Economic Theory (2005)
DOI 10.1007/s00199-005-0037-7

SPECIAL ISSUE ARTICLE

Yaw Nyarko · Andrew Schotter
Barry Sopher

On the informational content of advice: a theoretical and experimental study

Received: 15 November 2004 / Accepted: 2 September 2005

Abstract This paper examines the market for advice and the underlying perception that advice is useful and informative. We do this by first providing a theoretical examination of the informational content of advice and then by setting up a series of experimental markets where this advice is sold. In these markets we provide bidders with a demographic profile of the "experts" offering advice.

The results of our experiment generate several interesting findings. The raw bid data suggest that subjects bid significantly more for data than they do for advice. Second, in the market for advice there appears to be no consensus as to who are the best advisors although on average economists demand the highest mean price and women suffer a discount. In addition, we find that whether a subject suffers from a representativeness bias in the way he or she processes data has an impact on how he or she bids for advice and on his or her willingness to follow it once offered. Finally, we find that on average people impute a low level of informativeness onto advice, consistent with their bidding behavior for data versus advice.

Keywords Advice · Decision making · Risk aversion · Experiments

JEL Classification Numbers D81 · G11 · C91

This work was done under grant number SES-0425118 of the National Science Foundation. The authors would like to recognize the Center for Experimental Social Science at New York University for its additional support. We also acknowledge the help of Elizabeth Potamites for her research assistance.

A. Schotter (✉) · Y. Nyarko
Department of Economics, New York University, 269 Mercer Street,
New York, NY 10003, USA
E-mail: andrew.schotter@nyu.edu

B. Sopher
Department of Economics, Rutgers University, 75 Hamilton St., NJ Hall Room 202,
New Brunswick, NJ 08901-1248, USA

1 Introduction

It is commonly thought that a picture is worth a thousand words. If that is so, one might ask how much data is a piece of advice worth. In other words, if advice is important then we should be able to measure it in two ways: how much data would a rational decision maker be willing to give up in order to receive a piece of advice from a person who has just engaged in the decision problem he or she is about to engage in or, alternatively, how much would that person be willing to pay for such advice from a person with a given set of characteristics.

The fact that we expect that people will bid different amounts for advice from different types of people implies that in the "market for advice" certain types of people are likely to fare better than others. Such markets for advice and influence might function in a number of ways. Under one scenario there may be a perception that certain people or types of people are worth listening to. These perceptions amount to broad stereotypes that may bestow huge rents on some of the agents in the market.

Such stereotypes, if they persist, can lead to what we will call "perception rents", i.e. amounts paid for the advice of agents in excess of the expected informational content contained in their opinion. If such rents are substantial, they present us with a potentially large inefficiency. An alternative to perception rents is what we will call the "chauvinistic bias". Here people tend to believe that advice from people like themselves is the best and hence tend to bid higher amounts for people with characteristics like theirs whether or not those types give the best advice.

A related question deals with the impact of what are sometimes called "representativeness" and "conservative" biases, relative to Bayesian updating, and their implications for the process of advice giving and following. For example, when updating one's beliefs, a rational Bayesian decision maker is expected to place a certain amount of weight on his previous prior (or the base rate) and a certain amount on new information (the sample) as it arrives. How much weight is placed on the new information depends on the strength of his or her prior. If a decision maker places more than the Bayes-optimal weight on the prior (or base rate) he or she is called "conservative" while if excessive weight is placed on the sample he or she is considered to be subject to the "representative" bias, thinking, in the limit, that the sample received is in some sense representative of the population from which it was drawn. Such people fail to take base rates or priors sufficiently into account.

These concerns have wide ranging implications for our research on advice giving and following. For example, if we could measure the degree to which a decision maker is subject to one of the biases discussed above, could we correlate that characteristic to the decision maker's willingness to pay for and follow advice. More precisely, if conservatives are reluctant to update their priors on the basis of new information, are they therefore less inclined to pay for advice and also follow it once it is given? Also who are more persuadable, conservatives or representatives? All of these concerns can be summarized under the name of "advice bias".

In this paper we study an experimental market for advice in an attempt to measure both the informational content of advice and the market for it. To do this we create a set of "experts" by having some of our subjects get experience playing a simple 2×2 "game against nature" a large number of times. These experts are surveyed to obtain information about their gender, GPA, major, and year in school. Advice is elicited from these experts and is then sold to a new set of subject "clients" who play the game once and only once. The prices generated by this market for advice furnish us with an opportunity to measure potential perception rents in the market. In addition, by observing the decisions made by the client subjects, we are able to measure whether these perception rents are wasteful or not. By observing the way subjects update their priors in the experiment, we are able to categorize them along the conservative-representative spectrum according to how much weight they put on their prior and how much they place on the data sample they observe. We are then able to correlate their degree of conservatism with their behavior in the advice-following experiment. For example, are economists or scientists more likely to be Bayesians and hence process data correctly? If so, does that explain their increased value in the market. Do subjects impute a degree of representativeness for their advisors that is greater than their own, etc? Finally, the data generated by our experiment allow us to calculate what the informational content of advice is by imputing how many observations a subject would be willing to give up in order to receive a piece of advice.

The results of our experiment generate several interesting findings. First we find that, in general, our client subjects bid for data amounts that, on average, are approximately equal to the expected values of the information they might expect to receive. The raw bid data suggest that subjects bid significantly more for data than they do for either advice or beliefs. We also find a little evidence for perception rents for economics majors and a certain amount of support for what we call the "chauvinistic bias," meaning that subjects tended to bid more for advice from people sharing the same major as themselves than for people of other majors.

In answering these questions we find that the way a person processes data, i.e. how much he suffers from a base-rate bias affects his or her behavior in the market for advice quite dramatically. For example, if subjects place more weight on data than they should under perfectly rational Bayesian updating they tend to bid more for information in the form of data, advice or beliefs than those who place less than a perfect Bayesian weight on data. Finally, we find that people tend to impute a higher degree of base rate bias to others, i.e. those who are giving them advice, than they themselves have. This tends to make advice worthwhile since it gives an advisee more insight into the sample of observations than he would get if his expert were, say, a person who processes data as he does. In other words, he would rather get advice from an expert who was different than himself than one who was the same.

The rest of the paper is organized as follows. In Section 2 we will describe the experiment run to investigate the questions raised above. In Section 3, present the theory of decision making with advice that will be used as our guide in analyzing the data. In Section 4 we present the results of our experiments. Section 5 contains a brief review of some related literature, and conclusions are contained in section 6.

