Revising Beliefs About Advice Content, But Not Its Quality: A Laboratory Experiment on the Use and Demand for Advice (Preliminary and Incomplete). *

Jacopo Bregolin† Astrid Hopfensitz‡ Elena Panova§

Abstract

We experimentally test how the contents of advice, namely its alignment with common priors influence beliefs about its quality and future demand for advice. We reject the hypothesis that demand for advice can be increased by giving advice in alignment with common priors. Indeed our results show that participants of our experiment do not view such advice as a signal of quality regarding the advisor. However the given advice is influential by impacting participants in their guesses in an incentivized task. Hence, participants use advice, while largely foregoing inferences on the quality of advice itself.

Keywords: demand for information, belief updating.

JEL Code: D90, C91, D83.

*We thank Stéhane Cezera for his valuable help with organising the experiment. Elena Panova acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program).

†University of Liverpool, Management School, Department of Economics. E-mail: jacopo.bregolin@liverpool.ac.uk

‡emlyon business school, GATE, Lyon, France. E-mail: hopfensitz@em-lyon.com

§Toulouse School of Economics, University of Toulouse Capitole. E-mail: e_panova@yahoo.com

1 Introduction.

In many situations, private decision makers buy advice from experts. Examples include purchasing financial advice from intermediaries, paying physicians for medical guidance, or buying news from the media. Consumers of professional advice may remain uncertain about its quality. The high volatility of investment returns makes it difficult to evaluate the effectiveness of an investment strategy. Similarly, the lack of counterfactual knowledge makes it hard to assess whether a prescribed treatment was optimal, or which political candidate or policy would have been the best. Therefore, experts may bias their advice in their own interest.

An influential theoretical literature on reputational cheap-talk suggests that experts will confirm the common priors regardless of their true opinion so as to signal their competence. The underlying assumption is that consumers update their beliefs about an advisor's competence based on the content of the received advice using Bayes rule, and that their demand for advice increases in their posteriors. However, there is sizable evidence of errors in belief updating (for surveys see: Camerer (1998) and Benjamin (2019)). Furthermore, the empirical and experimental literature on demand for information (discussed below) has documented numerous deviations from optimal behavior. In this paper we present results from a laboratory experiment designed to test whether-or not the alignment of advice with common priors helps increasing demand for advice in the future.

The setup of our experiment is the following: participants of our experiment are guessing the color of a ball "drawn" by their computer from a jar with ten balls, most of which are green and the others are blue. Correct guesses are rewarded. The participants have computerized advisors which generate advises of uncertain quality. Indeed, an advisor is equally likely to be "perfect" or "defective". A perfect advisor generates correct advice, while a defective advisor generates advice which is equally likely to be correct or false. Participants interact with the same advisor(s) across two periods.[1] In the first period, the participants receive advice for free. In the second period, they decide whether- or not to buy advice before guessing.

We will consider six different experimental treatments. In the four baseline treatments, we vary the precision of *common priors*, measured by the number of green balls in the jar: either 6 out of 10 ("weak") or 7 out of 10 ("strong"); and the *market structure*, with the number of advisors set to either one ("monopoly") or two ("duopoly")[2] In these treatments, we did not elicit beliefs about the quality of the advisors, instead asking about these beliefs only in a post-experimental survey. Additionally, we ran one monopoly treatment for each prior (i.e. weak and strong), where we elicited beliefs by participants about the quality of the advisors after receiving advice in the first period.

We find that, across all treatments, the demand for advice in period two does not depend on the content of the advice received in period one. Hence, our subjects do not reward confirmatory advice with a higher demand.[3] In the same vein, the rate of following

[1]We will call these two period interactions a "round". Participants are observed across fifteen rounds in total, each time with a new set of advisors.

[2]This is inspired by the theoretical debate on the impact of competition on the quality of professional advice.

[3]Indeed, in the monopoly treatments, about three-quarters of participants either always or never buy advice. For the remaining subjects, the demand in a given round tends to match the demand from the previous round, but only if their guess in the previous round was correct. It's important to note that this "win-stay, lose-shift" rule cannot improve guessing, as the quality of advisors in different rounds is

advice in period two, conditional on buying, does not depend on the content of the advice received in period one.

This pattern matches the elicited beliefs on the quality of advisors. The beliefs elicited during the experiment, as well as those reported in the post-experimental survey, show only a minor tendency to view an advice confirming the common priors as a signal of quality. About half of the participants do not update their beliefs at all.

At the same time, we find clear evidence of participants using the received advises when guessing. The observed patterns of guessing suggest that subjects use both the priors and the advises. Indeed, in our monopoly treatments, participants tend to follow advice, either received or bought. They are more likely to follow advice recommending "green" (likely outcome) than those recommending "blue"(unlikely outcome), and this difference is larger in the strong prior treatments.

These observations suggest that advice induces participants to update their beliefs regarding the state of the world, while largely foregoing inferences regarding the quality of advice itself. Similarly, news coverage of controversial issues may shift beliefs about those issues, but not about the quality of the media outlet providing the information, which could create opportunities for propaganda.

Related literature. Our findings contribute to a sizable economic literature studying deviations from the Bayesian paradigm in learning and demand for information.

Mostly related are recent experimental studies investigating the choice over sources of

independent. Similarly, in the duopoly treatments, about eighty percent of period two demand comes from subjects who either never buy advice or always buy advice.

instrumentally valuable information. Charness et al. (2021) and Montanari and Nunnari (2023) consider neutral settings without a scope for motivated reasoning, reminiscent of our setting. Participants of their experiments aim at guessing the color of a ball drawn from the jar. They are restricted to select one of two computerized advisors which are biased either towards or against the common priors. In Charness et al. (2021) the participants tend to choose an advisor with a certain bias even though it may be optimal to consult either type of advisor depending on the nature of the bias. In Montanari and Nunnari (2023) the participants tend to (optimally) choose the advisor biased towards the common priors when advisors are equally biased. Otherwise, they pick the least biased advisor even if it is suboptimal. In both these experiments the participants know both the quality of their advisors and their biases. We consider unbiased advisors of uncertain quality, and we study how beliefs on the quality of advisors, demand for their advice, and the rate of following their advises depend on the alignment of previously received advises with common priors.

Our interest in the impact of advice contents on beliefs on its quality relates us to Schoar and Sun (2024). They show, using a randomized controlled trial, that participants rate financial advice significantly higher when it aligns with their reported priors regarding the best investment strategy (either passive or active). They elicit beliefs on quality making it salient. In their setting, there is a scope for belief confirmation motives, which where shown to be as important as accuracy concerns (Chopra et al., 2024) .[4]

Our study of demand for instrumentally useful unbiased information in non-strategic

[4]Chopra et al. (2024) quantify the trade-off between accuracy concerns and belief confirmation motives in the filed experiment studying demand for news by the US voters. They find a similar quantitative importance of both motives.

setting without a scope for belief confirmation motives relates us to Ambuehl and Li (2018). They elicit valuation of viewing an informative signal on a binary state of the world before guessing the state. They find overvaluation of low-quality information, undervaluation of high-quality information (compression effect) and a preference for information structures that may remove uncertainty. In their experiment there is no uncertainty about the quality of information, unlike in our experiment.

Our observations regarding the use of advises receive in period one and bought in period two relates us to Liang (2024) (xxx right citation?)) who shows a tendency to underreact to information of uncertain quality using both laboratory experimental data and observational data on stock price reactions to analyst earnings forecasts.

Theoretical predictions for our experiment are based on a reputational cheap talk model (Ottaviani and Sørensen, 2006a,b) extended to two advisors.[5] This connects us to the literature testing this model. Camara and Dupuis (2021) find, using structural estimation of this model, that professional movie reviewers bias their movie ratings towards the common priors, which is consistent with the model predictions. Meloso et al. (2023) (xxx right citation?) test the model in a laboratory experiment. In their experiment, the participants playing the role of reporters (or advisors in our terminology) are incentivized to signal a high quality of their private signals on a hidden binary state to either computerized- or human evaluators. They tend to confirm the common priors, more so in treatments with greater certainty about the state.[6] Human evaluators are incentivized to make precise

[5] We consider a nested information structure, so that the advisors' agreement can not be used as a signal of their quality.

[6] In this experiment, the reporters are free to choose the contents of their reports only when their signals

assessments of quality of signals available to their advisors' depending on the contents of advice. They know whether the received advice was correct or false but, naturally, they are uncertain about the advisor's strategy. Their assessments are consistent with Bayesian updating (in that posteriors are ordered in a Bayesian way) for some beliefs about that strategy. We study demand for unbiased computerized advice of uncertain quality without providing any incentives to update beliefs on the quality, and we find that confirmatory advice is not rewarded with a higher demand in low salience treatments (update upon new results). Comparison of behaviors by our subjects and human evaluators in Meloso et al (2017) suggests that humans update their beliefs on the quality of instrumentally relevant information upon its alignment with common priors only iff they are incentivized to make this exercise.

The rest of this paper is organized as follows: section 2 describes the underlying model and its predictions, section 3 describes the experimental design and the procedure, section 4 presents our findings, and section 5 concludes.

contradict the common priors, otherwise, they are restricted to report their signals. In treatments with probability of the most likely state being 0.6, 51% of reporters choose to confirm the priors instead of revealing the signal contradicting the priors. In treatments with that probability being 0.8, the figure is 63%.

