



# The Expert and The Charlatan: an Experimental Study in Economic Advice



Theodore Alysandratos, Nottingham Trent University  
Aristotelis Boukouras, University of Leicester  
Sotiris Georganas, City University London  
Zacharias Maniadis, University of Southampton

Working Paper No. 20 06  
July 2020

# The Expert and The Charlatan: an Experimental Study in Economic Advice

Theodore Alysandratos\*, Aristotelis Boukouras†, Sotiris Georganas‡and Zacharias Maniadis§

## *Preliminary*

### **Abstract**

How do people choose what economic advice to heed? We develop a set of validated multiple-choice questions on economic policy problems, to examine empirically the persuasiveness of expert versus populist advice. We define populism as advice that conforms to commonly held beliefs, even when wrong. Two (computerised) advisers suggest answers to each question, and experimental participants are incentivised to choose the most accurate adviser. Do participants choose the high-accuracy adviser ('the Expert'), or the low-accuracy one ('the Charlatan'), whose answers are designed to be similar to the modal participant's priors? Our participants overwhelmingly choose the Charlatan, and this is only slowly and partially reversed with sequential feedback on the correct answer. We develop Bayesian models to determine optimal choice benchmarks, but find that behaviour is best explained by a naive choice model akin to reinforcement learning with high inertia.

Keywords: Democracy, Economic Literacy, Expert Advice, Populism.

JEL Codes: C91, A11

---

\*Department of Economics, Nottingham Trent University, UK.

†**Corresponding Author.** Department of Accounting, Finance and Economics, University of Leicester, UK.  
Email: Aris.Boukouras@leicester.ac.uk

‡Department of Economics, City University London, UK.

§School of Economics, Social and Political Sciences, University of Southampton, UK.

# 1 Introduction

τότ' ἔφη τὰς πόλεις ἀπόλλυσθαι, ὅταν μὴ δύνωνται τοὺς φαύλους ἀπὸ τῶν σπουδαίων διακρίνειν

States fail when they cannot distinguish fools from serious men.

Antisthenes

Public debate is crucial for the functioning of modern democratic societies. However, as Pericles prominently pointed out in his Funeral Oration, democracy requires citizens who are informed about public affairs, not “idiots”, individuals who only have information about their private business. In modern democracies, newspapers and news channels offering curated content ensured some level of information among educated people. Recent technological developments, however, have upset information delivery. Citizens today have access to a wide range of opinions and advice via traditional media as well as via social media, podcasts and websites. In this cacophony, citizens may risk becoming systematically misinformed. Being presented with so much uncurated content, they may be attracted to ‘charlatans’; low quality sources of information that are pandering to the audience’s preconceptions. This form of populism is of high interest to social science today. The key questions are to what extent laypeople can be manipulated by such communication tactics, and whether they are able to distinguish ‘experts’ from ‘charlatans’ in the absence of perfect feedback.

In this paper we adopt the experimental approach, as especially suited to uncover causal relationships, and focus on a well defined question: whose advice do the public *choose* to heed? In particular, who does the public regard as an expert when different people juxtapose their views on an important topic? Experts are individuals who dedicate their lives to understanding a topic deeply, and outperform the laypersons in their ability to assess the relevant issues. In real life, scientific credentials, the exhibited confidence (conveyed in body language) and the ability to justify one’s point of view in a debate affect perceived expertise. However, in the era of social media some of these factors take a secondary role, since there is (usually) no physical presence and mostly no direct debate. Other metrics, such as the numbers of followers and likes, and the popular appeal of the message itself play a role. In this study we abstract from issues like acquired authority as

expressed in social media metrics, choosing to concentrate on another key factor that plays a role in digital communication: the public appeal of the suggested view regarding the ‘important topic’. In other words, when everyone has a voice, whose voice do people listen to and who do they elevate to the status of an ‘authority’?

We conduct a series of experiments to answer this question. Our findings confirm the hypothesis that a charlatan espousing popular beliefs can lead laypeople to choose to follow her advice rather than the advice of a genuine expert. This is true even in the face of increasing negative evidence regarding the accuracy of the charlatan. In particular, we introduce a questionnaire on economic policy issues and validate the correct answers. To do this, we use a sample of academic economists and select only questions that exhibit a high level of agreement on the correct answer. We also conduct a pilot study among laypeople in order to gauge the most popular answer to each question, i.e. the answer given by most people.

We then employ this questionnaire in a set of three-stage experiments with UK participants. In Stage 1, the participants answer each question on their own, receiving points for each correct answer. After answering each question, they see the suggested answers of two computerised ‘advisers’ on the same questions. One adviser is the ‘Expert’, who is designed to give the answer deemed correct by the academic consensus with probability 85% and a random answer otherwise. The other adviser is the ‘Charlatan’ who always proposes the most popular answer from the pilot study with layperson participants. In Stage 2, each participant sees a summary of the recommendations by both advisers from Stage 1 and they are asked to pick one adviser, who answers all of the questions on behalf of the participant. In Stage 3, participants go through the questionnaire one last time, with their selected adviser answering the questions for them. This time the participants are given immediate feedback on whether their adviser was correct and they can switch between advisers if they wish to do so.

Our design allows to examine the fundamental questions that we raised above. Note that we deliberately abstract from all other aspects affecting perceived expertise and assume that the only thing people know is the content of the advice. This is done in order to isolate the appeal of the message itself from the way it is conveyed. We demonstrate that, in economic matters, people have a strong tendency to follow the adviser who suggests similar answers to the people’s own priors (an insight consistent with [Gentzkow and Shapiro, 2006](#)). This leads them to incorrect choices at the cost of a significant foregone payoff, a loss of approximately 40% of the maximum possible earnings in our experiments.

The inability of the participants to choose the Expert does not result mechanically from their lack of knowledge. The majority of our observations come from an experimental environment where participants are fully informed of how the two advisers choose their answers. Participants knew that the Charlatan chooses answers that are popular, but these answers are to a large extent wrong. Sophisticated individuals, even if they do not know much about economics, should realise that the Charlatan’s answers most likely coincide with their own, and hence would pick the right adviser after seeing a large enough number of questions, even without feedback. Yet, the vast majority of participants do not. As [Kahneman \(2011\)](#) argued “Our comforting conviction that the world makes sense rests on a secure foundation: our almost unlimited ability to ignore our ignorance.” In our experiments, it seems that participants prefer to ignore their ignorance.

Even receiving feedback on an adviser’s past performance does not correct the ‘ignoring ignorance’ bias fully; updating is substantially slower than what a Bayesian model would suggest. While it is known that humans are not good in Bayesian updating ([Tversky and Kahneman, 1980](#)), our results are obtained in an environment that exhibits novel characteristics that expand the scope of the findings. First, we do not use an abstract numerical task, but policy related questions within a natural context. Second, feedback is very strong since correct answers are fully revealed. Third, in our questionnaire participants should be able to infer the appropriate answer using economic logic instead of a complicated mathematical formula. Our results are driven by a failure to disregard previously held beliefs and rationally assess the advisers on their merit given the feedback, not a failure to do maths.

The results are related to [Ronayne and Sgroi \(2018\)](#) and [Schotter \(2003\)](#) who examine how individuals respond to advice. However, we are interested in a different question, as we employ a design where advisers effectively compete for attention. [Chakraborty et al. \(2020\)](#) study theoretically the role of experts in electoral competition and contrast it to a populist alternative. Unlike their setting, the interests of the expert and the participants are aligned and there is no potential benefit from picking the charlatan. The literature on fake news is also relevant for our results. Previous studies show that fake news spread faster ([Vosoughi et al., 2018](#)) and people share these news even though they can tell they are likely not accurate ([Pennycook et al., 2019, 2020](#)). The authors propose that fake news are shared more because they are more interesting or surprising. This aspect is, arguably, absent in our case, suggesting another mechanism is at play.

In another strand of the literature, [Jerrim et al. \(2019\)](#) use the results from the PISA study to identify “bullshitters”, individuals who claim knowledge or expertise in an area where they have

little to none. This study finds that bullshitters have increased overconfidence, persevere at tasks, and are popular amongst their peers. While the experts and charlatans in our experiments are computerised, extensions with human advisers could yield insights as to how exactly the “bullshitters” persevere. In this sense our results also relate to [Pennycook et al. \(2015\)](#) who study receptivity to pseudo-profound bullshit.

Our results also inform the psychology literature. The latter teaches us that citizens who lack the capacity to think deeply about a topic are likely to believe theories and analyses that sound intuitive to their ears. In fact, the well-known Cognitive Reflection Test ([Frederick, 2005](#)) purports to distinguish between people who solve problems using the automatic and effortless ‘system 1’ vs. the slower and more costly and analytical ‘system 2’ ([Kahneman, 2011](#)). If laypeople address matters of expertise on the basis of system 1, populists who pander to them by giving intuitive, but wrong, answers are likely to be more successful. According to [Kahneman \(2011\)](#), when faced with an unknown domain, people are unable to ask the question “What type of information would I need in order to answer this question?” Instead, they use a System-1 heuristic: “Can I make a convincing-sounding story about this?” Our findings seem to validate the predictions of the above literature.

The paper is organised as follows. Section 2 elaborates further on the importance of economic expertise and the divergence of the views of the laypeople and economists. Section 3 briefly introduces the experimental design and the benchmark theoretical predictions, while Section 4 describes the conducted experiments in detail. Section 5 presents the empirical results and Section 6 concludes.

## 2 The Importance of Economic Expertise and Economics as a Science

In America, public distrust towards experts is considered a strong cultural characteristic, while in the UK conservative politician Michael Gove popularised the famous slogan that the public ‘has had enough of experts’. We argue that distrust towards economists both exemplifies the general problem of mistrust of experts and institutions, and at the same time it has unique features. It is important to first emphasize that the manner in which the public perceives economics problems can have enormous consequences for society. In general, economic thinking is central for key institutions and policies of the Western world, such as social insurance, trade policies, and independent central

banks. Their influence extends to issues such as policies to deal with the coronavirus pandemic, and as of April 2020, most economic experts agree that keeping quarantine measures in place makes economic sense (as per the IGM panel). However, in modern democracies it is the view of the general public that matters more, because without public support policies cannot be fostered and implemented.

There is a strong divergence between the views and ways of thinking of experts and the general public when it comes to economic affairs, and the difference is systematic (Leiser and Krill, 2017). In fact, this phenomenon is so old that more than a century ago Newcomb (1893) made a few arguments that could easily be made today: that the public seems to support detrimental policies for their self-interest, such as restrictions on immigration and trade and opposition to labour-saving technology. There are good reasons why this divide exists, and both economists and psychologists have conducted research in recent years in order to shed light to this phenomenon.

First of all, economics exhibits idiosyncratic characteristics as a science (Arthur, 2000). The economic way of thinking is often counter-intuitive and takes years of training (Colander, 2005). Despite this fact (and unlike many natural sciences) laypersons are prone to fostering personal simplistic theories of the economy without much technical training (Leiser et al., 2010; Dixon et al., 2014; Leiser and Krill, 2017). Moreover, economic issues might elicit emotional, rather than analytic responses (e.g. views on immigration), and ideology often exacerbates disagreements. The complexity of economic problems (especially in macroeconomics) raises the issue of whether knowledge of economic issues has the same meaning as expert knowledge in the physical or even biomedical sciences (Javdani and Chang, 2019). However, Caplan (2002) shows that economic expertise is meaningful, in the sense that training - rather than other factors (e.g. demographics) - accounts for differing beliefs between experts and laypeople.

Several cognitive biases have been proposed and empirically verified as responsible for the gap of economic understanding between laypeople and professional economists. Caplan (2011) categorizes four types of systematic biases in laypeople's economic reasoning: anti-market bias, anti-foreign bias, 'make-work' bias (the idea that work itself is valuable and not the fruit of it) and 'pessimistic bias' (the systematic tendency to see the economic future as grimmer than the past). The GBG-heuristic ('Good-Begets-Good', Leiser and Aroch, 2009) explains why laypeople believe in false causal links between pairs of economic variables when they are both perceived as 'good' (tax cuts and employment: Bartels, 2005) or 'bad' (inflation and unemployment: Dräger et al., 2016; inflation and interest rate increases: Andre et al., 2019). Metaphors are also powerful tools that laypeople



often rely on in order to transform the incomprehensible into accessible (Oberlechner et al., 2004), but over-reliance on them may lead to significant judgement errors (Krugman, 2010). Finally, one cannot ignore the ability of the human mind to perceive teleological forces behind random events or behind the outcomes of complex dynamic systems (Kelemen and Rosset, 2009; Kelemen et al., 2013). Although conspiracy-based explanations of economic events have not been studied as extensively as the other biases, there is evidence that people have a systematic tendency to ignore general equilibrium effects (Dal Bó et al., 2018), and that they are more likely to attribute economic crises to individual errors than institutional failings (Gangl et al., 2012; Aperia and Sappa, 2014).

