Using Other People's Opinions: An Experimental Study

Rudiger, Jesper

European University Institute

28 November 2013

Using Other People's Opinions: An Experimental Study*

Jesper Rudiger†
European University Institute

November 2013

Abstract

Expert opinions are often biased. To test how such bias affects the propensity to use opinions, we set up an experiment where subjects estimate the probability of an event that depends on (i) the subject's type, which is observable, and (ii) the unobserved state of the world. Before making their estimate, one group of subjects, the clients, observe the opinion (estimate) of another subject, the expert. The expert has private information about the state, but he may be of a different type than the clients, and therefore biased. Bias is observable and easily corrected. In spite of this, we find that clients' propensity to use expert opinions is decreasing in the size of the expert's bias. This aversion to use the opinions of biased experts is not explained by computational concerns, ex-post expert informativeness or reluctance to move away from the prior.

Keywords: Experiments; Probability Estimation; Biased Opinions; Naive Advice
JEL Codes: C91; D81; D82

*I would like to thank Praveen Kujal and my adviser Antonio Cabrales for valuable advice. Furthermore, I have received useful comments from Gary Charness, Piero Gottardi, Ángel Hernando-Veciana, Diego Moreno, Axel Ockenfels, Andrew Schotter and Matthias Sutter, as well as seminar audiences at Toulouse School of Economics BEE Seminar and the Alhambra Experimental Workshop in Granada.

†Department of Economics, European University Institute, Via delle Fontanelle 10, I-50014 San Domenico, Italy. *Telephone*: +39 055 4685 698. *E-mail*: jesper.rudiger@eui.eu

1 Introduction

We often base our decisions on the opinions of other people, although we may consider these opinions to be biased. Consider the following example. You are looking to buy a particular car, and before making the decision, you want to estimate its maintenance costs. These depend both on the quality of the car and how it is treated. As it turns out, a friend of yours already owns exactly the same model. He tells you that in his opinion, the car requires a great deal of maintenance. But you are aware that he, unlike you, is not very careful with his car. So whereas your friend's opinion contains information about the quality of the car, you also need to adjust for his carelessness when you estimate what your own maintenance costs would be. In this example, the friend is better informed than you, but has no vested interest in the decision made. Thus, the situation differs from the games of strategic information transmission that have been extensively studied in experiments.[1] Neither does it correspond to advice-giving, which is typically investigated in a sequential setting where subjects pass on advice about which actions to take.[2] Rather, the friend in our example states an opinion about the frequency with which his car breaks down, but you are interested in the frequency with which *your* car would break down. The friend's opinion is thus biased, in the sense that it concerns a variable that is different – but correlated – to the variable of interest.

In the present paper we investigate the following question: how do we use biased opinions in decision making? To this end we set up an experiment in which subjects estimate the probability of drawing a black ball from a cage that contains only black and white balls. Part of the cage is known and depends on the subject's *type* (the car owner's carefulness in the above example). The other part is unknown and depends on

[1] For instance, communication increases payoffs in coordination games (Cooper et al., 1989, 1992; Crawford, 1998; Blume and Ortmann, 2007). When subjects have partially aligned interests, there is often excessive information revelation (Cai and Wang, 2006) and aversion to lying (Gneezy, 2005; Sánchez-Pagés and Vorsatz, 2007; Hurkens and Kartik, 2009).

[2] In such intergenerational settings, advice is found to: increase coordination (Schotter and Sopher, 2003); facilitate backward induction in trust games (Schotter and Sopher, 2006); induce higher contributions in public goods games (Chaudhuri et al., 2006); cause lower offers and higher rejection rates in ultimatum games (Schotter and Sopher, 2007). Furthermore, observing advice seems to be better than observing actions in fostering social learning (Çelen and Kariv, 2010; Kocher et al., forthcoming).

the unobserved *state* (the car quality in the example). We can think of each subject as having his own cage, which depends on his type, and all the cages are correlated through the state. In stage 1 of the experiment, subjects observe a sample drawn from their own cage, and make an estimate. We refer to this estimate as the *opinion*. In stage 2, there is no sample. Instead, a subject (the *client*) observes the opinion of a subject from stage 1 (the *expert*) who was in the same state and observed a sample. The client and the expert may be of different types (this is the *bias*). This is fully observable. Hence, the expert's opinion is relevant to the client, since they are in the same state and the expert has private information in form of the sample. But the opinion is potentially biased since the expert might be of a different type. However, as this is observable, the bias can be corrected. The subjects are scored using a quadratic rule and payoffs depend only on the their own estimate, which eliminates any strategic incentives.

