, in our example, two agents updating in the same direction with the same valuation of the status quo might diverge in their opinions only when they are updating correctly. However, at the black part of the plot it is easy to find two points at which the agents diverge in their opinions.

E Appendix: Experimental Design and Procedure

E.1 Experimental Interface and Payment

Task 1 - Colorblind advisor game

If one round from this part is selected for the bonus payment, a subject receives the $15 bonus with the percentage probability equal to the number of points that she collected in that round. Since each line counts as a separate decision, one of which might be randomly drawn for payment, truthful revelation is strictly optimal. We constrain subjects to have at most one switching point for every advisor.

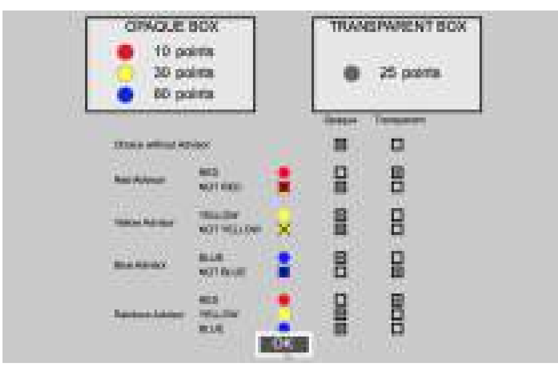

Task 1, Screen 1: Action choice

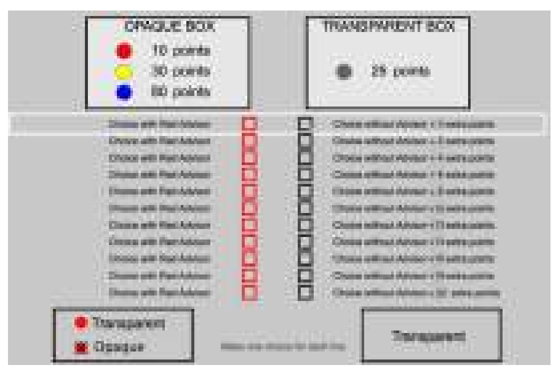

Task 2, Screen 2: WTA for each advisor

Figure 18: Task 1: Colorblind advisor game. Left: Subjects choose an action (box) contingent on the advisor and signal received. The possible values of each action are indicated on the top of the screen. Each state (ball color) is equally likely to occur. Right: Subjects indicate for each advisor the willingness to accept renunciation of its signal in a series of binary choices (BDM method). At most, one switch is allowed. Action choices selected in the previous stage are reported on the bottom of the screen.

Task 2 - Imprecise advisor game If one round from this part is selected for the bonus payment, subjects receive the $15 bonus with the percentage probability equal to the number of points that she collected in that round.

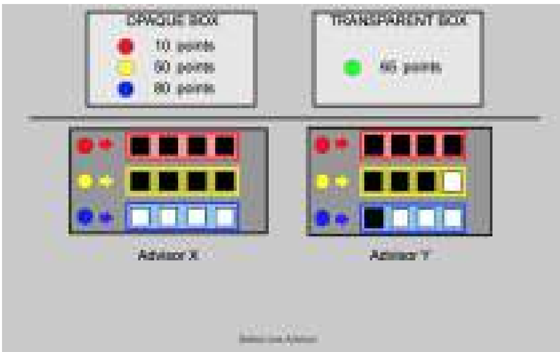

Task 1, Screen 1: Advisor choice

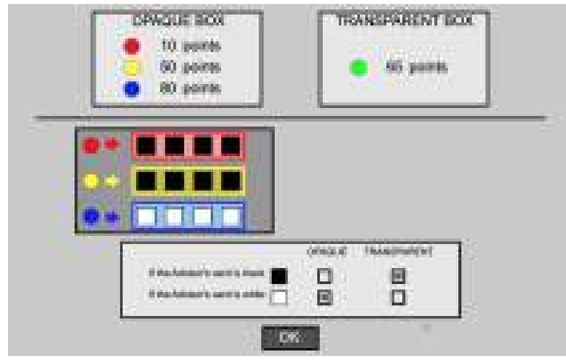

Task 2, Screen 2: Action choice

Figure 19: Task 2: Imprecise advisor game. Left: Subjects choose one signal structure (advisor) between the two options available. Each advisor is a triplet of state-contingent signal probabilities. Right: Subjects indicate the signal-contingent action for each signal (strategy method).

Task 3 - Card color prediction game If one round from this part is selected for the bonus payment, the computer randomly determines the state and realized signal, and subjects receive the \$15 bonus with the percentage probability determined by the quadratic loss scoring rule.

Task 4 - Ball color prediction game If one round from this part is selected for the bonus payment, the computer randomly determines the state and realized signal, and subjects receive the \$15 bonus with the percentage probability determined by the quadratic loss scoring rule.

E.2 Randomization

In all the tasks we randomize the order of the trials. For task 2 only, the first 3 trials are randomly drawn from the subset of trials where both the advisors provide certainty (in order to facilitate the transition from task 1 to task 2).

In task 1 we randomize the order of the four hiring screens within each trial.

In task 2 we randomize the positions of the two advisors on the screen.