It is very important for economists to further study the determinants of this discrepancy and how to reduce it. In the absence of a well-functioning communication between experts and laypeople, populists are likely to fill in the vacuum by pandering to the public’s prior beliefs. This, in turn, fosters a vicious cycle of reduced credibility for the economists and an increasing distrust from the public. Secondly, economists may run the risk of thinking in paternalistic terms. Zingales (2020) forcefully argues that economists might fall into the temptation of not making the implications of policies clear to citizens, from fear that if they fully understand those policies they might oppose it.<sup>1</sup>

We contribute to this literature by examining the robustness of the discrepancy in views between laypeople and economists in an incentivised environment with direct feedback. We find that the discrepancy is pretty sizeable and persistent. In addition, unlike previous studies, we examine experimentally how it affects people’s choice of representatives: it can be used to manipulate public support by a populist charlatan.

### 3 Overview of Experimental Design and Predictions

We have a simple three-stage laboratory experiment and each session has the following structure. Participants are seated in computer terminals and do not interact with each other. They are exposed to an economics questionnaire with eleven multiple-choice questions, as presented in Appendix A.<sup>2</sup> First, in Stage 1 they are asked to answer these eleven questions without any feedback. After answering each question, they observe the proposed answer by two ‘advisers’. They are informed that one of them is a high-accuracy adviser (whom we call the ‘Expert’ in the paper but not during

---

<sup>1</sup>For example, the true views of the American public are likely to be opposed to social insurance, but this domain is very difficult to fully understand, and conscious obfuscation likely prevents direct opposition.

<sup>2</sup>As we shall see, in essence participants answer these eleven questions three times.



the experiment), who answers on average 85 percent of questions correctly, and that the other one, the low-accuracy adviser (whom we call ‘the Charlatan’ in the paper but not in the experiment), only answers 55 percent of the answers correctly.

In Stages 2 and 3, participants have monetary incentives to detect who the high accuracy adviser is. In particular, in Stage 2 they make a choice of adviser once and for all, and the answers of their chosen adviser will apply to all the questions. Finally, in Stage 3, they have the opportunity (in every round) to update their choice of adviser to apply to the next round, after receiving feedback on the correct answer in the current round. For each correct answer in stages 1, 2 and 3 participants receive 4, 16 and 8 points, respectively. Each point counts for 0.05 British Pounds (approximately 0.06 US dollars).

In order to keep the incentives appropriate in Stage 3 of the experiments,<sup>3</sup> there is an independent draw for each question that determines whether the high-accuracy adviser gives the correct answer or not. The probability of a correct answer is fixed within each experiment to 85%. When it comes to the low-accuracy adviser, the answers are fixed ex ante. In particular, for each question the least accurate adviser gives the modal answer chosen in a pilot with laypersons run in Greece. These modal answers turned out to be correct for six out of the eleven questions, and the average percentage of subjects that chose the modal answer in each question was 63%. This is a critical design aspect: the less accurate adviser chooses by design the question most likely to be chosen by a person drawn randomly from a population of laypersons.

For each question, participants had to pick one out of four answers.<sup>4</sup> One of the answers is correct and the other three are wrong. Since these are not mere mathematical exercises, in order to determine the correct answer, we looked for expert consensus using a set of academic economists from university departments in Europe. We selected top departments of economics, and randomly chose members who were invited to participate. If they agreed, they simply answered a subset of our questionnaire without any feedback. In order to be validated, a correct answer required 70 percent consensus among these academic economists.

We have two main experiments. In the ‘low-information experiment’, our participants are informed about the structure of the experiment and the overall accuracy of the two advisers. However, we do not explain to them the process that we used to choose the answers of the low-

---

<sup>3</sup>Since in Stage 3 there is sequential feedback question-by-question on the correct answer, if the fraction of correct answers of the Expert was fixed to 85%, then after a long sequence of correct answers by the Expert, subjects would have an incentive to avoid using the Expert as their representative, since a wrong answer is due.

<sup>4</sup>Only one question diverged from this pattern, having five options instead of four.

accuracy adviser. This benchmark treatment corresponds to the natural setting in which the public simply observes prescriptions about the economy (in the media) without knowing how they were formed. In our ‘high-information experiment’ participants have all relevant information about the process with which the low-accuracy adviser chose their answers, including the average popularity of modal answers in the Greek pilot (63%). The key question is whether this comprehensive information suffices to make participants realise that they are more likely to have a common answer with the low-accuracy adviser, rather than with the high-accuracy one.

We are principally interested in participants’ choices in stages 2 and 3. Considering Stage 2 behaviour, we experimentally examine how the advisers’ suggestions alone drive laypersons’ choice of representatives/influencers. Thus, we examine how the opinions of laypeople about economic matters affect whom they choose to influence or represent them. Stage 3 examines the robustness of such potentially biased choice to feedback. How much feedback on the correct answer is needed for participants to select the high-accuracy adviser (if possible at all)?

To address such questions, we build an explicit framework of how participants may view the experimental interaction and choose to behave in our experiments. This framework guided our experimental design, especially with regards to the information that participants would need in the high-information experiment. We wish to predict behaviour in stages 2 and 3 using simple models of Bayesian learning. We construct and test two such models, which correspond to different levels of information from the participants’ side. The first one examines the case where participants do not know anything about the method through which the Charlatan selects her answers. We refer to this model as the ‘low information’ model or simply the ‘low info’ model. The second one analyses the case where participants are given full information on how the Charlatan selects her answers. We refer to it as the ‘high information’ model or simply the ‘high info’ model. The two models are presented below.

### **3.1 The Low Information Model**

First, we lay down the set of assumptions. Let us first assume a questionnaire with four possible answers for each question, and fix answer 1 as the correct one, without loss of generality. Consider three agents who provide answers to the questionnaire: the participant (p), the low-accuracy adviser (l), and the high-accuracy adviser (h). Let us make the simplifying assumption that the probability

of each individual answering correctly each question is fixed and independent across questions.<sup>5</sup> We assume that the participant has probability  $p_p \geq 1/4$  of giving a correct answer in every round, the low-accuracy adviser has analogous probability  $p_l > p_p$ , and the high-accuracy adviser has respective probability  $p_h > p_l$ . We also assume that wrong answers are equally likely (this is true for all three agents), and that the participant merely wishes to maximise experimental earnings by choosing the high-accuracy adviser as much as possible.

If a participant has this simple model of the world, which adviser should she choose at stages 2 and 3 of the experiment? We are interested in answering this question through the lens of the above model, keeping in mind that the participant does not know who the high-accuracy adviser is. What she knows is how many common answers with each adviser she has (in Stage 2) and which questions each adviser answered correctly (in Stage 3). Thus, from the participant’s perspective, there are two states of the world and she needs to form posteriors for each one of them. The two relevant states of the world are as follows.

- **S<sub>1</sub>**: Adviser 1 (A1) is the high-accuracy adviser, adviser 2 (A2) is the low-accuracy adviser.
- **S<sub>2</sub>**: Adviser 1 (A1) is the low-accuracy adviser, adviser 2 (A2) is the high-accuracy adviser.

We are ultimately interested in the participant’s posteriors  $P(S_1/h)$ , where  $h$  is her information set, namely the set of all of her answers and all the advisers’ suggestions for Stage 2, along with feedback on the correct answer for a subset of questions in Stage 3. But for now, let us focus on Stage 2 alone.

### 3.1.1 Bayesian Updating in Stage 2

Consider the case of a single question for the moment. Denote the participant’s answer as  $\alpha_p$ , the charlatan’s as  $\alpha_l$  and the expert’s as  $\alpha_h$ . Given that there are 4 possible answers, this gives  $4^3 = 64$  different possible answer configurations from the three agents. Formally, these configurations are denoted as  $\boldsymbol{\alpha} = \{\alpha_p, \alpha_l, \alpha_h\}$ . From the perspective of the participant, who does not know the correct answer, these 64 configurations form 5 distinct *events* on the basis of the coincidence of the answers across the agents. These events are as follows.

- *Event 1 (E1)*: All three agents give the same answer to the question.

---

<sup>5</sup>In reality, the probability of a correct answer may differ across questions. Moreover, note that the low-accuracy adviser’s answers are not random, but the participant does not know this. Finally, both correct and wrong answers could be correlated, since the answers to a block of questions may depend on which theory about the world is correct.

- *Event 2 (E2)*: The participant gives a common answer with A1 and a non-common with A2.
- *Event 3 (E3)*: The participant gives a common answer with A2 and a non-common with A1.
- *Event 4 (E4)*: The two advisers give the same answer, but the participant's answer is different.
- *Event 5 (E5)*: All three agents select different answers.

For each of the above events the participant calculates the probability of this event conditional on the state being  $S_1$  or  $S_2$ . The calculations of these probabilities for all events and states are provided in Appendix B, where we show that the formula for Bayesian updating of the probability of event  $S_1$  is:

$$P(S_1|h) = \frac{\pi_0}{\pi_0 + (1 - \pi_0) \cdot (OR)^K} \quad (1)$$

where

$$OR = \frac{p_p \cdot p_l \cdot (1 - p_h) + \frac{(1-p_p)(1-p_l)p_h}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}}{p_p \cdot p_h \cdot (1 - p_l) + \frac{(1-p_p)(1-p_h)p_l}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}} \quad (2)$$

In addition,  $\pi_0$  is the prior and  $K \equiv$  [number of times E2 has been observed]-[number of times E3 has been observed]. Selecting the natural prior  $\pi_0 = 1/2$ , we can easily see that if  $K > 0$  the posterior is greater than 1/2. This implies that a participant who believes the low info model would use the simple heuristic in Stage 2 of picking the adviser with whom she shares the most common answers. The key question is whether participants behave according to the predictions of this model or according to the high info model, which will be presented shortly. As we shall see, under the latter model, participants may choose in Stage 2 the adviser with the least common answers. However, note that in Experiment 1 the information that subjects have corresponds naturally to the low info model, whereas in Experiment 2 it corresponds to the high info model.

### 3.1.2 Bayesian Updating in Stage 3

Now, let us consider the implications of the model for Bayesian updating in Stage 3. Recall that in this stage the participant receives feedback for each question and decides whether to switch to the other adviser for answers or not. The feedback consists of the correct answer and the answer choice of the currently selected adviser. Thus, if the participant remembers the suggestions of both advisers from Stage 1, she can infer who answered it correctly and who did not.

Therefore, the information set of the participant gets updated as the stage progresses. Specifically, at an intermediate point of Stage 3, the participant evaluates her information on the basis of the correctness of the two advisers' responses for the questions for which she has received feedback. On the other hand, for the remainder questions, for which no feedback is yet available in Stage 3, she continues to evaluate her information on the sole basis of the number of common answers with the two advisers. We will formally derive the implications of this distinction soon.

With regards to the questions with Stage-3 feedback, there are four distinct events to consider for each question. Notation-wise, we distinguish these events from the events of Stage 2 by using  $\hat{E}$  instead of  $E$ :

- *Event 1* ( $\hat{E}1$ ): Only A1 gave the correct answer to the question.
- *Event 2* ( $\hat{E}2$ ): Only A2 gave the correct answer to the question.
- *Event 3* ( $\hat{E}3$ ): They were both correct.
- *Event 4* ( $\hat{E}4$ ): They were both wrong.

Let us consider the case of a single question with feedback. For each one of the above events, the conditional probabilities  $P(\hat{E}|S_1)$  and  $P(\hat{E}|S_2)$  can be explicitly derived. For example,  $P(\hat{E}1|S_1) = p_h(1 - p_l)$ . The full table of these probabilities is provided in Appendix B, where we show that Bayesian updating for any part of Stage 3, after observing a history  $h$  that contains both Stage-2 and Stage-3 events, can be described by the formula:

$$P(S_1/h) = \frac{\pi_0}{\pi_0 + (1 - \pi_0) \cdot (OR)^{K_{-\hat{N}}} \cdot \hat{OR}^{K_{\hat{N}}}} \quad \text{with} \quad \hat{OR} = \frac{p_l \cdot (1 - p_h)}{p_h \cdot (1 - p_l)} \quad (3)$$

where  $K_{\hat{N}}$  is the number of times that A1 gave a correct answer and A2 gave a wrong answer in Stage 3, minus the number of times that A2 gave a correct answer and A1 gave a wrong answer in Stage 3.  $K_{-\hat{N}}$  denotes the number of times that A1 alone had a common answer with the participant minus the number of times that A2 alone had a common answer with the participant in the questions that have not yet been included in Stage 3.