2 Experimental design

2.1 Experimental overview

In this paper we report the results of an experiment that involves, among other things, subjects playing an "Investment" game, where in each period a subject must choose to invest or play it safe and keep his money in a secure asset. This game is first run on a set of subjects in order to create a pool of "experts" who will be used to give advice to subjects ("clients") who later play the game.

We then run a game where we create "clients" or "advisees." We run two variants of these. In the first, the *Price-Elicitation game*, we auction off[1] the different types of advice from different types of advisors (those with different demographics, gender, academic major, GPA, etc.). In the second, the *Belief-Elicitation game*, after a subject arrives in the lab we read them instructions and then elicit their prior belief about the state of nature θ (using a proper scoring rule as found in Nyarko and Schotter 2002). We then offer them some data drawn from the distribution defined by the true θ, and elicit their posterior based upon this data. Finally after they state this posterior we offer them some advice from an expert who has performed the investment game many times and ask them to state yet another updated posterior. Hence, in the belief elicitation game subjects only report beliefs and do so three times and are paid for the accuracy of these reports.

Finally, a third variant, the *Belief-Price game,* was conducted where subjects played the belief elicitation experiment first followed by the price elicitation experiment. This gives us an ability to correlate behavior across these games for the same subjects. In each of the advisee experiments, after the advisee receives the information, the advisee will play the investment game. This enables us to study how the advisee uses the information. Since almost all of our data analysis only involves data from the Belief-Price game, and also due to space limitations, we will concentrate on that game most extensively in the discussion of our experimental design.

2.2 The investment game

In the basic investment game played in our experiment there are two actions, and in each period an individual is required to make a decision to either invest (the Investment Option) or not invest (the "Safe Option"). The financial market has two possible states, profitable or unprofitable, in each period. The investment option yields a total return which depends upon the state, while the return to the safe option is independent of the state. The payoffs are described in the matrix below (Table 1):[2]

The probability of the profitable state is equal to θ. The value of θ is unknown to subjects. If the subject chooses the safe option, the subject receives a return of 5. If she chooses the investment option, she will receive ten when the state is profitable and 0 otherwise. The value function is therefore given by

$$V = \max\{5, 10E\theta\}. \tag{1}$$

We set the prior of θ so that "rational" or Bayesian updating resembles the familiar "fictitious play" rule of thumb updating. In particular, we suppose that θ is drawn from a beta distribution (or equivalently the tuple $(\theta, 1-\theta)$ is drawn from a Dirichlet distribution). The beta distribution is parameterized by a constant α.

[1] More accurately, we elicit subjects maximum willingness to pay for different advisors using the Becker-DeGroot-Marschak mechanism.

[2] We denominate everything in the paper in US \$. Subjects performed the experiment in units denominated by "experimental" francs at the exchange rate of 10 francs to a \$1. For example, the stock-bonds game would have entries 100, 0, 50 and 50 francs.

Table 1 The payoff matrix

State	Profitable	Unprofitable
	Prob= θ	Prob= $(1-\theta)$
Investment option	\$10	\$0
Safe option	\$5	\$5

We use three different values of α . One of these results in the uniform distribution for $\theta(\alpha = 1)$; the remaining two will result in, respectively, a U-shaped density function ($\alpha = 1/4$) and an inverted U-shaped density function ($\alpha = 4$).

The subjects playing the investment game were told the value of α and it was explained to them in the instructions that θ is chosen from the beta (α, α) distribution. They were of course not told the value of θ. They were told that conditional on θ, the profitable and unprofitable states will be chosen with probabilities θ and $(1 - \theta)$ respectively. Under the beta prior over θ the mean of θ, and hence the prior probability of the profitable state, is 0.5.

Throughout this paper we assume that subjects are risk neutral. Risk aversion would imply subjects choose the safe option more often than predicted by the theory. We do not see this in the data; this is not presented in this paper due to due to space constraints. In the theory section we will describe what Bayesian updating implies for our model. There will also be a discussion there on deviations from this, and our empirical section will test some of the deviations.

2.3 Creating experts or advisors

As stated above in our experiment we have two types of subjects – experts and clients. To create our experts we ran sessions at the experimental laboratory at Rutgers University during the spring and summer of 2002, and at the Center for Experimental Social Science at New York University in the Fall of 2002. Subjects were recruited primarily from undergraduate courses. Before the students began the experiments we recorded information on their gender, age, class, major, grade point average. They were paid \$5.00 simply for showing up. The subjects were then given the instructions for the game. In particular, they were shown graphically the distribution function corresponding to the value of α for their game, and they were informed of how the θ would be chosen for their experiment, but not told its actual value. Payoffs in the expert games were denominated in experimental dollars and converted into U.S. dollars at a rate of 1 experimental dollar = \$ .05.

The subjects played the investment game over three rounds or 25 rounds, depending upon the session.[3] At the beginning of each round, each player chooses an action – the investment option or the safe option. The computer then generates for that round via the fixed θ a profitable/unprofitable realization. The subjects are then told the profitable/unprofitable state, and they are paid for that period. Then they go to the next round and repeat the process. The profitable/unprofitable state is drawn in each period independently, according to the fixed probability θ.

[3] We actually ran this particular game for 24 rounds (rather than three) or for 75 rounds (rather than for 25), with updating of beliefs every three or 25 periods. We only pass on beliefs or advice to clients based on the first three or 25 periods.

At the end of the game (round 25) subjects were asked what they believe the probability of the profitable state is. They were rewarded for this decision via a *quadratic scoring rule* (see Nyarko and Schotter 2002). We call these beliefs the *elicited beliefs*. After obtaining the elicited beliefs, subjects are then asked to give an "Investment option" or "Safe option" recommendation to be used by other future subjects. Note that we did not reward subjects for their recommendation. We refer to this as the subject's advice although sometimes instead of offering advice in the form of recommendations we simply offered the client the experts's beliefs over θ.

To summarize, the subjects received payment from three sources: \$5 for showing up, money from the belief elicitation, and money from playing the Investment Game. Subjects, on average, earned approximately \$15 for their participation which was paid to them at the end of the session. Each subject spent between 30 and 45 min on the game.

As we mentioned in the earlier section, we had three different values of α (1/4, 1 and 4). For each value of α, we randomly generate, from the beta(α) distribution, three independent draws of θ . This results in a total of $3 \times 3 = 9$ different combinations of α and θ for each of the two sets of expert games ($N = 25$ and $N = 3$). In each of the nine games there were between eight and 12 subjects. This information on the experiments is described in Table 2.

