Show and Tell?
An Experimental Examination of Customer Signals and Expert Requests in Conflict of Interest Situations

By
Ingrid O. Koch

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Business)

at the

UNIVERSITY OF WISCONSIN-MADISON

2019

Date of final oral examination: 06/06/2019

The dissertation is approved by the following members of the Final Oral Committee:

Tai S. Lim, Professor, Marketing

Teck-Hua Ho, Professor, Economics

Jan B. Heide, Professor, Marketing

Jordan Tong, Associate Professor, Operations & Information Management

Kevin Chung, Assistant Professor, Marketing

Acknowledgments

This dissertation would not have been possible without the massive amount of mentoring and support I've been fortunate enough to receive. Words do not allow me to express adequate thanks to my chair, Dr. Noah Lim, who believed in me always. My research would not be close to where it is today without your guidance and support, and neither would I. Thank you for everything – I am grateful always. Further, I cannot begin to express my gratitude for the help and direction provided by Dr. Teck Ho, whose extensive knowledge and quick-thinking is both inspiring and humbling. Thank you for always finding time to listen to my research ideas, and immediately improving them. I am also extremely grateful to Dr. Kevin Chung, who was the first person to ever teach me marketing research. Thank you for continuing to give advice and guidance, and for listening and understanding. My deepest gratitude goes to Dr. Jan Heide, who provided me with a fresh perspective and bestowed thoughtful and practical advice that greatly improved the positioning and scope of my research. And I am indebted to Dr. Jordan Tong, who was willing to help me even on a moment's notice. Your guidance illuminated the most worthwhile parts of my research and saved me from running into a brick wall down the line.

Beyond the members of my dissertation committee, I've received endless support and direction from other members of the academic community. My deepest thanks go to Dr. Jiabin Wu, who repeatedly lent me his economic expertise to look over and correct my equilibrium derivations. I am sincerely grateful to Dr. Jeeva Somasundaram, who was a mentor in the art of running field experiments and was always happy to patiently explain concepts to me (sometimes multiple times). Many thanks also go to Dr. Hua Chen, who provided thoughtful constructive criticism and useful suggestions that helped to shape this dissertation.

My research would not have been possible if not for the extremely generous funding provided by the Global Asia Institute (GAI) at NUS and the marketing department at the University of Wisconsin – Madison. I am forever grateful to the members of these communities who have assisted me along the way, including the faculty and doctoral students in the marketing department at UW-Madison. A special thanks go to Dr. Paul Hoban and Dr. Liad Weiss, who provided insightful feedback and suggestions on my dissertation proposal. And to Kym Aebly, who was always happy to share her administrative savviness with a smile. I am also forever grateful to have been hosted by GAI while completing research in Singapore and have received a huge amount of assistance from those at the institute. Specifically, many thanks go to Dr. Ta-Cheng Huang, who graciously shared his expertise in R and econometrics to improve the analysis of my data. I am also grateful to Dr. Aidas Masiliūnas, whose experimental knowledge helped greatly with the design and setup of my lab experiments. And a heartfelt thank you to Dr. Ashish Sachdeva, who reminded me to think big when I was focused on thinking small. Further, a special thanks goes to Min Si Lim, who assisted me in carrying out the taxi study. Thank you for keeping me (and the RAs) organized, and for your encouragement and friendship. To Lloyd Heng, Gee Lam Yau and Lini Kuang, thank you for always being willing to help out and participate in my research. Moreover, my sincerest gratitude goes to Sing Chik Tan and Duc Pham, who helped me manage countless rounds of subject recruitment and lab studies. Finally, a large thank you also goes to Teresa Nah, Jane Lim and Barathi Vadivelu for their never-ending administrative assistance.

Of course, I could not have completed this dissertation without the infinite love, support and understanding from my family and friends. Thank you for the endless letters and the early morning chats – your unyielding encouragement did not go unnoticed.

I hope, with this dissertation, I can make you all proud.

Table of Contents

ABSTRACT

Across many industries, consumers rely on professional advisors for information and expertise. Often these advisors have a conflict of interest (COI), such that they can potentially benefit from deceiving customers. In this paper, I present a series of 3 lab experiments and 1 field to delineate how customer signals and professional requests affect outcomes in COI situations. In two lab experiments, I demonstrate that customers often send trust signals to experts, and while weak trust signals (i.e., statements of trust) can improve outcomes, stronger trust signals (i.e., decision delegation) can backfire. Interestingly, however, the detrimental effect of delegation is dampened when experts can specifically request that customers delegate. Relatedly, in a field experiment in the context of taxi services, I find that passengers who signal misinformation are taken on longer and more expensive taxi rides. Moreover, consistent with the lab findings, evidence suggests that a strong signal of trust on top of misinformation can increase driver deception even further. The results have important implications for both customers and practitioners across various COI contexts.

1. Introduction

Conflicts of interest (COIs) permeate various industries worldwide, and consumers interact with experts who have COIs on a regular basis.[1] For instance, physicians often maintain close relationships with pharmaceutical companies, and may be offered kickbacks for prescribing certain medications.[2] Moreover, the incentive structure of financial advisors may lead them to bias their investment recommendations, while taxi drivers may take longer-than-necessary routes to earn a higher taxi fare. One specific type of COI arises when professionals sell credence goods. In a credence good market, advisors are better informed than advisees on which product or service the advisee requires, and thus advisees are often not able to determine whether they received the optimal good ex post. Examples of credence goods include repair services (e.g., vehicles and electronics), medical services, and taxi rides. Since professionals are better informed than customers on what product and/or service is optimal, they can use this information asymmetry to steer customers to purchase more expensive goods than necessary in order to reap personal rewards. Because of this, credence good markets are often plagued by problems of market inefficiency. In this paper, I examine COIs and credence goods, and evaluate factors that affect market efficiency as well as consumer and professional outcomes.[3]

---

[1] While the literature does not offer one standard definition of a COI, in this paper I will refer to a COI as arising when there is potential for personal considerations to compromise or bias professional judgment and objectivity.

[2] One research study conducted by Cegedim Strategic Data found that, in 2012, the pharmaceutical industry spent over $24 billion to promote their drugs directly to physicians (Cegedim Strategic Data, 2013). As I'll discuss in Section 1.2.2, this detailing can significantly affect the medications prescribed by physicians.

[3] COI situations are generally thought of as arising when a customer hires an expert to act as an advisor, and the expert is expected to provide objective advice. However, the expert may also have a personal incentive to inflate demand or steer demand to more costly than necessary goods and services. Common examples of professionals with COIs are physicians and financial advisors. Further, professionals who sell credence goods and services often have a COI, as their superior situational knowledge provides a foothold for advisor deception. While the ethical standard for professionals may be slightly less stringent in certain credence good contexts (e.g., car repair, taxi

Existing experimental research has evaluated how various influences affect outcomes in COI situations. For instance, field research has examined the effect of consumer traits such as price sensitivity (Kerschbamer et al. 2016; Lu 2014; Balafoutas et al. 2015), information (in)accuracy (Busse et al. 2017; Currie et al. 2011) and demographics (Balafoutas et al. 2013; Busse et al. 2017). Further lab experimental research has evaluated professional COIs in an advisor-advisee interaction, with a strong focus on the effects of COI disclosure (e.g., Cain et al. 2005, 2011; Sah et al. 2013; Church and Kuang 2009) and, to a lesser extent, second opinions (Sah and Loewenstein 2015; Mimra et al. 2016). While this work is certainly important, it is relatively silent on how consumers can strategically interact with experts to improve their outcomes or how behavioral influences can be used to reduce professional deception levels. Given this, the present paper has two focuses.

My primary focus is on customer signaling related to both trust and misinformation in COI contexts. I chose to examine these two factors because they arise often in advisee-advisor interactions – whether customers intend to signal them or not. For instance, customers may explicitly state their trust to professionals or, instead, imply their trust by asking advisors to make decisions on their behalf. Furthermore, consumers are gaining increasing access to information online and often choose to share this information with professionals. Given the abundance of misleading information available, it's likely that customers often share incorrect information with advisors. Thus, as both trust and misinformation are commonly signaled, a rigorous examination of their effect on decisions and outcomes in COI situations is warranted.

---

services), it is likely that customers in these situations still have an expectation of honest advice and service as they choose to utilize a professionals' services in the first place.

Second, I examine professional trust requests and how these can influence professional behavior. History shows that COIs are far-reaching and enduring. For instance, pharmaceutical companies are unlikely to stop detailing anytime soon, and commission-based pay for salespeople is a common and long-standing incentive structure across various markets. Because certain professionals will likely always face at least some incentive to act opportunistically, it is important to understand how experts themselves can be motivated to be honest. In my research, I study trust requests from an advisor to an advisee, as past research suggests that these requests may be able to increase advisor reciprocity and improve honesty levels (Ho and Weigelt 2005; Berg et al. 1995; Lim and Ham 2014). Further, these requests are easy to implement and common in practice. For instance, trust requests can be explicit statements (e.g., "Trust me, I've done this before") or more implicit appeals (e.g., doctor asking to perform patient's surgery). Thus, as trust requests have the potential to improve outcomes and are often utilized by professionals in practice, I empirically examine their effect in COI situations.

My research builds upon and extends several streams of existing work on professional COIs. First, the present research is related to the marketing and psychology literature on advisee-advisor interactions, which is largely focused on the effect of COI disclosure. For instance, Cain et al. (2005, 2011) find that disclosure can backfire and lead to worse outcomes for customers (see also Loewenstein et al. 2011; Sah et al. 2013). Competing work, however, has found that, in certain situations, disclosure can be beneficial (Sah and Loewenstein 2013; Church and Kuang 2009) or have a null effect altogether (Ismayilov and Potters 2013). Apart from work on COI disclosure, there is also limited research evaluating the effect of second opinions and advisor payments in COI situations (Sah and Loewenstein 2015; Angelova and Regner 2013). While the extant research is interesting and important, the literature could benefit from an examination of additional mechanisms that can be used to reduce advisor bias and restore efficient market outcomes in a COI context. In Studies 1 and 2, I evaluate how customer trust signals and professional trust requests influence professional bias levels and customer outcomes. Further, in Study 3, I offer an examination of how customer misinformation signals affect outcomes in a credence good context, both when customers also signal trust and when they do not.

Second, my research is related to the literature examining behavioral preferences, specifically trust, when parties may have reasons not to trust one another. Researchers have found that, even in one-shot and anonymous contexts, there is some degree of trusting behavior, even though the majority of people aren't trustworthy (Berg et al. 1995; Ho and Weigelt 2005). Moreover, related work asserts that this trust is driven by reciprocity concerns (McCabe et al. 2003; Berg et al. 1995). In Studies 1-3, I extend this research by evaluating not only derived measures of actual trust and trustworthiness, but also different types of customer trust signals and professional trust requests.

Third, field experimental research in marketing and economics has evaluated professional COIs, specifically those that arise from the sale of credence goods. For instance, in a taxi context, Balafoutas et al. (2013) find that foreigners are more likely to be cheated than locals, and that taxi drivers are more likely to overcharge passengers who signal that their ride will be expensed (Balafoutas et al. 2015; see also Kerschbamer et al. 2016 and Lu 2014). Further research has shown that mechanics consider future business potential and customer expertise when determining repair prices (Schneider 2012; Busse et al. 2017). I build upon this work in Studies 3 and 4 by showing that customers often share inaccurate information with professionals, who then take advantage of the mistaken advisee.

Fourth, the operations literature offers an examination of a channel relationship where a better-informed manufacturer has an incentive to bias information sent to a lesser-informed supplier. Conceptually, this context is similar to that found in an advisor-advisee interaction, and indeed, the game setup used in Studies 1 and 2 is borrowed from the literature. The existing research has found that, due to behavioral considerations, the parties cooperate resulting in effective information sharing that is not predicted by standard economic wisdom (e.g., Özer et al. 2011, 2014, 2018). The research also offers an understanding of how supplier trust and manufacturer trustworthiness levels change across different channel contexts. Studies 1 and 2 extend this work, not only by offering an evaluation that is focused on individual outcomes instead of channel efficiency, but also by examining how *signals* of trust (whether true or not) can influence players' outcomes. That is, my interest is on the impact of trust signals, rather than how actual levels of trust change based on market characteristics.

Finally, I build upon the theoretical literature on signaling games and agency theory. In traditional game theory, signaling games are used when a better-informed "sender" has private information, which can be signaled to a lesser-informed "receiver" (Spence 1973; Spence 2002; Crawford and Sobel 1982). The present research extends this work in one key way. In particular, in Studies 1 and 2, the advisor and advisee each play the role of both the sender and receiver. That is, each party has private information that can be shared with the other party. This is novel, as the past research assumes that the party with less information does not signal to the better-informed party. Relatedly, my research draws from the agency theory literature. In particular, the advisor in a COI context can be thought of as an "agent" who makes decisions on behalf of the advisee, the "principal" (Jensen and Meckling 1976; also see Bergen et al. 1992 and Eisenhardt 1989 for overviews of the literature). In Study 2, I examine whether it's better for an advisee to delegate decision power to the advisor ("agent") or to retain decision power for herself instead.[4] Further, my work is related to an additional agency relationship where a professional service firm (e.g., consulting firm, hospital, etc) is the principal that delegates work to its agents, the professional employees (e.g., financial advisors, doctors). In this case, the firm has a reputation-driven incentive to keep employees honest. In Study 2, I evaluate how professional trust requests can be used to encourage experts to be honest, even in the face of tempting COIs.

In this paper, I present both field and lab experiments that make important contributions to the literature. In particular, Study 1 evaluates how customer signals of trust and distrust influence advisor bias and advisee outcomes in a COI situation. I find that most advisees choose to send trust signals, rather than distrust signals, even when advisees don't actually trust the advisor. Further, I find some evidence that these trust signals can reduce advisor bias, while the effect on customer earnings is less clear. I extend this research in Study 2, where I take a closer look at different types of trust signals, examining both simple cheap talk trust statements and more extreme shows of trust (i.e., decision delegation). I find that trust statements reduce advisor bias and improve customer outcomes, but that acts of trust (decision delegation) can backfire and have a deleterious effect on outcomes. However, this harmful effect is dampened when professionals are able to send trust requests to advisees. Furthermore, in Study 3, I

[4] For ease of discussion in this paper, I use feminine pronouns when referring to advisees and masculine pronouns when referring to advisors.

examine when information signaling may be detrimental for consumers in the context of taxi services. I find that passengers who signal inaccurate knowledge of the fastest taxi route are taken on longer and more expensive rides than passengers who do not signal this misinformation. Surprisingly, the data suggests that taxi driver deception is increased even further when customers signal trust along with misinformation. Finally, I find that customers do often choose to share potentially inaccurate information with professionals in Study 4, and this willingness to share underscores the importance of understanding the effect of customer misinformation in COI situations.

1.2 Literature Review

In this section, I outline research related to COIs that has already been done. As the existing literature is vast, I focus on the work most relevant for the research presented in this paper.

1.2.1 Signaling Games

I begin with a discussion of theoretical signaling games, which have been the foundation of much of the existing experimental COI research related to information sharing from an advisor to an advisee. Specifically, this context closely resembles a cheap talk sender-receiver game, as in Crawford and Sobel (1982). In their seminal paper, the authors present a cheap talk model of advice, where a better-informed advisor shares information with a lesser-informed advisee. They show that, as the two players' interests become more aligned, the advisee's use of the information provided by the advisor increases, as does the expected utility to each player. However, if the players' incentives are too far misaligned, this leads to uninformative communication. The model assumes that the players' payoffs are common knowledge – that is, the advisor's COI is disclosed to the advisee (see also Spence 1973; Spence 2002). I build on this setup in my research but make one important modification. In particular, while only one party (the "sender") signals information in traditional signaling games, here I allow both parties, the advisor and advisee, to signal their private information. While the advisor is better informed on the value of a mutual decision variable (consistent with traditional setup), the advisee is better informed of her own beliefs. In reality, advisees' beliefs often drive their decisions, so knowledge of advisee beliefs could be beneficial for advisors. Further, as advisees generally hold a large degree of power in advisor-advisee interactions in reality, I introduce advisee signaling to capture this influence in my research.

1.2.2 Conflicts of Interest

Marketing and Psychology Literature. Much of the existing marketing and psychology research on COIs in advisee-advisor interactions builds upon and tests the theoretical work on information signaling presented above. For example, Cain et al. (2005) demonstrate that when an advisor's incentive is too far misaligned with his advisee's incentive, the advisor communicates information that is indeed uninformative. To examine this experimentally, the authors assigned participants to play the role of an advisor or advisee.[5] Advisors were given better information on the values of jars of coins than advisees and were asked to send estimates of the coin values to advisees. The advisees then submitted guesses of the coin jar values, and their guesses determined the payoff of both players. In particular, an advisee's payoff increased in the accuracy of her guess. And while some advisors were assigned to a condition where their payoff was also increasing in

[5] Here, I delineate the game setup in relative detail, as this setup is similar to others used in the literature. In this way, readers can get a better sense of the type of research that has been done in this area.

the advisee's accuracy (NO COI condition), others' payoffs were increasing in the level of the advisee's guess (COI condition). Not surprisingly, advisors in the NO COI condition gave lower estimates of the jar values than those in the COI condition. This provides experimental evidence that COIs can bias the information advisors share with advisees. Further, of those advisors who were assigned to the COI condition, some had their COI disclosed to their advisee (COI/DISCLOSED condition), while others did not (COI/UNDISCLOSED condition). In theory, disclosure is beneficial for advisees because the increased transparency encourages advisors to tell the truth. However, the findings suggest that disclosure can backfire in practice. Specifically, advisors in the COI/DISCLOSED condition gave more biased estimates than those in the COI/UNDISCLOSED condition. Moreover, as advisees did not sufficiently account for the bias in advisors' estimates, advisees in the COI/DISCLOSED condition earned less than advisees in the COI/UNDISCLOSED condition. Thus, disclosure of COIs was found to be detrimental for advisees. The authors propose two mechanisms, strategic exaggeration and moral licensing, that may drive the increased advisor bias with disclosure. In particular, the authors posit that advisors might strategically exaggerate their advice in order to counteract the increased discounting they expect advisees to have when COIs are disclosed. For example, car sellers might initially quote a high price in anticipation of car purchasers bargaining to bring the price down. Secondly, the authors propose a moral licensing effect, where disclosure of COIs leads to advisors feeling less guilty about biasing their advice, since advisees were forewarned about the advisor's incentives.

Building off their 2005 paper, Cain et al. (2011) attempt to understand these psychological mechanisms in more detail. In a setup similar to the coin jar experiment in their 2005 paper, the researchers asked advisors to self-report why they might give an estimate that differs from the truth, even if they know this could potentially hurt the advisee. Responses indicated heterogeneity in advisors' judgements: some reported strategically exaggerating their advice in anticipation of advisee discounting, whereas some responded that they'd strategically restrained their advice because the advisee might be suspicious of especially high estimates. Further, the authors find evidence of moral licensing, as advisors with a disclosed COI rated the ethicality of sending a high estimate as being higher than those whose COI was not disclosed.

Taken together, the above two papers not only suggest that advisor COIs can be detrimental for advisees, but that a mechanism commonly put in place to reduce advisor bias caused by COIs (i.e., disclosure) can actually make advisees even worse off due to advisors' strategic exaggeration and moral licensing effects. From an advisee standpoint, additional research has examined the psychological mechanisms that prevent the advice seekers from fully discounting biased advisor information. Sah et al. (2013) find that disclosure of COIs can impose a burden on advisees, such that advisees feel increased pressure to comply with advice when their advisor's COI has been disclosed, as not complying insinuates that the advisees expect their advisor to be influenced by their COI (termed "insinuation anxiety"). Moreover, the authors propose a second mechanism behind advisees' burden of disclosure, which they coin the "panhandler effect". The panhandler effect occurs when an advisor's COI is disclosed to advisees, who then feel the need to help the advisor satisfy his own personal interests. Indeed, in the paper, participants' self-reports indicate that, when their advisor's COI is disclosed to them, advisees trust the advice less, but feel greater pressure to help their advisors. Thus, disclosure not only drives advisors with COIs to bias their advice, but it also pressures advisees into consciously making decisions that are suboptimal for themselves, in order to help their advisors.

In contrast to the above research, there is some evidence that disclosure of COIs is not always detrimental for advisees. For instance, Sah et al. (2013) find that advisees' increased pressure to comply with advice can be mitigated when the advisee can change her mind later or can make her decision in private. Further, Sah and Loewenstein (2013) show that disclosure can be beneficial when advisors can avoid COIs in the first place, while Church and Kuang (2009) demonstrate that disclosure is useful when advisees can impose (costly) sanctions on advisors for bad advice. Still other research has found that disclosure has a null effect on advisor bias altogether (Ismayilow and Potters 2013).

Undoubtedly, the topic of disclosure is important as many healthcare and financial policies are put into place specifically to ensure professionals disclose their COIs to consumers.[6] However, the effect of disclosure seems to be quite sensitive to the context it is used in, and much of the literature suggests that disclosure does not reduce professional bias. Additionally, disclosure requirements are usually determined at a high level of policy making, which leaves little room to provide consumers with recommendations for how to interact in COI situations, or to provide service firms with strategies to improve worker honesty. Thus, the literature could benefit from a rigorous investigation of how bias can be decreased through alternative strategies, such as consumer and professional signaling, that can be implemented across a variety of COI settings. The following paragraph summarizes the important, yet scarce, marketing and psychology literature that examines factors outside of COI disclosure.

---

[6] Disclosure of accounting information is mandated for public corporations in the Sarbanes-Oxley Act of 2001, and 2010's health reform bill includes a "sunshine" provisions making it required to disclose various payments physicians receive from pharma companies and medical device manufacturers.

Sah and Loewenstein (2015) experimentally evaluate the effect of second opinions on advisor's advice. Interestingly, they find that merely being aware that advisees have access to a second opinion (even if it's costly for the advisees to obtain) can lead to more biased advice from advisors. This bias is especially strong when the advisors know the second opinion is of low quality. The authors use self-reports of participants to determine that the presence of a second opinion evokes a profit-maximizing frame for advisors. With this profit-maximizing frame, advisors feel legitimized in acting in a self-interested manner, where they care more about their profits and less about helping their advisees. This suggests that second opinions, like disclosure, can lead to perverse effects. However, this may not always be the case, as further experimental research by Mimra et al. (2016) find that introducing second opinions decreases advisors' propensity to recommend unnecessary treatments. Apart from second opinions, research by Angelova and Regner (2013) finds that advisors are more truthful when they receive an upfront payment from advisees, and advisees who volunteered a payment followed advice more often than those who were obligated to pay.

Overall, the existing marketing and psychology research on COIs offers important, yet oftentimes contradictory, insights. This overview of the extant literature highlights the need for a more thorough examination of the mechanisms that can be used to thwart professional bias caused by COIs. As COI situations are common across many industries, I now briefly summarize research in disciplines outside of marketing and psychology, including operations, finance, accounting and healthcare, in order to obtain a clearer picture of how COIs play out in various settings. While the context of this work is different, I believe the core insights gleaned from the research are easily generalizable to marketing situations.

Operations Literature. The operations literature examines a channel relationship between a supplier and manufacturer in a context that resembles a COI situation. Specifically, the manufacturer – who is closer in proximity to the consumer market – often holds better information about demand than the supplier. The manufacturer may share its information with the supplier, but the manufacturer is often incentivized to inflate the forecast of future demand, to ensure sufficient supply. This incentive conflicts with the supplier's incentive to choose an accurate capacity, an excess of which is costly. The literature is interested in coordinating mechanisms that can be used to improve channel efficiency. Specifically, the research most similar to mine is that which examines the role of trust and trustworthiness in information sharing (Özer et al. 2011, 2018). The research finds that, due to trust and reciprocity, the parties cooperate resulting in effective information sharing that is not predicted by standard economic wisdom. Moreover, of particular interest to my work is research by Özer et al. (2018), who find that the type of assistance provided by an advisor to an advisee matters. For instance, information sharing and advice provision lead to higher levels of trust and trustworthiness compared to decision delegation. In Study 2, I examine information sharing and delegation as well. However, my study makes the important distinction of allowing the lesser-informed party to *choose* whether to delegate to the better-informed party whereas in Özer et al. (2018), delegation is obligatory. Given this, delegation would not be seen as a signal of trust from the delegating party. Thus, Study 2 extends and contributes to this work by examining how decision delegation can be used as a signal of trust, and how its efficacy compares to another type of trust signal (i.e. trust statements).

Finance and Accounting Literature. In the finance literature, COI research is focused largely on analysts and their incentives for biasing recommendations. Various research has found that analysts give biased “buy” recommendations for reasons such as firm affiliations and management appeasement (Michaely and Womack 1999; Ke and Yu 2006). Further, when analysts’ brokerage firms are affiliated with the firms they’re analyzing, analysts are slower than unaffiliated analysts to downgrade their recommendations from a buy rating (O’Brien et al. 2005; Mehran and Stulz 2007). These studies provide additional evidence that COIs can result in professionals communicating biased information to consumers.

Relatedly, accounting research is primarily concerned with the COI that arises when auditors are hired and paid by the firms they audit, which creates an incentive for auditors to bias their reviews in favor of the audited company. Using self-reports of full-time auditors, Moore et al. (2003) find that auditors are, on average, 30% more likely to find that a firm’s accounting practices are acceptable when they are hired by the assessed firm, compared to when they are hired by an outside investor. Further, given that the assessed firm hires them to perform an audit, auditors are more likely to bias their reports when they have a relatively large incentive to maintain a long-term relationship with the client (Beeler and Hunton 2002; Libby et al. 2008). Finally, social factors may also influence auditors’ assessments, as Wilks (2002) finds that auditors who learn of their superiors’ views earlier in the audit process tend to bias their evaluations of a client’s going concern to match their superiors’ viewpoints.

Healthcare Literature. There is also a large stream of healthcare literature looking at COIs. In their review of biomedical studies, Bekelman et al. (2003) find that almost a quarter of healthcare researchers have industry ties, and Perlis and colleagues’ (2005) review of psychiatry clinical trial research found that 60% of researchers received funding from pharmaceutical companies. Importantly, both papers find a significant and positive relationship between industry sponsorship and pro-industry conclusions in the research. This signifies that healthcare researchers may bias their published results in favor of their financial sponsors, which, in turn, may lead physicians to bias the care they provide for patients. Indeed, Silverman et al. (2010) find that physicians pay more attention to the statistical findings of published research than the COIs that the researchers disclose for the work, suggesting that physicians may not properly account for researcher bias. Additionally, work by Caudill and colleagues (1996) finds that the more contact physicians have with pharmaceutical companies, the higher the cost of physicians' treatments, while Wazana (2000) similarly finds that pharmaceutical detailing leads to increased prescription costs (see also Dana and Loewenstein 2003). Moreover, doctors who attend pharma-sponsored events are more likely to prescribe the sponsor's medication, compared to physicians who do not attend (Bowman and Pearle 1988). Physician bias caused by COIs is especially troublesome, as there is evidence that patients aren't very concerned with their doctors' COIs. As Hampson et al. (2006) find through in-person interviews with patients in cancer-research trials, "more than 90% of patients expressed little or no worry about financial ties that researchers or institutions might have with drug companies." This lack of concern likely manifests in insufficient levels of patient discounting of biased advice, which may literally have life or death consequences in certain situations

As delineated above, there is striking evidence that COIs can induce professionals to provide biased information and advice to advisees. This speaks to the importance of understanding COIs and the mechanisms that can be used to reduce professional bias. Just as research on COIs crosses disciplines, it also spans research methodologies. While I have already discussed some COI research that has been done in a laboratory setting, there is also related research being done in the field. Below, I detail existing literature on field experiments related to COIs that specifically stem from selling credence goods.

1.2.3 Credence Goods

Credence goods and services have the property that, although customers can observe the utility they obtain from the good ex post, they cannot judge whether the type of good they received is the ex-ante needed one (Darby and Karni 1973; Dulleck et al. 2011). For instance, when a customer takes her car into the mechanic for a check-up, the mechanic might recommend she replace the car's brakes. After purchasing the new brakes, the customer doesn't know for sure whether her car would function just as well with the old brakes, or if the new brakes were necessary. As in much of the COI lab experiment literature discussed above, sellers of credence goods (e.g., mechanics, physicians) are better informed than buyers. In the case of credence goods, this information asymmetry is with respect to the product or service the customer needs, and it may give rise to market inefficiencies if advisors take advantage of their better information to reap personal rewards.