Our approach is most similar to Nyarko et al. (2006), who set up an investment game in which subjects must estimate the probability that an investment is profitable, and choose whether to invest. The authors create a set of experimental experts by having a group of subjects play the game first, and then elicit their beliefs and a piece of investment advice. This information is auctioned off to a new set of subjects who observe a set of personal characteristics of the expert (college major, gender, etc.), and the relation between the auction price and the expert characteristics is investigated. In our paper, on the other hand, experts have an exogenous bias, and we analyze how this bias affects clients' propensity to use expert opinions.

First, we investigate the usefulness of expert opinions. Opinions are found to underweight sample information and to be slanted toward the midpoint which gives equal probability to both events (*Result 1*). Underweighting of sample information is a well-known experimental phenomenon,[3] whereas slant toward the midpoint is most likely caused by risk-aversion. The consequence of this and general noisy updating by experts is that opinions have lower average score than the Bayesian prior.[4] However, opinions

[3]Although representativeness – overweighting of sample information – is normally more prevalent, both types of updating co-exist, and the propensity to use one or the other is context-specific (Grether, 1980; El-Gamal and Grether, 1995; Nyarko et al., 2006).

[4]By the score of the opinion and the prior, we refer to the score that a subject would obtain by

still contain significant information about the state. Next we investigate how clients use expert opinions. As it turns out, clients do incorporate expert opinions into their estimates, but evidence suggests that they are more likely to use the prior than to use the opinion (*Result 2*). The clients' average score is the same as that of the expert opinion, which is lower than the average score of the prior.

Second, we turn to the main question: does the propensity to use the opinion depend on the bias? In our setting, bias is exogenous and easily corrected for, and therefore should not affect clients' estimates. However, we find that the propensity to use opinions is decreasing in the size of the bias (*Result 3*). This finding is robust to the inclusion of controls such as computational costs,[5] failure to adjust for bias, aversion to make estimates that are different to the prior and heterogeneous informativeness of experts. Thus, our results are indicative of a type of *bias aversion* among clients: even though bias is observable and can be adjusted such that all experts are a priori equally informative, clients prefer less biased experts.[6]

Third, we investigate how the aforementioned control variables affect the propensity to use opinions and the clients' score. We show that the propensity to use opinions decreases the farther the opinion is from the prior, but is unaffected by ex-post expert informativeness and computational costs (*Result 4*). The first effect captures client aversion to changing beliefs, whereas the third effect suggests that clients are not very good at deducing the expert's ability. The clients' score is unaffected by bias size, as well computational costs and distance between opinion and prior (*Result 5*).

2 Experimental Design

Subjects were recruited from the undergraduate population at Universidad Carlos III de Madrid. They spent between 1/2 and 1 1/2 hours to complete the tasks, but were

setting his estimate equal to either of these.

[5]We test for computational costs by including a dummy that identifies bias/no bias.

[6]This has a very similar flavor to homophily, whereby clients choose experts that have personal characteristics similar to themselves (Nyarko et al., 2006). To the extent that lack of bias invokes feelings of similarity, the result can be related to studies which show that affect leads to more updating mistakes (Charness and Levin, 2005; Charness et al., 2007).

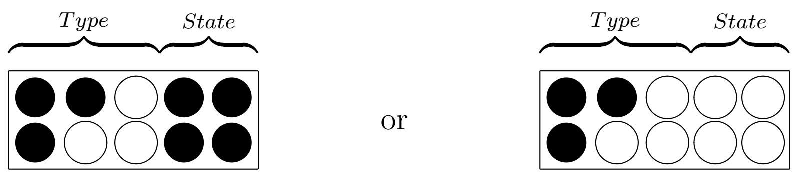

Figure 1 – Example with $Type = 3$. The left cage ($State = 4$) occurs with probability π, the right cage ($State = 0$) with probability $1 - \pi$.

not allowed to leave the experiment for the first 45 minutes. Average earnings were €12, including a €4 show-up fee. All experiments were computerized using common interactive instructions.[7] Four sessions were run in November and December of 2010, with a total of 66 subjects.

Subject Task. The experiment is constructed as a probability estimation game. Each subject must estimate the probability of drawing a black ball from a cage of 10 balls which are either black or white.[8] The first 6 balls are of known color, and these represent the type of the subject. Let $Type \in \{0, ..., 6\}$ denote how many of the 6 balls are black. The remaining 4 balls are not observed, and are either all black or all white. We denote by $State \in \{0, 4\}$ the number of these balls that are black. Let $\mathbb{P}(State = 4) = \pi$ and $\mathbb{P}(State = 0) = 1 - \pi$, with $\pi \in \{\frac{1}{2}, \frac{1}{3}\}$. $Type$ and π change throughout the experiment, but are always known to subjects. Thus, they can calculate the prior as $Prior = (ClientType + 4\pi)/10$.