2 Theoretical predictions of rational behavior

2.1 Underlying model

We start with describing the optimal demand for advice, as it is modeled in the reputational cheap-talk literature. An individual would like to guess the hidden state of Nature in two successive periods indexed with $t = 1, 2$. Period specific state x_t is drawn from the following distribution[7]

$$x_t = \begin{cases} G\text{, with probability } p \\ B\text{, with probability } 1-p, \end{cases} \tag{1}$$

where $p > \frac{1}{2}$. The states in different periods are independent.

Before guessing, the individual can get one or two period-specific signals on the prevailing state. These signals are termed "advice". If only one advice is available it is called "monopoly advice". If two types of advice are available, they are called "competing advice". We index advice with $i = l, r$.[8] Advice i in period t is denoted with a_t^i. In period 1 advice is received for free. In period 2, advice can be bought at price ε per advice.

Quality q^i of advice i is the same in either period. It is equally likely to be perfect ($q^i = 1$) or defective ($q^i = 0$).[9] The quality of advice l and r is independent. If an advice is perfect it is correct. If it is defective, its contents depend on a period specific flip of a fair coin, the outcome of which is denoted with z_t. If the coin comes up heads ($z_t = H$), a

[7] G stands for "green" and B stands for "blue".

[8] During the experiment the advisor indexed with l (r) appears on the left (right) side of the computer screen.

[9] Because in our experiment the advisors are represented by machines (robots), we use the term "perfect" instead of "smart" usually used in the literature on reputational cheap-talk, and "defective" instead of "dumb".

defective advice is correct, if the coin comes up tails ($z_t = H$), a defective advice is false:

$$a_t^i = \begin{cases} x_t, \text{ if } q^i = 1 \text{ or } q^i = 0 \text{ and } z_t = H; \\ \{G, B\} \setminus \{x_t\}, \text{ otherwise.} \end{cases}$$

Notice that the above information structure (commonly termed as "*nested*") guarantees that an agreement of advice received from two advisors does not signal their qualities. Indeed, advice of the same quality, either perfect or defective, always agree. At the same time, disagreement of advice implies that one of them is perfect and the other is defective.

2.2 Calibration

We limit the precision of the common priors to the following interval:

$$0.5 < p < 0.75. \tag{2}$$

The lower limitation tells us that the common prior is better information than a defective advice. The upper limitation tells us that advice "blue" is likely to be correct. For the simplicity of our experimental design, we let p be a multiplier of ten, consequently we study two values of p: with values 0.6 and 0.7.

In order to isolate risk and loss-aversion considerations in demand for advice as much as possible, we let the price of receiving advice be substantially smaller than the stake of the decision.[10] Specifically, we assume that a correct decision (i.e. guess) is rewarded with

[10]This is relevant in many applications: the price of financial advice is usually small relative to the stakes

$V = 500$ and the price of receiving an advice is $\varepsilon = 5$. Thus it is worth buying an advice iff it increases the probability to make the correct decision by at least one chance out of hundred.

2.3 Rational behavior

This section outlines the theoretical predictions regarding the optimal use of advice in period 1 and the optimal demand for advice in period 2 (technical details are presented in Appendix xxx).

Optimal behavior with monopoly advice. Suppose first that only one advice is available, say that by advisor l.

Prediction 1: guessing in period 1. *It is optimal to follow monopoly advice received in period* 1 *for either* $p \in \{0.6, 0.7\}$.

Indeed, the best alternative to following advice is guessing "G", because state "G" is more likely. This alternative is the more attractive the higher is the prior probability p. However, as the expected quality of monopoly advice is sufficiently high, following advice improves guessing for either p in set $\{0.6, 0.7\}$.

Prediction 2: demand for advice at a given price and guessing in period 2. *If* $p = 0.6$*, it is optimal to buy- and follow monopoly advice in period* 2 *regardless of whether advice received in period* 1 *was "*G*" or "*B*". If* $p = 0.7$*, it is optimal to buy and follow monopoly advice in period* 2 *if advice received in period* 1 *was "*G*". Otherwise, it is optimal*

of investment decisions, news is cheap as compared to the value of private decisions it may guide.

to guess "G" without buying advice.

Note that perfect advice endorses guess "G" with probability $p > \frac{1}{2}$ while defective advice is equality likely to endorse "G" or "B". Therefore, the probability of advice being perfect conditional on period 1 advice being "G" is above the prior $\frac{1}{2}$, while that conditional on period 1 advice being "B" is below $\frac{1}{2}$. These differences increase in p. They are sufficiently large to affect the demand if $p = 0.7$, but not if $p = 0.6$.[11]

Prediction 3 (willingness to pay for advice in period 2, WTP). *If $p = 0.6$, the WTP is equal to: 86.36 if period 1 advice is G and 61.11 if period 1 advice is "B". If $p = 0.7$, the WTP is equal to: 45.83 if period 1 advice is G and 0 if period 1 advice is "B".*

We compute WTP in Appendix A.3. as the difference between the posterior probability to receive correct advice and the prior probability p scaled by V.

Optimal behavior with competing advises. Suppose now that both, advice l and advice r, are available.

Prediction 4 (guessing in period 1). *The optimal period 1 guess is: "B" if both advice are "B", and "G" otherwise.*

Indeed, when advisors endorse the same guess it is optimal to follow them. When the advisors disagree there are two possible explanations: either (1) period 1 state is "B", the quality of advice recommending "B" is perfect and the other advice is defective or (2)

[11] It is optimal to buy and follow advice if $p < 0.68$. Otherwise, it is optimal to guess G without buying advice.

period 1 state is "G", the the quality of advice recommending "G" is perfect and the other advice is defective. Explanation (2) is more likely because it is more likely that the state is "G".

Prediction 5 (demand for advice and guessing in period 2). *If both advisors in period* 1 *say the same, it is optimal to buy from both advice in period* 2 *and guess: "G" if at least one period* 2 *advice is "G", and "B" otherwise. If period* 1 *advice differs, it is optimal to buy from the advisor who recommended "G" in period* 1 *and follow this advice in period* 2*.*

To determine the optimal decision making rule (DMR) in period 2, we have to compare the potential strategies. We consider the following three: (i) guess "G" independently, that is, without buying any advice; (ii) buy- and follow the advice which is most likely to be perfect (hereafter, *""the best" advice"*); (iii) buy both advises and guess: "G" if at least one recommends "G", and "B" otherwise.

DMR (ii) dominates DMR (i) iff the probability that "the best" advice is perfect is sufficiently high relative to the precision of common priors. For our calibration, this is true for any combination of advice received in period 1

It remains to compare DMR (ii) and DMR (iii). Note that an additional advice may affect guessing only when the two advisors are of different quality and the coin comes up tails. In this case, it corrects guessing when the state is "G", "the best" advice is defective and the other advice is perfect, and it creates an error when the state is "B", "the best" advice advice is perfect and the other advice is defective. Therefore, DMR (iii) is optimal when the posterior probabilities that advises are perfect are not too different. This condition is met if advises agree today, but not otherwise.

Appendix A.6 presents the table summarizing the willingness to pay for competing advises.[12]

3 Experimental design and data.

Overview. Based on the above described theoretical framework we developed a controlled laboratory task. In the following, we will describe the details of this task.

In the task, the participant (referred to as "she") was guessing the color of a ball "drawn" by her computer from a jar with ten balls, of which most were green with the remaining balls being blue.[13] All probabilities were objectively given. Correct guesses were rewarded with points converted into euros.

The advisors in the experiment were computerized (represented by different drawings of a robot) and provided advice of uncertain quality. Participants were informed that each advisor was equally likely to be perfect or defective. A perfect advisor always provided correct advice, while a defective advisor offered advice that was equally likely to be correct or false.

Guessing was done, across two periods. In period one, advice was provided for free,

[12]Our main finding is no demand response to the contents of advice received in period one and lack of updating of beliefs on advisors' quality upon the contents of this advice. We find support for this finding based on the data on WTP in monoply treatments. In duopoly treatments we dispose few observations to make conclusive statements about WTP for advice by a given advisor. EDIT

[13]Note that this task has no intrinsic relevance for the participants, which limits the scope for motivated reasoning. Indeed, a growing evidence shows that when information relates to personal traits, individuals tend to give more weight to positive news than negative news.

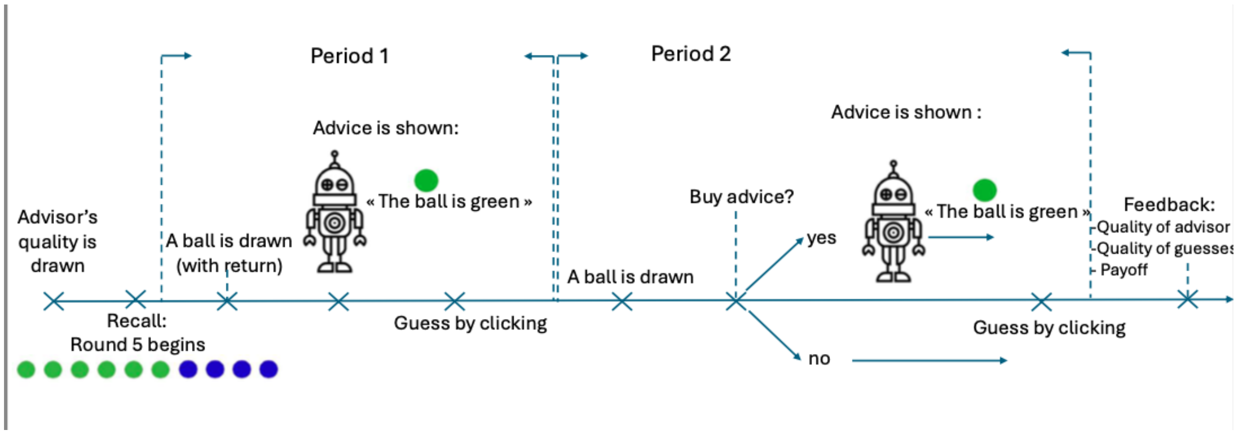

Figure 1: sequence of events in a round.

while in period two, participants had the option to purchase advice. We call these two periods a round. The sequence of events in a round is illustrated in Figure 1. After each two periods, a new round started with new advisor(s) who remained the same until the end of that round. Guessing was done across a total of 15 rounds, with one of these rounds randomly selected for payment.