Several theoretical papers have derived the conditions under which efficient and honest provisions of credence goods are obtained. Darby and Karni (1973) evaluate how market factors such as the presence of idle capacities and regulation along with reputational concerns affect the equilibrium amount of fraud in the market. Further, Pesendorfer and Wolinsky (2003) examine how customer search for multiple professional opinions and reputational concerns affects market organization, while Sülzle and Wambach (2005) study how the degree of insurance in a market impacts the amount of fraud.

The existing field research on credence goods tests the theory in various settings. To begin, two related studies look at the market for taxi services in Athens, Greece. In this context, taxi drivers are better informed than many passengers on the best route to take, as well as what fees can be applied to taxi fares. Interestingly, Balafoutas et al. (2013) find that taxi drivers take advantage of this information asymmetry by taking foreign (uninformed) riders on longer routes than they do local passengers. Further, Balafoutas et al. (2015) find that overcharging is more likely when passengers signal their ride will be expensed. Overall, this suggests that taxi drivers may strategically alter their behavior in order to receive higher fares from customers who have relatively low situational knowledge or low price sensitivity.

Perceived price sensitivity of customers has also been found to influence professional behavior in the context of repair work. Kerschbamer et al. (2016) analyze the market for computer repairs in Austria and find that the repair price was a substantial 80% higher when customers mentioned they have insurance to cover the cost. Further, in a car repair context, Schneider (2012) finds evidence of deception as car mechanics often tell customers that they require more repairs than they really do. Moreover, customers who have the potential to bring in long-term service are charged lower diagnosis fees than one-time customers, signifying that mechanics strategically price discriminate when offering their services. Also looking at how consumer characteristics affect car repair prices, Busse and colleagues (2017) find that the quoted price customers receive for a radiator replacement can be influenced by two factors. First, the authors find that, on average, females are quoted higher prices than males. Second, customers who communicate misinformation about the price of a radiator replacement are quoted higher prices for the repair, compared to those who are well informed or uninformed on the appropriate repair price. This provides further evidence that, in situations of information asymmetry, professionals may bias their information (and in some cases, prices) for personal rewards. This bias, in turn, is detrimental for consumers, who often end up paying more for goods and services than is necessary.

Consumer knowledge is also found to influence professional bias levels in a healthcare context. Whereas Busse et al. (2017) find that giving inaccurate price information to professionals may increase their bias, Currie et al. (2011) find that professional bias may be mitigated when consumers signal accurate situational knowledge. In particular, the researchers examine antibiotic prescriptions in China, a country where patients are prescribed nearly 1.5 antibiotics per doctor visit, on average (Dong et al. 1999). In their experiment, medical "patients" (who were study confederates, none were actually sick) either signaled accurate knowledge of appropriate antibiotic use to the physician, or not. The results show that patients who signaled knowledge were prescribed antibiotics 39% of the time, which is significantly lower than the 64% of patients who were prescribed antibiotics when they did not signal knowledge. Further, prescribed drug costs were lower by nearly 26% for those who signaled accurate information. A different field experiment on Chinese healthcare also examines physicians' prescribed drug costs. In particular, Lu (2014) finds that when physicians have a COI (defined here as receiving a proportion of patients' drug expenditures), patients are prescribed more and pricier drugs, and this effect is stronger when the patient is insured.

The existing field experimental research provides evidence that COIs can lead to market inefficiencies across industries. However, a large focus of the extant literature is on the effect of consumer price sensitivity. While there is a small stream of research evaluating how customer information affects professional deception and consumer outcomes (Busse et al. 2017; Balafoutas et al. 2015; Currie et al. 2011), I believe there is more to be explored in this area. This is especially true given the large amount of information available to consumers today, some accurate and some not. Thus, in Study 3, I examine how customer misinformation, both when coupled with trust signals and when not, affects professional bias and customer outcomes.

1.2.4 Agency Theory

Theory. The advisor-advisee interaction that I am interested in closely resembles an agency relationship. Agency relationships are classified by situations where one party (the agent) performs an action on behalf of another party (the principal) (Jensen and Meckling 1976; also see Bergen et al. 1992 and Eisenhardt 1989). In marketing, for instance, we observe agency relationships between managers and employees (Basu et al. 1985), as well as between firms and channel partners (Lal 1990). Agency relationships suffer from the issues of goal divergence and information asymmetry, and it is difficult for the principal to perfectly monitor the agent's actions. This framework does well to capture the issues that arise in COI situations. For instance, one can think of the relationships between consumers who rely on mechanics to fix their vehicle, and investors who count on financial agents to manage their investment funds. In both of these situations, the customers delegate some work to the professionals, who may have personal incentives to cheat the customers. However, because of information asymmetry, customers are often not able to say for sure whether they've been cheated or not. In Study 2, I examine whether customers should delegate their decision-making to advisors, who then act as agents on the customers' behalf, or instead retain decision power but utilize agent expertise in an information sharing capacity only.

Application to PSFs. One type of agency relationship that is related to my research context is that between Professional Service Firms ("PSFs") and their employees. Here, a PSF ("principal") employs a fleet of professional service workers ("agents") to interact with customers on the firm's behalf (Morris and Empson 1998; Greenwood et al. 2005). Common examples of PSFs include accounting and financial consulting firms, and healthcare entities such as hospitals. As discussed above, professionals in these industries often have personal incentives to act opportunistically. When employed by a PSF, however, professional dishonesty can not only damage a professional's own reputation but also that of the PSF he works for. This is a concern for firms, as existing research has found that reputational concerns can have large effects on long-term customer retention and earnings (Shapiro 1982; Herbig and Milewicz 1993). Indeed, Greenwood et al. (2005) show that, because clients are dependent on experts in situations of information asymmetry, this elevates the importance of reputation on firm performance (see also Rao et al. 2001). Given this, PSFs have a large incentive to develop training plans and policies that induce professionals to be honest, even in the face of COIs. My work is relevant as it offers a cost-free mechanism, namely, trust requests, that PSFs can introduce into their employee training programs to help restore honest professional behavior

1.2.5 Behavioral Economics & Trust

Finally, my work is also related to the behavioral economics literature (Kahneman and Tversky 1979; Thaler 1980; Ho et al. 2006), specifically that pertaining to trust behavior (see Evans and Krueger 2009 for an overview of the literature). Trust is interesting to study in the context of professional COIs, as its effect is unclear: while trust could improve outcomes if it is reciprocated with trustworthiness, it could also have a deleterious effect on outcomes if trusting advisees are perceived as being easier to take advantage of. Arrow (1974) refers to trust as an "important lubricant of a social system". And indeed, extant literature has shown that subjects trust one another, even in one-shot and anonymous settings (Ho and Weigelt 2005; Berg et al. 1995; Özer et al. 2011). However, researchers have also found that this trust breaks down towards the end of experimental games, which provides evidence of strategic trust (Ho and Weigelt 2005). That is, advisees may strategically signal trust in expectation that advisors will reciprocate this trust with more honest behavior (Berg et al. 1995; McCabe et al. 2003; Rousseau et al. 1998).

The present work builds on this literature in three ways. First, in Studies 1 and 2, advisees can choose to send (dis)trust signals to their advisors, who then send estimates of their private information to the advisees. This setup is different from much of the past literature on trust games, in that I am interested in trust signals rather than actual levels of trust. That is, in my studies, it is common knowledge that advisee signals are cheap talk and need not reflect the advisee's actual beliefs. Indeed, I do find that advisees strategically send trust signals to advisors, even when they don't actually believe the advisors will be trustworthy. Second, while cheap talk trust signals (i.e., statements of trust) are interesting to study, there are also other kinds of trust signals and some, such as decision delegation, may be perceived as stronger, more credible signals of trust. I evaluate how different types of trust signals affect outcomes in Study 2. Third, I also examine how trust requests from professionals to advisees can affect outcomes. Past research has found that salesperson reciprocity towards managers increases when requests for price delegation are allowed (Lim and Ham 2014; see also Fehr et al. 1997; 2007). While further research has found conceptually similar results when a firm requests that a third-party reviewer reports in its favor (Ham et al. 2019). Thus, it's possible that allowing advisors to ask for advisee trust may improve advisor honesty levels.

Above, I detailed how, why and when COIs culminate in biased professional information and advice across various literatures. The effects of this bias on consumer outcomes can potentially be devastating, as COIs permeate high-risk industries such as healthcare and finance. While there is a large body of work already done on COIs and credence goods, there is still much left to be explored. The rest of this paper is intended to address a number of the research gaps discussed above, while offering actionable recommendations for consumers, marketing practitioners and firms.

The paper proceeds as follows. In Section 2, I present a lab experiment on customer trust and distrust signals (Study 1). In Section 3, I delineate a lab experiment evaluating different types of customer trust signals as well as advisor trust requests (Study 2). A field examination of customer misinformation and trust signals (Study 3) is presented in Section 4, while Section 5 outlines a lab experiment on customer information sharing (Study 4). Section 6 discusses and concludes.

2. Customer Trust and Distrust Signals (Study 1)

My first experiment evaluates how customer signals of trust and distrust affect advisor honesty levels and customer outcomes in a COI context. In Study 1, I focus on simple, cheap talk messages of (dis)trust to see whether these signals can influence outcomes. My objective in Study 1 is to offer an introduction to advisee (dis)trust signals, and to gain insights into which type of signals – trust or distrust – advisees are more likely to send. Further, I am interested in not only the type of signal customers choose to send, but also whether their chosen signal aligns with their actual beliefs. In this way, I can evaluate whether customers act strategically in COI situations, and, if so, whether their strategy can be improved.

2.1 Theory & Equilibrium Prediction

Consider a better-informed advisor and lesser-informed advisee who interact in an advice-giving setting. The advisee makes an investment decision before the value of the investment, v, is realized. The investment value is given by $v = \mu + \lambda + \epsilon$, where $\mu > 0$ is a constant that denotes the average investment value and ϵ is market uncertainty. Both parties know μ, and they also know that ϵ is a zero-mean random variable uniformly distributed over the interval $[-\epsilon, \epsilon]$. The term λ represents the advisor's private knowledge of the value of the investment, and it is deterministically known to the advisor. One can think of λ as an advisor's expertise, which can be gained through training or experience. It is common knowledge that λ is a zero-mean random variable uniformly distributed over the interval $[-\lambda, \lambda]$. The advisor sends an estimate of λ, denoted $\hat{\lambda}$, to the advisee. Before the advisor sends $\hat{\lambda}$, however, the advisee signals whether she thinks the advisor's estimate will be truthful or not. The specific sequence of play is as follows.

First, the advisee sends the advisor a signal, s, which signifies how different, if at all, the advisee thinks $\hat{\lambda}$ will be from λ, where $s \in [-(2 \cdot \max(\lambda) + 1), 2 \cdot \max(\lambda) + 1]$. In particular, the advisee signals whether she thinks $\hat{\lambda}$ will be equal to, higher than, or lower than λ, and these signals equate to the values $s = 0, s > 0$ and $s < 0$, respectively. Moreover, if the advisee chooses to signal that $\hat{\lambda}$ will differ from λ, then the advisee also signals how many points different from λ that $\hat{\lambda}$ will be (i.e., signals $|\hat{\lambda} - \lambda|$). I refer to this difference as the "level" of the advisee's signal and use it to determine how high or low the value of s is. For instance, if the advisee signals that the advisor's estimate will be higher than the truth by, say, 5 points, then $s = 5$. Alternatively, if the advisee signals the advisor's estimate will be 3 points lower than the truth, then $s = -3$. Note that $s = 0$ indicates that the advisee believes the advisor's estimate will be equal to the true value of λ, and thus this can be thought of as a trust signal from the advisee to the advisor. On the other hand, $s \neq 0$ indicates that $\hat{\lambda}$ will differ from λ, and can therefore be thought of as a distrust signal from the advisee to the advisor. It is common knowledge that s is cheap talk, and as such carries no monetary consequences and need not reflect the advisee's actual beliefs.

Next, the advisor observes λ and s and then sends $\hat{\lambda}$ to the advisee. It is common knowledge that $\hat{\lambda}$ is cheap talk and can be any integer between $[-\lambda, \lambda]$. Then, the advisee is endowed with $\bar{i}$ experimental units and decides how many of these units to invest, $i \in [0, \bar{i}]$. The advisee makes her investment decision after observing $\hat{\lambda}$. The advisee's investment earnings are based on the realized value of v, and are equal to vi. The advisee also incurs a cost of investing, ci^2 where $c > 0$. One can think of c as the advisee's opportunity cost of investing. Further, the advisor earns ki, where $k > 0$ can be thought of as the commission or kickback that the advisor earns from the advisee's investment. In this way, the advisor's incentive in the game is to induce the advisee to invest as many units as possible (i.e., the advisor has a COI). Finally, both the advisor and advisee receive a fixed payment, $F > 0$. For the advisor, F represents a fixed fee for giving advice; for the advisee, F serves as a buffer in case the advisee's investment earnings less cost of investing is negative.

The advisor's and advisee's expected profits are shown below.

$$\pi_{Advisor} = F + kE[i] \tag{1}$$

$$\pi_{Advisee} = F + E[v]i - ci^2 \tag{2}$$

By taking the first order condition of equation 2 with respect to i, we observe that the advisee maximizes profits by investing

$$i = \frac{E[v]}{2c}. \tag{3}$$

However, in this setting, the advisor has an incentive to inflate $\hat{\lambda}$, since his profit is increasing in i and he does not incur a cost for sending an estimate that differs from λ. Knowing this, the advisee would not consider the advisor's estimate to be credible. Since the advisor anticipates this, he does not have an incentive to send any one estimate in particular. Moreover, since the advisor knows the advisee's signal is cheap talk, he will disregard s, and anticipating this, the advisee will choose to send any s. Thus, in equilibrium, the advisee selects any s, the advisor chooses any $\hat{\lambda}$ and the advisee does not update her belief about λ given $\hat{\lambda}$, resulting in $i^* = \frac{\mu}{2c}$.

Of note, regardless of whether the advisor observes s or not before sending $\hat{\lambda}$, the equilibrium prediction remains the same.

The advisor's COI in this game arises through his payoff function, $\pi_{Advisor} = F + ki$. Specifically, the advisor faces two competing interests: (1) helping the advisee to make an optimal investment decision for herself and (2) potentially personally benefitting from encouraging the advisee to choose a nonoptimal, relatively high investment amount, which directly increases the advisor's payoff.[7] Put differently, the advisor's payoff structure encourages him to send high estimates to the advisee, regardless of the true value of λ.

2.2 Behavioral Prediction

The equilibrium prediction outlined in Section 2.1 does not account for behavioral factors, such as trust and reciprocity, that might influence advisor and advisee decisions in the game. First, past research has shown that advisors believe their (cheap talk) estimates will affect advisee decisions, and consequently choose an estimate that is increasing in λ (e.g., Özer et al. 2011, 2018). Further, it is possible that, after receiving $s = 0$, the advisor might feel inclined to reciprocate the advisee's signal of trust by being honest, or at least reducing the level of bias in his estimate. Anticipating this, advisees may strategically send trust signals (i.e., $s = 0$) to induce advisors to be more honest, even if advisees actually believe advisors will send inflated estimates (Ho and Weigelt 2005). Alternatively, it is possible that sending $s > 0$ may be beneficial for advisees by the following reasoning: if the advisee sends $s > 0$, then the advisor might believe

[7] One can think of the advisor's COI as being similar to when a physician is incentivized (using money, vacations, "swag", etc) by pharmaceutical companies to prescribe their product, or when a mechanic can earn a higher payment by recommending more costly than necessary vehicle repairs.

that the advisee is wary of his advice and will monitor the advisor's estimate for signs of bias (Cain et al. 2011). In turn, the advisor might reduce his bias level. Thus, anticipating this, it may be optimal for advisees to send $s > 0$ instead of $s = 0$. Furthermore, as the advisor may experience disutility if his payoff is too different from the advisee's, bias levels may be lower than predicted (Fehr and Schmidt 1999). In this way, the equilibrium prediction that the advisor chooses any estimate and the advisee chooses any signal may not hold.

2.3 Experimental Manipulation

I manipulate whether the advisee sends a signal to the advisor or not. Specifically, I vary whether the advisee's signaling decision (i.e., to signal or not) is exogenously set by the experimenter (Treatments 1 and 2; "NO CHOICE" conditions), or whether the advisee can choose whether to send a signal to the advisor (Treatment 3; "CHOICE" condition). In the NO CHOICE conditions, the advisee either sends a signal to her advisor (Treatment 2: NO CHOICE/SIGNAL) or not (Treatment 1: NO CHOICE/NO SIGNAL). In the CHOICE condition, the advisee either chooses to send a signal to her advisor (Treatment 3b: CHOICE/SIGNAL) or chooses not to (Treatment 3a: CHOICE/NO SIGNAL). Recall that the advisee's signal will indicate whether she thinks her advisor will send an honest estimate or not, although it is common knowledge that the signal need not reflect the advisee's actual beliefs. Table 1 provides an overview of the experimental manipulations, while Figure 1 displays a decision tree for Study 1. The rationale behind varying whether an advisee's signaling decision is fixed or a choice (even though conventional economic wisdom predicts that doing so will not affect the results) is that, in the NO CHOICE conditions, I can analyze which signal, if any, is optimal for the advisee to send. Then in the CHOICE conditions, I can observe whether advisees can *anticipate* the optimal signaling decision (i.e., to send a signal or not).

Is Advisee's Cheap Talk Signal Sent to Advisor?

No	**Yes**	**Choice** (Advise chooses whether to send signal to Advisor)	
T1: Advisee does not send signal to Advisor	**T2:** Advisee sends signal to Advisor	**T3a:** Advisee chooses not to send signal to Advisor	**T3b:** Advisee chooses to send signal to Advisor

Table 1: Overview of Study 1's experimental manipulation

Figure 1: Study 1 Decision Tree[†]

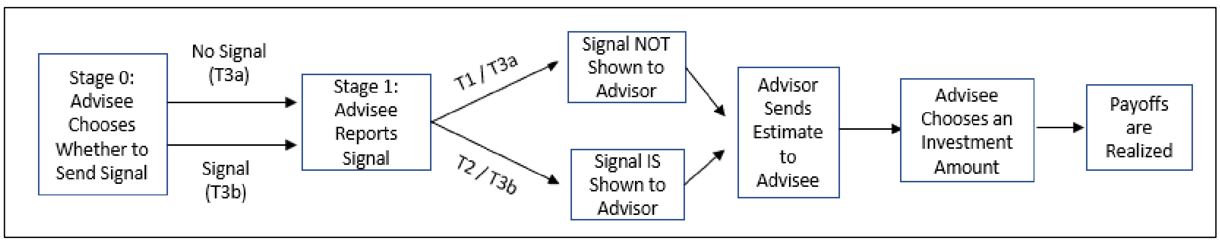

[†]Only advisees in the CHOICE condition enter Stage 0; advisees in the NO CHOICE conditions start at Stage 1.

2.4 Contribution

Study 1 contributes to the literature on COIs in the following ways. First, the study evaluates the effect of consumer (dis)trust signals on advisor bias and advisee outcomes. This has not yet been examined thoroughly in the literature and may raise important insights on how advisees should interact with advisors across a variety of COI contexts. Second, I measure advisee beliefs of advisor honesty and analyze the distribution of beliefs in the data. By comparing beliefs to chosen signals, I can understand whether advisees act strategically in COI situations. Third, I provide an investigation of how trustworthiness is influenced by advisee (dis)trust signals in a setup that has not yet been evaluated in the literature on advisor-advisee interactions.

2.5 Experimental Setup & Procedure

Participants were primarily undergraduate students at the University of Wisconsin-Madison. I ran 2 sessions each of T1 and T2, and 4 sessions of T3, in which participants would endogenously decide whether to send a signal (T3b) or not (T3a). A total of 86 subjects participated in the experiment. Between 10 and 14 subjects participated each session, and each session had a total of 10 decision rounds. Before the start of the 10 decision rounds, there were two practice rounds that carried no monetary consequences to familiarize participants with the experimental procedure. At the end of each experimental session, point earnings were converted to cash at a pre-specified rate, and payments were awarded privately. Participants earned $14 on average, with a range of earnings from $9 to $21. The parameter values used in the experiment were $k = 20, c = 1.2, \mu = 25, \bar{i} = 25$, and $F = 50$. Further, the range of λ and $\hat{\lambda}$ was exogenously set at [-20, 20], the range of ϵ was [-5, 5], and $v \in [0, 50]$. Moreover, given these parameters, $s \in [-41, 41]$. The standard economic model predicts, across all conditions, advisees will choose any s, advisors will choose any $\hat{\lambda}$, and advisees will invest $i = \frac{\mu}{2c}$, or 10.4 units. The experiment was coordinated using z-Tree software (Fischbacher 2007).

Below, I first describe the procedure used in the NO CHOICE/SIGNAL condition (T2) in detail, and then use it as a benchmark to explain the differences in the other three treatments.

NO CHOICE - SIGNAL (T2)

In T2, participants entered the lab, were seated at separate computer terminals, and received a set of instructions from the experimenter. At the beginning of the session, the experimenter went over the instructions as a group, and participants were encouraged to ask questions if they had them. Participants were randomly assigned to the role of Player A (the advisor) or Player B (the advisee), and their assignment was fixed for all 10 decision rounds. At the beginning of each round, each Player A was randomly and anonymously matched with a Player B. As everything in the experiment was anonymous, participants were asked to keep their eyes on their own computer and refrain from talking to anyone but the experimenter.

Player B makes the first decision. In particular, Player B decides on a report (Report 1; s) to send to Player A. Report 1 indicates whether Player B believes $Player\ A's\ Decision\ Number$ ($\hat{\lambda}$) will be equal to, higher than or lower than the true value of $Player\ A's\ Private\ Information$ (λ). If s indicates $\hat{\lambda}$ will be different from λ, Player B also indicates how many points different she thinks $Player\ A's\ Decision\ Number$ will be from $Player\ A's\ Private\ Information$ (i.e., $|\hat{\lambda} - \lambda|$). Recall that I refer to this as the "level" of Player B's signal, with more extreme levels indicating an estimate that is further from the truth. It is common knowledge that Player B can report anything for Report 1, and that the report has no monetary consequences and need not reflect her actual beliefs. After completing Report 1, Player B is asked to complete another report, Report 2, completely independent of Report 1. In Report 2, Player B reports on the same factors as in Report 1. However, it is common knowledge that Report 2 is *not* shown to Player A, and that Report 2 should reflect Player B's actual beliefs. By comparing Report 1 to Report 2, I can see whether Player B sent a strategic signal to Player A. Next, the computer randomly draws a value of $Player\ A's\ Private\ Information$ (λ) from the range -20 to 20, with each integer in this range having an equal chance of being chosen. Player A observes the realized value of λ and Player B's Report 1. Player A then chooses $Player\ A's\ Decision\ Number\ (\hat{\lambda}) \in [-20,\ 20]$, which will be observed by Player B. It is common knowledge that $Player\ A's\ Decision\ Number$ need not be equal to $Player\ A's\ Private\ Information$.

Next, Player B is endowed with 25 experimental units. Player B observes $Player\ A's\ Decision\ Number$, and then chooses $Player\ B's\ Decision\ Number, i \in [0, 25]$. It is common knowledge that Player B's payoff is increasing in the value of the $Final\ Number$, v, multiplied by $Player\ B's\ Decision\ Number$, and that Player B also incurs a corresponding $Decision\ Cost$, $1.2i^2$. Further, it is common knowledge that Player A earns the value of k (20) multiplied by $Player\ B's\ Decision\ Number$, where k can be thought of as Player A's commission rate in the game. Moreover, both players receive a fixed payment of 50 points in each decision round. The full set of $Player\ B's\ Decision\ Costs$ was provided to each participant in the format of "Player B's Decision Cost Table". After choosing $Player\ B's\ Decision\ Number$, Player B reports private guesses of λ and v. The former guess allows me to obtain a measure of Player B's discounting of Player A's estimate, while the latter guess is necessary to determine whether Player B chooses an optimal i given her beliefs of v. After Player B reports these guesses, the decision round ended, payoffs were calculated and the players were randomly re-matched with a new counterpart for the next round. Table 2 maps the decisions in the COI signaling game to the parameter values and terminology used in the experiment.

Table 2: Terminology and Parameter Values in the COI Signaling Game and the Experiment (Study 1)

COI Signaling Game	Experiment
$\lambda \in [-20,20]$	Player A's Private Information
$\hat{\lambda} \in [-20,20]$	Player A's Decision Number
$i \in [0,25]$	Player B's Decision Number
$c = 1.2$	Player B's Decision Cost Parameter
$F = 50$	Fixed Payment
$k = 20$	Player A's Decision Number Parameter
$s \in [-41,41]$	Player B's Signal (Report 1)
$v \in [0,50]$	Final Number
$\mu = 25$	Fixed Number
$\epsilon \in [-5,5]$	Random Number

In the experiment, variable distributions, the values of μ and F and payoffs were common knowledge. Participants were told that they would learn the value of the $Final\ Number$ (v), $Player\ B's\ Decision\ Number$ (i) and their own payoffs at the end of every decision round. The value of $Player\ A's\ Private\ Information$ (λ) was only known by Player A, as was the value of the $Random\ Number$ (ϵ), which Player A could back out after learning the value of v.

NO CHOICE – NO SIGNAL (T1)

T1 is identical to T2 except Player B is told that, for Report 1, she should report *as though* the report would be sent to Player A. However, it was common knowledge that Report 1 would not actually be shown to Player A. Thus, at the beginning of Player A's move, Player A was only told the value of λ (but was not told s).

CHOICE (T3a and T3b)

The CHOICE conditions are identical to their corresponding NO CHOICE treatments, except that, before Player B completes Report 1, Player B chooses whether to send Report 1 to Player A or not. If Player B chooses not to send Report 1 to Player A (T3a), then the game reverts to T1 and Player A is told that Player B has chosen not to send a report to him. If Player B chooses to send Report 1 to Player A (T3b), then the game reverts to T2 and Player A is told that Player B has chosen to send a report to him, along with Player B's chosen s.

2.6 Results

Next, I discuss the Study 1 findings. Note that, in T3, Player B chose to send a signal to Player A 73% of the time. That is, of the 210 Player B observations in T3, 73% of these were in T3b while 27% were in T3a.

2.6.1 Player A's Bias & Player B's Investment

To begin, I test the equilibrium prediction that, upon observing λ, Player A chooses any estimate, $\hat{\lambda}$. To evaluate this, I combine the data across treatments and regress $\hat{\lambda}$ on Player A's Private Information (λ). If the equilibrium prediction holds, then there should be no systematic relationship between λ and $\hat{\lambda}$, and the coefficient would not differ from 0. However, results of the OLS regression reveal that the coefficient is positive and significant. Specifically, $\hat{\lambda}$ increases by 0.46 points, on average, for every one-unit increase in λ ($se = 0.034$; $p \sim 0.000$). This provides evidence against the equilibrium prediction that Player A chooses any estimate.

Next, I calculate a measure of Player A's bias by taking $\hat{\lambda} - \lambda$. If the equilibrium prediction holds, then Player A's bias level is expected to be 0 on average. From Table 3, one can see that, across all conditions, Player A's mean bias level is significantly greater than zero (one-sample t-tests: p'$s \sim 0.000$). This speaks to the deleterious effect that COIs can have on professional honesty levels overall. Comparing between the columns for mean and median bias in Table 3, one can see the pattern of bias is fairly consistent between the two measures. One thing to note, however, is that the median bias levels are lower than the mean bias levels, suggesting that the distributions of bias are skewed to the right, with some Player As biasing by a substantial amount. Moreover, Player A's mean bias is directionally lower in T2 compared to T1 ($p = 0.114$) as well as in T3b compared to T3a ($p = 0.142$). Further, I also regress Player A's bias on a dummy variable equaling 1 if Player B sends a signal (T2, T3b) and 0 otherwise (T1, T3a), while also controlling for λ. I find that Player A's bias is, on average, 2.68 points lower when he receives a signal from Player B compared to when he does not ($se = 0.762$; $p \sim 0.000$). Interestingly, this suggests that there is a main effect of Player B's signaling to reduce Player A's bias. I explore this point further below.