For each θ and for each subject having that θ, we also compute the "pay rank" of that subject – the ranking in terms of performance in the money received by that subject among the cohort of subjects with the same θ. For example, as the table shows, there were 11 subjects having a $\theta = 0.586$. Each one of those 11 subjects will receive a unique number from 1 to 11 showing how well they did monetarily in their experiment.

3 Creating clients or advisees: the Belief-Price game

As stated earlier, we will focus almost exclusively here on the data of the Belief-Price game. In the Belief-Price Game, each subject took part in two different games. In the belief-elicitation part we set the distribution from which θ was drawn at $\alpha = 1$, the uniform case. Subjects then saw three data points from the $N = 3$ expert games. The realized observed data was (Profitable, Profitable, Profitable) for all subjects. In the second period all subjects received advice. This advice was

Table 2 Parameters for expert game

Sessions with N=25			Sessions with N=3		
α	θ	# Experts	α	θ	# Experts
4	0.59	11	4	0.29	10
4	0.35	9	4	0.37	10
4	0.65	10	4	0.58	10
1	0.78	8	1	0.14	10
1	0.44	10	1	0.29	10
1	0.52	8	1	0.81	10
1/4	0.11	10	1/4	0.07	10
1/4	0.87	11	1/4	0.33	10
1/4	0.15	10	1/4	0.96	9
		Total=87			Total=89

from a 21 year old male Senior Economics major with a GPA of 3.7 and payoff rank of fourth out of ten who played the investment game with the same θ. The advice given was, "INVEST." We elicited beleifs before each peirod. Finally, all subjects played the Investment game once after the belief elicitation process. The resulting market state in the single play of the Investment game that ends the Belief Elicitation game was always unprofitable.

After the subjects played the Belief Elicitation game they then played a version of the Price Elicitation game. Subjects were told that the θ for this experiment was drawn from the same prior distribution (the uniform or $\alpha = 1$ case), but was independent of the θ of the Belief Elicitation game they had just participated in. Subjects then bid on data, beliefs, and advice from advisors. Subjects bid on 12 different experts. For each expert (advisor) they were told some information (gender, age, GPA, year, major, and pay rank), and they placed bids on the advisor's data, beliefs and advice. They were told, as indeed was the case, that ten of the experts were "hypothetical" and that only two were actual subjects who had played the investment game. After the bids have been entered, we applied the Becker, DeGroot, Marschak (1964) (BDM) procedure to determine what information, if any, is made available to the advisee. In particular, since there are only two real advisors each with three types of information (data, beliefs, advice), there are six possible information choices; one of these is randomly chosen and the BDM mechanism is applied to the chosen one to determine whether that information is observed by the advisee. Subjects were told that only their bids on the real experts would be chosen for the actual BDM mechanism.

4 Some theory

4.1 Updating beliefs

In our experiment both experts and clients get an opportunity to observe a sample of independent realizations drawn with some unknown probability θ of which some N_1 are profitable and N_0 unprofitable. Application of Bayes' rule indicates that the posterior probability of θ will be beta with parameters $(\alpha + N_1, \alpha + N_0)$, with the posterior probability of the profitable state, which is equal to the posterior expectation of θ, given by

$$E\theta = \frac{\alpha + N_1}{2\alpha + N_1 + N_0}. \tag{2}$$

As is well-know, this updating also has a bounded rationality or "rule of thumb" interpretation. Ignoring the α and 2α term above, the expectation is nothing other than the simple average. The $\alpha's$ represents the strength of the updater's prior beliefs before seeing the data – she behaves as if she has seen $s = 2\alpha$ prior data points of which a fraction $\alpha/2\alpha = 1/2$ were profitable, and then applies the average to both the prior and current data. We will discuss later whether this Bayesian updating is followed and also the deviations from this observed in the data.

If the experts we recruit are rational in the above sense, then when asked to offer their posterior beliefs to future clients they should state the probabilities derived above and suggest investment whenever this posterior probability is greater than 0.5.

In the belief elicitation game a subject will state his or her prior, see a sample of observations, update his or her beliefs on the basis of this information and then be given either the beliefs of an expert or some invest/don't invest advice from their expert and update once more. In particular, in the belief elicitation game we observe three beliefs: the prior, the belief after observing data, and the belief after receiving either beliefs or advice. (In some experiments the order is reversed and the subject will get advice first.) In this section we ask how a Bayes-rational agent would behave in this experiment. We do this by describing how the subject would update his or her beliefs if in both stages he or she observes data (rather than data and either beliefs or advice) and then use this analysis to generalize it to the case of advice.

In the initial elicitation of beliefs, given the description of the distributions from which θ is drawn, we would expect the prior probability over the profitable state to be $b_0^* = 0.5$ for all subjects. Suppose that in each round the advisee sees N observations, with a fraction m_1 profitable in the first round, and m_2 in the second. Then it is easy to see that the rational probabilities of the profitable state after the two rounds of observations are respectively $b_1^* = (\alpha + Nm_1)/(2\alpha + N)$ and $b_2 = (\alpha + Nm_1 + Nm_2)/(2\alpha + 2N)$. Define

$$\psi_0^* \equiv \frac{N}{2\alpha + N} \text{ and } \psi_1^* \equiv \frac{N}{2\alpha + 2N}. \tag{3}$$

Then, if we define a function B as $B\,(b, m, \psi) \equiv (1 - \psi)\, b + \psi m$, we have

$$b_0^* = 0.5;\ b_1^* = B\left(b_0^*, m_1, \psi_0^*\right); \text{and } b_2^* = B\left(b_1^*, m_2, \psi_1^*\right). \tag{4}$$

Hence, we may think of B as the Bayesian Updating function, and we may think of ψ_0^* and ψ_1^* as the weight placed on data. Typically in discussing the updating formulas, the concept of strengths is often used. The prior strength s is related to ψ_0 and ψ_1 via

$$\psi_0 = \frac{N}{s + N} \text{ and } \psi_1 = \frac{N}{s + 2N}, \tag{5}$$

with rational strengths in round 1 and round 2 [see (3)], given by $s_0^* = 2\alpha$ and $s_1^* = 2\alpha + N$.

4.2 Representativeness/conservativeness biases

The analysis offered above is written from the point of view of a perfect Bayesian decision maker. However, various studies argue that people are subject to a number of biases when they make decisions such as the "base rate" bias (Kahneman and Tversky 1973). According to this bias people have trouble incorporating base rate frequencies into their updating and tend to use the representativeness of a sample as a proxy for the parameter they are attempting to estimate. Put differently, these studies argue that people tend to place more than the optimal Bayesian weight on data or information and tend to ignore the prior or posterior they have formed. Such people we will call "representative" while those who place too much weight on their prior will be called "conservative".