The price of advice in period two was fixed in rounds 1 to 10. In the final five rounds, the participant's willingness to pay (WTP) for advice was elicited. Participants were asked to propose a price for advice, and they would receive the advice only if their proposed price was at least as high as a hidden, randomly generated value. This procedure created the incentives for participants to propose a price equal to their WTP which we have explained in instructions.

To ensure that participants recalled that the quality of the advisors was the same across the two periods in a round and independent across rounds we presented advisors as a robot. The image remained the same in the first and the second period of the same round and changed at the beginning of each round. Advice provided by the advisor was visualized by a colored ball (either green or blue) presented by the robot.

Participants received information on their performance at the end of each round. After submitting their guess in period two, they were shown a screen containing information on the quality of the advisor(s), the color of the drawn ball in each period, a reminder of what the advisor(s) said, the participants' choices (guesses and purchasing decisions), and the points obtained in that round.

Treatments We conducted six different treatments. In the four *baseline* treatments, we vary the precision of the common priors, and the number of advisors. The precision of common priors was varied by adjusting the number of green balls in the jar. In treatments with a "weak prior", 6 out of 10 balls were green, in treatments with a "strong prior", 7 out of 10 balls were green. The number of advisors was varied by having either one advisor in the two periods (monopoly) or having two advisors (duopoly). We used a within-subject design regarding the two market structures. The order of the market structure was counterbalanced. Half of the participants began the experiment with the monopoly treatment, and the other half started with the duopoly treatment. Common priors were presented on a between-subject level.

In the above-described baseline treatments, we did not elicit beliefs about the quality of the advisors, inquiring into these beliefs only in a post-experimental survey. We conducted two additional monopoly treatments (one for weak and one for the strong prior), where we elicited beliefs about the quality of the advisors right after receiving advice in the first period.[14] As in the baseline treatment participants earnings were solely based on guessing correctly the number of the ball, taking into account the costs of buying advice. Beliefs about the quality of the advisor where not additionally incentivized.

[15] One of these had the weak prior, the other the strong prior. We ask about these beliefs without creating any incentives for giving an answer close to the Bayesian posteriors..

Procedure (incomplete). The experiment was conducted in the experimental economics laboratory of the Toulouse School of Economics. The four baseline treatments were implemented in June 2021 and May 2023, and the belief elicitation treatments in December 2024.

The experiment was programmed in oTree (Chen et al. 2016) and participants were recruited using a standard recruitment procedure.[16] After signing a consent form, each participant was randomly assigned a computer terminal.[17] Participants were aware that they were allowed to leave the experiment at any point, however, only a show-up fee would be paid in this case.

Participants in the baseline treatment were informed that there were two parts in the experiment (treatments). They received information about the second part only after having completed the first part. Participants in the belief elicitation treatments participated

[14]The exact wording we used is: *Compte tenu du message de votre conseiller ci-dessus, quelle est la probabilité qu'il soit parfait? Pour rappel, votre conseiller peut être parfait ou défectueux.* This sentence translates in English as *Given the message of your advisor above, what is the probability that it is perfect? We remind you that your advisor can either be perfect or defective.*

[15]We have found no effect of the order of treatments on the outcomes: demand for advice and the rate of following advice.

[16]The experimental protocol was approved by the TSE/IAST ethics committee under number xxxx

[17]We ran a total of xx sessions with on average xxx participants

in only one part (treatment).

In each part, participants first received some preliminary instructions and then read the game instructions. These were followed by a series of questions to verify their understanding. Participants then started with the task as described above. At the beginning of each round the participants were reminded of the number of green balls in the jar (see Figure B.4.1 in Appendix B.4).

Participants knew that at the end of the experiment, one round from each part would be randomly selected for final payout. To prevent negative earnings participants were endowed with 500 points at the beginning of each part. This amount could be used to pay for the purchasing of advice. When a round was selected for final payout, a correct guess in that round was rewarded with 500 points. The cost of advice in period two (if purchased) was deduced from the initial endowment. In rounds 1 to 10 the price of advice was fixed at 5 points. In rounds 11 to 15 the participants had to propose a price in a range from 0 to 250, and received advice if and only if this price was higher than a randomly drawn value from this interval. These points were converted into euros at a rate of 1 euro for 100 points.

At the end of the part, they received some post-experimental questions. Then the second part started.

Once participants were finished with both parts, they were informed about their total earnings which they could collect in an isolated room. Earnings were paid in cash.

The full instructions for each treatment are reproduced in Appendix B.2.As the participant advanced, the data on his choices were collected on the distant server, and then transferred into STATA data file.

Data. Table 1 gives an overview of the participant characteristics across the treatments conducted in May 2021 and July 2023. We had a total of 208 participants. Overall 48% of our participants were women and the average age was 21. Most participants were students and 58% reported to follow an economics-related field. In the survey after the experiment, participants were asked about their general tendency to ask for advice and about their risk aversion. Answers were given with integers between 1 to 10 (”Not at all” to ”Always”). Participants score an average of 6 on both scales, which implies that they have a tendency to ask for advice and that they have a tendency to take risk??

Participants did not consider the decisions to be particularly difficult to take and responded well to the control questions. The experiment lasted on average xxxh and the overall earnings from the experiment were 23 euros.

We observe no significant differences across the two sessions conducted in 2021 and 2023 and therefore report our findings for the pooled dataset in the next section.

4 Results.

This section is divided into three parts. The first part presents our findings regarding the demand for advice in period 2 (hereafter, *demand*). The following part describes guessing patterns. The last part reports the participants' beliefs about the quality of their advisors following the reception of period 1 advice (hereafter, *posteriors*).

	All sample	June 2021	May 2023
N. participants	208	122	86
N. participants (Prior 0.6)	107	61	46
N. participants (Prior 0.7)	101	61	40
Female	47.6%	41.8%	55.81%
Age (median)	21.00	21.00	21.00
Native French	82.84%	84.03%	81.18%
Education [years after highschool] (median)	3.00	3.00	2.50
Economics majors	57.86%	53.57%	64.29%
Color-blind	0.97%	0.83%	1.16%
Self-report: Advice seeking (0-10)	6.00	6.00	5.00
Self-report: Risk seeking (0-10)	6.00	6.00	6.00
Number of correct control questions (median)	19.00	20.00	17.00
Difficulty of decisions (0-10)	5.25	5.50	5.00
Time on instructions (seconds) (median)	115.25	119.25	109.50
Earnings (mean in euros)	23.16	23.36	22.88

Note: *Advice seeking* and *Risk seeking*, and *Difficulty for decisions* are self-declared propensity to ask for advice, propensity to take risks in real life, and difficulty in taking decisions during the experiment (Likert scale: 0 meaning "", 10 meaning ""). The table reports median values. Demographics are available for participants who provided the information as it was not mandatory. Education variables are manually coded and may be subject to noise.

Table 1: Overview of numbers of observations by treatment and time point and participant characteristics

4.1 Inelastic demand and its segments.

Recall that according to predictions 2 and 5, in the monopoly treatment with strong prior and in either duopoly treatment, demand for an advisor is supposed to be high when its period 1 advice confirms the common priors, that is, "green". We do not find such a pattern in our data.

Finding 1 (inelastic demand). *Demand for advice in period two does not depend on the content of advice(s) received in period one in any treatment.*

This holds for both demand at a fixed price (see Figure 2) as well as for the willingness to pay for advice elicited in rounds 11 to 15 (see Figure 3).

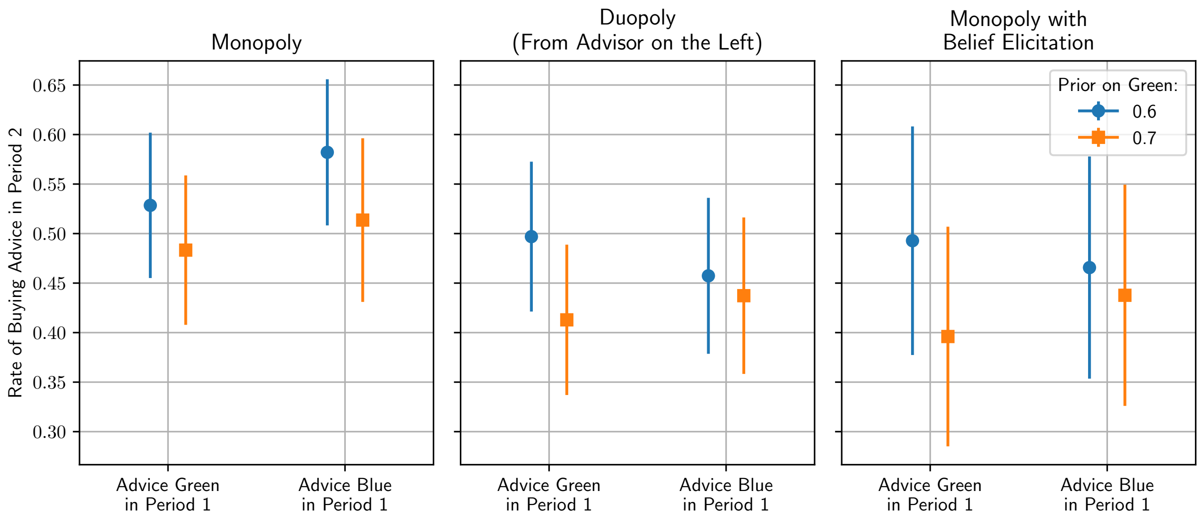

Figure 2: Inelastic demand for advice in different treatments. Figure 2 (left) depicts the average demand in baseline monopoly treatments. Figure 2 (middle) illustrates the average demand by the advisor situated on the left of the screen in duopoly treatments (figure 11 in Appendix B.6 shows a similar pattern for the advisor situated on the right of the screen). Figure 2 (right) illustrates the average demand in monopoly treatments with belief elicitation during the experiment.