Table 3: Player A's Mean and Median Bias Level

	Mean Bias	Median Bias
T1	10.1	8.0
T2	8.0	7.0
T3a	10.4	8.5
T3b	8.1	4.0

Next, I examine Player B's investment decision (Table 4). The equilibrium prediction is that, given Player A's estimate is uninformative, Player B does not update her belief of λ given $\hat{\lambda}$, and thus Player B invests $\frac{\mu}{2c}$. Given the experimental parameters, Player B's optimal investment amount is $\frac{25}{2*1.2}$, or 10.4 units. In T1 and T3b, Player B's mean investment level is not statistically different from 10.4 ($p = 0.513$ and $p = 0.193$, respectively). However, Player B invested significantly less than the equilibrium prediction in both T2 and T3a (both p'$s \sim 0.000$). Overall, the mean and median investment amounts are relatively consistent, suggesting that the distribution of investment amounts is relatively normal.

Table 4: Player B's Mean and Median Investment Level

	Mean Investment	Median Investment
T1	10.0	10.0
T2	8.5	8.5
T3a	7.0	6.0
T3b	9.8	10.0

2.6.2 Player B's Signals & Beliefs

Next, I examine the signals (s, Report 1) chosen by Player B, along with how they relate to Player B's actual beliefs (Report 2). Unsurprisingly, less than 2% of Player Bs reported $s < 0$. It is likely these signals appeared in the data because of subjects' bounded rationality, so I will exclude these observations from the analysis. For $s = 0$ and $s > 0$, I will refer to the signals as TRUST and HIGH signals, respectively. Note that a HIGH signal implies that Player B does not believe Player A will send an honest estimate and can therefore be thought of as a signal of distrust. Also note that in T1 and T3a, s is not observed by Player A. Thus, in the following analysis, the signals discussed in those treatments are the signals that Player B would have sent, if she were to send a signal to Player A.

Figure 2 presents the findings, broken down by condition. Recall that the equilibrium prediction is that Player B will choose any s. The fact that only 2% of advisees chose $s <$

0 provides evidence that this does not hold, since if it did, we would expect one-third of signals to be $s < 0$. I will examine whether the equilibrium prediction holds in more detail below.[8]

Figure 2: Player B's Signals and Beliefs

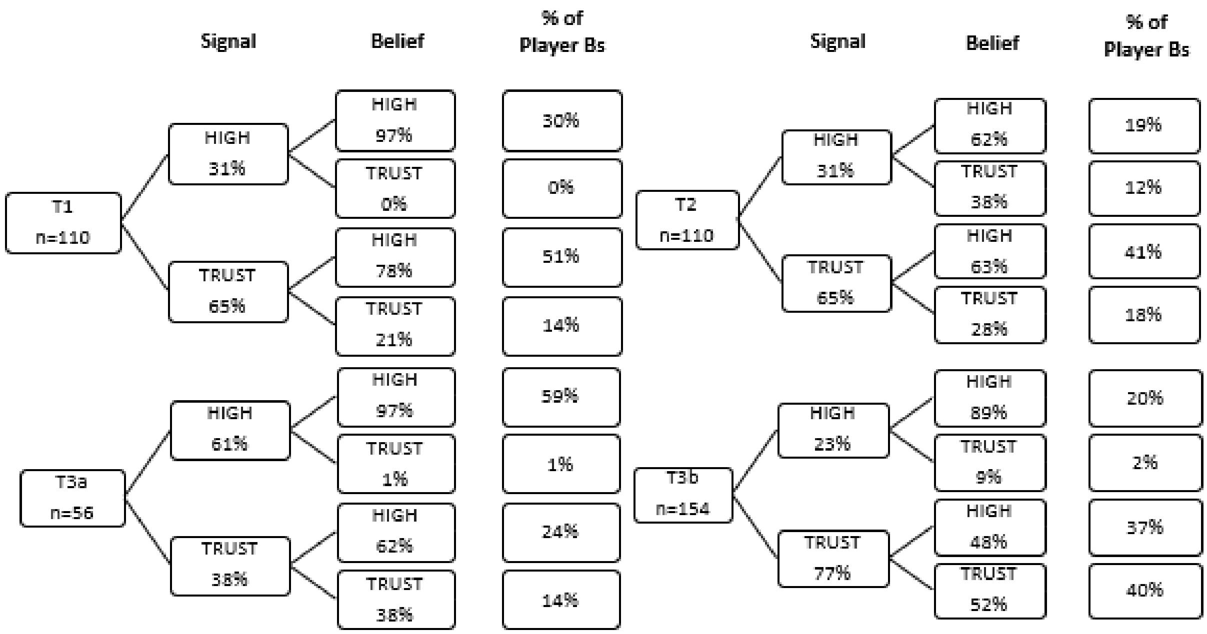

To begin, in T1, T2 and T3b, Player B sends significantly more TRUST signals than HIGH signals (all $p's \sim 0.000$). As the equilibrium prediction is that Player B will choose any signal (and thus the proportion of TRUST and HIGH signals should not differ), this suggests the prediction does not hold. Interestingly, in T1, T2 and T3a, given that Player B sends a TRUST signal, Player B is more likely to believe that Player A's estimate will be dishonest instead of truthful ($p's \sim 0.000$). This suggests that Player B strategically signals trust. Further, in T3b, given that Player B sends a TRUST signal, there is no difference in the proportion of Player Bs who believe Player A will be honest compared to dishonest ($p = 0.604$). This (lack of) finding is driven by the

[8] The p-values in the analysis are from 2-sample tests of proportions with Yates' continuity correction applied.

relatively high proportion of Player Bs who believe Player A will be honest after observing a TRUST signal in T3b. Given this, it appears that, even when advisors face large incentives to be dishonest, advisees do not always believe advisors will act opportunistically. Relatedly, when I compare the proportion of Player Bs who believe Player A will be truthful in T1 to the same proportion in T2, and also compare the proportions in T3a and T3b, I find that significantly more Player Bs believe Player A will be truthful when Player A can observe s (in T2 and T3b), compared to when Player A does not observe s (T1 vs. T2: $p = 0.011$; T3a vs. T3b: $p = 0.001$). This can be interpreted as Player B anticipating that her signal will influence Player A to be more honest.

2.6.3 Regression Analysis Looking at the Effect of Different Signals

Next, I delineate a series of regressions used to evaluate the effect of advisee signaling in more detail. I begin with an examination of Player A's bias level. In the analysis, I divide the data into two parts. The first part includes data from T1 and T2, and the second part includes data from T3a and T3b. In this way, I can cleanly evaluate the effect of Player A's bias given whether Player B's signaling is obligatory or voluntary.

Recall that if Player B reports $s > 0$, she is asked to report the *level* of signal (i.e., $|\hat{\lambda} - \lambda|$). For ease of interpretation, I classify the HIGH signals in T2 and T3b into two groups, based on signal level. The first group consists of the HIGH signals with a relatively low level, and the second group consists of the relatively high-level HIGH signals. To determine the cutoff, I looked at the distribution of HIGH signals in the data (Graph 1, Study 1 Appendix). Between T2 and T3b, there were 39 HIGH signals with a level below 10, and 33 with a level of 10 or higher. Since 10 was the modal level of HIGH signal observed, this provides a natural cutoff level. I classified the HIGH signals with a level less than 10 to the "HIGH_Small" signal group, and those with a level of 10 or more to the "HIGH_Big" group. The following regression analysis will treat these two groups as separate signals. Of course, by separating the observations into two groups, I do lose some explanatory power. However, the alternative is to control for the level of HIGH signal as a continuous variable. In this case, the results become more complex and it is difficult to interpret the effect of sending various HIGH signals in relation to other signals.[9]

Effect of Player B's Signal on Player A's Bias Level

Below, I examine how different advisee signals affect advisor bias levels. I first evaluate the data from the NO CHOICE conditions (T1 and T2), before investigating the CHOICE treatments (T3a and T3b).

NO CHOICE conditions (T1 and T2). Recall that I calculate Player A's bias by taking $\hat{\lambda} - \lambda$. To capture s, I create 4 dummy variables: "NO Signal", "TRUST", "HIGH_Small" and "HIGH_Big". All observations in T1 are coded as 1 for the NO Signal dummy, and 0 for the other 3 dummy variables. The intuition here is that, since Player A does not observe Player B's signal in T1, it does not matter what s Player B chose in that condition. For the observations in T2, if Player B chose to send a TRUST signal (i.e., $s = 0$), then the TRUST dummy is equal to 1; 0 otherwise. If Player B chose to send a HIGH signal with a level less than 10 (i.e., $0 < s < 10$), the HIGH_Small dummy variable is equal to 1, 0 otherwise. Finally, if Player B sent $s > 9$, the HIGH_Big dummy variable is equal to 1, 0 otherwise. I run a series of 4 OLS regressions using Player A's bias as the dependent variable. My main independent variables are the signaling dummies, and I also control for the

[9] Note that the findings are qualitatively similar across the two specifications.

level of Player A's Private Information (λ). I vary the baseline signal used across regressions, so that all comparisons can be clearly shown in Table 5. The baseline signal for each regression is indicated above the column number in the table.

Table 5: Regressions of Player A's Bias using T1 and T2 Data

	DV = Player A's Bias			
	NO Signal	TRUST	HIGH_Small	HIGH_Big
	(1)	(2)	(3)	(4)
NO Signal		4.113***	6.104***	0.448
		(1.106)	(1.814)	(2.004)
TRUST	-4.113***		1.991	-3.665*
	(1.106)		(1.867)	(2.052)
HIGH_Small	-6.104***	-1.991		-5.656**
	(1.814)	(1.867)		(2.492)
HIGH_Big	-0.448	3.665*	5.656**	
	(2.004)	(2.052)	(2.492)	
λ	-0.588***	-0.588***	-0.588***	-0.588***
	(0.045)	(0.045)	(0.045)	(0.045)
Constant	10.959***	6.846***	4.854***	10.511***
	(0.703)	(0.851)	(1.666)	(1.871)
Observations	212	212	212	212
R^2	0.473	0.473	0.473	0.473
Adjusted R^2	0.463	0.463	0.463	0.463
Residual SE (df = 207)	7.214	7.214	7.214	7.214
F Statistic (df = 4; 207)	46.439***	46.439***	46.439***	46.439***
Note:				*p<0.1; **p<0.05; ***p<0.01

From column 1 in Table 5, observe that, compared to not receiving a signal, Player A's bias is, on average, 4.11 points lower after receiving a TRUST signal and 6.10 points lower after receiving a HIGH_Small signal. The coefficient on the HIGH_Big variable, however, is relatively small (-0.448) and is not significantly different from the NO Signal estimate. That is, while Player

A's bias is reduced substantially when Player B sends a HIGH_Small signal, the effect is not found with a HIGH_Big signal. This suggests advisors respond differently to different levels of advisee distrust. This finding is strengthened by the results in column 4, where Player A's bias is shown to be lower after receiving a TRUST or HIGH_Small signal compared to a HIGH_Big signal. Further, I find no difference in Player A's bias after receiving a TRUST or HIGH_Small signal, suggesting that advisees should signal either trust or low levels of distrust in COI situations.

CHOICE conditions (T3a and T3b). Next, I evaluate the data from T3 using the same analysis. Regression results are shown in Table 6. The pattern of findings is quite consistent with that in the T1/T2 data, even though here, effect sizes and statistical significance are attenuated. One finding to note, however, is the difference in Player A's bias between a HIGH_Small signal and NO Signal (column 1). While not significant ($p = 0.175$), this result again suggests that signaling low levels of distrust may improve advisor honesty. However, as bias levels are not increased when Player B signals trust (modal signal), compared to low levels of distrust, it is difficult to say whether advisees could improve their outcomes by switching from signaling trust to distrust.

Table 6: Regressions of Player A's Bias using T3 Data

	DV = Player A's Bias			
	NO Signal	TRUST	HIGH_Small	HIGH_Big
	(1)	(2)	(3)	(4)
NO Signal		1.360	3.033	0.089
		(1.304)	(2.228)	(2.180)
TRUST	-1.360		1.673	-1.270
	(1.304)		(2.094)	(2.031)
HIGH_Small	-3.033	-1.673		-2.943
	(2.228)	(2.094)		(2.728)
HIGH_Big	-0.089	1.270	2.943	
	(2.180)	(2.031)	(2.728)	
λ	-0.530***	-0.530***	-0.530***	-0.530***
	(0.048)	(0.048)	(0.048)	(0.048)
Constant	10.089***	8.729***	7.056***	10.000***
	(1.073)	(0.739)	(1.954)	(1.896)
Observations	210	210	210	210
R^2	0.379	0.379	0.379	0.379
Adjusted R^2	0.367	0.367	0.367	0.367
Residual SE (df = 205)	8.030	8.030	8.030	8.030
F Statistic (df = 4; 205)	31.315***	31.315***	31.315***	31.315***
Note:		*p<0.1; **p<0.05; ***p<0.01		

Effect of Player B's Signal on Player B's Payoff

Next, I evaluate the effect of Player B's signal on Player B's earnings. As above, I first examine results in the NO CHOICE conditions, before examining the CHOICE condition findings.

NO CHOICE conditions (T1 and T2). I begin by looking at the T1 and T2 data. The regression analysis is similar to that employed above. However, here I use Player B's earnings as the dependent variable, and control for v instead of λ. From Table 7, observe that Player B earned

36.79 more points, on average, after sending a HIGH_Small signal compared to when Player B did not signal. Combined with the above results, one possible explanation is that sending a HIGH_Small signal reduces Player A's bias, which allows Player B to make more optimal investment decisions. I also find that Player B earns an average of 35.55 more points after sending a HIGH_Small signal compared to a TRUST signal. This suggests that it may be better for Player B to send a low-level distrust signal compared to a trust signal in certain contexts. This is particularly interesting because the majority of Player Bs chose to strategically send TRUST signals.

Table 7: Regressions of Player B's Payoff using T1 and T2 Data

	DV = Player B's Payoff			
	NO Signal (1)	TRUST (2)	HIGH_Small (3)	HIGH_Big (4)
NO Signal		-1.235	-36.788^{**}	-13.084
		(10.148)	(16.710)	(18.416)
TRUST	1.235		-35.553^{**}	-11.849
	(10.148)		(17.206)	(18.880)
HIGH_Small	36.788^{**}	35.553^{**}		23.703
	(16.710)	(17.206)		(22.909)
HIGH_Big	13.084	11.849	-23.703	
	(18.416)	(18.880)	(22.909)	
v	10.387^{***}	10.387^{***}	10.387^{***}	10.387^{***}
	(0.395)	(0.395)	(0.395)	(0.395)
Constant	-104.787^{***}	-103.552^{***}	-67.999^{***}	-91.703^{***}
	(12.334)	(12.601)	(17.197)	(19.077)
Observations	212	212	212	212
R^2	0.771	0.771	0.771	0.771
Adjusted R^2	0.767	0.767	0.767	0.767
Residual SE (df = 207)	66.318	66.318	66.318	66.318
F Statistic (df = 4; 207)	174.691^{***}	174.691^{***}	174.691^{***}	174.691^{***}
Note:				$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

CHOICE conditions (T3a and T3b). I run the same regression analysis using the T3 data, and results are shown in Table 8. Interestingly, compared to sending a TRUST signal or no signal at all, sending a HIGH signal, regardless of level, directionally reduces Player B's earnings. In particular, sending a HIGH_Small signal results in a payoff that is 26.01 points lower, on average, compared to sending a TRUST signal ($p = 0.169$). This is surprising, as HIGH_Small signals were found to reduce Player A's bias and improve Player B's earnings in the T1/T2 data. Thus, it seems when advisees are given the choice of whether to signal or not, they may be better off signaling trust compared to low levels of distrust. Given that 77% of advisees who chose to signal in T3 sent a TRUST signal, this suggests that advisees can anticipate their optimal signaling strategy.

Table 8: Regressions of Player B's Payoff using T3 Data

	DV = Player B's Payoff			
	NO Signal	TRUST	HIGH_Small	HIGH_Big
	(1)	(2)	(3)	(4)
NO Signal		-7.774	18.236	12.199
		(11.706)	(20.054)	(19.563)
TRUST	7.774		26.010	19.972
	(11.706)		(18.842)	(18.251)
HIGH_Small	-18.236	-26.010		-6.038
	(20.054)	(18.842)		(24.517)
HIGH_Big	-12.199	-19.972	6.038	
	(19.563)	(18.251)	(24.517)	
v	9.574***	9.574***	9.574***	9.574***
	(0.420)	(0.420)	(0.420)	(0.420)
Constant	-88.388***	-80.614***	-106.624***	-100.587***
	(14.320)	(12.937)	(19.635)	(20.380)
Observations	210	210	210	210
R^2	0.724	0.724	0.724	0.724
Adjusted R^2	0.718	0.718	0.718	0.718
Residual SE (df = 205)	72.165	72.165	72.165	72.165
F Statistic (df = 4; 205)	134.124***	134.124***	134.124***	134.124***
Note:				*p<0.1; **p<0.05; ***p<0.01

2.6 Summary & Discussion

Study 1 serves as an introduction to the topic of advisee (dis)trust signaling in COI situations. Findings reveal that, indeed, COIs can induce advisors to send biased information to advisees. Thus, advisees may want to use signals to try to reduce advisor bias and improve outcomes. In the study, a majority of advisees sent a trust signal, even though advisees often did not actually trust the advisor. This provides evidence that advisees strategically interact with advisors, and suggests that advisees believe signals of trust may have a beneficial effect on their outcome. Interestingly, while I find some evidence that trust signals can reduce advisor bias and increase advisee earnings, I also find competing evidence that low levels of distrust may actually be more beneficial in certain situations. The reason for this discrepancy is not immediately clear, and indeed, the statistical significance necessary to make any robust claims does not exist. However, this finding does speak to the importance of signaling environment when evaluating the effect of advisee (dis)trust signals.

Overall, the Study 1 findings suggest that there are behavioral factors that arise when advisees and advisors interact in COI situations, as parties are found to make decisions that differ from what is predicted by standard economic theory. However, the large amount of noise in the data makes it difficult to draw many concrete conclusions. Thus, to provide a deeper understanding of the behavioral factors at play, I now turn to my next experiment, Study 2. I designed Study 2 to build off the insights generated by Study 1. Specifically, Study 1 demonstrates that the signaling environment (i.e., whether signaling is obligatory or voluntary) matters. Thus, in Study 2, advisee signals are always voluntary, as this is most consistent with reality. Further, I focus on trust signals in Study 2, as the majority of advisees chose to signal trust in Study 1. Moreover, the Study 2 game is designed to reduce experimental noise, which should allow for sharper insights. Finally, Study 1 is intended to evaluate how customer signals can influence outcomes in COI situations. While this is certainly important, I design a more central role for the advisor in Study 2's game, in order to generate insights that are more relevant to marketing practitioners.

3. Customer Trust Types and Expert Trust Requests (Study 2)

Study 1 examines trust and distrust messages, both when customers can choose to send a (dis)trust message and when messaging is obligatory. In Study 2, I build on Study 1 by evaluating different types of customer trust signals. In particular, I examine both a cheap talk trust message ("statement of trust") as in Study 1, as well as a stronger signal of trust, decision delegation ("act of trust"). For instance, when a novice investor visits her financial advisor, she can tell her advisor she trusts him but still make the final investment decision on her own. Alternatively, the investor can show her trust by asking the advisor to make the investment decision on her behalf. As customers are able to signal trust in various ways, it's important to understand how experts respond to different types of signals. Study 2 speaks to the question, "Should customers always signal trust?". Furthermore, I also allow advisors to send a trust request to advisees. That is, advisors can request that advisees delegate their decision to them. In practice, professionals often ask for customers' trust and it is possible that doing so intensifies customer trust as well as increases professionals' trustworthiness. For instance, a financial advisor may explicitly ask a novice investor to "trust him" when he suggests an investment option or could implicitly ask for trust by requesting control of the investor's assets. By allowing advisors to take a more central role in the interaction, I can derive important insights for marketing practitioners. In the sections that follow, I first present the Study 2 setup, design and predictions before discussing the results in Section 3.6.

3.1 Game Setup & Equilibrium Prediction

The game setup used in Study 2 closely mirrors that in Study 1. However, I modified the basic setup to simplify the game and reduce noise. Specifically, in Study 2, I replace the advisee's investment decision with a guess of the investment value (consequently the advisee is not endowed with experimental units). Rather than investing the profit-maximizing number of units, the advisee is instead incentivized to guess the investment value as accurately as possible. Along with simplifying the decision space, this allows the results to better generalize to contexts outside of financial decisions. Additionally, I modify the investment value specification to remove the fixed constant, μ. This edit was made to simplify the experiment and make decisions more cognitively manageable for subjects. I also change the format of advisees' messages by allowing only trust statements to be sent (no distrust). I rationalize this change since the majority of advisees in Study 1 chose to send a trust message. Further, as I only allow trust messages to be sent, there is no signal "level" as there is in Study 1. By making this messaging space modification, I can test the robustness of Study 1's results, as well as simplify the game and reduce cognitive load on subjects.

The exact game setup of Study 2 is as follows. To begin, I use a similar specification for the "investment value" as is used in Study 1.[10] Here, the investment value is given by $v = \lambda + \epsilon$, where ϵ is still market uncertainty and the advisor and advise both know that ϵ is a zero-mean random variable uniformly distributed over the interval $[\underline{\epsilon}, \overline{\epsilon}]$. Further, the term $\lambda \sim U \in [\underline{\lambda}, \overline{\lambda}]$ represents the advisor's private knowledge of the value of the investment, and it is deterministically known to the advisor. Recall that, compared to Study 1, this specification

[10] For consistency sake, I will continue referring to the variable v as the "investment value", even though Study 2 captures a context beyond financial investment decisions.

removes one component (μ) of the investment value, which simplifies subjects' decisions. Given this change, the experimental parameters used (outlined later in the paper) are also slightly modified from Study 1. From here, I derive the game's equilibrium predictions by discussing two unique scenarios that correspond with different treatment conditions in Study 2.

Scenario 1 (statements of trust): This scenario is similar to the CHOICE condition in Study 1. The game begins with the advisee choosing whether to send a cheap talk statement of trust to the advisor. Here, the advisee can choose between sending a trust message or no message at all. After choosing whether to send a statement of trust to the advisor, the advisee privately reports her beliefs of the advisor's honesty, as in Study 1. Next, the advisor observes λ as well as the advisee's trust statement (if any). The advisor then chooses an estimate of λ, $\hat{\lambda}$, to send to the advisee. Finally, the advisee observes the advisor's estimate and submits a guess of the investment value, v. Here, the advisee's guess is denoted $\hat{v}^{Ae}$, and it determines the payoffs of both players. While the advisee is incentivized to guess the investment value as accurately as possible, the advisor's payoff is increasing in the advisee's guess of the investment value. Specifically, the expected payoffs for the advisor and advisee, denoted π_{Ar} and π_{Ae}, respectively, are given by equations (4) and (5).

$$\pi_{Ar} = F^{Ar} + \hat{v}^{Ae} \tag{4}$$

$$\pi_{Ae} = F^{Ae} - E[|\hat{v}^{Ae} - v|]. \tag{5}$$

Note that both the advisor and advisee earn a fixed payment, F^{Ar} and F^{Ae}, respectively. The advisor is given a fixed payment so that, even if λ is relatively low, he need not send an inflated estimate in an effort to secure a positive payment for himself. Further, the advisee's fixed payment is necessary given that her earnings are reduced by guess inaccuracy. The advisor's COI is similar to that in Study 1. That is, since his profit is increasing in $\hat{v}^{Ae}$, the advisor has an incentive to inflate his estimate, $\hat{\lambda}$, to the maximum value $\overline{\lambda}$. Knowing this, the advisee would not consider the advisor's estimate to be credible and does not update her belief of v given $\hat{\lambda}$. Instead, the advisee chooses $\hat{v}^{Ae}$ to maximize profits. From equation (5), the advisee earns a fixed payment but now, the advisee's profits are reduced by the absolute difference between the true investment value, v, and the advisee's guess of this ($\hat{v}^{Ae}$). Given this, the advisee minimizes $|\hat{v}^{Ae} - E[v]|$. To do so, she chooses $\hat{v}^{Ae} = E[v]$. Given $E[v] = \mathrm{E}[\lambda] + \mathrm{E}[\epsilon]$, the advisee chooses $\hat{v}^{Ae} = 0$. Thus, in expectation, the advisee's guess of the investment value will be off by $\frac{1}{4} * (\overline{v} - \underline{v})$. Moreover, since the advisor anticipates that the advisee will ignore his estimate, $\hat{\lambda}$, he does not have an incentive to send any one estimate in particular and thus chooses to send any estimate. Further, since the advisor knows the advisee's statement of trust is cheap talk, he will disregard the message if sent. Anticipating this, the advisee will randomly choose whether to send the trust statement. Thus, in equilibrium, the advisee randomly selects whether to send the statement of trust, the advisor selects any $\hat{\lambda}$ and the advisee guesses $\hat{v}^{Ae} = 0$. Given this, the expected payoff for the advisor is F^{Ar} and the expected payoff for the advisee is $F^{Ae} - \frac{1}{4} * (\overline{v} - \underline{v})$. Note that equilibrium behavior is consistent regardless of whether the advisee can send a trust statement to the advisor or not.

Scenario 2 (acts of trust): Scenario 2 is similar to Scenario 1, except, instead of sending a statement of trust, the advisee can choose whether to have the advisor submit a guess of the investment value on her behalf (i.e., the advisee can delegate her guess to the advisor). In particular, the game begins with the advisee choosing whether to delegate her guess of the investment value to the advisor (advisee also submits private beliefs of the advisor's honesty). If the advisee chooses not to delegate, then the game reverts to that in Scenario 1, where the advisor sends $\hat{\lambda}$ to the advisee, who then submits a guess of the investment value, $\hat{v}^{Ae}$, which determines both players' payoffs. Alternatively, if the advisee does choose to delegate, then the advisor observes λ and submits a guess of the investment value, $\hat{v}^{Ar}$, on the advisee's behalf. This guess then determines payoffs for both players. In particular, players' expected payoffs in Scenario 2 are specified as

$$\pi_{Ar} = F^{Ar} + [1 - I_D](\hat{v}^{Ae}) + I_D(\hat{v}^{Ar}) \tag{6}$$

and

$$\pi_{Ae} = F^{Ae} - [1 - I_D](E[|\hat{v}^{Ae} - v|]) - I_D(E[|\hat{v}^{Ar} - v|]). \tag{7}$$

I_D is an indicator variable equaling 1 when Player B delegates and 0 otherwise. Note that the payoffs in Scenario 1 are nested within the payoffs for Scenario 2. That is, I_D always equals 0 in Scenario 1, while it equals 0 in Scenario 2 when the advisee does not delegate but equals 1 when she does.[11] When the advisee delegates, the advisor faces two competing interests: (1) helping the advisee to submit an accurate guess of v and (2) personally benefitting by submitting a high guess of the investment value on the advisee's behalf. Thus, the advisor's incentive to act opportunistically is very straightforward, and he is predicted to choose $\hat{v}^{Ar}$ equal to $\overline{v}$, as this

[11] As the payoffs in Scenario 2 encompass both scenarios, I refer to the payoffs specified in equations (6) and (7) in the rest of Section 3.

maximizes the advisor's payoffs. Correspondingly, the advisee's expected payoff would be equal to $F^{Ae} - \frac{1}{2} * (\overline{v} - \underline{v})$, which is less than her expected payoff when she doesn't delegate. Given this, in Scenario 2, the advisee would never delegate her decision to the advisor. Rather, the advisee would choose to receive an estimate, $\hat{\lambda}$, from the advisor but submit the payoff-determining guess of v on her own. Thus, the equilibrium outcome in Scenario 2 is consistent with that in Scenario 1. Furthermore, note that the equilibrium prediction is invariant to the introduction of a non-binding trust request message sent from the advisor to the advisee. Specifically, consider the case where the advisor can ask the advisee to delegate her decision to him. The advisee is predicted to ignore this request and the game would revert to that in Scenario 2, where the advisee chooses not to delegate her guess to the advisor.

3.2 Behavioral Predictions

As in Study 1, advisees may strategically send statements of trust if they expect that advisors will reciprocate this trust signal by being more trustworthy. By this same reasoning, it's possible that advisees also choose to delegate their guess of v to the advisor with non-zero probability. That is, if advisees believe that their act of trust (i.e., delegation) will be perceived by the advisor as a relatively strong signal of trust, then it's possible that they also believe the advisor will reciprocate this trust with even stronger levels of trustworthiness, compared to the case where advisees just send trust statements. Given this, it may be beneficial for advisees to delegate, as the advisor could use his superior knowledge of the investment value to submit a more accurate guess of v than the advisee could. That is, advisors may not always choose the equilibrium behavior of submitting $\hat{v}^{Ar} = \overline{v}$ on the advisee's behalf when the advisee delegates. Thus, it is not clear whether delegation would lead to reduced payoffs for the advisee compared to when she chooses not to delegate. Further, past research has shown that price delegation requests from a salesperson to a manager and favorable reporting requests from a firm to a third-party reviewer can intensify the requestor's reciprocity towards the requestee (Lim and Ham 2014; Ham et al. 2019). Thus, it's possible that, given the advisor sends a trust request to the advisee (requesting that the advisee delegate), his feelings of reciprocity towards the advisee are increased and his bias is reduced. Consequently, when the advisor can send a trust request, advisee delegation rates may be higher and/or the level of advisor bias conditional on delegation may be lower.

