

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

Julie Agnew,^a Hazel Bateman,^b Christine Eckert,^c Fedor Iskhakov^d Susan Thorp^e

Keywords: vulnerability, experiment, credence, services, biases

^a Raymond A. Mason School of Business, William and Mary, Williamsburg, VA, 23187 USA, E-mail: julie.agnew@mason.wm.edu

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

^c (Corresponding Author) Technical University of Munich, TUM School of Management, Heilbronn, Germany, E-mail: christine.eckert@tum.de & UTS Business School, University of Technology Sydney, Sydney, NSW, Australia

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

^e The University of Sydney Business School, The University of Sydney, NSW, Australia, E-mail: susan.thorp@sydney.edu.au

Acknowledgements:

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

Funding Statement:

This research project was generously funded by Australian Research Council Grants [Grant DP1093842 and DP190101056]

Data Availability:

The data that support the findings of this article are not publicly available but are available from the corresponding author on reasonable request.

Policy Contribution Statement

This paper contributes to policy debates on how to protect consumers in markets for credence goods such as medical services and financial products. In such settings, information asymmetries expose consumers to the risk of relying on low-quality advisers and persisting with poor advice. This work addresses the ongoing policy conversation (e.g., *Treasury Australia Quality of Advice Review 2023*), that seeks to ensure advice markets function fairly and efficiently for consumers, particularly the vulnerable.

Our study advances understanding beyond existing marketing and public policy literature by integrating research streams on consumer learning, first impressions, trust formation, advice utilization, and consumer vulnerability into a unified framework. Using experimental survey data and a purpose-built model, we identify mechanisms that can lead consumers to use and persist with bad advisers. We detect behavioral biases (confirmation bias) and personal characteristics (impulsivity and low financial skills) which impede judgements of advice quality and are the first to estimate the costs vulnerable clients incur by being willing to pay too much for bad advice.

The results have direct implications for policymakers, regulators, and industry stakeholders. Regulators can use our findings to design interventions that reduce information asymmetry, such as adviser certifications or mandated checklists of advice quality. Policy interventions like mandatory cooling-off periods can mitigate impulsivity in advice uptake. Both policymakers and marketers can also invest in consumer education initiatives that enhance financial skills and improve the capacity to evaluate advice. Our model further enables stakeholders to prioritize interventions based on their relative effectiveness, thus providing an evidence base for policies that aim to protect vulnerable consumers and improve market outcomes in credence goods industries.

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

Abstract

Consumers of credence goods often rely on interpretations of seller signals to judge seller quality. We examine how consumers' interpretations of ambiguous signals explain why they fall for, and persist with, bad financial advisers. Using data from an experimental survey and a purpose-built model, we capture mechanisms that make consumers vulnerable. Survey participants ("clients") make subjective interpretations of financial advice on topics of varying difficulty, simultaneously learning about the quality of the advisers who give the advice. Most participants use ambiguous signals from advice on difficult topics to confirm prior beliefs about the advisers. Impulsive participants are prone to this confirmation bias. Impulsivity, when combined with low financial skills, impedes judgments of the quality of advice, and causes participants to be willing to pay too much to poor advisers – particularly to malevolent advisers who present advice strategically. Our model can identify strategies to help vulnerable consumers of financial advice and facilitate their prioritization based on relative effectiveness.

Keywords: vulnerability, experiment, credence, services, biases

Introduction

Getting advice can be hazardous. Advisers sometimes recommend unsuitable, or expensive products (e.g., Mullainathan et al. 2012). "Expert" advisers may have biases that they transfer to clients (Linnainmaa et al. 2021). Conflicting incentives may encourage advisers to sell second-rate products (Egan, Ge, and Tang 2020), while others may take advantage of less-savvy clients (Bucher-Koenen et al. 2021) or give advice based on social stereotypes (Baeckström, Marsh, and Silvester 2021). Worse still, some advisers repeatedly defraud or exploit clients, and then hide their misconduct by moving between firms or regulatory jurisdictions (Egan, Matvos, and Seru 2019).

At the same time, clients themselves may misjudge whose advice or what advice to take. Clients may reach opinions based on scant evidence (Yaniv and Kleinberger 2000). They might base their judgments on personal similarities (Stolper and Walter 2019), on agreement with their opinions (Soll, Palley, and Rader 2022), on adviser anxiety (Gerrans and Hershey 2017) 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 or pay high fees (Gennaioli, Shleifer, and Vishny 2015). These findings raise obvious questions. Why do clients persist with advisers who give poor advice? What is the cost of doing so? And how can public policymakers and industry prioritize initiatives that improve consumer welfare?

The difficulty of choosing a good adviser is typical for purchases of credence goods. Consumers of credence goods cannot be certain of the quality of the good either before or after purchase (Keh and Sun 2018). Consequently, sellers often provide the credence good, while also

determining consumers' needs and giving expert advice on which goods to purchase. Such expert services are provided by a wide range of sellers, including financial advisers, medical doctors, lawyers, car mechanics and appliance technicians. Unfortunately, such expert services are also prone to seller fraud: the information asymmetry between consumer and seller creates strong incentives for opportunistic seller behavior (Emons 1997). The Bernie Madoff case is a high-profile example of expert service fraud. Madoff, the now notorious but once highly regarded financial adviser, caused \$17 billion in losses to over 1,000 individuals and firms in a highly publicized Ponzi scheme.¹ However, other industries have 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 \$27 billion was paid by practitioners and providers paid over the past decade to settle allegations of fraud, bribery and patient harm (Berens 2023).