In tasks 2, 3, and 4 we randomize the advisors' card colors (black and white). This means that the signal-contingent choice in the second part of the round requires the subjects to analyze every advisor separately, since the colors do not convey any intrinsic message, and

Task 3: Beliefs over signal likelihood

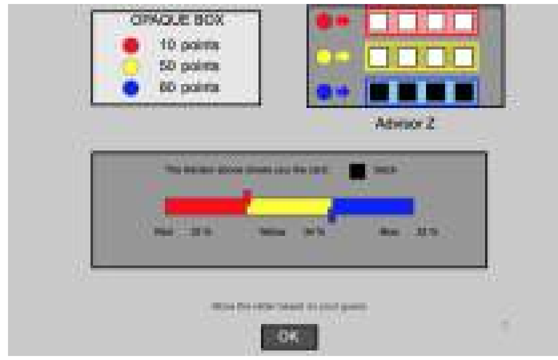

Task 4: Beliefs over state likelihood

Figure 20: Left: Task 3 (Card color prediction game). Subjects indicate the likelihood of observing each signal (card color) for the given advisor. Right: Task 4 (Ball color prediction game). Subjects indicate the likelihood of each state (ball color) given an advisor and signal. In both tasks subjects move the slider(s) and receive a number of points according to the quadratic loss scoring rule described in the instructions.

this procedure reduces the concern regarding inertia in the evaluation of the advisor and in actions.

E.3 Subject understanding

Instructions were provided on both the computer screen, as slides that can be browsed by each subject at the desired pace, and as a paper printout. The two versions of the instructions contained the same information verbatim. Before proceeding with every section of the experiment, subjects were required to correctly answer all the multiple-choice questions of the comprehension test to check understanding of the instructions. The number of questions ranged from two to four for every section, and subjects received a one-minute timeout before having a new attempt. Subjects were initially informed about the payment structure, the no-deception policy of the laboratory, and that choices in one section of the experiment did not affect any other section, or the questionnaire. A small number of subjects were recruited for each laboratory session (6 on average) in order to facilitate clarification of questions during the experiment.

F Appendix: Questionnaire

At the end of the four tasks, we have an additional section with a Holt and Laury test of risk aversion (Part 1), Raven matrices test of fluid intelligence (Part 2, five matrices of different difficulty), and a series of questions (Part 3), that we show here as they were presented to subjects.