The total number of black balls in the cage is hence $Type + State$, and the true probability of drawing a black ball is $(Type + State)/10$. Figure 1 illustrates the setup for the case where $Type = 3$. In this case the total number of black balls is either 3 or 7. Conditional on this information, the expected number of black balls is $7\pi + 3(1 - \pi)$.

Information. The experiment has two stages, and the subject's information depends on the stage. In stage 1, subjects observe a sample consisting of 3 balls drawn with

[7]The software used for the experiment was z-Tree (Fischbacher, 2007). The instructions are available online at the author's website.

[8]In the experiment, balls were either white or red, but for expositional purposes we use black here.

replacement from the cage, and make an estimate. Denote this by $Sample \in \{0, ..., 3\}$. In stage 2, we refer to the subjects as *clients*. Clients do not observe a sample, but instead observe the stage 1 estimate of another subject – the *expert* – who did the experiment at a previous date, and was in the same state. The expert has better information, since he observed a sample, but he may have a different type: $ExpertType \neq ClientType$. This is observed by clients. Thus, the expert's estimate is potentially biased in the view of the client, and therefore we refer to it as $BiasedOpinion$.

Since the client knows both his own type and that of the expert, he can calculate expert bias as $Bias = (ExpertType - ClientType)/10$. We refer to the bias-adjusted opinion as the $Opinion$ for short, and calculate it as $Opinion = BiasedOpinion - Bias$. The expert changes in each period and hence, clients cannot learn about the expert. This is made clear in the instructions. Thus, $Opinion$ would be the client's best estimate if all subjects were risk-neutral Bayesian updaters.

Let us recap the main features of the experiment. The expert's estimate is relevant to the client: both are in the same state, and the expert has private information. But if the expert and the client are of different types, they face different probabilities of drawing a black ball, even if they are in the same state. Therefore, from the point of view of the client, the expert's opinion is biased. Notice also that whereas the prior is straightforward to calculate, the posterior is more complicated. However, it is easy for subjects to come up with rules of thumb. For instance, letting $SampleProp$ be the proportion of black balls in the sample, a good estimate is $.8 \times Prior + .2 \times SampleProp$.[9] Table 1 summarizes the stages.

Table 1 – Stages

	Stage 1	Stage 2
Sample	Yes	No
Expert Opinion	No	Yes
Periods	20	30

[9]Actually, subjects who followed this rule would have scored higher (see scoring rule below) than Bayesian updaters.

Timing and Payoffs. The timing of the experiment is as follows. In each period of stage 1, subjects observe their type, their sample, and make an estimate. In each period of stage 2, clients observe their own type, the type of the expert, the expert's opinion, and then make an estimate.

Payoffs are calculated using a quadratic scoring rule, which elicits the true subjective probability of a subject under risk neutrality.[10] After each subject's cage has been determined by the subject's type and the state, a ball is drawn randomly from the cage by the computer. Let $Outcome = 1$ if the ball is black and $Outcome = 0$ if it is white. Denote the subject's estimate of the probability of drawing a black ball by $Estimate$. The subject's score is then

$$Score(Estimate) = \begin{cases} 10000 \times (1 - (1 - Estimate)^2) & \text{if } Outcome = 1, \\ 10000 \times (1 - Estimate^2) & \text{if } Outcome = 0. \end{cases}$$

After the experiment, the score is converted at the rate of 5000 points to one Euro, and subjects are paid for 5 randomly chosen rounds. Since subjects are paid only as a function of their own estimate, experts have no strategic incentives.

3 The Informational Content of Opinions

In this section we investigate the informational content of opinions and the propensity of clients to use them. Since the opinions are taken from stage 1 estimates, we first look at these. Let $Post$ denote the Bayesian posterior, conditional on the sample.[11] Figure 2 maps the distribution of $Estimate - Post$ in stage 1 and, as can be seen, it is single-peaked, almost symmetric around zero, but with thinner tails than the normal distribution. A t-test shows that the mean is not significantly different to zero (p-value = .61).[12] Next, we consider how estimates relate to the prior and the opinion. Subjects are in general "conservative", in the sense that they underweight the information contained in the sample, and skew their estimates toward the prior as well as the midpoint (i.e.

[10]This rule is extensively used in experiments and dates back to Brier (1950).