For example, say I have a prior for θ of 0.5 and see a sample of 25 profitable/unprofitable realizations of which 75% are profitable. A Bayesian will combine the information in the sample with his or her prior and move his prior somewhere in the interval between 0.5 and 0.75. with weights Ψ^B placed on the sample determined by the number of observations and α. A person who is subject to a representative bias will place a weight $\Psi > \Psi^B$ while one subject to a conservative bias will place a weight on the sample $\Psi < \Psi^B$. In our experiment we measure the degree of the representativeness bias by looking at the behavior of subjects in the first stage of their belief elicitation game, i.e. by observing how they update their prior given data.

4.3 Updating after observing advice or beliefs

The formulas above were stated for the situation where the advisee receives data in each period. We now discuss the updating after obtaining advice or beliefs from advisors. Here we will discuss the benchmark case where the advisee believes that the advisor is presenting the beliefs and advice without error – we show that the formulas stated earlier continue to hold with reinterpretations of the meanings of m_1 and m_2. In later sections, and again in the empirical results sections, we will discuss deviations from these assumptions.

Suppose that the advisor has initial beliefs $b_0^{\text{advisor}} = 0.5$, sees N realizations of the market state with fraction m^{advisor} being profitable and communicates to the advisor the beliefs $b_1^{\text{advisor}} = \left(b_0^{\text{advisor}} s_0 + m^{\text{advisor}} N\right) / (s_0 + N)$ where $s_0 = 2\alpha$. If the advisee knows the above perfectly then the advisee would use this to infer the precise value of m^{advisor}. In particular, upon observation of the beliefs of the advisor, the advisee can infer the value m^{advisor}, the average number of profitable states observed by the advisor. If the advisee receives beliefs of the advisor in round 1 (resp. round 2) and data in the other period, the updating formula will be the same as in (4) but where we replace m_1 (resp. m_2) with m^{advisor}.

Finally suppose that the advisor presents to the advisee invest/don't-invest advice. Suppose that the advisee believes that the advisor is presenting the advice rationally, and in particular that the advisor will present the advice Invest (resp. don't Invest) whenever the advisor sees data with fraction of profitable states $m^{\text{advisor}} > 1/2$ (resp. $m^{\text{advisor}} < 1/2$). Define $\hat{a}$ to be the expected value of m^{advisor} conditional on the advice: $\hat{a} \equiv E\left[m^{\text{advisor}} \mid advice\right]$. In particular, $\hat{a}$ will take two values, $\hat{a}^{invest}$ and $\hat{a}^{don't\text{ invest}}$, according to whether the advice is invest or don't invest. The values of $\hat{a}$ are a function of α and can easily be calculated.[4]

Our updating formulas (4) were for the situation where data is obtained in each of the two rounds. Suppose instead that the advisee receives invest/don't invest advice in round 1 (resp. round 2) and data in the other round. Then the updating formula will be the same as in (4), but where we replace m_1 (resp. m_2) with $\hat{a}$.

[4] When $N = 25$ the values are:

α	$\hat{a}^{don't\text{ invest}}$	$\hat{a}^{invest}$
4	0.3425	0.6575
1	0.2400	0.7600
0.25	0.1147	0.8853

4.4 Rational behavior in the Price-Elicitation game

In the price elicitation game subjects are asked to state how much they would be willing to pay for data and advice or beliefs from various different advisors. Let us begin with the price of data. If we let N_1 and N_0 be the profitable and unprofitable random realizations of the market state from the purchased data, then from (2) it is easy to see that rational price of data is

$$P = E_{N_1}\left[\max\left\{5, 10\left(\frac{\alpha + N_1}{2\alpha + N_1 + N_0}\right)\right\}\right] - 5, \tag{6}$$

where the expectation is taken over N_1 with respect to the distribution generated by the α (i.e., beta(α, α)). It is easy to compute explicitly the price of $N = 25$ observations of data for the different values of α we use: \$0.5964, \$1.2037 and \$1.8885 for $\alpha = 4$, 1 and 1/4 respectively. Notice that the price is decreasing in α. We indicated earlier that the strength of beliefs is 2α. An $\alpha = 4$ represents stronger beliefs and therefore a smaller price is paid for data than with weaker beliefs, say $\alpha = 1/4$.

To determine the rational price of an advisor's beliefs or advice, we have to decide what to assume about the advisee's beliefs about the advisor. Suppose we assume that the advisee believes that the advisor is rational and gives the advisee the correct beliefs or advice (we study deviations from this assumption later).

Let us price the advisor's advice first. The advisee can solve the advisor's problem and will realize that the advice "invest" will be announced whenever the advisor sees more than $N/2$ profitable states, otherwise the advice "do not invest" will be announced. The advisee will therefore choose the action invest (do not invest) when she gets the advice "invest" (resp. do not invest). This in turn means that the advisee will choose the action invest (resp. do not invest) whenever her advisor's data has more than (resp. less than) $N/2$ profitable states. This is, however, exactly the same rule that the advisee will follow if she receives the data directly. We therefore conclude that the price that the advisee should pay for the opportunity to see this advice is exactly the same as the price that the advisee should pay for data.

One should also see that the same prices hold for advisors' beliefs. In particular, assume that the advisor reports the correct belief conditional on the information she has seen. Then the advisor will report beliefs of value greater than (resp. less than) 0.5 whenever the advisor's data has more than (resp. than) $N/2$ profitable states, so the earlier arguments hold.

To summarize, the prices of advice and beliefs are the same as those for data as stated above.

4.5 The Informational content of advice

The above studied different methods of updating beliefs of the advisor, using assumptions on the strength of advisors updating function. Each of these earlier methods assume that the advisee treats data and advice with the same strengths. We now analyze the case where the advisee puts different weight on beliefs as opposed to data. In particular, we suppose that with beliefs, the agent assumes that one unit of "real" data is worth κ units of underlying data when that underlying data is transmitted through beliefs. Again, we focus on the case where the advisee receives the advisor's beliefs in the first period and data in the second period.