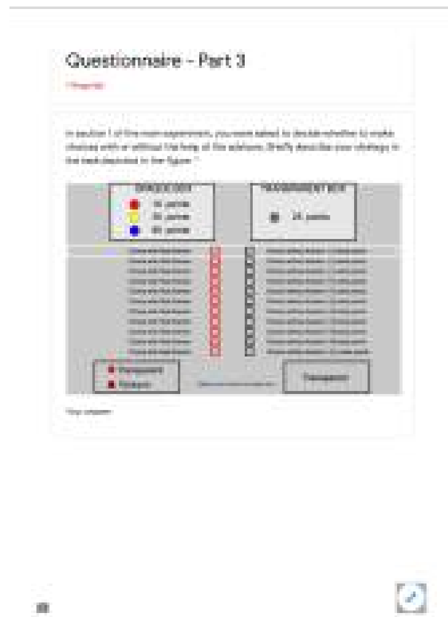

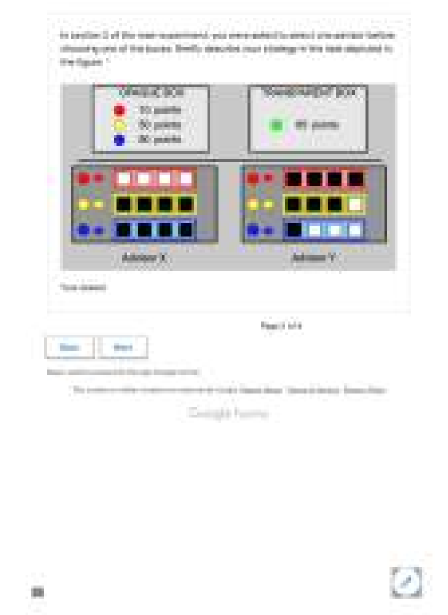

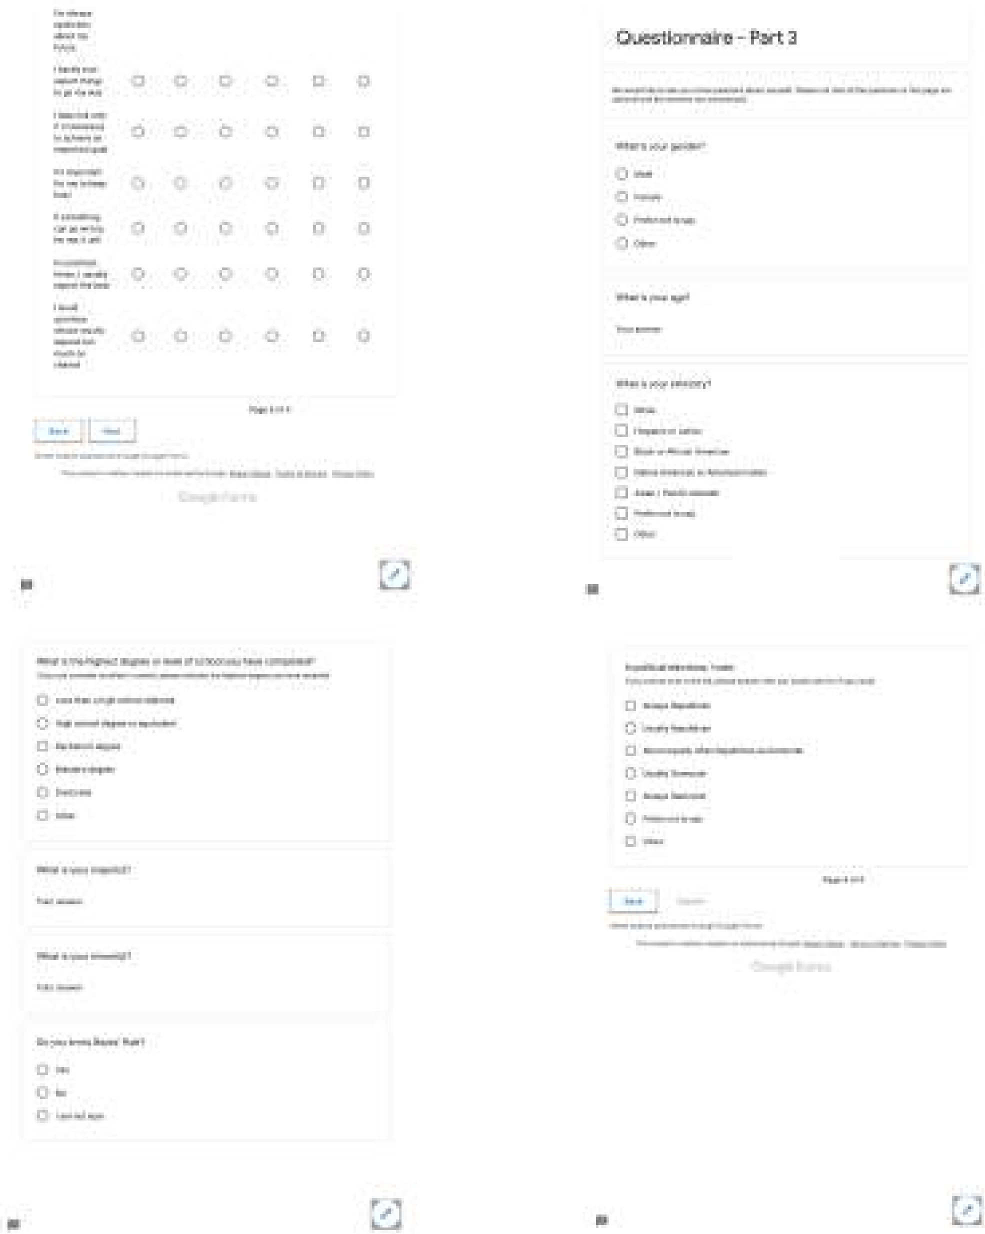
Questionnaire - Part 3

G Appendix: Further Analysis on Advisor Choice

Certain advisors provide an answer to a question of the kind: "Is the state red(/yellow/blue)?" and allow the subject to learn with certainty if a particular state is realized (with probaility one) or not realized (with probability zero).

Figure 23 shows advisor choice in the trials in which both advisors provide certainty. We display separately the trials with different status quo values. When the subjects have to

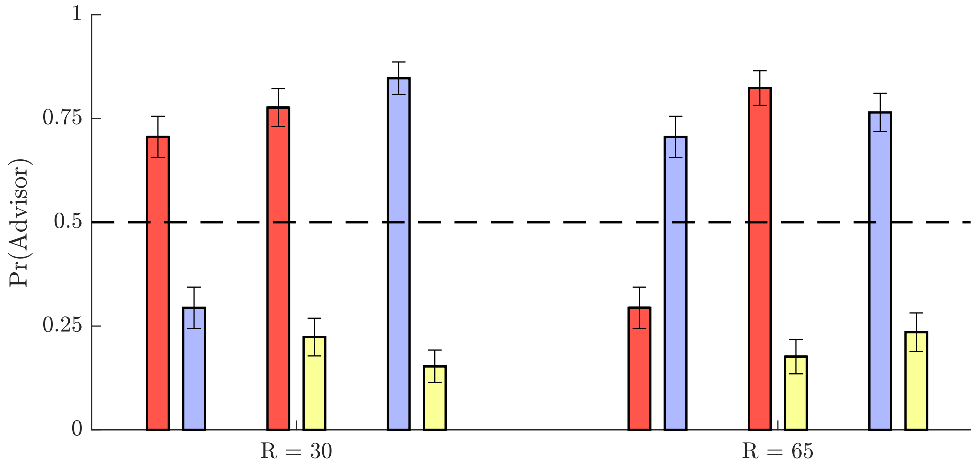

Figure 23: Comparison of advisors providing the ideal state pooling question: "Is the state red/yellow/blue ?" for two different status quo values. The color of the bar shows which state the question is about. The figure demonstrates the state pooling behavior, and also that participants do switch between advisors when it is valuable to do so.

choose between advisors that provide certainty and are also state poolers, that is, between an advisor providing information whether the state is blue and another advisor providing information whether the state is red (first couple of bars for $R = 30$ and $R = 65$), they significantly select the former for the high value of the status quo and latter for the low value of the status quo. This switch between advisors confirms our theoretically predicted state pooling effect. In particular, for a status quo value R the subject wants to learn whether the state-dependent payoff of the new policy is greater or lower than R. When subjects face a choice between a certainty state pooler and certainty advisors, they select on average the certainty state pooler in 74% of the trials.