[11]In particular, $Post = \frac{ClientType}{10} + \frac{4}{10} \cdot \frac{\pi \cdot \mathbb{P}(Sample, State=4)}{\pi \cdot \mathbb{P}(Sample, State=4) + (1-\pi) \cdot \mathbb{P}(Sample, State=0)}$.

[12]All standard errors reported and used for tests are clustered around subjects.

Table 2 – Score

	Stage 1	Stage 2
Mean $Score(Estimate)$	7,546	7,411
Mean $Score(Prior)$	7,723	7,599
Mean $Score(Opinion)$	-	7,415

1/2).[13] Since the midpoint gives constant payoff, we can see slanting toward it as an indication that subjects are risk-averse. Conservative updating is a common phenomenon, although it it is often found to be less prevalent than representativeness (overweighting data) and approximate Bayesianism (Grether, 1980; El-Gamal and Grether, 1995). In our setting, we speculate that conservatism is more prevalent due to the complexity of the task. The opinions are taken from stage 1 and will therefore in general reflect these patterns.

Table 2 breaks down the average score by stage, and compares it to how much subjects would have scored if they had "blindly" followed either the prior or the opinion. Notice that in stage 1 they score considerably lower than the prior. This may be due in parts to risk-averse behavior. The pattern is repeated in stage 2: opinions give lower average score than the prior, and a t-test says that this difference in mean scores is significant (p-value $<.001$). However, this does not imply that opinions do not contain information beyond the prior. Regressing the outcome on the prior and the opinion, we find that both the coefficient of the prior and of the opinion are positive and significant.[14] This leads us to our first result.

Result 1. *Opinions underweight sample information, are slanted toward the midpoint and yield lower mean score than priors. However, opinions still contain significant information beyond that contained in the prior.*

Since opinions carry significant information, a clever client might be able to combine

[13]Regression result: $Estimate = .16(.04)Midpoint + .40(.05)Prior + .44(.04)Post$. Standard errors in parenthesis. $Midpoint$ is a constant that takes the value 1/2. The positive coefficient on $Prior$ and the less-than-unit coefficient on $Post$ indicate that sample information is being underweighted.

[14]Regression result: $Outcome = .16(.04) + .47(.08)Prior + .20(.07)Opinion$.

Table 3 – Benchmarks

Prop. of estimates within 2 pts. of	
Opinion	.11
Prior	.23
BiasedOpinion	.05

Note: We include observations with no overlap between the measures, i.e. observations are not "double-counted".

these pieces of information to obtain a better score. However, as Table 2 shows, on average clients do about as well as the experts' opinions, and they have significantly lower average score than the prior.[15] Furthermore, there does not seem to be a learning effect: regressing score on period in stage 2 does not reveal a significant relationship.[16] To shed further light on how priors and opinions are used by clients, Table 3 shows the proportion of stage 2 client estimates that were within 2 points of the prior and the opinion, respectively. Twice as many estimates were close to the prior as were close to the opinion. Furthermore, regressing the estimate on the prior, the bias-adjusted opinion and the unadjusted opinion in stage 2, we find that the coefficient on the prior is significantly higher than the coefficient on the opinion.[17] We summarize this in our second result.

Result 2. *Clients have higher propensity to follow the prior than to follow the opinion. However, on average clients score the same as the opinion, and lower than the prior.*

An interesting empirical corollary to this result is found in Pogrebna (2008), who shows that game show contestants often do not follow the audience's advice, even when

[15] According to a t-test, the mean score of clients and opinions are not significantly different (p-value=.91), whereas the mean score of clients and priors are significantly different (p-value $< .001$).

[16] Regression result: $Score = 7451.6(105.7) - 2.6(6.0)Period$.

[17] Regression result: $Estimate = .12(.04)Midpoint + .56(.05)Prior + .09(.03)BiasedOpinion + .27(.05)Opinion$. A t-test reveals that the coefficient on *Prior* is significantly higher than both the coefficient on *Opinion* (p-value =.002), the coefficient on *BiasedOpinion* (p-value<.001), and the sum of the two (p-value=.02).

this is in general quite informative. In our setting, there is a great deal of uncertainty about the expert and the updating rule that he may have followed to arrive at his opinion. This may explain why expert opinions are used rather cautiously by clients.

4 The Effect of Bias

In this section we analyze the effect of expert bias on the propensity of clients to use expert opinions. Subsequently, we discuss alternative factors that might influence the propensity to use opinions. Lastly, we look at whether expert bias affects client score.

4.1 Bias and Opinions

To measure the effect of bias on clients' propensity to use opinions, we construct the variable $\Delta = Opinion - Prior$ and interact it with $BiasSize$. This interaction captures how clients shift weight from the prior to the opinion as the bias size increases. If it is negative, it implies that clients' propensity to use opinions is decreasing in the size of the bias.

