LEARNING AND CONFIRMATION BIAS: HOW FIRST IMPRESSIONS AND AMBIGUOUS SIGNALS INFLUENCE PERCEPTIONS OF FINANCIAL ADVICE

Julie Agnew
Mason School of Business, College of William and Mary, Williamsburg VA, USA

Hazel Bateman
UNSW School of Business, UNSW Sydney, Sydney, NSW, Australia

Christine Eckert
UTS Business School, University of Technology Sydney, NSW, Australia

Fedor Iskhakov
School of Economics, Australian National University, Canberra, ACT, Australia

Jordan Louviere
University of South Australia, Adelaide, SA, Australia

Susan Thorp
The University of Sydney Business School, The University of Sydney, NSW Australia

Abstract
One puzzling observation in the financial services sector is that even though consumers in general tend to distrust financial advisers, and that advisers have been shown to provide poor quality service in a sizeable number of cases, clients usually trust their own financial adviser and believe that the advice they are being given is good. A possible explanation for this discrepancy is confirmation bias, where consumers interpret ambiguous information in agreement with existing beliefs or first impressions. We test this explanation on responses to an online choice experiment. Respondents evaluated two financial advisers who gave good and bad advice on four topics and who did or did not display a certification. Our empirical results show that the sequence of difficult or easy advice topics combined with good or bad advice as well as the adviser's certification impacts respondents' evaluations of advisers. The choices of around 60% of respondents exhibit significant confirmation bias that also influences their willingness to pay (WTP) for subsequent advice. Our results suggest that first impressions and ambiguous signals have far reaching consequences that are likely to be underestimated by traditional learning models and have important implications for any situation where people update their preferences.

We thank Karen Cong, Chenghan Hou and Simon Yin for their expert programming assistance, and Gail Bradford for her excellent project co-ordination.

INTRODUCTION

Financial advisers can make their clients' lives much easier: As experts in their domain, they can provide clients economies of scale in information acquisition. Clients increasingly rely on advice as more and more products such as credit cards, mortgages and investments require sophisticated financial skills. At the same time, studies from around the world overwhelmingly show that many, if not most, consumers do not have enough financial literacy to make sound decisions (Lusardi and Mitchell 2011). Whilst financial advisers can thus provide their clients with guidance and alleviate the anxiety often associated with financial decisions, not all the advice they deliver helps. Expert evaluation of financial advice given in a shadow shopping rated only 3% as good (Australian Securities and Investment Commission 2012). It also seems that people generally distrust financial advisers: in one recent survey on the image of professions, financial advisers score 24 out of 100 on the ethics and honesty dimension which places them in the top half of the list of least trusted professions (Morgan and Levine 2015). Despite the bad reputation of financial advisers in general, the same shadow study mentioned above also found that 86% of the clients who participated thought they had received good financial advice and 81% said they trusted their own adviser "a lot".

One possible explanation for this misalignment may be first impressions and confirmation bias, which is defined as "the seeking or interpreting of evidence in ways that are partial to existing beliefs, expectations, or a hypothesis in hand" (Nickerson 1998). The importance of making a good first impression is well known to businesses (Evans et al. 2000). Financial advisers can create a good impression by displaying professional certifications or using smart catering strategies (Agnew forthcoming). Such a catering strategy may consist of starting the relationship with sound advice on an easy topic: studies show that advisers can cause a client to trust their advice about complicated or ambiguous matters if they confirm the client's current views or opinions (Agnew et al. 2016; Mullainathan et al. 2012; Zaleskiewicz et al. 2016). Experimental work also shows that advisers who confirm a client's views on straightforward questions early in an advice relationship are subsequently rated as more trustworthy and competent than advisers who contradict a client's views. Furthermore, clients are more likely to accept their later advice on complicated topics (Agnew et al. 2016).

In addition to the importance of a first impression per se, once clients have formed beliefs about the quality of the adviser, confirmation bias may influence the clients' subsequent evaluations of the adviser. More specifically, confirmation bias will lead any subsequent advice on difficult topics to be viewed as good advice, thus reinforcing the positive image the adviser created at the first impression. Establishing trust early on can thus make a financial adviser appear exceptional and lead to subsequent fruitful business relationships. More trusted advisers are likely to be able to charge higher fees and thus take a larger share of the benefits of the advice relationship (Gennaioli et al. 2015).

Confirmation bias violates the basic assumption of learning models that consumers learn in a conventional Bayesian way (Eckstein et al. 1988; Erdem and Keane 1996); Roberts and Urban (1988). Confirmation bias explains how two people can reach opposite opinions after they review common evidence (Darley and Gross 1983). The defining feature of confirmation bias is that additional information leads to the polarization, rather than the moderation, of prior opinions. In contrast to learning models that allow for consumers to give higher weight to new signals from specific sources (Camacho et al. 2011), confirmation bias may actually reverse the interpretation of the signal. Irrespective of the actual signal valence, a consumer with confirmation bias will treat an ambiguous signal as positive if their prior belief is positive. Experimental studies have found evidence for confirmation bias and polarization on topics as diverse as the deterrent effect of the death penalty, nuclear power generation, climate change, brand loyalty and sexual morality (see Fryer et al. 2015, Table 1, for a summary). Rabin and Schrag (1999) introduce a theoretical model that shows that confirmation bias induces overconfidence so that people may come to believe with near certainty in a false hypothesis despite receiving an infinite amount of information.

Confirmation bias is often founded on a first impression (Beattie and Baron 1988). Earlier studies have acknowledged the importance of first impressions *per se* (e.g. Hartley 1995; Muthukrishnan and Chattopadhyay 2007), but to our knowledge no studies have attempted to model empirically the impact of confirmation bias on how consumers learn about new product attributes. Traditional learning models will underestimate the impact of first impressions on consumers who display the overconfidence that confirmation bias induces. Chylinski et al. (2012) address the related but different question of how consumers resolve uncertainty over whether a product attribute is likely to help them achieve a utility goal. Their study uses an associative learning model (Lawson 1997) to explore how consumers form associations between *observed* attributes and compute the conditional probability that a goal will be achieved. While they allow for consumers to form prior beliefs and to use biased information processing (in the assimilation of disconfirming information) the point of interest in their study is the presence or absence of a product attribute. Here we study how people form perceptions of the "product attributes" themselves, in this case, perceptions of the quality of a financial adviser.

We designed and implemented a new incentivized online choice experiment using videos of advisers to test this explanation.[1] We showed the respondents video advice from two different advisers on four different topics. For each topic, one adviser gave good advice and the other gave bad advice. After both advisers had provided their advice for a topic, respondents had to choose the advice they would follow. Our experiment extends the work of Agnew et al. (2016) in two important ways. Agnew et al. showed that a client's first impression of an adviser affected later decisions about whether to follow their advice. Here, we choose specific sequences of advice that allow us to determine if clients' decisions to continue to follow advice is caused by confirmation bias. Agnew et al also related clients' evaluations of an adviser's trustworthiness and expertise to the client's first impressions of the adviser. In our experiment, after advice on all four topics is been given, we ask clients whether they would be willing to pay a certain fee to receive more advice from either of the advisers. Responses to these questions let us compute the monetary value of clients' good (or bad) opinions. We formed the experiment so as to limit the information respondents could use to update their beliefs, which simplified the model (see Ching et al. 2013 for a discussion of papers implementing different information sources).

To test whether clients exhibit confirmation bias in their choices of, and willingness to pay for, financial advice, we consider two ways that people can use ambiguous information about adviser quality: first, consumers can act "rationally", ignore ambiguous signals, and not use them to revise their opinion of the adviser, which is the common underlying assumption of traditional models of consumer learning (Eckstein et al. 1988; Erdem and Keane 1996; Roberts and Urban 1988). Second, clients can use an iterative Bayesian updating process with limited memory where they interpret ambiguous signals in favor of their current evaluation of the adviser (Fryer, Harms, and Jackson 2015). We treat advice on topics that are difficult to understand, or "hard topics", as ambiguous signals. When clients ignore ambiguous information their final evaluation of the adviser is independent of the sequence in which the advice was

[1] The experiments used videos made and tested in Agnew et al. (2016).

given. But if clients treat ambiguous signals as evidence in favor of their current prior belief, their final judgment will depend not only on the objective quality of advice, but also on the sequence of advice signals, in this case, the interaction between advice quality and the perceived difficulty, or ambiguity, of the advice topic.

We use a latent mixture choice model to test which of these approaches to ambiguous advice signals clients actually use. More specifically, we assume that the utility of choosing an adviser is also influenced by the client's opinion on whether the adviser is good. This opinion depends on initial prior beliefs about the quality of the adviser based on characteristics unrelated to their advice (such as the whether they display a creditable professional qualification), as well on the way the client treats ambiguous advice, and the sequence of advice quality and signal ambiguity. We use the information obtained from both the choice of advice data as well as the willingness to pay question to classify respondents into those likely to update "rationally" versus those likely to update with confirmation bias (which we call "Fryer updating" hereafter) and link this classification to observed characteristics of the respondents. We show that the importance of first impressions is underestimated if confirmation bias is not taken into consideration and that the true impact of a first impression might only show up after several more signals have been received by the client.

Our results provide important insights both for policy makers who want to protect and educate vulnerable advice clients and for financial service providers who want to measure not only the immediate but also the long-term monetary impact of an adviser who makes a good or bad impression. Given that the need to update beliefs formed with incomplete information is common to many situations and that the ubiquity of confirmation bias and its striking implications is a challenge so far not dealt with by traditional learning models, our paper also calls for a revision of standard assumptions about how people treat new, ambiguous, signals.

The article is structured as follows: we first briefly review the market for financial advice and show why we study confirmation bias in an advice setting. We then describe the standard Bayesian updating model and contrast it to the updating model with limited memory by Fryer et al. (2015). Next we describe our experimental design and propose an empirical model that allows us to use observed choices to identify whether and to what extent consumers are prone to confirmation bias. We present the results and conclude with a discussion of our model's limitations and directions for future research.