Specifically, Figure 2 depicts, across treatments, the average rate at which participants purchased advice across rounds 1 to 10.[18] There is no statistically significant difference between the rate of purchasing advice following advice "green" or "blue" in either the weak- or strong prior treatments.

Figure 3 depicts the average "bid" for advice in rounds 11 to 15 in the baseline monopoly treatments.[19] We observe similar patterns.

More generally (and perhaps not surprisingly in the light of the sizable literature), the above demand pattern is very different from the theoretical predictions 2, 3, and 5. For

[18]Figure B.6.1. in Appendix B.6 depicts demand in rounds 6 to 10. We find no "learning".

[19]In the duopoly treatments we dispose of only few observations to make conclusive statements about the WTP for advice by a given advisor (see Figure B.6.2). Can we make a figure for treatments with belief elicitation?

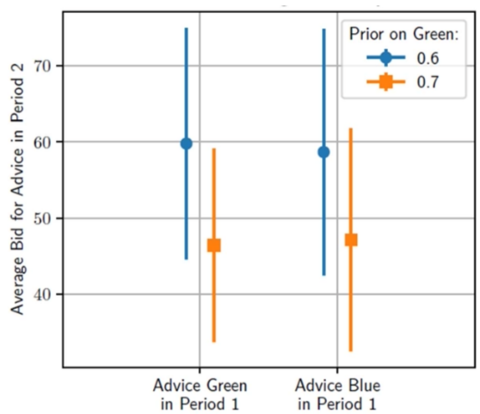

Figure 3: average "bid" for advice in period 2.

example, demand at a given price in the weak prior monopoly treatment (Figure 2, left) is almost twice as low as predicted, despite the price of advice being fixed to only 1% of the reward for a correct guess. The willingness to pay for advice in the baseline monopoly treatment (Figure 3) is "too low" following green advice and "too high" following blue advice, as compared to the values in prediction 3. However, one pattern comports nicely with rational behavior.

Remark 1. *Demand in the weak prior treatments is higher than in the strong prior treatments.*

Details on statistical significance. This difference may be explained by the fact that more informative priors provide a better guidance for independent guessing.[20]

Studying demand patterns in more detail, we find two sizable segments of inelastic demand.

[20]We find that indeed the participants guessing without any advice tend to guess "green" (see next section)

Table 2: Participants following rules (N) and (A) in all rounds from 1 to 10. Numbers in parenthesize indicate the number of participants in each treatment.

	Monopoly ()	Duopoly ()	Beliefs elicitation (99)
weak	A: 20% N: 14%	A: 9% N: 18%	A: ?% N: ?%
strong	A: 10% N: 17%	A: 0% N: 24%	A: ?% N: ?%

Finding 2 (two demand segments). *The two most prominent demand rules are: (A) "buy any available advice" and (N) "buy no advice".*

Indeed, between 24% to 34% of participants (depending on treatment), either *always* bought any available advice or *never* bought any advice (see Table 4).

Comparison of rows in Table 4 shows that a higher precision of common priors is associated with a greater share of participants fully adopting rule (N), reminiscent of Remark 1.

Figure 4 shows that the participants who do not fully adopt one of the rules (N) or (A) tend to alternate these rules. Specifically, most of the time in the duopoly treatments, participants either purchase no advice or both, even if advisors disagree in period one (which reveals that one of them is defective given our information structure).

The participants' choice between rules (N) or (A), and more generally their demand, is likely to be dictated by other considerations than the perceived quality of advice.

Remark 2. *The participants who do not fully adopt rule (N) or (A), tend to maintain the same demand as long as it is associated with successful guessing.*

Figure 5 illustrates this tendency.[21] Figure 5, left depicts the share of participants in the

[21] I suggest making three (at leat two) Figures akin to Figure 2. I suggest furthermore to exclude the

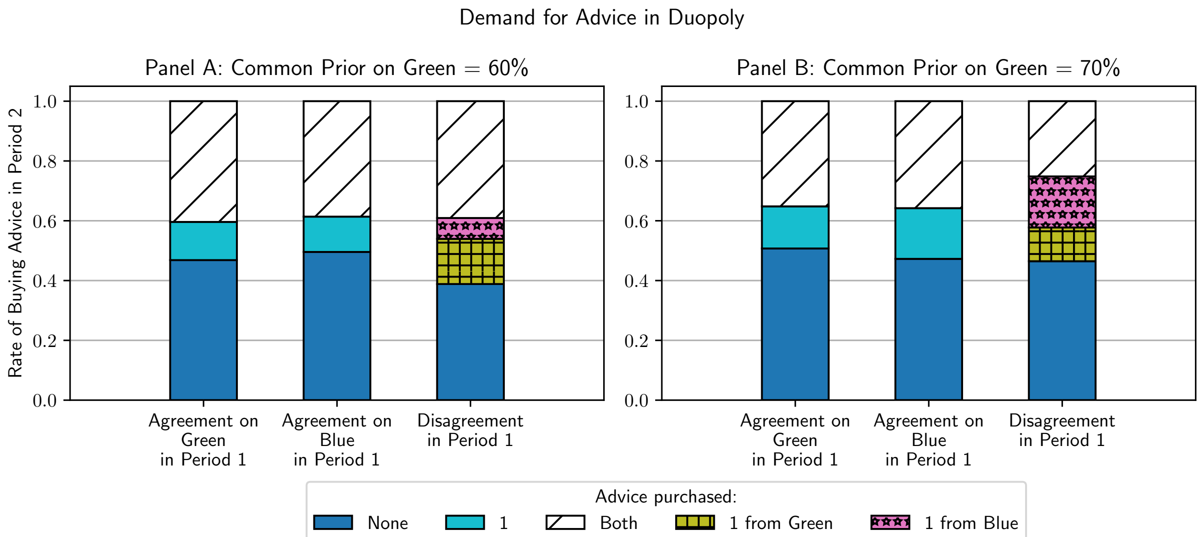

Figure 4: Demand in duopoly treatments.

baseline monopoly treatments keeping the same demand as in the previous round (excluding participants who always follow either rule A or N), depending on whether their period 2 guess in the previous round was correct or false. Figure 5, middle provides an equivalent figure for duopoly treatments, and Figure 5, left refers to treatments with belief elicitation.

Note that such "win-stay loose-shift" demand rule cannot help improve guessing because the quality of advisors in different rounds is independent.

4.2 Influential advice.

While the above patterns in demand for advice are very different from theoretical predictions 2, 3 and 5, the patterns of advice usage are broadly consistent with theoretical predictions 1 and 4.

participants fully adopting rules (A) and (N) in monopoly treatments this would lake figure in the left. Then, the same for duopoly treatments - perhaps diff set of excluded participants (middle). FInally, for belif elicitation treatments. Finally, I suggest using definition "win" for duccessful guess in period 2 and loos for errouneous guess (so that we conver the whole space of possibilities).

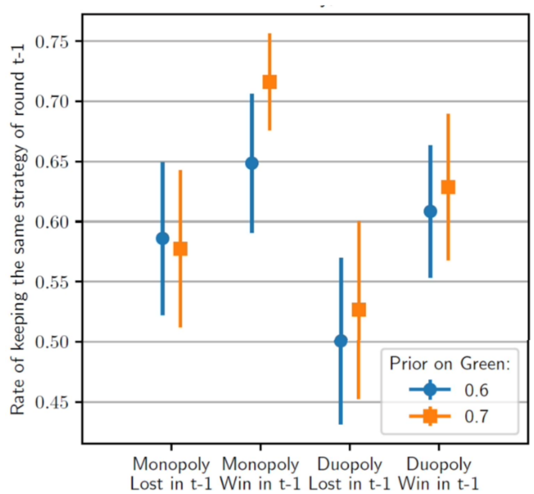

Figure 5: Win-stay loose-shift demand pattern (baseline treatments, periods 1 to 10)

Set of findings 3 (guessing upon advice and priors). *(i) The participants tend to follow advises in monopoly treatments, and (even more so) agreeing advises in duopoly treatments. (ii) These tendencies are stronger when advice(s) confirms the common priors, and this difference increases in the precision of common priors. (iii) Disagreement of advisors in duopoly treatments induces the subjects to rely on the common priors. (iv) The subjects guessing independently rely on common priors.*

Figure 6 illustrates our set of findings 3(i) to 3(iii), focusing on advice provided in period 1.