In summary, insights from existing behavioral economics research suggest that advisees may strategically send trust signals. Furthermore, these trust signals may reduce advisor bias, compared to what is predicted by standard economic theory. I conjecture that advisor reciprocity towards advisees may be increased with statements of trust, compared to no trust statements. Further, it's possible that this reciprocity is increased even more when advisees send stronger signals of trust via decision delegation. Finally, an even greater intensity of reciprocity could be achieved when the advisor can request that the advisee show her trust by delegating to the advisor. That is, the order of advisor reciprocity towards the advisee could be given as follows: No Trust Statements < Trust Statements < Delegation < Delegation + Trust Requests. However, past research has also shown that, even in the face of reciprocal preferences towards another party, one does not completely ignore his own incentives. Even though the advisor may feel increased reciprocity towards the advisee when the advisee delegates compared to when she just sends a statement of trust, the advisor also faces a decreased risk of bias in the former scenario. Consider the following: when there is no delegation and the advisor sends the advisee a relatively high estimate, $\hat{\lambda}$, the advisee may believe the likelihood of advisor dishonesty is higher than if the advisor's estimate was relatively low. Consequently, the advisee may discount higher estimates at greater levels. Given this, it may be risky for the advisor to choose a relatively high estimate, as the advisee's discounting may increase exponentially. However, when the advisee instead delegates to the advisor, he does not face a similar risk of submitting a high guess of v on the advisee's behalf. This is because the advisor's guess of the investment value (submitted on behalf of the advisee) determines both players' payoffs and the advisee is not able to change the advisor's guess once submitted. Thus, while delegation may lead to increased reciprocal preferences from the advisor to the advisee, it may also lead to a decreased risk of being biased for the advisor. Study 2 examines how the interplay of these motivations manifests with regards to advisor and advisee outcomes in a COI setting.

3.3 Experimental Manipulation

There are four conditions in Study 2. I manipulate both the type of advisee trust signal (statement versus act), as well as whether the advisor can send a trust request to the advisee. I present an overview of the conditions in Table 9. In Treatment 1 (T1), "NO SIGNAL", the advisee cannot signal any trust to the advisor. However, the advisee does privately respond yes or no the following belief measure: "I believe that [the advisor's] estimate will be truthful." This belief measure is collected in all treatments. The NO SIGNAL condition serves as a benchmark to evaluate how different trust signals can influence advisor bias and player earnings. In Treatment 2 (T2; "STATEMENT"), which is outlined in Scenario 1 above, the advisee can send a cheap talk statement of trust to the advisor. In particular, the advisee chooses whether to send the following message to the advisor: "I trust you to send me an accurate estimate of [λ]".[12] Results of the STATEMENT condition allow me to test the robustness of the Study 1 findings, as well as evaluate how a relatively weak signal of trust (i.e., cheap talk trust statements) affects decisions and outcomes. In Treatment 3 (T3; "ACT"), which is outlined in Scenario 2, I examine a stronger and more credible signal of trust, namely decision delegation. The rationale behind the ACT condition is that there are different ways for a customer to signal trust, and while some trust signals may help the customer, it's possible that other trust signals may do more harm than good. Thus, by comparing T2 to T3, I can examine what types of trust signals can help customers in COI situations. Finally, in Treatment 3b (T3b; "REQUEST"), the advisor can send an upfront trust request to the advisee, who then decides whether to delegate as in T3. The exact wording of the request is, "If you delegate your guess to me, I will help you out." It is common knowledge that this request is non-binding. Comparing between T3 and T3b will provide insights regarding in which situations customers should send strong signals of trust (e.g., via decision delegation) to an expert who has a COI.

Trust Environment

NO SIGNAL	**STATEMENT**	**ACT**	**REQUEST**
T1: No Trust Signal	T2: Statement of Trust Sent from Advisee to Advisor	T3: Delegation of Guess from Advisee to Advisor	T3b: Request for Delegation Sent from Advisor to Advisee

Table 9: Overview of Study 2's experimental manipulation

[12] The exact message will be outlined in Section 3.5. I replace [λ] with the name of the variable in the experiment.

3.4 Contribution & Relation to Existing Literature

In this section, I present the existing literature that is closest to Study 2 and discuss how the present research differs from what's been done in the past. First, the work relates to the behavioral economics research on trust and reciprocal preferences (e.g., McCabe et al. 2003; Ho and Weigelt 2005; Berg et al. 1995), as well as the research on advisee-advisor interactions in COI situations (e.g., Cain et al. 2005, 2011; Sah et al. 2013; Church and Kuang 2009). Study 2 builds on this literature by evaluating how different types of trust signals as well as trust requests affect trust and reciprocity levels between a better-informed advisor and lesser-informed advisee. Further, Study 2 is closely related to work in the operations literature that analyzes how information sharing in a supply channel can affect channel coordination and efficiency (Özer et al. 2011, 2018). Specifically, Özer et al. (2018) examine a context where a manufacturer has an incentive to inflate the demand forecast sent to the supplier. The authors present an examination of how different types of assistance processes (information sharing, advice provision and delegation) affect levels of trust and trustworthiness found in the channel. My research on advisee-advisor interactions differs from this work because (1) my work is related to individual outcomes rather than channel efficiency and (2) my primary interest is not to evaluate how actual levels of trust change in different situations, but rather how optimal trust *signals* change. Indeed, as I find that customers often strategically signal trust even when they don't actually trust, it's important to understand both advisees' chosen behavior and actual beliefs. Study 2 is also related to past research on price delegation from sales managers to salespeople (e.g., Lim and Ham 2014; Joseph 2001; Bhardwaj 2001). For instance, Lim and Ham (2014) find that price delegation occurs frequently, even when not predicted by standard economic theory, and salespeople reciprocate delegation by acting in the interest of the manager. While interesting, this literature is silent on when, if ever, customers should use decision delegation when interacting with professionals who have COIs. Relatedly, Lim and Ham (2014) also find that reciprocity is intensified when the salesperson can request that the manager delegate the pricing decision to her. A similar outcome has been demonstrated in the context of a firm requesting for favorable reviews from a third-party reviewer (Ham et al. 2019). Study 2 extends this work by evaluating decision delegation requests from a better-informed advisor to a lesser-informed advisee in a COI situation.

3.5 Experimental Setup & Procedure

One-hundred twenty-two subjects, primarily undergraduate students at the National University of Singapore (NUS), participated in the experiment. I ran 2 sessions of each condition, and between 12 and 20 subjects participated each session. The experiment had 10 decision rounds, and, before the start of the 10 decision rounds, there were two practice rounds that carried no monetary consequences to familiarize participants with the experimental procedure. At the start of the first round, subjects were assigned to be either the advisor (“Player A”) or the advisee (“Player B”), and these roles were fixed for all rounds. Fixed roles add realism to this COI setting, as advisors and advisees typically maintain their position as either an advice-giver or advice-receiver. At the end of each experimental session, point earnings were converted to cash earnings at a pre-specified conversion rate, and participants were paid privately in cash. Participants earned $14 on average, with a range of earnings from $10 to $19. The parameter values used in the experiment were $F^{Ar} = 80$ and $F^{Ae} = 100$. Further, variable ranges were set at $v \in [-50, 50]$, $\lambda \in [-40, 40]$ and $\epsilon \in [-10, 10]$. The standard economic model predicts advisees will randomly choose whether to send a statement of trust (T2) and will never delegate (T3 and T3b). Similarly, advisors will randomly choose whether to send a trust request (T3b). Further, advisors will select any $\hat{\lambda}$, and advisees will submit an investment value guess of 0. Thus, the equilibrium payoffs are 80 for Player A and 75 for Player B, across all conditions. The experiment was coordinated using z-Tree software (Fischbacher 2007). Below, I first describe the procedure used in T3b (where Player A can send a trust request), and then use it as a benchmark to explain the differences in the other three conditions.

REQUEST condition (T3b)

In T3b, participants entered the lab, were seated at separate computer terminals, and received a set of instructions from the experimenter. At the beginning of the session, the experimenter went over the instructions as a group, and participants were encouraged to ask questions if they had them. Participants were randomly assigned to the role of Player A (the advisor) or Player B (the advisee), and their assignment was fixed for all 10 decision rounds. At the beginning of each round, Player A was randomly and anonymously matched with a Player B. The specification of v, parameter ranges and payoffs were common knowledge. As everything in the experiment was anonymous, participants were asked to keep their eyes on their own computer and refrain from talking to anyone but the experimenter.

Player A moves first by deciding whether to send a trust request to Player B. Player A's request asks Player B to delegate her guess to Player A and, in return, Player A will "help her out". It is common knowledge that this request is non-binding. Next, Player B observes Player A's request (if any) and decides whether to allow Player A to submit a guess of the $Final\ Number\ (v)$ on her behalf. In this case, Player A's guess would determine payoffs for both players. After making her delegation decision, Player B privately reports whether she believes Player A will be truthful or not. If Player B does delegate her decision to Player A, then upon observing the random draw of $Player\ A's\ Private\ Information\ (\lambda)$, Player A chooses a $Guess\ of\ the\ Final\ Number\ (\hat{v}^{Ar})$ to submit on Player B's behalf. It is common knowledge that Player A can submit any guess within the range of possible values for v. Player A is also asked to privately report his $Best\ Guess\ of\ the\ Final\ Number\ (\hat{v}_{BG}^{Ar})$. Consequently, I can create a measure of Player A's bias, given Player B delegated, by taking the difference between $\hat{v}^{Ar}$ and $\hat{v}_{BG}^{Ar}$. Note that neither Player A nor Player B observe the random draw of ϵ ("$Random\ Number$") during the round. Once Player A submits his decisions, payoffs are automatically calculated based on $\hat{v}^{Ar}$ and the decision round ends. Before starting the subsequent round, subjects observe the $Final\ Number$, along with $Player\ A's\ \ Guess\ of\ the\ \ Final\ Number\ (\hat{v}^{Ar})$ and the point earnings for both players.

If, instead, Player B chooses not to delegate, then she will also report her beliefs as above. Next, upon observing the random draw of λ, Player A chooses $Player\ A's\ Estimate\ (\hat{\lambda})$ to send to Player B. It is common knowledge that $Player\ A's\ Estimate$ can be any number within the range of possible values for λ, and need not reflect the truth. Player A also submits $\hat{v}_{BG}^{Ar}$ for the sake of consistency. Upon observing $\hat{\lambda}$, Player B reports $Player\ B's\ Best\ Guess\ of\ Player\ A's\ Private\ Information\ (\hat{\lambda}_{BG}^{Ae})$. By comparing $\hat{\lambda}_{BG}^{Ae}$ to $\hat{\lambda}$, I can obtain a clean measure of advisee discounting. Next, Player B submits $Player\ B's\ Guess\ of\ the\ Final\ Number\ (\hat{v}^{Ae})$, which determines point earnings for both players. Payoffs are automatically calculated and the decision round ends. Again, subjects observe the $Final\ Number$ along with $Player\ B's\ Guess\ of\ the\ Final\ Number$ and both players' payoffs. See Table 10 for a mapping of the COI game decisions to the parameter values and terminology used in the experiment.

Table 10: Terminology and Parameter Values in the COI Game and the Experiment

COI Game	Experiment
$\lambda \in [-40,40]$	$Player\ A's\ Private\ Information$
$\epsilon \in [-10,10]$	$Random\ Number$
$v \in [-50,50]$	$Final\ Number$
$\hat{\lambda} \in [-40,40]$	$Player\ A's\ Estimate$
$\hat{v}_{BG}^{Ar} \in [-50,50]$	$Player\ A's\ Best\ Guess\ of\ the\ Final\ Number$
$\hat{v}^{Ar} \in [-50,50]$	$Player\ A's\ Guess\ of\ the\ Final\ Number\ Made\ on\ Behalf\ of\ Player\ B$
$\hat{v}^{Ae} \in [-50,50]$	$Player\ B's\ Guess\ of\ the\ Final\ Number$
$\hat{\lambda}_{BG}^{Ae} \in [-40,40]$	$Player\ B's\ Best\ Guess\ of\ Player\ A's\ Private\ Information$
$F^{Ar} = 80$	$Player\ A's\ Fixed\ Payment$
$F^{Ae} = 100$	$Player\ B's\ Fixed\ Payment$

ACT condition (T3)

T3 is identical to T3b except Player A cannot send an upfront trust request to Player B. Instead, the game begins by Player B choosing whether to delegate her guess of the $Final\ Number$ to Player A, and subsequently reverts to the T3b game.

STATEMENT condition (T2)

T2 is identical to T3 except Player B cannot choose to delegate her guess of the $Final\ Number$ to Player A. Instead, the game begins by Player B choosing whether to send a cheap talk statement of trust to Player A, and privately reporting her beliefs. The exact statement Player B can send is: "I trust you to send me an accurate estimate of Player A's Private Information". The game then reverts to that found in T3 when Player B chooses not to delegate.

NO SIGNAL condition (T1)

T1 is identical to T2, except Player B cannot send a statement of trust to Player A. Instead, the game begins by Player B privately reporting beliefs of Player A's honesty, and subsequently reverts to Player A's decision of $\hat{\lambda}$ as in T2. See Figure 3 for a diagram showing the flow of game decisions across treatments.

Figure 3: Flow Chart of Study 2 Moves and Decisions

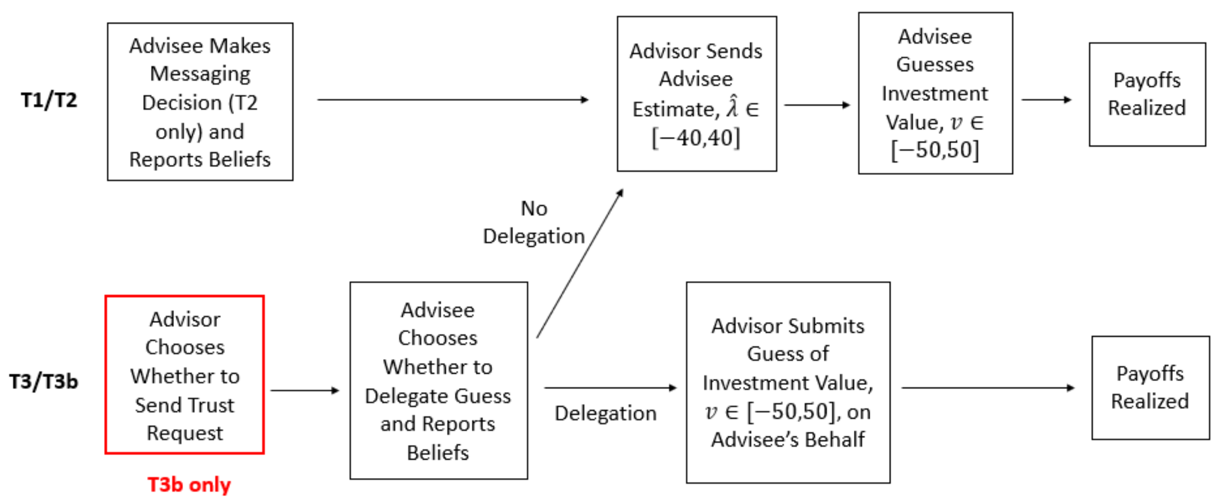

3.6 Results

I now present an examination of the Study 2 results. I begin with an analysis of the findings across treatments, before moving to a within-treatment analysis.

3.6.1 Cross-Treatment Results

First, I present the average levels of decision variables across treatments. In particular, Table 11 displays the mean levels of Player A's bias, Player B's discounting, and payoffs (both individual and combined).[13] Note that, if Player B delegates in T3 or T3b, I calculate Player A's bias as $\hat{v}^{Ar} - \hat{v}_{BG}^{Ar}$ instead of $\hat{\lambda} - \hat{\lambda}_{BG}^{Ae}$. Further, if Player B delegates then there is no measure of Player B's discounting. Thus, numbers in the table reflect discounting conditional on no delegation.

Table 11: Average Decision Variable Levels

	Player A's Bias $\hat{\lambda} - \lambda$	**Player B's Discounting** $\hat{\lambda} - \hat{\lambda}_{BG}^{Ae}$	**Player A's Payoffs**	**Player B's Payoffs**	**Total Payoffs**
T1 (NO SIGNAL) n=36	**29.9** (24.5)*	**24.8** (20.1)	**79.3** (21.5)	**74.0** (19.5)	**153.3** (29.0)
T2 (STATEMENT) n=30	**21.8** (21.1)	**21.4** (19.9)	**79.6** (17.8)	**78.0** (16.2)	**157.5** (24.7)
T3 (ACT) n=26	**24.8** (22.8)	**21.5** (21.0)	**92.0** (22.9)	**74.2** (19.1)	**166.3** (25.5)
T3b (REQUEST) n=30	**19.0** (22.5)	**19.0** (20.0)	**85.2** (23.0)	**77.7** (19.9)	**162.9** (26.4)

*Standard deviations shown in parentheses

Next, I evaluate how often Player B chooses to signal trust in T2-T3b, as well as the frequency of Player A sending a trust request in T3b. In T2, Player B chooses to send a statement of trust 72% of the time, which a proportion test reveals is higher than the 50% predicted by standard theory ($p{\sim}0.000$). This finding is consistent with that found in Study 1. Also consistent with Study 1, I find that, given Player B sent a trust statement, only 30.6% of Player Bs actually trust Player A to be honest. This provides further evidence that advisees strategically signal trust in COI situations. In T3, Player B chooses to delegate her guess of the $Final\ Number$ to Player A 21.5% of the time. This suggests that, while advisees strategically send statements of trust, they are more hesitant to initiate an act of trust. Still, 21.5% is significantly greater than the

[13] The results for Player A's bias in Table 11 are qualitatively consistent when I specify bias as $(\hat{\lambda} - \lambda)^2$ or $|\hat{\lambda} - \lambda|$ instead. Given this, I use the bias specification of $\hat{\lambda} - \lambda$ in the paper for ease of explanation. A similar reasoning is used for the specification of Player B's discounting, $\hat{\lambda} - \hat{\lambda}_{BG}^{Ae}$.

equilibrium prediction of 0% ($p{\sim}0.000$), suggesting that Player B may anticipate Player A will reciprocate her relatively strong signal of trust by being more trustworthy. Finally, in T3b, Player A sends a trust request message 83.3% of the time, which is significantly higher than the 50% predicted by theory ($p{\sim}0.000$). Interestingly, this has no effect on the proportion of Player B's who choose to delegate in T3b (18.0%) compared to T3 (21.5%) ($p = 0.554$).

Player A's Estimates ($\hat{\lambda}$).[14] Standard economic theory predicts that, given λ, Player A will choose any estimate $(\hat{\lambda})$ to send to Player B. To test whether this holds, I regress Player A's estimate on Player A's information (λ). If Player A chooses any $\hat{\lambda}$ given λ, then there should be no correlation between the two variables and the coefficient on λ would be zero. I run a separate OLS regression for each of the 4 conditions and cluster errors as the subject-level to account for within-subject correlation (Table 12).[15] I find that, in all four regressions, the coefficient on λ is significant and positive. This suggests that Player A does not choose any estimate after observing λ, but rather chooses an estimate that increases in the level of λ.

Table 12: OLS Regressions of Player A's Estimate on Player A's Private Information

OLS Regression: $\hat{\lambda} = b_0 + b_1\lambda$	T1: NO SIGNAL	T2: STATEMENT	T3: ACT	T4: REQUEST
$\boldsymbol{b_0}$	28.54** (1.20)+	22.58** (0.76)	24.28** (1.71)	18.80** (1.75)
$\boldsymbol{b_1}$	0.08** (0.03)	0.27** (0.003)	0.25** (0.05)	0.36** (0.07)
R^2	0.020	0.184	0.163	0.184
Number of Clusters	18	15	13	15

+ Standard errors are shown in the parentheses.

** Significant at the 5% level.

[14] When discussing T3 and T3b in this section, I use only those observations where Player B did not delegate, as these are the only observations where Player A sends an estimate to Player B.

[15] Unless otherwise noted, I cluster standard errors at the subject-level for the remaining Study 2 analysis.

Next, I present treatment-level scatterplots of Player A's estimate by Player A's information in Figure 4 below. In the plots, dots on the 45° line represent unbiased estimates. Further, dots above the line represent estimates that are biased upwards, whereas those below the line are estimates biased downwards. As expected, very few observations lie below the line suggesting that Player A either sends honest or inflated estimates. Note that, as dots appear higher on the graph, this indicates that Player A's estimate was relatively higher. Interestingly, one can see that, in T1, many of the dots are nearing the top of the graph, suggesting estimates with relatively high amounts of bias. However, in the conditions where Player B can signal trust (T2-T3b), the dots are more dispersed between the diagonal line and top of the graph, suggesting less biased estimates than in T1. Moreover, looking at dots lying on the diagonal line (non-biased estimates), one can see that there appears to be a higher proportion of these observations in T2 and T3b compared to T1 and T3. This provides suggestive evidence that statements of trust and trust requests may dampen Player A's bias rates. Even so, one can see that many of the unbiased estimates are for observations where the value of λ was relatively high, suggesting that Player A is more likely to be honest when doing so is potentially beneficial for him.

Figure 4: Scatterplots of Player A's Estimate by Player A's Private Information

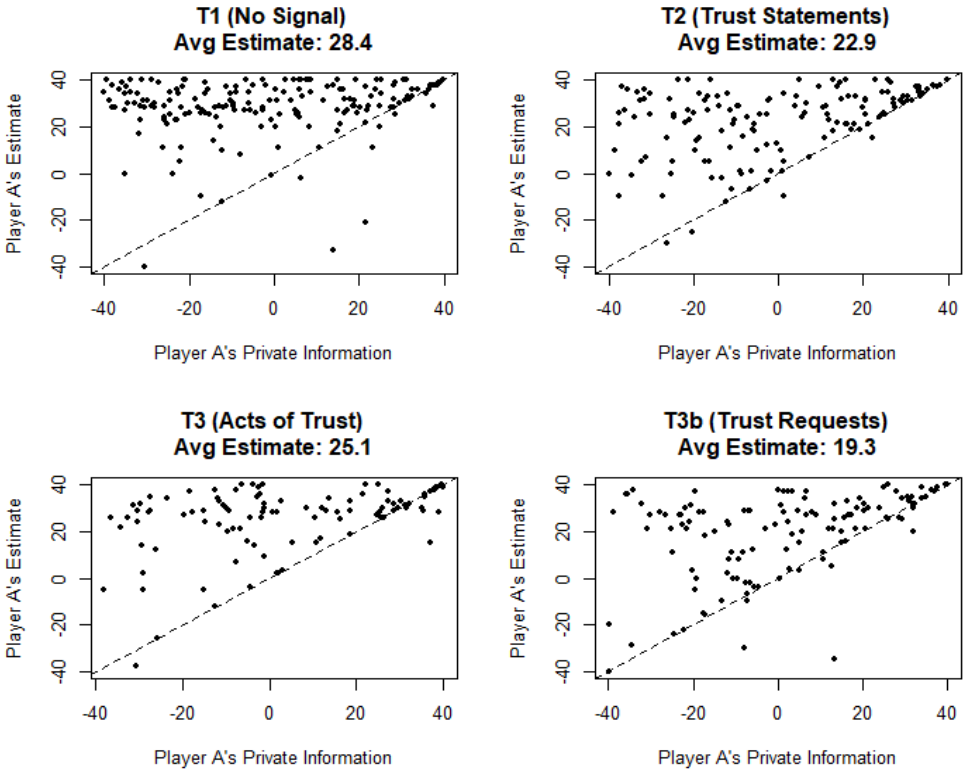

Player A's Bias.[16] Next, I evaluate Player A's bias level across treatments in more detail. Recall that standard economic theory predicts Player B does not delegate and Player A will choose any $\hat{\lambda}$. In expectation, then, the average bias level would be 0.[17] Not surprisingly, I find that bias levels are significantly higher than predicted in all four conditions (all p's ~ 0.000).[18] Thus, I conclude that Player A does not choose any $\hat{\lambda}$ given λ.

[16] In this section, I calculate bias as $\hat{\lambda} - \lambda$ whenever Player B does not delegate. Alternatively, when Player B does delegate in T3 and T3b, I calculate bias as $\hat{v}^{Ar} - \hat{v}^{Ar}_{BG}$. I combine all observations in this section, regardless of whether Player B delegates or not. I present an examination of how Player A's bias changes with and without delegation later in the paper.

[17] If instead I used a measure of Player A's bias as $|\hat{\lambda} - \lambda|$, then the expected level of bias is 25. In this case, Player A's bias level is higher than predicted in T1 (29.9; $p = 0.008$), lower than predicted in T2 (21.8; $p = 0.066$) and T3b (19.0; $p = 0.001$), and no different than predicted in T3 (24.8; $p = 0.927$).

[18] In T3 and T3b, results are qualitatively and statistically consistent if I analyze bias using only the observations where Player B did not delegate.

Next, I perform regression analysis to test for differences across conditions. To analyze Player A's bias across treatments, I regress Player A's bias on a categorical variable for treatment, and control for λ. I run a series of three OLS regressions and change the baseline treatment category so that all treatment comparisons can be observed. Results of the analysis are shown in Table 13. First, there is no statistical difference in Player A's bias level between T1 and T3, suggesting that acts of trust (without Player A's trust request) do not reduce Player A's bias compared to the situation where Player B cannot signal trust. However, compared to T1, Player A's bias is lower in both T2 (trust statements) and T3b (trust requests), while there is no difference between these two treatments. This is interesting because it shows: (1) statements of trust, which can be used across various contexts, can substantially reduce professional bias caused by COIs; and (2) allowing advisors to request for advisee trust, which can easily be introduced into service professionals' interactions with customers, can also improve advisor honesty levels.

Table 13: OLS Regressions Showing Differences in Player A's Bias Across Treatments

	DV = Player A's Bias		
	Base = T1	Base = T2	Base = T3
	(1)	(2)	(3)
λ	-0.778***	-0.778***	-0.778***
	(0.027)	(0.027)	(0.027)
T1		6.087***	1.563
		(1.387)	(2.029)
T2	-6.087***		-4.524**
	(1.387)		(1.804)
T3	-1.563	4.524**	
	(2.029)	(1.804)	
T3b	-9.053***	-2.966	-7.491***
	(2.047)	(1.806)	(2.316)
Constant	28.741***	22.654***	27.179***
	(1.178)	(0.739)	(1.652)
Observations	610	610	610
R^2	0.596	0.596	0.596
Adjusted R^2	0.593	0.593	0.593
Residual Std. Error (df = 605)	14.779	14.779	14.779
F Statistic (df = 4; 605)	222.910***	222.910***	222.910***
Note:			*p<0.1; **p<0.05; ***p<0.01

Player B's Discounting $(\hat{\lambda} - \hat{\lambda}_{BG}^{Ae})$.[19] Next, I turn to examining Player B's discounting of Player A's estimate. In Figure 5 below, I plot the level of Player B's discounting (y-axis) by Player A's estimate (x-axis) across treatments. In the figure, the horizonal lines indicate a discounting level of zero. Dots above the line indicate Player B thought Player A inflated his estimate, while those below the line represent observations where Player B thought Player A's estimate was lower than the truth. The equilibrium prediction is that Player B will ignore Player A's estimate. As such, the

[19] Results in this section are for observations where Player B did not delegate, as these are the only observations where Player B has a discounting measure.

distribution of observations above and below the line would be similar. However, the figure shows that the majority of observations lie above the line, indicating Player B thought Player A's estimate was inflated and suggesting the equilibrium behavior does not hold. Across all conditions, there are also some observations which lie on the horizontal line, suggesting that Player Bs do not always discount. This pattern is particularly noticeable in T2 and T3b, which provides suggestive evidence that Player B was more likely to trust Player A in those treatments. Next, note that there is large variation in the amount of within-condition discounting. That is, given a certain estimate, Player B did not always discount by the same amount within the same treatment group. This underscores the inherent heterogeneity in the data. However, one can see that, as Player A's estimate increases, the level of Player B's discounting tends to increase as well. Unsurprisingly, this suggests that Player B is more suspicious of high estimates. Finally, when I use regression analysis to evaluate whether Player B's discounting varies across treatments, I find no significant differences across treatments after controlling for Player A's estimate (all p's > 0.1).[20] Indeed, from Figure 5, the distribution of observations across treatments looks relatively consistent. For simplicity sake, I do not report the full regression results here. However, the lack of differences in discounting across treatments suggest that any differences in Player B's earnings across conditions may be driven by a reduction in Player A's bias, rather than by an improvement in Player B's discounting. Relatedly, I provide an examination of the players' earnings next.

