

Who Pays the Price for Bad Advice? The Role of Consumer Vulnerability, Learning and Confirmation Bias

Julie Agnew*, Hazel Bateman†, Christine Eckert‡, Fedor Iskhakov§, Susan Thorp¶

Abstract

We study how consumers learn about the quality of expert services for credence goods. For these goods, sellers provide the service, and also act as experts who determine consumers' needs. Due to the nature of credence goods, consumers often must rely on their interpretation of ambiguous seller signals to judge seller quality. We focus on the context of financial advice and examine how consumers interpret ambiguous signals to infer a financial adviser's quality. Using data from an experimental survey, we propose and estimate a structural model of learning by advice clients. Survey participants ("clients") make subjective interpretations of advice on topics of varying difficulty, simultaneously learning about the quality of the advisers giving the advice. Most participants fail to follow Bayesian updating rules and use ambiguous signals from difficult topics to confirm their prior beliefs about the advisers. Impulsive participants are more prone to such confirmation bias. Vulnerable participants who combine impulsiveness with low financial skills, and thus an inability to judge the quality of advice, will pay too much to poor advisers – particularly to malevolent advisers who present advice strategically. Our model can identify strategies to help vulnerable consumers in the context of financial advice and beyond.

Keywords: Consumer vulnerability, expert services, credence goods, confirmation bias, willingness to pay

*Raymond A. Mason School of Business, William and Mary, Williamsburg, VA, 23187 USA; E-mail: Julie.Agnew@mason.wm.edu

†UNSW Business School, UNSW Sydney, Sydney, NSW, Australia, E-mail: h.bateman@unsw.edu.au

‡Corresponding author, UTS Business School, University of Technology Sydney, NSW, Australia, E-mail: christine.eckert@uts.edu.au and EBS Business School, EBS University, Oestrich-Winkel, Germany, E-mail: christine.eckert@ebs.edu

§Research School of Economics, Australian National University, Canberra, ACT, Australia, E-mail: fedor.iskhakov@anu.edu.au

¶The University of Sydney Business School, The University of Sydney, NSW, Australia, E-mail: susan.thorp@sydney.edu.au; Tel: +61 2 9036 6354.

We dedicate this paper to the memory of Jordan Louviere, loved mentor and colleague.

This project was approved by the College of William and Mary Protection of Human Subjects Committee (Phone 757-221-3966) on 2013-03-08.

INTRODUCTION

Getting financial advice can be hazardous. Advisers sometimes recommend under-diversified or under-performing investments, as well as over-trading, return-chasing, or expensive products (Bergstresser et al., 2009; Mullainathan et al., 2012; Hackethal and Inderst, 2013; Hoechle et al., 2017). "Expert" advisers may have biases that are transferred to clients (Linnainmaa et al., 2021). Conflicting incentives may encourage advisers to sell second-rate products (Egan et al., 2020), while others offer poor advice to less-savvy clients (Bucher-Koenen et al., 2021) or give advice based on social stereotypes (Baekström et al., 2021). Worse still, some advisers repeatedly defraud or exploit clients then hide their misconduct by moving firms or regulatory jurisdictions (Egan et al., 2019; Honigsberg et al., 2022).

At the same time, clients themselves may misjudge whose, or what, advice to take. Clients may reach opinions based on scant evidence (Yaniv and Kleinberger, 2000). They might base their judgements on personal similarities (Stolper and Walter, 2019), on agreement with their opinions (Soll et al., 2022) or on mistaken beliefs that advisers always act in their clients' best interests (Yaniv and Kleinberger, 2000). Research also documents that established trust may persist despite poor first service encounters (Mullainathan et al., 2012). Misplaced trust can be costly as trusting clients are more likely to take advice, invest in risky assets or pay high fees (Gennaioli et al., 2015; Germann et al., 2018). These findings raise obvious questions. Why do clients stick with advisers who give poor advice? And what is the cost of doing so?

The difficulty of choosing a good adviser is emblematic for purchases of credence goods. For these goods consumers can neither *ex ante* nor *ex post* be certain of their quality (Darby and Karni, 1973). Consequently, sellers of these goods often provide the goods and also act as experts who determine consumers' needs and give advice on which goods to purchase. Such expert services are provided by financial advisers, medical doctors, lawyers, car mechanics and appliance service persons, to name a few. Unfortunately, such expert services are also prone to seller fraud, as the information asymmetry between consumer and seller creates strong incentives for opportunistic seller behavior (Emons, 1997). The Bernie Madoff case is probably the most high-profile example of expert service fraud. Madoff, the now notorious but once highly regarded, financial adviser single-handedly caused \$17 billion in losses to over 1,000

individuals and firms in a highly publicized Ponzi scheme¹. Other industries have however been equally affected over the years, such as the U.S. auto repair business where consumers lost \$20 billion annually on faulty auto repairs (Brown, 1995), or the U.S. medical industry, where practitioners and providers paid, over the past decade, \$27 billion to settle federal allegations including fraud, bribery and patient harm (Berens, 2023).

Consumers who lack the ability to discern good from bad expert advice are thereby most likely to fall for bad sellers. Yet, "while consumer vulnerability is often invoked in consumer research, it is usually discussed informally" (Hill and Sharma, 2020). As such, we lack a clear understanding of the mechanisms that lead consumers to fall for and stick with bad sellers. This paper is the first to present and estimate a structural model that describes how consumers learn about sellers' quality, identifies personality traits that impact this learning, quantifies the price difference vulnerable clients are willing to pay for bad sellers as compared to their resilient counterparts, and suggests strategies to protect these consumers. The remainder of this introduction discusses our empirical approach, gives a snapshot of our results, and outlines our contributions. To facilitate readability, we will call the sellers who provide expert services from here on "advisers" and the consumers "clients".

While trust in an adviser builds up over time and can depend on the consequences of advice taking (e.g., portfolio performance, health outcomes, durability of a car), past research in marketing has also stressed the importance of first impressions on future customer interactions (Muthukrishnan and Chattopadhyay, 2007; Verhoef et al., 2009). We therefore construct a scenario similar to an initial meeting between a client and an adviser and model how this encounter may impact future advice taking. We use financial advice as the context of our study because choosing a good financial adviser is not straightforward, and there are costly consequences if a wrong choice is made, as recently highlighted in an article in the Wall Street Journal (Zweig, 2023). We collect data from a purpose-built, pre-tested, incentivized, online choice experiment and survey, completed by 2003 participants. Model-free analysis of sequences of responses shows that when advice is difficult for participants to evaluate, most will follow the adviser they chose previously, even if that adviser's advice is bad. When participants have low financial skills and high impulsiveness, this "stickiness" is more common. We then propose

¹A complete list of individuals and corporations that lost money can be found at this link https://s.wsj.net/public/resources/documents/st_madoff_victims_20081215.html

our structural model of biased belief formation that can explain why some clients persist with poor quality advisers.

In our model, clients learn about an adviser's quality from the advice she or he delivers on a short sequence of financial topics, as might be covered in a first advice meeting. Clients will rate advice about a topic that they find complex as an ambiguous signal of adviser quality, and advice about a topic they find easy as a clear signal. These ratings are subjective and vary between clients according to their personal understanding of financial topics.² Learning relies on memory, so the model allows clients to remember these signals and update their beliefs about the quality of the adviser. While all agents have some recall, it is implausible to assume that all agents have perfect recall (Nagel and Xu, 2021). We propose that clients fall into two types. The first type has perfect recall that allows them to conform to standard Bayesian belief updating. The second type has imperfect recall and follows limited memory updating that does not conform to Bayesian rules (Fryer et al., 2019). Bayesians remember whether past signals are ambiguous or clear, they set aside ambiguous signals and use only clear signals to update their beliefs. Limited memory learners do not remember the ambiguity of signals. They use every signal as it arrives to update beliefs, treating ambiguous signals as confirming their prior beliefs. Both types of agents update after clear signals of adviser quality, whether good or bad. And as limited-memory clients get new ambiguous signals they update again, so that their opinions become polarized, a typical feature of confirmation bias (Darley and Gross, 1983; Rabin and Schrag, 1999).³ Confirmation bias may lead clients to choose bad advice when good advice is available, and to pay for a poor adviser. The estimated structural model identifies the types of clients whose learning leads to biased beliefs. Simulations of the structural model show the extra fees that such vulnerable and biased clients will pay to poor advisers.

Our analysis yields several new insights. First, we find that confirmation bias is common. We estimate that close to two-thirds of participants use limited memory, rather than Bayesian, learning. We also find that impulsive consumers are more likely to rely on limited memory. While participants with better financial knowledge and numeracy can usually tell good advice

²Lusardi and Mitchell (2011) report wide variation in financial literacy in many countries. This variation is consistent with some advice topics being ambiguous to potential advice clients. See also Agnew et al. (2018)

³See Benjamin (2019) for discussion of instances where richer models of Bayesian updating can lead to polarization, such as where agents have ancillary information (Benoit and Dubra, 2018).

from bad advice, most participants are still unclear about one or more topics. It follows that the combination of limited memory and limited financial knowledge means that participants often use their (biased) prior beliefs about an adviser to guide choices.

Second, we find that participants' responses to an adviser differ by learning type. We show that, before an adviser gives any advice, a client will rate them as good if the client is generally predisposed to trust advisers and if the adviser shows a professional credential. Then we show how ambiguity of advice matters when confirmation bias is at play. In the presence of ambiguous advice, limited memory learners, in contrast to their Bayesian counterparts, will gradually converge to an almost certain belief that the credentialed adviser is good. Initial beliefs and learning process thus combine to establish trust in the adviser.

Third, we use simulations to show that segments of "vulnerable" and "resilient" clients differ by what they will pay advisers. By our definition, vulnerable clients are more predisposed to trust advisers. As a result, vulnerable clients have higher prior beliefs of adviser quality before beginning the experiment. They also show evidence of lower financial literacy which leads them to find more advice topics ambiguous. Vulnerable clients also reveal high impulsiveness which makes them more likely to use limited memory learning. We show that the additional fees that vulnerable clients are willing to pay, relative to resilient clients, increase as advisers give more bad advice. In fact, vulnerable clients bear costs from bad advice even when advisers do not aim to exploit their vulnerability. Our simulations also show that advisers can charge even greater fees, beyond the non-manipulative case, if they structure their advice to exploit clients' confirmation bias. Further, the characteristics that predict confirmation bias are not difficult to see, making this vulnerable group a potential target for bad advice.

Our findings add to the literature on the impact of first impressions and confirmation bias, the formation of consumer trust, and the vulnerability of consumers in contexts of asymmetric information. Previous research stresses the importance of trust in service encounters in general (Garbarino and Johnson, 1999; Sirdeshmukh et al., 2002) and in the investment advice industry in particular (Gennaioli et al., 2015; Gurun et al., 2018). Studies also document the effect of first impressions on trust formation (Konya-Baumbach et al., 2019; Evans et al., 2000) as well as the additional price clients are willing to pay a trusted service provider (Germann et al., 2018). We contribute insights from consumer learning to explain the foundation and evolution

of client trust in this context.

We also add to studies in marketing, economics and finance that test alternative psychological models of belief updating (e.g., Chylinski et al., 2012; Camacho et al., 2011; Benjamin, 2019). Our model is somewhat similar to that of Charness and Dave (2017), who test for, and find, differences in the way agents update when new information confirms or contradicts their beliefs. They use a confirmatory bias model from Rabin and Schrag (1999), where confirmation-biased agents misperceive disconfirming signals but update using Bayes' rule. By contrast, we base our model on Fryer et al. (2019). This model allows both Bayesian and limited memory clients to receive ambiguous signals, with confirmation bias arising from the updating rule used by limited memory clients. Unlike Rabin and Schrag (1999), our model allows individual differences in signal ambiguity, consistent with well-documented variation in financial literacy (Lusardi and Mitchell, 2011). Our study also relates to Augenblick and Rabin (2021), who set out theory and statistical tests for violations of Bayesian updating: our model detects confirmation bias in particular, rather than violations of Bayesian updating in general.

Some studies in financial advice settings do not set up a specific psychological model of updating but still find empirical evidence of confirmation bias. Experimental subjects in Zaleskiewicz and Gasiorowska (2018) rate financial advisers who gave "action advice" as more authoritative, but this rating was moderated by whether advice catered to clients' prior beliefs. Similar *prima facie* empirical evidence of confirmation bias is present in Agnew et al. (2018): polarization of clients' beliefs about advisers' trustworthiness followed a manipulated first impression, where advisers won favor with good advice on easy topics. Soll et al. (2022) show that receiving advice that agrees with a clients' initial opinion can exacerbate overconfidence. Our study builds evidence for a mechanism that can explain the impact of catering to beliefs and first impressions.

Our most important contribution is to the understanding of consumer vulnerability (Hill and Sharma, 2020), the identification of subpopulations that are at disadvantage (e.g. Bone et al., 2014) and the marketer manipulations that may affect vulnerable consumers' decision making (Langenderfer and Shimp, 2001). Langenderfer and Shimp (2001) report that officers at the Better Business Bureau - a private, non-profit organization that settles disputes between consumers and businesses - perceive victims of business scams to be significantly more trusting,

less educated and less self-controlled than consumers who did not fall for scams. Past research in the context of financial services (e.g., Stolper and Walter, 2019) highlights key factors that make financial service consumers more vulnerable. Egan et al. (2019) show that some advice firms target an unsophisticated clientele for exploitation. Similarly, Bucher-Koenen et al. (2021) confirm that advisers target less-financially-literate women for unsuitable products, while Georgarakos and Inderst (2014) show that less capable clients are more likely to follow advice if they trust their adviser. These studies identify vulnerable types of clients; our structural model for the first time provides insights into the mechanisms that make them vulnerable and identifies ways to help them.