To check the robustness of the effect of bias size, we test for a number of alternative explanations of the propensity to use expert opinions. First, it is possible that bias matters only due to computational concerns, such that clients place less weight on opinions when the bias-adjustment is harder to compute. Since the bias-adjustment always follows the same formula, there is only a computational difference between the cases where the bias is zero, and those where it is non-zero. Thus, we construct the variable $Zero = \mathbb{I}(BiasSize = 0)$, where $\mathbb{I}(\cdot)$ is the identity function, and interact it with Δ. This allows us to check whether clients react to bias/no bias rather than the size of the bias. Second, clients may have a psychological bias in the sense that they are reluctant to use opinions that are far away from the prior. To test for this, we interact Δ with $|\Delta|$. Third, it is possible that the informativeness of the expert's opinion depends on his type.[18] Therefore, we interact Δ with $ExpertType(k) = \mathbb{I}(ExpertType = k)$.

[18]Suppose, for instance, that subjects believe that $Opinion = 0$ if $ExpertType + ExpertSample < 5$ and $Opinion = 1$ otherwise. If $ExpertType < 2$, the opinion is uninformative since $Opinion = 0$ always, but for $ExpertType \geq 2$ some information is transmitted. Therefore, the expert type may

Fourth, it is conceivable that the client perceives the ability of the expert through the estimate (for instance, if the estimate is clearly not Bayesian), and adjusts for this. Therefore we interact *ScoreOpinion* with Δ, to see if opinions that ex-post lead to a higher score are more likely to be used. Finally, we allow for the possibility that subjects forget to adjust for the bias, and use *BiasedOpinion*.

The regression equation is then

$$\begin{aligned} Estimate = {} & \beta_0 Midpoint + \beta_1 Prior + \beta_2 Opinion + \beta_3 BiasedOpinion \\ & + \Delta \times [\delta BiasSize + \delta_0 Zero + \delta_D |\Delta| + \delta_S ScoreOpinion] \\ & + \textstyle\sum_k \gamma_k \Delta ExpertType(k) + \epsilon, \end{aligned} \tag{M1}$$

where ϵ is a normally distributed error term. We are interested in the coefficient δ which measures how the weight given to the opinion changes with the size of the bias.

We estimate M1 using all stage 2 observations and present the results in Table 4. The coefficient δ is negative and significant, and thus the bias size is negatively correlated with the weight placed on the opinion. To test the robustness of our approach, Appendix A estimates two alternative models to test the influence of bias size on clients' propensity to use expert opinions. Both models confirm the above result on δ. We thus have the following result.

Result 3. *Bias size has a negative effect on clients' propensity to use expert opinions.*

The smaller the bias of the expert, the more likely it is that his opinion is used. This finding has the same flavor as Nyarko et al. (2006), but their finding of homophily was related to personal characteristics, whereas our types are exogenously assigned. We have controlled for a number of alternative factors which could influence the weight given to opinions versus priors, and none of them seem to take away the effect of bias size. Our best explanation is that some clients are simply bias averse. Although they realize that they can adjust for the bias (because they do so when the bias is small), they are reluctant to do so when the bias is large. Furthermore, this effect is economically significant: for each unit of bias, the coefficient on the opinion decreases by 0.1.[19]

influence the propensity to use opinions.

[19] Notice that the maximum value the coefficient on the opinion can take is .85, since it is decreasing in all the interaction terms. Hence, a change of 0.1 constitutes more than 10% of the initial value.

Table 4 – M1: Client Estimates

	Estimate	
Prior	-.01	(.16)
Opinion	.85***	(.17)
$\Delta\times$ BiasSize	-.10**	(.03)
$\Delta\times$ Zero	-.10	(.10)
$\Delta \times \|\Delta\|$	-.79**	(.23)
$\Delta\times$ ScoreOpinion	-.04[a]	(.06)[a]
BiasedOpinion	.08*	(.03)
Midpoint	.10*	(.04)
TypeControls	Yes	
R^2	.54	
Observations	1980	

Estimated by OLS. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors are clustered around subjects.

[a] The coefficient and standard error have been multiplied by $10,000$.

4.2 Other Determinants of the Propensity to Use Opinions

Above we consider three alternative hypotheses: the propensity to use opinions depends on either the presence rather than the size of bias, on the distance of the opinion to the prior, or on the ex-post score of the opinion. Of these hypotheses, only one has bite: since δ_D is negative and significant (Table 4), clients have lower propensity to use expert opinions the farther they are from the prior.[20] On the other hand, δ_0 is insignificant, and hence clients do not simply pay attention to unbiased experts and discard the rest. Similarly, δ_S is insignificant, implying that the propensity of clients to use opinions is not correlated with the ex-post realized score of the opinion.