Similar patterns arise for guessing in period 2 (see Fig xxx in Appendix B.6), with a stronger impact of advice on guessing (conditional on purchasing).

Remark 3 (cost of advice and its impact). *The rate of following costly advice in period 2 (conditional on buying) is higher than that provided for free in period 1.*

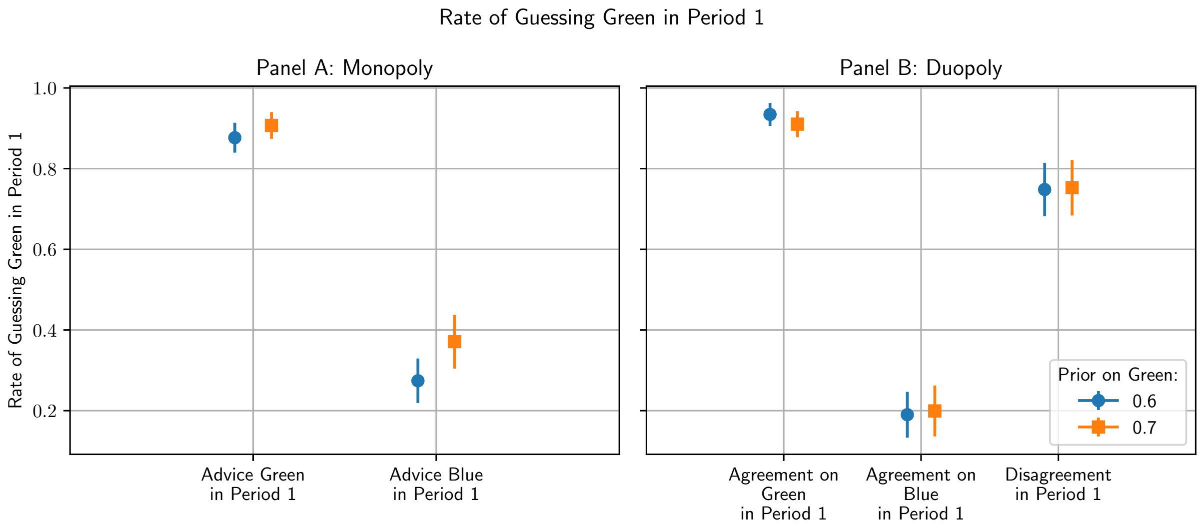

Figure 6: Guessing patterns in period 1 (all rounds)

In weak prior treatment, this difference amounts to: 6.2% for advice “green” and 3.8% for advice “blue”, while in strong prior treatment it amounts to 3.8% and 7.8%, respectively. These differences are statistically significant (??). Remark 3 is consistent with the literature reporting that people are more inclined to rely on information which was costly to get (refs).

In absence of advice, from X to Y percent of participants depending on the treatment guess “green” when buying no advice in period 2, relying thereby on common priors.

More generally, set of findings 3 is consistent with participants using the common priors and advice to update their beliefs about the correct guess (see Benjamin 2019 and references therein).

An alternative interpretation of finding 3, (ii) is participants’ viewing advice “green” as a signal of quality and thereby following this advice at a higher rate. However, this interpretation is not aligns with finding 1 and it is not confirmed by our findings reported in the following section. Furthermore, the content of advice received in period 1 does not

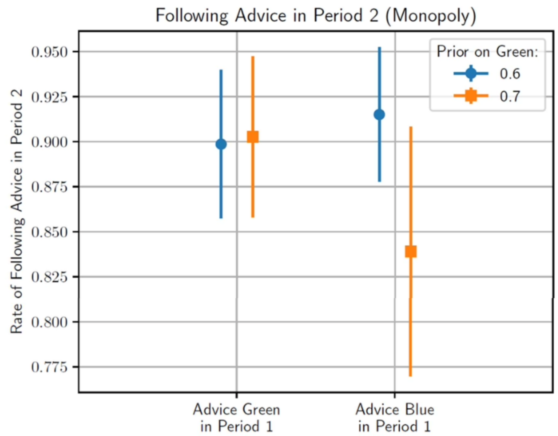

Figure 7: Rate of following advice in period 2 depending on the contents of advice in period 1 (baseline monopoly treatments)

affect the rate of following advice in period 2.

Remark 4 (influence independent on contents of previous advice). *The rate of following advice in period two (conditional on buying) does not depend on the contents of advice received in period 1.*

4.3 Beliefs on the quality of advisors.

While finding 3 is consistent with participants' using advice to revise their prior beliefs about correct guesses, findings 1, 2 and 4 suggest that they may not revise their beliefs on the qualities of advice itself upon the content of the advice received in period 1. Indeed, we find, using the data from treatments with beliefs elicitation during the experiment that the participants' posteriors[22] are similar to their prior beliefs.

[22]Recall that "posteriors" refers to beliefs about the quality of their advisors following the reception of period 1 advice.

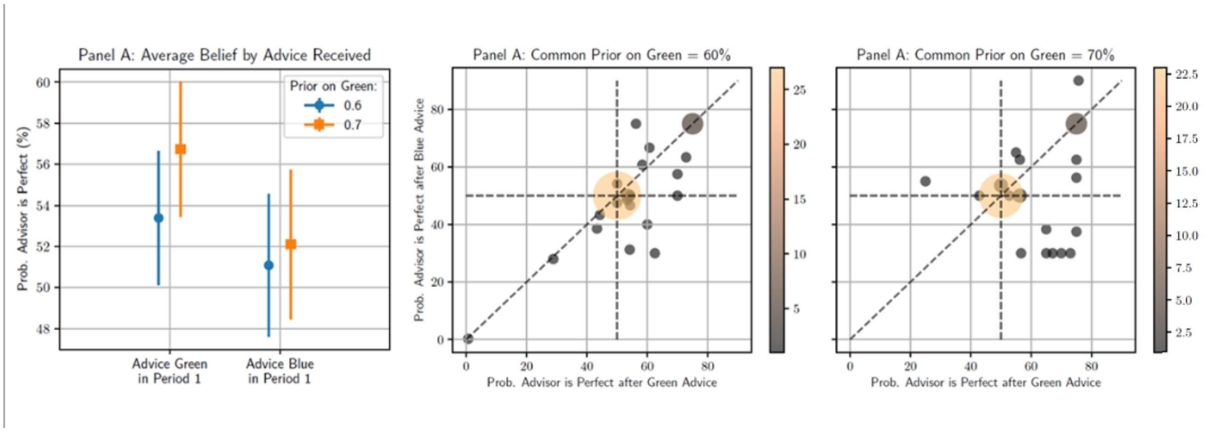

Figure 8: beliefs on the quality of advisor depending on the content of advice received in period 1.

Finding 4 (rigid beliefs on quality of advisors). *Participants tend to maintain their prior beliefs about the quality of their advisors regardless of the contents of advice received in period 1.*

While the average posteriors in treatments with belief elicitation during the experiment are generally higher when period one advice is "green"(Figure 8, left) this difference is minor (less than 5%) and it is significant only in the strong prior treatment.

Figure 8, middle and right show that about half of participants do not update their beliefs on the quality of their advisors based on the content of period one advice *at all*. Indeed, as many as X% of participants in weak prior treatment and Y% in strong prior treatment report posteriors of 50%.[23]

Furthermore, the kernel density estimation of the distribution of reported posteriors in Figure 9 brings no evidence of the distribution being shifted in either the strong- or weak

[23] We observe similar reports in the post-experimental survey from the baseline treatments, in which we asked the participants to report their beliefs about the quality of an advisor who advised to guess "green": X% of participants in the weak prior treatment and Y% in strong prior treatment answered 50%.

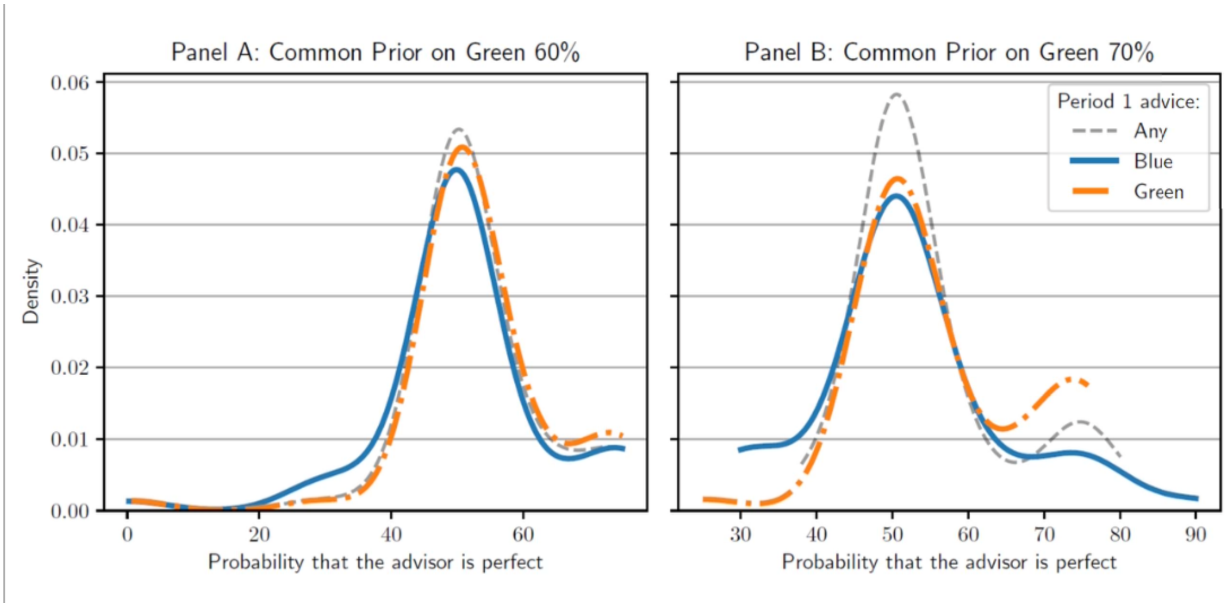

Figure 9: distribution of posteriors

prior treatment.