This paper is structured as follows. In the next section, we explain the learning process that shapes our experimental design, describe the survey and report preliminary model-free evidence for confirmation bias. We then set out the structural model and estimation strategy. We present estimation results and simulations in the section that follows. The final section discusses the implications of the findings and concludes the paper.

EXPERIMENTAL APPROACH

In this section, we explain the learning processes used to design our data collection. Next, we outline the experimental survey that collects the data needed to identify two latent characteristics of consumers. These are their learning types and their subjective assessment of the ambiguity of quality signals. We then report descriptive statistics from our experiment in the context of financial advice and show model-free results as preliminary evidence that some participants have confirmation bias.

Learning In The Presence Of Ambiguous Signals

Learning can be understood as a hypothesis-testing process where new information is encoded and integrated with existing beliefs (Hoch and Deighton, 1989). When clients decide whether or not to follow financial advice in an area where they have little experience, they usually rely on signals about the quality of an adviser, such as past experience, or whether the adviser displays a qualification, to form a prior belief of the adviser's expertise (Agnew et al., 2018). As the client continues to interact with the adviser, new signals and further experience help them update these beliefs until they can make better-informed decisions.

A problem arises when new signals of adviser quality are ambiguous, that is, when the client is unsure whether the adviser is giving good or bad advice. When rational Bayesian learners get an ambiguous signal, they will recognize its ambiguity and act according to their prior beliefs of whether the adviser is a source of good or bad advice, in the absence of clear new information. They then ignore the ambiguous signal when forming their posterior belief of adviser quality, meaning that they will not use ambiguous signals to update their views on adviser quality (Roberts and Urban, 1988; Erdem and Keane, 1996). It follows that the future choices of Bayesian learners are guided by beliefs based only on unambiguous signals. Limited memory learners, however, do not update their beliefs in this way. Instead they make a "double update" (Fryer et al., 2019). First, they interpret ambiguous new information in line with prior beliefs, and then go on to update their posterior beliefs of adviser quality using this interpretation. These learners only recall their interpretation of an ambiguous signal and do not recall its ambiguity. As such, limited memory updating generates confirmation bias. Irrespective of the actual signal valence, for the purposes of updating posterior beliefs about adviser quality, people with confirmation bias will treat an ambiguous signal as positive if their prior beliefs are positive and will treat it as negative if their prior beliefs are negative. Such biased updating in turn leads to overconfidence, where people may come to believe with near certainty in a false hypothesis despite receiving an infinite amount of information (Rabin and Schrag, 1999).

We cannot directly observe whether a client uses Bayesian or limited memory learning. Instead, these are latent traits that we can only infer by observing the client's choices in a suitable setting. Further complicating the identification of these latent traits is the fact that we do not know when a signal is subjectively clear or ambiguous to a client. We therefore design an experiment and an associated model to enable us to make these inferences.

Survey And Choice Experiment

To identify confirmation bias, we need to observe learning from sequences of signals that are more or less subjectively ambiguous to the learners. Our incentivized, online experiment uses pre-tested components from the video advice experiment in Agnew et al. (2018) in a new design. As in Agnew et al. (2018), participants in our experiment watched videos of advice given by two advisers on four topics and chose the advice they would follow each time. The observed adviser pairs were selected from four possible advisers who differed by age, gender,

and displayed qualification. In addition, the order of the advice topics and the quality of advice given by each adviser varied by condition. Otherwise, the videos were identical (e.g., by shot, lighting, clothing). One adviser provided objectively correct, and the other objectively incorrect, advice on each topic. In the experiment, participants chose the wrong advice often, and at different rates for different topics. Based on the average number of wrong choices, participants (subjectively) found advice on two topics low on ambiguity, and advice on the other two topics high on ambiguity.

We purposefully modified the Agnew et al. (2018) design to fit our research questions. Since we focus on client traits rather than on adviser characteristics, we only include four adviser pairs, that mirror each other in gender, age and certification. We also neutralize the objective overall quality of the advisers; each adviser provides two objectively good, and two objectively bad, recommendations. Limited memory learning can best be detected when the two most ambiguous topics follow each other because only limited memory learners will update beliefs about the advisers after an ambiguous signal.⁴ So our design uses conditions that cover all possible combinations of good versus bad advice and high and low ambiguity topics.⁵ From the choice experiment data we can distinguish between Bayesian and limited memory learners. The survey also captures participants' willingness-to-pay for the advisers they see, so that we can newly estimate the marginal economic cost of a biased learning process. We use other survey responses to measure how these costs are moderated by clients' prior beliefs and abilities to discern good from bad advice, and how client characteristics matter.

Survey structure

The survey had four sections.⁶ Part one measured participants' financial literacy using standard questions on interest rates, inflation and diversification (Lusardi and Mitchell, 2011). This was followed by questions on numeracy skills (Lipkus et al., 2001), aimed at understanding the four advice topics covered later in the choice experiment, familiarity with a range of financial

⁴The experimental design of Agnew et al. (2018) shows participants a restricted topic sequence that does not allow for identification of the learning type employed by participants.

⁵When we combine the four possible pairs of advisers with the six possible sequences of objective advice quality and the six possible sequences of financial topics, we get a design with $6*6*4 = 144$ conditions.

⁶In parts one, three and four we replicate the instrument administered by Agnew et al. (2018) and in part two we include a choice task that we designed specifically to support estimation of our structural model. Web Appendix A shows screen shots of the survey.

products, and attitudes towards financial advisers. Part two delivered an incentivized, four-part choice task that used the video components developed by Agnew et al. (2018) but following a new experimental design. At the end, participants answered questions about their willingness to pay for financial advice. Part three collected 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. This section included two instructional manipulation (attention) checks (IMC1 and IMC2) (Oppenheimer et al., 2009). The last part of the survey comprised a debriefing that encouraged participants to consult a genuine adviser for financial decisions, identified the correct advice from the choice task and administered an incentivized quiz on the choice task topics.

Choice task components

The choice task showed videos of two advisers who gave financial advice on four common consumer finance topics: credit card debt repayment, retirement savings account consolidation, diversification in equity investments, and index fund fees.⁷ At the start of the experiment, individuals viewed a narrator in an introductory video. The narrator welcomed them to the study, noted the IRB (Ethics Board) approval, and explained the task. Following the introduction, the survey offered participants an incentive to choose the adviser giving the correct recommendation for each advice topic. Participants learned that for each correct choice, they would receive one entry for a \$A50 draw. A series of financial advice videos followed. For each advice topic, the video narrator introduced the situation and then two advisers appeared on the participant's screen. Adviser L was on the left and Adviser R was on the right. The advisers differed by age, gender, and whether or not they had a professional credential. The credential was displayed in text on the video for several seconds while the adviser video was played. Participants first watched Adviser L, followed by Adviser R. Importantly, participants saw the same two advisers for all advice topics and the advisers' positions on the screen did not change.⁸ The advice scripts had both advisers giving an identical introduction to each topic, followed by different recommendations for action, one objectively correct and the other objectively incorrect. After viewing the videos for the advice topic as many times as they liked,

⁷Agnew et al. (2018) gives a detailed justification for the choice of these topics.

⁸To view an example of the video advice from a treatment in Agnew et al. (2018), please follow this link: <https://www.youtube.com/watch?v=nCFBq5LIoQ0>.

participants chose the advice recommendation they would follow.

The four advice topics must satisfy two conditions to be suitable for our experiment. First, they need to have objectively correct answers, so that one adviser can deliver objectively good advice and the other objectively bad advice in the choice task. Second, they should be sufficiently difficult that at least some participants would treat them as subjectively ambiguous signals of the quality (expertise) of the video advisers. In our sample, and consistent with evidence from other studies, the chosen topics are difficult for some participants. Overall, participants chose correct advice slightly less than 80% of the time, with substantial variation by topic. Participants chose correct advice about retirement account consolidation 86% of the time, about stock diversification 79% of the time, about minimizing index fund fees 65% of the time, and about debt repayment 88% of the time. It follows that index fund fees and stock diversification are topics of higher average ambiguity and retirement account consolidation and debt repayment are of lower ambiguity. Table 1 reports the scripts for each topic.

By using the video components produced by Agnew et al. (2018), we guaranteed that the advice topics, advisers, environment and mode of advice delivery were uniform. The actors and their fictional names were pretested to ensure that participants would view each of them as equally credible. The production company was instructed to make sure that the actors were dressed, made-up and filmed in the same way. This design allows us to control the content of the videos, the order of advice topics, the quality of advice given and the adviser attributes.

Table 1: Advice Scripts

Notes: This table shows the scripts read by the narrator and the advisers for each of four advice topics. The narrator introduces each topic. Each adviser begins with the same (italicized) introduction then one adviser gives the good advice and the other gives the bad advice.

Narrator Introduction	Advice	Narrator Introduction	Advice
<p>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 two financial advisers will recommend what you should do about it.</p>	<p><i>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.</i> Good Advice: 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.</p> <p>Bad Advice: ... 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.</p>	<p>Choosing an Index Fund: In this scenario, you are thinking about investing in a managed share index fund. The next two financial advisers will recommend what you should do about it.</p>	<p><i>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?</i> Good Advice: These various share index funds provide an almost identical product 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.</p> <p>Bad Advice: ... 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.</p>
<p>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 two financial advisers will recommend what you should do about it.</p>	<p><i>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?</i> Good Advice: As a result, I recommend that you roll all of these accounts together so you are not paying extra fees.</p> <p>Bad Advice: ...Despite that, I recommend that you not roll all of these accounts together so you are diversified across different superannuation funds.</p>	<p>Diversifying a Stock Portfolio: In this scenario, you are thinking about investing in the share market. The next two financial advisers will recommend what you should do about it.</p>	<p><i>I understand you need help regarding how to invest your superannuation money. Did you know money invested in shares can go up and down?</i> Good Advice: 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.</p> <p>Bad Advice: ..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.</p>

Choice task design

We designed the choice task to identify participants’ learning mechanisms and then to measure participants’ willingness to pay the advisers. The key design feature for identifying learning type is the sequence of financial topics and their relative ambiguity. Limited memory learning, and therefore confirmation bias, can be best detected when the two most ambiguous topics follow each other. This is because, while both Bayesian and limited memory learners will rely on their prior beliefs to guide their advice choice when the topic is ambiguous, only limited memory learners will also update and form a new posterior evaluation of the advisers after their first choice.

Usually, a full factorial design is the most desirable choice when designing an experiment. However, in our case this is infeasible because the 28 $\binom{8}{2}$ possible combination of advisers, 16 possible sequences of good (G) and bad (B) advice, and 24 possible sequences of topics would require a total of 9,216 conditions. From this complete design, we chose a subset that was most informative for the identification of learning type.

Since we are focusing on client traits rather than on adviser characteristics, we only included four adviser pairings, where each adviser was the mirror image of the other (i.e., the older, female, certified adviser was paired with the younger, male, non-certified adviser). We also rendered the overall objective quality of the advisers to be neutral by only including six sequences of good and bad advice orders — those where each adviser gives two objectively good, and two objectively bad, recommendations.⁹ Thus, in contrast to advice orders where one adviser would be clearly superior to the other (e.g., by providing three or four pieces of objectively good advice), these six sequences should result in similar posterior beliefs about both advisers for participants who subjectively rate all topics as clear (unambiguous). As such, variations in response patterns for these advice orders are powerful in detecting participants’ subjective perceptions of topic ambiguity as well as their learning mechanism. Ideally, we would use all 24 possible financial topic sequences in our experimental conditions. Due to feasibility constraints, we include the six topic sequences that we rate as most informative for identifying learning and subjective ambiguity: those that present an equal number of high ambiguity (index fund fees and stock diversification) and low ambiguity (account consolidation and debt repayment)

⁹Agnew et al. (2018) found that clients considered advisers trustworthy and competent even if giving two pieces of bad advice.

financial topics.¹⁰ When we combined the four possible pairs of advisers with the six possible sequences of objective advice quality and the six possible sequences of financial topics, we get a design with $6 \cdot 6 \cdot 4 = 144$ conditions. We randomly assigned at least 10, and up to 14, participants to each condition (see Table 2).

We also newly measured the financial loss clients could incur when using limited memory updating on top of the financial consequences borne from bad decisions. To assign a dollar value to this loss, after the choice tasks we asked participants whether they are willing to pay $p \in \{\$50, \$100, \$150, \$250, \$500, \$750\}$ for a one-hour session with both, one, or none of the advisers they saw in the experiment. We randomly assigned fixed fee values X to participants to minimize the predictability of responses from the other manipulated characteristics of the experiment.

Survey Descriptive Statistics And Preliminary Evidence Of Confirmation Bias

We conducted the resulting four-part online survey that included the incentivized choice experiment in Australia in December 2014. Australia offers an ideal setting for this experiment, where at the time, and subsequently, regulators have been reviewing the qualification standards, fee structures and responsibilities of financial advisers (Bateman and Kingston, 2012). Australians face similar financial planning problems as consumers in other developed economies, and financial literacy levels are also similar, but participation in risky asset markets is more common than in, say, the US, since a large majority of adults also hold mandatory defined contribution savings accounts (Agnew et al., 2013).