THE MARKET FOR FINANCIAL ADVICE

"Do-it-yourself" finance is the term Ryan et al. (2011) use to describe the increased responsibility consumers have been given for financial decisions. Consumers confront difficult choices over new and more complicated products for credit, mortgages, investments and retirement plans. However, there is overwhelming evidence that people frequently make poor financial decisions (Campbell et al. 2011). Reasons for these errors include low levels of financial literacy (Lusardi and Mitchell 2011), issues of trust in markets and financial products (Christelis et al. 2010), behavioral biases (Thaler and Benartzi 2004) and limited cognition (Lusardi and Mitchell 2006). Opportunities to learn from experience in financial contexts are limited because many consequential decisions, such as the choice of a mortgage or retirement account investment, are made infrequently, and feedback from outcomes is often delayed. Financial firms may have incentives to strategically raise the complexity of products to impede consumer learning and maintain rents (Carlin 2009; Carlin et al. 2010). Financial choices, and the mistakes that often follow, have serious implications for financial welfare and stability, individually and in aggregate (Agarwal et al. 2009; Bar-Gill and Warren 2008; Campbell 2006; Campbell et al. 2011). Consumers can delegate difficult decisions to financial advisers to compensate for low financial literacy and lack of expertise. Hackethal et al. (2012) point out that, in theory, the use of financial advisers can ameliorate a lack of financial literacy. Advisers can provide clients economies of scale in information acquisition and guide them to superior practice.

However theory predicts that an adviser's willingness to de-bias and educate clients can be diluted by incentive structures (Inderst and Ottaviani 2009). Empirical studies likewise show that advisers can exploit the biases of clients (Hackethal et al. 2012; Mullainathan et al. 2012). In a similar vein, an audit study by the Australian regulator (Australian Securities and Investment Commission 2012) concludes that only 3% of the advice given to auditors was objectively good quality while the majority of the advice (58%) was adequate and the remainder was poor. Despite these low evaluations, most participants (86%) rank the quality of the advice they receive as high. In addition, 81% trust the advice they receive from their adviser 'a lot'. In a recent study, Agnew et al. (2016) illustrate how much first impressions in the client-adviser relationship matter, which complements research that shows that clients form opinions of their financial adviser rapidly (Yaniv and Kleinberger 2000). Together, these findings explain how some advisers can successfully use strategies to build and maintain client trust while also providing unhelpful advice (Anagol et al. 2013; Mullainathan et al. 2012).

Despite these hazards, citizens of many countries use advice services. For example, Chater et al. (2010) report that 58% of individuals' stock purchases were influenced by an adviser in a survey of 6,000 consumers across eight EU countries. Holden et al. (2013) find that individuals choose to work with advisers because the adviser has expertise in an area that the consumer does not. Other studies emphasize that the personal qualities of advisers matter; for example, clients must decide if an adviser is trustworthy and competent before acting on advice. Georgarakos and Inderst (2014) show that clients with limited financial capability are more likely to follow advice if they trust their adviser, but trust depends on many factors, including the client's capability, the accuracy and quality of information provided, and a belief that the adviser and client's incentives are aligned (Sniezek and Van Swol 2001; Yaniv and Kleinberger 2000). In summary, consumers cannot always rely on financial advisers to provide the best advice possible, but have to form beliefs about advisers' quality based on the signals they receive from them. Prior beliefs about advisers' quality will depend on client and adviser characteristics; these beliefs could also influence a client's interpretation of subsequent signals from the advisers.

A SIMPLE MODEL OF CONFIRMATION BIAS

Several studies show that advisers build trust early and rapidly in their relationships with clients, and that confirmation bias can assist in creating trust. Psychological studies show that, in general, people form impressions of the trustworthiness of a person or entity early and quickly, and these early impressions have a "cascading" effect on interpretations of subsequent interactions (Holtz 2015; Reinhard et al. 2006). A good first impression can combine with confirmation bias to make the client overconfident in their favored (subjective) views (Rabin and Schrag 1999). In financial advice relationships, clients are thus more willing to trust in, and act on, advice that confirms their own opinions (Agnew et al. 2016; Zaleskiewicz et al. 2016).

While confirmation bias and the polarization of opinion that follows is at odds with standard Bayesian updating models, it can be explained by updating models with limited memory (Fryer et al. 2015). Consider a rational Bayesian decision maker who may or may not be able to discern the quality of a financial adviser. Assume that the client forms an expectation over two states of the world: A (the adviser is good) and B (the adviser is bad). The client holds an initial prior (or starting) belief that $P(A)=\lambda_0$, which he or she updates as they receive a sequence of clear or ambiguous signals. We can interpret this starting belief as the client's first impression of the adviser. The clear good (a) or bad (b) signal of the adviser's quality might be a correct or incorrect recommendation from the financial adviser that either agrees with or contradicts what a client thinks is factual or sound. The client uses clear signals to update their prior belief and form a posterior expectation of the quality of the adviser. Let $s>1/2$ denote the probability of receiving a clear, good signal conditional on the advisor being good, $P(a \mid A)=s$, and assume that the probability of receiving a good clear signal from a bad advisor is $P(a \mid B)=1-s$. Then, the parameter s can be interpreted as signal strength.

However, the client may also receive an ambiguous signal ab. An ambiguous signal might be a recommendation on a subject where the client is inexperienced or uninformed. Ambiguous signals create an opportunity for confirmation bias to operate. A rational Bayesian updater ignores ambiguous signals and forms a posterior over the sequence of clear signals. He or she thus gradually uncovers the true state of the world. However, when the decision maker has a confirmation bias, he or she will not overlook an ambiguous signal. Rather the client will interpret the ambiguous signal in line with their current belief, either as a or b, and thus reinforce their prior view of the adviser's quality. Fryer et al. show that limited memory – the need to form a posterior belief at each signal rather than wait to the end of the sequence - forces an interpretation of ambiguous signals and generates confirmation bias and polarization of opinions.

More formally, in the rational model, beliefs are updated according to

$$(1)\ \lambda_{t+1} = P(A \mid \lambda_t, \sigma_{t+1}) = \begin{cases} \dfrac{s\lambda_t}{s\lambda_t + (1-s)(1-\lambda_t)}, & \text{if } \sigma_{t+1} = a, \\ \dfrac{(1-s)\lambda_t}{(1-s)\lambda_t + s(1-\lambda_t)}, & \text{if } \sigma_{t+1} = b, \\ \lambda_t, & \text{if } \sigma_{t+1} = ab, \end{cases}$$

where σ_t is the advice received, that is, the signal. When the Fryer updater holds a prior belief that the adviser is good quality, he or she interprets an ambiguous signal as a good signal and in the reverse way when he believes the adviser is poor quality. In the Fryer updating model, the beliefs are updated according to:

$$(2)\ \lambda_{t+1} = P(A \mid \lambda_t, \sigma_{t+1}) = \begin{cases} \dfrac{s\lambda_t}{s\lambda_t + (1-s)(1-\lambda_t)}, & \text{if } \sigma_{t+1} = a, \text{or } \sigma_{t+1} = ab \text{ and } \lambda_t > \dfrac{1}{2}, \\ \dfrac{(1-s)\lambda_t}{(1-s)\lambda_t + s(1-\lambda_t)}, & \text{if } \sigma_{t+1} = b, \text{or } \sigma_{t+1} = ab \text{ and } \lambda_t < \dfrac{1}{2}, \\ \lambda_t, & \text{if } \lambda_t = \dfrac{1}{2}. \end{cases}$$

In our empirical application, discussed in more detail below, clients could perceive advice as ambiguous or clear, and as good or bad. The experimental design included four advice topics and six sequences of good/bad advice combinations, totaling 96 possible ways clients could revise their beliefs about an adviser by either rational or Fryer updating. Figure 1 illustrates all possible paths of beliefs for the respondents with different initial priors and updating strategies (in rows). As well as initial prior beliefs and updating strategy, the paths depend on how many advice topics were perceived by the clients as clear (in columns, ranging from all topics perceived as clear to all topics perceived as ambiguous) The size of the dots reflects the number of paths that pass through the respective point. Signal strength s is arbitrarily set to 0.75. Rows 1, 3 and 4 reflect Fryer updating with starting priors of 0.6, 0.5, and 0.4, respectively. Row 2 is based on rational updating with starting prior 0.5.

INSERT **Figure 1** ABOUT HERE

Consider, for example, the first column of graphs, where we assume that all topics are clear to the client. The only difference in path direction thus comes from differences in the sequence of good and bad advice that one adviser gives the client over four topics. If all topics are clear, and if all clients start with the same starting prior, the rational and Fryer updating lead to the same posterior beliefs (see rows 2 and 3). When moving to the next column of graphs, we can see that if one topic is ambiguous, Fryer updating leads to polarization that shows up in deviations from the rationally updated priors. The polarization becomes more pronounced when clients perceive more topics to be ambiguous (columns further right), and becomes extreme when all four topics are perceived as ambiguous: With Fryer updating, clients reach an almost certain belief that the adviser is good if they start with an initial prior larger than 0.5 (row 1) and reach an almost certain belief that the adviser is bad if they start with a prior smaller than 0.5

(row 4). Rational clients do not update their priors if all signals are ambiguous (row 2) or when their starting prior is equal to 0.5 (row 3).

EXPERIMENTAL DESIGN

We designed our experiment to achieve two goals. First, we aimed to collect data to estimate and compare a model of client confirmation bias (Fryer) with a model of rational (standard Bayesian) choice. Second, we aimed to measure the willingness to pay for advice of respondents whose choices are best matched by either the biased or rational model. Since clients who lack experience or financial literacy are probably more susceptible to manipulation, we also collect an array of demographics, preferences, financial capability measures and psychological inventories to help identify these clients.

We fielded a four-part online survey that included an incentivized choice experiment in December 2014.[2] While similar to the experiments found in Agnew et al. (2016), these new experiments captured an important variable, willingness to pay, which enables us to measure the impact of confirmation bias in dollar terms. They also incorporated different signal sequences from Agnew et al. (2016) as defined by the order of the advice quality (good or bad) and topic difficulty (clarity or ambiguity) of the advice given. Members of a nationally representative online panel were invited to complete the survey. Those who responded to the invitation had to pass two screening questions to meet age and gender quotas. This resulted in 2,003 "clients". To ensure incentive compatibility, we compensated participants who completed the experiment for their time and rewarded them if they chose correct advice in each choice set and in a post experiment quiz. Respondents first answered a set of questions that measured their general financial literacy and numeracy (Lusardi and Mitchell 2011; Lipkus et al. 2001), and that evaluated their understanding and experience of the subjects of the four advice topics covered by advisers in the subsequent discrete choice experiment (DCE).

[2] The survey offered participants who completed all questions a small compensation for their time (around $4) and one entry in a prize draw for $A50 for each correct choice of advice. View the full survey at http://survey.us.confirmit.com/wix/p3070864270.aspx

Our DCE offers "clients" a sequence of videos of advisers who give financial advice on four common and important consumer finance topics.[3] The topics include credit card debt repayment, retirement savings account consolidation, diversification in equity investment and index fund fees. Table 1 records the scripts for the good and bad advice for each topic.