Naturally, there is some heterogeneity in belief updating among the participants. Indeed, xxx% of participants update their beliefs positively upon reception of confirmatory advice (notice the second largest "hill" in Figure 10 middle and right). Their demand for advice differs from demand by participants with rigid beliefs, as illustrated in Figure 10. However, this difference is not statistically significant.

5 Conclusion

Our findings suggest that an advice on correct guess induces the participants to update their beliefs on that guess while largely foregoing inferences on the quality of advice itself.

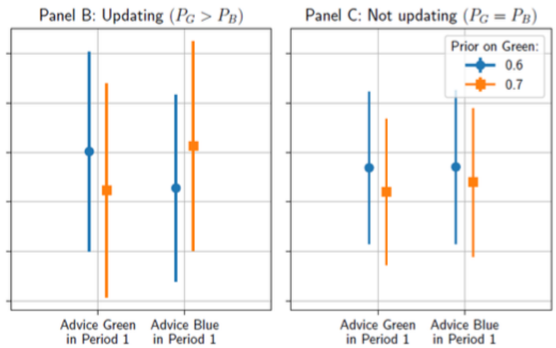

Figure 10: Beliefs and demand.

Similarly, news on controversial issues may change beliefs on these issues, but not on the quality of the sponsoring media outlet, creating a potential for biased or even fake reports.

We have studied inferences on only one aspect of advice contents, namely its alignment with common priors. Future experimental research may consider other on sources of inferences about the quality of information source which can be deduced from information contents.

References

Ambuehl, S. and S. Li (2018): "Belief updating and the demand for information," *Games and Economic Behavior*, 109, 21–39.

Benjamin, D. J. (2019): "Errors in probabilistic reasoning and judgment biases," *Handbook of Behavioral Economics: Applications and Foundations 1*, 2, 69–186.

Camara, F. and N. Dupuis (2021): "Structural Estimation of Expert Strategic Bias: The Case of Movie Critics," *Working Peper*.

Camerer, C. (1998): "Bounded rationality in individual decision making," *Experimental economics*, 1, 163–183.

Charness, G., R. Oprea, and S. Yuksel (2021): "How do People Choose Between Biased Information Sources? Evidence from a Laboratory Experiment," *Journal of the European Economic Association*, 19, 1656–1691.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Chopra, F., I. Haaland, and C. Roth (2024): "The demand for news: Accuracy concerns versus belief confirmation motives," *The Economic Journal*, 134, 1806–1834.

Liang, Y. (2024): "Learning from unknown information sources," *Management Science*.

Meloso, D., S. Nunnari, and M. Ottaviani (2023): "Looking into Crystal Balls: A Laboratory Experiment on Reputational Cheap Talk," *Management Science*, 69, 4973–5693.

Montanari, G. and S. Nunnari (2023): "Audi alteram partem: An experiment on selective exposure to information," .

Ottaviani, M. and P. N. Sørensen (2006a): "Reputational cheap talk," *The Rand journal of economics*, 37, 155–175.

——— (2006b): "The strategy of professional forecasting," *Journal of Financial Economics*, 81, 441–466.

Schoar, A. and Y. Sun (2024): "Financial Advice and Investor Beliefs: Experimental Evidence on Active vs. Passive Strategies," *NBER Working Paper*.

A Appendix

A.1 Proof of statements in Section "rational behavior".

A.1.1 Proof of prediction 1.

If $q^l = 1$, advice l is correct, if $q^l = 0$, advice l is correct with probability $\frac{1}{2}$. Hence, the probability of making correct guess by following advice l is equal to

$$\Pr\left(q^l = 1\right) + \tfrac{1}{2}\Pr\left(q^l = 0\right) = \tfrac{1}{2}\left(1 + \Pr\left(q^l = 1\right)\right).$$

The best alternative is guessing "G", which is correct with probability p. Therefore, following the advice l is optimal iff its expected quality is sufficiently high as compared to the precision of the common priors p:

$$\Pr\left(q^l = 1\right) > 2p - 1. \tag{3}$$

Given that $\Pr\left(q^l = 1\right) = \frac{1}{2}$, this condition is met for any $p < 0.75$, hence for either p in set $\{0.6, 0.7\}$.

An alternative argument is that by the upper limitation in set (2) advice "B" is likely to be correct:

$$\Pr\left(x_1 = G \mid a_1^l = B\right) = \tfrac{p}{3(1-p)+p} < \Pr\left(x = B \mid a_1^l = B\right) = \tfrac{3(1-p)}{3(1-p)+p} \tag{4}$$

iff $p < 0.75$, hence for either p in set $\{0.6, 0.7\}$.

A.1.2 Proof of prediction 2.

It is optimal to buy and follow the advice iff its posterior quality $\Pr\left(q^l = 1 \mid a_1^l\right)$ is sufficiently high as compared to the precision of common priors (and the price of advice):

$$\Pr\left(q^l = 1 \mid a_1^l\right) + \tfrac{1}{2}\Pr\left(q^l = 0 \mid a_1^l\right) - \varepsilon > p, \text{ where } \varepsilon = 0.01 \tag{5}$$

which is equivalent to

$$\Pr\left(q^l = 1 \mid a_1^l\right) > 2\left(p + \varepsilon\right) - 1. \tag{6}$$

By Bayes rule,

$$\Pr\left(q^l = 1 \mid a_1^l = G\right) = \tfrac{2p}{2p+1}, \tag{7}$$

$$\Pr\left(q^l = 1 \mid a_1^l = B\right) = \tfrac{2(1-p)}{2(1-p)+1}. \tag{8}$$

By equation (7), when $a_1^l = G$ inequality (6) is equivalent to

$$4p^2 - 2\left(1 - 2\varepsilon\right)p - \left(1 - 2\varepsilon\right) < 0, \text{ where } \varepsilon = 0.01$$

which is true iff $p < 0.79$. Hence, inequality (6) holds for either p in set $\{0.6, 0.7\}$

By equation (8), when $a_1^l = B$ inequality (6) is equivalent to

$$4p^2 - 2p\left(5 - 2\varepsilon\right) + \left(5 - 6\varepsilon\right) < 0, \text{ where } \varepsilon = 0.01$$

which is true iff $p < 0.68$. Hence, inequality (6) holds for either $p = 0.6$ but not for $p = 0.7$.

A.1.3 Willingness to pay for monopoly advice.

We compute WTP as the difference between the posterior probability to receive correct advice and the prior probability p scaled by V.

Period 2 advice is correct with probability $\Pr(q = 1 \mid a_1) + \frac{1}{2}\Pr(q = 0 \mid a_1)$. The best alternative to buying- and following advice is guessing "green". Such guess is correct with probability p. Therefore, the decision maker's willingness to pay for advice is equal to

$$\left(\left[\Pr(q = 1 \mid a_1) + \tfrac{1}{2}\Pr(q = 0 \mid a_1)\right] - p\right) V, \text{ where } V = 500. \tag{9}$$

Straightforward calculus using equations (9), (7) and (8) gives the predicted figures.

A.1.4 Proof of prediction 4.

In period 1 the decision maker receives two advises. By inequality (4), it is better to follow advice by one advice (either of the two, say l) than to guess "G". Trivially, advice r does not alter decision making if the advices agree. Suppose they disagree: $a_1^i = B$ while $a_1^{-i} = G$, where $i \in \{l, r\}$ and $-i = \{l, r\} \setminus \{i\}$. The decision maker can explain it in two ways: (1) With probability $\frac{1}{2}p\Pr(q^i = 0, q^{-i} = 1)$, the state is "$B$", advce i is perfect while advice $-i$ is defective and false; (2) With probability $\frac{1}{2}(1-p)\Pr(q^i = 1, q^{-i} = 0)$, the state is "$G$", advice $-i$ is perfect while advice i is defective and false. *A priori* the advices have the same qualities, that is,

$$\Pr\left(q^i = 0, q^{-i} = 1\right) = \Pr\left(q^i = 1, q^{-i} = 0\right).$$

At the same time, it is more likely that the state is "G": $p > \frac{1}{2}$. Hence, explanation (2) is more plausible than (1) and guess "G" is the best. This paragraph shows that it is optimal to guess "G", iff at least one advice reccomends it. Notice that additional advice improves the probability of making correct guess by:

$$\tfrac{1}{2}\left(p \Pr\left(q^i = 0, q^{-i} = 1\right) - (1-p)\Pr\left(q^i = 1, q^{-i} = 0\right)\right) > 0. \tag{10}$$

A.1.5 Proof of prediction 5.

Step 1. The optimal DMR[24] in period 2 is among the following three: (i) guess "G" ; (ii) buy- and follow advice

$$i^* = \arg\max_{i=l,r} \Pr\left(q^i = 1 \mid a_1^l, a_1^r\right), \tag{11}$$

(iii) buy both advices and guess: "G" if at least one advice reccomends it, and "B" otherwise.