[20] I use an OLS regression that controls for Player A's estimate (with heteroscedastic robust standard errors) to see whether the level of Player B's discounting varies across treatments.

Figure 5: Scatterplots of Player B's Discounting by Player A's Estimate

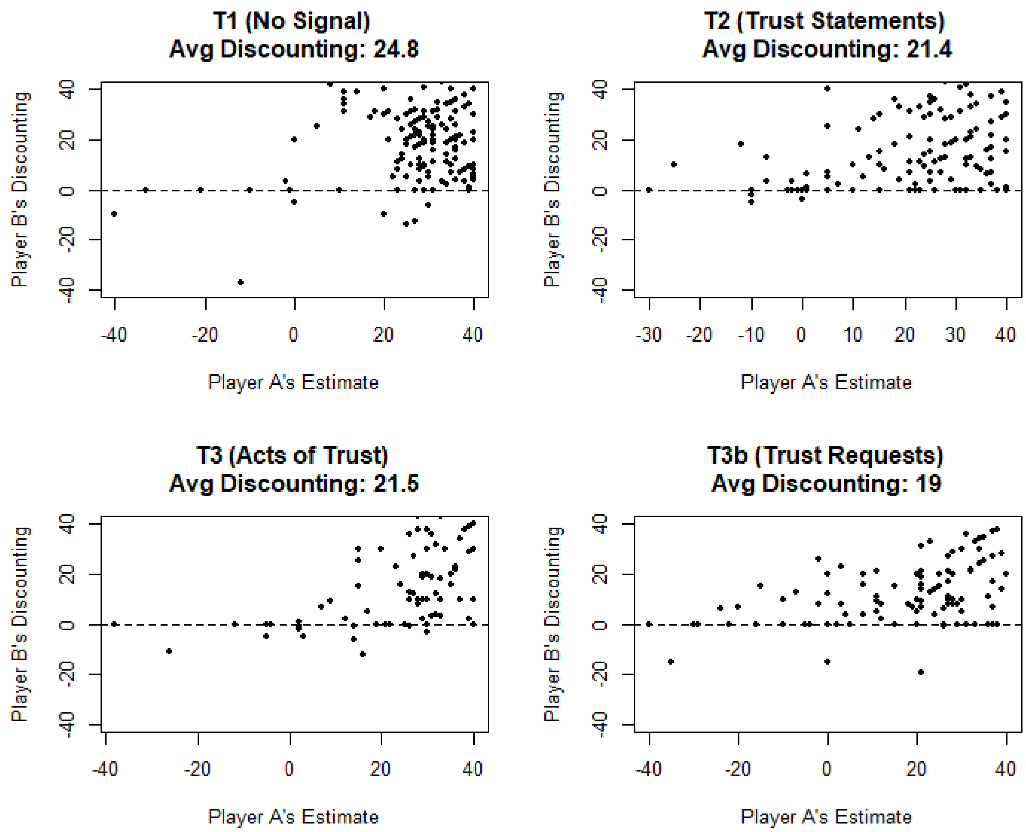

Point Earnings. Next, I evaluate Player B's earnings as well as the combined payoffs of both Player A and Player B. I examine the combined earnings to check for evidence of win-win situations.[21] I begin with Player B's earnings, which are predicted to be 75. I find that the average earnings for Player B in T1 and T3 are, 74.0 and 74.2, respectively, and they do not differ from 75 (T1: $p = 0.488$; T3: $p = 0.653$). In T2 and T3b, Player B's average earnings (78.0 and 77.7, respectively) are at least marginally higher than predicted (T2: $p = 0.027$; T3b: $p = 0.096$). Consistent with the insights from the analysis of Player A's bias level above, this suggests that customers may

[21] Win-win situations are important to study in COI situations. On one hand, improving customer outcomes is important, as customers are often the party being hurt by professional COIs. However, it's also important to delineate contexts where advisors can also improve their outcomes, as this can help motivate professionals to change their behavior, even when they have incentives to act opportunistically.

obtain improved outcomes when they can send trust statements to advisors, or when advisors can send trust requests to customers.

Next, Figure 6 presents scatterplots for each condition, showing Player B's earnings by the investment value, v. Note that, in the plots for T3 and T3b, I indicate the observations where Player B delegated her guess to Player A with an "×" symbol. Further, the horizontal lines indicate Player B's expected earnings, 75. First, the plots reveal an inverted-U shaped relationship between Player B's earnings and v. That is, Player B's earnings tend to be relatively high with moderate values of v, but decrease as v becomes higher or lower. One possible explanation for this could be that, when v is relatively high, Player A's estimate, $\hat{\lambda}$, tends to be high as well (since Player A observes that $Player\ A's\ Private\ Information$ is relatively high). As Player B discounts higher estimates more heavily, the advisee may end up over-discounting high estimates which leads to relatively low earnings when v is truly a high value. Further, when v is relatively low, Player A may bias by a lot to make up for this and Player B may not sufficiently account for this high level of bias. Consequently, Player B also faces relatively low earnings when v is low. While this inverted-U relationship still exists in T3 and T3b, it is interesting to note that there are many delegation observations (indicated by "×") lying on the 45° axis. These are observations where Player A submitted a guess of the investment value on Player B's behalf that was equal to the maximum value, 50. Thus, the pattern shown in the graphs indicates that delegating may not always be beneficial for Player B, a point I discuss in more detail below. Next, I turn to regression analysis to evaluate whether there are statistical differences in Player B's point earnings across conditions.

Figure 6: Scatterplots of Player B's Earnings by the Investment Value

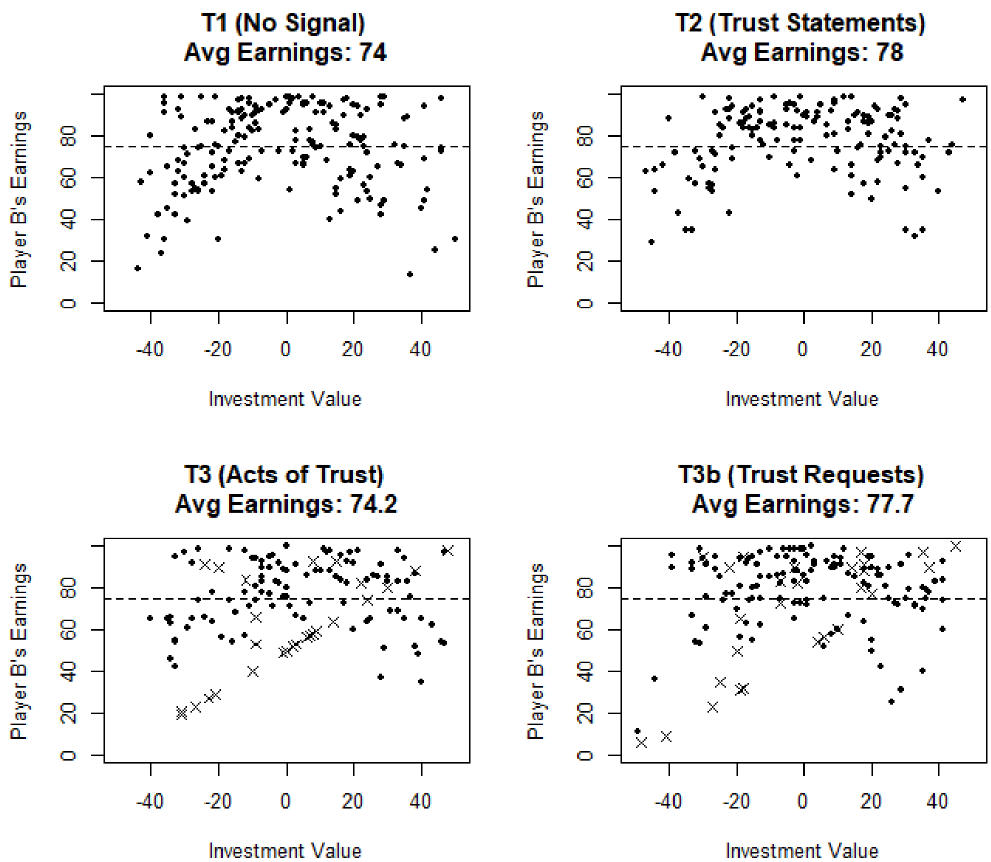

In particular, I regress Player B's point earnings on a categorical variable for treatment. Further, given the inverted-U relationship delineated above, I control for the square of the investment value, v^2. Results are displayed in Table 14 below. First, note the negative and significant coefficient on v^2, confirming the inverted-U relationship between the investment value and Player B's earnings. Further note that, as in Study 1, although there are relatively large differences across treatments, they are not always statistically significant at traditional levels. This may be due to Study 2's relatively large standard errors, which I'll discuss further in Section 3.7. However, I do find that Player B's point earnings are marginally higher in T2 compared to T1 ($p = 0.067$). Coupled with the findings presented above, this suggests that, when Player B sends a statement of trust to Player A, this reduces Player A's bias which results in higher earnings for Player B. Even though I find that Player A's bias was lower in T3b compared to T1, I do not find that Player B's earnings were significantly higher in the former ($p = 0.287$). However, when I run the same regression analysis again, but only include those observations in T3b where Player B did not delegate, then Player B's point earnings in T3b are marginally higher than in T1 ($p = 0.061$). This suggests that, in some cases, customers' outcomes can be improved when advisors ask customers to trust them.

Table 14: OLS Regressions Showing Player B's Earnings Across Conditions

	DV = Player B's Earnings		
	Base = T1	Base = T2	Base = T3
	(1)	(2)	(3)
v^2	-0.013***	-0.013***	-0.013***
	(0.002)	(0.002)	(0.002)
T1		-3.688*	0.017
		(2.012)	(2.726)
T2	3.688*		3.705
	(2.012)		(2.380)
T3	-0.017	-3.705	
	(2.726)	(2.380)	
T3b	2.983	-0.705	3.000
	(2.797)	(2.460)	(3.054)
Constant	80.989***	84.677***	80.972***
	(1.955)	(1.464)	(2.137)
Observations	610	610	610
R^2	0.142	0.142	0.142
Adjusted R^2	0.136	0.136	0.136
Residual Std. Error (df = 605)	17.48	17.48	17.48
F Statistic (df = 4; 605)	25.010***	25.010***	25.010***
Note:			*p<0.1; **p<0.05; ***p<0.01

Next, I evaluate total point earnings for both players. Note that theory predicts Player A earns 80 points and Player B earns 75 points. Thus, total earnings are 155 in expectation. Average total earnings in each treatment are as follows: T1 = 153.3, T2 = 157.5, T3 = 166.3, T3b = 162.9. In T1 and T2, total point earnings do not differ from 150 (T1: $p = 0.430$; T2: $p = 0.208$). However, in T3 and T3b, total earnings are significantly higher than predicted (both $p's \sim 0.000$). This means that either Player A, Player B or both players obtain higher than expected earnings in T3 and T3b. Below, I plot total point earnings by the investment value, v, for each condition (Figure 7). The dotted horizontal line indicates the equilibrium prediction of total earnings equaling 155. First, note that there are points above and below the line in each condition, indicating that subjects do not always perform at or above equilibrium. Again, this highlights the heterogeneity in the data. From the figure, a striking line of observations can clearly be seen going upwards across the top of the dots in all conditions. Given how payoffs are specified in the game, these observations represent those where the payoff-determining investment value guess (either $\hat{v}^{Ae}$ or $\hat{v}^{Ar}$) was at least as high as the true investment value (v).[22] Thus, observations below this line are those where the investment value guess was lower than the true value. This helps to explain why the frequency of dots below the line increases as v increases. Finally, the four plots look relatively consistent. This suggests that, together, Player A and Player B earn similar amounts across treatments. However, it is not clear how the relative proportion of total earnings for each player changes across conditions. I present an evaluation of this later in the section.

[22] In this case, total point earnings are equal to $F^{Ar} + F^{Ae} + v$, or $180 + v$.

Figure 7: Scatterplots of Total Earnings by the Investment Value

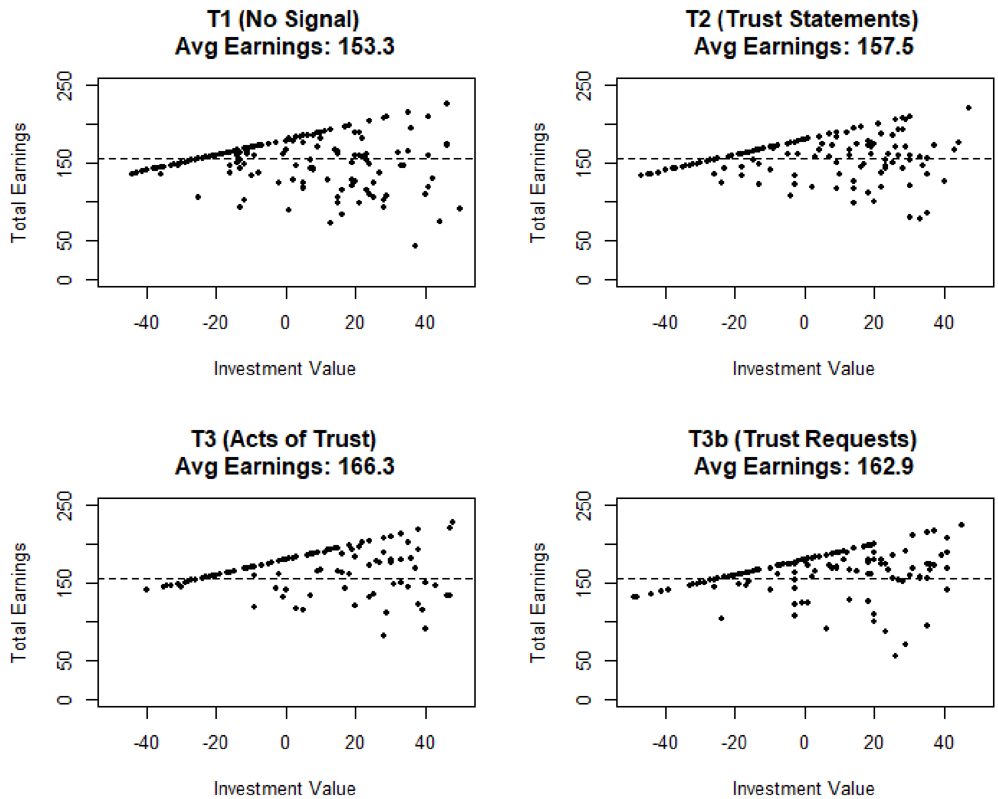

First, however, I run regression analysis to check for statistical differences in total earnings across conditions, controlling for the investment value, v. Here, cluster robust standard errors are necessary to account for the heteroskedasticity in the data. From Table 15, I find that, compared to T1 and T2, total earnings are higher in T3. This is interesting, since Player B's earnings were at least directionally lower in T3 compared to T1 and T2. This suggests that relatively high earnings for Player A in T3 (delegation without trust requests) may be driving the relatively high total earnings in the treatment. Further, I note that the p-value for the difference between T1 and T3b is 0.101, providing some suggestive evidence that total earnings tend to be higher when Player A can send a trust request to Player B. As Player B's individual earnings are also directionally higher in T3b compared to T1, this suggests that there could be potential for win-win situations when advisors can send advisees a trust request. I take a closer look at this now, by comparing the two players' point earnings across treatments.

Table 15: OLS Regressions Comparing Total Point Earnings Across Conditions

	DV = Total Earnings		
	Base = T1 (1)	Base = T2 (2)	Base = T3 (3)
v	0.165**	0.165**	0.165**
	(0.065)	(0.065)	(0.065)
T1		-4.013	-12.159**
		(5.276)	(5.666)
T2	4.013		-8.146**
	(5.276)		(4.091)
T3	12.159**	8.146**	
	(5.666)	(4.091)	
T3b	9.242	5.228	-2.918
	(5.626)	(4.133)	(4.554)
Constant	153.505***	157.518***	165.664***
	(4.606)	(2.582)	(3.120)
Observations	610	610	610
R^2	0.054	0.054	0.054
Adjusted R^2	0.048	0.048	0.048
Residual Std. Error (df = 605)	26.34	26.34	26.34
F Statistic (df = 4; 605)	8.636***	8.636***	8.636***
Note:			*p<0.1; **p<0.05; ***p<0.01

I regress each player's earnings on a categorical variable for treatment, controlling for v^2. A side-by-side comparison of the results are shown in Table 16.[23] The regression in column 1 is for Player B's earnings, while column 2 is the regression for Player A's earnings. From the table,

[23] I only present the results with T1 as the baseline condition for parsimony sake.

one can see that, in T2 and T3, only one player earns more than they do in T1. Specifically, Player B earns more in T2 than in T1, while Player A earns more in T3 than in T1. This suggests that allowing statements of trust and acts of trust (via delegation) does not result in win-win outcomes in COI situations. Interestingly, though, both players earn directionally higher in T3b compared to T1. Although the lack of statistical significance makes it difficult to make any robust claims here (Player B: $p = 0.287$; Player A: $p = 0.138$), it seems there could be a win-win situation when Player A requests Player B's trust. I evaluate this further when I look at the within-treatment results in Section 3.6.2.

Table 16: OLS Regressions Showing Comparison of Players' Earnings

	DV = Earnings	
	Player B	Player A
	(1)	(2)
v^2	-0.013***	0.005***
	(0.002)	(0.002)
T2	3.688*	0.394
	(2.012)	(3.870)
T3	-0.017	12.830***
	(2.726)	(4.486)
T3b	2.983	6.154
	(2.797)	(4.144)
Constant	80.989***	76.759***
	(1.955)	(3.563)
Observations	610	610
R^2	0.142	0.065
Adjusted R^2	0.136	0.059
Residual Std. Error (df = 605)	17.48	21.23
F Statistic (df = 4; 605)	25.01***	10.59***
Note:	*p<0.1; **p<0.05; ***p<0.01	

Up to this point, I have evaluated the data across treatments. I've found that, when Player B can choose to send a statement of trust to Player A and when Player A has the option of sending a trust request to Player B, this reduces Player A's bias. While it seems this reduction in bias with statements of trust drives higher earnings for Player B, the outcome of Player A's trust request on Player B's earnings is less clear. Moreover, I find that Player B is not better off, on average, when she is allowed to delegate her guess of the investment value to Player A (T3), but that Player A can benefit greatly from this. Interestingly, however, when compared to T3, Player A's earnings are lower in T3b ($p = 0.078$), suggesting that, contrary to what standard economic theory predicts, advisor trust requests may affect decisions and outcomes in a COI context.

3.6.2 Within-Treatment Results

Next, I turn to evaluating the data within treatments T2, T3 and T3b, to evaluate the effect of different types of trust signals, both when Player A can send a trust request and when he cannot.

T2: Do Statements of Trust Help?

First, I examine the results of T2, where Player B sent a trust statement 72% of the time. I evaluate how results differ when Player B chose to send a trust statement compared to when she chose not to. Mean results, with 95% confidence intervals, are displayed in Figure 8. Starting with the leftmost panel (Player A's bias), I find that, when Player B sent a trust statement, Player A's average bias was 20.0, which is directionally lower than Player A's bias when Player B did not send a trust statement, 26.4 ($p = 0.114$). While not statistically significant (likely due to noise in the data) this does provide suggestive evidence that observations where Player A received a trust statement from Player B are driving the lower bias levels in T2 compared to T1 (discussed in Section 3.6.1 above).

Figure 8: T2 Averages With and Without Statements of Trust

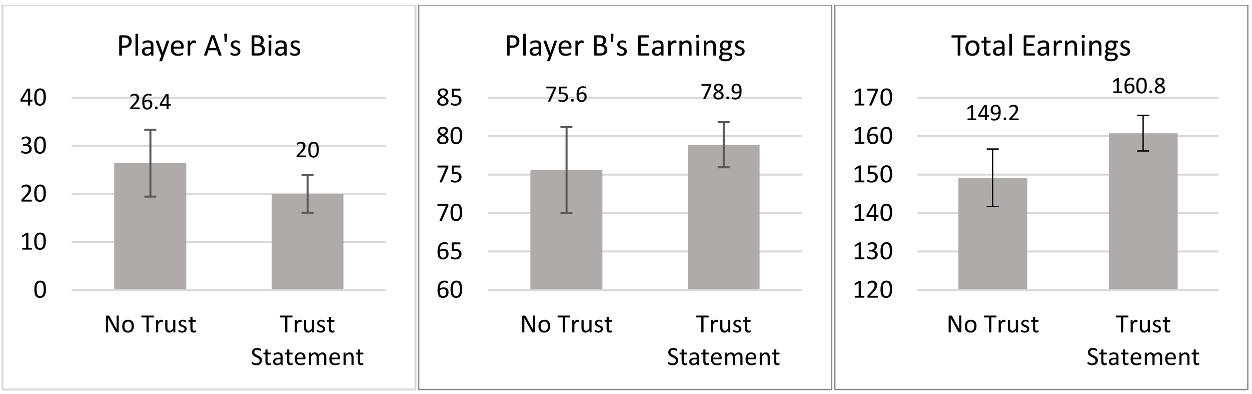

Next, I examine Player B's earnings. When Player B sends a trust statement, her average earnings are 78.9, compared to 75.6 when Player B does not send a trust statement. Again, the difference is not statistically significant ($p = 0.296$). However, as I found Player B's earnings are higher in T2 compared to T1, this finding suggests that the observations where Player B sent a trust statement are driving the higher T2 earnings. I also look at Player A's earnings briefly and find that they are significantly higher when Player B sends a statement of trust (81.9) compared to when Player B does not send a trust statement 73.6 ($p = 0.011$). Given both players earn (directionally) higher earnings when Player B signals trust, this suggests that, when an advisee is given the chance to signal trust to her advisor via a trust statement, doing so can result in a win-win situation for both advisor and advisee. In line with this, I also evaluate total earnings, with and without a trust statement. When Player B sends a trust statement, the combined earnings of Players A and B are 160.8, which is significantly higher than the combined earnings when Player B does not send a message, 149.2 ($p = 0.010$). In summary, I find that Player A is more honest and both players potentially obtain higher earnings when Player B sends Player A a statement of trust.

T3: Do Acts of Trust (Delegation) Help?

Next, I turn to T3, where Player B has the option of delegating her guess of the $Final\ Number$ to Player A (without trust requests from Player A). Recall that Player A's point earnings are equal to $80 + \hat{v}$, where, if Player B delegates, then Player A's submitted guess of the investment value ($\hat{v}$) determines payoffs. In the case of delegation, Player A is expected to choose $\hat{v}^{Ar}$ equal to $\overline{v}$, or 50. Note that Player A's COI is very strong in this case, as his payoff is exactly determined by his investment value guess. Consequently, Player B is predicted to never delegate, and Player A's earnings would instead be determined by Player B's guess of the $Final\ Number$, $\hat{v}^{Ae}$. As mentioned above, Player B delegated 21.5% of the time in T3, significantly more than predicted. I now evaluate decisions and outcomes when Player B did and did not delegate (Figure 9).

Figure 9: T3 Averages With and Without Acts of Trust (Delegation)

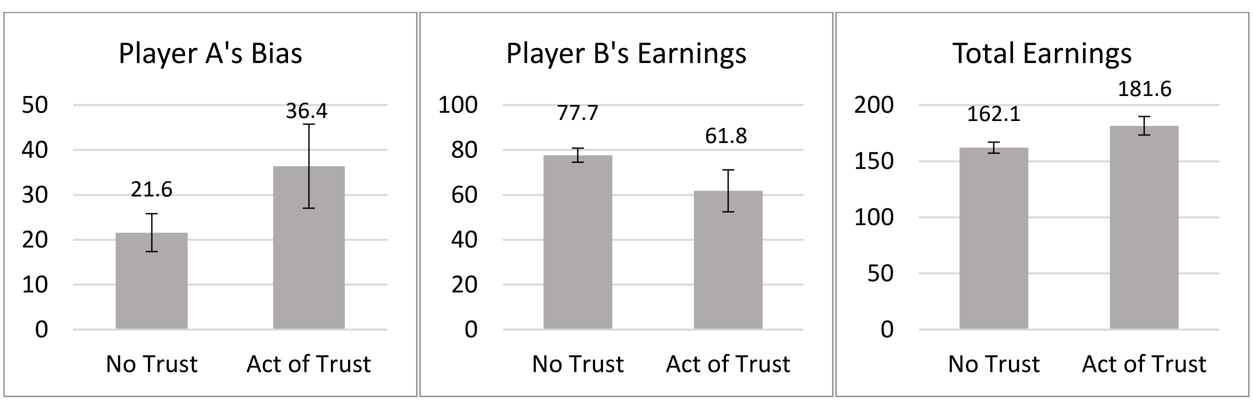

First, when Player B delegates, Player A's average bias is 36.4, compared to 21.6 when Player B does not delegate. This difference is statistically significant ($p = 0.006$). Thus, it seems that Player A is substantially more biased when Player B delegates compared to when she does not. Moreover, there is evidence to suggest that delegating can also have a deleterious effect on Player B's earnings. Specifically, Player B earns an average of 61.8 points when she delegates and 77.7 when she does not ($p = 0.002$). Together, these results suggest that Player B is better off not delegating. Not surprisingly, Player A's earnings when Player B delegates are significantly higher than when she does not (119.9 vs. 84.4; $p \sim 0.000$). Given this, it seems that an increase in Player A's earnings alone drive the higher total earnings with delegation, 181.6, compared to no delegation, 162.1 ($p \sim 0.000$). Thus, while advisee delegation is beneficial for advisors with COIs, it can have detrimental effects for advisees.

T3b: Do Acts of Trust (Delegation) Help When Player A Can Send a Trust Request?

Finally, I examine results within T3b, where Player A can send an upfront, non-binding trust request to Player B, who then decides whether to delegate to Player A. The equilibrium prediction in T3b is identical to that in T3 (i.e., trust requests are not expected to affect behavior). Recall that, in T3b, Player A chose to send a trust request 83.3% of the time, and Player B delegates 18% of the time. That is, overall, Player A's trust requests do not increase the proportion of Player Bs who delegate compared to T3. Indeed, I find that Player B is not more likely to delegate when she receives a trust request from Player A compared to when she does not (19.2% vs. 12.0%, $p = 0.569$).

Since I find that the vast majority of Player As send a trust request, I do not break down the results by whether Player A sent a trust request, but rather I examine the findings based on whether Player B delegated or not.[24] Results are displayed in Figure 10. First, Player A's average bias level is 24.3 when Player B delegates, which, surprisingly, is not significantly higher than the level when Player B does not delegate, 17.9 ($p = 0.268$). This result is interesting, considering that Player A's bias was substantially higher when Player B delegated in T3 compared to when Player B did not delegate. This suggests that Player A's trust requests can have a large effect on reducing Player A's bias, especially when Player B delegates to Player A. I explore this point further below. Next, looking at Player B's earnings, I find that she obtains marginally lower earnings when she delegates, 68.1, compared to when she does not, 79.8 ($p = 0.050$). Still this difference in payoffs when Player B does and does not delegate, 11.7 points, is smaller than the difference found in T3 (15.9 points), providing further evidence that delegating after receiving a trust request from Player A is not as bad for Player B. Again, I find that Player A's earnings are significantly higher when Player B delegates (108.8) compared to when she does not (78.0; $p \sim 0.000$). As in T3, this increase in Player A's earnings likely drives the higher total earnings when Player B delegates, 176.9, compared to when she does not, 159.8 ($p = 0.002$). Overall, the results suggest that Player B should not signal trust of Player A via acts of delegation, even when Player A requests that Player B do so. However, I also find evidence that delegation is not as detrimental for Player B when Player A sends her a trust request first. Next, I examine this in more detail.

[24] The analysis in this section includes all observations in T3b, not only those where Player A sent a trust request. However, I note that the results are qualitatively consistent if I only use the subset of data where Player A sent a request.