¹⁰While Agnew et al. (2018) used a total of 256 experimental conditions, they only exposed participants to four out of the possible 24 topic sequences. Their sequences included only two combinations of high and low ambiguity topics, rather than the six possible combinations that we include in our design. The Agnew et al. (2018) design restriction to only two relative ambiguity sequences prevents the identification of participants' learning types. Note that our definition of topic ambiguity here refers to the aggregate ambiguity of the topic and is based on the number of incorrect answers to these topics. Individual participants' perception of topic clarity may deviate from these aggregate classifications - a possibility that we capture in our structural model (equations (5) and (6)). While we acknowledge the possibility that participants may confidently hold wrong beliefs about a topic (i.e., perceive a topic to be clear while at the same time identifying bad advice incorrectly as good and vice versa), our analysis in this section, as well as our structural model presented in section , assumes that if a topic is clear to a participant, they will hold correct beliefs about the advice related to that topic.

Table 2: Experimental Design

(a) Design of adviser pairs

Pair	Adviser L			Adviser R		
	Gender	Age	Certification	Gender	Age	Certification
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

(b) Sequence of advice topics

Sequence	Choice 1	Choice 2	Choice 3	Choice 4
1	Diversification	Fees	Consolidation	Debt
2	Consolidation	Debt	Diversification	Fees
3	Diversification	Consolidation	Fees	Debt
4	Consolidation	Diversification	Debt	Fees
5	Diversification	Consolidation	Debt	Fees
6	Consolidation	Diversification	Fees	Debt

(c) Design of sequence of advice quality

Quality Sequence	Advice from Adviser L				Advice from Adviser R			
	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: Each participant in the experiment makes four choices of financial advice where the design of the four choice sets consists of: one row from Panel (a) (adviser characteristics); one row from Panel (b) (sequence of advice topics); and one row from Panel C (sequence of delivery of good or bad advice from Adviser L and Adviser R). Panel (a) shows the combinations of adviser characteristics: each pair of advisers consisted of an adviser with three characteristics (gender, age, certification) and an adviser with the reverse. Adviser L appeared on the left-hand side of the choice set screen, and Adviser R appeared on the right-hand side. Each participant saw the same two advisers for the entire experiment, and each adviser stayed on the same side of the screen. Panel (b) shows the sequence of advice topics for each condition in the experiment. Panel (c) shows sequences of advice quality for each condition where “G” stands for objectively good advice and “B” stands for objectively bad advice.

Overview of survey responses

We invited members of PureProfile, an Australian nationally representative online panel, to participate. Respondents had to pass two screening questions to meet our age and gender quotas, resulting in 2,003 participants who completed the survey. Table 3 reports the characteristics of the sample compared with the Australian population, showing that the experiment sample is slightly more likely to be married and better educated than the population but otherwise is very similar. To ensure incentive compatibility, we compensated participants who completed the survey for their time (approximately \$A4) and rewarded them by giving one entry in a drawing for a \$A50 prize for each correct choice of financial advice in each of four choice sets and for each correct answer in a post-experiment quiz as discussed above. The majority of participants completed the survey in under 30 minutes, and the entire data collection process took less than three weeks. Table 4 defines the variables used in analysis and estimation.

At the aggregate level (see Table 5), participants chose good over bad advice 79% of the time, and, as noted earlier, on average were more likely to choose the objectively correct advice about debt repayment and account consolidation than about index fund manager fees and stock diversification. Participants chose the young female adviser's advice more often, despite each adviser giving equal numbers of good and bad recommendations for each topic in total. They were also slightly more likely to choose the advice delivered by the adviser who showed a qualification next to their name.¹¹ We found that more than half of participants (55%) were willing to pay a non-zero amount for at least one of the two advisers although every adviser had given bad advice twice. As such, on top of the financial loss participants could incur by following such bad advice, many are also willing to pay for bad advisers.

¹¹These results are consistent with the findings in Agnew et al. (2018, Table 4).

Table 3: Sample Demographics

	Survey Sample	18-79 yrs Aust. Pop'n		Survey Sample	18-79 yrs Aust. Pop'n
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 stated	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 labor force	18%	29%	Not stated	0%	10%
Retired	12%	n.a.			
Not stated	0%	5%			

Notes: This table presents the demographics of the sample of 2003 participants drawn from a nationally representative online panel by email invitation in 2014 with the 2011 (most recent at the time of the experiment) Australian census.

Table 4: Variable Descriptions

Variable Name	X0	X3	X4	X5	Description
Constant	x	x	x	x	Constant; topic specific for X4.
<i>Adviser Characteristics</i>					
Displays NO credential	x				Indicator variable that equals 1 if only the adviser's name was displayed and -1 when "Certified Financial Planner" and adviser's name were displayed.
Price				x	Mean-centered 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.
<i>Advice</i>					
Good advice $q_{c,R}$					Indicator variable that equals 1 if adviser R gives objectively good advice g in choice set c and -1 if adviser R gives objectively bad advice b .
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.
<i>Participant Characteristics</i>					
Participant female			x		An indicator variable that equals 1 if the participant is female, -1 otherwise.
Participant older than 39 years			x		An indicator variable that equals 1 if the participant is older than 39 years, -1 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 ability to make financial decisions, -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.

Continued

Table 4 – *Continued*

Variable Name	X0	X3	X4	X5	Description
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, -1 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, -1 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, -1 otherwise. Questions test topics used in the advice experiment: credit card debt, index funds, account consolidation, diversification.
Conscientiousness		x			An indicator variable that equals 1 if the participant's conscientiousness is above the sample median, -1 otherwise. Participants rated themselves as organized, responsible, hardworking or careless (reverse coded) on a four-point scale. Ratings are averaged.
Impulsivity		x			An indicator variable that equals 1 if the participant's impulsivity is above the sample median, -1 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, -1 otherwise. Participants reported whether they owned a credit card (debt), units in an index fund (fees), a superannuation account (consolidation), and stocks (diversification).

Notes: This table reports definitions of variables used in estimation, where X_i are vectors of explanatory variables for the components of the model (consisting of elements marked with an “x” in the corresponding column). Variables are computed from responses to an online survey of a representative sample of 2003 Australian adults conducted in December 2014.

Table 5: Summary of Survey Responses

Good advice chosen	% of total choices	Participant characteristic	% of participants
<i>All topics</i>	79.48	Trust in financial advisors	
<i>Topic: Account consolidation</i>	85.87	<i>No trust</i>	38.19
<i>Topic: Stock diversification</i>	79.38	<i>Neutral</i>	29.36
<i>Topic: Index fund fee</i>	64.95	<i>Trust</i>	32.45
<i>Topic: Debt repayment</i>	87.72	Responsible for financial decision making	
Advisor Chosen		<i>Someone else</i>	6.99
<i>Younger male</i>	25.04	<i>Together with someone else</i>	28.91
<i>Older male</i>	24.01	<i>Alone</i>	64.1
<i>Younger female</i>	26.01	Paid for financial advice before	27.31
<i>Older female</i>	24.94	Passed IMC 1	91.56
Displays credential	51.75	Passed IMC 2	92.96
Willing to pay stated price			Median score
<i>For none of the advisors</i>	45.23	High confidence in own decisions	3
<i>For only one advisor</i>	30.71	Risk tolerance	2.31
<i>For both advisors</i>	24.06	High financial literacy	0.75
		High numeracy	0.33
		High product knowledge	0.5
		High conscientiousness	3.4
		High impulsiveness	2.5
		High market experience	0.44

Notes: This table summarizes survey responses for variables used in the estimated model. The sample consisted of 2003 participants. Variables are defined in Table 4.

Preliminary evidence of confirmation bias

Preliminary analysis of the choice patterns indicates that participants' beliefs about the advisers guide their subsequent choices. We focus on participants' persistence, that is, their tendency to stick with a previously chosen adviser. There are two possible reasons why participants would choose an adviser twice in sequence: i) the topic covered at choice 2 is unambiguously clear to the participant and the same adviser, who the participant chose at choice 1, then gives the correct advice at choice 2; and ii), the topic covered at choice 2 is ambiguous to the client so they stick with the adviser they chose at choice 1 because the participant rates that adviser to be higher quality. To rule out i), we focus on instances where an adviser first gives good advice and was chosen by the participant, but then provides bad advice. We then count the number of times the participant chooses that adviser next.

Table 6, Panel A, shows the percentage of persistent (and thus bad) choices by financial topic ambiguity sequence. Respondents persist with an adviser, and thus choose the wrong advice,

in 22% of observations. Yet, this number varies markedly depending on the ambiguity of the topics in the two choice sets. The percentages of persistent (wrong) choices are lowest if the topic at choice 2 is low ambiguity. Correspondingly, the higher percentage of persistent choices when the financial topic at choice 2 is high ambiguity shows that participants depend on their prior opinion of the adviser’s quality. Furthermore, the table shows that a less-well-understood topic occurring first also increases the percentage of persistent choices coming next. While this finding could suggest that participants reward correct answers on difficult topics by forming positive opinions of the adviser, it could also indicate limited memory updating. Even if the participant thought advice at the first choice was ambiguous and could not evaluate it objectively, they treat it as correct and update posterior beliefs accordingly. After updating with their interpretation of an ambiguous signal, the participant makes a more favourable rating of the adviser and chooses that adviser again at the next, ambiguous, but wrong choice.

Table 6: Persistent Adviser Choices by Topic Sequence and Participant Characteristics

	Persistent choices (%)	Total observations
<i>All</i>	21.94	3195
Panel A		
Financial topic sequence		
<i>low ambiguity</i> → <i>low ambiguity</i>	7.8	551
<i>low ambiguity</i> → <i>high ambiguity</i>	28.81	1,149
<i>high ambiguity</i> → <i>low ambiguity</i>	12.95	973
<i>high ambiguity</i> → <i>high ambiguity</i>	38.51	522
Panel B		
Participants		
<i>capable</i>	19.7	2,612
<i>non-capable</i>	34.8	583
<i>non-capable</i> & <i>impulsive</i>	38.2	387
<i>non-capable</i> & <i>non-impulsive</i>	28.1	196

Notes: Column 2 shows the percentage of observations where participants first choose an adviser who gives correct advice and then continues to choose this adviser at the next topic where the adviser gives incorrect advice. Column 3 shows the total number of observations where the respective topic sequence co-occurred with an adviser giving correct advice first and bad advice second in the experimental design (Panel A), or the total number of observations the respective participant type faced such an adviser. Panel A shows the percentage of persistent choices by financial topic ambiguity sequence, where account consolidation and debt payment are classified as low ambiguity topics and index fund fee minimization and stock diversification are classified as high ambiguity topics. Panel B shows persistence by participant characteristics where ‘capable’ (non-capable) participants have financial literacy and numeracy above (below) the median and ‘impulsive’ (non-impulsive) divides the non-capable group into those who score above (below) the median in their responses to the impulsiveness scale.

Table 6, Panel B, provides further evidence of limited memory learning among groups of

participants. Here we simply classify participants into "capable" and "non-capable" of making sound financial decisions, and focus again on persistent choices. (Later, when estimating the structural model, we use a more comprehensive selection of covariates to explain participants' decision making.) We define participants with above median financial literacy and above median numeracy as capable, and those with below median financial literacy and below median numeracy as non-capable. Around 19% of capable participants persist with the previously chosen adviser when that adviser gives bad advice at choice 2. But almost twice as many (34.8%) of the non-capable participants persist, reflecting the fact that non-capable participants are more likely to find topics ambiguous and rely on their prior beliefs about the advisers. We further split the non-capable respondents by above- or below-median impulsivity score. Previous research (e.g., Glimcher et al., 2007; Franken et al., 2008; Hou et al., 2011; Yang et al., 2016) has linked impulsivity to poorly-considered decisions. Table 6, Panel B, shows that non-capable but impulsive people are 10 percentage points more likely to persist with the adviser than their non-impulsive counterparts, indicating that these respondents may indeed use limited memory and update posterior beliefs about the advisers with their interpretation of the previous, ambiguous signals.¹²

This first analysis shows that participants vary by their ability to recognize signals of good and bad quality advice and that they rely on previous ambiguous signals to make judgements. Our main goal is to test a mechanism of an alternative model of belief updating that we hypothesize can explain these patterns. We conduct the test by setting up, estimating, testing and simulating a structural model based on the experimental data and participant survey responses.

¹²We note that inattention is a facet of impulsivity (Weafer et al., 2013). In our survey we find that impulsive respondents are less likely to pass the instructional manipulation checks IMC1 and IMC2. Excluding those respondents who failed the IMCs, however, does not change the substantive findings of this section. We also note that impulsive respondents are more likely to fail the debrief quiz, which may be either due to inattentiveness or pronounced confirmation bias, leading them to not override their well established opinions. In the structural model in the next section we therefore only include the personality trait of impulsivity to account for inattentiveness.

STRUCTURAL MODEL OF CONSUMER LEARNING AND CONFIRMATION BIAS

We now set out the structural model that allows us to identify two latent processes underlying participants' advice choices and willingness to pay. The first process is the participant's perceptions that signals of adviser quality, delivered as topic-specific advice, are (subjectively) either clear or ambiguous. These perceptions of clarity or ambiguity vary within a participant by the topic of the advice in each choice set. The second latent process is the participant's learning type, that is constant within a participant over choice sets.