By inequality (6), DMR (ii) is more efficient than DMR (i) iff

$$\Pr\left(q^{i^*} = 1 \mid a_1^l, a_1^r\right) > 2\left(p + \varepsilon\right) - 1. \tag{12}$$

[24]Recall that abbreviation DMR is used for "decision making rule".

By inequality (6), DMR (iii) is more efficient than DMR (ii) iff

$$\tfrac{1}{2}\left(p\Pr\left(q^{i}=0,q^{-i}=1\right)-(1-p)\Pr\left(q^{i}=1,q^{-i}=0\right)\right)>\varepsilon. \tag{13}$$

Step 2 shows that if $a_1^l \neq a_1^r$, then it is optimal in period 2 to buy- and follow the advice which has recommeded "G" in period 1 if $p < 0.98$, and guess "G" without buying advice otherwise. Formally,

$$i^{*}=\left\{i\in\{l,r\}\,|\,a_1^i=G\right\}$$

where i^* is defined by equation (11), and the optimal DMR is: (ii) if $p < 0.98$, and (i) otherwise.

Let $a_1^i = G$ and $a_1^{-i} = B$ where $i \in \{l,r\}$ and $-i = \{l,r\} \setminus \{i\}$. By Bayes rule, the decision maker infers that the coin has tailed, that is $z_1 = T$, and the qualities of advices l and r are different:

$$\Pr\left(q^{l}=q^{r}=0\mid a_1^l\neq a_1^r\right)=\Pr\left(q^{l}=q^{r}=1\mid a_1^l\neq a_1^r\right)=0. \tag{14}$$

With probability p, advice i is perfect and advice $-i$ is defective:

$$\Pr\left(q^{i}=1\mid a_1^i=G,a_1^{-i}=B\right)=\Pr\left(q^{i}=1,q^{-i}=0\mid a_1^i=G,a_1^{-i}=B\right)=p, \tag{15}$$

$$\Pr\left(q^{-i}=1\mid a_1^i=G,a_1^{-i}=B\right)=\Pr\left(q^{i}=0,q^{-i}=1\mid a_1^i=G,a_1^{-i}=B\right)=1-p. \tag{16}$$

By equations (14) to (16), DMR (ii) dominates DMR (iii) for any p:

$$\tfrac{1+p}{2} - 2\varepsilon < \tfrac{1+p}{2} - \varepsilon \text{ for any } p.$$

Furthermore, DMR (ii) dominates DMR (i) for any $p < 1 - 2\varepsilon = 0.98$:

$$\tfrac{1+p}{2} - \varepsilon > p \text{ for any } p < 1 - 2\varepsilon = 0.98.$$

Step 3 shows that if $a_1^l = a_1^r = G$, then i^* defined by equation (11) is either element of set $\{l, r\}$ and the optimal DMR in period 2 is: (ii) if $p < 0.5579$, (iii) if $0.5579 \leqslant p < 0.8323$, and (i) otherwise. Indeed, by Bayes rule, advices l and r are equally likely to be perfect:

$$\Pr\left(q^l = q^r = 1 \mid a_1^l = a_1^r = G\right) = \tfrac{2p}{1+4p}, \tag{17}$$

$$\Pr\left(q^l = q^r = 0 \mid a_1^l = a_1^r = G\right) = \tfrac{1}{1+4p}, \tag{18}$$

$$\Pr\left(q^i = 1,\ q^{-i} = 0 \mid a_1^l = a_1^r = G\right) = \tfrac{p}{1+4p}, \text{ where} \tag{19}$$

$i \in \{l, r\}$ and $-i = \{l, r\} \setminus \{i\}$. By equations (17) and (19),

$$\Pr\left(q^i = 1 \mid a_1^l = a_1^r = G\right) = \tfrac{3p}{1+4p}, \text{ where } i \in \{l, r\}. \tag{20}$$

By equations (17), (18) and (20), inequality (13) which is necessary and sufficient for

DMR (iii) to dominate DMR (ii) is equivalent to inequality

$$(2p-1)\Pr\left(q^{i}=1,q^{-i}=0 \mid a_1^l=a_1^l=G\right)>2\varepsilon,$$

$$\text{which is equivalent to } 2p^2-p\,(1+8\varepsilon)-2\varepsilon>0,$$

which is true iff $p>0.5579$.

It remains to compare DMRs (iii) and (i). DMR (iii) improuves guessing in state "B", unless at least one advice is defective and the fair coin tails ($z_2=T$) and creates an error in state "G" when both advices are defective ($q^l=q^r=0$) and the fair coin tails ($z_2=T$). Hence, it dominates DMR (i) iff

$$\begin{aligned}(1-p)\left(\Pr\left(q^l=q^r=1 \mid a_1^l,a_1^r\right)+\tfrac{1}{2}\left(1-\Pr\left(q^l=q^r=1 \mid a_1^l,a_1^r\right)\right)\right)\\ >p\tfrac{1}{2}\Pr\left(q^l=q^r=0 \mid a_1^l,a_1^r\right)+2\varepsilon,\text{ where } \varepsilon=0.01.\end{aligned} \tag{21}$$

By equations (17) and (18), inequality (21) is equivalent to

$$6p^2-4p\,(1-4\varepsilon)-(1-4\varepsilon)<0,$$

which is true iff $p<0.8323$.

Hence, if $0.5579\leqslant p<0.8323$ DMR (iii) is optimal. We now describe the optimal DMR for probability p outside interval $[0.5579,0.8323)$. To this goal, we compare the efficiency of DMRs (i) and (ii). By equation (20), DMR (ii) dominates DMR (i) iff

$$8p^2 - p(5-8\varepsilon) - (1-2\varepsilon) < 0,$$

which is true iff $p < 0.7733$.

Step 4 shows that if $a_1^l = a_1^r = B$, then i^* defined by equation (11) is either element of set $\{l, r\}$ and the optimal DMR in period 2 is: (ii) if $p < 0.5628$, (iii) if $0.5628 \leqslant p < 0.7096$, and (i) otherwise. Indeed, by Bayes rule,

$$\Pr\left(q^l = q^r = 1 \mid a_1^l = a_1^r = B\right) = \tfrac{2(1-p)}{5-4p}, \tag{22}$$

$$\Pr\left(q^l = q^r = 0 \mid a_1^l = a_1^r = B\right) = \tfrac{1}{5-4p}, \tag{23}$$

$$\Pr\left(q^i = 1,\ q^{-i} = 0 \mid a_1^l = a_1^r = B\right) = \tfrac{1-p}{5-4p}, \text{ where} \tag{24}$$

$i \in \{l, r\}$ and $-i = \{l, r\} \setminus \{i\}$. By equations (22) and (24),

$$\Pr\left(q^i = 1 \mid a_1^l = a_1^r = G\right) = \tfrac{3(1-p)}{5-4p}, \text{ where } i \in \{l, r\} \tag{25}$$

By equations (22), (23) and (25), DMR (iii) dominates DMR (ii) iff

$$2p^2 - p(3+8\varepsilon) + (1+10\varepsilon) < 0,$$

which is equivalent to $0.5629 < p < 0.9771$.

By equations (22) and (23)), DMR (iii) dominates DMR (i) iff

$$6p^2 - p\,(14 - 16\varepsilon) + (7 - 20\varepsilon) > 0,$$

which is true iff $p < 0.7097$.

By equation (25), DMR (ii) dominates DMR (i) the expected payoff (**??**) iff

$$8p^2 - p\,(17 - 8\varepsilon) + (8 - 10\varepsilon) > 0,$$

which is true iff $p < 0.6958$.

A.6. Willingness to pay for competing advices.

Prediction A.1. (WTP for competing advices). *The following table summarizes WTP depending on advises received in period 1. "Best" refers to advice which is the most likely to be perfect,"both" refers to both advises.*

	$p = 0.6$		$p = 0.7$	
WTP for:	*best*	*both*	*best*	*both*
period 1 advices differ	100	100	0.15	0.15
period 1 advices are"G"	82.35	91.18	38.16	56.58
period 1 advices are"B"	65.38	73.68	2.23	15.91

The decision maker's willingness to pay for one advice by advisor with the highest reputation in period 2 is equal to the difference between the probability of this advice being correct and the prior probability p scaled by V:

$$\left(\left[\Pr\left(q^{i^*}=1 \mid a_1\right)+\tfrac{1}{2}\left(1-\Pr\left(q^{i^*}=1 \mid a_1\right)\right)\right]-p\right)V, \tag{26}$$

where $V = 500$.

Similarly, the decision maker's willingness to pay for two advises is equal to the difference between the probability of making correct guess based on these advises (that is, "green" if at least one advisor reccomends so and "blue" otherwise) and the prior probability p scaled by V

$$\begin{gathered}(p\left(1-\tfrac{1}{2}\Pr\left(q^l=q^r=0 \mid a_1^l, a_1^r\right)\right)+(1-p)\left(\Pr\left(q^l=q^r=1 \mid a_1^l, a_1^r\right)+\right.\\ \left.+\tfrac{1}{2}\left(1-\Pr\left(q^l=q^r=1 \mid a_1^l, a_1^r\right)\right)\right)-p)V.\end{gathered} \tag{27}$$

Straightforward calculus using equations (26), (27), (15), (16), (17), (18), (20), (22), (23) and (25) gives figures in Table 1.

B Experimental procedure and data.

B.1 Consent form.

Formulaire de consentement et renseignements concernant la confidentialité des données.