Recall the definitions of the weights on signals as a function of a fixed strength s in (5). Given fixed initial strength, s, optimal or otherwise, these two equations imply a mapping from the weight on signal in period 0 to the weight in period 1:

$$\psi_1 = \frac{N}{s + 2N} = \frac{N}{\left(\frac{N(1-\psi_0)}{\psi_0}\right) + 2N} = \frac{\psi_0}{1 + \psi_0} \equiv \Psi_1 (\psi_0) \tag{7}$$

Inverting the above, we obtain $\psi_0 = \psi_1/(1 - \psi_1) \equiv \Psi_0 (\psi_1)$. The weight on initial period information should be $\psi_0 = N/(s + N)$. Suppose instead that since the information is obtained via an advisor's beliefs, the advisee thinks that one observation by the advisor is worth κ units of direct observation, so that the weight on information is actually $\psi_0^\kappa \equiv \kappa N/(s + \kappa N)$ so that Bayesian updating is: $b_1 = B\left(b_0, \hat{m}, \psi_0^\kappa\right)$. For example, if the advisee thinks that her advisor is prone to errors, she may set $\kappa < 1$; alternatively, if she thinks that her advisor is better able to process information than she is, she may set $\kappa > 1$.

Now, in the second period, the updating is using data so there is no discounting applied to this period (i.e., $\kappa = 1$) . If ψ_1 is the signal strength in the second period, then as mentioned earlier the updating in the first period, if there was no discounting via kappa, would be $\Psi_0 (\psi_1) \equiv \psi_1/(1 - \psi_1)$. In particular, $\Psi_0 (\psi_1)$ is the no discounting (or $\kappa = 1$) period 1 strength, while ψ_0^κ is the actual discounted strength. One would suspect that the ratio of these strengths, $\psi_0^\kappa / \Psi_0 (\psi_1)$ is related to κ. Indeed, if we set $\psi_1 = N/(s + \kappa N + N)$, which follows from our assumptions on updating, then $\Psi_0 (\psi_1) \equiv \psi_1/(1 - \psi_1) = N/(s + \kappa N)$ so

$$\kappa = \frac{\psi_0^\kappa}{\Psi_0 (\psi_1)}. \tag{8}$$

Kappa therefore has a nice interpretation: it is the the ratio of the discounted weight on signals over the implied "real" or $\kappa = 1$ weight on signals.

Next, as empirical measures of the strengths, we compute the advisee strength in the period in which data is obtained, and use (7) to compute the advisee strength in the other period. In particular, define the following "experimental" values: $\psi_0^{\text{emp}} \equiv B_\psi^{-1} (b_1, b_0, m_1)$ and $\psi_1^{\text{emp}} \equiv B_\psi^{-1} (b_2, b_1, m_2)$. Note that ψ_0^{emp} (resp. ψ_1^{emp}) can unambiguously be determined if data is received in period 1 (resp. 2), since in that case case m_1 (resp. m_2) does not involve inverting the beliefs or advice of the advisor. Hence when data is received in the second period, we use $\left(\Psi_0 \left(\psi_1^{\text{emp}}\right), \psi_1^{\text{emp}}\right)$.

5 Questions and answers: empirical results

Since there are three main categories of question that we are interested in asking we will divide up our analysis accordingly: 1) the market for advice, 2) belief bias, and 3) the informational content of advice. We now address these questions in turn.

Table 3 Client bids for data, beliefs and advice by major Belief-Price Elicitation game

Client major	Means	Expert major Econ	Sci	Hum	Social Sci	Total
Econ	Data	1.49	0.871	1.02	1.18	1.20
	Beliefs	1.14	0.550	0.609	0.678	0.815
	Advice	0.902	0.388	0.495	0.625	0.655
	N	99	59	43	60	261
Sci	Data	2.32	2.51	2.14	1.74	2.18
	Beliefs	1.48	1.41	1.38	0.806	1.28
	Advice	1.31	1.40	1.05	0.665	1.12
	N	49	24	28	32	133
Hum	Data	0.998	1.01	1.98	1.17	1.20
	Beliefs	0.600	0.404	0.640	0.553	0.544
	Advice					
	N	63	50	32	47	192
Social Sci	Data	1.83	1.88	1.69	2.62	1.97
	Beliefs	0.997	0.939	0.484	0.484	0.761
	Advice	0.657	0.889	0.428	0.348	0.581
	N	83	38	50	45	216
Other	Data	1.42	1.89	0.889	1.57	1.45
	Beliefs	1.34	1.98	1.02	1.65	1.49
	Advice	1.41	2.01	1.05	1.67	1.52
	N	36	23	197	18	96
Total	Data	1.60	1.43	1.56	1.62	1.56
	Beliefs	1.07	0.865	0.750	0.713	0.865
	Advice	0.893	0.800	0.637	0.637	0.767
	N	330	194	172	202	898

5.1 Question 1: the market for advice

5.1.1 In the market for advice are there perception rents or a chauvinistic bias? In other words, is there a consensus as to who is the best advisor and hence do those people enjoy a perception rent?

The raw data from the bids that client subjects made for expert data, beliefs and advice are contained in Table 3. On the face of it it appears that there are perception rents in the advice markets we set up since, on average subjects bid more for the advice of economics students followed by scientists humanists and social scientists.

More precisely, in terms of means, subjects tended to bid 0.893, 0.800, 0.637, and 0.637 to hear the advice of economics students, science students, humanists and social science (other than economics) students, respectively and 1.07, 0.865, 0.75, and 0.71. to observe the beliefs of these same types of subjects. From these means it would seem as if economics students are perceived to be the best followed by scientists, humanists and social science majors. Also it appears that people are willing to pay more for beliefs than the binary advice invest/not invest. This seems intuitive, to the extent that clients believe that experts are not infallible, or that clients for some reason do not trust experts to communicate properly what they know. Recall that a perfectly rational Bayesian account predicts equal amounts to be paid for data, beliefs and advice.

But perception rents only exist if there is a consensus by all types of subjects that certain types of people are best at a given task. When we disaggregate the data we find the interesting phenomenon that subject types were sometimes more willing to pay for the advice of their own type than for that of others. This is particularly true of data (Table 3). For beliefs and advice, we see some tendency for economics and/or science majors to receive higher bids (e.g., from social science majors) yet humanists persist in bidding more for advice from other humanists than from any other type of advisor. Of course, there really is no reason to bid more for data resulting from the play of one type of expert than another, since the expert has no role in processing this information, unlike the case of beliefs and advice.

As we see, there is certainly no strong consensus among the subject types as to which expert types (major) are the best advisors. This apparently refutes the perception rent hypothesis. On the other hand, there is some evidence pointing to economics majors being favored in the bidding for beliefs and advice. We investigate this question more formally with regression analysis of bidding behavior in the next section.