Figure 10: T3b Averages With and Without Acts of Trust (Delegation)

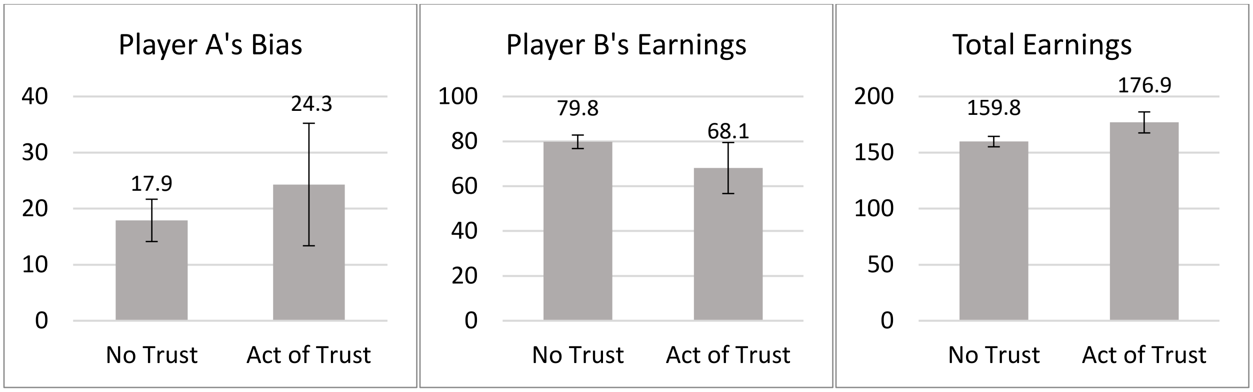

T2-T3b: How Do Different Trust Signals Stack Up?

I conclude my analysis of the Study 2 data by evaluating the effect of trust signals across treatments, to obtain a clearer picture of the type of trust signal advisees should send, and when. For this, I only consider the subset of observations in T2, T3 and T3b where Player B signaled trust (either through statements of trust or delegation). I consider two main outcomes, Player A's bias and Player B's earnings.

Player A's Bias. I begin by examining Player A's bias conditional on receiving a trust signal across conditions (Figure 11). First, comparing between T2 and T3, I find that acts of trust result in higher levels of Player A's bias, 36.4, compared to statements of trust, 20.0 ($p = 0.002$). That is, Player A is more dishonest when Player B sends a strong signal of trust by allowing Player A to submit a guess of the investment value on her behalf. Moreover, comparing between T3 and T3b, I find that, surprisingly, Player A's bias is marginally lower when Player B delegates in T3b, 24.3, compared to T3, 36.4 ($p = 0.089$) In fact, Player A's bias is reduced by nearly one-third in T3b, which suggests that Player B is better off showing her trust by delegating when Player A requests that Player B do so. This is a surprising result, given standard economic theory predicts that Player

A's bias should not depend on whether Player A is allowed to send a trust request. Finally, comparing between T2 and T3b, I find no significant difference between Player A's bias with a statement of trust, 20.0, compared to that with an act of trust, 24.3 ($p = 0.461$). This suggests that signaling trust via statements and signaling trust via decision delegation (given Player A can sent trust requests) can have similar effects on Player A's bias level.

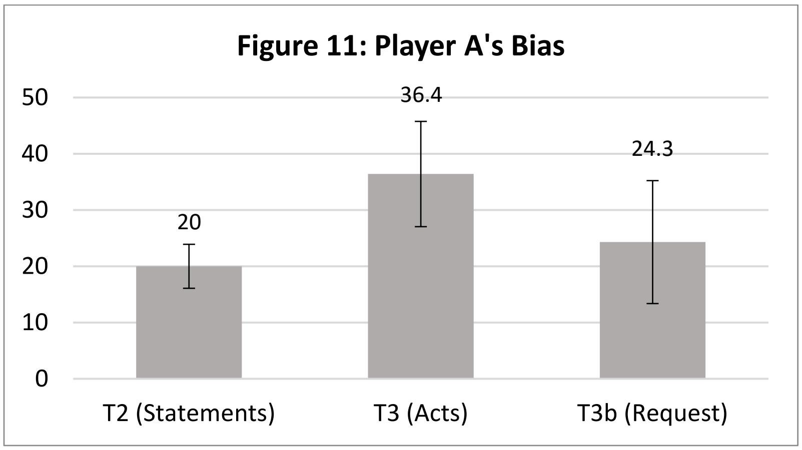

To provide a more visual display of the distribution of Player A's bias, I present a series of boxplots and cumulative distribution functions (CDFs) in Figure 12. The figure indicates that the bias levels in T2 and T3b are similar, although there is somewhat higher variance in the T3b distribution. One can also see that bias levels in T3 are consistently higher than in T2 or T3b, which is consistent with the results presented in Figure 11 above.

Figure 12: Distributions of Player A's Bias Conditional on Trust Signal

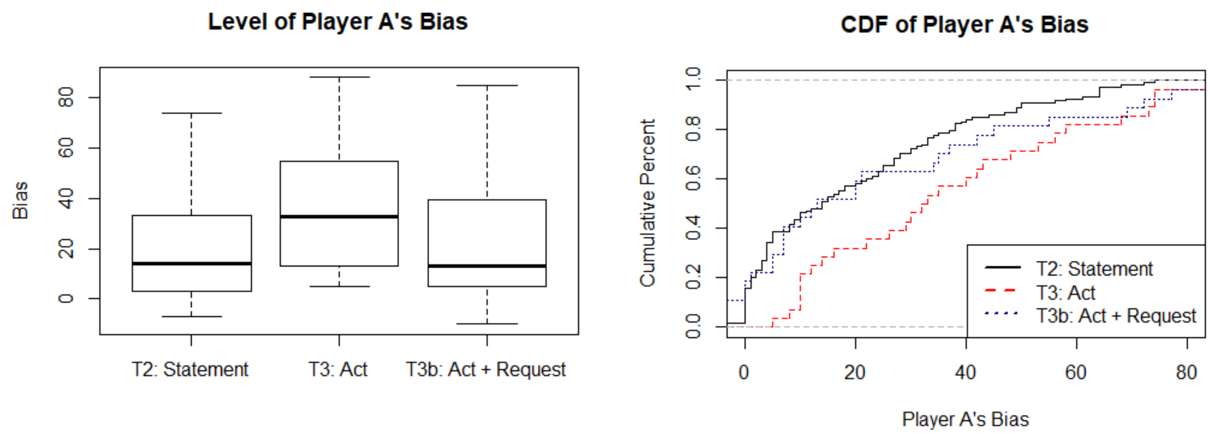

Player B's Earnings. I now evaluate Player B's earnings across T2, T3 and T3b, conditional on Player B signaling trust. I display the average levels in Figure 13. First, comparing between T2 and T3, I find that acts of trust (i.e., delegation) result in lower earnings for Player B, 61.8, compared to statements of trust, 78.9 ($p = 0.001$). That is, Player B's earnings are lower when she sends a strong signal of trust and allows Player A to submit a guess of the investment value on her behalf. This result is consistent with the findings of Player A's bias above, suggesting that higher bias levels in T3 also drive the lower earnings for Player B in T3, compared to T2. Comparing between T3 and T3b, I find that Player B's earnings are directionally higher in T3b, 68.1, compared to T3, 61.8, but this is not statistically significant ($p = 0.380$). This suggests that, while Player A is somewhat less biased when he can send a trust request and Player B delegates, compared to when he cannot send a request, this does not necessarily result in higher earnings for Player B. Finally, comparing between T2 and T3b, I find that Player B's point earnings are marginally higher in T2, 78.9, compared to in T3b, 68.1 ($p = 0.070$). This provides additional evidence that, even if an advisor asks an advisee to trust him, the advisee should not always do so.

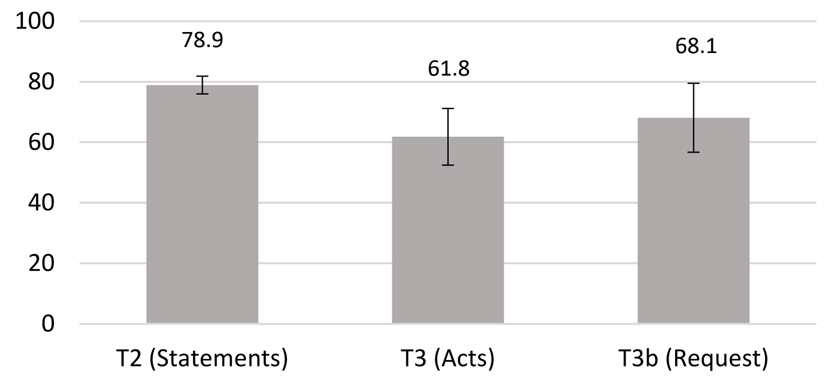

Figure 13: Player B's Earnings

Finally, I present a series of boxplots and CDFs of Player B's earnings in Figure 14. One can see that, overall, earnings are less disbursed in T2 compared to T3 and T3b. Comparing T2 to T3b, the median level of earnings is relatively consistent, although there is higher variance in the T3b distribution. Further, the plots provide suggestive evidence that earnings in T3 are generally lower than in T2 and T3b, however again there is a large amount of heterogeneity in the data.

Figure 14: Distributions of Player B's Earnings Conditional on Trust Signal

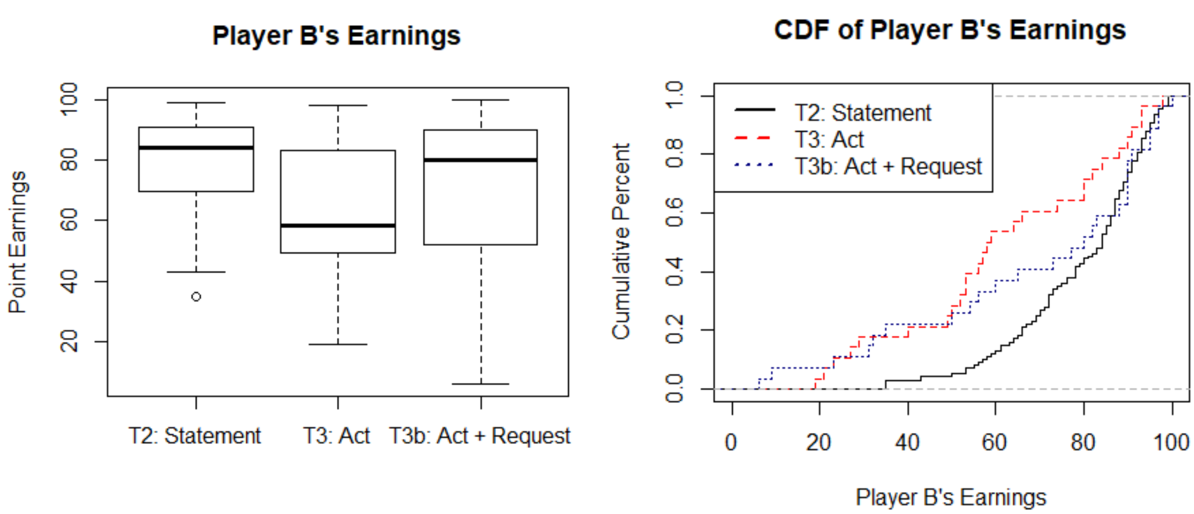

3.7 Summary & Discussion

Study 2 demonstrates that the type and context of an advisee's trust signal matters in a COI situation. Consistent with Study 1, Player A's bias is relatively high and Player B's earnings are relatively low when Player B cannot signal any trust to Player A (T1). This suggests that professionals' COIs can have a deleterious effect on expert honesty and customer outcomes and speaks to the importance of understanding both how advisees can improve outcomes for themselves and also how service professionals can be motivated to provide honest advice in the face of COIs. The benefit of more honest advice for customers is straightforward, as this should result in customers making more optimal decisions overall. But the positive outcomes extend beyond those for just the customer herself. For instance, as discussed in the introduction, improving advisor honesty in a COI situation can increase overall market efficiency and, from a PSF's perspective, can even result in long-term revenue gains due to reputational improvements.

The findings presented in Study 2 contribute to our understanding of how these beneficial outcomes can be achieved. First, the data suggests that customers often choose to strategically send statements of trust to advisors. Indeed, doing so can reduce bias and improve customer earnings. This is interesting because standard economic theory predicts that statements of trust should have no effect on expert bias or advisee earnings. As delineated in the discussion of my behavioral hypotheses (Section 3.1), one reason why statements of trust might improve outcomes in a COI situation is that advisors may reciprocate this signal of trust by being more honest. Anticipating this, customers may choose to strategically send trust statements to obtain more accurate information.

Next, I find that, in the absence of a trust request from their advisor, customers should not delegate their decision to the advisor. I find that doing so greatly increases advisor bias and reduces advisee earnings. Interestingly, this suggests that trust signals are not always beneficial for customers, and even though weak signals of trust (i.e., trust statements) can help, stronger signals (i.e., delegation) can backfire. For instance, amateur investors should tell financial advisors that they trust their advice but should not ask the advisor to make an investment decision for them. As only 21.5% of advisees chose to delegate in T3, it seems that this perverse result may be anticipated. However, this still leaves over one fifth of customers who face detrimental consequences. If my behavioral conjecture holds, and experts do feel stronger reciprocity towards advisees when they receive stronger signals of trust from advisees, then this perverse result of delegation may be seen as contradictory. However, I also presented the possibility that, when advisees delegate, this reduces an advisor's risk of acting opportunistically. Thus, even though relatively strong trust signals may intensify advisor reciprocity towards the advisee, it's possible that this increase is outweighed by the reduced risk that the advisor faces when biased (given the advisee delegates). Consequently, advisee delegation leads to higher bias levels and reduced customer earnings.

Finally, when the advisor can send a trust request to the advisee (asking advisee to delegate), I find that this does not increase advisees' propensity to delegate (18.0% vs. 21.5%). Interestingly, though, I do find that trust requests reduce advisor bias given the advisee delegates. That is, bias is lower and customer earnings are higher when advisees delegate after receiving a request to do so, compared to when advisees delegate without a trust request from the advisor. However, regardless of whether the advisor sent a trust request, advisee earnings are still lower with delegation compared to without. This suggests that customers should not delegate, even when their advisor asks them to. Still, the findings do offer evidence that, given a customer wants or needs to delegate (e.g., low risk situation, low situational knowledge), then she should do so upon eliciting a trust request from her advisor. This finding is consistent with my behavioral conjecture. To elaborate, in T3b, the advisor still faces increased reciprocity from receiving a relatively strong signal of trust (i.e., delegation), as in T3. However, in T3b, the advisor's reciprocity may be intensified even further, given that he sent a trust request. This increased reciprocity may be driven by feelings of guilt because the advisor specifically asks the advisee to delegate and, in exchange, the advisor will help the advisee out. Thus, given the relatively high level of reciprocity the advisor feels towards the advisee when the advisor sends a trust request which the advisee accepts, this may outweigh or at least counterbalance the reduced risk of bias that the advisor faces when the advisee delegates. Consequently, while delegating without a request to do so may have a relatively large deleterious effect on customer outcomes, delegating when asked to do so by an advisor can result in less detrimental outcomes for customers.

Overall, the findings suggest that customers should state their trust to advisors but should be wary of showing their trust through acts of delegation. However, given a customer does feel the need to delegate, it is better to do so when it is requested by the advisor. Indeed, when advisors send trust requests to advisees, this can have a substantial effect on reducing advisor bias. Thus, training programs for professional service workers who have COIs (e.g., doctors, financial advisors, etc) may want to encourage professionals to ask for their customers' trust.

While I believe the present research provides valuable insights for both consumers and marketing practitioners, there are some limitations. First, I do not study how decisions in long-term relationships play out. It's possible that, in this case, advisors are less biased and customers enjoy higher earnings. Future research could examine how repeated interaction between an advisor and advisee affects outcomes. Next, in Study 2, I do not look at the case where Player A sends a trust request when Player B is only allowed to send statements of trust. I chose not to study this as I believe the influence of Player A's trust requests is more appropriately studied in the context of decision delegation. Consequently, it's not clear how results would change if advisors could send a request asking advisees to trust their estimate (when no delegation is available). Furthermore, while I do outline a behavioral theory that could possibly explain the data, I do not offer a behavioral model to test this theory. Future research could build an empirical model of how reciprocity and trust affect outcomes in a COI situation, given different types of trust signals. Finally, the data presented in Study 2 is quite noisy, which reduces my ability to find statistically significant differences. There are a few reasons for this. First, there is a large amount of decision heterogeneity in the data. For instance, some advisors tend to be more biased than others, and some advisees tend to be more trusting. Not only does this lead to large variance within each player's decision variables, but it also makes it difficult for subjects to determine whether their counterpart is, say, being honest or whether they should discount by a lot. In fact, there are several instances in the data where, for instance, Player A biases by 60 points, but Player B does not discount at all. These large mismatches greatly increase the data noise. Next, the game itself is quite complex and requires that subjects make decisions in a context where their private information may not be so detailed. On top of this, when there are multiple parties interacting, an error by one subject can affect several decisions. While Study 2 did reduce the data noise significantly compared to Study 1, future work that aims to improve this even further could be a worthwhile endeavor.

4. Customer Misinformation Signals (Study 3)

Like Studies 1 and 2, Study 3 evaluates how customer signals affect decisions and outcomes in a COI situation. Specifically, Study 3 examines how customer *misinformation* signals affect the route a taxi driver takes in a field setting, and the subsequent ride fare and duration for the taxi passenger. Today's connected consumers have access to virtually limitless amounts of information, some correct and some not. It is not surprising, then, that customers sometimes share inaccurate information when interacting with professionals.[25] For instance, a sick patient may share erroneous information found on WebMD with her physician or a novice investor may share her friend's mistaken advice with a financial advisor. Moreover, professionals may believe that misinformed customers have a relatively low likelihood of detecting advisor bias. Thus, customer misinformation signals may result in opportunistic professional responses, especially in COI situations. In my context, many taxi drivers have a COI because they can potentially earn higher fares by taking an out-of-the-way route that is longer and more expensive than necessary, even though doing so is detrimental for passengers who hired the taxi driver with an expectation that the driver would take them to their destination via the fastest route possible.[26] In this case, drivers may be more willing to cheat passengers who signal that they are misinformed on the fastest route to take. Thus, in Study 3, I evaluate the effect of customer misinformation, both when signaled on its own and when accompanied by trust.

---

[25] In Section 5, I present a single-condition lab experiment (Study 4). The results suggest that customers tend to overshare information with professionals, even when there's a chance this information may not be accurate.

[26] Here, I consider more traditional taxi markets, where fares are variable and determined by ride distance and/or duration. This is in contrast to fixed fare taxi services provided by ride sharing companies such as Uber and Lyft.

While the main focus of Study 3 is customer misinformation signals, my rationale for incorporating trust signals is as follows. First, in Studies 1 and 2, I find that most advisees choose to send trust signals. Given this, it is important to understand how these signals play out in a field setting. Further, it's possible that, when customers know their information may not be correct, they even strategically signal trust as a sort of "safety net", to deter professional bias if they are in fact mistaken.[27] In my setting, when taxi passengers signal trust, taxi drivers may reciprocate this trust by being more trustworthy, even if passengers are perceived as being misinformed. If this is the case, it offers a more relaxed stance on how customers should share information with professionals, as even if customers' information is wrong, they might not be cheated as long as they supplement their information with a trust signal. However, as shown in Study 2, not all trust signals are beneficial for customers and it's possible that trust signals actually increase drivers' propensity to cheat passengers. Consider the case where drivers believe passengers are less likely to monitor their behavior (e.g., verify on Google Maps that the driver is taking the fastest route) when passengers signal trust. Given this, it'd be best if passengers refrained from seeming too trusting of drivers. Thus, ex-ante, the effect of a trust signal in addition to a misinformation signal is not so clear.

In the market for taxi services, a COI problem arises because drivers are often better informed than passengers about what route is optimal given current traffic patterns. Further, there may be many available routes the driver could take to the passenger's destination, with

[27] Note that trust signals differ from misinformation signals in one key way. Namely, customers generally have no incentive to appear misinformed when they are actually correctly informed. In fact, customers who communicate misinformation likely believe that they are communicating correct information. Trust signals are different in that customers can easily and knowingly state they trust a professional, regardless of their actual beliefs. Thus, trust signals are particularly interesting to study because they can be so widely and effortlessly used.

certain routes earning the driver a higher fare than others. Drivers may have an incentive to take a longer and more expensive route to earn a higher fare, especially if they believe the passenger would not be able to detect when the driver took an out-of-the-way route. Given the information asymmetry found in taxi contexts, taxi rides often fall under the classification of credence goods (Darby and Karni 1973).[28] Credence goods have the important property that, while customers can observe the utility they gain from a purchase, they are not able to verify whether the purchased good is the ex-ante needed one. Beyond taxi markets, credence goods and services can be found across a variety of fields including healthcare, finance and vehicle repair. For instance, mechanics can often diagnose vehicle issues that customers can neither detect nor identify. While this information asymmetry is the exact reason why customers seek out professional assistance in the first place, it also means that experts can potentially take advantage of customers and mislead them to purchase goods and services that are more expensive than necessary. This deception not only manifests in higher costs for consumers but also results in decreased market efficiency.

Dulleck and Kerschbamer (2006) offer a model of a credence good market that classifies how market inefficiencies arise. The authors posit that the sale of credence goods can give rise to two treatment-related inefficiencies: (1) *undertreatment*, where the customer requires a high quality good but receives a good of lesser quality and (2) *overtreatment*, where the customer

---

[28] Note that taxi rides may not necessarily be a credence good if the passenger is familiar with the possible taxi routes and knows which is the fastest/least expensive. This may happen for local passengers, or those who repeatedly use taxis to travel between two locations. In these cases, the taxi drivers may not be better informed than the passenger on which is the optimal route to take. However, in Study 3, I am interested in taxi rides taken by passengers who are perceived as not being familiar with the area, and thus it is almost certain that the taxi driver is better informed than the passenger.

requires a low quality good but receives a good of higher quality. For overtreatment, the additional benefits that the high-quality good offer beyond the low-quality good do not outweigh the additional costs associated with the high-quality product. In the present context, one can think of a taxi driver who takes a longer-than-necessary (and consequently, more expensive) route as overtreating a customer. Dulleck and Kerschbamer (2006) also posit that, when customers are relatively well-informed on the good or service they require, there is little room for experts to overtreat, since these customers are able to detect if the expert is trying to sell them a good that is more expensive than necessary. This may explain the finding by Balafoutas and colleagues (2013) that foreign taxi passengers are overtreated more than local passengers. Further, since it is unlikely that passengers would be dropped off before their desired destination in a taxi context, the incidence of undertreatment is not so relevant in Study 3, and rather my focus will be on the inefficiency that arises from overtreatment. This leads me to Study 3's research question: How do customer signals of misinformation affect drivers' propensity to overtreat in a COI context? Both when customers also signal trust, and when they do not.

4.1 Experimental Design & Hypotheses

Study 3 presents a field experiment in the context of taxi services. There are three conditions in the study. In the Control condition, the passenger enters a taxi and simply asks to be taken to the destination. The "Misinfo" condition is identical, except the passenger also (mistakenly) states that [Route X] seems to be the fastest route to take. Finally, the "Misinfo + Trust" condition is identical to the "Misinfo" treatment, but the passenger additionally indicates s/he trusts the driver to take the fastest route.

The experiment is designed such that there are two main routes from the place of departure to the destination. One route (Route Y) is typically the fastest and least costly route. The second route (Route X) is, under normal conditions, longer and more costly to take. The driver has an incentive to take Route X in order to earn a higher fare, even though doing so causes market inefficiency by forcing the passenger to pay a higher-than-necessary fare. The study's key dependent variable is which route, Route X or Route Y, the driver chooses to take, and I also evaluate ride fare and duration.

The intuition for the experimental design is as follows. In the Control condition, passengers enter a taxi and simply state their destination. I chose this Control condition because it is the simplest communication between the passenger and driver, and maps well to a common taxi context. Furthermore, even passengers in the Control condition may be perceived as being unfamiliar with the city (i.e., uninformed about the best route to take), as they board a taxi from a hotel going to the airport, carry luggage and a map, and speak in a foreign accent. However, since passengers in the Control do not signal that they are misinformed (just uninformed), it is not clear whether drivers will overtreat these passengers. Dulleck and Kerschbamer (2006) suggest that being perceived as lesser-informed in a certain area (which, arguably, foreign passengers are) should lead to overtreatment. And indeed, past research by Balafoutas et al. (2013) suggests that drivers do overtreat passengers who are perceived as being uninformed foreigners (in Athens, Greece). In the present study, I test the robustness of this finding in a different, Asian country by looking at the incidence of overtreatment in the Control. Furthermore, while passengers in the Misinfo condition are similarly perceived as foreigners, they also explicitly state their misinformation about which is the best route to take. This signal likely leads the driver to perceive a larger information asymmetry between himself and the passenger (compared to a Control passenger) and, thus, I expect a higher incidence of overtreatment in the Misinfo condition compared to the Control. Finally, as discussed above, it is unclear how signaling trust alongside misinformation will affect driver deception. Consequently, I present competing hypotheses for whether the incidence of overtreatment in the Misinfo + Trust condition will be higher or lower than in the Misinfo treatment. I specify the Study 3 hypotheses below.

> ***H1**: Passengers in the Control condition are prone to overtreatment since they are perceived as being uninformed (foreigners).*
>
> ***H2**: Since they are perceived as being misinformed, passengers in the Misinfo condition are more prone to overtreatment than passengers in the Control condition.*
>
> ***H3a**: Passengers who signal that they trust the driver, on top of being misinformed, are less prone to overtreatment than passengers in the Misinfo condition.*
>
> ***H3b**: Passengers who signal that they trust the driver, on top of being misinformed, are more prone to overtreatment than passengers in the Misinfo condition.*

4.2 Experimental Procedure & Manipulations

The implementation of Study 3 is similar to that used in Balafoutas et al. (2013, 2015). Ten research assistants (RAs), all of whom were students at NUS, assisted with data collection in Singapore. The RAs worked over the 11-month period from December 2017 through October 2018. All of the RAs were foreigners (i.e., non-Singaporeans) who were either from Vietnam or China, and who had been living in Singapore for less than a year at the time of data collection.[29]

---

[29] I used foreign RAs to make the study as realistic as possible, as foreigners are most likely to be misinformed on the best route to take.

This is important because all RAs still had very strong foreign accents, which would clearly signal to any taxi driver that the passenger was not local. Out of the five RAs from China, 3 were male and 2 were female. The five RAs from Vietnam consisted of 2 males and 3 females.

To help ensure the drivers viewed the RAs as foreign passengers, the RAs all carried a suitcase and a map of the city while taking rides, which originated at a hotel's taxi stand (center of city) and ended at Singapore's Changi Airport (outskirts of city).[30] The RAs took separate, sequential taxi rides one right after the other. I will refer to each set of two rides as a "trip".[31] In each trip, rides were randomly assigned to be in one of the three conditions. Table 17 shows an example of how the taxi trips are broken down.[32]

Table 17: Example of Trip Breakdown in Study 3

	Ride	RA	Condition
Trip 1	1	RA 2	Misinfo
	2	RA 3	Control
Trip 2	1	RA 7	Control
	2	RA 4	Misinfo + Trust
Trip 3	1	RA 3	Misinfo + Trust
	2	RA 5	Misinfo

[30]Most taxi rides in Singapore originate from a taxi stand. The RAs entered the taxi stand queue and, once they were at the front of the line, boarded whichever taxi was next to arrive. The RAs were instructed not to interact with one another while at or around the taxi stand.
[31] Within each trip, RAs were the same nationality. This is so I can better control for locality effects when analyzing the data.
[32] Although Study 3 has three conditions, the RAs worked in pairs and took rides two at a time. This is due to logistics, as it was difficult to find non-local RAs from the same country who could work at the same time. However, since rides were randomly assigned to a condition, this should not affect the results.

Upon entering the taxi, while looking at a map of the city, the RA would communicate his/her destination to the driver based on condition. The exact wordings are as follows.[33]

Control: "Hello, I'd like to go to Changi Airport Terminal 3, please."

Misinfo: "Hello, I'd like to go to Changi Airport Terminal 3, please. From the map, it looks like the fastest route to take is East Coast Parkway. So, please take East Coast Parkway, unless you know a faster route." [34]

Misinfo + Trust: "Hello, I'd like to go to Changi Airport Terminal 3, please. From the map, it looks like the fastest route is to take East Coast Parkway, but I trust you to take the fastest route."