Our paper is the first to provide a holistic picture of the mechanisms that can lead consumers to fall for, and persist with, bad sellers. We unify previously separate streams in service research on first impressions and consumer learning about service provider quality (Dagger and Sweeney 2007; Lantzy et al. 2020), formation of consumer trust (Sirdeshmukh, Singh, and Sabol 2002), advice provision and utilization (Wittkowski et al. 2020; Johnson et al. 2021), and consumer vulnerability (Cenopht et al. 2004). Moving beyond a simplistic treatment of consumer vulnerability (see discussion in Hill and Sharma 2020), we investigate if and how combinations of characteristics associated with vulnerability increase harm. We model the mechanisms connecting vulnerability and harm, describing how consumers learn about sellers'

¹ A 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.

quality, identifying personality traits that affect learning, quantifying the extra fees that bad sellers can impose on more vulnerable consumers. This enables us to suggest strategies to protect these consumers as well as rank them by their expected effectiveness. The remainder of this introduction discusses our empirical approach, previews our results and outlines our contributions to research and practice. From here on, we will call the sellers who provide expert services "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 client interactions (Muthukrishnan and Chattopadhyay 2007). We construct scenarios of initial meetings between clients and advisers, and model how these meetings impact subsequent advice taking. We use financial advice as the context because choosing a good financial adviser is difficult (Hansen 2012). We collect data through a purpose-built and pre-tested online choice experiment and survey, which was incentivized and completed by 2003 participants. Model-free analysis 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 impulsivity, this "stickiness" is more common. We then propose a model of biased belief formation that can explain why some clients persist with poor quality advisers, and estimate the cost of doing so. Our model captures an important pattern that has been previously documented in the context of consumers' learning about service provider quality, namely that customers "tend to evaluate service performance in line with their prior expectations" (Habel et al. 2016).

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 simple as a clear signal. Clients' ratings of ambiguity and clarity are subjective and vary with clients' understanding of financial topics.² Learning about the quality of the adviser then relies on clients recalling these signals and updating their beliefs. While all clients have some recall, it is implausible to assume that all clients 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, Harms, and Jackson 2019). Bayesians remember whether past signals are ambiguous or clear, they set aside ambiguous signals and use only clear signals to update their beliefs.³ 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 clients 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 (Rabin and Schrag 1999). Confirmation bias may lead clients to choose bad advice when good advice is available, and to pay for a bad adviser. The estimated

² Lusardi 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, while clear to others.

³ An alternative explanation for ignoring ambiguous signals is presented by Kominers et al. (2018), who argue that Bayesians may not update if the cost of processing the signal outweighs the perceived benefits of understanding the signal.

model identifies the types of clients whose learning leads to biased beliefs. Simulations of the model then show the extra fees that such vulnerable⁴ and biased clients will pay to bad advisers.

Our findings enhance the literature on the impact of first impressions and consumer learning about service provider quality, 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 (Sirdeshmukh, Singh, and Sabol 2002). Studies also document the effect of first impressions on trust formation (Konya-Baumbach et al. 2019) and the additional price clients are willing to pay a trusted service provider (Gennaioli, Shleifer, and Vishny 2015). We specifically contribute to research in the context of advice provision and utilization in expert services (Wittkowski et al. 2020). Previous research has documented that not adhering to expert advice from the service provider can harm customers (Chan, Yim, and Lam 2010). Adhering to bad expert advice, however, may be similarly detrimental (Egan, Ge, and Tang 2020). Several factors, including trust in the service provider, determine whether customers utilize advice (Seiders et al. 2015). We draw on insights from consumer learning to explain the foundation and evolution of client trust in this context.

We also add to studies in marketing that discuss different psychological models of belief updating (e.g., Erdem and Keane 1996), and extend the existing body of literature on information processing as it pertains to credence goods (Keh and Sun 2018). Specifically, we base our model on Fryer, Harms, and Jackson (2019). This model allows both Bayesian and limited memory clients to receive ambiguous signals, with confirmation bias arising from the updating rule used

⁴ Many definitions of consumer vulnerability exist (see Hill and Sharma 2020) and we are not proposing another one here. We use ‘vulnerability’ as a description of our segmentation of survey participants.

by limited memory clients. Our model explains how consumers process information, including ambiguous signals of quality, when evaluating service providers, and adds new insights into how customers construct opinions about service quality (Dagger and Sweeney 2007). We thereby integrate previous findings (Habel et al. 2016) that show that confirmation bias is particularly pronounced if customers lack the motivation or ability to evaluate the service performance – a typical feature of credence goods.

We contribute to advancing knowledge about consumer vulnerability (Hill and Sharma 2020), the identification of disadvantaged subpopulations (e.g. Bone, Christensen, and Williams 2014) and the marketer manipulations that may affect vulnerable consumers' decision making. Langenderfer and Shimp (2001), for example, report that victims of business scams are likely to be significantly more trusting, less educated and less self-controlled than consumers who did not fall for scams. Past research in financial services (e.g., Stolper and Walter 2019) also highlights key consumer vulnerabilities. We provide insights into the mechanisms that make consumers vulnerable and identify ways to help them.

Finally, in times of tight public policy budgets, our model enables us to evaluate the relative effectiveness of different initiatives to make consumers more resilient to bad advice. We thus contribute to research aiding policymakers evaluating the effectiveness of interventions (e.g., Edelen, Fong, and Han 2025).

Conceptual Framework

Learning In the Presence of Ambiguous Signals

Learning about an adviser's quality 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 to follow advice in an area where they have little experience, they usually rely on signals about the quality of an adviser, such as 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 receive 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. 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). It follows that the future choices of Bayesian learners are guided by beliefs based only on unambiguous signals. Limited memory learners, however, make a "double update" (Fryer, Harms, and Jackson 2019). First, they interpret ambiguous new signals 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.⁵ Irrespective of the actual signal valence, for the purposes of updating posterior beliefs about adviser quality, people who use limited memory updating 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

⁵ We note that confirmation bias may also result from biased selection of information sources (see Chaxel, Russo, and Kerimi 2013).

hypothesis despite receiving an infinite amount of information (Rabin and Schrag 1999). Figure 1 illustrates how prior beliefs, clients' learning type, and signal ambiguity lead to differences in interpretations of the signal as well as to differences in posterior beliefs about adviser quality.




Figure 1: Conceptual Framework. This figure illustrates how clients interpret signals and form posterior beliefs about adviser quality based on their prior beliefs, their learning type, and whether the signal is clear or ambiguous to them.