INSERT Table 1 ABOUT HERE

To identify the effect of confirmation bias in the DCE, we need advice topics, advisers, the environment and the mode of advice delivery to be uniform. Agnew et al. (2016, section 3.3 and supplementary materials) reports in detail how the videos were made and tested to achieve maximum uniformity with minimum extraneous variation, and we do not repeat that here. In this section we outline the experiment and explain how our experimental design differs from this earlier research. Figure 2 shows the advisers from the videos and their "names". The videos allow us to control the two advisers shown to each participant, the order of advice topics, the quality of advice given by each adviser for each topic and the attributes of advisers giving advice. Notably, we extend the advice experiments of Agnew et al. (2016) over a wider range of signal sequences to allow a test of confirmation bias.

INSERT ABOUT HERE

In the DCE videos, two advisers give a recommendation on each of the four topics: each respondent received four pairs of advice; the two advisers were the same across the four advice topics for each respondent; and in each case one adviser provided a correct recommendation, while the other provided an incorrect recommendation. Correct and incorrect advice hereafter is termed "good" and "bad". The videos systematically varied adviser factors - the adviser's gender (2 options: male or female) and age (2 options: young or old), professional certification (2 options: certification presented or not) - the order of the advice topics (4 options: first, second, third, or fourth) and the quality of the advice (2 options: correct or incorrect).

The experiment used a between-subjects design. As noted, the advice viewed by any one respondent is provided by the same two advisers; hence, variation in adviser factors (age, gender, certification) is a between-subjects manipulation. To minimize the between-subjects treatment groups we used a fold-over design in which we created the 2^3 complete factorial of possible advisers and paired each of them with their "mirror image" (that is, the exact opposite level, so

[3] The topics were selected by Agnew et al. (2016) based on their relevance for people around the world, that they had unequivocally right and wrong answers and were based on the mistakes often made in these areas.

that a younger woman adviser was matched with an older male adviser). This produced pairs of advisers who were orthogonal in the differences in factor levels. The resulting design is optimally efficient under the assumption that a conditional multinomial logit choice model underlies the respondent choices (Street et al. 2005; Street and Burgess 2007). This design approach produced four between-subject treatment groups and is shown in panel A of Table 2.

Further variation in the DCE relates to between-subject manipulation of a) topic sequence and b) order in which good and bad advice is given by each adviser. Variation in these orders is essential to test hypotheses about formation of persistent respondent preferences for advisers. The fold-over design used to create the between-subjects manipulations ensures variation in quality of advice. We also maximized variation in adviser attributes by ensuring that both financial advisers gave advice on the same topic in each pair. Thus, we combined the between-subject treatment groups (4) with a design to vary the orders of topics (4 levels) and good and bad advice (2 levels). A full factorial design would have required a very complex survey program and a very large sample, since it implies 16 possible sequences of good (G) and bad (B) advice and 24 possible sequences of topics. To maximize variation and to enable a test of confirmation bias we used six sequences of good and bad advice orders – those where each adviser gives two good and two bad recommendations (see panel C of Table 2) - and topic sequences with an equal number of hard and easy financial topics (see panel B of Table 2).[4]

When we combined the four possible pairs of advisers with the six possible sequences of topics and the six possible sequences of advice quality, we obtained a design with 6*6*4 =144 conditions. We randomly assigned at least 10 and up to 14 participants to each condition. After the DCE task, respondents rated the trustworthiness, competence, attractiveness, understanding, professionalism, financial expertise, genuineness and persuasiveness of the advisers they saw, and stated their willingness to pay $X for a one-hour session with both, one or none of the advisers. Fixed values $X \in \{50, 100, 150, 250, 500, 750\}$ were assigned so as to minimize their predictability from the other manipulated characteristics of the experiment condition.

[4] We rely on the results by Agnew et al (2016) who find that two topics (debt repayment and retirement account consolidation) are relatively easy (E) and that the other two topics (diversification and index fund fees) are relatively hard (H)

After they had answered questions about demographics (e.g., marital status, household size and number of dependents, education, labor market status, income, gross assets and debts/liabilities) and personal characteristics, including personality traits and risk attitudes, respondents read debriefing information that explained correct advice. The survey closed with an incentivized quiz on the debriefing material. Table 3 compares the characteristics of the sample with the Australian Census data from 2011. Our sample reports slightly higher educational attainments and a higher probability of being married than the census shows, but otherwise are representative of the population.

JOINT MODEL OF CONFIRMATION BIAS AND CHOICES

Model Description

The DCE described above captures two pieces of information that help to identify confirmation bias: first, a client's posterior belief about an adviser at each choice set determines the advice they choose when the advice topic is ambiguous; and second, the client's final posterior belief about an adviser determines, in combination with the price of future advice and some client characteristics, whether clients are willing to pay for an adviser. Next, we provide a mathematical description of these links.

We first review how the posterior belief, that is, the belief about the adviser's quality, is formed. Equations (1) and (2) describe how clients update their beliefs depending on whether they are Bayesian rational decision makers or whether they are prone to confirmation bias, respectively. Which updating scheme a particular respondent uses is not known by the researcher and needs to be inferred from the choices the respondent makes. Similarly, the researcher does not know whether a signal is clear or ambiguous to the respondent and must make inferences about this from the choice data. In both updating schemes, the posterior belief (or updated prior) does however depend on the initial prior belief λ_{k0}^{r} (or starting prior) about adviser $r \in \{R, L\}$ of the respondent *k=1,…, K*, which in turn is a function of adviser and respondent characteristics X_0 (see Table 4 for a description of the variables used) and an unknown vector of parameters β_0. We set the starting prior to be

(3) $$\lambda_{k0}^{r} = \frac{\exp(\beta_0 X_0)}{1 + \exp(\beta_0 X_0)}$$

When combined with a value for signal strength, s, we can calculate λ_{kj}^{r}, the updated prior about adviser $r \in \{R, L\}$ of the respondent k after choice set j=1,.., 4, conditional on the respondent's updating scheme and signal clarity based on Equations (1) and (2).[5]

We infer which type of updating scheme the respondent uses and which topics are ambiguous or clear to them. Recall that in our DCE, a respondent k makes four choices of advice delivered by the same pair of advisers L and R. As discussed above, the experiment is designed in such a way that the characteristics of the advisers, namely gender, age and display of certification, are a mirror image of each other. For example, if adviser L is an older male with certification, adviser R is then a younger female without certification. Moreover, in each topic one adviser gives good advice, and another one gives bad advice. Also, if the respondent chooses adviser L, she necessarily does not choose adviser R. This means we can use effects coding, where for example we assign a value of 1 if adviser R possesses a particular characteristic (and adviser L does not), and a value of -1 if adviser R does not have the particular characteristic (and adviser L does).

If the respondent perceives the topic to be easy, and consequently a clear signal of adviser quality, we assume that the respondent will choose the adviser who gives correct advice. That is, if the topic is easy, the respondent chooses the advice based on its quality alone. However, if the respondent thinks the topic is hard, the respondent perceives the advice to be an ambiguous signal. In this case, we assume that the respondent will choose the advice according to their posterior belief about the adviser.

We code the choice data as

$$(4) \qquad y_{kj} = \begin{cases} 1, & \text{if advisor } R \text{ was chosen at choice } j \text{ by respondent } k, \\ 0, & \text{if advisor } L \text{ was chosen at choice } j \text{ by respondent } k, \end{cases}$$

and define q_k^j as equal to *1* if the quality of advice adviser R gives to respondent k in choice set j is good, -1 otherwise. If we assume that respondents still make some errors when making their choices and that this error is extreme value distributed with scale $1/\beta_1$ and $1/\sigma$, respectively, we obtain

[5] To enable us to identify parameters, we set s to an arbitrary value greater than 0.5 and check sensitivity of estimation to alternative choices. The results we report below use $s=0.75$.

(5) $$P(y_{kj}=1\,|\,\text{topic=clear})=\frac{\exp(\beta_1 q_k^j)}{1+\exp(\beta_1 q_k^j)}$$

and

(6) $$P(y_{kj}=1\,|\,\text{topic=ambiguous})=\frac{\exp(\sigma\cdot(\lambda_{kj}^R-\lambda_{kj}^L))}{1+\exp(\sigma\cdot(\lambda_{kj}^R-\lambda_{kj}^L))}.$$

Thus both β_1 and σ are scale parameters: as σ (β_1) approaches infinity, the expression in the right hand side approaches 1 for $\lambda_{kj}^R > \lambda_{kj}^L$ (q_k^j =1), and 0 otherwise.

Next we turn to the respondent's willingness to pay. The respondent can choose to pay for additional advice from both, one, or none of the advisers. We model this choice as follows:

(7) $$P(\text{willing to pay for advisor } r)=\frac{\exp(\beta_2^0+\beta_2^1\lambda_{k4}^r+X_2\beta_2^2)}{1+\exp(\beta_2^0+\beta_2^1\lambda_{k4}^r+X_2\beta_2^2)},$$

where X_2 are attributes of the respondent and the adviser, including the price of future advice, β_2^0 is a constant, β_2^1 captures the impact of the posterior on the willingness to pay, and β_2^2 is a vector of unknown parameters.

In this specification, the respondent's choice of advice for each of the four topics at a given price and thus his or her willingness to pay for future advice are functions of posterior beliefs about the adviser. This posterior belief is a function of the way the respondent updates his or her beliefs (either rationally or with confirmation bias) and of which signals the respondent perceives to be clear (which topics are easy). Our model assigns respondents to latent classes distinguished by clarity of topic and updating scheme. In the interest of parsimony, we assume that the probability that a respondent is an updating type and the probability that a respondent treats any topic as clear are independent, conditioning on the characteristics of the individual respondent, so that[6]:

(8) $$P_k(\tau)=P_k(\tau_{\text{clarity}},\tau_{\text{rationality}})=P_k(\tau_{\text{clarity}})P_k(\tau_{\text{rationality}}),$$

and

$P_k(\tau_{\text{clarity}})=P_k(\text{topic}_1=\text{clear})P_k(\text{topic}_2=\text{clear})P_k(\text{topic}_3=\text{clear})P_k(\text{topic}_4=\text{clear})$, where topic_1 refers to the same topic for all respondents rather than to the first topic seen and so on.

[6] Note that without any further assumptions there exist two types of updaters as well as a classification of clear or ambiguous for each topic, which results in 2^5=32 different combinations.