5.1.2 What determines the price of advice or beliefs for different experts?

In the market for advice all experts can be represented as a bundle of characteristics (age, major, gpa, gender, rank on experts experiment, etc). Under this interpretation of experts as a bundle of characteristics, we can expect that each of these factors could contribute to the price a given subject would be willing to pay for advice form such an agent. To explore this relationship we ran a hedonic type of regression, in which we estimate the contribution of expert characteristics to the bid price. We report results for the pooled regression (with all bidder types) and with disaggregated regressions (one for each bidder type). The results of these regression are contained in Table 4.[5]

The estimation results substantiate much of our previous analysis and add some new dimensions as well. Note that the bids in these regressions are denominated in the experimental currency, or ten times the actual dollar value. In the pooled regression, economists get a significant premium over other majors. Neither scientists nor humanists garner bids that are significantly different from those of social scientists (the default category in the regression). The dummy variables for bids on beliefs and advice are both negative and significant, confirming the fact that less is bid for beliefs and advice than for data. There is a significant penalty to female experts. Neither the expert's gpa nor the expert's age are significant, but the ranking of the expert based on his or her performance relative to others who played the same investment game is an important factor. This variable ranges from roughly 0.1 (for the highest ranked) to 1 (for the lowest ranked). Thus, there is a significant penalty for being a low-earning expert.

In the disaggregated regressions, if we find that all bidder types pay a significant premium for one expert type category, this can be interpreted as support for

[5] The regression treats the data as a panel (repeated observations of a cross section of bidders), with a GLS random effects error structure. Note that this structure, since it takes into account individual-specific effects, should control for the effect of individuals who bid systematically higher than others. Thus, the effect of outliers here should not be a concern, as it was in examining the raw bid data.

Table 4 Determinants of bid price

Belief-Price Elicitation game Dep. var.: bid	Pooled reg.	Econ. reg.	Sci. reg.	Hum. reg.	Soc.Sci. reg.
Expert female	−1.27[a]	−0.99	−3.71[a]	−0.03	−0.69
	0.51	0.98	1.21	0.70	1.09
Expert economist	1.57[a]	2.52[a]	4.19[a]	1.12	0.01
	0.68	1.28	1.49	0.91	1.50
Expert scientist	0.89	1.31	1.27	0.82	0.34
	0.77	1.49	1.77	0.98	1.77
Expert humanist	−0.19	−0.05	2.85[b]	0.20	−0.67
	0.78	1.57	1.69	1.08	1.64
Expert GPA	0.29	−1.11	4.36[a]	0.82	1.62
	0.53	1.00	1.31	0.73	1.12
Expert age	−0.14	0.20	−0.72[b]	−0.09	−0.00
	0.18	0.32	0.41	0.27	0.40
Expert payoff rank	−7.16[a]	−8.89[a]	−14.77[a]	−3.80[a]	−10.64[a]
	0.86	1.65	2.05	1.18	1.94
Belief bid dummy	−6.77[a]	−3.89[a]	−8.95[a]	−6.66[a]	−12.14[a]
	0.59	1.14	1.32	0.80	1.26
Advice bid dummy	−7.95[a]	−5.42[a]	−10.59[a]	−7.09[a]	−13.95[a]
	0.59	1.14	1.32	0.80	1.26
Constant	21.29[a]	15.96[a]	27.94[a]	12.82[a]	20.64[a]
	4.22	7.84	9.76	6.09	9.07
N	2693	782	399	576	648
R^2	0.06	0.04	0.17	0.09	0.16
Rho	0.50	0.49	0.56	0.41	0.40

Numbers on the bottom in each cell is the standard error.
[a] 5% significance
[b] 10% significance

the perception rent hypothesis. The results do not support this. Both economist and scientist bidders pay a significant premium for economists experts, but humanists and social scientists do not. There is a consensus that the payoff ranking of the expert is important. All types also agree in paying less for beliefs and advice than for data. With respect to gender, we find that only scientists pay significantly less for female experts which would indicate that the result in the pooled regression was driven by this sub group. These separate or disaggregated regressions then tend to cast doubt on the perception rent hypothesis without providing compensating strong evidence for the chauvinistic hypothesis (Table 4).

5.1.3 Is the market for advice efficient in that those who are in fact the best advisors receive the highest bids for their advice?

Now that we know how much subjects are willing to pay for the advice of various types, it is of interest to know how rational these bids are, i.e. are the bids consistent with how accurate our expert types really are. We provide some simple computations to answer these questions based on the result of the expert game.

To characterize the performance of our experts we use three different performance measures. In Metric 1 our accuracy measure is simply the fraction of subject experts who provided the correct advice (i.e., the advice that a Bayesian decision

Table 5 Performance of advisors by major in expert games

Major		Expert game (N=25) Metric 1	Metric 2	Metric 3	Expert game (N=3) Metric 1	Metric 2	Metric 3
Econ.	Mean	0.87	−0.07	0.83	0.68	−0.03	0.64
	Sd	0.34	0.11	0.38	0.48	0.27	0.49
	N	53	53	53	22	22	22
Sci./Math	Mean	1	−0.02	0.94	0.67	−0.02	0.78
	Sd	0	0.03	0.24	0.50	0.31	0.44
	N	17	17	17	9	9	9
Other	Mean	0.82	−0.08	0.59	0.67	0.06	0.66
	Sd	0.39	0.09	0.51	0.47	0.25	0.48
	N	17	17	17	58	58	58
Total	Mean	0.89	−0.06	0.80	0.67	0.03	0.66
	Sd	0.32	0.10	0.40	0.47	0.26	0.48
	N	87	87	87	89	89	89

maker would have provided, given the observations) in round $N = 3$ or $N = 25$.[6] Metric 2 asks how different are the actual beliefs provided from the "correct" beliefs that should be provided given the observations of the individual? Finally Metric 3 is slightly more complex and deals with the case where the client receives the expert's beliefs as advice. If you were to receive an advisor's beliefs, and if you believe that your advisor has updated her beliefs correctly, then you would choose the investment option when the beliefs exceed 0.5, and the no investment option when beliefs are less than 0.5. One can then compute that number of times the actual beliefs of advisors, given their observed data, results in the same answer a Bayesian decision maker would have arrived at. This measure of accuracy of beliefs just checks whether the announced beliefs agree with Bayesian beliefs. We score the accuracy as 1/2 when elicited beliefs are exactly 0.5 while rational beliefs are either strictly above or below 0.5. Table 5 contains the results of these calculations.

As we see, if we were to lump experts by major into three categories, Economics majors, Science/Math majors and other we can see that Science/Math majors appear to make the best advisors. For example, for the N=25 expert game they outperform each of the other types on all three metrics. In fact, they are the only group that offered the correct advice all of the time in the sense of suggesting investing when the state of the world was profitable. There is a significant difference in the means of Economics and Math/Science majors at the 6 and 5% levels for Metrics 1 and 2, respectively. There is no significant difference for Metric 3. In the $N = 3$ expert game their advantage is less clear, and there is essentially no difference between the Economic majors and the Other major, except for Metric 2. None of the differences are significant.