Linking Prior Beliefs, Learning Type, and Signal Updating to Client Characteristics

We cannot directly observe whether a client uses Bayesian or limited memory learning, which prior beliefs a client holds, or whether a signal is clear or ambiguous to a client. Instead, these are latent traits that we can only infer by observing the client's choices in a suitable setting. Yet, our conceptual framework shows that these latent traits can compound, and lead clients to posterior beliefs about the adviser which do not match reality.

Previous studies on the foundations of trust in marketing and management offer valuable insights into the determinants shaping a client's *prior belief about an adviser*. One key antecedent is the adviser's ability, that is, the adviser's expertise and experience in delivering high-quality service (Schumann et al. 2010). Given the challenge of evaluating this ability at the beginning of a relationship, clients often turn to third-party information to help their judgment (Keh and Sun 2018). Agnew et al. (2018) show that information from a third party in the form of professional credentials can improve clients' first impressions of an adviser. While credentials can explain some differences in clients' prior beliefs, differences in clients' perceptions of advisers in general contribute variation: Hansen (2012) shows that clients' general trust in advisers can influence trust in a specific adviser. So, we expect that both credentials and general trust in advisers influence clients' prior beliefs about an adviser.

Clients' information processing style, the likelihood that they belong to a particular *learning type*, depends on their cognitive style (Oreag and Bavazit 2009). We allow clients' conscientiousness and impulsivity to predict their learning type. Conscientiousness is a Big Five personality trait related to orderliness, industriousness and effort (e.g., Roberts et al. 2014). 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 (e.g., Glimcher, Kable, and Louie 2007). Impulsive people tend to rely on information that is close to hand to form beliefs, without first thinking about whether it is correct. Because

people can more easily remember their interpretations of signals than the actual past signals, we hypothesize that impulsive people are more likely to be limited memory learners.⁶

Clients' perception of advice as *clear or ambiguous* depends on their ability to evaluate the advice (Habel et al. 2016). As this ability is context specific, we use insights from the finance domain, and expect financial literacy, numeracy, product knowledge, and market experience, as well as gender and client age (as a proxy for their experience) to impact signal ambiguity (Agnew et al. 2018).

Combined, client and adviser characteristics may both impact how clients learn about an adviser's quality and may explain why some clients fall for, and persist with, a bad adviser. We propose an experiment and an associated model to detect clients' prior beliefs about an adviser, their learning style, and which signals they find ambiguous, and link these traits back to the characteristics of the client, thus enabling us to infer who is vulnerable and why.

Method

Survey and Choice Experiment

To identify clients' beliefs as well as their learning style, 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 two financial advisers giving financial advice on four topics. For each topic, one adviser provided objectively correct advice and the other objectively incorrect advice. After

⁶ Impulsivity is a facet of the Big Five personality factors, but its factors remain a subject of debate (Borghans et al. 2008). Following previous research, we use a finance domain-specific measure of impulsivity (Tsukayama et al. 2012).

watching the two advisers' videos on a topic, participants had to choose the adviser whose advice they would follow, before then moving to the next topic. As such, participants received a sequence of signals from the two advisers that they could use to update their beliefs about the advisers. After having watched the eight videos on the four different topics and having chosen the adviser whose advice they would follow in each topic, we asked participants about their willingness to pay for each of the two advisers. Participants' choices of a particular adviser in each topic, combined with their willingness to pay for the advisers at the end of the choice task, allow us to infer prior as well as posterior beliefs about the adviser, participants' learning type, as well as which signals are ambiguous versus clear to them (see Figure 2 for experimental setup).




Figure 2: Experimental Setup. This figure illustrates the experimental input, i.e., the signal (advice) an adviser provides on each topic, the learning process which is latent and varies between participants based on their learning type and the perceived ambiguity of the signal, and the observed output, namely the choice of the adviser's advice on a particular topic and the willingness to pay for the adviser.

Survey structure

The task of choosing advisers was embedded in a larger survey. This survey had four sections.⁷ Part one measured participants' financial literacy using standard questions on interest rates, inflation and diversification (Lusardi and Mitchell 2011), numeracy skills (Lipkus, Samsa, and Rimer 2001), understanding of the four advice topics covered later in the choice experiment, familiarity with a range of financial 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, and also asked about 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 as per Agnew et al. (2018). This section included two instructional manipulation (attention) checks (IMC1 and IMC2) (Oppenheimer, Meyvis, and Davidenko 2009). The last part of the survey 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 the four topics.⁸ 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

⁷ Parts one, three and four replicate the instrument by Agnew et al. (2018). Part two includes a choice task designed specifically to support estimation of our model. Web Appendix A shows screenshots of the survey.

⁸ By using video components from Agnew et al. (2018), we guaranteed that advice topics, advisers, environment and mode of advice delivery were uniform. Actors and fictional names were pretested to ensure that participants viewed each of them as equally credible.

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 (young vs. old), gender (female vs. male), and whether or not they had a professional credential (see Web Appendix B for the experiment design). 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.⁹ 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. Table 1 reports the scripts for each topic. After viewing the videos for the advice topic as many times as they liked, participants chose the advice recommendation they would follow. We also 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 were 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 one fixed fee value p to each participant.

⁹ To view an example of the video advice from a treatment in Agnew et al. (2018), please follow this link: https://researchbox.org/3425&PEER_REVIEW_passcode=RHAZSA.

TABLE 1: *Advice Scripts*

Narrator Introduction	Advice	Narrator Introduction	Narrator Introduction
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.	<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. 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.	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.	<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. 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.
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	<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. Bad Advice: ...Despite that, I recommend that you not roll all of these accounts together so you are diversified across different superannuation funds.	Diversifying a Stock Portfolio: In this scenario, you are thinking about investing in the share market. The next two financial advisers will recommend what you should do about it.	<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. 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.

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. *In Australia, retirement accounts are called ‘superannuation accounts’ and retirement plans are called ‘super funds’.

To be suitable for our experiment, the advice topics had to satisfy two conditions. First, they needed to have objectively correct answers, so that one adviser could deliver objectively good advice and the other objectively bad advice in the choice task. Second, the topics had to be sufficiently difficult that at least some participants would treat them as subjectively ambiguous signals of the quality (expertise) of the video advisers. We chose four common consumer finance topics: credit card debt repayment, retirement savings account consolidation, diversification in equity investments, and index fund fees.¹⁰ 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. 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.