Dependence between the latent classes for any respondent k is captured by allowing class membership probabilities to depend on respondent-specific covariates X_4 and X_5 and associated parameter vectors β_4, β_5, and topic specific constants β_5^m:

(9) $$P_k(\tau_{rational}) = \frac{\exp(\beta_4 X_4)}{1+\exp(\beta_4 X_4)},$$

and

(10) $$P_k(\text{topic}_\text{m} = \text{clear}) = \frac{\exp(\beta_5^m + \beta_5 X_5)}{1+\exp(\beta_5^m + \beta_5 X_5)}.$$

We estimate the parameters $\theta = \{\beta_0, \beta_1, \beta_2, \beta_4, \beta_5, \sigma\}$ by jointly maximizing the likelihood of choices and willingness to pay decisions. Conditional on the respondent belonging to one of the 32 clarity-rationality classes τ_t, $t=1,\ldots,32$, the likelihood of respondent k's sequence of choices is

(11)

$$l_k(\theta \mid \tau_t) = \prod_j^4 P(y_{kj}=1 \mid \tau_t)^{y_{kj}=1} P(\text{willing to pay for advisor } R \mid \tau_t)^{y_k^R=1} P(\text{willing to pay for advisor } L \mid \tau_t)^{y_k^L=1},$$

with y_k^L (y_k^R) indicator variables, taking the value 1 if the left (right) adviser was chosen.

The unconditional likelihood of respondent k's sequence of choices is thus:

(12) $$l_k(\theta) = \sum_{t=1}^{32} P_k(\tau_t) l_k(\theta \mid \tau_t).$$

Parameter identification

A formal analysis of identification is not feasible for the complex, non-linear learning model discussed above (see also the discussion in Ching et al. 2013). In the following we sketch our identification strategy for the key model parameters.

First, consider the initial prior belief about the adviser, which is the starting prior λ_{k0}^r. The starting prior belief is the basis for the updated posterior belief and thus influences both the choices of advice as well as willingness to pay for future advice. The starting prior belief itself also influences directly the choices made in choice set 1, as in this set we make the assumption that (up to uncertainty) the adviser with the higher initial prior is chosen if the topic is ambiguous. Since the design of the experiment ensured that respondents face both easy as well as hard (ambiguous) topics in choice set 1 (refer to panel B in Table 2), we thus obtain sufficient information to estimate the starting prior as well as how it depends on both adviser and respondent characteristics.

Next we discuss the signal strength *s*, that is, the probability that a good (bad) signal comes from a good (bad) advisor. We set *s=0.75* to allow the probability that a good adviser delivers good advice to be greater than 0.5 but less than one to ensure that updating can occur. We tested for the sensitivity of results at *s=0.60* and results remained largely unchanged.

The parameter σ is in turn identified via the starting prior λ_{k0}^{r} and *s*. These two parameters jointly define the updated beliefs and can thus be considered as pre-determined covariates when respondents face an ambiguous topic. Choices made over ambiguous topics can thus identify σ.

Choices made over the four different topics allow us to identify the latent clarity classes. More specifically, our assumption about the choice process can (up to uncertainty in the choice process) be summarized as follows. If we observe that the respondent chooses the adviser who gives bad advice, we can conclude that the topic is hard for that respondent. We cannot make that inference if the respondent chooses the adviser who gives good advice as this could imply either that the topic was easy for that respondent or that the adviser was chosen because of a higher posterior belief. The combined information of updated prior beliefs about the advisers and incorrect choices of advice thus allows us to identify the clarity classes.

Since the starting prior belief about the adviser can be inferred from the data without any assumptions about the way the respondents update their beliefs and since signal strength *s* is fixed, we can calculate the posteriors for both updating schemes. The posterior associated with the higher likelihood then helps to pin down rationality classes.

Discontinuity of the likelihood function and estimation method

The estimation of our model is complicated by the fact that the likelihood function is discontinuous for those cases where respondents update their beliefs according to the Fryer updating scheme. The kinks in the likelihood appear along the dimensions of the parameters of prior beliefs. Even in the simplest case when the starting prior belief is represented by a single constant (as illustrated in), as this constant moves from zero to one and crosses particular thresholds dependent on other parameters, the values and the counts of the possible updated prior beliefs change discontinuously.

Figure 3 ABOUT HERE

Figure 3 compares the updated prior beliefs λ_4 (after all choices are made) under the rational and Fryer updating for three values of the signal strength parameter *s* as the starting prior λ_0 changes from zero to one. For each value of the starting prior λ_0 we draw all values of posterior beliefs that are possible in the model, with the size of the circles indicating the number of theoretical paths leading to that belief.

The discontinuities in the likelihood function cause implementation problems due to the inherent computational difficulty for maximum likelihood estimators (see also Chernozhukov and Hong 2004). We overcome these difficulties with Bayesian estimation methods. More specifically, we use Sequential Adaptive Bayesian Learning (SABL) proposed by Durham and Geweke (2014). SABL is an extension of sequential Monte Carlo methods that additionally exploits the benefits of parallel computing environments. SABL does not require the modeler to specify conjugate priors and it is also robust to multimodal posteriors which can arise in high dimensional problems (Jasra et al. 2007) such as ours. When used for Bayesian inference, SABL is a posterior simulator. Our interest only lies in the latter and we thus focus the following basic description of SABL on this whilst at the same time ignoring the aspects that make SABL an efficient tool to address very complex problems.

As with any Bayesian estimation approach SABL requires the user to specify the likelihood function $l(\theta)$ as well as prior distributions $p^{(0)}(\theta)$ for the parameters to be estimated. SABL then produces draws from the posterior $p^{*}(\theta)$ as follows:

- Draw parameters from the prior distributions. To do this SABL represents initial information by $\theta_{gh}^{(0)} \sim_{iid} p^{(0)}(\theta)$, organized into H groups[7] of G draws each (SABL defaults to H=16 and G=192). Let $p^{(0)}(\theta)$ be a very flat version of the likelihood function, i.e. $p^{(0)}(\theta) = l(\theta)^{r_0}$ with r_0 very small.

- For a sequence of cycles *n=1, 2,..*

  a. Correction (C) phase: Determine $p^{(n)}(\theta)$ by raising the likelihood function to a higher power, i.e. $p^{(n)}(\theta) = l(\theta)^{r_n}$ with $r_n > r_{n-1}$. Calculate for each draw a weight $w^{(n)}(\theta_{gh}^{(n-1)}) = p^{(n)}(\theta_{gh}^{(n-1)}) / p^{(n-1)}(\theta_{gh}^{(n-1)})$, *h=1,…,H, g=1,…,G*;

  b. Selection (S) phase, applied independently to each group *h=1,..,H*: Use multinomial residual resampling (e.g. Douc and Cappé 2005) based on $\{w^{(n)}(\theta_{gh}), g = 1,..,G\}$ to select $\{\theta_{gh}^{(n,0)}, g = 1,...,G\}$ out of $\{\theta_{gh}^{(n-1)}, i = 1,...,I\}$.

  c. Mutation (M) phase, applied independently to each group *h=1,..,H*: The M phase is a Metropolis random walk. In each step *o (o>0)* of the random walk obtain for each *g=1,…,G* a proposal $\theta_{gh}^{(n,o)*}$ is drawn from $N(\theta_{gh}^{(n,o-1)}, \Sigma^{(n,o-1)})$, where $\Sigma^{(n,o-1)}$ is proportional to the sample variance of the particles $\{\theta_{gh}^{(n,o)}, g = 1,...,G\}$. Accept $\theta_{gh}^{(n,o)*}$ with probability α where α is defined as $\alpha = \min\{1, p^{(0)}(\theta_{gh}^{(n,o)*}) l(\theta_{gh}^{(n,o)*}) / p^{(0)}(\theta_{gh}^{(n,o-1)}) l(\theta_{gh}^{(n,o-1)})\}$ and set $\theta_{gh}^{(n,o)} = \theta_{gh}^{(n,o)*}$, otherwise set $\theta_{gh}^{(n,o)} = \theta_{gh}^{(n,o-1)}$. The proportionality factor is thereby increased when the rate of accepting the proposal draws is higher than a particular threshold (the default in SABL is 0.25), and decreases otherwise. The random walk terminates once the dependence among particles has been sufficiently broken, that is when the particles are sufficiently independent (note that the S-phase introduces dependence via

---

[7] SABL organizes the draws into groups to exploit the parallel processing possibilities of the algorithm.

repeated sampling of the same $\theta_{gh}^{(n-1)}$). SABL assumes sufficient independence of the particles when the variance (calculated across the H group means) falls below a certain threshold. The last set of $\theta_{gh}^{(n,o)}$ is then denoted as $\theta_{gh}^{(n)}$.

- If $p^{(n)}(\theta) = l(\theta)$ then $N=n$ and the algorithm terminates with draws $\theta_{gh}^{(N)}$ from the posterior distribution $p^{*}(\theta)$.

We chose uninformative priors. We assumed that the prior for each parameter of interest is independent normal with a mean of zero and a standard deviation of five. We evaluated the sensitivity of prior influence by a careful visual examination of the posterior distribution against the prior distribution.

The advantage of using SABL (or a Bayesian approach in general) is that the posterior distribution of draws can help in assessing the identification of the model parameters (see also discussion in the previous section). More specifically, a high correlation between the posterior draws of two parameters may suggest that these are not separately identified by the choice data. In addition to including different covariates in the different model parts (see also Table 4), we used this correlation matrix check to further assess the identification of our model.[8]

EMPIRICAL RESULTS

In aggregate, respondents chose correct recommendations in 79% of the time. The percentages of correct choices by topic are 86% for retirement account consolidation, 88% for credit card debt repayment, 79% for stock diversification and 64% for index fund fees. Respondents chose the advice of the young female adviser two percentage points more often than the advice of the older male who appears alongside her in the experiment. Respondents chose the younger male and older female equally often.

We estimated our model using data from 1,903 of the 2,003 respondents and held back the remaining responses to assess hold out fit. In-sample fit was satisfactory and hold out fit did

[8] We run SABL using its MATLAB interface. SABL itself can be downloaded from http://www.quantosanalytics.org/garland/mp-sps_1.1.zip. The time to estimate our model using SABL is approximately 60 minutes.

not deviate very much from in-sample fit, which shows that our model does not over-fit the data. The model predicted an average (over all choice sets) probability of 0.69 for the estimation sample that the adviser who was in fact chosen would be chosen, and predicted a related probability of 0.69 for the hold out sample. When the adviser was not chosen, the predicted choice probability decreased to 0.29 for the estimation sample whereas it decreased to 0.28 for the hold-out sample. The predicted probabilities were less discriminating in the willingness to pay choice probabilities. When the respondent actually chose to pay the adviser, the model's average predicted probability was 0.48 for the estimation sample data and 0.44 for the hold out data. When the respondent chose not to pay the adviser, the average predicted probability of being paid was 0.28 for the estimation sample data and 0.34 for the hold out sample data. Thus, the model underestimates the probability that a respondent is willing to pay the proposed fee for the adviser.