5.1.4 Do people bid too much too little or just the right amount for advice and information?

It is of interest to ask whether the prices determined in the market are rational, i.e. are their levels too high or too low. Remember that the optimal bid for advice,

[6] Six out of the ten advisors who gave the incorrect advice had data for which the correct beliefs are between 0.44 and 0.56. (For the entire data, only 16 out of 87 had data resulting in beliefs in the same region). In particular, the six out of ten advisors who gave the incorrect advice may not have been too far off in their decisions.

beliefs and data are exactly the same and are a function of the α parameter of the distribution they are drawn from. As we derived in section 3, the optimal bid for data, beliefs and advice are \$ 0.596 if $\alpha = 4$, \$1.203 if $\alpha = 1$ and \$1.888 if $\alpha = 1/4$. Since in our Belief-Price experiment we only used $\alpha = 1$ we expect to see \$1.20 bid for information, beliefs and advice.

As we saw in Table 4, it is clear that subjects bid more for data than they do for either advice or beliefs. Put differently, ex ante they believe that data will be more informative. The amounts they bid for data are above the expected value of such information (mean bid on data = \$1.56), and this is significantly larger than \$1.20. Of course, as mentioned before, there is quite a bit of variance in bids.

The reason why data appears to be more valuable might be because data is unfiltered by any expert analysis. Hence, if subjects fear that their advisor is something less than fully rational they may decide that observing data is more informative than listening to the poorly processed advice or beliefs they receive. In light of this argument, it is not surprising that the bids of subjects for advice and beliefs are below the expected informational content of such advice. In other words, if subjects pay, on average, the expected value of information for data they can be expected to pay less for advice and beliefs. The same argument can be applied to the comparison between beliefs and advice. Subjects are willing to pay more for beliefs than advice because they feel it is a sharper instrument being continuous instead of binary. The mean bid for beliefs is \$0.89, and the mean bid for advice is \$0.77, and both are significantly less than \$1.20.

5.2 Question 2: belief biases

5.2.1 Do subjects tend to update their beliefs in a Bayesian manner? If not, do they tend to give more or less weight to data than is optimal, i.e., do they suffer from a representative bias?

To answer this question we focus on the Belief-Price experiments where subjects observe three observations taken from a distribution where α =1. We focus on the Ψ_0 they use to update their prior after observing three positive observations in a row indicating that the state is profitable. We consider this a good game to focus on since it offers subjects a limited number of observations (after seeing, say, 25 observations, the optimal weight to put on the sample is practically 1 so it is hard to differentiate who is Bayesian from who is representative). The fact that all subjects saw three positive observations is also beneficial since a Bayesian would update his belief by placing a weight of 0.6 on the sample and 0.4 on his or her prior leading to a posterior of 0.8 while a totally representative agent would place all of his or her weight on the sample and update his or her prior to 1.

Table 6 presents the results of our calculations. It presents the person-by-person weights that subjects used in their updating . As we can see, the median weight that subjects used was 0.6 so the median subject updated using perfectly Bayesian weights. (The mean was 0.55 which is not far off.) In addition, note that subjects are almost equally split between being conservative, using a weight less than 0.6 on data and representative, using a weight greater than 0.6. Interestingly, 21 subjects proved themselves to be perfectly representative and placed a weight of 1 on the data they saw. This was the modal choice. However, the next two most frequently

Table 6 $\Psi_0's$ In the Belief-Price Elicitation game

Ψ_0	Frequency	Ψ_0	Frequency
−1.64	1	0.48	1
−0.5	2	0.5	9
−0.34	1	0.52	1
−0.33	1	0.555	1
0	4	0.6	8
0.085	1	0.625	1
0.142	1	0.657	1
0.2	1	0.666	1
0.249	2	0.7	3
0.285	2	0.749	1
0.3	1	0.75	1
0.333	1	0.8	4
0.375	2	1	21
Mean = 0.550, Median = 0.6			

used weights were the Bayesian weight, which was chosen by eight subjects, and the perfectly conservative weight of 0 chosen by four subjects. So while Bayesian weighting seems to be a mode for the subject population there is a very wide distribution of subjects with some acting as if they were conservative and others being representative (with a large number totally representative).

5.2.2 Are advisees who are more representative in their updating more likely to bid more for advice and data?

In general the answer to this question is yes. The Ψ_0 that subjects exhibited in the belief-elicitation part of the Belief-Price game is a measure of the weight given to new information (whether from data, beliefs or advice). Table 7 contains the average bids made for data, advice and beliefs conditional on whether a subjects is categorized as being Bayesian $\Psi_0 = 0.6$, conservative $\Psi_0 < 0.6$ or representative $\Psi_0 > 0.6$.

Note that while there is not strictly monotonicity exhibited in the table, it is clearly true that subjects who are representative, $\Psi_0 > 0.6$ bid more for all types of information, data, advice or beliefs, than do subject who are conservative $\Psi_0 < 0.6$. It is only in the case of data that Bayesians bid more that representative types. In general, representative types seem to respect data more and bid more for it. The bids of conservatives are significantly lower than those of representative types for all types of information according to a t-test on the means. The distributions of bids by conservatives also differ from those by representatives according to a nonparametric Wilcoxon rank sum test. The bids of Bayesians are not significantly

Table 7 Mean bids by type conservative, Bayesian, representative (Belief-Price Elicitation game)

	Type of information		
Type of subject	Data	Beliefs	Advice
Conservative $N = 27$	0.99	0.69	0.57
Bayesian $N = 8$	2.04	1.07	0.62
Representative $N = 33$	1.91	1.09	1.02

different from representatives for data or beliefs, nor from conservatives for advice, by either type of test. The significant tests are all highly significant (1% level or better).

5.2.3 Are advisees who are more representative in their updating more likely to follow advice?

In the data we used here to answer this question, subjects have seen three positive observations from the relevant distribution and have updated their priors once. The weight they used in this initial updating was their Ψ_0. We then offered them advice in the form of a binary suggestion about whether they should invest or play it safe. In fact all the advice they received was positive, i.e. all were told to invest. We then measured the change in their belief after receiving this advice and now take that as a proxy of the impact of advice on them. Those who change more are more affected by the advice. (We do not look at whether subjects followed the advice they were given since there is little variance in this, as most people chose the invest option.)