¹⁰ Agnew et al. (2018) give a detailed justification for the choice of these topics.

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.

Overview of survey responses

We invited members of PureProfile, an Australian nationally representative online panel, to participate. Respondents had to pass two questions to meet our age and gender quotas, resulting in 2,003 participants who completed the survey. Web Appendix C reports 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. Most participants completed the survey in under 30 minutes, and the data collection took less than three weeks. Table 2 defines the variables used in analysis and estimation.

TABLE 2: *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.

TABLE 2 *Continued: Variable Descriptions*

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, and 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). To 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.

At the aggregate level (see Table 3), 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 the 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.

¹¹ These results agree with Agnew et al. (2018, Table 3).

TABLE 3: *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 impulsivity	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. IMC1 and IMC2 are instructional manipulation checks. All other variables are defined in Table 2

Preliminary evidence of confirmation bias

Preliminary analysis of the choice patterns indicates that – as expected – not only the quality of advice, but also participants’ beliefs about the advisers guide their subsequent choices. We focus on participants’ persistence, that is, their tendency to persist 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, whom 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 persist with the adviser they chose at choice 1 because the participant rates that adviser to be higher quality.¹² 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 4, 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 21.9% 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 (7.8% and 13.0%). Correspondingly, the higher percentage of persistent choices when the financial topic at choice 2 is high ambiguity (28.8% and 38.5%) 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 treated it as correct and updated posterior beliefs accordingly. After updating with their interpretation of an ambiguous signal, the participant makes a more favorable rating of the adviser and chooses that adviser again at the next, ambiguous, but wrong choice. Table 4, Panel B, provides further evidence of limited memory learning among groups of participants. Here we simply classify participants into

¹² Participants made active choices, ruling out state dependence. Incentives for correct choice of advice reduce the chance that choices reflect preference for adviser rather than beliefs about adviser quality.

"capable" and "non-capable" of making sound financial decisions and focus again on persistent choices. (Later, when estimating the 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 20% 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.

Table 4, 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.¹³

¹³ Inattention is a facet of impulsivity (Weafer et al. 2013). In our survey impulsive respondents are less likely to pass the manipulation checks. Excluding respondents who failed the IMCs does not change the substantive findings of this section.

TABLE 4: *Persistent Adviser Choices by Topic Sequence and Participant Characteristics*

	Persistent choices (%)	Total observations
All	21.9	3195
Panel A		
Financial topic sequence		
<i>low ambiguity → low ambiguity</i>	7.8	551
<i>low ambiguity → high ambiguity</i>	28.8	1,149
<i>high ambiguity → low ambiguity</i>	13.0	973
<i>high ambiguity → high ambiguity</i>	38.5	522
Panel B		
Participants		
<i>capable</i>	19.7	2,612
<i>non-capable</i>	34.8	583
<i>non-capable & impulsive</i>	38.2	387
<i>non-capable & 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 continue 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 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 ($\chi^2(3) = 225.49, p < .001$). Panel B shows persistence by participant characteristics where ‘capable’ (non-capable) participants have financial literacy and numeracy above (below) the median ($\chi^2(1) = 69.07, p < .001$). The label ‘impulsive’ (non-impulsive) divides the non-capable group into those who score above (below) the median in their responses to the impulsivity scale ($\chi^2(1) = 5.94, p < .05$).

Model of Consumer Learning and Confirmation Bias

According to our conceptual framework, participants’ latent prior beliefs about the advisers, their latent learning process type, as well as their latent perception of advice as ambiguous or clear, determine their posterior beliefs after each choice set and thus their advice

choices and willingness to pay for the advisers. We now set out the model that allows us to identify these latent factors.

Linking Client Characteristics to 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.

We assume that at choice set c , participant i holds a prior belief 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. 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 the adviser they rate lower as bad. Bayesian learners will form an updated belief $\lambda_{i,c,a}$ according to the Bayes rule and ignore their subjective interpretation of the advice if the signal is ambiguous, whereas limited memory learners update using their readily available interpretation of ambiguous signals. As such, and in line with Figure 1, Bayesian learners only update their beliefs when the advice is clear, whereas limited memory updaters also update their beliefs when the advice is ambiguous (see Web Appendix D for details on the calculation of $\lambda_{i,c,a}$ for each learning type).

We define a series of logit functions that link participants' beliefs about the advisers to the choices they make in the experiment as well as to participant and adviser characteristics. First, in the special case of choice set $c = 1$, we estimate the initial ($c - 1 = 0$) prior belief 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 (Hansen 2012) (see also Table 2), and an unknown vector of parameters, β_0 :

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

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 . When participant i interprets the advice as clear, then they should choose to follow—up to some random error with scale β_1 —the good advice and not the bad advice. 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 ($\lambda_{i,(c-1),R}$) versus adviser L ($\lambda_{i,(c-1),L}$), the more likely the participant should be to follow adviser R (and vice versa). We again allow for a random error with scale β_2 in this choice and write the probability as:

$$P(y_{i,c} = 1 | \tau_{i,c} = amb) = \frac{\exp(\beta_2 (\lambda_{i,(c-1),R} - \lambda_{i,(c-1),L}))}{1 + \exp(\beta_2 (\lambda_{i,(c-1),R} - \lambda_{i,(c-1),L}))}, \quad (2)$$

Both types of learners' interpretations of signals, and therefore choices of advisers, thus 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 $\lambda_{i,(c-1),a}$.

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.¹⁴ 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. 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 2).

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

¹⁴ This 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.

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 p_i of an additional unit of advice, and the estimated posterior beliefs about the adviser's quality, $\lambda_{i,4,a}$ (Gennaioli, Shleifer, and Vishny 2015), and parameters β_5 (see Table 2).