### 3.2 The High Information Model

Recall that the ‘most popular answer’ in the Greek pilot questionnaire was on average selected by about 63% of these laypersons, and this answer was used as the low-accuracy adviser’s recommendation in our experiments. In Experiment 1, subjects know nothing about this fact, but only know

the accuracy of the two advisers. In Experiment 2 participants have full information about the manner in which the two advisers choose their answers. Accordingly, the low info model applied to Experiment 2 could be considered as a prediction of the behaviour of a bounded-rational participant in Experiment 2, who fails to realise that the correlation of their answers with the Charlatan’s answers is higher than the correlation with the Expert’s answers. We will now provide a model which describes the reasoning of a more sophisticated participant, one who takes into account the information on how the low-accuracy expert’s recommendations were selected and adjusts her posteriors accordingly. Since the participant knows how the Charlatan selects answers, we term this model the high information model.

Note that behaviour according to the high info model is opposite to a powerful psychological tendency to maintain a positive view of oneself (Akerlof and Dickens, 1982). In particular, believing that the Expert is the one providing the most similar answers to them allows participants to maintain the view of themselves as being smart and making correct decisions. In particular, a participant with motivated reasoning will want to believe that they are more similar to the Expert than to the Charlatan. We will comment in later chapters on how these motivated reasoning models can illuminate aspects of our participants’ behaviour.

Once more, assume four possible answers for each question, and that answer 1 is the correct one. Let  $\pi > 1/4$  be the fixed probability, same for each question, that the low-accuracy adviser gives the same answer as the participant. This captures the manner in which the low-accuracy adviser’s answers were selected. Here we implicitly treat our participants as members of the same population of laypersons as the participants to the Greek pilot. On the other hand, as in Model 1, the high-accuracy adviser and the participant have probability  $p_h$  and  $p_p$ , respectively, of selecting a correct answer, with  $p_h > p_p \geq 1/4$ .

### 3.2.1 Bayesian Updating in Stage 2

Calculating the posteriors in Stage 2 is very similar to the low info model. As in section 3.1, there are two states,  $S_1$  and  $S_2$ , and the same five events described in page 10. The configurations of possible answers  $\alpha = \{\alpha_p, \alpha_l, \alpha_h\}$  fall into these five events and the probability of each event can be calculated once the correlation between the answers of  $p$  and  $l$  is taken into consideration.

For example, the probability of event E1 in state  $S_1$  is  $P(E1|S_1) = p_p \cdot p_h \cdot \pi + (1 - p_p) \cdot (1 - p_h) \cdot \pi / 3$ . This is derived as follows. E1 is the event where the three agents give all the same answer, namely the possible configurations in this event are  $\{1, 1, 1\}$ ,  $\{2, 2, 2\}$ ,  $\{3, 3, 3\}$ , and  $\{4, 4, 4\}$ . Configuration

$\{1, 1, 1\}$  happens when all three give (the correct) answer 1, which happens with probability  $p_p \cdot p_h \cdot \pi$ . Note that  $p_p \cdot \pi$  is the probability that  $p$  answers correctly and  $l$  gives the same answer. Each one of the other configurations corresponds to a wrong answer and this has probability  $3 \cdot \frac{(1-p_p) \cdot \pi}{3} \cdot \frac{(1-p_h)}{3}$ . Summing up the two numbers gives the expression above for  $P(E1|S_1)$ . The probabilities of all other events in all states are computed similarly and they are provided in Appendix B.

With the above calculations at hand, Bayesian updating in Stage 2 is similar as in the low info model. Given any history  $h$  of events,  $P(S_1|h)$  is computed by equation (4) below with the corresponding odds ratio given in (5):

$$P(S_1|h) = \frac{\pi_0}{\pi_0 + (1 - \pi_0) \cdot (OR_{II})^K} \quad (4)$$

$$OR_{II} \equiv \frac{p_p \cdot \pi \cdot (1 - p_h) + (1 - p_p) \cdot \pi \cdot p_h + 2/3[(1 - p_p) \cdot \pi \cdot (1 - p_h)]}{p_p \cdot (1 - \pi) \cdot p_h + [(1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h)]/3} \quad (5)$$

Note that the low info model predicts that a participant with prior  $1/2$  should *always* pick the adviser with the most answers common with her. However, the high info model is richer, in that it allows  $p$  to choose the adviser with the least common answers. When could this happen? Let us assume that  $K > 0$ , so that Event 2 has been observed more times than Event 3. Accordingly,  $A1$  is the adviser with the most common answers with the participant. In order for the participant to prefer  $A2$  (the adviser with the least common answers) it has to be the case that the posterior for  $S_1$  is less than half. Assuming the natural prior  $\pi_0 = 1/2$ , we have that the relevant inequality is  $P(S_1/h) = \frac{1}{1+(OR_{II})^K} < 1/2$ . This holds whenever  $OR_{II} > 1$ . By substituting in the value of  $OR_{II}$ , we can find the range of values of  $\pi$  for which this inequality holds. Doing so gives us the inequality:

$$\pi > p_p p_h + (1 - p_p)(1 - p_h)/3 \quad (6)$$

For instance, let us set  $p_h = 0.85$  (as in most of our high-information experiments) and  $\pi = 0.63$ . This value of  $\pi$  is the average popularity of the modal answer in the Greek pilot (the low-accuracy recommender's answer) which subjects know in the high-information treatments. Let us also assume that  $p_p = 0.47$ , meaning that the participant believes that he is as accurate as the average participant in our experiments. In this case, the left-hand side of inequality (6) is equal to 0.63 and the right-hand side is 0.426. As a result, we see that for parameters corresponding



with our experimental conditions, this model predicts that rational and fully-informed participants should choose the adviser with the least amount of common answers. The condition in (6) is very intuitive, since  $\pi$  is the expected fraction of answers common with the low-accuracy adviser and  $p_p p_h + (1 - p_p)(1 - p_h)/3$  is the expected share of answers that are common with the high-accuracy adviser. Whenever the former is greater, the participant rationally infers that the adviser with the most common answers with herself is most likely the low-accuracy adviser, hence choosing the other adviser.

Notice that in our experiments the only parameter that differs across subjects is  $p_p$ , the perceived accuracy of the participant. Subjects with different confidence in their economic knowledge could choose differently. Define  $T \equiv p_p p_h + (1 - p_p)(1 - p_h)/3$ . For values of  $p_p$  between 0.05 to 0.6,  $T$  is less than 0.63 (meaning that the participant should choose the adviser with the least common answers), while for values of  $p_p$  greater than 0.6,  $T$  is greater than 0.63, so the subject should choose the adviser with the most common answers. In other words, overconfident subjects could still choose the low-accuracy adviser, despite using the correct model of the world to analyse the problem. This is interesting, because the model teaches us that overconfident people are easier targets for populists. We delve more deeply into this issue in our data analysis.

### 3.2.2 Bayesian Updating in Stage 3

Stage 3 Bayesian updating in the high info model is slightly different than in the low information model. This is because the probability of the low-accuracy adviser getting a correct answer in this model depends on whether the participant got the correct answer or not, and so there are now eight possible cases with regards to the feedback that the participant can receive in a given period of Stage 3. Before we present these cases, notice that the probability of the low-accuracy adviser being right, conditional on the participant being wrong, is  $(1 - \pi)/3$ . This is because  $l$  picks a different answer from  $p$  with probability  $1 - \pi$  and, conditional on  $p$  being wrong,  $l$  picks the correct answer from the remaining three with probability  $1/3$ . Also, the probability of  $l$  being wrong, conditional on  $p$  being wrong, is  $1 - (1 - \pi)/3 = (2 + \pi)/3$ . Now, the eight possible events of Stage 3 in the high info model are as follows.

- *Event 1* ( $\tilde{E}1$ ): Everyone (participant,  $A1$  and  $A2$ ) give the correct answer.
- *Event 2* ( $\tilde{E}2$ ): Only the participant is correct.
- *Event 3* ( $\tilde{E}3$ ): The participant and  $A1$  are correct,  $A2$  is not.

- *Event 4* ( $\tilde{E}4$ ): Only  $A1$  is correct.
- *Event 5* ( $\tilde{E}5$ ): The participant and  $A2$  are correct,  $A1$  is not.
- *Event 6* ( $\tilde{E}6$ ):  $A1$  and  $A2$  are correct, the participant is not.
- *Event 7* ( $\tilde{E}7$ ): Only  $A2$  is correct.
- *Event 8* ( $\tilde{E}8$ ): They are all wrong.

The probability  $P(\tilde{E}|S)$  of each event in each state is provided in Appendix B. Moreover, events  $\tilde{E}1$ ,  $\tilde{E}2$ ,  $\tilde{E}6$ , and  $\tilde{E}8$  give the same probability for both states, and so their odds ratio is equal to one. Thus, the only informative events are  $\tilde{E}3$ ,  $\tilde{E}4$ ,  $\tilde{E}5$ , and  $\tilde{E}7$ , which are paired in terms of odds ratios.  $\tilde{E}5$  has the inverse odds ratio of  $\tilde{E}3$  and  $\tilde{E}7$  has the inverse odds ratio of  $\tilde{E}4$ . The odds ratios for  $\tilde{E}3$  and  $\tilde{E}4$  are derived in Appendix B and are given below.

$$OR_{III} \equiv \frac{P(\tilde{E}3|S_2)}{P(\tilde{E}3|S_1)} = \frac{\pi(1-p_h)}{p_h(1-\pi)} \quad (7)$$

$$OR_{IV} \equiv \frac{P(\tilde{E}4|S_2)}{P(\tilde{E}4|S_1)} = \frac{(1-\pi)(1-p_h)}{p_h(2+\pi)} \quad (8)$$

Intuitively,  $OR_{III}$  applies to the case where only one adviser and the participant are correct, while  $OR_{IV}$  applies to the case where only one adviser is correct. Let us now define  $k^1 \equiv$  [number of times  $\tilde{E}3$  has occurred in Stage 3] – [number of times  $\tilde{E}5$  has occurred in Stage 3], and  $k^2 \equiv$  [number of times  $\tilde{E}4$  has occurred in Stage 3] – [number of times  $\tilde{E}7$  has occurred in Stage 3]. Again,  $\hat{N}$  denotes the set of rounds for which feedback has been provided in Stage 3 and  $-\hat{N}$  denotes the set of rounds for which feedback has not been provided in Stage 3. Therefore,  $k_{\hat{N}}^i$  denotes the measure  $k^i$ ,  $i \in \{1, 2\}$  applied to the rounds with Stage-3 feedback, and  $K_{-\hat{N}}$  denotes the measure  $K$  (as defined for the low-info model) applied to the rounds without Stage-3 feedback. Overall, the final updating formula for Stage 3 of the high info model is:

$$P(S_1/h) = \frac{\pi_0}{\pi_0 + (1-\pi_0) \cdot (OR_2)^{K_{-\hat{N}}} \cdot (OR_{III})^{k_{\hat{N}}^1} \cdot (OR_{IV})^{k_{\hat{N}}^2}}. \quad (9)$$

## 4 The Experimental Sessions

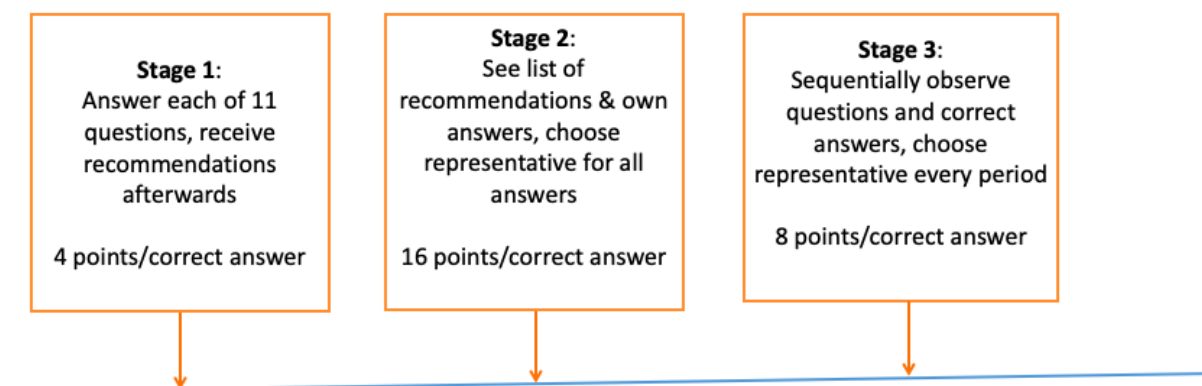
We run a series of experiments, with the key manipulation pertaining to the information provided. We also conducted some additional checks, such as varying the accuracy of the high-accuracy adviser from 70 to 85 percent, and examining the role of financial incentives. The timeline of all the experimental sessions is the same (Figure 1) and let us describe it in detail. In Stage 1 participants answered eleven questions sequentially. Each correct answer carried a prize of 4 points (there was no punishment for wrong answers at any stage). Participants were informed of the existence of two advisers labelled ‘Adviser X’ and ‘Adviser M’. In each round, after choosing their own answer, participants were informed of the answers that the two advisers suggested for the particular question (although they could not use this knowledge any more for this stage). After they answered all eleven questions Stage 1 ended.