Result 4. *Clients' propensity to use expert opinions is decreasing in the distance between the opinion and the prior. However, the propensity is unaffected by the score of the opinion and is equal for experts with, respectively, no bias and strictly positive bias.*

The lower propensity to use opinions that are farther from the prior may simply express reluctance of the clients to move their estimate away from their prior. But we can also speculate that clients use the distance of opinions to priors to derive information about the expert, and feel less compelled to trust experts who have opinions that are very different to the prior. However, if this is the explanation, the clients fail the objective: as illustrated by the second part of the result, clients are very bad at picking high-scoring experts. How could they do this? Since learning is disabled in the experiment, the only manner in which clients can evaluate experts is by their opinion. These may be informative about the expert in the following sense: if the estimate is outside the interval in which the true probability lies (e.g. the interval $[0.3, 0.7]$ in Figure 1), the expert must be either risk-loving, of low ability or have misunderstood the exercise. This should lower the client's propensity to use the opinion (unless he is himself risk-loving). Clients do not seem to have benefited from such considerations. Finally, the finding that – conditional on other covariates – clients do not have higher propensity to use opinions of unbiased experts, implies that we can rule out that clients are motivated by simple computational concerns.

[20] We also run model M1 including $|BiasedOpinion - Prior| \times \Delta$ to check whether it is the distance between the prior and the unadjusted opinion that matters, but find no effect.

Table 5 – M2: Score and Bias Size

	Score	
BiasSize	-57.6	(74.3)
Zero	-105.1	(256.5)
$\lvert\Delta\rvert$	-137.2	(481.3)
Constant	7623.7***	(279.2)
TypeControls	Yes	
R^2	0.05	
Observations	1980	

Estimated by OLS. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors are clustered around subjects.

4.3 Bias and Score

We have shown above that clients react to expert bias, and also that clients are not very good at picking high-ability experts. We now want to investigate whether there is an interaction between the two: does bias size affect client score? As earlier, we control for the expert type. The regression equation is then.

$$Score = \beta + \delta BiasSize + \delta_0 Zero + \delta_D|\Delta| + \sum_{k=0}^{5} \gamma_k ExpertType(k) + \epsilon, \tag{M2}$$

where ϵ is a normally distributed error term. δ captures the effect of bias size on score.

We estimate model M2 using all observations in stage 2 and present the results in Table 5. The coefficient δ is not significant, and neither are δ_0 or δ_D. We summarize this in our final result.

Result 5. *There is no effect of bias size on score. Furthermore, the score is unaffected by the distance between the opinion and the prior and is equal for experts with, respectively, no bias and strictly positive bias.*

Thus, although clients react to bias by shifting away from the opinion toward the prior, this has no effect on their score on average. In our case, this may be a consequence of the low informativeness of opinions, which implies that although giving less weight to the opinion means ignoring the expert's private information, it also means incorporating less noise.

5 Conclusion

We have explored how opinions are used in a setting where clients must simultaneously handle the fact that experts are biased and of unknown ability. In our setup, priors are easy to compute but posteriors less so. However, experts can do very well by using simple heuristics, and therefore their estimates should contain useful information. The results shed light on how subjects react to bias in these situations. We have set up the experiment to make the bias-adjustment easy, in order to make it as easy as possible for clients to use expert opinions. Furthermore, we control for a range of factors which might influence the clients' propensity to use opinions. We find that clients have significantly smaller propensity to use the expert opinion the greater the bias.

The results are indicative that bias matters, even when it should not. Whereas homophily based on personal characteristics may represent a certain type of statistical discrimination (or at least, a belief that this can be done), in our case no such explanation is available. Neither does our setup invoke emotional responses in the subjects. Bias is purely exogenous and refers simply to the composition of a cage.

The findings are relevant for the study of experts, and especially for situations where bias is a choice variable (such as in the study of commercial news media). Clients prefer "straight talk", rather than having to extrapolate information by considering expert bias. Thus, the choice of bias has an effect beyond the informativeness of the opinions and the emotional responses they elicit.

A Alternative Models for Bias and Opinions

In this appendix we provide two alternative models that test how clients' propensity to use opinions are affected by bias size.