Model likelihood and estimation

The likelihood function of our 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 two learning classes, $(\tau_{Clarity}, \tau_{Learner}) = \tau_j \in j = 1, \dots, 32$, 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:

$$l_i(\theta|\tau_j) = \prod_{c=1}^4 P(y_{i,c} = 1|\tau_j)^{y_{i,c}=1} \cdot P(w_{i,R} = 1|\tau_j)^{w_{i,R}=1} \cdot P(w_{i,L} = 1|\tau_j)^{w_{i,L}=1}. \quad (3)$$

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 :

$$l_i(\theta) = \sum_{j=1}^{32} P_i(\tau_j) l_i(\theta|\tau_j). \quad (4)$$

We use the Sequential Adaptive Bayesian Learning (SABL) algorithm developed by Durham and Geweke (2014) to estimate the model. Web Appendix E outlines the estimation procedure, while Web Appendix F provides a brief sketch of parameter identification.

Estimation and Results

Model Fit and Parameter Estimates

We estimated the model using data from 1,903 of the 2,003 participants and reserved the remaining participants' responses to assess hold-out fit. Table 5 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 5: *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 impulsivity	-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. Estimates are based on the 1,903 randomly selected participants from 2,003 responses collected in December 2014. Variables are defined in Table 3. 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 E for detail

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 $\tau_{\text{Learner}} = \text{Bayesian}$ for all participants), consistent with conventional learning models. The log marginal density for this restricted model is -5887.17, compared to -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 (Kass and Raftery 1995).

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 (see Web Appendix G). 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 in Table 5 and rates of correct choices in the raw data in Table 3.

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

We use estimates from Table 5 to show how covariates affect 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.¹⁵

Client's belief about an adviser – and consequently their willingness-to-pay for this adviser depends on the belief they hold about the adviser before they receive the first piece of advice (initial prior belief). On average, participants believe that an adviser is good with a

¹⁵ See Web Appendix H for graphs of the marginal effects associated with estimated parameters.

probability of 85% (determined by Equation 1: $\exp(1.728)/(1+\exp(1.728))$). 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. 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. 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 find that participants who are not impulsive are 8.31 percentage points more likely to be Bayesian updaters. 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, whether the participant has previously paid for financial advice is 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. 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.

Summarizing, our model supports findings from previous research into first impressions, trust in advisers, consumer learning, and willingness to pay for advice. Yet, it is the first that allows identification of how these mechanisms interact to lead vulnerable clients to overpay bad advisers. Furthermore, our model enables comparison of intervention effectiveness: Consider a client with above median trust in advisers and impulsivity, but below median financial literacy and numeracy (all other characteristics identical to the median client). Such a client is, on average, willing to pay \$104.80 for one hour with a previously unknown, certified adviser. If the adviser gives bad advice on a random topic, willingness to pay drops to \$93.51. Reducing the client's impulsivity or trust below median lowers willingness to pay by \$0.96 (-1.02%) and \$4.07 (-4.35%), respectively. Increasing the client's financial literacy or numeracy above median reduces willingness to pay by \$6.65 (-7.11%) and \$6.41 (-6.85%), respectively. Removing the adviser's certification in turn only reduces willingness to pay by \$0.47 (-0.50%).

Vulnerable Clients' Willingness to Pay: Simulation Results

We simulate the model to measure the costs potentially 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 include younger people in the vulnerable group. We identify these participants by their responses to survey questions answered before the choice task.

Figure 3, Panel A shows how the resilient, vulnerable and mean reference groups divide into Bayesian and 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 3: 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.

Figure 4 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.




Figure 4: 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).

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 participants (red line). Therefore, in a simplistic world where advisers only

give good advice, at average topic clarity, resilient and vulnerable participants 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-topic advice given is good. However, this first impression translates into a large difference between resilient and vulnerable participants' willingness to pay for an adviser. At each price point until just after \$A150, the vulnerable participant is likely to pay more than the resilient participant. 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. Unfortunately, the vulnerable participant will pay more for this adviser than the resilient participant, who are more able to distinguish good from bad advice and mark the adviser down.

Figure 5 shows 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 4 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), the difference is close to zero. When the last piece of advice is bad (sequence GGGB), differences begin to emerge. Once the adviser gives an equal amount of good and bad advice (GGBB), the difference becomes more apparent and lies above zero with a high probability.




Figure 5: Statistical difference between vulnerable and resilient participants' willingness-to-pay. Panel A plots the differences between the vulnerable and resilient participants' probability lines for each quality sequence shown in Figure 6. 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.

Most striking are the cases in which the adviser sets a good first impression and then delivers bad advice (GBBB) and the worst-case world (BBBB). For these instances, vulnerable participants are much more willing to pay advisers for their services than resilient participants. This finding demonstrates that in a realistic world where advisers do not always deliver good advice, vulnerable clients will be 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 clients with ways to differentiate between the quality of advisers could help decrease this divergence in willingness-to-pay.

To best illustrate the possible strategic exploitation of vulnerable clients, we focus only on the quality sequence that emphasizes the first impression (GBBB). In Figure 5, Panel B, we contrast the probabilities for GBBB shown in Panel A with the probabilities obtained for GBBB if the topics are presented in a strategic order. The strategic topic order is debt, consolidation, diversification and fees. We label this the ‘strategic scenario’ because debt is the clearest topic for most participants and so this topic sequence allows advisers to make the strongest positive first impression when they give good advice first. The dashed line in Panel B shows the differences between vulnerable and resilient participants’ willingness to pay for further advice, after receiving the *strategic* advice scenario. The solid line shows the differences between vulnerable and resilient participants’ willingness to pay from the non-strategic case, already illustrated in Panel A. Panel B therefore demonstrates that if an adviser chose to use observable characteristics to identify that a client was vulnerable, and to then strategically manipulate the sequence of advice topics, giving clearest first, 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 (expert) and identify a group of consumers who are 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 consistent with confirmation bias. It can explain polarized opinions despite individuals receiving the same information signals. Our model identifies a consumer's learning process and the characteristics associated with it.