In Stage 2, participants were first presented with an overview of the answers provided by themselves and the two advisers for each question. They then had to choose a representative among the two advisers, whose answers would be used for all questions and apply to the participant’s payoffs. Each correct answer at this stage carried a prize of 16 points. In other words, in Stage 2, the selected adviser answered the same eleven questions from Stage 1 and the participant earned points based on the performance of the adviser. Hence, participants were financially incentivised to pick the adviser whom they thought had the highest accuracy among the two.

Finally, in Stage 3 participants had the chance of using the advisers in order to answer the same set of questions one last time. However, in this stage the correct answers were revealed sequentially after each question. In particular, each participant was informed about the correct answer for the last question and the chosen answer of the adviser whose answer applied for that question. After this feedback on each question, participants chose the adviser who would answer the next question. Participants could switch advisers as many times as they wished. Each correct answer in this stage was worth 8 points. The selected adviser from Stage 2 was used to answer the first question of Stage 3, after which the participant could choose between her current or the other adviser.

After the main part of each experimental session, participants answered a questionnaire. We collected additional information on several dimensions, namely cognitive reasoning (CRT), strategic sophistication (using the ‘undercutting game’ from [Georganas et al., 2015](#)), psychological indicators (measures of confidence and ego-centricity), political attitudes (especially towards redistribution) and demographics. Our three-stage design maximises the information that can be elicited from the

given set of questions. First we get the the participants’ personal answers, then elicit the choice of representative without feedback, and finally we are able to examine the role of feedback.



**Figure 1:** *The timeline of the experiments.*

#### 4.1 Experiment 1

Our first experiment emulates a natural setting in a strong sense. In real life, the public observes the opinions and suggestions of influencers regarding important matters, without knowing the process underlying these views and suggestions. This means that if there are populists pandering to the public to gain its support, this is not overt. Accordingly, in our first experiment participants were only informed of the overall accuracies of the two advisers (they did not know which one of X or M was the high-accuracy adviser, of course). No information regarding the exact manner in which answers were chosen for the Charlatan was provided. We argue that this information environment captures realistically the chances that the Charlatan has of gaining public support in the current state of affairs in the public sphere. In this setting, we expect that our simple low information model makes reasonable predictions for the behaviour of participants, since it is natural that they consider that this model describes the actual interaction.

Let us recall that our Stage 1 has 11 experimental choices, Stage 2 has only one choice, and Stage 3 has 10 choices. Our Experiment 1 took place at the University of York (EXEC Lab) in June 2019. 69 participants took part in six sessions. The experiments were designed using z-tree (Fischbacher, 2007), and each session lasted about 40 minutes, with average payoffs equal to £11.5 or approximately \$14. In half of the sessions the accuracy of the Expert was 70% and in the other half this accuracy was 85%. Although our preferable specification was the one where the Expert had 85% accuracy (to maximise the treatment effect), we were sceptical about whether it would

lead to uniform behaviour in Stage 3. Therefore, we also used Expert accuracy equal to 70% in some sessions in order to make sure that participants do not choose exclusively the Expert in Stage 3. This concern did not turn out to be relevant, as we will show in the results part. The exact instructions for all experiments can be found in the online appendix.

## 4.2 Experiment 2

The second experiment took place at the University of Southampton (Southampton Social Sciences Experimental Laboratory) and the University of York (EXEC Lab) during the Fall 2019 and Winter of 2020. In this experiment, participants were provided with complete information regarding the manner in which the two advisers chose their answers to the questionnaire. In particular, we provided several examples to illustrate the concept of modal answers, and there was also a quiz to test for participants' understanding. In addition, we elicited their beliefs regarding how many answers they believed they answered correctly in Stage 1 and the probability they chose the high accuracy adviser in Stage 2. In all other aspects, this experiment was identical to Experiment 1 (except the incentives in the Southampton sessions, as explained below). This environment allowed us to test whether complete information on how advisers select answers affects the popularity of an influencer who gives similar answers to themselves. Apart from testing a rigorous Bayesian model, this experiment informs us about the strength of populism and fake news in a controlled environment. In particular, how can we de-bias receivers of news who have strong home-made theories about the economy? Does providing strong evidence about the existence of influencers/politicians who follow populist strategies change who the public chooses as a representative?

Moreover, in Experiment 2 we wanted to examine the role on financial incentives. In particular, [Caplan \(2011\)](#) summarises very strongly the point that people may have motivated reasoning regarding economic problems. Models of motivated reasoning argue that people attribute real value to holding certain personal beliefs ([Bénabou and Tirole, 2002](#)) and that as long as maintaining them is not too costly, they are unlikely to try and correct wrong beliefs. Varying the level of financial incentives in our Experiment 2 may provide us with insights on the importance of such motivations, and in particular regarding the trade-off between holding cherished beliefs and financial incentives. Moreover, there is a long standing debate about whether financial incentives make a difference in the behaviour of experimental participants, with many scholars arguing that participants are intrinsically motivated and therefore extrinsic financial incentives often make a small difference ([Camerer and Hogarth, 1999](#)). Our Southampton laboratory sessions of Experiment 2 did not

provide financial incentives to participants,<sup>6</sup> while in the York sessions incentives were identical as in Experiment 1 (20 points translated to £1).

Finally, in Experiment 2 we conducted an additional check, examining whether the order of the questions matters. This is of particular concern in Stage 3, where path dependence may matter significantly in how participants form their beliefs. Accordingly, in the York experiments we examined two different sequences (both randomly generated), and we shall also comment on the importance of the order when we discuss our results. Table 1 below summarizes the two experiments and the checks carried out in each one of them.

	<b>Experiment 1</b>	<b>Experiment 2</b>
<b>Treatment</b>	Low-Info	High-Info
<b>Expert Accuracy</b> (baseline)	85%	85%
<b>Location</b>	York (incentivised)	York (incentivised)  Southampton (non-incentivised)
<b>Additional Checks</b>	70% accuracy of Expert	No incentives (Ch. 1); Alternative Question Order (Ch. 2)
<b>Sessions</b>	6 (3 for add. check)	10 (4 for Ch. 1; 2 for Ch. 2)
<b>No of Participants</b>	69 (31 for add. check)	203 (54 for Ch. 2) (70 for Ch. 1)

**Table 1:** *Characteristics of different sessions in the two experiments.*

<sup>6</sup>In the Southampton experiments, participants gathered points per correct answer as in all other experiments, but these points did not translate into cash.

## 5 Results

### 5.1 Descriptives

Let us start with some simple descriptive statistics. On aggregate, over both experiments, we recruited 272 participants across both institutions (University of York and University of Southampton). The participants were registered students, mainly undergraduate, although there was substantial participation from postgraduate students and some life-long learning students. Table 2 below gives a precise break-down of the participants from the two experiments across gender, age and study area.

	Experiment 1	Experiment 2
<b>By Subject Area</b>	<b>No of participants</b>	
Economics	12	89
Finance and Accounting	7	2
Other Business	3	6
Other Social Sciences and Humanities	20	64
Positive Sciences	27	42
<b>By Gender</b>		
Female	39	102
Male	27	98
Other	3	3
<b>By Age Group</b>		
18-21	28	139
22-29	32	41
30 and over	9	23

**Table 2:** *Demographics of participants across the two experiments.*

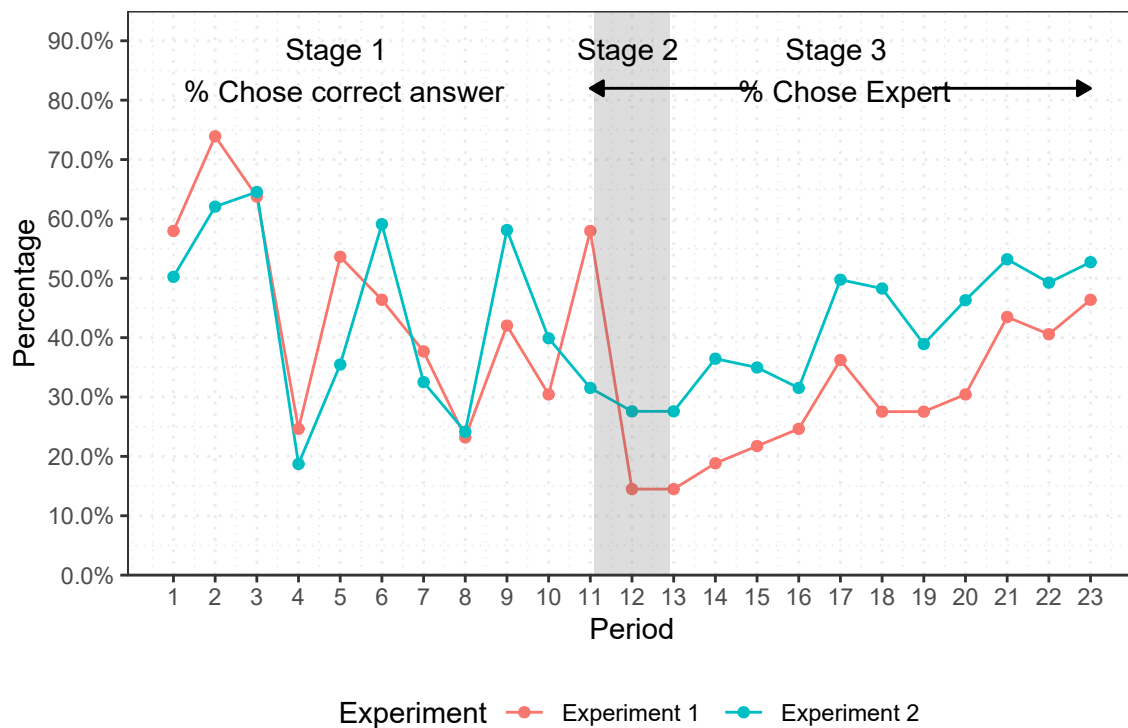
We observe that there is adequate variance regarding the course of study in the subject pool. Although more than a third of the participants (101 in total) came from economics background, a substantial fraction (84 in total) came from social sciences or humanities, and another sizeable fraction (69 in total) came from positive sciences. Hence, there is considerable representation of students from non-economic backgrounds. Similarly, we see that gender representation is balanced, while age representation is predominantly between 18 and 21 (167 out of 272), indicating a mainly undergraduate level of study. Overall, the subject pool is quite balanced and with some variation



for both experiments across all dimensions.

## 5.2 Choices

Now let us look at the frequency of correct answers across stages. This information is summarised in Figure 2 below. In the first stage, similar to the experiments we run in the pilot stage to set up the questionnaire, participants answer less than half of the questions correctly in the absence of feedback. This reflects the fact that the questions are not trivial, the answers are often counter-intuitive, and this seems to hold for various subject groups. On average, participants get 47% of the answers right in Experiment 1 and 43% in Experiment 2.



**Figure 2:** Subjects choosing the right answer in Stage 1 and the Expert in Stages 2 and 3.

In Stage 2, when the participants have to choose an adviser, the lack of knowledge is amplified. Recall that participants choose an adviser in that stage without feedback regarding the correctness of the answers. Given that most participants give answers that are similar to the Charlatan’s, uninformed choice goes overwhelmingly in favour of the Charlatan. Only 15% of the participants choose the Expert in Experiment 1 and about 28% choose her in Experiment 2 (Wilcoxon test for the difference yields no significance, with a p-value=0.07).

In Stage 3 we observe some learning, but it is very slow. Even after receiving feedback on 11

questions, close to half the participants choose the Charlatan in both experiments. In particular, in Experiment 1, only 46% of participants choose the Expert for the last question, and in Experiment 2 the percentage rises slightly to 52%.

### 5.3 Stage 3 Updating

In this section we will have a closer look at updating behaviour in Stage 3. Assuming a prior of 50%, our Bayesian models yield exact updating behaviour for each participant, given the answer revelations she observes. The other parameters that need to be calibrated are  $\pi$ , i.e. the correlation between the answers of the average participant and the Charlatan, and  $p_p$ , i.e. the accuracy of the participant in answering the questions. One can set  $p_p$  equal to the actual one or the believed one (according to our belief elicitation). It turns out it does not make a difference, as both specifications for  $p_p$  give very similar results, so in this section the results presented use the actual precision.

The parameter  $\pi$  is conceptually more complicated, as it is the probability that the participant gives an identical answer to the Charlatan. It is related to how close the subject believes she is to the median member of the population of laypersons. To start, we set this equal to the empirical fraction of modal choice using the actual pilot data, i.e. 0.63, which the participants explicitly know in Experiment 2. But note that it is not clear that the participants' perceived probability is actually equal to this; the Kahneman quote from the introduction actually suggests that few people understand the extent of their own lack of knowledge. Put differently, a participant may not wish to believe that she belongs to the population of laypersons.