L'Institut d'Etudes Avancées de Toulouse et l'École d'Economie de Toulouse (1 Esplanade de l'Université, 31080 Toulouse, Cedex 06 France) réalise aujourd'hui une expérience laboratoire au Laboratoire d'Economie Expérimentale de TSE. Vous avez été invité car vous êtes inscrit dans notre système de recrutement. Les procédures de l'expérience vous seront expliquées avant le début de l'expérience.

Au cours de l'expérience, on vous demandera d'accomplir certaines tâches et vos réponses seront enregistrées dans notre système informatique. Les informations enregistrées pendant l'expérience ne permettront pas de faire de conclusion sur la participation ou le comportement de personnes individuelles. L'analyse des données et la présentation des résultats de cette expérience auront lieu exclusivement sous forme anonyme. Les données anonymes seront archivées et pourraient être mises à la disposition d'autres chercheurs à des fins de recherche.

Il n'y aura pas de lien entre les données générées pendant l'expérience et les données dans le système de recrutement. Les reçus remplis pendant le paiement seront toujours conservés séparément. Votre participation aujourd'hui est entièrement volontaire. Si vous ne participez pas, il n'y aura aucun inconvénient pour vous. Cependant, notez que dans ce cas vos gains de l'expérience seront ajustés de manière appropriée. Vous pouvez à tout moment annuler votre participation à l'expérience.

Pour votre participation et l'utilisation des données, nous vous demandons votre consentement. Le consentement peut être révoqué à tout moment, par exemple par courrier à : jacopo.bregolin@gmail.com. Ce consentement est la base juridique de toute utilisation de données.

Vous pouvez trouver une copie de cette information au laboratoire expérimental.

J'ai lu les informations sur la confidentialité des données et suis d'accord avec la participation à l'expérience et l'utilisation des données connectées décrites ci-dessus:

☐ Oui ☐ Non

________________________ ______________________________

Date Signature

Sans consentement, vous ne pouvez pas participer à l'expérience aujourd'hui. Veuillez en informer la personne en charge de l'expérience dans ce cas. Merci.

B.2 Instructions (french).

B.3 General Instructions. Treatment variations between treatment 0.6 and 0.7 are marked by brackets

Merci de participer à notre expérience!

Dans cette expérience, vous pouvez gagner de l'argent en fonction de vos décisions. Nous vous demandons donc de lire attentivement les instructions. Vos gains pendant l'expérience seront calculés en points. Vos points seront échangé en euros à la fin de l'expérience au taux de change suivant:

100 points = 1 euro

Dans l'expérience d'aujourd'hui, on vous demandera de deviner à plusieures reprises la couleur d'une balle tirée d'un pot. Vous gagnerez des points en fonction de vos réponses.

Vous n'êtes pas autorisé à communiquer avec d'autres participants pendant l'expérience. Nous vous demandons également d'éteindre vos téléphones mobiles maintenant. Si vous avez des questions, veuillez lever la main et quelqu'un viendra à votre bureau pour vous répondre. Ce que vous allez faire.

Vous allez jouer une sessions de 15 tours. L'un de ces tours sera choisi pour déterminer votre compensation à la fin de la session. Chaque tour se compose de deux périodes. Dans chaque période, une balle est tirée de- (et retournée dans) un pot avec 10 balles, dont (6)(7) sont vertes et (4)(3) sont bleues, comme sera montré sur l'écran de votre ordinateur. Vous devinerez la couleur de cette balle. Vous gagnez 500 points si votre réponse est correcte.

B.3.1 Monopoly.

La première période Avant de donner votre réponse dans la première période, vous obtiendrez automatiquement un conseil de votre conseiller informatisé. Autrement dit, votre conseiller vous montrera soit un cercle bleu, suggérant de répondre bleu ou vert, suggérant de répondre vert . Votre conseiller est soit parfait soit défectueux, avec la même probabilité. La qualité de votre conseiller est déterminée à nouveau au début de chaque tour et elle reste la même pendant les deux périodes du tour. Vous ne serez pas informé de la qualité de votre conseiller. Si le conseiller est parfait, son conseil est correct (c.-à-d. qu'il montre un cercle vert si la balle sélectionnée est verte, et elle montre un cercle bleu si la balle sélectionnée est bleue).

Si le conseiller est défectueux, le contenu de son conseil dépend du résultat d'un jeu à pile ou face. Après avoir reçu le conseil, vous devez deviner si la couleur de la balle sélectionnée est bleue ou verte. Vous apprendrez si votre estimation est correcte ou non à la fin du tour. Vous gagnez 500 points si votre supposition est correcte. Vous apprendrez également la qualité de votre conseiller à la fin de chaque tour.

La seconde période Votre conseiller est le même que dans la première période. Vous ne recevez pas son conseil automatiquement. Vous pouvez le demander en cliquant sur le bouton conseil . Le conseil est coûteux (le coût est détaillé ci-dessous). Alternativement, vous pouvez répondre sans conseil. Une fois que votre réponse est soumise, vous apprendrez si elle est correcte ou fausse. Vous gagnez 500 points si votre réponse est correcte. En outre, à la fin de chaque tour, vous apprendrez également si votre conseiller était parfait ou défectueux.

Le coût du conseil au cours de la deuxième période. *Tours 1 à 10:* Si vous demandez conseil en période 2, vous payez 5 points.

Tours 11 à 15: Si vous demandez conseil en période 2, vous devez proposer une somme entre 0 et 250 (en points) que vous êtes prêt à payer pour le conseil. L'expérience est conçue de telle sorte qu'il est préférable pour vous de proposer le prix le plus élevé que vous êtes prêt à payer. En effet, nous allons tirer au sort un prix seuil X pour le conseil. Si la somme que vous proposez de payer est supérieur à ce prix seuil, vous recevrez alors le conseil demandé et vous payez le prix seuil X (pas la somme que vous avez proposée). Sinon, vous devrez deviner sans conseil.

B.3.2 Duopoly

Figure B.2.1 for nested info structure:

B.4 Understanding.

The questions verifying understanding of the instructions had multiple-choice format. There were X questions in the monopoly treatments and Y questions in the duopoly treatments. These questions are presented below. If a participant did not select the correct answer to a question, (s)he was notified that her answer was false, and the correct answer and the corresponding part of the instructions were displayed on the screen.

B.4.1 Monopoly

Rounds 1 to 10 text (list of questions)

Figures on performance

Rounds 11 to 15 Figure (part of screenoshot) on which they were spending lots of time. Distribution of time (proxy for skills).

B.4.2 Duopoly

Rounds 1 to 10 text (list of quetsions)

Figures on performance

Rounds 11 to 15 Figure (part of screenoshot) on which they were spending lots of time. Distribution of time (proxy for skills).

B.5 Details on events in a round of experiment.

Screenshot "recall" Figure B.4.1.

Screeshot "feedback" Figure B.4.2.

Screenshot for Fig B.4.3

Screenshot for belief elicitation Fig B.4.5

B.5. Post-experimental surevey.

Fig B.5.1: lottery

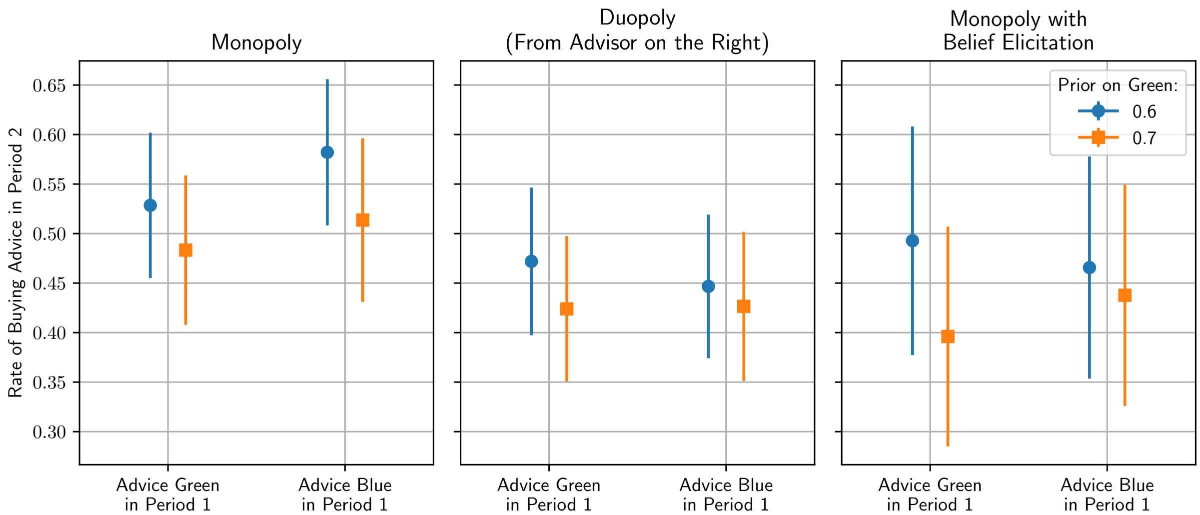

Figure 11: Inelastic demand for advice in different treatments. Figure 11 (left) depicts the average demand in baseline monopoly treatments. Figure 11 (middle) illustrates the average demand by the advisor situated on the right of the screen in duopoly treatments. Figure 11 (right) illustrates the average demand in monopoly treatments with belief elicitation during the experiment.

Data on personal characteristics

B.6 Additional findings.

Figure B.6.1.: equivalent of Figure 2 for rounds 6 to 10.

Figure B.6.3.: WTP by advisors located on different sides of the screen (??)