Prior Beliefs Of Adviser Quality, Perceptions And Interpretations Of Advice Signals, And Belief Updating

In the experiment, participant i receives a sequence of adviser quality signals delivered by two different advisers, $a \in (L, R)$, appearing on the left or right of the participant's screen. The signals of adviser quality come in the form of advice on four different financial topics, one for each choice set, $c = 1, \dots, 4$. Participants (subjectively) perceive quality signals as either clear ($\tau_{i,c} = clr$) or ambiguous ($\tau_{i,c} = amb$), depending on the topic of the advice.¹³

Let $q_{i,c,a} \in (\bar{g}, \bar{b})$ indicate the **objective** quality of the advice given at choice set c by adviser a where \bar{g} indicates objectively good advice and \bar{b} indicates objectively bad advice. Let $\sigma_{i,c,a} \in (g, b, gb)$ indicate participant i 's **subjective** interpretation of the signal at choice set c by adviser a : g indicates an interpretation that the advice is good, b indicates an interpretation that advice is bad; and gb indicates an undefined interpretation.

We assume that at choice set c , participant i holds a prior probability that adviser a is good quality denoted by $\lambda_{i,(c-1),a}$. We assume that if i perceives advice to be clear ($\tau_{i,c} = clr$) their interpretation of the quality of the advice will be consistent with the objective quality, $q_{i,c,a} \in (\bar{g}, \bar{b})$ ¹⁴. If i thinks the advice is ambiguous ($\tau_{i,c} = amb$) their interpretation of the quality of the advice will be consistent with their prior belief about the adviser $\lambda_{i,(c-1),a}$, interpreting advice from the adviser they hold in higher esteem as good and the advice from

¹³Our definition of ambiguity thereby differs from the one used in Fryer et al. (2019): The signal in our experiment has an objective quality, whereas in their model ambiguous signals are even objectively up to interpretation. A more general model would allow for degrees of ambiguity. We use a binary classification to help identification of the estimated parameters in the model.

¹⁴Note that this assumption precludes participants from confidently holding an *incorrect* belief about a signal's quality.

the adviser they rate lower as bad. As such, $\sigma_{i,c,a}$ depends on both subjective perceptions of clarity or ambiguity and prior belief about the advisers, so that,

$$\sigma_{i,c,a} = \begin{cases} g, & \text{if } \tau_{i,c} = clr \text{ and } q_{i,c,a} = \bar{g}, \text{ or } \tau_{i,c} = amb \text{ and } \lambda_{i,(c-1),a} > \lambda_{i,(c-1),\neg a}, \\ b, & \text{if } \tau_{i,c} = clr \text{ and } q_{i,c,a} = \bar{b}, \text{ or } \tau_{i,c} = amb \text{ and } \lambda_{i,(c-1),a} < \lambda_{i,(c-1),\neg a}, \\ gb, & \text{if } \tau_{i,c} = amb \text{ and } \lambda_{i,(c-1),a} = \lambda_{i,(c-1),\neg a}, \end{cases} \quad (1)$$

where $\neg a$ denotes the other adviser presented to the participant. We further assume that the advice given by an adviser carries a signal strength s ¹⁵ of the adviser's quality, that is the same for all participants.

Conditional on signal strength, s , updated beliefs depend on the participant's learning process and their perception of whether a signal is clear or ambiguous. Bayesian learners will update beliefs according to the Bayesian rule and ignore their subjective interpretation of the advice if the signal is ambiguous:

$$\lambda_{i,c,a} = \begin{cases} \frac{s\lambda_{i,(c-1),a}}{s\lambda_{i,(c-1),a} + (1-s)(1-\lambda_{i,(c-1),a})}, & \text{if } \tau_{i,c} = clr \text{ and } \sigma_{i,c,a} = g, \\ \frac{(1-s)\lambda_{i,(c-1),a}}{(1-s)\lambda_{i,(c-1),a} + s(1-\lambda_{i,(c-1),a})}, & \text{if } \tau_{i,c} = clr \text{ and } \sigma_{i,c,a} = b, \\ \lambda_{i,(c-1),a} & \text{if } \tau_{i,c} = amb, \end{cases} \quad (2)$$

That is, Bayesian learners will update beliefs of adviser quality only when they perceive the signal as clear. But limited memory learners will use all signals—both clear and ambiguous—to update their beliefs about the quality of the advisers. Limited memory learners update using their readily available interpretation of ambiguous signals. Based on Fryer et al. (2019), we assume that a limited memory learner updates beliefs according to:

$$\lambda_{i,c,a} = \begin{cases} \frac{s\lambda_{i,(c-1),a}}{s\lambda_{i,(c-1),a} + (1-s)(1-\lambda_{i,(c-1),a})}, & \text{if } \sigma_{i,c,a} = g, \\ \frac{(1-s)\lambda_{i,(c-1),a}}{(1-s)\lambda_{i,(c-1),a} + s(1-\lambda_{i,(c-1),a})}, & \text{if } \sigma_{i,c,a} = b, \\ \lambda_{i,(c-1),a}, & \text{if } \sigma_{i,c,a} = gb. \end{cases} \quad (3)$$

¹⁵The signal strength s , ($s > 1/2$) equals the probability that a good adviser gives objectively good advice, $P(\bar{g}|G) = s$, and that a bad adviser gives objectively bad advice, $P(\bar{b}|B) = s$. Thus the probability of receiving an objectively bad (good) signal from a good (bad) adviser is $P(\bar{b}|G) = 1 - s$ ($P(\bar{g}|B) = 1 - s$). To enable us to identify parameters, we set s to an arbitrary value greater than 0.5 and check the sensitivity of our estimation to alternative choices. The results we report below use $s = 0.75$.

Equations (2) and (3) show how ambiguous signals create an opportunity for confirmation bias to operate. Bayesian learners ignore ambiguous signals and form posterior beliefs using clear signals that generally will align with the objective quality of the financial advice. Thus, they gradually uncover the true quality of the adviser. However, limited memory learners use their interpretation of both clear and ambiguous signals to update their beliefs. Since their interpretation of the ambiguous signal can diverge from objectivity, being guided by prior beliefs, limited memory learners tend to confirmation bias.¹⁶

Table 7: Illustration of Belief Updating by Bayesian and Limited Memory Learner

Choice set (1)	Objective quality (2)	Clarity of advice (3)	Learning Process			
			Bayesian		Limited Memory	
			Beliefs updated (4)	Posterior beliefs (5)	Beliefs updated (6)	Posterior beliefs (7)
1	Good	Clear	Up	+	Up	+
2	Bad	Ambiguous	No	+	Up	++
3	Good	Ambiguous	No	+	Up	+++
4	Bad	Clear	Down	±	Down	++

Notes: The illustrative example should be read top to bottom, following the sequence of advice choices in column 1. Column 2 shows the objective quality of the advice given by an advisor, while column 3 shows the clarity of the topic and correspondingly whether advice signal is perceived as clear or ambiguous. Columns 4 and 6 show how the individual beliefs are updated, and columns 5 and 7 use graphical symbols to represent the probability assigned by the individual for the adviser to be high quality, for the Bayesian and limited memory learner respectively.

Table 7 gives a simplified example that compares the learning processes for one sequence of four choice sets for participants starting with neutral beliefs about the advisers. After four sets of advice, where adviser R delivered two signals that were objectively good and two that were objectively bad, but where both learners’ subjectively perceive the signals at choice sets 2 and 3 as ambiguous, the Bayesian learner has not changed their overall assessment of the adviser’s quality and the limited memory learner holds a more strongly positive posterior belief of the adviser’s quality.

¹⁶Equations (2) and (3) imply that beliefs about the two advisers are negatively correlated, irrespective of learning type.

Estimation Strategy

Linking observed and latent features of the model

We define a series of logit functions to estimate observed and latent features of the model. First, in the special case of choice set $c = 1$, we estimate the initial ($c - 1 = 0$) prior probability that adviser a is good quality to be a logit function of X_0 , where X_0 is a vector that includes information on whether the adviser displays a professional credential (Agnew et al., 2018) and on the participant’s general trust in advisers (Georgarakos and Inderst, 2014; Gennaioli et al., 2015) (see also Table 4), and an unknown vector of parameters, β_0 ¹⁷.

$$\lambda_{i,a,0} = \frac{\exp(\beta_0 X_0)}{1 + \exp(\beta_0 X_0)} \quad (4)$$

Second, we model the probability of choosing advice in each choice set. Let the binary variable $y_{i,c} = 1$ if participant i chooses R at choice c , and 0 if participant i chooses adviser L . We also assign a value of 1 to $q_{i,c,R}$ if adviser R gives objectively good advice \bar{g} and assign a value of -1 if adviser R gives bad advice \bar{b} . When participant i interprets the advice as clear, then they should choose to follow—up to some random error—the good advice and not the bad advice. As such, if the advice is clear to i , the probability that i chooses $q_{i,c,R}$ is given by:

$$P(y_{i,c} = 1 | \tau_{i,c} = clr) = \frac{\exp(\beta_1 q_{i,c,R})}{1 + \exp(\beta_1 q_{i,c,R})}, \quad (5)$$

where β_1 can be interpreted as the scale of the random error. Similarly, when the participant interprets the advice as ambiguous, we assume that the probability of choosing the advice from a particular adviser depends on their prior beliefs about the relative quality of both advisers. The higher the difference in beliefs about adviser R versus adviser L , the more likely the participant should be to follow adviser R (and vice versa). We again allow for a random error

¹⁷To assist with model identification, we define the X_i , $i \in \{0, 3, 4, 5\}$ using non-overlapping subsets of adviser and participant characteristics. Allocation of characteristics to the different subsets was informed by prior research.

in this choice and write the probability as:¹⁸

$$P(y_{i,c} = 1 | \tau_{i,c} = amb) = \frac{\exp(\beta_2(\lambda_{i,(c-1),R} - \lambda_{i,(c-1),L}))}{1 + \exp(\beta_2(\lambda_{i,(c-1),R} - \lambda_{i,(c-1),L}))}. \quad (6)$$

Equations 5 and 6 model the fact that both types of learners' interpretations of signals, and therefore choices of advisers, depend on the same factors: their perceptions of the signals as clear or ambiguous $\tau_{i,c}$ and, when the signal is ambiguous, their prior beliefs of adviser quality, $\lambda_{i,(c-1),a}$. Thus differences in interpretations (choices) for Bayesian and limited memory learners arise from differences in the evolution of $\lambda_{i,(c-1),a}$. Consistent with Equations 2 and 3, Bayesian learners recall if $\tau_{i,c} = amb$ when updating their posterior beliefs and do not change $\lambda_{i,c,a}$, but limited memory learners "forget" the ambiguity and update anyway. Hence Bayesian and limited memory learners can have divergent values of $\lambda_{i,(c-1),a}$ entering Equation 6.

Third, we estimate two unobserved latent classes of learner types indicated by $\tau_{Learner} \in$ (Bayesian, limited memory) and 16 unobserved latent classes of subjective clarity/ambiguity interpretations, $\tau_{Clarity} \in (1, \dots, 16)$. The 16 latent clarity classes span the space of clarity/ambiguity possibilities for all four advice topics.¹⁹ The probability that i is a member of any latent clarity class is the product of the probabilities that each topic is subjectively clear to the participant:

$$P_i(\tau_{Clarity}) = \prod_{c=1}^4 P_i(\tau_{i,c} = clr). \quad (7)$$

In the interest of parsimony, we assume the joint probability of being a Bayesian or limited memory learner, and of finding a sequence of advice topics clear or ambiguous is $P_i(\tau_{Clarity}, \tau_{Learner}) = P_i(\tau_{Clarity})P_i(\tau_{Learner})$. However, we accommodate participant-level dependence between the latent classes by modelling class membership probabilities, $P_i(\tau_{Learner})$ and $P_i(\tau_{Clarity})$, as logit functions of participant-specific covariates, X_3 and X_4 , and associated parameter vectors, β_3 and β_4 , respectively (see Table 4).

We allow participants' conscientiousness and impulsivity to predict their learning type. Con-

¹⁸In both cases, we estimate the scale parameters $\frac{1}{\beta_1}$ and $\frac{1}{\beta_2}$ of the extreme value distribution of the random components. As $\beta_1(\beta_2)$ approaches infinity, the expression on the right-hand side of Equation (5) (Equation (6)) approaches 1 for $q_{i,c,R}$ ($\lambda_{i,c,R} > \lambda_{i,c,L}$).

¹⁹This is an important difference between the treatment of advice ambiguity in the study of Agnew et al. (2018) where topics were defined as clear or ambiguous for all participants based on observed mistake rates by participants.

scientiousness is a Big Five personality trait related to orderliness, industriousness and effort (Roberts et al., 2014; Bidjerano and Dai, 2007; Chamorro-Premuzic and Furnham, 2003; Richardson and Abraham, 2009; von Stumm and Furnham, 2012). We hypothesize that conscientious individuals will be more likely to apply cognitive effort to evaluating new information, consistent with Bayesian learning, and less likely to use memory shortcuts, consistent with limited memory learning. Impulsivity is the tendency to act without thinking (Glimcher et al., 2007; Franken et al., 2008; Hou et al., 2011; Yang et al., 2016). Impulsive people tend to rely on information that is close to hand to form beliefs, without first thinking about whether it is actually correct. Because people can more easily remember their interpretations of signals than the past actual signals, we hypothesize that impulsive people are more likely to be limited memory learners.²⁰ We model participants’ probability of perceiving advice about one of the four topics as dependent on participants’ financial literacy, numeracy, product knowledge and market experience (Agnew et al., 2018).

Fourth, posterior beliefs formed over the four choice sets ($\lambda_{i,c,a}|c = 4$) directly influence participants’ willingness-to-pay for the advice provided by each adviser. We offer the option to purchase additional advice from each adviser at a fixed price p_i . Let $w_{i,a}$ be an indicator variable taking the value of 1 if participant i is willing to pay price p_i for advice from adviser a . We estimate a logit function of the probability that i is willing to pay for adviser a as dependent on a vector of covariates, X_5 , that includes attributes of the participant and the adviser (Stolper and Walter, 2019), the price of an additional unit of advice, and the estimated posterior beliefs about the adviser’s quality, $\lambda_{i,4,a}$ (Gennaioli et al., 2015), and parameters β_5 (see Table 4).