To examine this question we ran a regression where the left hand variable was the change in a subjects belief after receiving advice. We regressed this on the person's Ψ_0 and a dummy variable which took the value of 1 if the advice offered was counter to the best response choice of the subject to his updated beliefs (updated after seeing three observations). The results of this regression are contained in Table 8.

This regression points to an interesting phenomenon. Note that Ψ_0 has a negative sign. While this might lead one to think that subjects who are highly representative are less likely to follow advice, it is interesting to note why this is true. To explain take the extreme case. Say that I am completely representative and observe three plus signals during my observation of data. In this case I would update my prior to 1. Now when I get a piece of advice to invest, as we have indicated before, unless I also think that my advisor is as perfectly representative as I am, I would invert his advice for its informational content and assume that he probably saw a less optimistic sample than I did. In this case, ironically, his advice to invest is actually treated as bad news since it implies that the sample he saw was less positive than mine. Decreasing one's posterior after such advice is therefore not as perverse as one might think and explains the negative coefficient on the Ψ_0 variable. Note that 21 subjects were fully representative so that this result applies to many people in the sample.

5.2.4 Question concerning the informational content of advice

How many data points is a piece of advice based on N observations worth? Do people over value the informational content of advice or discount it? As our introduction

Table 8 The impact of advice on change in beliefs

Belief-Price Elicitation game (N = 73, R^2 = 0.54)				
Variable	Coef	Std err	t	$P >\| t \|$
Ψ_0	−0.212	0.045	−4.68	0.00
Dummy	0.374	0.073	5.15	0.00
constant	0.096	0.030	3.20	0.002

indicates one piece of information we are interested in is knowing how informative advice is when compared to data. Note that our analysis about the informational content of advice has no role for expertise. When experts offers advice the client simply assumes that the expert processed the data he saw optimally (or with some type of representative or conservative bias) and then inverts the advice he gets to extract its informational content. In other experiments where there is a task where skill is involved it may be that advice takes on even greater content then it does here.

Despite this caveat we will see here that some subjects find advice to be quite informative. Advice based on three observations is, on average seen to be the equivalent of 3×1.44 observations in the Belief-Price Elicitation game. This may seem to contradict our previous finding that people pay less for advice and beliefs than they do for data. This may be because we have restricted ourselves here to the behavior of what we call "twice-rational" subjects. More precisely, in order for our analysis in section 3.3 to work subjects must satisfy some minimal rationality requirements. Their updating should be monotonic in that if they are offered beliefs which are more positive that their posteriors, then those beliefs should lead them to increase their beliefs and not reduce them. Put differently, we require that both Ψ_0, Ψ_1 be in the interval [0,1]. A Ψ_i outside of this bound leads to either negative updating (if $\Psi_i < 0$) or overshooting (if $\Psi_i > 0$). In what we do below we restrict our sample to only those subjects whose updating weights in both their first and second updating periods of the Belief elicitation portion of the Belief-Price game were in this interval i.e. $0 < \Psi_0, \Psi_1 < 1$. There are 35 such individuals in the Price-Belief game, 17 in the Belief Elicitation game (N = 3), and 27 in the Belief Elicitation game ($N = 25$), and for each of them we calculate the κ as outlined in Section 3.3. These results are presented in Table 9 where we break the data down by whether the value of kappa is equal to 0, between 0 and 1, or greater than 1.

As we can see, by restricting our sample to twice-rational subjects we have no negative κ's, by construction. While the average kappa is greater than 1 in the Belief-Price game, the mean is less than 1 in the other two Belief Elicitation games. Striking is the number of subjects with a kappa=0. Thus, a majority of these subjects, in fact do not value advice very highly, though a significant subset of them do. Clearly there is a lot of heterogeneity in behavior here, warranting further investigation.

6 Conclusions

In this paper we have investigated the impact of advice on decision making. Using an experimental design where "subject experts" gain experience in playing a game and then pass on advice to their "subject clients" about a relevant parameter in

Table 9 Informational content of advice (κ)

	Belief-Price (3 obs.)			Belief Elicit. (3 obs.)			Belief Elicit. (25 obs.)		
Range	Mean κ	S.d.	N	Mean κ	S.d.	N	Mean κ	S.d.	N
$\kappa = 0$	0	0	23	0	0	9	0	0	12
$0 < \kappa <= 0$	0.43	0	2	0.49	0.29	5	0.31	0.31	11
$\kappa > 1$	4.95	6.08	10	3.79	1.25	3	4.84	5.74	4
Total	1.44	3.86	35	0.81	1.51	17	0.84	2.60	27

the experiment, we have gained a variety of insights into the way people process advice, its informativeness, and the market for advice.

To begin, there seems to be no broad consensus as to what types of experts are perceived to be the best. More precisely, in one part of our experiment we set up a market for advice given by subjects of different demographic characteristics which include gender, age, major, GPA, and how well the expert did (his rank in payoff terms) in his experimental trials. While economics students received the highest mean prices, followed by scientists, there was no commonly shared belief that these types of subjects made the best advisors. In fact, there was a tendency for students to bid more for advice from their own types. These results allowed us to reject what we called the "perception rent" hypothesis which implies that there will be a consensus amongst students as to who makes the best advisor.

In addition to analyzing the market for advice, we also looked at the impact of decision biases on the processing of advice. Here we ran a set of experiments to categorize subjects into three broad types; those that update information in a basically Bayesian manner, those that give too much weight to their prior (conservatives) and those that give too much weigh to the sample of observations they see (representative types). We find that these characteristics have an impact on how willing subjects are to pay for advice and follow it once received. More precisely, it appears that the more representative you are in your updating habits the more willing you are to pay for sample information and advice. In addition, people tend to assume that others are more representative than they are in their updating procedures, a fact that would make advisors worth listening to since their advice would encode more information about the sample they just saw. It ironically appears to be the case that representative types may be less willing to follow advice.

Finally, we attempted to measure the informational content of advice by observing how greatly subjects change their priors after receiving a piece of advice. From our calculations it appears that advice is generally thought to contain less information than the pure sample information upon which it is based. This is consistent with the bidding data, which showed that subjects bid significantly less for advice (and beliefs) than for raw data.

References

Becker, G.M., DeGroot, M.H., Marschak, J.: Measuring utility by a single-response sequential method. Behav Sci **9**, 226–232 (1964)

Kahneman, D., Tversky, A.: On the psychology of prediction. Psychol Rev **80**, 237–251 (1973)

Nyarko, Y., Schotter, A.: An experimental study of belief learning using elicited beliefs. Econometrica **70**, 971–1005 (2002)