Under normal conditions, taking the Pan Island Expressway (PIE) is the fastest and least costly route from the hotel taxi stand to the airport. Drivers can also take the East Coast Parkway (ECP), but this route is longer and results in more expensive fares for passengers.[35]

During the ride, the RA mapped the driver's route discretely, using a GPS application on his or her phone. The app tracked the route, distance, duration and time of the ride. At the completion of a ride, the RA noted the trip fare and requested a receipt for the trip (as a second form of data collection). As controls, the RAs also recorded the gender and approximate age of the taxi driver, taxi company and license plate number (to track repeat drivers), weather

[33] RAs communicated with the drivers in English, as nearly all Singaporeans can speak the language.
[34] Of note, in the treatment conditions, the passenger states that the driver should take the fastest route. Without this, one could argue that drivers in the treatment conditions take a longer route because that is what they perceive to be the passenger's preferred route. Therefore, I would not necessarily be able to attribute rides on the East Coast Parkway to driver deception.
[35] Most Singaporeans refer to the East Coast Parkway as the "ECP", so the RAs further indicate their unfamiliarity of the city by saying "East Coast Parkway" in full.

conditions (rainy or fair), whether there was a traffic jam, and any extra surcharges or discounts in the fare. Additionally, while the RAs were instructed not to engage the drivers in unnecessary conversation, they discretely recorded any questions or conversation the drivers initiated during the ride.

4.3 Contribution & Relation to Existing Literature

Study 3 contributes to the literature by evaluating how customer misinformation regarding the optimal product/service available affects professional deception and consumer outcomes in a COI context, specifically in a relatively low cost and high use credence good market. I evaluate this both when the consumer explicitly signals that she trusts the professional, and when she does not. To the best of my knowledge, neither customer misinformation regarding the optimal good nor customer trust signals has been studied using a field experiment in a COI context. The study provides a clean, one-shot setting (i.e., taxi rides) where reputational concerns are limited. In this way, I can obtain pure estimates of professional deception.

Next, I discuss the two papers most similar to the present study. Conceptually, Study 3 is similar to research by Busse et al. (2017), who evaluate the effect of customer price misinformation on car repair quotes. While certainly important, my work differs from this past research in several important ways. First, I examine customer misinformation regarding the optimal good instead of price misinformation. In many contexts, consumers don't communicate price information but rather communicate knowledge of what product or service is best for them. For instance, a customer may ask her doctor if a certain medication is right for her, but it's less likely that the customer would ask if the medication had a price of $X. Relatedly, in many cases, professionals may not even be able to change the offered price, which could be fixed by outside forces (such as pharmaceutical companies or the stock market). Third, when shopping around for car repair quotes, consumers can always choose to walk away if the price quote is too high, and thus professionals must balance giving a high quote (and potentially earning a higher revenue) with losing the business altogether. With taxi services, customers typically do not receive a price quote upfront, so professionals may feel they can more easily get away with cheating. Finally, it is not clear whether the professionals providing the price quotes in Busse et al. (2017) are the ones who would reap the benefits of an inflated price, whereas the taxi drivers in my study directly receive the gains from their deception. Differences aside, Busse et al. (2017) make the important distinction of misinformed versus uninformed customers and find that outcomes can vary significantly between the two. This point is germane to the discussion of how the current study differs from another experiment by Balafoutas et al. (2013).

In Balafoutas et al. (2013), the authors experimentally evaluate the taxi market in Athens, Greece. In terms of implementation, Study 3 is very similar to that used by Balafoutas and colleagues. However, the research questions evaluated differ between studies. In Balafoutas et al. (2013), the authors are interested in the effect of perceived passenger income and familiarity with the city (passengers could be uninformed foreigners or informed locals) on the incidence of overtreatment. They find that uninformed foreign passengers are more likely to be overtreated than local passengers.[36] However, the authors do not evaluate the effect of passenger *mis*information or trust on overtreatment propensity. As Busse et al. (2017) show that it is

[36] In their study, uninformed passengers told the driver their destination (in English) and then asked if the driver knew the location as the passenger claimed to not be familiar with the area. Local passengers just told the driver their destination, in Greek.

important to distinguish between uninformed and misinformed customers, Study 3 contributes knowledge beyond that offered by Balafoutas et al. (2013).

4.4 Results

In total, ten RAs took 312 taxi rides in 156 trips over the span of 11 months.[37] For the analysis, I excluded 26 rides where the taxi driver had already been a part of the study, resulting in 286 rides used.[38] All except seven taxi drivers were male, which is consistent with the demographics of Singapore's taxi drivers. As expected, all of the rides ended at the requested destination, so there is no evidence of undertreatment in the data. Below, I verify that taking the ECP is indeed a form of overtreatment. That is, I check that rides on the ECP took longer and were more costly than their PIE counterparts in the data. Subsequently, I evaluate whether the incidence of overtreatment (i.e., ECP rides) varies by treatment.

4.4.1 Suggestive Evidence that ECP Rides are Longer and More Expensive than PIE Rides

Before delving into a more robust statistical analysis, I provide some model-free evidence that ECP rides are longer and more expensive than PIE rides. First, Table 18 provides the average ride statistics, broken down by condition and route. Across all three conditions, ECP rides are longer and more expensive than PIE rides, on average.

[37] All rides were taken on a weekday between 9:00am and 4:30pm to avoid heavy traffic times.
[38] I used license plate numbers to determine whether the driver had already taken part in the study. In the case where a driver provided multiple rides, I used the first ride in the analysis but dropped subsequent rides.

Table 18: Average Duration and Fare of PIE and ECP Rides

	Duration (Minutes)		Fare (S$)	
	PIE	**ECP**	**PIE**	**ECP**
Control	22.7	23.0	18.5	20.6
Misinfo	22.7	26.6	18.4	19.9
Misinfo + Trust	22.4	25.6	18.5	21.6
Overall	**22.6**	**25.8**	**18.5**	**20.9**

Next, I plot the distributions of ride duration and fare, broken down by route, in Figure 15. As seen from the figure, the distributions of ECP rides for duration and fare both lie to the right of the distributions for PIE rides. Further, a t-test of the data shows that ECP rides are significantly longer than PIE rides (25.84 vs. 22.61 min.; $p \sim 0.000$). Likewise, ECP rides are more expensive than PIE rides ($20.87 vs. $18.45; $p \sim 0.000$). This provides suggestive evidence that, overall, rides on the ECP were longer and more costly than PIE rides. However, it's likely that other factors such as cab company, RA gender and weather affect ride fare and duration, and thus should be taken into account. Consequently, below I control for these variables using regression analysis.

Figure 15: Distribution of Ride Durations and Fares by Route

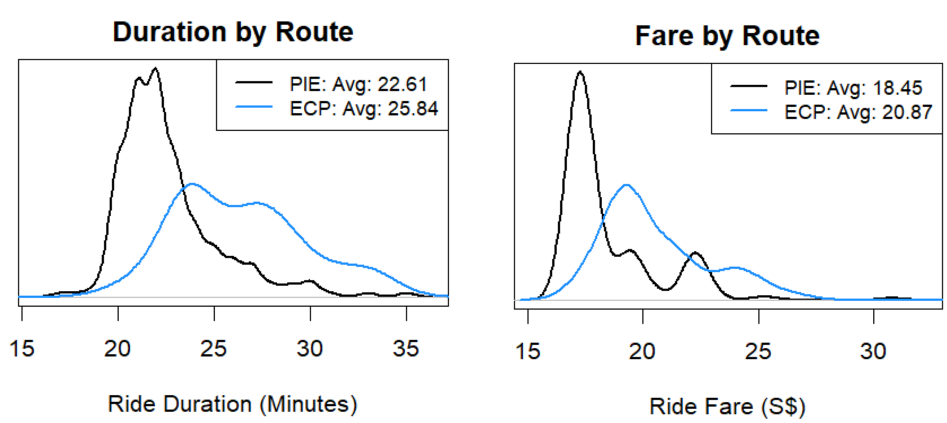

4.4.2 Regression Evidence that ECP Rides are Longer and More Expensive than PIE Rides

I begin with an examination of ride duration. I perform OLS regression analysis on the data, using ride duration (in minutes) as the dependent variable. The main independent variable of interest is a dummy variable for ride route (equaling 1 if driver took ECP; 0 if driver took PIE). I also control for the following observable factors: condition (Control, Misinfo or Misinfo + Trust), weather (rainy or not), traffic (heavy or normal), time of day, cab company (largest cab company in Singapore, "Comfort", or other) and approximate driver age (above or below 45-years-old). I cluster the errors at the RA-level to control for within-RA correlation and include fixed effects for RA gender and nationality. The regression results are shown in column 1 of Table 19. I find that, controlling for condition, ECP rides are longer than PIE rides by an average of 3.3 minutes ($p \sim 0.000$). Similarly, I examine ride fare using the same regression except using fare as the dependent variable. Again, I find ECP rides are more expensive than PIE rides (column 2, Table 19). In particular, rides on the ECP were $2.25 more expensive than rides on the PIE, on average ($p \sim 0.000$). In Tables A1 and A2 of the Study 3 Appendix, I display results of the same analysis, except looking within condition. I find that, in the Misinfo and Misinfo + Trust conditions, ECP rides are longer than PIE rides by an average of 3.91 and 2.80 minutes, respectively (both p's < 0.01). Similarly, ECP rides are more expensive than PIE rides by an average of $1.57 and $2.76 in the Misinfo and Misinfo + Trust conditions, respectively (both p's < 0.01).

Table 19 | OLS estimation results showing ECP rides are longer and more expensive than PIE rides

	Dependent Variable	
	Ride Duration (Minutes) (1)	**Ride Fare (S$)** (2)
ECP	**3.295***** **(0.428)**	**2.251***** **(0.406)**
Misinfo	0.302 (0.316)	-0.132 (0.191)
Misinfo/Trust	0.183 (0.475)	0.297 (0.390)
Rainy	1.462** (0.721)	-0.180 (0.433)
Traffic Jam	0.262 (0.701)	0.323 (0.413)
Driver Age 45+	0.218 (0.306)	-0.161 (0.278)
Comfort	-0.805*** (0.279)	-0.466** (0.202)
10:00-11:29	-0.089 (0.462)	0.740** (0.303)
11:30-12:29	-0.389 (0.567)	-0.017 (0.353)
12:30-14:29	0.770 (0.561)	1.242* (0.659)
14:30-16:30	2.229*** (0.407)	0.710** (0.304)
Female RA	0.261 (0.329)	0.153 (0.331)
Vietnamese RA	-0.385 (0.248)	0.018 (0.374)
Constant	22.010*** (0.292)	18.141*** (0.460)
Observations	286	286
R^2	0.271	0.174
Adjusted R^2	0.237	0.135
F Statistic	7.797*** (df = 13; 272)	4.409*** (df = 13; 272)

*Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors shown in parentheses. Errors clustered at RA-level. Reference categories are PIE, Control, fair weather, no traffic jam, driver age < 45, not Comfort cab company, ride time 9:00-9:59, male RA, Chinese RA.*

Taken together, the above results provide evidence that drivers who take ECP are overtreating passengers. In turn, ride fare and duration are higher than necessary and market efficiency decreases. Having established how drivers overtreat in my context (i.e., by taking ECP), below, I examine whether drivers' propensity to overtreat varies across condition.

4.4.3 Customer Misinformation Signals Increase Overtreatment, With and Without Trust

Table 20 shows the breakdown of which route – PIE or ECP – drivers took across conditions.[39] In the Control, only 3 drivers (4%) took ECP. A test of proportions reveals that this proportion does not differ from zero ($p > .1$). Thus, H1 does not hold and I conclude that, even when passengers are uninformed foreigners, they are not overtreated by taxi drivers in Singapore. This finding differs from the results in Balafoutas et al. (2013), and I will raise a possible explanation for this in the discussion section. In the Misinfo condition, drivers took ECP 17% of the time, which is significantly higher than the 4% of drivers who took ECP in the Control ($p = 0.01$). This suggests that H2 holds, and drivers do overtreat customers who signal misinformation more than customers who do not signal misinformation. Furthermore, drivers took ECP 28% of the time in the Misinfo + Trust condition. This proportion is higher than that in the Control ($p \sim 0.000$), and marginally higher than that in the Misinfo condition ($p = 0.086$). Interestingly, this suggests that H3b holds, and signaling trust alongside misinformation increases drivers' propensity to overtreat, or at the very least, does not decrease the likelihood of overtreatment. While this analysis provides suggestive evidence that professionals cheat customers more when customers signal misinformation, below I turn to a more robust regression analysis of the data.

---

[39] All drivers took either ECP or PIE to reach the destination.

Table 20: Driver Route by Condition

	Driver took ECP	Driver took PIE	Total
Control	**3 (4%)**	77 (96%)	80
Misinfo	**18 (17%)**	90 (83%)	108
Misinfo + Trust	**27 (28%)**	71 (72%)	98

The following regression analysis explores whether there are differences in drivers' propensity to overtreat across treatments, after controlling for observables. I use a linear regression to explore this question.[40] The dependent variable is a binary variable equaling 1 if the driver took ECP; 0 otherwise. My main independent variable of interest is a categorical variable for condition. In Table 21 below, I present the results using the Control condition as a baseline.[41] The analysis controls for the same variables as in Table 19 above, and I cluster the standard errors at the RA-level. Finally, I estimate a series of nested models which allows me to examine the robustness of the estimation results in relation to different model specifications.

[40] The corresponding logistic regression output can be seen in Table A3 of the Study 3 Appendix, and it is consistent with the linear analysis. As I am not using my analysis for prediction purposes, I present the OLS results in the main text because they are simpler to interpret.
[41] In Table A4 of the Study 3 Appendix, I present the results when I instead use the Misinfo condition as a baseline.

Table 21: OLS Regression Showing Overtreatment Higher in Treatment Conditions

	DV: Driver Took ECP = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
Misinfo	0.130***	0.128***	0.130***	0.118***	0.117***	0.119***
	(0.026)	(0.025)	(0.025)	(0.024)	(0.025)	(0.026)
Misinfo + Trust	0.239***	0.238***	0.255***	0.251***	0.250***	0.246***
	(0.052)	(0.054)	(0.053)	(0.059)	(0.059)	(0.068)
Rainy		0.039	0.041	0.027	0.026	0.030
		(0.093)	(0.094)	(0.094)	(0.095)	(0.099)
Traffic jam			0.149	0.147	0.150	0.151
			(0.111)	(0.111)	(0.107)	(0.111)
Driver 45+				-0.085**	-0.086**	-0.077**
				(0.035)	(0.036)	(0.035)
Comfort					0.013	0.003
					(0.046)	(0.044)
10:00-11:29						-0.021
						(0.065)
11:30-12:29						-0.100
						(0.097)
12:30-14:29						0.036
						(0.063)
14:30-16:30						-0.070
						(0.071)
Female RA	-0.007	-0.006	-0.020	-0.021	-0.021	-0.016
	(0.028)	(0.027)	(0.025)	(0.026)	(0.026)	(0.028)
Vietnamese RA	-0.008	-0.007	-0.019	-0.018	-0.018	0.0001
	(0.025)	(0.025)	(0.025)	(0.027)	(0.027)	(0.026)
Constant	0.044**	0.039*	0.028	0.096**	0.090**	0.108**
	(0.019)	(0.023)	(0.023)	(0.046)	(0.044)	(0.043)
Observations	286	286	286	286	286	286
R^2	0.063	0.064	0.077	0.088	0.088	0.101
Adjusted R^2	0.049	0.047	0.057	0.065	0.062	0.061
F-Statistic	4.70*** (4, 281)	3.821*** (5, 280)	3.880*** (6, 279)	3.811*** (7, 278)	3.335*** (8, 277)	2.55*** (12, 273)

Note: *p<0.1; **p<0.05; ***p<0.01

To begin, note that the results across the different regression specifications are relatively consistent (Table 21). Thus, in my discussion below, I refer to the results using the full model (column 6). One can see that the coefficient on the Misinfo variable is positive and significant. In particular, I find that, all else equal, the probability that drivers in the Misinfo condition take ECP is 11.9% higher than for drivers in the Control. Furthermore, the coefficient on the Misinfo + Trust variable is also positive and significant. Its interpretation is as follows: the probability that drivers in the Misinfo + Trust condition take ECP is 24.6% higher than for drivers in the Control. These findings are consistent with the Table 20 statistics and provide evidence in support of H2. Specifically, customers who signal misinformation on the optimal route to take are more likely to be overtreated than passengers who do not signal misinformation. Finally, I compare between the Misinfo and Misinfo + Trust conditions to see whether overtreatment varies when misinformed customers signal trust versus when they do not (Table A4, Study 3 Appendix). I find that the coefficient on the Misinfo + Trust variable is positive and marginally significant ($p = 0.059$).[42] Results can be interpreted as follows: the probability that drivers in the Misinfo + Trust condition take ECP is 12.6% higher than for drivers in the Misinfo condition. This suggests that signaling trust does not help customers who signal misinformation (i.e., H3a does not hold). Further, this provides some evidence that signaling trust on top of misinformation can actually backfire and increase overtreatment beyond the level present when customers just signal misinformation (i.e., support for H3b). I discuss the implications of the findings in Section 4.5 below.

4.5 Summary & Discussion

[42] Note that for three of the regression specifications in the table (those in columns 3-5), the p-value is significant at the 5% level.

Study 3 demonstrates that, when passengers signal misinformation in a taxi setting, this increases drivers' propensity to overtreat whereby drivers take longer and more expensive routes. Additionally, the data suggests that signaling trust, in addition to misinformation, does not reduce driver bias and, in fact, may even increase overtreatment rates. Overall, the Study 3 findings suggest that customers should not share information regarding which product or service they need with professionals unless they know for sure the information is correct. Note that the significant overtreatment effects found in the field context arise after passengers simply added one extra phrase when communicating with their drivers. This speaks to the power of consumer signaling in COI situations. Drawing on insights presented in Dulleck and Kerschbamer (2006), I posit that drivers perceive a reduced risk of their deception being detected when the passenger is inaccurately informed and consequently, respond by increasing their propensity to act opportunistically. Indeed, professionals may perceive an even greater reduction in risk if customers signal trust, as this may be interpreted as a lack of consumer monitoring. Thus, customers should be cognizant of how their communications will be interpreted by professionals, as professionals' beliefs are what ultimately drive their behavior.

Recall that in Balafoutas et al. (2013) the authors find that uninformed passengers are overtreated more than local passengers. I do not find a consistent result in my data, as passengers in the Control are not overtreated. One possible explanation for this discrepancy is that the nature of cab drivers in Singapore differs from that of drivers in Athens, Greece. For example, in 2018, Transparency International ranked Singapore as the 3rd least corrupt country in the world; Greece was ranked 67th in this corruption index. It's possible that Singaporean drivers are more honest overall, and that is why I found they do not take advantage of uninformed passengers.

This highlights the importance of understanding how cultural norms and geographical differences can affect field study outcomes and suggests Study 3's findings are relatively conservative estimates that may be larger if the study were ran in a different country.

Finally, while I believe the present research presents important insights and implications, there are several limitations of the work. First, some might argue that the increase in ride duration and fare due to overtreatment is trivial. While I do examine a low-cost context, one should keep in mind the huge volume of taxi rides being taken on any given day in any given location, meaning that, in aggregate, overtreatment can lead to considerable market inefficiencies. Second, because I study a one-shot context, it is not clear how my results would generalize to other contexts where professionals and advisees have more long-term relationships. In this case, it's possible that reputational motivations may help to deter professional deception in COI situations. However, since past research has demonstrated that COIs do induce professional bias across contexts more commonly thought of as producing long-term relationships, such as healthcare and financial advising, I believe there is a good chance that Study 3's results can be generalized outside of one-shot interactions.[43] Third, I am not able to completely rule out an alternative explanation for the increase of ECP rides found in the two treatment conditions. One could argue that, just by having the passenger mention "ECP" in

[43] The market for taxi rides provides an ideal context to evaluate the effect of customer misinformation on driver deception, as it is common for taxi passengers to be foreigners who may not be correctly informed, and it is also common for drivers to cheat customers. Further, the context is easily accessible given the ubiquity of taxis and the relatively low cost and short duration of taxi rides in general. However, I acknowledge that the results of Study 3 may not hold in other contexts more commonly thought of as COI situations, such as healthcare or financial advising. In these situations, customers generally build up a relationship with their advisor over time and the expectation of ethical professional behavior is stronger. Of course, it is important to understand how customer misinformation affects outcomes across various COI and credence good contexts. I believe taxi markets are a good context to begin an exploration of customer misinformation, given the points specified above. However, future research that examines how the effect of misinformation changes across contexts is certainly worthwhile.

his/her communication to the driver, the driver might take that route to avoid arguing with the passenger. However, in both treatment conditions, the passenger's communication was phrased in such a way that the driver should understand the passenger is just looking to take the fastest route, regardless of whether that's the ECP or an alternative route. Thus, I believe the driver could always choose to take PIE in the treatment conditions, without appearing to argue with the taxi passenger. Finally, at first glance it might appear that the effects of misinformation found in Study 3 are not particularly surprising. That is, why is it surprising or interesting that professionals cheat customers who are misinformed? My response is that the research is interesting because customers *do* choose to share information with professionals, even when they are not certain that this information is correct. Because of this, it's important to understand how professionals respond to customer misinformation. To empirically demonstrate that customers tend to overshare information, I run a single-condition lab experiment (Study 4) in the context of COIs. I present Study 4 in Section 5 below.

5. Customer Information Sharing (Study 4)

In Study 4, I provide evidence that customers share their information with professionals, even when they aren't sure that the information is correct. That is, customers overshare their information with experts who have COIs. For instance, patients often reveal their findings from a WedMD search when they visit a doctor and customers can be quick to share a friend's idea of what might be wrong with their vehicle when visiting the mechanic. Due to the large amounts of information available to customers nowadays and a tendency for customer expertise to be relatively low in a certain area when they seek out professional assistance, it comes as no surprise that customers often stumble across inaccurate information that they mistake as being credible. They might share this information with professionals in an effort to accelerate the speed of diagnosis or even to (mistakenly) signal expertise to deter professional bias. To provide empirical evidence that advisees often overshare their information with advisors in a COI context, I run Study 4, a single-condition lab experiment.

5.1 Theory & Equilibrium Prediction

Consider a better-informed advisor and lesser-informed advisee who interact in a COI setting. The advisor will ultimately make a decision on behalf of his client, the advisee, and this decision (i.e., choosing between Option A or Option B) will determine payoffs for both players. Payoffs also depend upon the game state, which is randomly determined. There are two state types, $t \in \{l, h\}$. Furthermore, there are three unique states, $\theta \in \{\theta^l, \theta_1^h, \theta_2^h\}$, which occur with probability q, $(1-p)(1-q)$ and $p(1-q)$, respectively. Note that the probability of state type h occurring (either θ_1^h or θ_2^h) is $(1-p)(1-q) + p(1-q) = 1-q$.

The advisor observes state type t with probability one. Given $t = l$, the unique state (θ^l) is deterministically known to the advisor. However, conditional on $t = h$, the advisor does not know the unique state (θ_1^h or θ_2^h) for sure. The advisee also observes t, however the advisee's information is probabilistic. That is, the advisee is told the correct type t with probability $p > 0$, and is thus incorrectly informed with probability $(1-p)$. For instance, assume $t = l$. Then, the advisee is (correctly) informed that $t = l$ with probability p. Conversely, the advisee is (incorrectly) informed that $t = h$ with probability $(1-p)$. The advisee never knows whether her information is accurate or not. Note that, given $t = h$, the advisee's information regarding t

determines whether the unique state is θ_1^h or θ_2^h (see Figure 16). That is, if the advisee is correctly informed that $t = h$, then the unique state is θ_2^h. Conversely, if the advisee is incorrectly informed that $t = l$, then the unique state is θ_1^h.

Figure 16: Determination of Unique State

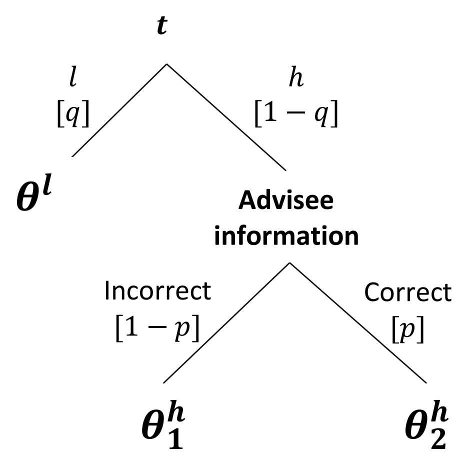

The context can be thought of as a taxi driver (advisor) choosing which route (Option A or B) to take his passenger (advisee). In some instances, the route that earns the driver a relatively high fare is also the route that the passenger would prefer to take, regardless of whether the passenger is a local or non-local. This can be thought of as unique state θ^l, in which the driver does not have a COI as both parties' incentives are aligned (more details below). However, in other instances, the route that earns the driver a relatively high fare could be different from the passenger's optimal route. In this case, the two parties' incentives are misaligned, representing a COI situation similar to that found in Study 3. Here, the driver may prefer to take a longer route only if he thinks he can do so without being detected (e.g., if the passenger is foreign and not knowledgeable about local taxi routes), and this is represented by unique state θ_1^h. If the passenger is instead a local who would likely be able to detect driver cheating, then the driver may prefer to take the most direct route to avoid the possibility of being caught cheating (state θ_2^h).

Recall that both the unique state and the advisor's choice between Option A and Option B determine players' payoffs. I elaborate on this now. First, consider state θ^l, which exists with probability q. The players' payoffs from each of the two options in θ^l can be seen in Table 22. Advisor payoffs are indicated by "α" and advisee payoffs are indicated by "β". Assume that $\alpha^X > \alpha^Y > 0$ and $\beta^X > \beta^Y > 0$. Thus, when the state is θ^l, the players' incentives are aligned and both earn a higher payoff from Option A. Referring to the taxi analogy above, one can think of state θ^l as representing a situation where the most direct taxi route (optimal for passenger) also maximizes the driver's expected earnings, and this is true regardless of whether the passenger is a local or foreigner. As such, given state θ^l, advisee information (whether accurate or not) does not impact payoffs.

Table 22: Payoffs in State $\boldsymbol{\theta^l}$

	Option A	**Option B**
$\pi_{Advisor}$	α^X	α^Y
$\pi_{Advisee}$	β^X	β^Y

Next, I discuss θ_1^h. Recall that state θ_1^h arises when $t = h$ and the advisee is incorrectly informed (i.e., advisee told $t = l$ instead of $t = h$), which occurs with probability $(1 - q)(1 - p)$. See Table 23 below for players' payoffs in state θ_1^h. Here, the players' incentives are no longer aligned: while the advisor still earns a higher payoff from Option A, the advisee now earns a higher payoff from Option B. Referring to the taxi analogy, one can think of θ_1^h as representing a situation similar to that found in Study 3, where the passenger is a non-local and the taxi driver can choose between 2 routes, 1 of which is direct and preferred by the passenger (Option B) and one of which is longer but earns the driver a higher fare (Option A). As the advisee is incorrectly informed, she can be thought of as a non-local who is not able to detect when the advisor takes a longer and more expensive route (Option A). That is, the advisor does not incur a detection cost or suffer any penalties for choosing Option A given θ_1^h, even though Option A is not the advisee's best option. Thus, in both state θ^l and θ_1^h, the advisor would maximize his profits by choosing Option A, as $\alpha^X > \alpha^Y$. However, recall that the advisor only knows t (not the unique state) with probability one, and thus, given $t = h$, the advisor cannot distinguish between θ_1^h and θ_2^h using only his default level of information in the game.