Model likelihood and estimation

The likelihood function of our structural model is based on participants’ adviser selections, replies to the willingness-to-pay questions and other survey information. Conditional on participant i being a member of a unique pair of (16) latent clarity and (2) learning classes,

²⁰Impulsivity is a facet of the Big Five personality factors, but its factors remain a subject of debate. Some argue it is a facet of conscientiousness, while others view it as a facet of neuroticism. Still others see it as a blend of several factors (Borghans et al., 2008). Given this debate, we believe that impulsivity may have explanatory power over and above conscientiousness. Based on Tsukayama et al.’s (2011) findings, we use a finance domain-specific measure of impulsivity, which should provide greater predictive power in our context.

$(\tau_{Clarity}, \tau_{Learner}) = \tau_j \in j = 1, \dots, 32$, we can write the likelihood of i 's sequence of four advice choices and two decisions over willingness to pay for advisers R and L as:

$$\begin{aligned}
 l_i(\theta|\tau_j) = & \prod_{c=1}^4 P(y_{i,c} = 1|\tau_j)^{y_{i,c}=1} \\
 & \cdot P(w_{i,R} = 1|\tau_j)^{w_{i,R}=1} \\
 & \cdot P(w_{i,L} = 1|\tau_j)^{w_{i,L}=1}.
 \end{aligned} \tag{8}$$

The resulting *unconditional* likelihood of participant i 's sequence of choices is then given by the probability-weighted sum over all possible latent class pairs τ_j :

$$l_i(\theta) = \sum_{j=1}^{32} P_i(\tau_j) l_i(\theta|\tau_j). \tag{9}$$

Because the likelihood function exhibits discontinuities along the dimensions of the parameters of beliefs, it presents estimation problems for standard maximum likelihood methods. We thus use the Sequential Adaptive Bayesian Learning (SABL) algorithm developed by Durham and Geweke (2014) to estimate the model. The SABL Bayesian estimation algorithm that we use does not rely on direct maximization of the likelihood function and therefore allows us to overcome this challenge. SABL is an extension of sequential Monte Carlo methods and additionally exploits the benefits of parallel computing environments. SABL does not require the researcher to specify conjugate priors, and it is also robust to multimodal posteriors that can arise in high-dimensional problems (Jasra et al., 2007), such as ours. In Web Appendix B we provide a discussion of the discontinuity problem and outline the estimation procedure.

Sketch of Parameter Identification

A formal analysis of identification is not feasible for our complex, nonlinear learning model. Here, we sketch our identification strategy for the key model parameters.

First, consider participant i 's initial prior probability of adviser a 's quality, that is, $\lambda_{i,(c-1),a}|c = 1$. This prior influences the choices i makes at $c = 1$: we assume that (up to a random error) if participant i subjectively interprets the signal from a at $c = 1$ to be ambiguous, $\tau_{i,c} = amb$, they will choose the adviser with the higher value of $\lambda_{i,(c-1),a}|c = 1$. Since the experimental design shows participants any of the four possible topics at $c = 1$,

some participants will find their assigned topic difficult, and thus view the advice signal as ambiguous (Panel B in Table 2), and interpret the signal using their initial prior. Variation in topic ambiguity at the first choice set thus identifies initial priors and how the prior depends on advisers' certification and participants' trust (β_0).

We identify the parameter β_2 via the initial prior $\lambda_{i,(c-1),a}|c = 1$ and assumed signal strength s . These two parameters jointly define the updated probabilities $\lambda_{i,c,a}$, so we treat them as predetermined covariates when participants face an ambiguous topic $\tau_{i,c} = amb$. Choices made in choice sets with ambiguous signals can thus identify β_2 . Choices made in choice sets with clear signals $\tau_{i,c} = clr$ identify β_1 .

Participants' choices of financial advice allow us to identify membership of the 16 latent classes $\tau_{Clarity}$. Specifically (up to uncertainty in the choice process), if the participant selects adviser R ($y_{i,c} = 1$), when R gives an objectively bad quality signal $q_{i,c,R} = \bar{b}$, we can conclude that the signal was ambiguous for that participant, $\sigma_{i,c,R} = gb$. (We cannot make a similar inference if the participant selects R when R gives an objectively good quality signal, $q_{i,c,R} = \bar{g}$, as this could imply either that i treated the signal as subjectively clear, $\tau_{i,c} = clr$, or that i treated the signal as ambiguous, $\tau_{i,c} = amb$, but $\lambda_{i,(c-1),R} > \lambda_{i,(c-1),L}$.) The combined information of posterior beliefs of adviser quality and objectively incorrect choices of advice thus allows us to identify $\tau_{Clarity,i}$.

Since $\lambda_{i,(c-1),a}|c = 1$ can be inferred from the data without any assumptions about participants' learning type, and since signal strength s is fixed, we can calculate the posterior beliefs $\lambda_{i,c,a}$ for both learning types. The posterior belief associated with the higher likelihood then helps to determine $\tau_{Learner}$.

ESTIMATION AND RESULTS

Model Fit And Parameter Estimates

To begin, we assess the fit of our model. We estimate the model in SABL using data from 1,903 of the 2,003 participants and reserved the remaining participants' responses to assess hold-out fit. (Table 4 provides the variables and associated definitions used in the estimation.) Table 8 reports the parameter 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, denoted with an asterisk. Of the 26 parameters we estimate, 20 satisfy this condition, and we concentrate on these effects in the interpretation below. The interpretation of the CIs is largely analogous to a frequentist 95% confidence interval.

Table 8: Estimated Parameters

	Mode	2.5 Percentile	97.5 Percentile
<i>Initial Prior Belief about Adviser: β_0</i>			
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
<i>Choice of Advice: β_1, β_2</i>			
Quality (R gives correct advice)	4.296*	3.599	5.060
Prior belief	2.510*	1.494	3.622
<i>Bayesian and Limited Memory Updating: β_3</i>			
Constant	-0.454*	-0.994	-0.185
High conscientiousness	0.243	-0.046	0.485
High impulsiveness	-0.344*	-0.724	-0.154
<i>Topic Ambiguity: β_4</i>			
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
<i>Willingness-to-Pay: β_5</i>			
Constant	-7.782*	-9.687	-6.228
Price	-0.085*	-0.124	-0.043
Posterior	18.309*	14.808	22.23
Paid for advice	0.466*	0.348	0.57
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

Notes: This table reports statistics from the posterior belief distributions of estimated parameters of the choice model (Equation 9). Estimates are based on the 1,903 randomly selected participants from 2,003 responses collected in December 2014. Variables are defined in Table 4. For each parameter, we report the mode of its posterior distribution, as well as the 2.5 and 97.5 percentiles of this distribution, i.e., the equi-tailed credible interval (CI). There is a 95% probability that the parameter is not zero if zero does not fall in the CI. The mode includes an asterisk in those cases. Estimation was conducted using SABL; see Online Appendix B for details.

Overall, in-sample fit is satisfactory, and hold-out sample fit is close to in-sample fit, which

shows that our model does not overfit the data. The model also has discriminatory predictive power: for the estimation sample, it predicts an average (over all choice sets) probability of 0.69 that the adviser who is in fact chosen in the data would be chosen and predicts an average probability of 0.69 for the hold-out sample. When the adviser is not chosen in the data, the average model-predicted choice probability that the adviser would be chosen decreased to 0.29 for the estimation sample and 0.28 for the hold-out sample.

The predicted probabilities are less discriminating in the willingness-to-pay probabilities. When a participant chose to pay the adviser, the model’s average predicted probability that the participant would choose to pay is 0.48 for the estimation sample and 0.44 for the hold-out sample. When a participant chose not to pay the adviser, the average model-predicted probability that the participant would pay is 0.28 for the estimation sample data and 0.34 for the hold-out sample data. Thus, the model slightly underestimates the probability that a participant will pay the assigned fee for the adviser.

We also compared our model to a restricted model that allows for only rational Bayesian updating (where $\tau_{Learner} = \text{Bayesian}$ for all participants), consistent with conventional learning models. The log marginal density for this restricted (rational) model is -5887.17, compared to a log marginal density of -5827.47 for our model for in-sample fit and -305.62 versus -298.35 for hold-out sample fit. These log marginal densities translate into Bayes factors that suggest that there is strong evidence against the restricted model based on the in-sample fit and substantial evidence against the restricted model based on the hold-out sample fit.

Latent Learning And Topic-Clarity Classes

A key result from our model is the prevalence of limited memory updaters who are susceptible to confirmation bias. Table 9 presents the percentages of participants assigned by the model to latent classes. We estimate that a significant majority – 63% – of the participants make choices consistent with limited memory updating. The model also assigns participants to 16 latent topic clarity classes. The largest group of participants (21.9%) belongs to the class that treats every topic as clear except for index fund fees, followed by 18.2% who see all the topics as clear. The third largest class (14%) of participants perceived only debt and account consolidation as clear. This latent class assignment implicitly ranks topic clarity – with fund fees as the

most unclear topic and debt and consolidation as the clearest. This ranking is consistent with estimates of topic-specific constants (related to expected correct choices) in Table 8 and with rates of correct choices in the raw data in Table 5.

Table 9: Proportion of Participants in Latent Classes

Latent Class		Segment Size (%)				
<i>Learning process</i>						
Bayesian updater		37.11				
Limited memory updater		62.89				
<i>Clarity of topics</i>						
Clarity class	Consolidation	Diversification	Fees	Debt		
1	1	1	1	1	18.2	
2	1	1	1	0	2.2	
3	1	1	0	1	21.9	
4	1	1	0	0	4.5	
5	1	0	1	1	6.9	
6	1	0	1	0	1.4	
7	1	0	0	1	14	
8	1	0	0	0	4.8	
9	0	1	1	1	3.1	
10	0	1	1	0	0.6	
11	0	1	0	1	6.4	
12	0	1	0	0	2.3	
13	0	0	1	1	2	
14	0	0	1	0	0.7	
15	0	0	0	1	7.2	
16	0	0	0	0	3.8	

Notes: This table shows the estimated posterior belief percentage of participants assigned to 2 latent classes differentiated by the learning process (Bayesian or limited memory) and 16 latent classes differentiated by the clarity or ambiguity of the four advice topics. A “1” indicates that participants in that class treated the topic as clear, and “0” indicates that they treated the topic as ambiguous. For example, the model assigns 18.2% of participants to latent class 1 (row 1), which treats all topics as clear, and assigns 3.8% of participants to latent class 16 (row 16), which treats all topics as ambiguous. We infer latent classes from estimation of the choice model (Equation 13) – see Table 8 for estimation results.

Initial Prior Beliefs, Latent Class Membership And Willingness-To-Pay: Marginal Effects

We now use estimates from Table 8 to show how covariates affect the outcomes predicted by the model. These outcomes are the probability that participants hold an initial prior that an adviser is good quality, that participants follow one of the two latent learning processes, or perceive topics as clear or ambiguous, and the probability that the client is willing to pay for another advice session with the adviser. We set all other variables in the respective equations

at their median values and calculate the probabilities at the minimum and maximum values for the variable of interest.²¹

Client's belief about an adviser – and consequently their willingness-to-pay for this adviser – depends on the belief they hold about the particular adviser before they receive the first piece of advice (initial prior belief). On average, participants believe that an adviser is good with a probability of 85% (determined by the constant – 1.728). This probability is very close to the number of respondents who reported that they trusted the advice from their adviser "a lot" in the ASIC shadow shopper study (81%, Australian Securities and Investments Commission, 2012) and thus provides further face validity to our results. The prior belief about an adviser's quality also depends on the participant's general trust in financial advisers and whether the adviser displays certification in our model. In Table 8, CIs exclude zero for both the trust and certification parameters. We find that the difference in initial prior beliefs between a participant reporting no prior trust in advisers to one reporting trust is large: participants hold a 13.45 percentage point higher likelihood of the adviser being good if they display a general tendency to trust advisers. Similarly, certification matters, which is consistent with previous studies (Guiso et al., 2008; Agnew et al., 2018). An adviser certification increases the mean participant's prior by 2.22 percentage points.

The model uses personal characteristics, financial literacy and knowledge to estimate the probability that participants follow one of the two latent learning processes, or perceive topics as clear or ambiguous. We report on parameters where the CIs in Table 8 exclude zero. The probability that a participant is a rational Bayesian updater is influenced by personality traits. We find that participants who are not impulsive are 8.31 percentage points more likely to be Bayesian updaters. This is consistent with our earlier hypothesis. The probability that a participant finds an advice topic to be clear increases significantly with prior product knowledge (6.54 percentage points), increased age (13.75 percentage points), higher financial literacy (9.10 percentage points) and higher numeracy (8.75 percentage points).

In addition, the estimation shows that the only participant characteristic that influences willingness-to-pay for advice, apart from price and participants' posterior beliefs about the adviser, which we will discuss in greater detail in the next section, is whether the participant

²¹See Web Appendix C for graphs of the marginal effects associated with estimated parameters.

has previously paid for financial advice. Participants who had previously paid for financial advice were 5.92 percentage points more likely to pay for advice than those who had not previously paid for advice. Since we condition on other covariates such as income, confidence in one's own financial capabilities, risk aversion and whether the participant is the household's financial decision maker, this result provides further evidence that, overall, clients evaluate their interactions with financial advisers as positive.