In the two scenarios of choice between the red-advisor and blue-advisor, we see that each option is chosen by more than one quarter of the participants, and this is at odds with the model's prediction, especially in a trial with a relatively simple problem. This can be interpreted as a general signal of noise in the participants' actions, or as a systematic preference towards information about low or high states. Figure 24 suggests that the latter interpretation can partially explain the pattern. Starting from all the trials, we calculate for every subject the probability of choosing the advisor that is best under low or high status quo value.

If choices are just noisy, we should observe most of the subjects to be clustered around the coordinates 0.5-0.5, which would be consistent both with optimal behavior (always pick the best advisor), and with completely erratic choices (pick randomly). If participants have non-instrumental preferences over skewed sources of information (as shown by Masatlioglu, Orhun and Raymond (2017)), and such preferences are heterogeneous, we would expect a distribution of subjects that systematically deviate towards 1-0 (reveal information about the low state) and 0-1 (about the high state). In fact, we observe that participants deviate in both directions, and some of them also deviate towards lower probabilities in both dimensions - this can be the case when the chosen advisor is the worse choice under both status quo scenarios.

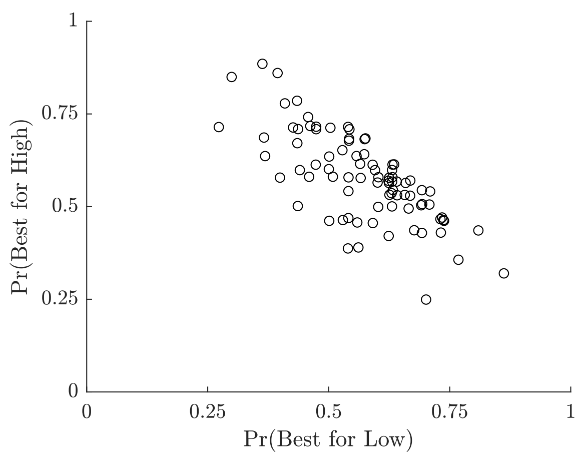

Figure 24: Distribution of participants' advisor choices: probability of choosing the advisor that provides more information about the low or high state in different types of trials.

H Appendix: Further Analysis on Risk Attitude

In Section 5.2 we discussed how risk preferences could represent an explanation for the subjects' deviation from the model's predictions. We provide here further details about how subjects' actions are, on average, not characterized by a significant deviation from risk neutrality.

Figure 25a shows the realized probability of selecting the risky option as a function of the difference in the EVs between the actions. Trials are grouped based on the x-axis value for visualization purposes. The optimal agent would have a sharp jump in probability from 0 (when the value difference is negative) to 1. We observe a smoother transition in our data, suggesting that action probability is modulated by the cost of mistakes, similarly to our discussion in Figure 8b about the choice between advisors. Such a sigmoid curve is normally found in experiments involving choice under risk (Mosteller and Nogee, 1951; Khaw, Li and Woodford, 2019).The indifference point appears close to the trials in which both actions have the same values, suggesting that the participants are overall close to risk neutrality.

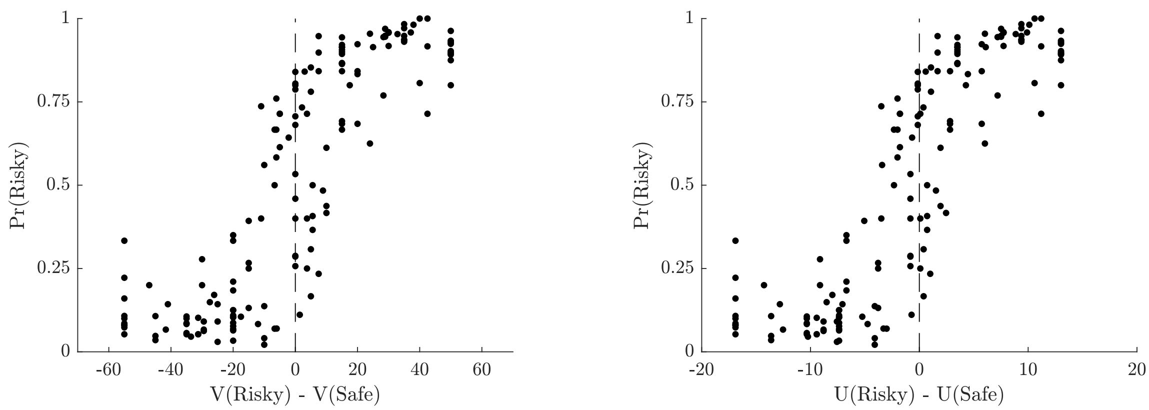

(a) (b)

Figure 25: Task 2: Action selection probability. Left: Action choice under risk neutrality. Observed probability of choosing the risky action in task 2. 6,800 observations unequally divided across 160 cases (2 cases per trial, conditional on the advisor choice). Right: Action choice under risk aversion (best fit). The expected values for each action is replaced with the expected utility, with CRRA utility and the MLE coefficienty $\hat{\alpha} = 0.34$ estimated from the dataset.