Our research informs multiple areas of literature. First we add to the body of work on trust in service encounters (e.g., Sirdeshmukh, Singh, and Sabol 2002), as well as advice provision and utilization in expert services (Wittkowski et al. 2020) by explaining the foundation and evolution of client trust in this context. Second, we contribute to literature in marketing that discusses alternative psychological models of belief updating (e.g., Chylinski, Roberts, and Hardie 2012), and extend literature on how information is processed in service settings in the context of credence goods (Keh and Sun 2018) by proposing a model that detects confirmation bias from individual choice data. In line with Habel et al. (2016), we explicitly acknowledge that customers are prone to behavioral biases and thus go beyond existing models that explain how customers construct service quality perceptions (Dagger and Sweeney 2007). Further, we present a model that explains the mechanisms that make consumers vulnerable, and thus answer a recent research call for a more nuanced and complex treatment of consumer vulnerability (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 model enables us to identify probable markers for limited memory updaters. 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 and other instances of intentional misdirection. Knowledge of these mechanisms 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 refuses to pay for it, vulnerable participants continue to pay. The results show that the failure of vulnerable clients to discern bad quality offerings is likely to hold even in the absence of malevolent actors, where, say, advisers do not use strategic manipulation or catering, as suggested by earlier studies. Further analysis predicts that in the presence of strategic manipulation, vulnerable clients pay significantly higher fees.

In the context of financial advice, we identify vulnerable and resilient participants in our sample as differing by age, financial literacy, numeracy, prior trust of advisers and impulsivity. We demonstrate that impulsivity 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 participants' ability to discern good from bad advice. Taken together, these characteristics significantly impact participants' willingness-to-pay for advisers and our model enables us to identify the contribution each of these characteristics has on associated welfare loss. 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: which initiatives help vulnerable consumers make better choices? And in times of tight budgets, which initiatives should be prioritized?

Implication

Our model suggests three different means to help vulnerable consumers make better choices. First, marketers and policy makers need to ensure that prior beliefs match reality. Misplaced trust may be corrected by providing risk information (Menon, Block, and Ramanathan 2002), such as checklists that aid consumers in identifying high quality advisers or requiring objective quality signals of trustworthiness such as professional certification. If certifications are signals of superior quality, they can provide helpful information to the consumer. However, as it stands, many certifications of varying quality are available worldwide and in different contexts including not only financial advice (e.g. Goetz et al. 2011), but ranging from food labelling (Kuchler, Sweitzer, and Greene 2020) to educational attainment (Schneider 2009). Consumers need guidance to identify reliable certifications; marketers should only use certifications that are clear signals of superior quality, and regulators could consider standardizing the certifications available. The benefit of certifications for the service provider thus goes beyond the effect of reduced intangibility discussed in previous literature (Laroche et al. 2004).

Second, marketers and policy makers should aim to reduce limited memory updating. Our study showed that impulsivity is a major contributor to this kind of learning process and we therefore suggest that policy should focus on curbing consumer impulsivity. Marketers may use visceral stimuli (Thomas, Desai, and Seenivasan 2011), such as the pain of paying with cash, as a useful intervention to regulate impulsivity. Moreover, the time gap between service use and its assessment affects how that experience is evaluated (Pizzi et al. 2015), public policy makers, on the other hand, may stipulate cool-down periods between getting and acting on advice.

Third, marketers and policy makers need to help consumers correctly identify the quality of signals. Industry or government initiatives may improve 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 to make a difference. Combined with previous research that shows that enhancing customers' knowledge can strengthen trust and serve as an important service differentiator (Eisingerich and Bell 2008), our research thus also makes a strong case for individual service providers to invest in customer education.

In our experiment, we find that the third initiative, i.e. helping consumers to correctly identify the quality of signals by increasing their financial literacy and numeracy, has the highest impact on consumer welfare, whereas taking away the certification of an adviser has the least impact, yet still shows an improvement of 0.5% in customer welfare. Depending on the costs of the different initiatives, the associated elasticities may provide further guidance on which initiatives policymakers should prioritize.

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 clients' learning style.

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 (Brüggen et al. 2025). This is particularly relevant given that with advances in AI, consumers become increasingly easy to read and thus easy to manipulate by malevolent actors.

Further, 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 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 may 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 the impact of events that are not direct interactions between adviser and consumer once prior beliefs have been formed.

Gurun, Stoffman, and Yonker (2018) however, demonstrate that, in the wake of the notorious Madoff fraud, trust disruptions propagate through social networks, leading clients to broadly lose confidence in financial advisers after the fraud is exposed. Future research might investigate what kind of events have the potency to influence consumers' learning.

References

- Agnew, Julie R., Hazel Bateman, Christine Eckert, Fedor Iskhakov, Jordan Louviere, and Susan Thorp (2018), "First Impressions Matter: An Experimental Investigation of Online Financial Advice," *Management Science*, 64 (1), 288-307.
- Australian Securities and Investments Commission (2012), "Report Guide 279—Shadow shopping study of retirement advice," Canberra, Australia: ASIC.
- Baeckström, Ylva, Ian W. Marsh, and Joanne Silvester (2021), "Variations in Investment Advice Provision: A Study of Financial Advisors of Millionaire Investors," *Journal of Economic Behavior & Organization*, 188, 716-735.
- Berens, Michael. (2023), "How Doctors Buy Their Way Out of Trouble," *Reuters*. May 24. <https://www.reuters.com/investigates/special-report/usa-healthcare-settlements/>
- Bone, Sterling A., Glenn L. Christensen, and Jerome D. Williams (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, Lex, Angela Lee Duckworth, James J. Heckman, and Bas Ter Weel (2008), "The Economics and Psychology of Personality Traits," *Journal of Human Resources*, 43 (4), 972-1059.
- Brown, W. (1995). "Wrenching Frustrations," *The Washington Post*. December 2. <https://www.washingtonpost.com/archive/business/1995/12/02/wrenching-frustrations/8f188854-6d34-4f39-9d55-13628b43c80b/>
- Brüggen, Lisa, Robert Gianni, Floris de Haan, Jens Högrevé, Darian Meacham, Thomas Post, and Minou van der Werf (2025), "AI-Based Financial Advice: An Ethical Discourse on AI-Based Financial Advice and Ethical Reflection Framework," *Journal of Public Policy & Marketing*, 44(3), 436-456.
- Bucher-Koenen, Tabea, Andreas Hackethal, Johannes Koenen, and Christine Laudenbach (2023), "Gender Differences in Financial Advice," SSRN Scholarly Paper ID 2572961, Social Science Research Network, Rochester, NY.
- Cenophtat, Sadrac, Martin Eisend, Tomás Bayón, and Alexander Haas (2024), "The role of customer relationship vulnerability in service recovery," *Journal of Service Research*, 27(2), 283-301.
- Chan, Kimmy Wa, Chi Kin (Bennett) Yim, and Simon S. K. Lam (2010), "Is Customer Participation in Value Creation a DoubleEdged Sword? Evidence from Professional Financial Services Across Cultures," *Journal of Marketing*, 74 (3), 48-64.
- Chaxel, Anne-Sophie, J. Edward Russo, and Neda Kerimi (2013), "Preference-Driven Biases in Decision Makers' Information Search and Evaluation," *Judgment and Decision Making*, 8 (5), 561-576.