Comparison Of Adviser Choice By Learning Type

The previous section discusses how the certification of advisers influences initial prior beliefs about advisers. In this section, we examine how the impact of certification on prior beliefs about advisers indirectly influences participants' choice of advice. Our model allows us to do so while simultaneously accounting for participants' learning strategies (Bayesian or limited memory), as well as their ability to subjectively discern good from bad advice (i.e., clarity of topics). To illustrate, consider two participants, A and B, who update their beliefs according to the standard Bayesian and limited memory updating processes, respectively. Let the adviser who appeared on the right-hand side of the choice screen – the right adviser (R) – display a certification (+ 0.085 from Table 8), and the left adviser (L) not display a certification (-0.085 from Table 8). We arbitrarily assume that the participants both *distrust* financial advisers (the negative of the trust mode of 0.52 from Table 8), set other characteristics at the medians of the survey sample distributions, and fix estimated parameters at the mode of the posterior distributions. Both participants thus hold the same initial beliefs about the two advisers, with both judging the right (R) adviser only slightly better than the left (L) adviser (0.785 vs. 0.755) with the difference likely caused by adviser R possessing a certification, while adviser L does not. Assume that both advisers give (from both participants' perspectives) ambiguous advice in all four choice sets.

Table 10 shows the evolution of beliefs and choice probabilities for the advisers in this scenario. Both A and B have choice probabilities of 0.519 and 0.481 for Adviser R and L, respectively, in the first choice set (see Equation (6)). Yet, since the topic in this first choice set is ambiguous, participants differ in how they update their beliefs afterwards. Participant A's beliefs about the advisers, as well as the associated choice probabilities, remain the same throughout the later three choice sets, as this participant does not update their posterior with interpretations

of the ambiguous information. Participant A thus ends the experiment still only slightly favoring adviser R. In contrast, participant B updates their posterior using interpretations of all ambiguous advice signals. Since participant B begins with an initial prior that is slightly higher for the certified adviser R, they treat all subsequent ambiguous information as evidence that adviser R is good and adviser L is bad. Limited memory updating leads to a choice probability in the fourth choice set that is very close to one for adviser R (0.921) and close to zero for adviser L (0.079), while the corresponding probabilities stay at 0.519 (adviser R) and 0.481 (adviser L) for the rational updater. These calculations show the difference that initial priors (here stemming from the display of a certification) make, and how these are intensified in each period by confirmation bias. Even a small difference has a strong influence on the limited memory updater, whose opinion approaches certainty over a few choices.

Table 10: Evolution of Beliefs with Four Ambiguous Topics

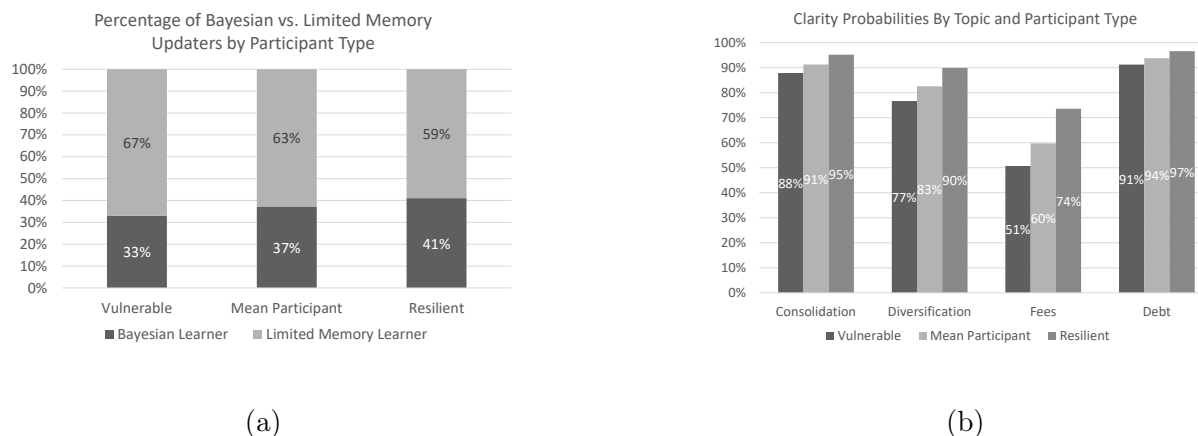
	$\lambda_{i,0,a}$	$Pr(y_{i,1} = 1)$ $\tau_{i,1} = amb$	$\lambda_{i,1,a}$	$Pr(y_{i,2} = 1)$ $\tau_{i,2} = amb$	$\lambda_{i,2,a}$	$Pr(y_{i,3} = 1)$ $\tau_{i,3} = amb$	$\lambda_{i,3,a}$	$Pr(y_{i,4} = 1)$ $\tau_{i,4} = amb$	$\lambda_{i,4,a}$
<i>Bayesian Updater A</i>									
Adviser R	0.785	0.519	0.785	0.519	0.785	0.519	0.785	0.519	0.785
Adviser L	0.755	0.481	0.755	0.481	0.755	0.481	0.755	0.481	0.755
<i>Limited Memory Updater B</i>									
Adviser R	0.785	0.519	0.916	0.737	0.970	0.858	0.990	0.903	0.997
Adviser L	0.755	0.481	0.506	0.263	0.255	0.142	0.102	0.097	0.037
$\lambda_{i,0,a}$ prior belief about adviser quality at choice set i									
$Pr(y_{i,1} = 1)$ probability of choosing to follow advice of adviser at choice set i									

This table presents the sequence of choices when clients use either a standard Bayesian (Participant A) or limited memory (Participant B) process to update beliefs about adviser quality. In the example, we assume that both participants are initially distrusting of financial advisers, and that otherwise both participants have characteristics at the medians of the sample distributions. Parameters are set to the modes of the posterior distributions. Adviser R shows a professional certification, and Adviser L does not. Both participants thus have the same prior beliefs that the right (R) and the left (L) advisers are of good quality, $\lambda_{i,0,a}$. All topics are ambiguous to both clients, i.e., choice probabilities are calculated based on Equation (6). Bayesian Updater A’s beliefs about the advisers and choice probabilities remain constant because the rational client does not update using ambiguous signals (see Equation (2)). Limited memory Updater B treats ambiguous information as evidence in favor of their priors and continues to update in favor of Adviser R (see Equation (3)).

Vulnerable Clients’ Willingness To Pay: Simulation Results

The previous section showed how participants’ prior beliefs about advisers, learning strategies and ability to discern good from bad advice (i.e., subjective ambiguity of advice topics) can impact the choice of adviser. We now simulate the model to measure the costs potentially

Figure 1: Breakdown of Participants by Learning and Topic Clarity Types



Panel (a) reports the simulated percentage of Bayesian versus limited memory updaters for each type of participant (vulnerable, mean and resilient). Panel (b) reports the simulated probabilities that the different types of participants perceive each advice topic as clear.

carried by vulnerable clients.

Specifically, we divide the sample into “resilient” and “vulnerable” participants and a benchmark reference group with mean characteristics. First, we define vulnerable (resilient) clients as those who show above (below) sample median impulsivity and therefore are more (less) likely to use limited memory processing, who are more (less) predisposed to trust advisers and thus have higher (lower) prior beliefs of adviser quality, and who score below (above) the sample median for financial literacy and numeracy. Based on estimates of topic ambiguity class membership, we also include younger people in the vulnerable group. We can identify these participants by using their responses to survey questions answered before the choice task.

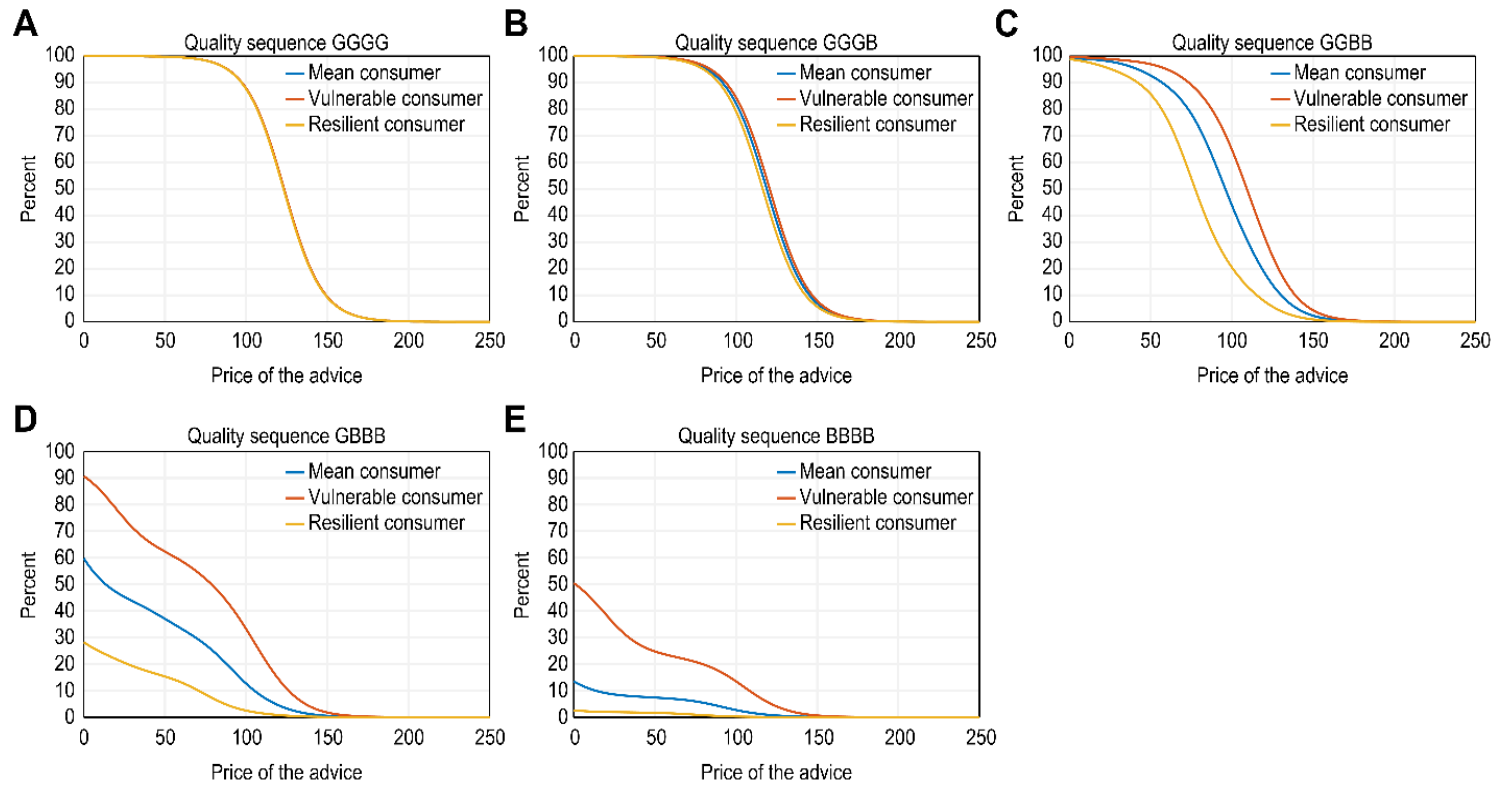
Figure 1, Panel A shows how the resilient, vulnerable and mean reference groups break down into Bayesian versus limited memory updaters. We observe that the percentage of limited memory learners is much higher for the vulnerable group than for the resilient group (67% versus 59%). In Panel B, the probabilities that resilient and vulnerable groups understand advice topics are also very different. For instance, for the fee topic, which is considered the most difficult topic, the difference between the two groups is 23 percentage points (vulnerable participants have a 51% probability of perceiving this topic as clear compared to resilient participants who have a 74% probability).

Figure 2 plots the probability distribution of each group's willingness-to-pay various adviser fees for different sequences of good versus bad advice. Each panel represents a different sequence of advice quality. Going from Panel A to Panel E, we move from a "best-case world" where all the advice given is good (GGGG) to a "worst-case world" where all the advice given is bad (BBBB). The probabilities are averaged over all possible matchings of clear and ambiguous topics at each of the four advice points in the sequence. Thus, while the probabilities in these figures are averaged over topics and therefore are not impacted by strategic manipulation by advisers, for example, by offering good advice on a clearly understood topic during the first meeting, they still account for the fact that, on average, vulnerable clients will perceive more topics as ambiguous. The figure shows that for all groups of participants, willingness-to-pay for an adviser decreases as more bad advice is given, as shown by the shift of the lines from Panel A to Panel E towards the horizontal axis.

Notice in the best-case world (GGGG) in Panel A, the lines lie on top of each other. This shows that the probability of paying for an adviser at each price is the same for resilient (yellow line) and vulnerable clients (red line). Therefore, in a simplistic world where advisers only give good advice, at average topic clarity, resilient and vulnerable clients are willing to pay the same costs. This changes as bad advice is added. In Panel B, the lines separate from one another when one piece of bad advice is introduced at the end of the sequence (GGGB). The separation widens as additional bad advice is added in Panels C through E.

The penultimate Panel D shows the power of a good first impression reinforced by confirmation bias. In this advice sequence (G BBB), only the first advice given is good. This first impression translates into a large difference between what the resilient and vulnerable participants are willing to pay for an adviser. At each price point until just after \$A150, the vulnerable client is likely to pay more than the resilient client. Finally, Panel E displays the worst-case world, where the adviser only gives bad advice (BBBB). In this case, clients should be unwilling to pay anything for further advice from the adviser. Unfortunately, vulnerable clients are much more willing to pay for this adviser than resilient clients who find it easier to distinguish good from bad advice and are thus better able to recognize the adviser's lack of value.