Table 5**: Empirical Results** reports the model estimates. For each parameter we report the mode of its posterior distribution as well as the 2.5 and 97.5 percentiles of this distribution, that is, the corresponding equi-tailed credible interval (CI). There is a 95% probability that the parameter is not zero if zero does not fall in the CI. Next we discuss each of the model components.

Prior belief about adviser

We allow the starting prior belief about an adviser's quality to depend on the trust that a respondent has in financial advisers (Gennaioli et al. 2015) and also whether the adviser displayed a professional qualification certification (Agnew et al. 2016). Both factors have been shown by earlier studies to influence whether people will take financial advice. The mode of the distribution for the trust parameter equals 0.520 and the 95% credible interval does not contain zero. This shows that respondents who rate financial advisers as trustworthy hold a higher prior belief that the adviser is good as we would expect. The mode of the distribution for the non-certification parameter equals -.085 and the 95% CI again does not contain zero. Based on the posterior draws of the parameters, we can infer the distribution of the difference in prior beliefs for certified versus uncertified advisers. From the perspective of respondents who trust financial advisers already, this posterior distribution has a mean of 0.016 with an associated 95% CI of [0.001, 0.036]. For respondents who generally distrust financial advisers, the impact of certification is double: this posterior distribution has a mean of 0.032 with a 95% CI of [0.003, 0.073]. The mean of the posterior for respondents who are neutral about financial advisers is 0.023 [0.002, 0.054]. So we conclude that if an adviser displays a professional certification, respondents form a significantly higher prior belief that he or she will give good advice and this higher expectation will be subsequently reflected in higher choice probabilities in the case of ambiguous topics and higher willingness to pay for additional advice from this adviser. Certification has a stronger influence on respondents who are generally skeptical of adviser quality. This is in line with the findings by Agnew et al. (2016) who show that displaying a certification significantly increases an adviser's likelihood of being chosen.

Choice of advice

Respondents' ability to choose good advice depends on the clarity of the topic. In the case of easy topics, the respondent will choose the good advice (up to some error, Equation (5)). In the case of hard topics, Equation (6) posits that (up to some error) the respondent will choose the adviser they rate as better according to the respondent's updated (posterior) belief. The parameters associated with both the quality of the advice and the belief about the adviser are positive equaling 4.138 and 2.663, respectively, and the 95% credible intervals based on their posterior distributions do not include zero. This translates into the following choice probabilities: the probability that the respondent chooses good advice if the topic is easy equals $\exp(4.296)/(1+\exp(4.296)) = 0.99$; if the topic is ambiguous, the probability of an adviser R with associated belief $\lambda^R = 1$ being chosen when being evaluated against an adviser L with associated belief $\lambda^L = 0$ is $\exp(2.510 \cdot (1-0))/(1+\exp(2.510 \cdot (1-0))) = 0.92$. These results confirm that ambiguous signals are related to more uncertainty and variability in respondents' choices.

Willingness to pay for advice

Our model assumes that the willingness to pay a particular price for an additional hour with the adviser depends on the actual price to be charged, several characteristics of the respondent, as well as the respondent's posterior belief about this adviser. Table 5 shows that parameters here have the expected signs but some have credible intervals that include zero. The impact of price is negative with a mode of -0.085 and a 95% CI interval that does not include zero. The impact of the posterior belief about the adviser is on the other hand positive (18.309) with the associated 95% CI also not including zero. Of the remaining respondent characteristic, the only parameter with a CI that does not include zero is the indicator for whether the respondent has paid for financial advice in the past. The mode of this parameter is positive at 0.466 with a 95% CI of [0.348, 0.570] and we conclude that respondents who have paid for advice in the past are more willing to pay than those who have not.

Based on these parameters and Equation (7) it is possible to calculate the associated price difference $\Delta\text{price} = \text{price}_{new} - \text{price}_{old}$ that a respondent is willing to pay for a specific difference in posterior beliefs $\Delta\text{belief} = \text{belief}_{new} - \text{belief}_{old}$, namely

$$\Delta\text{price} = -\frac{\beta_2^{\text{posterior}}}{\beta_2^{\text{price}}} \cdot \Delta\text{belief} \cdot 100, \tag{13}$$

where the multiplication with 100 is necessary because the price was divided by 100 before entering the estimation. We can use this formula to calculate the additional dollar amount that respondents are willing to pay for their preferred adviser. Based on the posterior distribution of the estimates, we obtain additional willingness to pay estimates that have a mean of $1722 with the lower bound of the 95% CI equaling $189 and the upper bound equaling $4639.

Rational versus Fryer updating

Our model shows that a 62.9% of respondents exhibit confirmation bias. In our setup, we use respondents' conscientiousness and impulsiveness to explain which respondents are more likely to display confirmation bias. Table 5 shows that respondents with high impulsiveness are less likely to be rational updaters (mode of -0.3445, the 95% CI does not include zero). This parameter implies that more impulsive respondents are more likely to interpret ambiguous signals as a confirmation of their prior belief. In contrast, high conscientiousness has a positive mode (0.243) but the 95% CI does include zero.

Clarity of topics

In our model we assume that whether a topic is perceived as clear or ambiguous by a respondent depends on the respondent's characteristics as well as on the topic itself. More specifically, we find that respondents with more expertise are more discerning. Results show that respondents with high knowledge of the financial products related to the advice, high financial literacy and high numeracy are more likely to perceive a topic as clear. For all these variables the modes of the posterior distributions are positive and the 95% CIs do not contain zero. Gender and respondents' age also significantly impact whether a topic is perceived as clear versus ambiguous: Female respondents are significantly more likely to perceive a topic as clear and so are respondents who are 40 years or older. In addition, the size and sign of the topic-specific constants is in line with the share of correct answers for these topics. The advice related to index fund fees is perceived as significantly more difficult than all other topics since the associated 95% CI does not overlap with the CI of any other topic.

Table 6 reports the percentage of respondents who belong to each of the 16 clarity classes: 18.2% of respondents perceive all topics to be clear; 3.8% of respondents perceive all topics to be ambiguous; and 21.9% of respondents struggle to understand advice on index fund fees even though all other topics are clear to them.

Illustration of Model Implications

Our model allows us to compare the impact of the two different updating strategies on consumers' choices. It additionally allows us to measure the impact of first impressions on subsequent choices. To illustrate, consider two respondents A and B who update their beliefs according to the rational and biased updating scheme, respectively. Let the right adviser (*R*) display a certification and the left adviser (*L*) not display a certification. For the sake of simplicity we assume that the respondents distrust financial advisers. Based on the estimation results regarding the prior beliefs (and focusing on their mode again for simplicity), both respondents will thus have the same prior belief about the right (*R*) and the left (*L*) adviser of

$\lambda_{A0}^{R} = \lambda_{B0}^{R} = \dfrac{\exp(1.728 + 0.085 - 520)}{1 + \exp(1.728 + 0.085 - 520)} = 0.785$ and

$\lambda_{A0}^{L} = \lambda_{B0}^{L} = \dfrac{\exp(1.728 - 0.085 - 0.520)}{1 + \exp(1.728 - 0.085 - 0.520)} = 0.755$. Assume that Adviser R gives good advice on a clear topic in the first choice set and that both advisers give (from the client's perspective) ambiguous advice in the remaining three choice sets.

Table 7 shows how the updated prior beliefs, choice probabilities for the advisers in each choice set evolve in this scenario (again all calculations are based on the mode of the posterior distributions for the sake of simplicity). Both clients update their beliefs in the same way at the first choice because they get clear information about Adviser quality. Respondent A's beliefs about the advisers, as well as the associated choice probabilities, remain the same throughout the later three choice sets as this respondent simply ignores the ambiguous information and ends the experiment still favoring the Advisor R. In contrast, respondent B interprets all new information in line with current beliefs, so this respondent will treat all ambiguous information as evidence that adviser R is good and that adviser L is bad. Thus, respondent B's updated beliefs about adviser R rise steadily and so does his or her probability of choosing adviser R.

The table thus shows that Fryer updating leads to a choice probability that is very close to one adviser R and close to zero for adviser L while the same probabilities are 0.9 (Adviser R) and 0.1 (Adviser L) for the rational updater. It also shows the difference a first impression makes. The early clear signal has a stronger influence on the Fryer updater, whose opinion approaches certainty over few choices.

DISCUSSION

We study confirmation bias, a behavioral bias that leads to new information being interpreted in support of prior beliefs. Learning about a product or service occurs in situations where people have to form opinions under incomplete information and can be found in a wide variety of situations ranging from purchases in unknown product categories to evaluation of candidates in job interviews or reviews of academic papers. Learning often starts with a first impression and can then be distorted by confirmation bias, leading to a polarization of beliefs and an exaggerated importance of first impressions.

We conduct a discrete choice experiment to measure the impact of confirmation bias both on choices of, and willingness to pay for, financial advice. We the experiment via an online survey to infer the proportion of respondents who are actually prone to this bias when they evaluate an unobserved product attribute – in this case the quality of a financial adviser.

Our results show that to assume that all learners apply rational Bayesian methods when they receive and process signals of attribute quality is too strong. A majority of subjects in our experiment did not use the commonly assumed rational method but made choices consistent with a limited memory updating where people use unclear signals to confirm and reinforce their current belief. Our experiment shows that people who are unsure how to interpret the signals they receive and who do not ignore them not only end up with strongly biased beliefs but will spend accordingly.

We also show how differences in first impressions manifest themselves in differences in willingness to pay and how polarization induced by confirmation bias can lead to dramatically different outcomes. Thus, our model puts a dollar value on differences in first impressions as well as on differences in updating schemes and shows that traditional learning models that ignore confirmation bias will underestimate the impact of first impressions. By using respondent characteristics to predict the probability of using rational versus Fryer updating, our model allows us to identify those consumers who are most likely to suffer financially from the information processing bias.

Our results are in line with findings reported in the context of brand equity research that is based on signaling theory (Erdem and Swait 1998) and that shows that clear brand signals can increase a brand's credibility, which in turn reduces perceived risk, and thus increases expected utility. However, whilst this research reasons that consumers use perceived clarity to infer firms' willingness and ability to offer the promised products, we show that signal clarity can also impact utility by leading to distorted perceptions of the signal itself.