Alternative Model 1: Propensity to Use Opinions. As a first alternative measure, we look at "behavior types" who use either the opinion or the prior (within two percentage points) directly. Such behavior accounts for 34% of the total estimates (Table 3). We then construct the variables $UseOpinion = \mathbb{I}(|Estimate - Opinion| \leq 0.02)$ and $UsePrior = \mathbb{I}(|Estimate - Prior| \leq 0.02)$, and analyze how the propensity to use priors or opinions depends on the bias.

$$UseOpinion = \delta BiasSize + \delta_0 Zero + \sum_{k=0}^{5} \gamma_k ExpertType(k) + \epsilon, \tag{A1}$$

where ϵ is an error term, and the model is estimated by Probit. Again we are interested in the coefficient δ which measures how bias size affects the propensity to use the opinion.

Table 6 – Alternative Model 1: Using the Benchmarks

	UseOpinion		UsePrior	
BiasSize	-0.20***	(0.04)	0.22***	(0.05)
Zero	-0.21	(0.13)	0.73***	(0.18)
Constant	-0.31	(0.20)	-1.36***	(0.22)
TypeControls	Yes		Yes	
Pseudo R^2	0.030		0.073	
Observations	1837		1837	

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Only includes stage 2 observations where $UseOpinion \neq UsePrior$.
Standard errors are clustered around subjects.

Table 6 presents the results of model A1, which confirm that bias size is negatively and significantly correlated with client propensity to use the opinion. The table also shows that, conversely, bias size is positively correlated with the propensity to use the prior. Model A1 thus corroborates Result 3.

Alternative Model 2: Direction of Estimates. As a second alternative measure, we look at the cases in which the prior and the opinion are on different half-intervals, and analyze the propensity of clients to choose one half-interval or the other. Let $EstimateHigh = \mathbb{I}(Estimate > 0.5)$, $OpinionHigh = \mathbb{I}(Opinion > 0.5)$ and $PriorHigh = \mathbb{I}(Prior > 0.5)$. This measure is very robust, since the half-interval of the client's estimate should not depend on his risk attitude.[21] The model is then

$$EstimateHigh = \delta OpinionHigh \times BiasSize + \beta_0 + \beta_1 OpinionHigh + \epsilon, \tag{A2}$$

where ϵ is an error term. We estimate this by Probit, using only observations for which the opinion and the prior are on different half-intervals: $OpinionHigh + PriorHigh = 1$. In this way, δ captures the effect of the bias size on the propensity of the client's estimates to be on the half-interval of the opinion rather than that of the prior.

Table 7 – Alternative Model 2: Half-intervals

	EstimateHigh	
OpinionHigh	1.87***	(0.54)
OpinionHigh× BiasSize	-0.82***	(0.25)
OpinionHigh × Zero	-1.56**	(0.55)
Constant	-0.20	(0.16)
Pseudo R^2	0.054	
Observations	234	

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Only includes stage 2 observations where $OpinionHigh \neq PriorHigh$. Standard errors are clustered around subjects.

As revealed by Table 7, δ is negative and significant: the propensity of clients to

[21] However, it may depend on his beliefs about the expert's risk attitude.

choose the same half-interval as the expert's opinion is decreasing in bias size. Hence, Result 3 is confirmed by Model A2.