We replicate the analysis for the choices made in task 1. An advantage of this dataset is that we observe two types of action choice scenario.

First, when the advisor confirms the color of the hidden ball, the decision maker faces a choice between two degenerate lotteries with different values, for example 80 points for sure (risky action if you know the color is blue) or 60 points for sure (safe action). In these cases, participants pick the best option 90% of the times, confirming the small amount of noise in the action implementation in these simple choices.

Second, all the participants encounter choices with full uncertainty about the color (decision without any hint) or with a hint about two possible colors (e.g. red or yellow with equal chance, but not blue). The participants pick the option with the highest expected value 84% of the times, and we encounter again the sigmoid curve discussed above. The MLE of the risk aversion parameter under CRRA utility is $\hat{\alpha} = 0.52$.

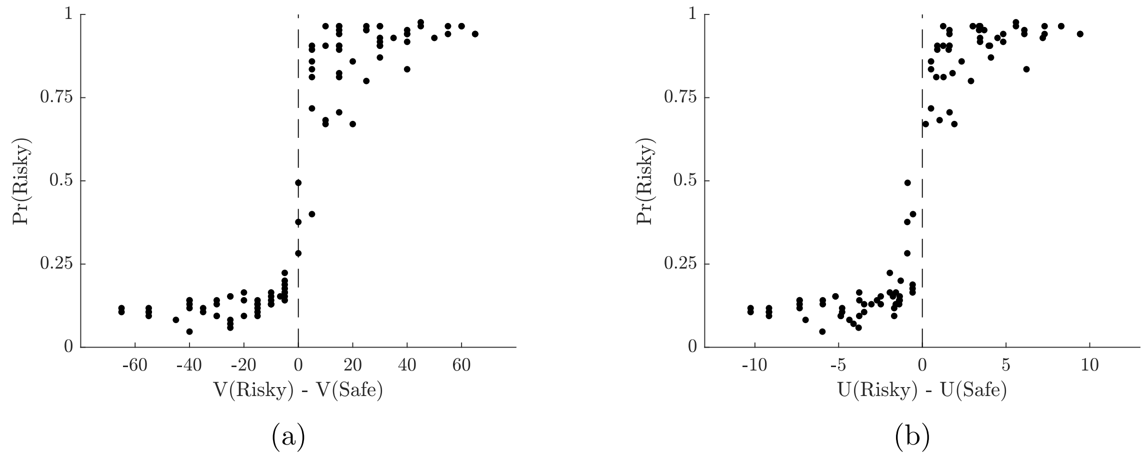

Figure 26: Task 1: Action selection probability. Left: Action choice under risk neutrality. Observed probability of choosing the risky action in task 1. 100 questions (10 questions for each of the 10 trials), 85 observations per question. Right: Action choice under risk aversion (best fit). The expected values for each action is replaced with the expected utility, with CRRA utility and the MLE coefficient $\hat{\alpha} = 0.52$ estimated from the dataset.

I Appendix: Further Analysis of Beliefs

In Section 5.2 we discussed how subjective beliefs could represent an explanation for the subjects' deviation from the model's predictions. We provide here further details about how subjects' subjective beliefs elicited in tasks 3 and 4 display, on average, only mild evidence of conservatism.

In both tasks we observe accurate probability estimates, close to the predictions of an optimal Bayesian agent. Figure 27a shows the subjective estimate of a signal realization (y-axis, averaged across participants) compared to the optimal estimates (x-axis). Similarly, Figure 27b shows the subjective estimate of each of the three possible states in the posterior compared to the unbiased posterior, with different colors in the figure matching the state. In both plots, the 45 degree lines represent our theoretical benchmark and we