Our results have important implications that can stimulate research into further development of similar models. Service providers could leverage insights from this model to have consumers choose their products and services by establishing a very clear positive impression very early on and before any ambiguous information is provided. Similarly, since our model segments consumers by their updating strategies, companies can gain valuable insight into which customer segment should be the focus of education campaigns. Such research should use covariates related more closely to the product category under consideration, but it is conceivable that previously introduced scales such as the Lay Rationalism Scale (Hsee et al. 2015) or the Emotion-Based Decision Making Scale (Barchard 2001) as covariates for the rationality segments could be particularly valuable in predicting consumer's updating strategies.

In the context of financial advice our model provides useful insights both for public policy and for financial advisers. For the latter, we show that displaying recognizable professional certification has a significant positive impact on first impressions which then filters through to a higher chance that clients will accept advice and a higher willingness to pay for additional advice. Thus, advisers should assess the costs of gaining a qualification in the light of these possible future gains. The implications for public policy are even more interesting: While current research emphasizes improvements to financial literacy to encourage sensible financial decisions, our model shows that less impulsiveness can also increase consumer welfare. Thus, our model proposes another potentially important lever.

A possible modification concerns the updating strategies: We assume that respondents are either purely Fryer or purely rational updaters; they interpret ambiguous signals as exactly confirming their prior belief (Fryer updating) or as not being informative at all (rational updating). It is possible that respondents interpret ambiguous signals as only part confirmation of their prior, meaning that respondents sit on a continuum between extreme updating processes. We leave such a modification of the model to future research.

References:

Agarwal, Sumit, John C Driscoll, Xavier Gabaix, and David Laibson (2009), "The age of reason: Financial decisions over the life cycle and implications for regulation," Brookings Papers on Economic Activity, 2009 (2), 51-117.

Agnew, J., Bateman, H., Eckert, C., Ishkanov, F., Louviere, J., Thorop, S. (forthcoming), "First Impressions Matter: An Experimental Investigation of Online Financial Advice," Management Science.

Agnew, Julie R, Hazel Bateman, Christine Eckert, Fedor Iskhakov, Jordan Louviere, and Susan Thorp (2016), "First Impressions Matter: An Experimental Investigation of Online Financial Advice," Management Science.

Anagol, Santosh, Shawn Allen Cole, and Shayak Sarkar (2013), "Understanding the advice of commissions-motivated agents: Evidence from the Indian life insurance market," Harvard Business School Finance Working Paper (12-055).

Australian Securities and Investment Commission, ASIC (2012), "Shadow shopping study of retirement advice, Report 279, Sydney.," (accessed September 14, 2012, [available at http://www.asic.gov.au/asic/pdflib.nsf/LookupByFileName/rep279-published-27-March-2012.pdf/$file/rep279-published-27-March-2012.pdf].

Bar-Gill, Oren and Elizabeth Warren (2008), "Making credit safer," University of Pennsylvania Law Review, 1-101.

Barchard, Kimberly Anne (2001), "Emotional and social intelligence: Examining its place in the nomological network," University of British Columbia.

Beattie, Jane and Jonathan Baron (1988), "Confirmation and matching biases in hypothesis testing," The Quarterly Journal of Experimental Psychology, 40 (2), 269-97.

Camacho, Nuno, Bas Donkers, and Stefan Stremersch (2011), "Predictably non-Bayesian: quantifying salience effects in physician learning about drug quality," Marketing Science, 30 (2), 305-20.

Campbell, John Y (2006), "Household finance," The Journal of Finance, 61 (4), 1553-604.

Campbell, John Y, Howell E Jackson, Brigitte C Madrian, and Peter Tufano (2011), "Consumer financial protection," The Journal of Economic Perspectives, 25 (1), 91-113.

Carlin, Bruce I. (2009), "Strategic price complexity in retail financial markets," Journal of Financial Economics, 91(3), 278-287.

Carlin, Bruce I, and Gustavo Manso (2010), "Obfuscation, learning, and the evolution of investor sophistication," The Review of Financial Studies, 24(3), 754-785.

Chater, Nick, Steffen Huck, and Roman Inderst (2010), "Consumer decision-making in retail investment services: A behavioural economics perspective," Report to the European Commission/SANCO.

Chernozhukov, Victor and Han Hong (2004), "Likelihood estimation and inference in a class of nonregular econometric models," Econometrica, 72 (5), 1445-80.

Ching, Andrew T, Tülin Erdem, and Michael P Keane (2013), "Invited Paper—Learning models: An assessment of progress, challenges, and new developments," Marketing Science, 32 (6), 913-38.

Christelis, Dimitris, Tullio Jappelli, and Mario Padula (2010), "Cognitive abilities and portfolio choice," European Economic Review, 54 (1), 18-38.

Chylinski, Mathew B, John H Roberts, and Bruce GS Hardie (2012), "Consumer Learning of New Binary Attribute Importance Accounting for Priors, Bias, and Order Effects," Marketing Science, 31 (4), 549-66.

Darley, John M and Paget H Gross (1983), "A hypothesis-confirming bias in labeling effects," Journal of Personality and Social Psychology, 44 (1), 20.

Douc, Randal and Olivier Cappé (2005), "Comparison of resampling schemes for particle filtering," in Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on: IEEE.

Durham, Garland and John Geweke (2014), "Adaptive sequential posterior simulators for massively parallel computing environments," in *Bayesian Model Comparison*: Emerald Group Publishing Limited.

Eckstein, Zvi, Dan Horsky, and Yoel Raban (1988), "An Empirical Dynamic Model of Brand Choice," Working Paper No. 88, University of Rochester."

Erdem, Tülin and Michael P Keane (1996), "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets," Marketing science, 15 (1), 1-20.

Erdem, Tülin and Joffre Swait (1998), "Brand equity as a signaling phenomenon," Journal of consumer Psychology, 7 (2), 131-57.

Evans, Kenneth R., Robert E. Kleine, Timothy D. Landry, and Lawrence A. Crosby (2000), "How First Impressions of a Customer Impact Effectiveness in an Initial Sales Encounter," Journal of the Academy of Marketing Science, 28 (4), 512-26.

Fryer, Roland G, Philipp Harms, and Matthew O Jackson (2015), "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization," Available at SSRN 2263504.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny (2015), "Money doctors," The Journal of Finance, 70 (1), 91-114.

Georgarakos, Dimitris and Roman Inderst (2014), "Financial advice and stock market participation."

Guiso, L. , P. Sapienza, and L. Zingales (2008), "Trusting the stock market," The Journal of Finance, 63 (6), 2557-600.

Hackethal, Andreas, Michael Haliassos, and Tullio Jappelli (2012), "Financial advisors: A case of babysitters?," Journal of Banking & Finance, 36 (2), 509-24.

Hartley, Robert F. (1995), Marketing Mistakes (6 ed.). New York: John Wiley & Sons.

Holden, Sarah A, SM Olivia, and K Smetters (2013), When, why, and how do mutual fund investors use financial advisers?: Oxford University Press Oxford, UK.

Holtz, Brian C (2015), "From first impression to fairness perception: Investigating the impact of initial trustworthiness beliefs," Personnel Psychology, 68 (3), 499-546.

Hsee, Christopher K, Yang Yang, Xingshan Zheng, and Hanwei Wang (2015), "Lay rationalism: Individual differences in using reason versus feelings to guide decisions," Journal of Marketing Research, 52 (1), 134-46.

Inderst, Roman and Marco Ottaviani (2009), "Misselling through agents," The American Economic Review, 99 (3), 883-908.

Jasra, Ajay, David A Stephens, and Christopher C Holmes (2007), "On population-based simulation for static inference," Statistics and Computing, 17 (3), 263-79.

Lawson, Robert (1997), "Consumer decision making within a goal ‐ driven framework," Psychology & Marketing, 14 (5), 427-49.

Lusardi, Annamaria and Olivia S Mitchell (2006), "Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth," National Bureau of Economic Research.

---- (2011), "Financial literacy around the world: an overview," Journal of Pension Economics and Finance, 10 (04), 497-508.

Morgan, Gary and Michele Levine (2015), "Roy Morgan Image of Professions Survey 2015."

Mullainathan, Sendhil, Markus Noeth, and Antoinette Schoar (2012), "The market for financial advice: An audit study," National Bureau of Economic Research.

Muthukrishnan, A. V. and Amitava Chattopadhyay (2007), "Just Give Me Another Chance: The Strategies for Brand Recovery from a Bad First Impression," Journal of Marketing Research, 44 (May), 334-45.

Nickerson, Raymond S (1998), "Confirmation bias: A ubiquitous phenomenon in many guises," Review of general psychology, 2 (2), 175.

Rabin, Matthew and Joel L Schrag (1999), "First impressions matter: A model of confirmatory bias," The Quarterly Journal of Economics, 114 (1), 37-82.

Reinhard, Marc-André, Matthias Messner, and Siegfried Ludwig Sporer (2006), "Explicit persuasive intent and its impact on success at persuasion—The determining roles of attractiveness and likeableness," Journal of Consumer Psychology, 16 (3), 249-59.

Roberts, John H and Glen L Urban (1988), "Modeling multiattribute utility, risk, and belief dynamics for new consumer durable brand choice," Management Science, 34 (2), 167-85.

Ryan, Andrea, Gunnar Trumbull, and Peter Tufano (2011), "A brief postwar history of US consumer finance," Business History Review, 85 (03), 461-98.

Sniezek, Janet A and Lyn M Van Swol (2001), "Trust, confidence, and expertise in a judge-advisor system," Organizational behavior and human decision processes, 84 (2), 288-307.

Thaler, Richard H and Shlomo Benartzi (2004), "Save more tomorrow™: Using behavioral economics to increase employee saving," Journal of political Economy, 112 (S1), S164-S87.

Washington Post (2015), "California protesters take on Nestle over drought."

Yaniv, Ilan and Eli Kleinberger (2000), "Advice taking in decision making: Egocentric discounting and reputation formation," Organizational behavior and human decision processes, 83 (2), 260-81.

Zaleskiewicz, Tomasz, Agata Gasiorowska, Katarzyna Stasiuk, Renata Maksymiuk, and Yoram Bar-Tal (2016), "Lay Evaluation of Financial Experts: The Action Advice Effect and Confirmation Bias," Frontiers in Psychology, 7.