B Figures

Figure 2 – Deviation from Bayesian Posterior

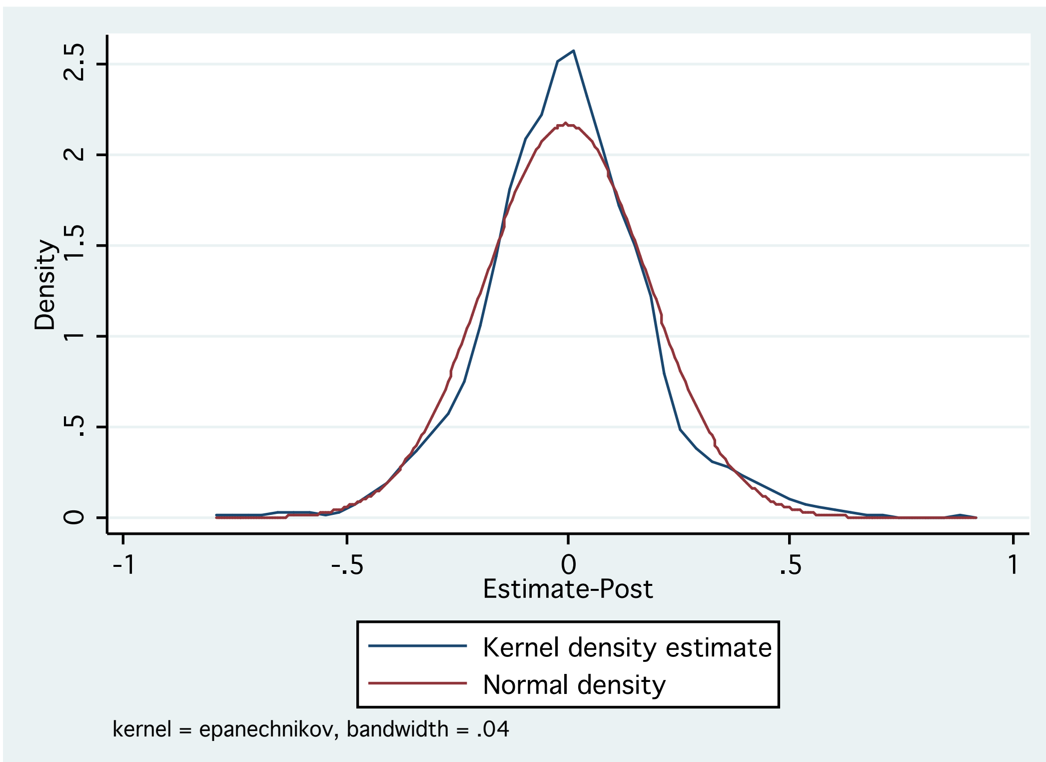

References

A. Blume and A. Ortmann. The Effects of Costless Pre-play Communication: Experimental Evidence from Games with Pareto-ranked Equilibria. *Journal of Economic Theory*, 132(1):274–290, 2007.

G. W. Brier. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78(1):1–3, 1950.

H. Cai and J. T. Wang. Overcommunication in Strategic Information Transmission Games. *Games and Economic Behavior*, 56(1):7 – 36, 2006.

B. Çelen and A. Kariv, S.and Schotter. An Experimental Test of Advice and Social Learning. *Management Science*, 56(10):1687–1701, 2010.

G. Charness and D. Levin. When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. *American Economic Review*, 95(4): 1300–1309, 2005.

G. Charness, E. Karni, and D. Levin. Individual and Group Decision Making under Risk: An Experimental Study of Bayesian Updating and Violations of First-order Stochastic Dominance. *Journal of Risk and Uncertainty*, 35(2):129–148, 2007.

A. Chaudhuri, S. Graziano, and P. Maitra. Social Learning and Norms in a Public Goods Experiment with Inter-generational Advice. *Review of Economic Studies*, 73 (2):357–380, 2006.

R. Cooper, D. V. DeJong, R. Forsythe, and T. W. Ross. Communication in the battle of the sexes game: Some experimental results. *RAND Journal of Economics*, 20(4): 568–587, 1989.

R. Cooper, D. V. DeJong, R. Forsythe, and T. W. Ross. Communication in coordination games. *Quarterly Journal of Economics*, 107(2):739–771, 1992.

V. Crawford. A Survey of Experiments on Communication Via Cheap Talk. *Journal of Economic Theory*, 78(2):286–298, 1998.

M.A. El-Gamal and D.M. Grether. Are People Bayesian? Uncovering Behavioral Strategies. *Journal of the American Statistical Association*, 90(432), 1995.

U. Fischbacher. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10(2):171–178, 2007.

U. Gneezy. Deception: The Role of Consequences. *American Economic Review*, 95(1): 384–394, 2005.

D.M. Grether. Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *Quarterly Journal of Economics*, 95(3):537–557, 1980.

S. Hurkens and N. Kartik. Would I Lie to You? On Social Preferences and Lying Aversion. *Experimental Economics*, 12(2):180–192, 2009.

M. Kocher, M. Sutter, and F. Wakolbinger. Social Learning in Beauty-contest Games. *Southern Economic Journal*, forthcoming.

Y. Nyarko, A. Schotter, and B. Sopher. On the Informational Content of Advice: A Theoretical and Experimental Study. *Economic Theory*, 29(2):433–452, 2006.

G. Pogrebna. Naïve Advice When Half a Million Is at Stake. *Economics Letters*, 98(2): 148 – 154, 2008.

S. Sánchez-Pagés and M. Vorsatz. An Experimental Study of Truth-telling in a Sender–Receiver Game. *Games and Economic Behavior*, 61(1):86–112, 2007.

A. Schotter and B. Sopher. Social Learning and Coordination Conventions in Intergenerational Games: An Experimental Study. *Journal of Political Economy*, 111(3): 498–529, 2003.

A. Schotter and B. Sopher. Trust and Trustworthiness in Games: An Experimental Study of Intergenerational Advice. *Experimental Economics*, 9(2):123–145, 2006.

A. Schotter and B. Sopher. Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach. *Games and Economic Behavior*, 58(2):365–393, 2007.