Figure 2: Willingness-to-Pay from a “Best-Case World” to a “Worst-Case World” by Participant Type



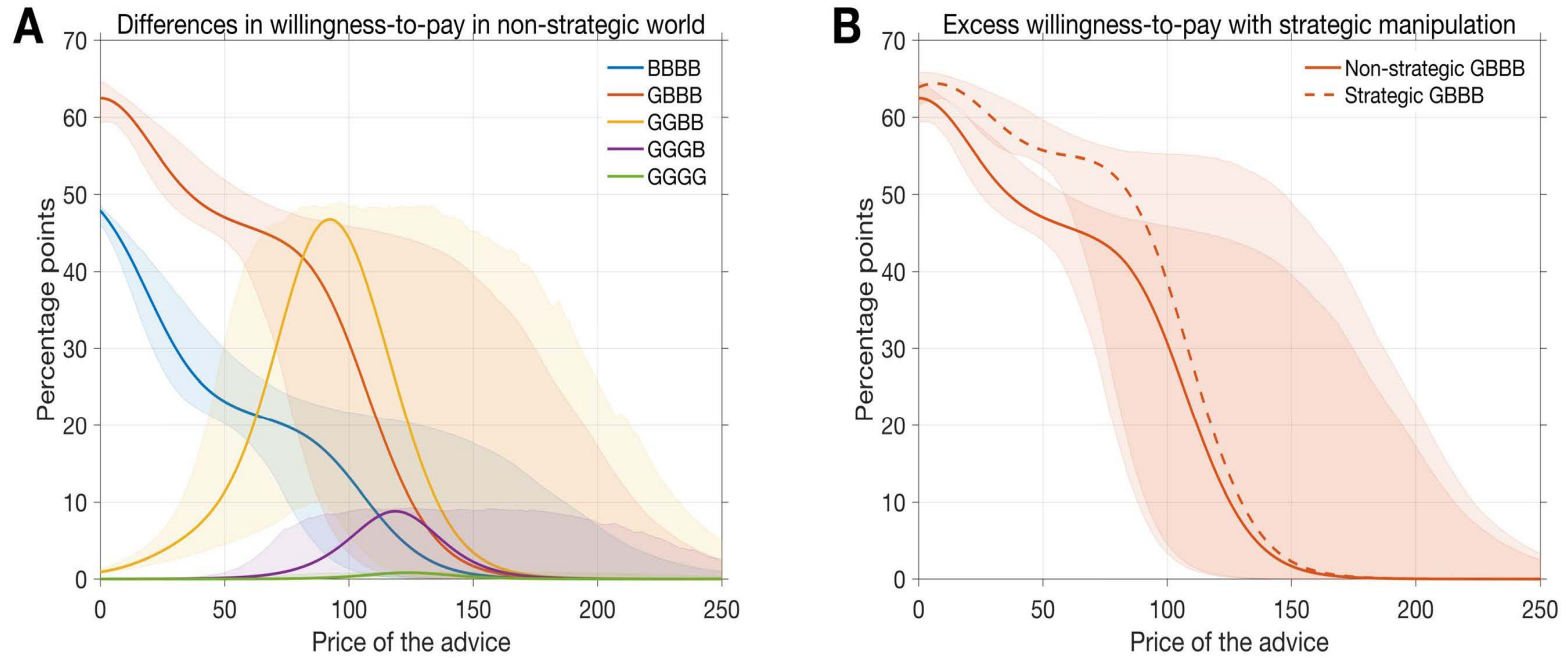
These figures plot the simulated probability distribution of each group’s willingness-to-pay various adviser fees for different sequences of good versus bad advice. The “Best-Case World” is where all the advice given is good (Panel A- Quality Sequence GGGG) and the “Worst-Case World” is where all the advice given is bad (Panel E- Quality Sequence BBBB).

Figure 3 allows us to determine whether the differences between the vulnerable and resilient participants are significant. In Panel A of this figure, the differences between the two lines for each quality sequence shown in the panels of Figure 2 are plotted with the 95% CIs displayed in the shaded areas. If these shaded areas do not overlap the horizontal axis, the differences between vulnerable and resilient participants are 95% or more likely to be different from zero. Not surprisingly, in the best-case world, when all advice is good (GGGG, green line), the difference is close to zero. When the last piece of advice is bad (sequence GGGB, purple line), differences begin to emerge. Once the adviser gives an equal amount of good and bad advice (GGBB, yellow line), the difference becomes more apparent and lies above zero with a high probability.

Most striking are the cases in which the adviser sets a good first impression and then delivers bad advice (GBBB, red line) and the worst-case world (BBBB, blue line). For these instances, vulnerable consumers are much more willing to pay advisers for their services than resilient consumers. This finding demonstrates that in a realistic world where advisers do not always deliver good advice, vulnerable clients are willing to pay a higher price for poorer quality advisers, and first impressions and confirmation bias play even greater roles. Again, this is the case when there is no strategic manipulation of topics. This suggests that providing vulnerable consumers with ways to differentiate between the quality of advisers could help decrease this divergence in willingness-to-pay.

We now consider the case in which advisers strategically manipulate advice delivery. Agnew et al. (2018) demonstrate how advisers can manipulate clients by strategically pairing the sequence of advice on clear or ambiguous topics with the (good or bad) quality of their advice. We investigate the cost of such strategic manipulation for consumers.

Figure 3: Statistical Difference between Vulnerable and Resilient Clients' Willingness-to-Pay



Panel A plots for each quality sequence shown in Figure 6 the differences between the vulnerable and resilient participants' probability lines. Advisers are not strategically manipulating advice. The shaded areas provide the CIs. If the shaded areas representing the CIs do not overlap the horizontal axis, the differences between the two groups are considered different from zero at the 95% level or more. Panel B compares the differences between when there is no strategic manipulation for the quality sequence GBBB (the line is the same as is shown in Panel A for that quality sequence) and when there is strategic manipulation of the topics for that same quality sequence. In the strategic case, the adviser provides advice topics in the following order to take advantage of a strong first impression: debt, consolidation, diversification and fees. The shaded areas provide the CI.

To best illustrate the possible exploitation of vulnerable clients, we focus only on the quality sequence that exploits the first impression (GBBB). In Figure 3, Panel B, we contrast these probabilities with the probabilities obtained if the topics are presented in the following order (i.e., a strategic scenario): debt, consolidation, diversification and fees. Debt is the clearest topic and thus allows for the strongest positive first impression to be set if good advice is delivered. We calculate the differences between the probabilities associated with the strategic scenario for vulnerable versus resilient consumers in the GBBB quality sequence. In addition, we show the differences from the nonstrategic case for the GBBB already illustrated in panel A as a point of comparison. Our figure demonstrates that if an adviser chose to use the participant’s observable characteristics to identify vulnerable participants and strategically offer advice to them, the probability of vulnerable clients paying more than resilient clients is greater.

CONCLUSION

Contributions

This study focuses on consumer learning about expert service quality in the context of credence goods, where ambiguous signals are profuse. We explain why some consumers are more likely to ignore bad signals about an adviser and identify those consumers most vulnerable to manipulation by malevolent actors. That is, those consumers who are more likely to follow and suffer from a learning process that is consistent with confirmation bias. Given that learning processes are latent and often unconscious, we employ an online, large-scale experiment to separate consumers into standard Bayesian learners and limited memory learners. In contrast to Bayesian learners, who do not use ambiguous new information to update their beliefs, limited memory learners interpret ambiguous information consistent with their priors and update their beliefs using these interpretations. This limited memory learning results in behavior heavily influenced by initial and early interactions, and is consistent with confirmation bias. In addition, this learning approach can explain polarized opinions despite individuals receiving the same information signals. We use a structural model to identify a consumer’s learning process and the characteristics associated with it.

Our work contributes to several literature streams. First, we add to the literature on trust in service encounters (Garbarino and Johnson, 1999; Sirdeshmukh et al., 2002) by explaining the

foundation and evolution of client trust in this context. Second, we contribute to literature in marketing, economics and finance that test alternative psychological models of belief updating (e.g., Chylinski et al., 2012; Camacho et al., 2011; Benjamin, 2019) by proposing a model that enables to detect confirmation bias based on individual choice data. Finally, we are the first to present a structural model that explains the mechanisms that make consumers vulnerable and thus move beyond the mostly "informal" treatment of consumer vulnerability in past consumer vulnerability research (Hill and Sharma, 2020). Our results thus provide new insight into consumer decision making and have direct public policy implications.

One of our most notable findings is that nearly two-thirds of our experiment's participants make choices that conform to a limited memory updating process. These participants update their prior beliefs based on their interpretation of ambiguous information. This result is important not only in the context of financial advice but also for other types of decision making where consumers are confronted with information that is open to interpretation (e.g., the formation of political opinions, trust in medical advice). Our experiment is a possible model for tests of limited memory updating in these contexts given its capacity to identify probable markers for limited memory updaters through observable characteristics and personality traits revealed through responses to survey questions. Whereas previous literature showed that different characteristics distinguish scam victims from non-victims (Langenderfer and Shimp, 2001), our model explains the mechanisms by which these characteristics increase susceptibility to scams. Knowledge of these mechanisms in turn provides opportunities to limit consumer vulnerability. For the first time, we estimate the economic cost limited memory updaters are willing to pay relative to standard rational Bayesian updaters. We find a significant divergence in willingness-to-pay adviser fees between two segments of participants that we call "vulnerable" and "resilient". Vulnerable participants are more likely to be limited memory updaters, not Bayesians, to have higher prior beliefs of adviser quality, and are less able to correctly identify the objective quality of signals. While resilient participants can pick out the lowest quality offering, in this case, an adviser who only gives objectively bad advice, and refuse to pay for it, vulnerable participants continue to pay. This failure by vulnerable clients to discern bad quality offerings holds even in the absence of malevolent actors, where, say, advisers do not use strategic manipulation or catering, as suggested by earlier studies. Further analysis reveals that in the

presence of strategic manipulation, vulnerable clients pay significantly higher fees.

In the context of financial advice, we identify vulnerable and resilient clients in our sample as differing by age, financial literacy, numeracy, prior trust of advisers and impulsiveness. We demonstrate that impulsiveness increases the likelihood of being a limited memory updater, that prior trust in financial advisers and the presentation of professional certifications by advisers lead to higher prior beliefs, and that financial literacy affects clients' ability to discern good from bad advice. Taken together, these characteristics significantly impact participants' willingness-to-pay for advisers. In sum, we find that vulnerable participants need help selecting a qualified, high-quality adviser even when advisers are not purposefully targeting them. These results raise the following question: how can we help vulnerable consumers make better choices?

Implications

Our model suggests three different means to help vulnerable consumers make better choices. First, policy makers need to ensure that prior beliefs match reality. Misplaced trust may be corrected by providing risk information (Menon et al., 2002) or requiring objective quality signals such as professional certifications that signal that the offer is trustworthy. If certifications are signals of superior quality, then certifications can provide helpful information to the consumer. However, as it stands, many different certifications of varying quality are available around the world and in different contexts including not only financial advice (e.g. Goetz et al., 2011), but ranging from food labeling (Kuchler et al., 2020) to educational attainment (Schneider, 2009). Consumers need guidance to identify reliable certifications and regulators could consider whether to standardize the certifications available.

Second, policy makers should aim to reduce limited memory updating. Our study showed that impulsiveness is a major contributor to this kind of learning process and we therefore suggest that policy could focus on curbing consumer impulsivity. Public policy makers could, for example, stipulate cool-down periods between getting and acting on advice. Past research (Thomas et al., 2011) has also documented that visceral stimuli, such as the pain of paying with cash, can be useful interventions to regulate impulsiveness.

Third, policy makers need to help consumers correctly identify the quality of signals. Examples of such help include improving financial literacy to give clients the skills to evaluate financial

advice, news literacy to identify fake news, or food literacy to make the right food choices. Given the low levels of literacy in each of these domains, there is ample potential for policy makers to make a difference.

Our final suggestion on how to help consumers does not, however, directly derive from our model. We propose that regulators or professional associations enforce industry standards rather than only relying on voluntary participation. Egan et al. (2020), for example, demonstrate that even proposing a fiduciary standard can have positive effects for consumers. Properly enforcing standards is likely to have an even stronger effect.

Future Research

Our work provides important contributions to extant research and its limitations can serve as opportunities for future work. To begin, we study learning in the context of financial advice and extrapolate managerial implications from the insights discovered in our model. Yet, expert services abound, ranging from medical doctors to appliance service persons. Since trust formation has been shown to be context-specific (Song et al., 2007), it is likely that, in these other contexts, future research may identify additional covariates that impact prior beliefs, the ability to distinguish good from bad signals, or the learning style consumers employ. Next, our study focuses on beliefs about human actors (here: financial advisers). Given the surge in automated advice in different areas (e.g., robo financial or health advice), future research could focus on consumer learning about robo-advice quality. Lourenço et al. (2020) point out that consumers' perceptions of the firm providing the automated advice are important drivers of advice acceptance. Our model could be used to investigate how consumers update their behavior after having received the first round of advice from the robot.