Chylinski, Mathew B., John H. Roberts, and Bruce G.S. Hardie (2012), "Consumer Learning of New Binary Attribute Importance Accounting for Priors, Bias, and Order Effects," *Marketing Science*, 31 (4), 549-566.

Dagger, Tracey S., and Jillian C. Sweeney (2007), "Service Quality Attribute Weights: How do Novice and Longer-Term Customers Construct Service Quality Perceptions?," *Journal of Service Research*, 10(1), 22-42.

Durham, Garland, and John Geweke (2014), "Adaptive Sequential Posterior Simulators for Massively Parallel Computing Environments," in *Bayesian Model Comparison*, Emerald Group Publishing Limited, 1-44.

Edelen, R. M., Fong, K. Y., & Han, J. (2025). Regulating inattention in fee-based financial advice. *Journal of Financial Economics*, 164, 103985.

Egan, Mark, Shan Ge, and Johnny Tang (2022), "Conflicting Interests and the Effect of Fiduciary Duty: Evidence from Variable Annuities," *The Review of Financial Studies*, 35 (12), 5334-5386.

Egan, Mark, Gregor Matvos, and Amit Seru (2019), "The Market for Financial Adviser Misconduct," *Journal of Political Economy*, 127 (1), 233-295.

Eisingerich, Andreas B., and Simon J. Bell (2008), "Perceived Service Quality and Customer Trust: Does Enhancing Customers' Service Knowledge Matter?," *Journal of Service Research*, 10 (3), 256-268.

Emons, Winand (1997), "Credence Goods and Fraudulent Experts," *The RAND Journal of Economics*, 28 (1), 107-119.

Erdem, Tülin, and Michael P. Keane (1996), "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15 (1), 1-20.

Fryer Jr., Roland G., Philipp Harms, and Matthew O. Jackson (2019), "Updating Beliefs When Evidence Is Open to Interpretation: Implications for Bias and Polarization," *Journal of the European Economic Association*, 17 (5), 1470-1501.

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

Gerrans, Paul, and Douglas A. Hershey (2017), "Financial adviser anxiety, financial literacy, and financial advice seeking," *Journal of Consumer Affairs*, 51(1), 54-90.

Glimcher, Paul William, Joseph Kable, and Kenway Louie (2007), "Neuroeconomic Studies of Impulsivity: Now or Just as Soon as Possible?" *American Economic Review*, 97 (2), 142-147.

Goetz, Joseph W., Dandan Zhu, Vickie L. Hampton, Swarn Chatterjee, and John Salter (2011), "Integration of Professional Certification Examinations with the Financial Planning Curriculum: Increasing Efficiency, Motivation, and Professional Success," *American Journal of Business Education*, 4 (3), 35-46.

Gurun, Umit G., Noah Stoffman, and Scott E. Yonker (2018), "Trust Busting: The Effect of Fraud on Investor Behavior," *The Review of Financial Studies*, 31 (4), 1341-1376.

Habel, Johannes, Sascha Alavi, Christian Schmitz, Janina-Vanessa Schneider, and Jan Wieseke (2016), "When do Customers Get What They Expect? Understanding the Ambivalent Effects of Customers' Service Expectations on Satisfaction," *Journal of Service Research*, 19 (4), 361-379.

Hansen, Torben (2012), "Understanding Trust in Financial Services: The Influence of Financial Healthiness, Knowledge, and Satisfaction," *Journal of Service Research*, 15 (3), 280-295.

Hill, Ronald Paul and Eesha Sharma (2020), "Consumer Vulnerability," *Journal of Consumer Psychology*, 30 (3), 551-570.

Hoch, Stephen J., and John Deighton (1989), "Managing What Consumers Learn from Experience," *Journal of Marketing*, 53 (2), 1-20.

Johnson, Joseph. M., Gerrard J. Tellis, Noah VanBergen (2021), "Fooled by Success: How, Why, and When Disclosures Fail or Work in Mutual Fund Ads," *Journal of Public Policy & Marketing*, 41 (1), 54-71.

Kass, Robert E., and Adrian E. Raftery (1995), "Bayes Factors," *Journal of the American Statistical Association*, 90 (430), 773-795.

Keh, Hean Tat, and Jin Sun (2018), "The Differential Effects of Online Peer Review and Expert Review on Service Evaluations: The Roles of Confidence and Information Convergence," *Journal of Service Research*, 21(4), 474-489.

Kominers, Scott Duke, Xiaosheng Mu, and Alexander Peysakhovich (2018), "Paying (for) Attention: The Impact of Information Processing Costs on Bayesian Inference," Available at SSRN 2857978.

Konya-Baumbach, Elisa, Monika C. Schuhmacher, Sabine Kuester, and Victoria Kuharev (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, Fred, M. Bowman, M. Sweitzer, and Catherine Greene (2020), "Evidence from Retail Food Markets That Consumers Are Confused by Natural and Organic Food Labels," *Journal of Consumer Policy*, 43, 379-395.