Given these parameters, we calculate this optimal updating behaviour for each participant individually and can then infer which adviser they should be choosing. We present the average adviser choice (pooling together all sessions, including additional checks) given the two models, compared to the actual behaviour, in Figure ??.

Participants chose the Expert substantially less often than theory would suggest. Even after 11 occasions of feedback, a bit more than half the participants are choosing the Expert, while theory suggests this should be done in 70% to 75% of the cases. Excluding aberrant sessions, where the Charlatan actually got in total (weakly) more questions right than the Expert, the difference is even starker. Model 1 predicts that everyone should choose the Expert in the last two questions, while less than 60% actually do so. The particular comparisons pertaining to our additional checks can be found in Appendix C.

Overall, the most sophisticated of the two models, the high info model, gives a fit of 50.9%

across all subjects and periods, meaning it predicts about half the participants’ choices correctly, while the low info model does better, at 56.08%, but is still far from perfect. Clearly some elements of participants’ behaviour cannot be captured by the Bayesian models, so in the next section we build and calculate an alternative behavioural model to explain actual choices.

### 5.3.1 Alternative behavioural model

One possible behavioural model in the spirit of fictitious play would have the participants simply count the correct answers each adviser gives and choose the one with the most correct ones. On top of this count, the participant can give additional weight,  $\omega_{init}$ , to the initial choice they made in Stage 2. We allow this weight to be different for different participants, reflecting the strength of their beliefs (or stubbornness). We also allow for a common extra weight  $\omega_{cur}$  on the current observation the participant is making. That is, the Expert’s attractiveness for participant  $i$  after  $t$  questions is simply:

$$E_t^i = N_{correct}^{exp} + \omega_{init}^i \chi_{exp} + \omega_{cur} \chi_t^{exp}$$

where  $\chi_{exp}$  is an indicator function, equal to one if the participant chose the Expert in Stage 2, and  $\chi_t^{exp}$  is an indicator function equal to one if the Expert is giving the correct answer to the current question. The Charlatan’s attractiveness is equivalently:

$$C_t^i = N_{correct}^{ch} + \omega_{init}^i \chi_{ch} + \omega_{cur} \chi_t^{ch}$$

Participant  $i$ , after receiving feedback for question  $t$ , selects the Expert if  $E_t^i > C_t^i$ , the Charlatan if  $E_t^i < C_t^i$ , and randomises otherwise.

Notice that this model has a close correspondence to a simple reinforcement learning (RL) model (Erev and Roth, 1998). In such a model, players have a propensity  $p$  for each option or strategy they can play. In our setup the advisers are the options, and after each correct answer is revealed, the player would update the propensity of choosing an adviser by one unit if the adviser got it right, zero otherwise. Propensities are then converted to probabilities, usually setting the probability equal to the relative propensity. In our case, choice is simpler: participants just choose the adviser with the highest propensity.<sup>7</sup> The strength of the initial propensities is the free parameter in the basic reinforcement learning model, as in ours. We are also allowing for an extra weight on the

<sup>7</sup>This is similar to versions of RL that allow for cut-off parameters below which the probability of choosing an option becomes zero. We could allow for mixing as in the standard reinforcement learning model, but that does not improve the fit and would reduce the comparability to the Bayesian models.

current observation, which is related to versions of RL that allow for *recency* or *forgetting*, although it turns out not to be an important part of the model.

We estimate  $\omega_{init}^i$  and  $\omega_{cur}$  using an exhaustive grid search up to the second decimal. The estimation yields  $\omega_{cur} = 0.2$  and a relatively wide, declining distribution of  $\omega_{init}^i$  from 0 to 1. This means that about half the participants have low inertia, and switch relatively fast to the adviser with the best record, while the other half have high inertia, with a small percentage even sticking to their initial choice until the end, despite all evidence that the chosen adviser might be the Charlatan.

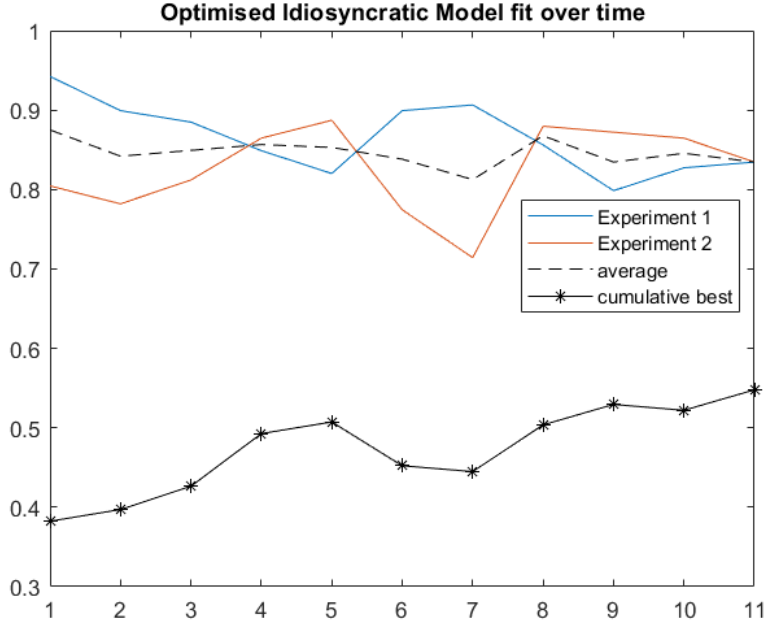
The fit of the model is 0.844, meaning it predicts correctly more than 8 out of 10 participants’ decisions across all periods. On the other hand, forcing  $\omega_{init}$  to be common for all subjects yields an estimated value of 0.6. The impact on the model’s fit is not high, as it falls to 0.7332. However there is a qualitative difference between the two specifications: with a common  $\omega_{init}$  the model fit falls with time, while the idiosyncratic  $\omega_{init}^i$  yields a consistent fit across all periods/questions, indicating individual differences in learning speed: some individuals switch fast while others slowly.

Figure 3 shows the model fit over time, for the two different experiments. The ‘*cumulative best*’ line represents the fit of the simplest possible counting model, where participants would only choose the adviser with the best record, without putting any weight on the initial choice. This line should be compared with the ‘average’ line, to give us an overall comparison between the simplest model and our behavioural model. The simplest counting model starts with a low fit and improves with time, but its fit is always substantially below that of the behavioural model. This shows that the increased complexity of our behavioural model comes with a substantial benefit.

## 5.4 Model Comparison

In this Section we compare the fit of the different models. Just looking at the percentage of choices explained by the models, in Figure 6, the simple behavioural model fits the data consistently best, explaining 84.7% in Experiment 1 and 84.9% in Experiment 2. The low info model is somewhat better than the high info model in both experiments (62.13% and 50.44% against 48.86% and 47.92% respectively).

A more sophisticated approach is to calculate a noisy model, where participants make mistakes with probability  $\epsilon$  and play what is predicted by each model with probability  $1 - \epsilon$ . We estimate the model parameters for this specification using a maximum likelihood procedure. The ranking of the models is similar to the one using just the fractions of explained choice as a measure of goodness




**Figure 3:** Model fit over time in the two Experiments.

of fit, but we can now compare the fit formally using likelihood ratio tests and the Bayesian Information Criterion. Fitting the behavioural model with idiosyncratic inertia weights brings a jump in the fit, but the Bayesian Information Criterion punishes the high number of parameters in this case. Table 3 reveals that the Behavioural model with homogeneous  $\omega$  is the best according to this criterion. Interestingly, a likelihood ratio test between the two nested behavioural models chooses the unrestricted model as best.

Model	Low Info	High Info	Beh. CW	Beh. IW
estimated error rate	0.45	0.52	0.28	0.15
LL	-2052	-2074	-1738	-1284
free parameters	1	1	1	273
BIC	4111.4	4155.1	3484.6	4745.4

**Table 3:** Estimation and goodness of fit of the various models, using all data. Low Info refers to the Bayesian low info model, High Info refers to the Bayesian high info model, Beh. CW refers to the behavioural model with parameter  $\omega$  common across subject, and Beh. IW refers to the behavioural model with idiosyncratic  $\omega$  for each subject.



**Figure 4:** Comparing the fit of the three models in the two Experiments.

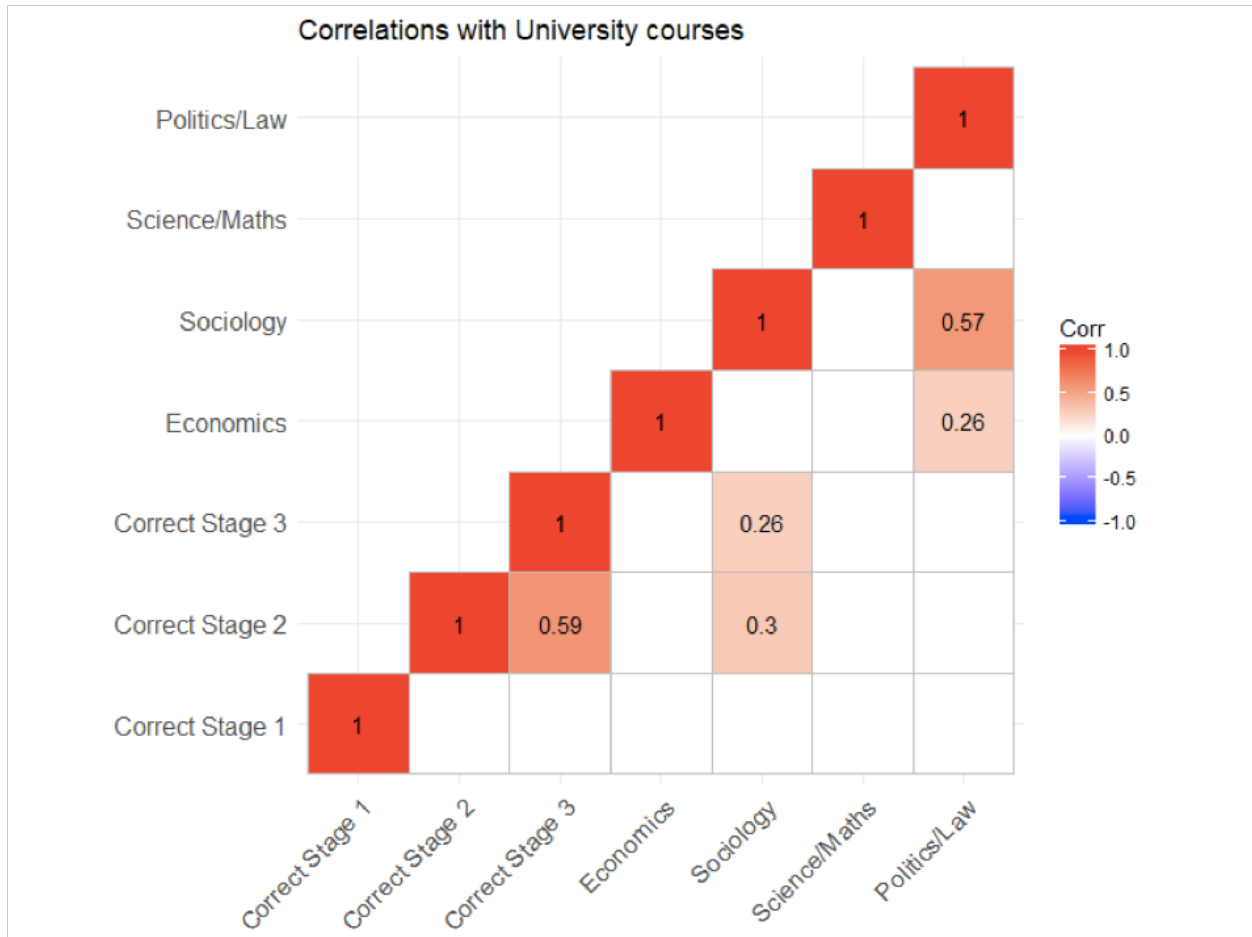
## 5.5 Who gets it?

We now present results on the importance of several participant characteristics, when looking for the right adviser. The next figure presents correlations between subject of study and correct choice in the different stages. Studying sociology seems to correlate positively ( $p\text{-value} < 0.05$ ) with correct choices in Stage 2 and 3. Economics, unexpectedly, has a lower correlation and hard sciences do not have any significant correlation with correct answers. Obviously, the numbers of these participants are low, so more robust results will have to await a wider study. Looking at confidence, the median participant believes they got 6 answers correct, and that is not different for participants choosing the Charlatan vis-à-vis those choosing the Expert in Stage 2.

As an overconfidence measure, we can look at the number of questions subjects answer correctly in Stage 1 minus the number they *think* they answered correctly.<sup>8</sup> Median overconfidence across Experiment 2 is 1, meaning that the median participant believed that they gave one more correct answer than they actually did. Please also note that both Charlatan-choosers and Expert-choosers believed that they had the same number of correct answers. Accordingly, overconfidence is significantly higher for Charlatan-choosers.