Furthermore, our experimental setup mimics the first encounter with an adviser and consumers' learning about the adviser during this first encounter. In many contexts, consumers will, however, have longer lasting relationships, allowing them to observe advice itself, and outcomes from acting upon this advice, over a longer time. Clients may, for example, invest in a particular stock recommended by their financial adviser and observe stock performance; patients may follow a particular dietary advice from their physician and observe changes in their health; and car owners may follow their mechanic's suggestion to replace a brake early and observe

the outcomes of that decision. While these outcomes are often very noisy and not necessarily the result of advice-taking, consumers tend to interpret them as evidence for or against the quality of the adviser. Future research could therefore incorporate these outcomes and their attribution into our model of belief updating.

Finally, our experimental setup and model do not allow for impact of events that are not direct interactions between adviser and consumer once prior beliefs have been formed. Gurun et al. (2018) however show, in the context of the infamous Madoff fraud, that trust shocks are transmitted through social networks, such that clients lose confidence in advisers in general after the fraud is exposed. Future research might investigate what kind of events have the potency to influence consumers' learning.

REFERENCES

- Agnew, J. R., Bateman, H., Eckert, C., Iskhakov, F., Louviere, J., and Thorp, S. (2018). First impressions matter: An experimental investigation of online financial advice. *Management Science*, 1:288–307.
- Agnew, J. R., Bateman, H., and Thorp, S. (2013). Financial literacy and retirement planning in Australia. *Numeracy*, 6(2):7.
- Augenblick, N. and Rabin, M. (2021). Belief movement, uncertainty reduction, and rational updating. *Quarterly Journal of Economics*, 136(2):933–985.
- Australian Securities and Investments Commission (2012). *Report Guide 279—Shadow shopping study of retirement advice*. Canberra, Australia: ASIC.
- Baeckström, Y., Marsh, I. W., and Silvester, J. (2021). Variations in investment advice provision: A study of financial advisors of millionaire investors. *Journal of Economic Behavior & Organization*, 188:716–735.
- Bateman, H. and Kingston, G. (2012). The future of financial advice and MySuper. *JASSA*, 3:49–53.
- Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:69–186.
- Benoît, J.-P. and Dubra, J. (2018). *When do populations polarize? An explanation*. Documentos de Trabajo/Working Papers 1801, Facultad de Ciencias Empresariales y Economía. Universidad de Montevideo.
- Berens, M. (2023). How doctors buy their way out of trouble. Reuters. <https://www.reuters.com/investigates/special-report/usa-healthcare-settlements/>.
- Bergstresser, D., Chalmers, J. M., and Tufano, P. (2009). Assessing the costs and benefits of brokers in the mutual fund industry. *Review of Financial Studies*, 22(10):4129–4156.

- Bidjerano, T. and Dai, D. Y. (2007). The relationship between the Big Five model of personality and self-regulated learning strategies. *Learning and Individual Differences*, 17(1):69–81.
- Bone, S. A., Christensen, G. L., and Williams, J. D. (2014). Rejected, shackled, and alone: The impact of systemic restricted choice on minority consumers’ construction of self. *Journal of Consumer Research*, 41(2):451–474.
- Borghans, L., Duckworth, A. L., Heckman, J. J., and Weel, B. t. (2008). The economics and psychology of personality traits. *Journal of Human Resources*, 43(4):972–1059.
- Brown, W. (1995). Wrenching frustrations. The Washington Post. <https://www.washingtonpost.com/archive/business/1995/12/02/wrenching-frustrations/8f188854-6d34-4f39-9d55-13628b43c80b/>.
- Bucher-Koenen, T., Hackethal, A., Koenen, J., and Laudenbach, C. (2021). Gender differences in financial advice. SSRN Scholarly Paper ID 2572961, Social Science Research Network, Rochester, NY.
- Camacho, N., Donkers, B., and Stremersch, S. (2011). Predictably Non-Bayesian: Quantifying Salience Effects in Physician Learning About Drug Quality. *Marketing Science*, 30(2):305–320.
- Chamorro-Premuzic, T. and Furnham, A. (2003). Personality predicts academic performance: Evidence from two longitudinal university samples. *Journal of Research in Personality*, 37(4):319–338.
- Charness, G. and Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104:1–23.
- Chylinski, M. B., Roberts, J. H., and Hardie, B. G. (2012). Consumer learning of new binary attribute importance accounting for priors, bias, and order effects. *Marketing Science*, 31(4):549–566.
- Darby, M. R. and Karni, E. (1973). Free competition and the optimal amount of fraud. *The Journal of Law and Economics*, 16(1):67–88.
- Darley, J. M. and Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology*, 44(1):20.
- Durham, G. and Geweke, J. (2014). Adaptive sequential posterior simulators for massively parallel computing environments. In *Bayesian Model Comparison*, pages 1–44. Emerald Group Publishing Limited.
- Egan, M., Ge, S., and Tang, J. (2020). Conflicting Interests and the Effect of Fiduciary Duty — Evidence from Variable Annuities. SSRN Scholarly Paper ID 3652115, Social Science Research Network, Rochester, NY.
- Egan, M., Matvos, G., and Seru, A. (2019). The market for financial adviser misconduct. *Journal of Political Economy*, 127(1):233–295.
- Emons, W. (1997). Credence goods and fraudulent experts. *The RAND Journal of Economics*, pages 107–119.
- Erdem, T. and Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20.

- Evans, K. R., Kleine, R. E., Landry, T. D., and Crosby, L. A. (2000). How first impressions of a customer impact effectiveness in an initial sales encounter. *Journal of the Academy of Marketing science*, 28(4):512–526.
- Franken, I. H. A., van Strien, J. W., Nijs, I., and Muris, P. (2008). Impulsivity is associated with behavioral decision-making deficits. *Psychiatry Research*, 158(2):155–163.
- Fryer, Jr, R. G., Harms, P., and Jackson, M. O. (2019). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. *Journal of the European Economic Association*, 17(5):1470–1501.
- Garbarino, E. and Johnson, M. S. (1999). The different roles of satisfaction, trust, and commitment in customer relationships. *Journal of Marketing*, 63(2):70–87.
- Gennaioli, N., Shleifer, A., and Vishny, R. (2015). Money doctors. *Journal of Finance*, 70(1):91–114.
- Georgarakos, D. and Inderst, R. (2014). Financial advice and stock market participation. SSRN Scholarly Paper ID 1641302, Social Science Research Network, Rochester, NY.
- Germann, M., Loos, B., and Weber, M. (2018). Trust and delegated investing: A Money doctors experiment. SSRN Scholarly Paper ID 3198134, Social Science Research Network, Rochester, NY.
- Glimcher, P. W., Kable, J., and Louie, K. (2007). Neuroeconomic studies of impulsivity: Now or just as soon as possible? *American Economic Review*, 97(2):142–147.
- Goetz, J. W., Zhu, D., Hampton, V. L., Chatterjee, S., Salter, J., et al. (2011). Integration of professional certification examinations with the financial planning curriculum: Increasing efficiency, motivation, and professional success. *American Journal of Business Education (AJBE)*, 4(3).
- Guiso, L., Sapienza, P., and Zingales, L. (2008). Trusting the stock market. *Journal of Finance*, 63(6):2557–2600.
- Gurun, U. G., Stoffman, N., and Yonker, S. E. (2018). Trust busting: The effect of fraud on investor behavior. *Review of Financial Studies*, 31(4):1341–1376.
- Hackethal, A. and Inderst, R. (2013). *How to make the market for financial advice work*. Oxford: Oxford University Press.
- Hill, R. P. and Sharma, E. (2020). Consumer vulnerability. *Journal of Consumer Psychology*, 30(3):551–570.
- Hoch, S. J. and Deighton, J. (1989). Managing what consumers learn from experience. *Journal of Marketing*, 53:1–20.
- Hoechle, D., Ruenzi, S., Schaub, N., and Schmid, M. (2017). The impact of financial advice on trade performance and behavioral biases. *Review of Finance*, 21(2):871–910.
- Honigsberg, C., Hu, E., and Jackson Jr, R. J. (2022). Regulatory arbitrage and the persistence of financial misconduct. *Stanford Law Review*, 74:737.
- Hou, R., Mogg, K., Bradley, B. P., Moss-Morris, R., Peveler, R., and Roefs, A. (2011). External eating, impulsivity and attentional bias to food cues. *Appetite*, 56(2):424–427.
- Jasra, A., Stephens, D. A., and Holmes, C. C. (2007). On population-based simulation for

static inference. *Statistics and Computing*, 17(3):263–279.

- Konya-Baumbach, E., Schuhmacher, M. C., Kuester, S., and Kuharev, V. (2019). Making a first impression as a start-up: Strategies to overcome low initial trust perceptions in digital innovation adoption. *International Journal of Research in Marketing*, 36(3):385–399.
- Kuchler, F., Bowman, M., Sweitzer, M., and Greene, C. (2020). Evidence from retail food markets that consumers are confused by natural and organic food labels. *Journal of Consumer Policy*, 43:379–395.
- Langenderfer, J. and Shimp, T. A. (2001). Consumer vulnerability to scams, swindles, and fraud: A new theory of visceral influences on persuasion. *Psychology & Marketing*, 18(7):763–783.
- Linnainmaa, J. T., Melzer, B. T., and Previtero, A. (2021). The misguided beliefs of financial advisors. *Journal of Finance*, 76(2):587–621.
- Lipkus, I. M., Samsa, G., and Rimer, B. K. (2001). General performance on a numeracy scale among highly educated samples. *Medical Decision Making*, 21(1):37–44.
- Lourenço, C. J., Dellaert, B. G., and Donkers, B. (2020). Whose algorithm says so: The relationships between type of firm, perceptions of trust and expertise, and the acceptance of financial robo-advice. *Journal of Interactive Marketing*, 49(1):107–124.
- Lusardi, A. and Mitchell, O. S. (2011). Financial literacy around the world: An overview. *Journal of Pension Economics and Finance*, 10(04):497–508.
- Menon, G., Block, L. G., and Ramanathan, S. (2002). We’re at as much risk as we are led to believe: Effects of message cues on judgments of health risk. *Journal of Consumer Research*, 28(4):533–549.
- Mullainathan, S., Noeth, M., and Schoar, A. (2012). The market for financial advice: An audit study. Report, National Bureau of Economic Research.
- Muthukrishnan, A. V. and Chattopadhyay, A. (2007). Just give me another chance: The strategies for brand recovery from a bad first impression. *Journal of Marketing Research*, 44 (May):334–345.
- Nagel, S. and Xu, Z. (2021). Asset pricing with fading memory. *Review of Financial Studies*, 35(5):2190–2245.
- Oppenheimer, D. M., Meyvis, T., and Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4):867–872.
- Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82.
- Richardson, M. and Abraham, C. (2009). Conscientiousness and achievement motivation predict performance. *European Journal of Personality*, 23(7):589–605.
- Roberts, B. W., Lejuez, C., Krueger, R. F., Richards, J. M., and Hill, P. L. (2014). What is conscientiousness and how can it be assessed? *Developmental Psychology*, 50(5):1315–1330.
- Roberts, J. H. and Urban, G. L. (1988). Modeling multiattribute utility, risk, and belief dynamics for new consumer durable brand choice. *Management Science*, 34(2):167–185.

- Schneider, S. L. (2009). *Confusing credentials: The cross-nationally comparable measurement of educational attainment*. PhD thesis, Oxford University, UK.
- Sirdeshmukh, D., Singh, J., and Sabol, B. (2002). Consumer trust, value, and loyalty in relational exchanges. *Journal of Marketing*, 66(1):15–37.
- Soll, J. B., Palley, A. B., and Rader, C. A. (2022). The bad thing about good advice: Understanding when and how advice exacerbates overconfidence. *Management Science*, 68(4):2949–2969.
- Song, J. et al. (2007). Trust in health infomediaries. *Decision Support Systems*, 43(2):390–407.
- Stolper, O. and Walter, A. (2019). Birds of a feather: The impact of homophily on the propensity to follow financial advice. *Review of Financial Studies*, 32(2):524–563.
- Thomas, M., Desai, K. K., and Seenivasan, S. (2011). How credit card payments increase unhealthy food purchases: Visceral regulation of vices. *Journal of Consumer Research*, 38(1):126–139.
- Tsukayama, E., Duckworth, A. L., and Kim, B. (2012). Resisting everything except temptation: Evidence and an explanation for domain-specific impulsivity. *European Journal of Personality*, 26(3):318–334.
- Verhoef, P. C., Lemon, K. N., Parasuraman, A., Roggeveen, A., Tsiros, M., and Schlesinger, L. A. (2009). Customer experience creation: Determinants, dynamics and management strategies. *Journal of Retailing*, 85(1):31–41.
- von Stumm, S. and Furnham, A. F. (2012). Learning approaches: Associations with typical intellectual engagement, intelligence and the Big Five. *Personality and Individual Differences*, 53(5):720–723.
- Weafer, J., Baggott, M. J., and de Wit, H. (2013). Test–retest reliability of behavioral measures of impulsive choice, impulsive action, and inattention. *Experimental and Clinical Psychopharmacology*, 21(6):475–481.
- Yang, Y., Zhong, X., Wu, D., Li, H., and Li, M. (2016). Positive association between trait impulsivity and high gambling-related cognitive biases among college students. *Psychiatry Research*, 243:71–74.
- Yaniv, I. and Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83(2):260–281.
- Zaleskiewicz, T. and Gasiorowska, A. (2018). Tell me what I wanted to hear: Confirmation effect in lay evaluations of financial expert authority. *Applied Psychology*, 67(4):686–722.
- Zweig, J. (2023). You don't know what you don't know about your financial adviser. The Wall Street Journal. <https://www.wsj.com/articles/financial-advisor-disclosure-regulations-adv-11667576312>.