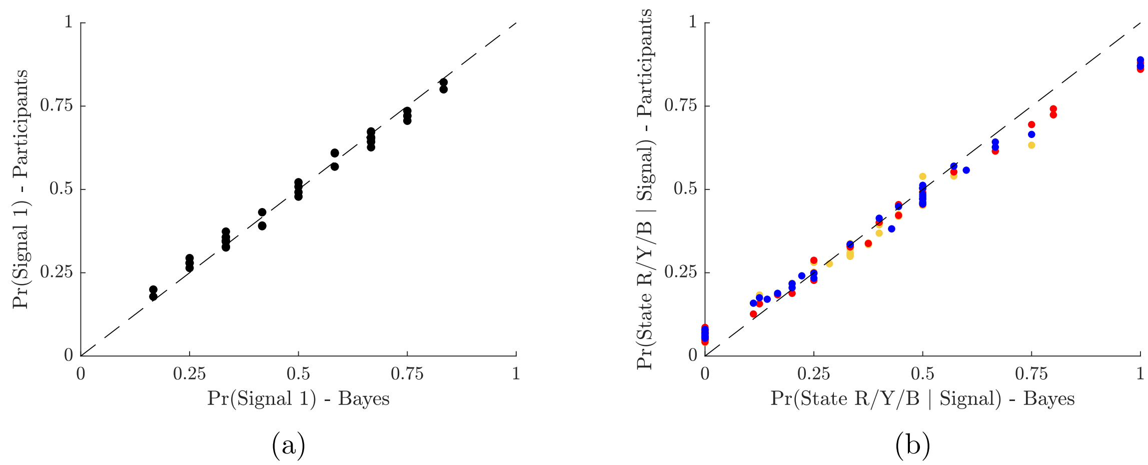

Figure 27: Average subjective beliefs. Left: Estimated probability of receiving a signal realization in Task 3. The plot compares the average of the subjective estimates collected with the optimal estimates of a Bayesian decision maker. 1,700 observations across 20 trials (85 observations per point). Right: Estimated posterior probability of each state in task 4 conditional on the realized signal. Colors indicate which state was estimated (red, yellow, blue). The plot compares the average of the subjective estimates collected with the optimal estimates of a Bayesian decision maker. 20,400 observations across 40 trials (6 observations per trial, 85 observations per point in the plot).

can see that 1) participants are on average accurate in the estimate of probabilities, 2) we do not observe a systematic difference between estimates involving different states (i.e., we do not have evidence of motivated beliefs, Bénabou (2015)), and 3) both tasks show mild evidence of conservatism (central tendency of judgement), as vastly reported in experiments with subjective estimates (Hollingworth, 1910; Anobile, Cicchini and Burr, 2012).

For the signal probability (task 3) a linear fit of the subjective estimates $\hat{p}$ over the true probabilities p returns the coefficients $\hat{p} = 0.041 + 0.918 \cdot p$ with $R^2 = 0.991$. For the state probability (task 4) the linear fit for the whole dataset returns $\hat{p} = 0.058 + 0.825 \cdot p$ with $R^2 = 0.993$. The slopes are not significantly different across the three types of states: $\beta_{red} = 0.831$, $\beta_{yellow} = 0.805$, $\beta_{blue} = 0.827$.

J Appendix: Further Analysis of Willingness to Accept

In this section we add further results from the analysis of the willingness to accept renunciation of an advisor in task 1. We reported in Figure 12 and Table 3 that WTA is characterized by compression, a conservatism in the evaluation of the instrumental value of an advisor that leads to overpayment for the advisor with little or no informative value.

This result is robust across subjects, as displayed in Figure 28. For every subject, we estimate the sensitivity to the instrumental value by using a simple OLS regression of the subjective evaluation $V^j(i)$ over the instrumental value $V^{\text{Bayes}}(I)$. The graph shows the cumulative distribution of the fitted slopes, where 0 indicates no response to the true value and 1 indicates full alignment between the two variables. 82% of the participants show values between 0 and 1.

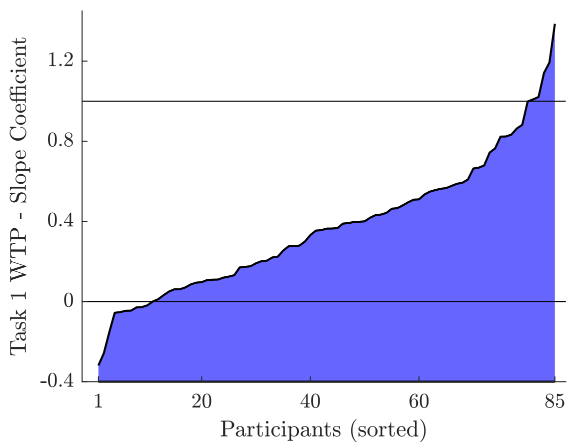

Figure 28: Distribution of participants' responses to the instrumental value of the advisors in Task 1.

Another effect discussed in the paper is the excessive value assigned to blue advisors (that reveal the high-payoff state) and rainbow advisors (that provide full disclosure). Using Equation 6 we can easily recognize that the value of the rainbow advisor is equal to the highest value among the other three advisors. Participants do not seem to follow this rule, as they tend to pay much less for the rainbow advisor. Figure 29 shows the distribution of the differences, within each trial, between the WTA for the rainbow advisor and the maximum of the other three WTA. Participants are willing to pay strictly less 39% of the times, and they are willing to pay strictly more only 17% of the times.

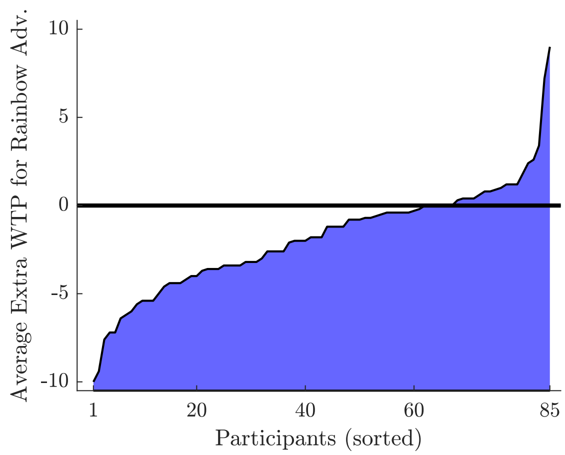

Figure 29: Distribution of participants' average excessive WTP for the rainbow advisor with respect to the highest WTP among the three simple advisors.