Confidence in an adviser is weakly linked to performance, at best. The Pearson correlation between the number of correct answers and belief in the adviser is 0.029 for participants choosing

<sup>8</sup>We do not have data on this from Experiment 1.



**Figure 5:** *Correlation of courses and correct answers.*

the Charlatan and 0.1 for those choosing the Expert (both are insignificant though). However, the number of answers that participants think they answered correctly is significantly related to their confidence in the adviser ( $\rho = 0.4$  and  $0.28$  for those choosing the Charlatan and the Expert respectively, p-values  $< 0.001$  and  $< 0.027$ ).

## 6 Conclusions

We ran a series of experiments with financial incentives to investigate whether the discrepancy of opinions between laypersons and economists is persistent, and its implications for the appeal of populist influencers on the general public. By developing a novel questionnaire on economic policy, comprised of 11 questions, we were able to address the above questions. As is frequently the case in economics, most of the questions had counter-intuitive answers, while the popular answer



is, following [Mencken \(2012\)](#), neat, plausible and *wrong*.<sup>9</sup> Our main general result is that our experimental participants overwhelmingly chose the Charlatan, who offered popular but frequently wrong advice, over the more accurate Expert. In some sessions less than 10% of subjects chose the Expert, which resulted in significant foregone experimental profits.

Our second general result is that providing information about how the Charlatan selects his answers hardly helps. Providing all necessary information about the Charlatan's *modus operandi*, short of actually telling the participants who the Charlatan actually is, does not significantly reduce the chances of choosing him over the simple experiment where subjects do not know anything about him. This is a strong indication of *confirmation bias*. Subjects followed a simple heuristic, believing naively that the person with the most common answers with them must be the Expert. Of course, an alternative interpretation is that participants derive direct utility from believing that the adviser with the most common answers with them is the high-accuracy adviser (*motivated reasoning*). This may also explain part of the results of Stage 3.

Perhaps more surprisingly, providing feedback on the correct answers and allowing for learning does not eliminate the choice of Charlatan. On the contrary, almost half of the participants stuck with the Charlatan in spite of the strong evidence against this choice. Even after 11 occasions where the right answer to the questions was revealed, only about 55% of participants chose the Expert. This number is way below the benchmarks set by the canonical model of learning, i.e. the Bayesian, which predicts between 69% to 75%, depending on the model details. In cases where the Expert was actually more frequently correct than the Charlatan, the theory predicts that the Expert would be chosen 90% to 100% of the time, while participants only chose her in about 60% of the cases. Note that in real life such perfect and direct feedback is very rare in public policy matters. For instance, even among the academic community, the effect of a minimum wage reform takes decades to be measured properly and in a way acceptable to most economists. In more complicated issues one needs to be even more pessimistic. For instance, how many years will it take for the economic effects of Brexit to be cleanly demonstrated and measured? This means that feedback in our experiment was, if anything, too swift and strong.

What accounts for our results, especially in the stage where feedback is provided? It seems that the two Bayesian models significantly overestimate the speed with which subjects switch between the two experts once they observe them giving out wrong answers. Indeed, both the low info

---

<sup>9</sup>We did not choose only questions with counter-intuitive answers, to avoid the possibility that our participants may recognise the pattern and consciously choose counter-intuitive answers because of this. Clearly from their answers in Stage 1, this did not happen at any substantial rate.

and the high info model predict that it would take 3 to 5 questions with feedback for the average subject to identify the Expert. This clearly did not happen in our experiments. On the other hand, the behavioural model, which seems to be much better in explaining our data, suggests that the Charlatan gets a head-start of about 0.6 questions on average simply by being selected over the Expert at Stage 2 (which is the case for the vast majority of participants). This explains why feedback was a weak instrument in improving participants' choices and why, in fact, many ended up choosing the Charlatan until the very end.

Our results have several implications. First, economic expertise, as demonstrated by traditional academic titles, may not be by itself convincing enough for the public. The problem of economic education that [Newcomb \(1893\)](#) lamented on more than a century ago does not seem to have been addressed. The challenge is still to provide as much feedback about the accuracy of economic advice that sounds counter-intuitive to the untrained public. To this end, social media need to be used competently by the experts, at the same time as regulation regarding "fake news" is being contemplated. Our evidence also suggests that information campaigns to help the public identify charlatans in a decentralized way can hardly help. Most of our participants are unwilling or unable to process this information properly. In Kahneman's words, they are unable to accept their ignorance. Or, to put it differently, they do not seem to come to terms with the fact that what sounds right to them in technocratic issues probably is incorrect. Finally, our experimental design can be applied to other disciplines and particularly medicine that suffers from its own populist movements, such as the coronavirus-deniers. Several other disciplines are also experiencing a rising distrust in experts and a topic of future research is whether a common mechanism is behind these phenomena.

## Acknowledgements

Maniadis and Boukouras are supported by BA/Leverhulme small research grant SG162362. We are grateful to John Hey, Roberto Weber, seminar participants at the University of Southampton, City University, King's College London, Leicester University, University of Maastricht, University of Durham, ESA Global Meeting in Los Angeles and CRETE 2019 in Tinos, Greece.

## Appendix A: The Questionnaire

Below are the 11 questions we used in the experiments. We tried more questions (available upon demand), which we deemed prima facie appropriate, but they did not meet our criterion of academic consensus when tested by academic economists. The correct answers are denoted by CO, the Charlatan's answer in red by CH.

1. A prestigious, merit-based, scholarship for graduate studies from a private institution is awarded to 5% of the applicants. Assume that an undergraduate student is chosen at random and applies for the scholarship. What is the likelihood that (s)he will be a recipient of this scholarship?

a) 0% b) Less than 5% c) 5% d) More than 5%

CO: b

CH: c

2. In an attempt to address its housing crisis, manifested through rapidly rising rents and house prices, the mayor of Bigcitia, a burgeoning capital in a high income country, announced that the city will impose a freeze for existing renters and restrict rent increases to 1% of the average price in the neighbourhood for new renters. Will this policy alleviate the housing crisis and result to more people finding a house in the next 5 years?

a) Yes b) Maybe Yes c) No d) All previous options are equally likely to be correct

CO: c

CH: b

3. In Richland at first no Value Added Tax (VAT) or other tax is imposed on fruits. The price of a kilo of apples is 100 Richland Pounds. The government is thinking of imposing a VAT of 24% on all fruits. What will be the price of apples after the market adjusts to the tax?

a) 100 b) Between 100 and 124 c) 124 d) More than 124

CO: b

CH: c

4. The government of Freeland, a multiethnic, without a dominant ethnic group, free market, high income economy announces a new law according to which all workers of ethnicity K

must receive a 50% higher wage than comparable employees. What do you expect to be the effect of the new law on the probability of finding a job for a random member of ethnicity K that is now entering the labour market for the first time?

a) Positive b) Neutral c) Negative d) All previous options are equally likely to be correct

CO: c

CH: c

5. After several successful trials a start up from California has announced the commercial licencing of its eagerly awaited autonomous car technology. Market analysts expect it will take 6 months for taxi companies to obtain regulatory approval and another 6 months to fully deploy the technology. Assuming the analysts' timeline is accurate, what do you expect to be the effect on the employment rate of current taxi drivers 12 months from now?

a) It will increase b) It will be unaffected c) It will decrease d) All previous answers are equally likely to be correct

CO: c

CH: c

6. Hobson Plc and Thornbush Plc announced on Friday, after the stock market had closed, an unexpected merger of equals. During the weekend the majority of economic analysts and financial media, who were surprised by the news, have spoken against it on the basis that it will be unprofitable. What is the most likely price movement for the stock prices of the two companies over the coming week (Monday to Friday) if they are allowed to continue trading their stocks on Monday and no additional news on the value of the two companies arrives to the markets?

a) Both up b) Hobson up, Thornbush down c) Hobson down, Thornbush up d) Both down  
e) No change in either company

CO: d

CH: d

7. Following its commitment to cut global warming emissions, the Prime Minister of Richland announced a 10-year guaranteed price scheme, significantly above current market prices, for buying electricity from new installations of wind and solar power farms in the country. Five

years after the implementation of this policy, the percentage of electrical power produced from renewable sources will be:

- a) Lower b) The same c) Higher d) It cannot be determined

CO: c

CH: c

8. The government of Rainland borrows £100 billion from private investors at a 5% interest rate. If it pays back to investors £5 billion per year, how many years will it take to repay its debt?

- a) 20 b) 40 c) It depends on the borrowing terms d) For ever

CO: d

CH: c

9. Robert won a free ticket to see Justin Bieber. But Beyonce is performing on the same night and he can only attend one of the two events. He likes Beyonce and he would pay up to £50 to see her perform, and the tickets for Beyonce's event cost £40. What must be the minimum value of Bieber's performance to Robert so that Robert chooses Bieber over Beyonce?

- a) £0 b) £10 c) £40 d) £50

CO: b

CH: a

10. The previous government of Girtonia, a developed country, invested \$100 million last year in building a regional airport. The airport is now ready to open its doors and it is expected to generate a total net profit of \$75 million for the duration of its use. The current government is reconsidering the project and has found a new location for the airport. The new airport would yield earnings of \$150 million for the duration of its use and it would also cost \$100 million to build. If the old airport is abandoned it would have \$0 value to the government. Should the government go ahead with the new project?

- a) No b) Yes c) Both projects are equally profitable d) Insufficient information to answer

CO: a

CH: a

11. Kate is a successful investment portfolio manager. In each one of the previous three years her investments in stocks and bonds had 5% average net returns. During the same time period,

the investments of her competitors had 1% average net returns. What is the most likely level of returns for Katie's portfolio in stocks and bonds the coming year?

Less than 5% b) 5% c) More than 5% d) Insufficient information to answer.

CO: d

CH: d

## Appendix B: Theoretical Models

### Low Information Model: Stage 2

Events	Configurations	Total
E1	{1,1,1}, {2,2,2}, {3,3,3}, {4,4,4}.	4
E2	{1,1,2}, {1,1,3}, {1,1,4}, {2,2,1}, {3,3,1}, {4,4,1}, {2,2,3}, {2,2,4}, {3,3,2}, {3,3,4}, {4,4,2}, {4,4,3}.	12
E3	{1,2,1}, {1,3,1}, {1,4,1}, {2,1,2}, {3,1,3}, {4,1,4}, {2,3,2}, {2,4,2}, {3,2,3}, {3,4,3}, {4,2,4}, {4,3,4}.	12
E4	{2,1,1}, {3,1,1}, {4,1,1}, {1,2,2}, {1,3,3}, {1,4,4}, {3,2,2}, {4,2,2}, {2,3,3}, {4,3,3}, {2,4,4}, {3,4,4}.	12
E5	{1,2,3}, {1,2,4}, {1,3,2}, {1,3,4}, {1,4,2}, {1,4,3}, {2,1,3}, {2,1,4}, {3,1,2}, {3,1,4}, {4,1,2}, {4,1,3}, {2,3,1}, {2,4,1}, {3,2,1}, {3,4,1}, {4,2,1}, {4,3,1}, {2,3,4}, {2,4,3}, {3,2,4}, {3,4,2}, {4,2,3}, {4,3,2}.	24

**Table 4:** *Low Info Model, Stage 2: Configurations of answers belonging to each event. The last column of the table gives the total number of configurations belonging to that event. Recall that a configuration  $\{\alpha_p, \alpha_{A1}, \alpha_{A2}\}$  denotes the answers provided by the participant, adviser A1 and adviser A2 respectively. Each question has four possible answers. The full list of events is provided in page 10.*

Events	$\mathbf{P(E S_1)}$
E1	$p_p \cdot p_h \cdot p_l + \frac{(1-p_p)(1-p_h)(1-p_l)}{9}$
E2	$p_p \cdot p_h \cdot (1-p_l) + \frac{(1-p_p)(1-p_h)p_l}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}$
E3	$p_p \cdot p_l \cdot (1-p_h) + \frac{(1-p_p)(1-p_l)p_h}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}$
E4	$p_h \cdot p_l \cdot (1-p_p) + \frac{(1-p_h)(1-p_l)p_p}{3} + \frac{2 \cdot (1-p_p)(1-p_h)(1-p_l)}{9}$
E5	$\frac{2}{3} \cdot \left[ p_h \cdot (1-p_l) \cdot (1-p_p) + p_p(1-p_h)(1-p_l) + p_l(1-p_h)(1-p_p) + \frac{(1-p_p)(1-p_h)(1-p_l)}{3} \right]$
Events	$\mathbf{P(E S_2)}$
E1	$p_p \cdot p_l \cdot p_h + \frac{(1-p_p)(1-p_l)(1-p_h)}{9}$
E2	$p_p \cdot p_l \cdot (1-p_h) + \frac{(1-p_p)(1-p_l)p_h}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}$
E3	$p_p \cdot p_h \cdot (1-p_l) + \frac{(1-p_p)(1-p_h)p_l}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}$
E4	$p_l \cdot p_h \cdot (1-p_p) + \frac{(1-p_l)(1-p_h)p_p}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}$
E5	$\frac{2}{3} \cdot \left[ p_l \cdot (1-p_h) \cdot (1-p_p) + p_p(1-p_l)(1-p_h) + p_h(1-p_l)(1-p_p) + \frac{(1-p_p)(1-p_l)(1-p_h)}{3} \right]$