Table 1: Financial Advice Script

Narrator Introduction	Advice	Narrator Introduction	Advice
Paying Down Debt In this scenario, you have accumulated some large outstanding credit card debt with a high associated interest rate. Recently, you have inherited some money unexpectedly and would like to know what to do with it. The next 2 financial advisers will recommend what you should do.	Good Advice: <u>I understand that you have some large credit card debt but recently inherited money. It is important to think about your overall financial position when making a decision about what to do.</u> *It is easy to simply save this big sum of money in a savings account to achieve a savings goal, but the interest gained is far smaller than the high interest expense of not paying down your credit card debt. Therefore, I recommend you pay off your credit card debt to eliminate the high interest charges.* Bad Advice: [Insert underlined above] *It is hard to save big sums of money so it is important to think about your special savings goals when making this decision. Therefore, I recommend you ignore your credit card debt for now and put your inheritance in a separate savings account.*	Choosing an Index Fund In this scenario, you are thinking about investing in a managed share index fund. The next 2 financial advisers will recommend what you should do about it.	Good Advice: <u>I understand you need help regarding your choice of share index fund. Did you know that all share index funds invest with the aim of matching the overall share market return? These various share index funds provide an almost identical product</u> *so why pay a fund manager more than the others for the same thing. Therefore, I recommend that you choose the share index fund with the lowest management fees.* Bad Advice [Insert underlined above] *but some fund managers have better reputations than others and you get what you pay for. Therefore, I recommend that you avoid the share index funds with low management fees.*
Consolidating Retirement Accounts In this scenario, suppose you have just changed jobs and started a new superannuation account. Currently, you already have two other superannuation accounts from past jobs. The next 2 financial advisers will recommend what you should do about it.	Good Advice: <u>I see that you have three superannuation accounts with different super funds. Did you know that people are typically charged regular fixed administration fees on all of these superannuation accounts?</u> *As a result, I recommend that you roll all of these accounts together so you are not paying extra fees.* Bad Advice: [Insert underlined above] *Despite that, I recommend that you not roll all of these accounts together so you are diversified across different superannuation funds.*	Diversifying a Stock Portfolio In this scenario, you are thinking about investing in the share market. The next 2 financial advisers will recommend what you should do about it.	Good Advice: <u>I understand you need help regarding how to invest your superannuation money. Did you know money invested in shares can go up and down?</u> *It is good to try to balance out the shares that go up with the shares that go down. Therefore, I recommend that you spread your money across a variety of shares in different types of companies and industries.* Bad Advice: [Insert underlined above] *That is why it is good to invest in something you know and can easily monitor. Therefore, I recommend that you invest your money in one blue chip company.*

Notes: Table 1 provides the scripts for the four advice topics. Each video begins with the narrator's introduction. The two advisers provide identical advice (the underlined advice) at the beginning of their talk and then depart from one another at the end (the italicized part).

Table 2: Experimental design

Panel A. Design of advisers pairs

	Adviser 1 (Shown on the left)			Adviser 2 (Shown on right, mirror image)		
Pair	Gender	Age	Accreditation	Gender	Age	Accreditation
1	Female	Young	Yes	Male	Old	No
2	Female	Old	No	Male	Young	Yes
3	Male	Young	No	Female	Old	Yes
4	Male	Old	Yes	Female	Young	No

Panel B. Sequence of advice topics

Sequence	Choice 1	Choice 2	Choice 3	Choice 4	Clarity
1	Diversification	Fees	Consolidation	Debt	HHEE
2	Consolidation	Debt	Diversification	Fees	EEHH
3	Diversification	Consolidation	Fees	Debt	HEHE
4	Consolidation	Diversification	Debt	Fees	EHEH
5	Diversification	Consolidation	Debt	Fees	HEEH
6	Consolidation	Diversification	Fees	Debt	EHHE

Panel C. Design of the sequence of advice quality

	Advice from adviser 1 (shown on the left)				Advice from adviser 2 (shown on the right, mirror image)			
Quality Sequence	1st topic	2nd topic	3rd topic	4th topic	1st topic	2nd topic	3rd topic	4th topic
1	G	G	B	B	B	B	G	G
2	G	B	G	B	B	G	B	G
3	G	B	B	G	B	G	G	B
4	B	G	G	B	G	B	B	G
5	B	G	B	G	G	B	G	B
6	B	B	G	G	G	G	B	B

Notes: Panel A shows the combination of adviser attributes using a foldover design for each possible adviser. Each participant to the survey viewed only one of the eight rows. Thus, they saw the same two advisers for the entire experiment and each adviser stayed on the same side of the screen throughout the experiment. Panel B shows sequence of advice topics for each condition in the experiment. Each participant viewed one of the four columns, interacted with the rows in Panel C, where "E" stands for one of the easy topics (debt and account consolidation) and "H" stands for one of the hard topics (fees and diversification). Panel C shows the eight sequences of advice quality for each condition in the experiment. Each participant viewed one of the eight rows. G stands for good advice, while B stands for bad advice.

Table 3: Demographics, survey sample and Australian population (18-79 years), 2011 Census.

	Survey Participant Sample	18-79 yrs Australian Population		Survey Participant Sample	18-79 yrs Australian Population
Gender			**Marital Status**		
Male	50%	49%	Never Married	26%	30%
Female	50%	51%	Divorced/ Separated	10%	13%
Age			Widowed	2%	3%
18-24 years	8%	10%	Married or long term relationship	62%	54%
25-29 years	8%	10%	**Personal Income**		
30-34 years	12%	10%	$1-$20,799 (i.e. less than $399 a week)	24%	25%
35-39 years	12%	10%	$20,800-$51,999 (i.e. $400-$999 a week)	35%	32%
40-44 years	12%	10%	$52,000-$103,999 (i.e. $1,000-$1,999 a week)	25%	23%
45-49 years	9%	10%	$101,000 (i.e. $2,000 a week) or more	7%	7%
50-54 years	12%	10%	Negative or Nil Income	9%	6%
55-59 years	12%	9%	Not Started	0%	7%
60-64 years	13%	8%	**Highest Level of Education**		
65-69 years	2%	6%	High School or Less	26%	40%
70-79 years	0%	8%	Vocational/Technical certificate	21%	20%
Work Status			Tertiary diploma	11%	9%
Employed	62%	63%	Bachelor degree	23%	15%
Unemployed	8%	3%	Graduate certificate, diploma or degree	19%	6%
Not in the labour force	18%	29%	Not stated	0%	10%
Retired	12%	not broken out			
Not stated	0%	5%			

Table 4: Variable description

Variable Name	X_0	X_2	X_4	X_5	Description
Constant	x	x	x	x	Constant; topic specific for X_5
Adviser characteristics					
Displays NO credential	x				Indicator variable that equals 1 if only adviser's name was displayed and -1 when "Certified Financial Planner" and adviser's name was displayed.
Price		x			Price in $ (divided by 100) for one additional hour with this adviser
Posterior		x			Posterior belief about adviser after advice on all four topics has been provided – estimated within the model
Advice					
Good advice					Indicator variable that equals 1 if the wrong advice was given in the particular choice set, -1 otherwise. Enters the model via the choice specification in Equation (5)
Topic: Account consolidation				x	Indicator variable that equals 1 if the topic was account consolidation, 0 otherwise.
Topic: Stock diversification				x	Indicator variable that equals 1 if the topic was stock diversification, 0 otherwise.
Topic: Index fund fee				x	Indicator variable that equals 1 if the topic was index fund management fees, 0 otherwise.
Topic: Debt repayment				x	Indicator variable that equals 1 if the topic was debt repayment, 0 otherwise.
Participant characteristics					
Participant female				x	An indicator variable that equals 2 if the participant is a female, 1 otherwise.
Participant older than 39 years				x	An indicator variable that equals 1 if the participant is a older than 39 years, 0 otherwise.
Trust in advisers	x				An indicator variable that equals 1 if the participant reported general trust in financial advisers, -1 if distrust, 0 otherwise
Paid for advice		x			Indicator variable that equals 1 if the participant has ever paid for financial advice, -1 if they have not
Household income		x			Household income ($'000, mean centered)
Confidence in financial decisions		x			Indicator variable that equals 1 if participant has high confidence in their own ability to make financial decision, -1 if low
Financial risk tolerance		x			Indicator variable that equals 1 if participant's risk tolerance is high and -1 if low
Decision maker		x			Indicator variable that equals 1 when the participant is most responsible for financial decisions, 0 when jointly responsible and -1 when someone else is responsible.
Financial literacy				x	An indicator variable that equals 1 if the participant's correct percentage on four financial literacy questions is above the sample median, 0 otherwise. Questions test simple interest, inflation, diversification, and compound interest.
Numeracy				x	An indicator variable that equals 1 if the participant's correct percentage on three numeracy questions is above the sample median, 0 otherwise. Questions test fractions, percentages and probabilities.
Product knowledge				x	An indicator variable that equals 1 if the participant's correct percentage on four financial product questions is above the sample median, 0 otherwise. Questions test topics used in advice experiment: debt, index funds, account consolidation, diversification.
Conscientiousness			x		An indicator variable that equals 1 if the participant's conscientiousness is above the sample median, 0 otherwise.

			Participants rated themselves as organized, responsible, hardworking and careless (reverse coded) on a four-point scale. Ratings are averaged.
Impulsiveness	x		An indicator variable that equals 1 if the participant's impulsiveness is above the sample median, 0 otherwise. Participants rated themselves as buying too much, buying impulsively, buying without planning, and/or buying unnecessarily on a five point scale. Ratings are averaged.
Market experience		x	An indicator variable that equals 1 if the participant's percentage on owning four financial securities is above the sample median, 0 otherwise. Participants reported whether they owned a credit card (debt), units in an index fund (fees), a superannuation account (consolidation) and stocks (diversification).