Finally, we want to show that the extra WTA for blue and rainbow advisors is due to a fixed premium that participants are willing to add, and not because of different elasticity to the instrumental values. In Table 3 we assumed that advisors' values can have different intercepts but share the same slope. We relax the assumption in a series of regressions displayed in Table 6. Compared to the benchmark model (column 1, same slope and intercepts for all), we notice a major improvement in the fit when we add the advisor-specific intercepts, which is not as much as for advisor-specific slopes.

Method: OLS, Dependent variable: $V^i(I)$

	(1)	(2)	(3)	(4)
Constant	6.66	5.62	6.80	5.65
	(0.166)	(0.231)	(0.167)	(0.257)
Red - constant		-0.193		-0.496
		(0.353)		(0.447)
Blue - constant		3.41		3.42
		(0.349)		(0.422)
Rainbow - constant		2.74		2.96
		(0.373)		(0.505)
$V^{\text{Bayes}}(I)$ - slope	0.372	0.269	-0.123	0.222
	(0.027)	(0.030)	(0.164)	(0.172)
Red - slope			0.252	0.099
			(0.163)	(0.181)
Blue - slope			0.631	0.038
			(0.164)	(0.180)
Rainbow - slope			0.550	0.009
			(0.162)	(0.181)
Trials	All	All	All	All
Observations	2,520	2,520	2,520	2,520

Table 6: Aggregate valuations of information structures in task 1.

K Appendix: Further Analysis of Questionnaire

Method: OLS, Dependent variable: Accuracy of advisor choice in task 2

	Baseline (1)	Full (2)	Baseline (3)	Full (4)
Risk attitude (Holt and Laury)	−0.36	−0.36	−0.28	−0.28
	(0.09)	(0.10)	(0.15)	(0.16)
Fluid intelligence (Raven test)	0.19	0.16	0.22	0.13
	(0.07)	(0.10)	(0.08)	(0.10)
Familiar with Bayes rule	0.09	0.11	0.17	0.20
	(0.05)	(0.06)	(0.06)	(0.06)
Analytical studies	−0.03	−0.03	−0.06	−0.06
	(0.05)	(0.05)	(0.05)	(0.06)
LOT-R scale		−0.01		−0.04
		(0.03)		(0.03)
SUPERSTITION scale		0.02		0.03
		(0.03)		(0.03)
RISK scale		−0.00		−0.03
		(0.03)		(0.03)
Observations	63	63	63	63
Demographic Controls			✓	✓

Method: OLS, Dependent variable: Probability of selecting the simple advisor in task 2

	Baseline (1)	Full (2)	Baseline (3)	Full (4)
Risk attitude (Holt and Laury)	−0.07	−0.08	−0.17	−0.18
	(0.11)	(0.11)	(0.15)	(0.15)
Fluid intelligence (Raven test)	−0.06	−0.08	−0.08	−0.11
	(0.08)	(0.10)	(0.09)	(0.13)
Familiar with Bayes rule	−0.02	−0.01	−0.04	−0.03
	(0.04)	(0.05)	(0.05)	(0.06)
Analytical studies	−0.01	−0.01	0.00	−0.00
	(0.04)	(0.04)	(0.04)	(0.04)
LOT-R scale		0.01		0.01
		(0.02)		(0.03)
SUPERSTITION scale		0.02		0.04
		(0.02)		(0.03)
RISK scale		−0.01		−0.00
		(0.03)		(0.03)
Observations	63	63	63	63
Demographic Controls			✓	✓

Method: OLS, Dependent variable: Probability of selecting the risky action in task 1

	Baseline (1)	Full (2)	Baseline (3)	Full (4)
Risk attitude (Holt and Laury)	0.03	0.03	−0.00	0.00
	(0.08)	(0.08)	(0.10)	(0.10)
Fluid intelligence (Raven test)	0.05	0.06	0.03	0.05
	(0.05)	(0.07)	(0.06)	(0.07)
Familiar with Bayes rule	−0.00	−0.02	−0.05	−0.06
	(0.04)	(0.04)	(0.04)	(0.04)
Analytical studies	−0.01	0.00	0.00	0.01
	(0.03)	(0.03)	(0.03)	(0.03)
LOT-R scale		0.01		0.02
		(0.02)		(0.02)
SUPERSTITION scale		−0.02		−0.02
		(0.02)		(0.02)
RISK scale		−0.01		−0.01
		(0.02)		(0.02)
Observations	63	63	63	63
Demographic Controls			✓	✓

Method: OLS, Dependent variable: Accuracy in beliefs elicitation (tasks 3 and 4)

	Baseline (1)	Full (2)	Baseline (3)	Full (4)
Risk attitude (Holt and Laury)	−0.59	−0.60	−0.40	−0.51
	(0.68)	(0.63)	(0.50)	(0.53)
Fluid intelligence (Raven test)	0.45	0.45	0.34	0.48
	(0.39)	(0.41)	(0.35)	(0.35)
Familiar with Bayes rule	0.26	0.35	0.37	0.42
	(0.20)	(0.20)	(0.26)	(0.30)
Analytical studies	−0.15	−0.21	−0.30	−0.30
	(0.19)	(0.19)	(0.20)	(0.20)
LOT-R scale		−0.03		0.07
		(0.13)		(0.12)
SUPERSTITION scale		0.13		0.11
		(0.14)		(0.15)
RISK scale		0.10		0.14
		(0.09)		(0.11)
Observations	63	63	63	63
Demographic Controls			✓	✓