**Table 5:** *Low Info Model, Stage 2: Probability of event E in state S.*

To illustrate the calculation of the entries of Table 5, let us suppose that the state of the world is  $S_1$  and consider the probability  $p(E1|S_1)$ . With probability  $p_p \cdot p_l \cdot p_h$  all three agents select answer 1, the correct answer, while the probability of all of them selecting answer 2, which is wrong, is equal to  $\frac{(1-p_p)}{3} \cdot \frac{(1-p_l)}{3} \cdot \frac{(1-p_h)}{3}$ . The latter probability also applies for answers 3 and 4. Thus:

$$p(E1|S_1) = p_p \cdot p_l \cdot p_h + 3 \cdot \frac{(1-p_p)}{3} \cdot \frac{(1-p_l)}{3} \cdot \frac{(1-p_h)}{3} = p_p \cdot p_l \cdot p_h + \frac{(1-p_p)(1-p_l)(1-p_h)}{9}$$

One can easily show that the same probability applies for  $p(E1|S_2)$ . Using Table 5 we can obtain the odds ratio for any event. Clearly, for events E1, E4, and E5,  $P(E|S_1) = P(E|S_2)$  and so  $OR(E) = 1$ . Intuitively, events where the two advisers agree have zero diagnostic value for determining who the Expert is. Only events E2 and E3 generate different probabilities under the two states. Moreover,  $\frac{p(E2|S_2)}{p(E2|S_1)} = \frac{p(E3|S_1)}{p(E3|S_2)}$ . The OR of E2 in equation (11) below is derived by dividing  $P(E2|S_2)$  by  $P(E2|S_1)$ . It is easy to verify that  $OR(E3)$  is the inverse of  $OR(E2)$ .

### Deriving the Formula for Bayesian Updating

To calculate the posterior  $P(E|S_1)$  after the participant observes a *single* event  $E \in \{E1, E2, E3, E4, E5\}$ , let  $\pi_0$  denote the prior of the participant for state  $S_1$ . Then:

$$P(S_1|E) = \frac{\pi_0 p(E|S_1)}{\pi_0 p(E|S_1) + (1-\pi_0)p(E|S_2)} = \frac{\pi_0}{\pi_0 + (1-\pi_0)\frac{p(E|S_2)}{p(E|S_1)}} = \frac{\pi_0}{\pi_0 + (1-\pi_0)OR(E)} \quad (10)$$

where  $OR(E) = \frac{p(E|S_2)}{p(E|S_1)}$  is the ‘Odds Ratio’ of event  $E$ . Under events E1, E4, and E5,  $OR(E) = 1$  and so  $P(S_1|E) = \pi_0$ . By replacing the Odds Ratio for events E2 and E3 in equation (10) one obtains the corresponding posteriors. Under E2 the Odds Ratio is given below, while for E3 we have that  $OR(E3) = [OR(E2)]^{-1}$ .

$$OR(E2) = \frac{p_p \cdot p_l \cdot (1-p_h) + \frac{(1-p_p)(1-p_l)p_h}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}}{p_p \cdot p_h \cdot (1-p_l) + \frac{(1-p_p)(1-p_h)p_l}{3} + \frac{2 \cdot (1-p_p)(1-p_l)(1-p_h)}{9}} \quad (11)$$

A *history*  $h$  is a sequence  $\{E_1, E_2, \dots\}$  of events. Since the correct answers to the questions are assumed to be uncorrelated, the events are independent of each other and the posterior can be calculated recursively in the usual way.

$$P(S_1|h) = \frac{P(S_1|h-1)P(E|S_1)}{P(S_1|h-1)P(E|S_1) + P(S_2|h-1)P(E|S_2)}$$

However, the calculations are simplified by the observation that for any two independent events  $E_1$  and  $E_2$ , the posterior of the joint events is equal to:



$$P(S_1|E_1, E_2) = \frac{\pi_0}{\pi_0 + (1 - \pi_0)OR(E_1)OR(E_2)}$$

and so

$$P(S_1|h) = \frac{\pi_0}{\pi_0 + (1 - \pi_0) \prod_{E \in h} OR(E)}$$

Taking into account that  $OR(E) = 1$  for events E1, E4, and E5 and  $OR(E3) = [OR(E2)]^{-1}$ , the expression simplifies further to:

$$P(S_1|h) = \frac{\pi_0}{\pi_0 + (1 - \pi_0) \cdot (OR)^K} \quad (12)$$

where OR is simply  $OR(E2)$ , and  $K \equiv [\text{number of times E2 has been observed}] - [\text{number of times E3 has been observed}]$ .

### Low Information Model: Stage 3

Events	$P(\mathbf{E} \mathbf{S}_1)$	$P(\mathbf{E} \mathbf{S}_2)$
$\hat{E}1$	$p_h \cdot (1 - p_l)$	$p_l \cdot (1 - p_h)$
$\hat{E}2$	$(1 - p_h) \cdot p_l$	$(1 - p_l) \cdot p_h$
$\hat{E}3$	$p_h \cdot p_l$	$p_l \cdot p_h$
$\hat{E}4$	$(1 - p_h) \cdot (1 - p_l)$	$(1 - p_l) \cdot (1 - p_h)$

**Table 6:** *Low Info Model, Stage 3: Probability of event  $\hat{E}$  in state  $S$ . The full list of events in Stage 3 is provided in page 12.*

Given the probabilities expressed in Table 6, the odds ratio  $\hat{OR}(\hat{E}) \equiv \frac{P(\hat{E}|\mathbf{S}_2)}{P(\hat{E}|\mathbf{S}_1)}$  can be calculated in the usual manner. Similarly to Stage 2, the odds ratio for events  $\hat{E}3$  and  $\hat{E}4$  is equal to one, while  $\hat{OR}(\hat{E}2) = [\hat{OR}(\hat{E}1)]^{-1}$ . Dividing  $P(\hat{E}1|\mathbf{S}_2)$  with  $P(\hat{E}1|\mathbf{S}_1)$  gives  $\hat{OR}$  used in equation (13) below. The posterior of  $S_1$  after feedback on a single question and with no other information is given by replacing  $OR$  with  $\hat{OR}$  in equation (10), but with  $\pi_0$  interpreted as the prior at the beginning of Stage 3.

Now, a history  $\hat{h}$  in Stage 3 is a collection of events  $\{\hat{E}_1, \hat{E}_2, \dots, \hat{E}_k, \}$  (for questions with feedback) and a collection  $\{E_1, E_2, \dots, E_{11-k}, \}$  (for questions without feedback). The posterior of  $S_1$  after a history  $\hat{h}$  can be derived by modifying equation (12) accordingly. Let us define as  $\hat{N}$  the set of questions for which feedback has been provided in Stage 3. Let  $K_{\hat{N}}$  measure the number of times that A1 gave a correct answer and A2 gave a wrong answer in  $\hat{N}$ , minus the number of times that

A2 gave a correct answer and A1 gave a wrong answer in  $\hat{N}$ . Let  $K_{-\hat{N}}$  measure the number of times that A1 alone had a common answer with the participant minus the number of times that A2 alone had a common answer with the participant in the set of questions that do not belong to  $\hat{N}$ .  $K_{-\hat{N}}$  is defined in the same way as measure  $K$  in Stage 2, but for only a subset of questions. Given the above definitions, the formula for Bayesian updating for Stage 3 can be expressed as:

$$P(S_1/h) = \frac{\pi_0}{\pi_0 + (1 - \pi_0) \cdot (OR)^{K_{-\hat{N}}} \cdot \hat{OR}^{K_{\hat{N}}}} \quad \text{with} \quad \hat{OR} = \frac{p_l \cdot (1 - p_h)}{p_h \cdot (1 - p_l)} \quad (13)$$

Note that the above equation implies that for questions with feedback in Stage 3, the events of Stage 2 (whether or not the participant had common answers with the advisers) are not relevant for the calculation of the posterior. Only the events of Stage 3 (whether the advisers were correct or not) matter. Let us explain why this is the case. Suppose two histories,  $h$  and  $\hat{h}$ , which correspond to Stage 2 and Stage 3 respectively, and which concern the same questions, say questions 1 to 3.

For instance, assume  $h$  specifies that in questions 1 to 2 the participant agreed only with A1 and then in question 3 he only agreed with A2. On the other hand, let  $\hat{h}$  specify that in questions 1 to 3, A2 was always right (so that A1 was always wrong). Under the information of only Stage 2, the participant infers that she has more common answers with A1 than A2 and so  $S_1$  is more likely than  $S_2$ . However, once the feedback of Stage 3 is made available in  $\hat{h}$ , this conclusion is overturned since A1 was wrong in all of them. Intuitively, the information on who is correct on each question from Stage 3 generates a finer partition of the participant's information set, so that the coarser partition from Stage 2 is obsolete for these questions. Continuing our example, for questions beyond question 3, namely 4 to 11, the participant has no feedback from Stage 3 and so she can utilize only the coarse partition of Stage 2. Since the correct answers across questions are independent, this gives this convenient formula.

## High Information Model: Stage 2

The list of events and the configurations of answers per event are the same as in the low info model and they are provided by Table 4. The table below gives the probabilities of these events in each state. These are computed by estimating the probability of each configuration as in the example provided in the main text.

Table 7 allows us to calculate the odds ratios for each event. Since  $P(E|S_1) = P(E|S_2)$  for events E1, E4, and E5, these events have odds ratios equal to one. For E2, the odds ratio is given by

Events	$P(E S_1)$
E1	$p_p \cdot \pi \cdot p_h + (1 - p_p) \cdot \pi \cdot (1 - p_h)/3$
E2	$p_p \cdot (1 - \pi) \cdot p_h + (1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h)/3$
E3	$p_p \cdot \pi \cdot (1 - p_h) + (1 - p_p) \cdot \pi \cdot p_h + (2/3) \cdot (1 - p_p) \cdot \pi \cdot (1 - p_h)$
E4	$(1/3) \cdot (1 - p_p) \cdot (1 - \pi) \cdot p_h + (2/9) \cdot (1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h) + (1/3) \cdot p_p \cdot (1 - \pi) \cdot (1 - p_h)$
E5	$(2/3) \cdot p_p \cdot (1 - \pi) \cdot (1 - p_h) + (4/9) \cdot (1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h) + (2/3) \cdot (1 - p_p) \cdot (1 - \pi) \cdot p_h$
Events	$P(E S_2)$
E1	$p_p \cdot \pi \cdot p_h + (1 - p_p) \cdot \pi \cdot (1 - p_h)/3$
E2	$p_p \cdot \pi \cdot (1 - p_h) + (1 - p_p) \cdot \pi \cdot p_h + (2/3) \cdot (1 - p_p) \cdot \pi \cdot (1 - p_h)$
E3	$p_p \cdot (1 - \pi) \cdot p_h + (1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h)/3$
E4	$(1/3) \cdot (1 - p_p) \cdot (1 - \pi) \cdot p_h + (2/9) \cdot (1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h) + (1/3) \cdot p_p \cdot (1 - \pi) \cdot (1 - p_h)$
E5	$(2/3) \cdot p_p \cdot (1 - \pi) \cdot (1 - p_h) + (4/9) \cdot (1 - p_p) \cdot (1 - \pi) \cdot (1 - p_h) + (2/3) \cdot (1 - p_p) \cdot (1 - \pi) \cdot p_h$

**Table 7:** High Information Model, Stage 2: Probability of event  $E$  in state  $S$ .

dividing  $P(E|S_2)$  by  $P(E|S_1)$  and this gives the odds ratio in (5) in page 14. It is straightforward to verify that the odds ratio for E3 is the inverse of E2.

### High Information Model: Stage 3

The full description of the possible events in questions with feedback for the high info model is given in page 15. The following table gives the probabilities of each event for each state.