Langenderfer, Jeff, and Terence A. Shimp (2001), "Consumer Vulnerability to Scams, Swindles, and Fraud: A New Theory of Visceral Influences on Persuasion," *Psychology & Marketing*, 18 (7), 763-783.

Lantzy, Shannon, Rebecca W. Hamilton, Yu-Jen Chen, and Katherine Stewart (2020), "Online Reviews of Credence Service Providers: What Do Consumers Evaluate, Do Other Consumers Believe the Reviews, and Are Interventions Needed?" *Journal of Public Policy & Marketing*, 40 (1), 27-44.

Laroche, Michel, Gordon H. G. McDougall, Jasmin Bergeron, and Zhiyong Yang (2004), "Exploring How Intangibility Affects Perceived Risk," *Journal of Service Research*, 6(4), 373-389.

- Linnainmaa, Juhani T., Brian T. Melzer, and Alessandro Previtero (2021), "The Misguided Beliefs of Financial Advisors," *The Journal of Finance*, 76 (2), 587-621.
- Lipkus, Isaac M., Greg Samsa, and Barbara K. Rimer (2001), "General Performance on a Numeracy Scale Among Highly Educated Samples," *Medical Decision Making*, 21 (1), 37-44.
- Lusardi, Annamaria and Olivia S. Mitchell (2011), "Financial Literacy Around the World: An Overview," *Journal of Pension Economics & Finance*, 10 (4), 497-508.
- Menon, Geeta, Lauren G. Block, and Suresh Ramanathan (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, Sendhil, Markus Noeth, and Antoinette Schoar (2012), *The Market for Financial Advice: An Audit Study*, National Bureau of Economic Research.
- Muthukrishnan, A. V. and Amitava Chattopadhyay (2007), "Just Give Me Another Chance: The Strategies for Brand Recovery from a Bad First Impression," *Journal of Marketing Research*, 44 (2), 334-345.
- Nagel, Stefan and Zhengyang Xu (2022), "Asset Pricing with Fading Memory," *The Review of Financial Studies*, 35 (5), 2190-2245.
- Oppenheimer, Daniel M., Tom Meyvis, and Nicolas Davidenko (2009), "Instructional Manipulation Checks: Detecting Satisficing to Increase Statistical Power," *Journal of Experimental Social Psychology*, 45 (4), 867-872.
- Oreg, Shaul, and Mahmut Bayazit (2009), "Prone to bias: Development of a bias taxonomy from an individual differences perspective," *Review of General Psychology* 13(3), 175-193.
- Pizzi, Gabriele, Gian Luca Marzocchi, Chiara Orsingher, and Alessandra Zammit (2015), "The Temporal Construal of Customer Satisfaction," *Journal of Service Research*, 18(4), 484-497.
- Rabin, Matthew and Joel L. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114 (1), 37-82.
- Roberts, Brent W., Carl Lejuez, Robert F. Krueger, Jessica M. Richards, and Patrick L. Hill (2014), "What Is Conscientiousness and How Can It Be Assessed?" *Developmental Psychology*, 50 (5), 1315-1330.
- Roberts, John H. and Glen L. Urban (1988), "Modeling Multiattribute Utility, Risk, and Belief Dynamics for New Consumer Durable Brand Choice," *Management Science*, 34 (2), 167-185.
- Schneider, Silke (2009), "Confusing Credentials: The Cross-Nationally Comparable Measurement of Educational Attainment," PhD Thesis, Oxford University, UK.
- Schumann, Jan H., Florian von Wangenheim, Anne Stringfellow, Zhilin Yang, Sandra Praxmarer, Fernando R. Jimenez, Vera Blazevic, Randall M. Shannon, Shainesh G., and Marcin Komor (2010), "Drivers of trust in relational service exchange: understanding the importance of cross-cultural differences," *Journal of Service Research*, 13(4), 453-468.

Seiders, Kathleen, Andrea Godfrey Flynn, Leonard L. Berry, and Kelly L. Haws (2015), "Motivating Customers to Adhere to Expert Advice in Professional Services: A Medical Service Context," *Journal of Service Research*, 18(1), 39-58.

Sirdeshmukh, Deepak, Jagdip Singh, and Barry Sabol (2002), "Consumer Trust, Value, and Loyalty in Relational Exchanges," *Journal of Marketing*, 66 (1), 15-37.

Soll, Jack B., Asa B. Palley, and Christina A. Rader (2022), "The Bad Thing About Good Advice: Understanding When and How Advice Exacerbates Overconfidence," *Management Science*, 68 (4), 2949-2969.

Song, Jaeki and Others (2007), "Trust in Health Infomediaries," *Decision Support Systems*, 43 (2), 390-407.

Stolper, Oscar and Andreas Walter (2019), "Birds of a Feather: The Impact of Homophily on the Propensity to Follow Financial Advice," *The Review of Financial Studies*, 32 (2), 524-563.

Thomas, Manoj, Kalpesh Kaushik Desai, and Satheeshkumar Seenivasan (2011), "How Credit Card Payments Increase Unhealthy Food Purchases: Visceral Regulation of Vices," *Journal of Consumer Research*, 38 (1), 126-139.

Tsukayama, Eli, Angela Lee Duckworth, and Betty Kim (2012), "Resisting Everything Except Temptation: Evidence and an Explanation for Domain-Specific Impulsivity," *European Journal of Personality*, 26 (3), 318-334.

Weafer, Jessica, Matthew J. Baggott, and Harriet de Wit (2013), "Test-Retest Reliability of Behavioral Measures of Impulsive Choice, Impulsive Action, and Inattention," *Experimental and Clinical Psychopharmacology*, 21 (6), 475-481.

Wittkowski, Kristina, Jan F. Klein, Tomas Falk, Jeroen J. L. Schepers, Jaakko Aspara, and Kai N. Bergner (2020), "What Gets Measured Gets Done: Can Self-Tracking Technologies Enhance Advice Compliance?" *Journal of Service Research*, 23(3), 281-298.

Yaniv, Ilan and Eli Kleinberger (2000), "Advice Taking in Decision Making: Egocentric Discounting and Reputation Formation," *Organizational Behavior and Human Decision Processes*, 83 (2), 260-281.