Method: OLS, Dependent variable: WTP slope in task 1

	Baseline (1)	Full (2)	Baseline (3)	Full (4)
Risk attitude (Holt and Laury)	−0.42	−0.39	−0.12	−0.15
	(0.20)	(0.21)	(0.27)	(0.28)
Fluid intelligence (Raven test)	0.27	0.31	0.31	0.33
	(0.18)	(0.24)	(0.16)	(0.21)
Familiar with Bayes rule	0.20	0.17	0.14	0.15
	(0.10)	(0.11)	(0.12)	(0.13)
Analytical studies	−0.20	−0.16	−0.20	−0.16
	(0.10)	(0.11)	(0.09)	(0.10)
LOT-R scale		0.02		0.03
		(0.07)		(0.06)
SUPERSTITION scale		−0.06		−0.03
		(0.05)		(0.06)
RISK scale		−0.02		−0.01
		(0.05)		(0.06)
Observations	63	63	63	63
Demographic Controls			✓	✓

L Appendix: Prediction of the status quo type

We now look at our framework from the perspective of a platform that wants to infer the type (status quo value) of the decision maker and has access to a dataset with some observable activities. We can divide these activities into two groups: final actions, like *voting* or the choice between the status quo and a new policy, and information acquisition, like *reading newspapers* or selecting an advisor in our design. For a more concrete example, imagine a social media platform like Facebook or Twitter, that has access to a dataset of actions performed by its users. These actions include publicly observable actions (likes, list of friends or followers), but also a series of additional actions (clicks, searches) that involve the process of information acquisition. Are these search activities helpful in improving the prediction of the type of the user, on top of the observable actions?

	Prediction	Data
No information	50.0%	50.0%
Choice only	69.7%	62.6%
Search only	100.0%	68.0%
Search+Choice	100.0%	68.4%
Search+Signal+Choice	100.0%	72.9%

Table 7: Inference of the agent's status quo: predicted and realized accuracy (pairs of trials with expected advisor switch only). The table indicates the accuracy of the prediction of the type (status quo) of the decision maker based on the data available. The model's predictions are based on rational and unbiased agents. The accuracy realized refers to the dataset collected in the laboratory experiment.

We consider separate scenarios in which the platform has access to choices only (opaque or transparent box), searches only (advisor X or advisor Y), or both, under the assumptions of our model (rational decision makers) and in the dataset collected in the laboratory experiment. Table 7 shows the results of this exercise: having access to the search data guarantees a much higher accuracy with respect to the action data, with minor improvements when both datasets are available. When we consider the trials in which we expect to observe an advisor switch (column 2), the type prediction accuracy with advisor choices is 68%, but it is only 62.6% when we observe only the final actions. When both datasets are available, the accuracy increases marginally to 68.4%, with a further improvement if the signal realization (that occurs between search and choice) is also observed.

We can conclude that, in this simple setup, the data about the choice over sources of information is more valuable than the final action from the perspective of an observer who wants to infer the type (status quo value) of the decision maker.

M Appendix: Timeline of the Problem

Our experimental design allows us to estimate how agents evaluate an informative signal structure (advisor), and measure how the subjective evaluation depends on the properties of the signal structure, including instrumental value (expected improvement in the choice process) and non-instrumental properties (ease of interpretation).

The timing of the problem (as in task 2) can be summarized as follows:

1. The agent is informed of the prior $\mathcal{P}(s) = \frac{1}{3}\ \forall s$ and the state-contingent returns $\{v_s\}_S,\ R$.
2. One state is realized, but the agent is unaware of it.
3. The agent is offered two sources of information (advisors) I_1 and I_2.
4. The agent chooses one advisor and discards the other.
5. The selected advisor observes the realized state (ball in the opaque box).
6. The selected advisor returns a binary signal, whose likelihood depends on the realized state.
7. The agent observes the realized signal.
8. The agent chooses one action (opaque or transparent box) and receives the payoff π.
9. The agent plays a lottery and receives the final prize k with probability $\frac{\pi}{100}$.

The problem presented in task 1 is similar up to a change in steps 3 and 4:

3'. The agent is offered one single source of information (advisor) I

4'. The agent indicates how much she is willing to accept renunciation of the advisor.

In the *Colorblind advisor game* (task 1), we elicit the probability w_I such that the agent is indifferent between making a choice after observing the realization of a known signal structure I, and choosing without additional signals but receiving additional w_I tickets to win the prize. In the *Imprecise advisor game* (task 2), we offer pairs of signal structures, and collect binary choices between advisors. If the valuation and choices differ from those a Bayesian expected utility maximizer would display, we would like to pinpoint the source of the deviation. For this reason, we add two control tasks to elicit a subjective signal of beliefs' realization (*Card color prediction game*, task 3) and subjective posterior beliefs (*Ball color prediction game*, task 4). We collect posteriors only after eliciting preference over advisors, so we do not nudge the subjects towards thinking about information valuation in a specific fashion.