Events	$P(\tilde{E} S_1)$	$P(\tilde{E} S_2)$
$\tilde{E}1$	$p_p p_h \pi$	$p_p \pi p_h$
$\tilde{E}2$	$p_p (1 - p_h) (1 - \pi)$	$p_p (1 - \pi) (1 - p_h)$
$\tilde{E}3$	$p_p p_h (1 - \pi)$	$p_p \pi (1 - p_h)$
$\tilde{E}4$	$(1 - p_p) p_h \frac{(2+\pi)}{3}$	$(1 - p_p) \frac{(1-\pi)}{3} (1 - p_h)$
$\tilde{E}5$	$p_p (1 - p_h) \pi$	$p_p (1 - \pi) p_h$
$\tilde{E}6$	$(1 - p_p) p_h \frac{(1-\pi)}{3}$	$(1 - p_p) \frac{(1-\pi)}{3} p_h$
$\tilde{E}7$	$(1 - p_p) (1 - p_h) \frac{(1-\pi)}{3}$	$(1 - p_p) \frac{(2+\pi)}{3} p_h$
$\tilde{E}8$	$(1 - p_p) (1 - p_h) \frac{(2+\pi)}{3}$	$(1 - p_p) \frac{(2+\pi)}{3} (1 - p_h)$

**Table 8:** High Information Model, Stage 3: Probability of event  $\tilde{E}$  in state  $S$ .

From the above table we observe that the odds ratios are equal to one for events  $\tilde{E}1$ ,  $\tilde{E}2$ ,  $\tilde{E}6$ , and  $\tilde{E}8$ . Dividing  $P(\tilde{E}|S_2)$  by  $P(\tilde{E}|S_1)$  for events  $\tilde{E}3$  and  $\tilde{E}4$  one obtains the odds ratios in equations (7) and (8) respectively in page 16. Finally, one can easily verify that the odds ratio for  $\tilde{E}5$  is the inverse of  $\tilde{E}3$  and the odds ratio for  $\tilde{E}7$  is the inverse of  $\tilde{E}4$ .

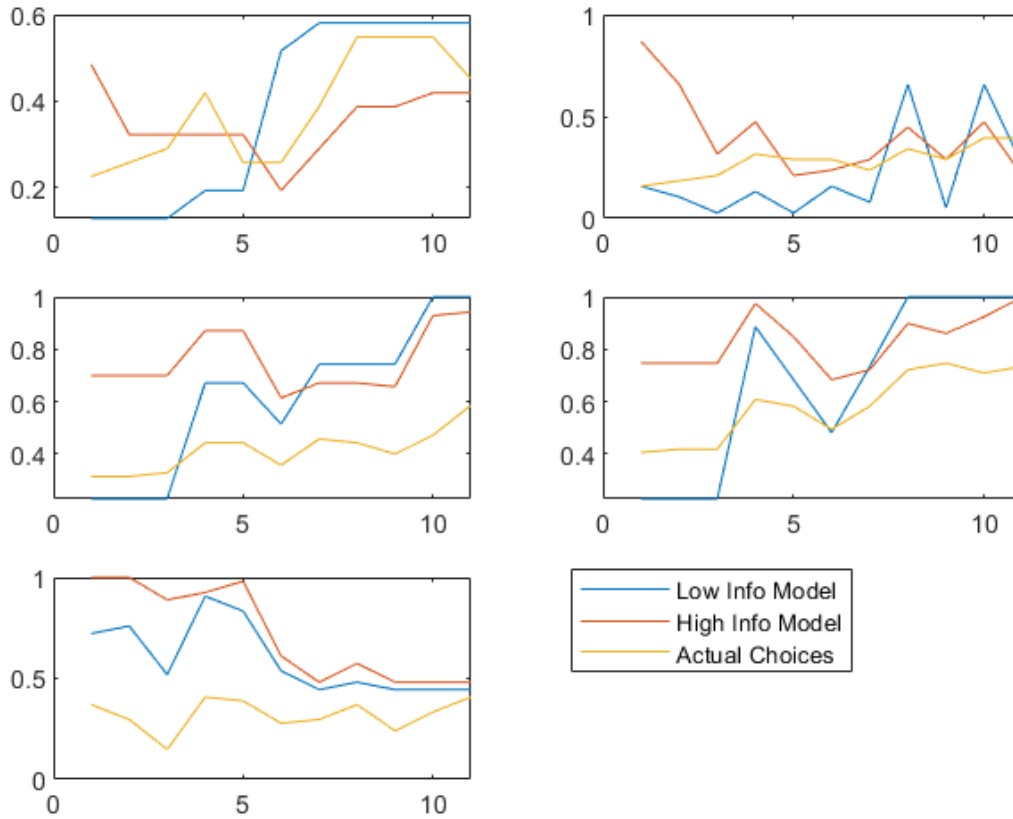
## Appendix C: Graphs by Condition

Our experiments had a total of 5 conditions, including the additional checks. Conditions 1, 2 and 3 are low info conditions corresponding to Experiment 1. Conditions 4 and 5 are high info conditions, corresponding to Experiment 2. In Condition 1 we allowed the expert to be correct 70% of the time, against 85% in all other four conditions. Condition 3 was high-info and it was the only condition without financial incentives. Conditions 4 and 5 differed in the order of the questions, in that in Condition 4 the order was harder than in 5 for the participants to distinguish between the two advisers. The table below summarises the characteristics of each condition.

Condition	Description
1	low-info, 70% expert accuracy, incentivised
2	low-info, 85% expert accuracy, incentivised
3	high-info, 85% expert accuracy, not incentivised
4	high-info, 85% expert accuracy, incentivised
5	high-info, 85% expert accuracy, incentivised, reversed order

**Table 9:** *Summary of the five conditions*


The graphs below present the predictions of the two models broken down by condition and juxtapose it with actual behaviour. As we can see, in Condition 1 convergence to optimal choice (choosing the Expert) is expected to be relatively slow by the two Bayesian models. The low-info model in particular does not seem to capture the data well. In Condition 2 the low-info model seems to follow the data relatively closely. In the remaining three Conditions it is the high-info model that captures the environment theoretically. In Condition 3 the high-info model consistently overpredicts the tendency of participants to choose the Expert. In Condition 4 the prediction of this model is closer to the participants' behaviour. The stronger tendency of participants to choose the Expert in late rounds of Condition 4 relative to Condition 3 indicates that incentives may have played a role in Condition 4. Finally, in Condition 5 the model predicts that a high fraction of participants chooses the Expert at the beginning of Stage 3, which is not what happens in the data. However, the model seems to converge to actual behaviour in late rounds, in terms of its prediction of the fraction of the participants that choose the Expert.



**Figure 6:** *Rational Models Over Time.* The horizontal axis shows the number of rounds (questions) of stage three and the vertical axis the fraction of participants who chose the ‘Expert’. The top left panel depicts condition 1, then top right is condition 2 and so on.

## References

- AKERLOF, G. A. AND W. T. DICKENS (1982): “The Economic Consequences of Cognitive Dissonance,” *The American Economic Review*, 72, 307–319.
- ANDRE, P., C. PIZZINELLI, C. ROTH, AND J. WOHLFART (2019): “Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample,” *CESifo Working Paper*.
- APREA, C. AND V. SAPPA (2014): “Variations of Young Germans’ Informal Conceptions of Financial and Economic Crises Phenomena,” *Journal of Social Science Education*, 13, 57–67.
- ARTHUR, W. B. (2000): “Cognition: The Black Box of Economics,” in *The Complexity Vision and the Teaching of Economics*. D. Colander (Ed.), Northampton, MA: Edward Elgar Publishing.



**Figure 7:** Comparing the fit of the three models in the various conditions.

BARTELS, L. M. (2005): “Homer Gets a Tax Cut: Inequality and Public Policy in the American Mind,” *Perspectives on Politics*, 3, 15–31.

BÉNABOU, R. AND J. TIROLE (2002): “Self-Confidence and Personal Motivation,” *The Quarterly Journal of Economics*, 117, 871–915.

CAMERER, C. F. AND R. M. HOGARTH (1999): “The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework,” *Journal of Risk and Uncertainty*, 19, 7–42.

CAPLAN, B. (2002): “Systematically Biased Beliefs about Economics: Robust Evidence of Judgmental Anomalies from the Survey of Americans and Economists on the Economy,” *The Economic Journal*, 112, 433–458.

——— (2011): *The Myth of the Rational Voter: Why Democracies Choose Bad Policies-New Edition*, Princeton University Press.

CHAKRABORTY, A., P. GHOSH, AND J. , ROY (2020): “Expert Captured Democracies,” *American Economic Review*, 110, 1713–51.

- COLANDER, D. (2005): “The Making of an Economist Redux,” *Journal of Economic Perspectives*, 19, 175–198.
- DAL BÓ, E., P. DAL BÓ, AND E. EYSTER (2018): “The Demand for Bad Policy when Voters Underappreciate Equilibrium Effects,” *The Review of Economic Studies*, 85, 964–998.
- DIXON, R., W. GRIFFITHS, AND G. LIM (2014): “Lay People’s Models of the Economy: A Study Based on Surveys of Consumer Sentiments,” *Journal of Economic Psychology*, 44, 13–20.
- DRÄGER, L., M. J. LAMLA, AND D. PFAJFAR (2016): “Are Survey Expectations Theory-Consistent? The Role of Central Bank Communication and News,” *European Economic Review*, 85, 84–111.
- EREV, I. AND A. E. ROTH (1998): “Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review*, 848–881.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments,” *Experimental Economics*, 10, 171–178.
- FREDERICK, S. (2005): “Cognitive Reflection and Decision Making,” *Journal of Economic Perspectives*, 19, 25–42.
- GANGL, K., B. KASTLUNGER, E. KIRCHLER, AND M. VORACEK (2012): “Confidence in the Economy in Times of Crisis: Social Representations of Experts and Laypeople,” *The Journal of Socio-Economics*, 41, 603–614.
- GENTZKOW, M. AND J. M. SHAPIRO (2006): “Media Bias and Reputation,” *Journal of Political Economy*, 114, 280–316.
- GEORGANAS, S., P. J. HEALY, AND R. A. WEBER (2015): “On the Persistence of Strategic Sophistication,” *Journal of Economic Theory*, 159, 369–400.
- JAVDANI, M. AND H.-J. CHANG (2019): “Who Said or What Said? Estimating Ideological Bias in Views Among Economists,” *SSRN No 3356309 Working Paper*.
- JERRIM, J., P. PARKER, AND D. SHURE (2019): “Bullshitters. Who Are They and What Do We Know about Their Lives?” *IZA Discussion Paper*.

- KAHNEMAN, D. (2011): *Thinking, Fast and Slow*, Macmillan.
- KELEMEN, D. AND E. ROSSET (2009): “The Human Function Compunction: Teleological Explanation in Adults,” *Cognition*, 111, 138–143.
- KELEMEN, D., J. ROTTMAN, AND R. SESTON (2013): “Professional Physical Scientists Display Tenacious Teleological Tendencies: Purpose-based Reasoning as a Cognitive Default.” *Journal of Experimental Psychology: General*, 142, 1074.
- KRUGMAN, P. (2010): “Block Those Metaphors,” *The New York Times*, 12.
- LEISER, D. AND R. AROCH (2009): “Lay Understanding of Macroeconomic Causation: The Good-Begets-Good Heuristic,” *Applied Psychology*, 58, 370–384.
- LEISER, D., S. BOURGEOIS-GIRONDE, AND R. BENITA (2010): “Human Foibles or Systemic Failure—Lay Perceptions of the 2008–2009 Financial Crisis,” *The Journal of Socio-Economics*, 39, 132–141.
- LEISER, D. AND Z. KRILL (2017): “How Laypeople Understand the Economy,” *Economic Psychology*, 139–154.
- MENCKEN, H. L. (2012): *Mencken Chrestomathy*, Vintage.
- NEWCOMB, S. (1893): “The Problem of Economic Education,” *The Quarterly Journal of Economics*, 7, 375–399.
- OBERLECHNER, T., T. SLUNECKO, AND N. KRONBERGER (2004): “Surfing the Money Tides: Understanding the Foreign Exchange Market Through Metaphors,” *British Journal of Social Psychology*, 43, 133–156.
- PENNYCOOK, G., J. A. CHEYNE, N. BARR, D. J. KOEHLER, AND J. A. FUGELSANG (2015): “On the reception and detection of pseudo-profound bullshit,” *Judgment and Decision Making*, 10, 549–563.
- PENNYCOOK, G., Z. EPSTEIN, M. MOSLEH, A. A. ARECHAR, D. ECKLES, AND D. G. RAND (2019): “Understanding and reducing the spread of misinformation online,” .
- PENNYCOOK, G., J. MCPHETRES, Y. ZHANG, J. G. LU, AND D. G. RAND (2020): “Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy nudge intervention,” .



- RONAYNE, D. AND D. SGROI (2018): “Ignoring Good Advice,” *Competitive Advantage in the Global Economy (CAGE)*, Working Paper No. 359.
- SCHOTTER, A. (2003): “Decision Making with Naive Advice,” *American Economic Review*, 93, 196–201.
- TVERSKY, A. AND D. KAHNEMAN (1980): “Causal Schemas in Judgments under Uncertainty,” *Progress in Social Psychology*, 1, 49–72.
- VOSOUGHI, S., D. ROY, AND S. ARAL (2018): “The spread of true and false news online,” *Science*, 359, 1146–1151.
- ZINGALES, L. (2020): “The Political Limits of Economics,” *Paper Presented at the 2020 American Economic Association Meetings*.