Table 5: Empirical Results

	Mode	2.5 Percentile	97.5 Percentile
Prior belief about adviser			
Trust in financial advisers	0.520	0.421	0.610
Displays NO credential	-0.085	-0.199	-0.009
Constant	1.728	1.511	1.903
Choice of Advice			
Quality	4.296	3.599	5.060
Sigma	2.510	1.494	3.622
Willingness to pay			
Constant	-7.782	-9.687	-6.228
Price	-0.085	-0.124	-0.043
Posterior	18.309	14.808	22.230
Paid for advice	0.466	0.348	0.570
Household income	0.094	-0.021	0.163
Confidence in financial decisions	-0.088	-0.186	0.051
Financial risk tolerance	0.055	-0.047	0.156
Decision maker	0.034	-0.125	0.186
Rational vs Fryer updating			
Constant	-0.454	-0.994	-0.185
High Conscientiousness	0.243	-0.046	0.485
High Impulsiveness	-0.344	-0.724	-0.154
Clarity of Topics			
High Market Experience	0.073	-0.046.	0.151
High Product Knowledge	0.267	0.171	0.358
Participant older than 39	0.554	0.441	0.646
Participant female	0.138	0.053	0.237
High Financial Literacy	0.372	0.244	0.458
High Numeracy	0.357	0.278	0.482
Consolidation	1.405	1.148	1.632
Diversification	0.615	0.395	0.814
Fees	-0.545	-0.794	-0.358
Debt	1.768	1.511	1.995
Size of rational updating segment: 37.11%			

Table 6: Share of respondents in clarity classes

Consolidation	Diversification	Fees	Debt	Segment Size
1	1	1	1	18.2
1	1	1	0	2.2
1	1	0	1	21.9
1	1	0	0	4.5
1	0	1	1	6.9
1	0	1	0	1.4
1	0	0	1	14.0
1	0	0	0	4.8
0	1	1	1	3.1
0	1	1	0	0.6
0	1	0	1	6.4
0	1	0	0	2.3
0	0	1	1	2.0
0	0	1	0	0.7
0	0	0	1	7.2
0	0	0	0	3.8

Note: 1 implies that the topic is perceived as clear, 0 otherwise

Table 7: Evolution of choice probabilities and beliefs with one clear and three ambiguous topics.

	λ_0	$\Pr(y_1=1)$	λ_1	$\Pr(y_2=1)$	λ_2	$\Pr(y_3=1)$	λ_3	$\Pr(y_4=1)$	λ_4
Adviser R, client A	0.785	0.987	0.916	0.886	0.916	0. 886	0.916	0. 886	0.916
Adviser R, client B	0.785	0.987	0.916	0.886	0.970	0.913	0.990	0.921	0.997
Adviser L, client A	0.755	0.013	0.098	0.114	0.098	0. 114	0.098	0. 114	0.098
Adviser L, client B	0.755	0.013	0.098	0.114	0.035	0.087	0.012	0.079	0.004

Notes: Example assumes client A uses rational updating and client B uses Fryer updating. Both clients are initially distrusting of financial advisers. Adviser R shows a professional certification and Adviser L does not. Adviser R delivers good advice on a clear topic at choice 1. Topics 2-4 are ambiguous to both clients.

Figure 1: Updating of beliefs

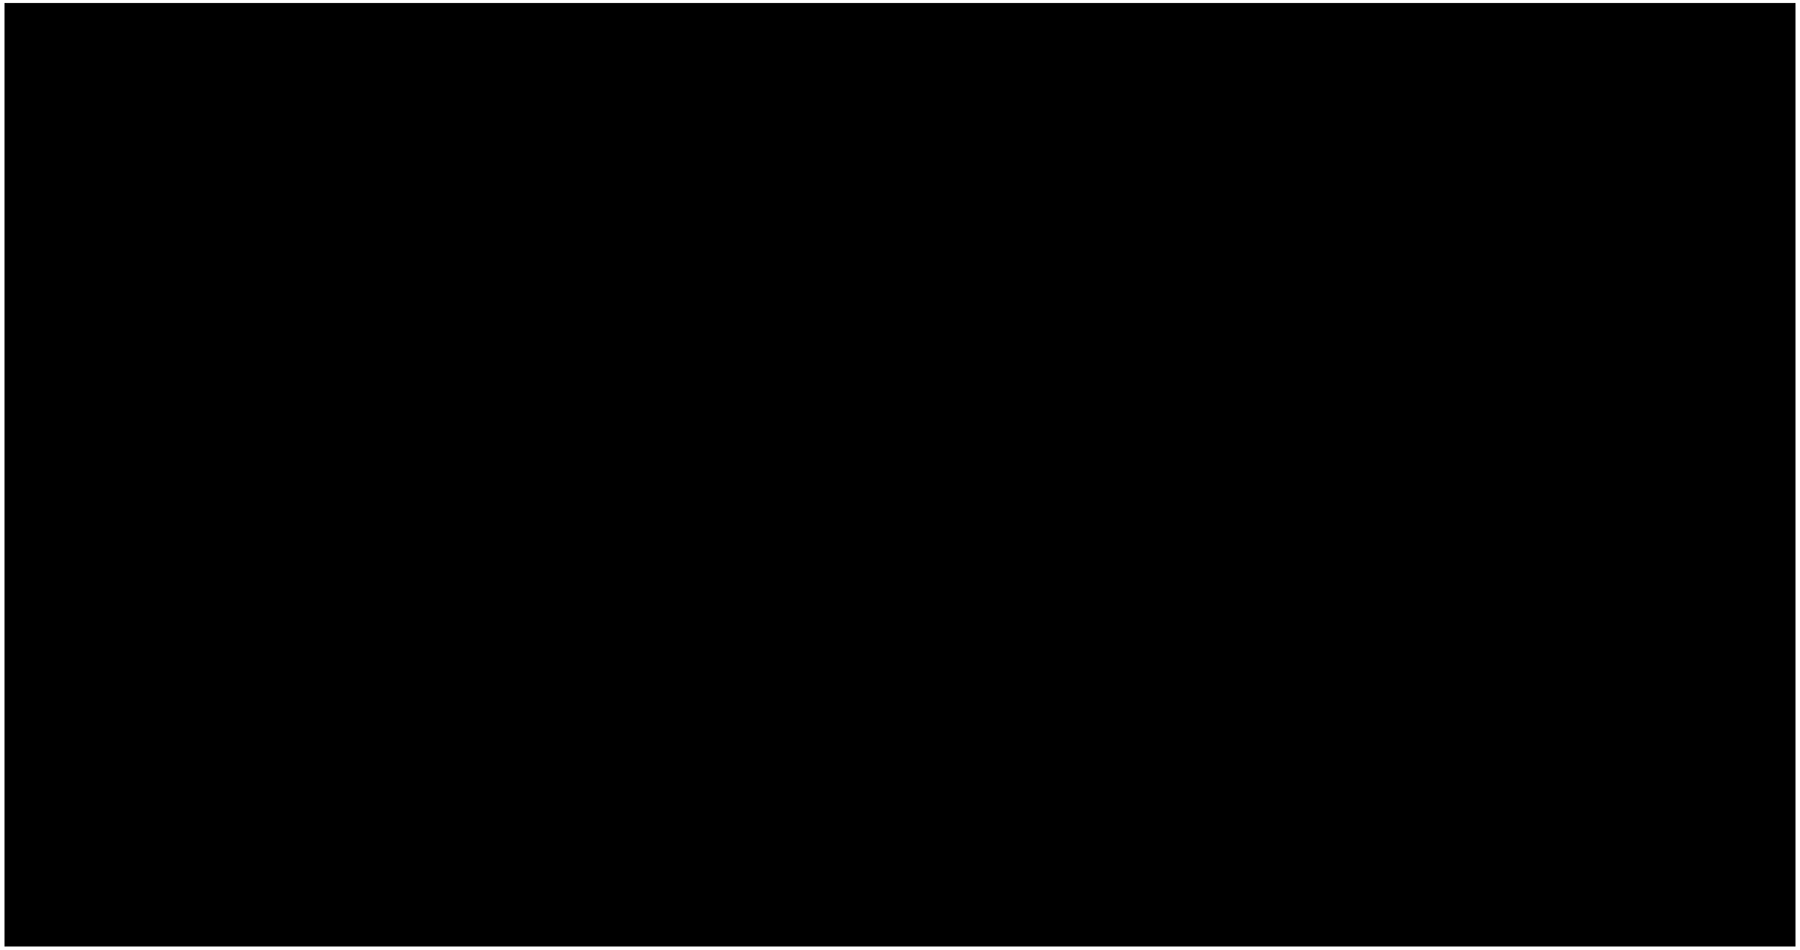

Figure 2: Advisers

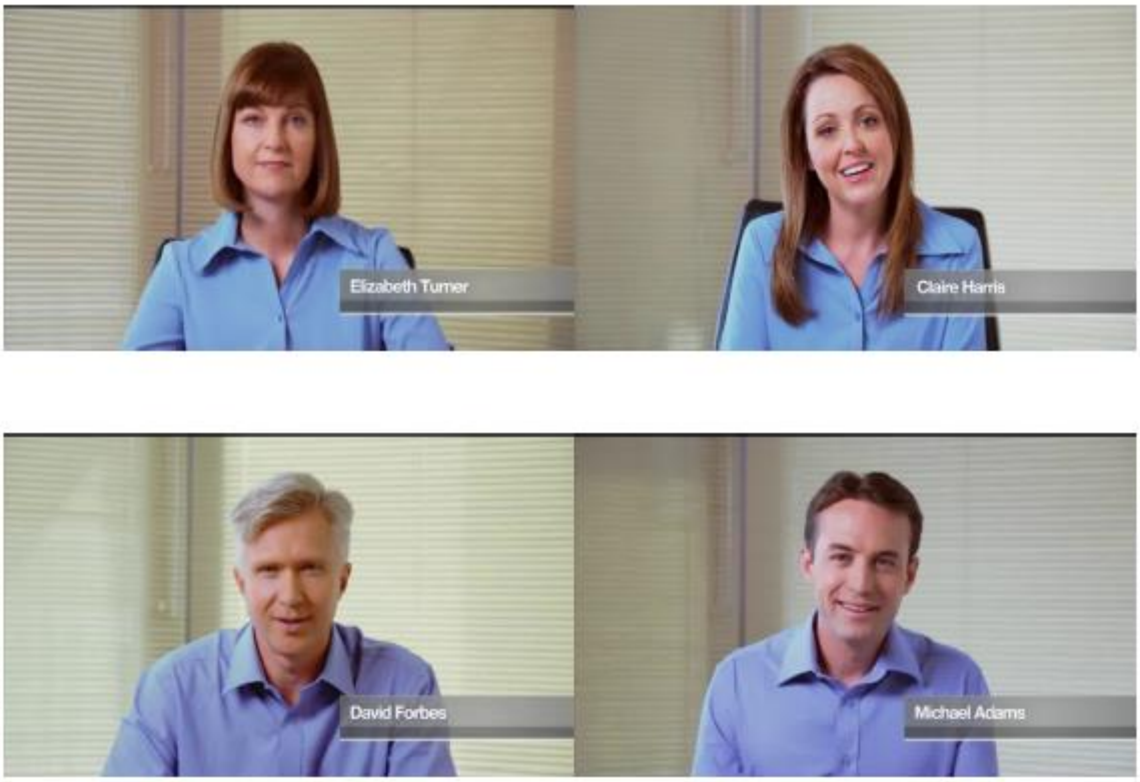

Figure 3: Impact of prior on posterior