Table 23: Payoffs in State θ_1^h

	Option A	**Option B**
$\pi_{Advisor}$	α^X	α^Y
$\pi_{Advisee}$	β^Y	β^X

Finally, consider state θ_2^h. Recall that state θ_2^h occurs when $t = h$ and the advisee is correctly informed of this. That is, θ_2^h occurs with probability $p(1-q)$. See Table 24 below for the players' exact payoffs given θ_2^h. From the table, one can see that the payoffs are identical to those found in θ_1^h, except for one difference. Specifically, when the state is θ_2^h, the advisor now incurs a cost $k > 0$ when he chooses the advisee's worse option, Option A. Thus, as long as k is sufficiently large, i.e. $k > (\alpha^X - \alpha^Y)$, the advisor would choose Option B, which is also Player B's best option. Referring to the taxi analogy, one can think of θ_2^h as representing a situation similar to that found in Study 3, where the driver makes a decision between two routes. However, here the passenger can be thought of as a local who would be able to detect when the driver took the longer route. Thus, if the driver chooses to take a longer route (Option A), he will still earn the higher fare (α^X) but now he will incur a cost, k, of being detected by the local passenger. This detection cost can be thought of as a reputational penalty from cheating the passenger or a monetary fee (e.g., driver forced to refund passenger a portion of the fare). In this game, assume that k is sufficiently large, i.e. $k > (\alpha^X - \alpha^Y)$. In this case, the driver prefers to choose Option B in state θ_2^h in order to avoid incurring the detection cost, k.[44]

Table 24: Payoffs in State $\boldsymbol{\theta_2^h}$

	Option A	**Option B**
$\pi_{Advisor}$	$\alpha^X - k$	α^Y
$\pi_{Advisee}$	β^Y	β^X

As stated above, the advisor will choose between Option A or Option B, which will determine payoffs for both players. When type $t = l$, the advisor always earns more from Option A, and the same goes for the advisee. In this case there is no COI, and the advisor's decision is straightforward – he chooses Option A. However, consider the case where $t = h$, and thus the advisor cannot distinguish between unique state θ_1^h and θ_2^h. This makes the advisor's decision more difficult, as he prefers to choose Option A in state θ_1^h and Option B in state θ_2^h. Thus, it is in the advisor's best interest to figure out whether the advisee is correctly informed or not. This is

[44] The payoffs in state θ_2^h help to illustrate why professionals with COIs do not cheat all the time. That is, even though they could earn short-term gains from acting opportunistically, there are potentially costs that outweigh these benefits.

because, if the advisor knows whether the advisee is correctly informed, then he can determine the unique state with probability 1.

In this game, I allow the advisee to share her information with the advisor.[45] The advisee can share her information for cost $c > 0$ and, if the advisee chooses to share, she is not able to lie about her information.[46] If the advisee shares her information, the advisor can perfectly determine the unique state. For instance, consider the case where the advisor observes $t = h$ (which he knows is accurate). Now assume that the advisee shares her information with the advisor. If the advisee's information reveals she was (correctly) told that $t = h$, then the advisor knows with certainty that the state is θ_2^h and he is predicted to choose Option B. This is because the advisor can compare his information with the advisee's, and if it matches then he knows the advisee is correctly informed (see Figure 16 above). Alternatively, if the advisee shares information that reveals she was (incorrectly) told that $t = l$, then the advisor knows with certainty that the unique state is θ_1^h and he is predicted to choose Option A.

While the advisor benefits from knowing the advisee's information, the advisee may also benefit from sharing her information. Consider the case where, given the advisee does not share her information, the advisor will always choose Option A. Then the advisee may have incentive to reveal her information if doing so will induce the advisor to choose Option B, which is the advisee's best option given $t = h$. However, it is also risky for the advisee to share her

[45] If the advisee shares her information with the advisor, one can think of this as a taxi passenger revealing to the driver, say, whether she is a local or non-local or whether she knows the fastest route to take.

[46] I include this information sharing cost to obtain conservative estimates of how often advisees share their information with advisors, as in reality, it is virtually costless for advisees to do so.

information, as in some cases, it might reveal she is misinformed which could lead the advisor to choose the option that is not optimal for the advisee.

In reality, this can be thought of as a standard COI situation, where the advisor's optimal decision depends on the expertise of the advisee. That is, if the advisee is well-informed and would thus be able to detect when the advisor is deceiving her, then the advisor is better off being honest and choosing the option that's in the best interest of the advisee. However, if the advisee is not well informed and thus may not be able to detect advisor deception, then the advisor prefers to choose his best option, even though this may not be preferred by the advisee.

Table 25 shows a breakdown of the game states and probabilities. From the table, one can see that the advisee observes $t = l$ with probability $1 - p - q + 2qp$, and $t = h$ with probability $q + p - 2qp$. The exact move sequence in the game is described in Section 5.2 below, where I begin to refer to the advisor as "Advisor" and the advisee as "Advisee".

Table 25: Breakdown of Game States and Probabilities

$\boldsymbol{t}$	Advisee's Information Accuracy	Advisee Information	Unique State	Probability
$\boldsymbol{l}$ $[q]$	Correct (p)	$t = l$	θ^l	qp
	Incorrect $(1 - p)$	$t = h$	θ^l	$q(1 - p)$
$\boldsymbol{h}$ $[1 - q]$	Correct (p)	$t = h$	θ_2^h	$(1 - q)p$
	Incorrect $(1 - p)$	$t = l$	θ_1^h	$(1 - q)(1 - p)$

5.2 Move Sequence

1. Nature randomly determines type t and Advisee information accuracy (thus, unique state is also randomly determined)

2. Advisee observes t with probability of accuracy = p
    a. Chooses whether to share private information with Advisor for cost $c > 0$
3. Advisor observes t with probability of accuracy = 1 and also observes Advisee's information if Advisee shares this
    a. Chooses between Option A or B
4. Payoffs are determined

5.3 Equilibrium Prediction

The extensive form of the game is shown in the Study 4 Appendix (Figure A1).[47] To solve for the game equilibrium, I assume both players are self-interested. I note that Advisor's decision between Option A or Option B will depend on Advisor's expected payoff from each option which, in turn, depends on the unique state. To solve for equilibrium behavior, I check 4 different cases of Advisee's information sharing to see whether each is an equilibrium.

Case 1: Advisee never shares $\boldsymbol{t = l}$ and never shares $\boldsymbol{t = h}$. To begin, note that, given Advisor observes $t = l$, his optimal decision does not depend on Advisee's information. Rather, Advisor always chooses Option A. This is because, conditional on $t = l$, Advisor's payoff from Option A (α^X) is higher than his payoff from Option B (α^Y). Next, I evaluate Advisor's decision when he is instead told $t = h$. Recall that, in this situation, Advisor's payoff will depend on whether Advisee is correctly informed. In Case 1, Advisee doesn't share any information with Advisor, so Advisor is not able to deduce whether the unique state is θ_1^h or θ_2^h (because Advisor cannot tell whether

[47] In the figure, Advisor decision nodes are marked as "A" and Advisee decision nodes are marked as "B", for simplicity sake.

Advisee is correctly informed or not). Thus, Advisor's best response will depend on his expected payoff from each of the two options, as shown below in equations (8) and (9).

$$E\left[\pi_{Advisor}^{\boldsymbol{Option\ A}} | t = h, No\ sharing\right] = p(\alpha^X - k) + (1-p)(\alpha^X) = \alpha^X - pk \qquad (8)$$

$$E\left[\pi_{Advisor}^{\boldsymbol{Option\ B}} | t = h, No\ sharing\right] = p(\alpha^Y) + (1-p)(\alpha^Y) = \alpha^Y \qquad (9)$$

Advisor's expected payoff from choosing Option A is greater than his expected payoff from choosing Option B when $k < \frac{\alpha^X - \alpha^Y}{p}$. Intuitively, Advisor would act opportunistically (i.e., choose Option A) as long as he would incur a sufficiently small detection cost if his cheating is detected. Assume this condition holds. Then, given Advisee's strategy in Case 1, Advisor would always choose Option A, resulting in the following expected payoffs.

$$\pi_{Advisor} = q(\alpha^X) + (1-q)(p)(\alpha^X - k) + (1-q)(1-p)(\alpha^X) \qquad (10)$$

$$\pi_{Advisee} = q(\beta^X) + (1-q)(\beta^Y) \qquad (11)$$

To check whether Case 1 is an equilibrium, I evaluate whether Advisee has an incentive to deviate. If Advisee has any incentive to deviate, then Case 1 is not an equilibrium. Consider the alternate strategy where Advisee does not share $t = l$ but does share $t = h$. In this case, Advisor's off-equilibrium beliefs kick in when Advisee shares $t = h$. Specifically, given Advisee shares $t = h$, then Advisor's belief of the probability that Advisee observed $t = l$ is 0 (this is because Advisee is not permitted to lie about private information). Then, when Advisor knows $t = h$ and Advisee shares $t = h$, Advisor would choose Option B. This is because Advisee is correctly informed, so Advisor knows the unique state is θ_2^h and thus Advisor's payoff from

Option B [α^Y] is higher than his payoff from Option A [$\alpha^X - k$]. Alternatively, when Advisor knows $t = l$ and Advisee shares $t = h$, Advisor still chooses Option A (because Advisor always chooses Option A when state is θ^l).

Consider the right-hand side information set of Advisee (where Advisee observes $t = h$) in the extensive form game shown in Figure A1 of the Study 4 Appendix. If Advisee does not deviate and instead plays the strategy as outlined in Case 1 (no information sharing), then Advisee's expected payoff is shown in equation (12).

$$\left[\beta^X * \frac{q(1-p)}{q(1-p)\ +(1-q)p}\right] + \left[\beta^Y * \frac{p(1-q)}{q(1-p)\ +(1-q)p}\right] \tag{12}$$

Next, I calculate Advisee's expected payoff from deviating to sharing $t = h$. In this case, Advisee's expected payoff is

$$\left[(\beta^X - c) * \frac{q(1-p)}{q(1-p)\ +(1-q)p}\right] + \left[(\beta^X - c) * \frac{p(1-q)}{q(1-p)\ +(1-q)p}\right]. \tag{13}$$

When β^X is sufficiently larger than β^Y and the cost of signaling c (specifically, when $\beta^X > \beta^Y + 2c$), Advisee's expected payoff from deviating is higher than her expected payoff from not deviating. Thus, Case 1 is not an equilibrium.

Case 2: Advisee always shares $t = l$ and never shares $t = h$. Again, given Advisor observes $t = l$, he always chooses Option A. Next, I evaluate Advisor's decision given he observes $t = h$. In Case 2, I assume that Advisee only shares her information when she observes $t = l$. Thus, Advisor can back out Advisee's information given her shared information (or lack thereof). In particular, if Advisor observes $t = h$ and Advisee shares $t = l$, then Advisor knows Advisee is misinformed and can deduce the exact state is θ_1^h. Thus, Advisor would choose Option A. Alternatively, if Advisor observes $t = h$ and Advisee does not share any information, then Advisor can infer Advisee observed $t = h$ and is thus correctly informed. In this situation, the unique state is θ_2^h, and Advisor chooses Option B. The outcome of the game is that Advisor only chooses Option B when he observes $t = h$ and Advisee does not share her information; otherwise Advisor chooses Option A. Expected payoffs for each player are shown below.

$$\pi_{Advisor} = q(\alpha^X) + (1-q)(p)(\alpha^Y) + (1-q)(1-p)(\alpha^X) \quad (14)$$

$$\pi_{Advisee} = q(\beta^X) + (1-q)(p)(\beta^X) + (1-q)(1-p)(\beta^Y) - c(1-p-q+2qp) \quad (15)$$

Note that Advisee's payoff is reduced by $c(1-p-q+2qp)$ because she incurs a cost of c to share her information, which she does whenever she observes $t = l$ (see Table 25).

Next, I check whether Advisee has an incentive to deviate. Consider the strategy where Advisee never shares her information. In this case, Advisor will still believe that Advisee's lack of sharing can be inferred as Advisee observed $t = h$. Then, Advisor will still always choose Option A when he observes $t = l$, but now Advisor believes that Advisee is always correctly informed when Advisor observes $t = h$. Consequently, Advisor will always choose Option B when he observes $t = h$ and Option A when he observes $t = l$.

Consider the left-hand side information set of Advisee in the extensive form game (information set when Advisee observes $t = l$). If Advisee plays the strategy as outlined in Case 2 (only sharing $t = l$), then Advisee's expected payoff is given by

$$\left[(\beta^X - c) * \frac{qp}{qp + (1-q)(1-p)}\right] + \left[(\beta^Y - c) * \frac{(1-q)(1-p)}{qp + (1-q)(1-p)}\right]. \quad (16)$$

However, if Advisee deviates and instead does not share information with Advisor, then her expected payoff is

$$\left[\beta^X * \frac{qp}{qp + (1-q)(1-p)}\right] + \left[\beta^X * \frac{(1-q)(1-p)}{qp + (1-q)(1-p)}\right]. \quad (17)$$

Since $c > 0$ and $\beta^X > \beta^Y$, then, by definition, Advisee's expected payoff is higher when she deviates to not sharing any information. Thus, Case 2 is not an equilibrium.

Case 3: Advisee never shares $t = l$ and always shares $t = h$. As in the other cases, given Advisor observes $t = l$, he always chooses Option A. Next, I evaluate Advisor's decision given he observes $t = h$. In Case 3, Advisee only shares information when she observes $t = h$. Knowing this, Advisor can back out Advisee's information given her shared information (or lack thereof). In particular, if Advisor observes $t = h$ and Advisee shares $t = h$, then Advisor knows Advisee is correctly informed and thus the state is θ_2^h. In this case, Advisor would choose Option B. Alternatively, if Advisor observes $t = h$ and Advisee does not share any information, then Advisor can infer Advisee observed $t = l$ and is thus incorrectly informed. In this situation, the unique state is θ_1^h, and Advisor chooses Option A. As such, Advisor only chooses Option B when he observes $t = h$ and Advisee shares her information; otherwise Advisor chooses Option A. Players' expected payoffs are presented in equations (18) and (19).

$$\pi_{Advisor} = q(\alpha^X) + (1-q)(p)(\alpha^Y) + (1-q)(1-p)(\alpha^X) \quad (18)$$

$$\pi_{Advisee} = q(\beta^X) + (1-q)(p)(\beta^X) + (1-q)(1-p)(\beta^Y) - c(q + p - 2qp) \tag{19}$$

Next, I check whether Advisee has an incentive to deviate. First, consider the case where Advisee deviates to the strategy of always sharing her information, both $t = l$ and $t = h$. In this case, Advisor will still always choose Option A when he observes $t = l$. However, now when Advisor observes $t = h$ and Advisee shares $t = h$, Advisor will choose Option B. Otherwise, when Advisor observes $t = h$ and Advisee shares $t = l$, Advisor will choose Option A. The outcome will be the same as in Case 3, except here Advisee will always incur cost c to share information (instead of only incurring this cost when she observes $t = h$), so Advisee does not have an incentive to deviate.

Second, I consider the case where Advisee deviates to the strategy of never signaling. Again, Advisor always chooses Option A when he observes $t = l$. But now when Advisor observes $t = h$ and Advisee doesn't share information, Advisor believes that Advisee observed $t = l$ and is thus misinformed. In this case, Advisor will choose Option A when he observes both $t = l$ and $t = h$. Advisee's expected payoff will be $q(\beta^X) + (1-q)(\beta^Y)$. Assuming c is sufficiently small and β^X is sufficiently large, then Advisee's expected payoff from deviating is lower than Advisee's expected payoff from playing the strategy in Case 3. Thus, Advisee does not have an incentive to deviate.

Finally, I consider the case where Advisee deviates to always sharing $t = l$ and never sharing $t = h$. Advisor will always choose Option A if he observes $t = l$. Now when Advisor observes $t = h$ and Advisee doesn't signal, Advisor assumes Advisee observed $t = l$ and thus deduces that Advisee is misinformed, and the state is θ_1^h. Given this, Advisor will choose Option

A. Similarly, when Advisor observes $t = h$ and Advisee signals $t = l$, Advisor knows the state is θ_1^h and chooses Option A. Consequently, Advisor always chooses Option A and Advisee's expected payoff will be $q(\beta^X) + (1-q)(\beta^Y) - c(1-p-q+2qp)$. Since this is lower than Advisee's expected payoff from her strategy in Case 3, Advisee doesn't have an incentive to deviate. Thus, as Advisee has no incentive to deviate from her strategy in Case 3, I confirm that there is an equilibrium where Advisee only shares $t = h$ and Advisor chooses Option B when he observes $t = h$ and Advisee shares her information; otherwise Advisor chooses Option A.

Case 4: Advisee always shares $t = l$ and always shares $t = h$. When Advisor observes $t = l$, he always chooses Option A. Next, when Advisor observes $t = h$ and Advisee shares $t = l$, Advisor knows Advisee is misinformed and the state is θ_1^h. Thus, Advisor chooses Option A. When Advisor observes $t = h$ and Advisee shares $t = h$, then Advisor knows the state is θ_2^h and he chooses Option B. Thus, Advisor only chooses Option B when he observes $t = h$ and Advisee shares $t = h$; otherwise Advisor chooses Option A. Expected payoffs are shown in equations (20) and (21).

$$\pi_{Advisor} = q(\alpha^X) + (1-q)(p)(\alpha^Y) + (1-q)(1-p)(\alpha^X) \tag{20}$$

$$\pi_{Advisee} = q(\beta^X) + (1-q)(p)(\beta^X) + (1-q)(1-p)(\beta^Y) - c \tag{21}$$

Next, I check whether Advisee has an incentive to deviate to the strategy of never sharing $t = l$ and always sharing $t = h$. Advisor will still always choose Option A when he observes $t = l$. But now when $t = h$ and Advisee doesn't share her information, Advisor's off-equilibrium beliefs kick in and Advisor will choose the option that maximizes his expected payoffs (Option A). When $t = h$ and Advisee shares $t = h$, Advisor chooses Option B. Next, I calculate Advisee's expected payoff from deviating to not sharing $t = l$.

Consider the left-hand side of the extensive form of the game tree. If Advisee plays the strategy as outlined in Case 4 (always sharing information), then Advisee's expected payoff is given by

$$\left[(\beta^X - c) * \frac{qp}{qp + (1-q)(1-p)}\right] + \left[(\beta^Y - c) * \frac{(1-q)(1-p)}{qp + (1-q)(1-p)}\right]. \quad (22)$$

However, if Advisee deviates and instead does not share $t = l$ but shares $t = h$ with Advisor, then her expected payoff is given by

$$\left[\beta^X * \frac{qp}{qp + (1-q)(1-p)}\right] + \left[\beta^Y * \frac{(1-q)(1-p)}{qp + (1-q)(1-p)}\right]. \quad (23)$$

Since Advisee's expected payoff is, by definition, higher when she deviates to only sharing $t = h$, Case 4 is not an equilibrium.

Thus, the equilibrium behavior of the game is that Advisee always shares her information when she observes $t = h$ but never shares her information when she observes $t = l$. In turn, Advisor chooses Option B when he observes $t = h$ and Advisee shares her information; otherwise, Advisor chooses Option A. The expected payoffs of the players are below.

$$\pi_{Advisor} = q(\alpha^X) + (1-q)(p)(\alpha^Y) + (1-q)(1-p)(\alpha^X) \quad (24)$$

$$\pi_{Advisee} = q(\beta^X) + (1-q)(p)(\beta^X) + (1-q)(1-p)(\beta^Y) - c(q + p - 2qp) \quad (25)$$

5.4 Experimental Design & Procedure

Participants were primarily undergraduate students at NUS. I ran 5 sessions, with 58 total subjects. In the experiment, subjects were randomly and anonymously paired at the start of each of 12 decision rounds. Before the start of the 12 decision rounds, there were two practice rounds that carried no monetary consequences to familiarize participants with the experimental procedure. At the end of each experimental session, point earnings were converted to cash earnings at a pre-specified conversion rate, and participants were paid privately in cash. Participants earned $14 on average. The experiment was coordinated using z-Tree software (Fischbacher 2007).

Next, I describe the procedure used for the experiment. Upon entering the lab, participants were seated at separate computer terminals and received a set of instructions from the experimenter. At the beginning of the session, the experimenter went over the instructions as a group, and participants were encouraged to ask questions if they had them. Participants were randomly assigned to the role of Player A (Advisor) or Player B (Advisee) at the beginning of each of the 12 decision rounds. Additionally, I randomly and anonymously matched participants into pairs each round, such that each pair had one Player A and one Player B. As everything in the experiment was anonymous, participants were asked to keep their eyes on their own computer and refrain from talking to anyone but the experimenter.

The game progresses as follows. First, the computer randomly assigns a $Payoff\ Set$ ("PS"), PS1 or PS2. Here, the $Payoff\ Set$ equates to the state type t, with PS1 equating to $t = l$ and PS2 equating to $t = h$. Next, Player B observes the $Payoff\ Set$ – PS1 or PS2 – with 60% accuracy. The probability that Player B is correctly informed is common knowledge. Then, Player B decides whether to share her information with Player A for a cost of 1 point. If Player B decides to share her information, Player A is told the $Payoff\ Set$ that Player B observed along with the true $Payoff\ Set$. Given this, Player A can deduce the exact $Decision\ Set$: DS1, DS2a or DS2b. The $Decision\ Set$ equates to the unique state with DS1, DS2a and DS2b corresponding to θ^l, θ_1^h and θ_2^h, respectively. If Player B does not share her information, Player A only knows whether the $Decision\ Set$ is DS1 or one of either DS2a or DS2b. Finally, Player A chooses between Option A or Option B, and payoffs are determined.

The equilibrium prediction is that Player B always shares her information when she observes PS2, but never shares her information when she observes PS1. In turn, Player A chooses Option B when he observes PS2 and Player B shares this as well; otherwise Player A chooses Option A. Given this, along with the experimental parameter values shown in Table 26 below, the equilibrium payoffs are 37.48 for Player A and 26.66 for Player B. However, since the roles were randomly assigned in each round, the expected payoffs across the 12 rounds are identical for each subject. After every round, participants observed the true $Decision\ Set$ and Player A's option decision, along with point earnings for each player.

Table 26: Terminology and Parameter Values in the COI Game and the Experiment

COI Game	Experiment
$p = 0.6$	Probability Player B is Correctly Informed of $Payoff\ Set$
$q = 0.3$	Probability of $Payoff\ Set\ 1$
$\alpha^X = 40, \alpha^Y = 34$	Player A's Earnings
$\beta^X = 30, \beta^Y = 20$	Player B's Earnings
$k = 9$	Player A's Detection Cost
$c = 1$	Player B's Cost of Information Sharing

Note that it is equilibrium behavior for Player B to share information that may not always be accurate. Specifically, Player B is always predicted to share PS2, even though sometimes the true $Payoff\ Set$ will be PS1. Intuitively, it is okay for Player B to share this because, if her information is inaccurate, Player A will still choose her best option as the players' incentives are aligned. In reality, while advisees may share information with experts for various reasons (e.g., to speed up diagnosis or satisfy curiosity), they may not be aware of the risk of doing so. That is, some advisees might not be aware of their expert's COI in the first place. This is just one reason why advisees may share inaccurate information in the real world. However, as players' incentives are common knowledge in the Study 4 game, I now delineate reasons why Player B may overshare in the experiment.

5.5 Behavioral Hypotheses

The equilibrium behavior derived from standard economic theory does not account for trust, reciprocity or irrational behavior of any kind. Without these things, Player B has no incentive to share her information upon observing PS1 (as predicted by theory). However, consider the fact that Player B may share her information as a way to elicit increased reciprocity from Player A. That is, Player B may believe that sharing her information could be perceived as a signal of trust, which Player A may reciprocate by being more trustworthy. Indeed, the Study 2 results show that certain trust signals can induce greater levels of advisor honesty. Alternatively, Player B may actually believe that Player A will be honest and choose Player B's best option, even in the face of personal incentives not to do so. This is possible, as past research has shown people typically show some degree of trust, even in one-shot settings (e.g., Berg et al. 1995; Ho and Weigelt

2005). That is, trust and expectations of reciprocity may drive Player B to overshare her information. Thus, my primary behavioral hypothesis is that, after observing PS1, Player B will share her information significantly more often than 0% of the time. By the same logic, I believe Player A may also display behavioral preferences towards Player B. Specifically, I evaluate whether Player A always chooses Option A when he observes PS2 and Player B doesn't share her information. While theory predicts this is true (even though this behavior is off equilibrium), I believe Player A will choose Player B's best option in this case (Option B) with non-zero probability.

5.5 Results

Note that I observe no significant differences across all five sessions, so I combine the sessions and analyze the data altogether. Over the five sessions, I collected 348 observations for each Player A and Player B. The following analysis of the Study 4 data is brief, as the game results are relatively straightforward, and the main purpose of the study is to demonstrate that advisees overshare their information in COI situations.

5.5.1 Player B's Decisions & Earnings

I begin with Player B's information sharing decision. First, Player B observed "PS2" 54.7% of the time, which is consistent with the 54% predicted by theory. Recall that the equilibrium prediction is that Player B will always share her information in this case. The results show that Player B shares her information 76.6% of the time, which is significantly lower than the prediction of 100% ($p \sim 0.000$). I also find that Player B obtains significantly higher earnings when she shares her information compared to when she does not (28.2 vs. 23.9; $p \sim 0.000$). This suggests that advisees do not always optimally share their information, and non-optimal information sharing can have a deleterious effect on advisee outcomes. Next, when Player B observes PS1, she is predicted to never share her information. Interestingly, I find that Player B shares her information 27.6% of the time, which is significantly higher than the prediction of 0% ($p{\sim}0.000$). Thus, the data supports my behavioral hypothesis that Player B overshares her information in some situations. Moreover, I confirm that oversharing information can be detrimental for Player B, as she obtains lower average earnings after sharing PS1 compared to not sharing PS1 (22.8 vs. 25.3; $p = 0.004$). This result speaks to the importance of understanding how advisors respond to inaccurate advisee information. While I presented an analysis of this in Study 3 above, I can also evaluate this in the Study 4 game. I do so next by examining Player A's decisions.

5.5.2 Player A's Decisions

First, I find that, when Player A observes PS1, he chooses Option A 97.8% of the time. This is consistent with the equilibrium prediction of 100% and suggests that the players understood the nature of the game. Next, I analyze Player A's option decision conditional on Player A observing PS2 (74% of observations). In this case, when Player B shares PS2 (so Player A knows Player B is correctly informed and the $Decision\ Set$ is DS2b), Player A chooses Option B 91.1% of the time, which is consistent with the theory prediction of 100%. Further, when Player B shares information that is incorrect and Player A can back out that the $Decision\ Set$ is DS2a, Player A chooses Option A 84.5% of the time. This finding is consistent with the Study 3 results. Namely, it suggests that advisors with COIs will act opportunistically, especially when advisees signal that they are misinformed. Finally, when Player A observes PS2 but Player B does not share information, then the equilibrium prediction is that Player A will choose Option A as his expected earnings are higher than from choosing Option B (Player B's best option). I find that Player A chooses Option A 73.9% of the time. That is, 26.1% of the time, Player A chooses Option B ($p \sim 0.000$). This finding is consistent with my behavioral hypothesis that Player A may be driven, at least to some extent, by reciprocal preferences towards Player B. That is, even though advisors may sometimes take advantage of their COIs, they do not always do so.

5.6 Summary & Discussion

Results of Study 4 confirm that lesser-informed customers do not always optimally share their information in COI situations. I find that sometimes when advisees should share their information, they do not. And, more importantly, the opposite is also true: there are times when customers overshare their mistaken information, even when they know advisors may have incentives to act deviously. While there are certainly situations where sharing misinformation may not be detrimental for advisees, the results presented in Studies 3 and 4 suggest that advisors do take advantage of advisee misinformation by acting opportunistically. Taken together, these studies suggest that customers should be wary of oversharing with professionals, especially if customers are not certain of their advisors' incentives or are unsure of the accuracy of their own information. In this case, customers may be better off not sharing their information with professionals, at all.

6. Discussion & Conclusion

This paper examines the relationship between customer signals, advisor requests and the two parties' outcomes in COI settings. Study 1 shows that advisees strategically signal that they trust their advisors to give them honest advice, even when advisors have no incentive to be honest and advisees actually believe that advisors will be dishonest. Further, the results suggest that both trust and low-level distrust signals may be beneficial in certain contexts.

I extend this research in Study 2, where I manipulate the type of trust signal the advisee can send, as well as advisor trust requests. I find that simple statements of trust can reduce advisor bias and improve customer outcomes. However, stronger trust signals (i.e., decision delegation) can backfire and greatly increase advisor bias and decrease advisee earnings. I posit that, while advisors feel greater reciprocity towards advisees under delegation compared to trust statements, they also face a reduced risk of bias which outweighs the increase in reciprocity. Interestingly, I find that the detrimental effect of delegation is suppressed when advisors can request that the advisee delegate to them in the first place. My conjecture is that trust requests increase advisor reciprocity towards advisees even further, which helps to counterbalance the reduced risk of bias under delegation.

Studies 3 and 4 provide a further investigation into customer signals, with a focus on misinformation signals. In particular, Study 3 finds that, when customers share inaccurate information with professionals who have COIs, the professionals are more likely to act opportunistically. This detrimental effect of customer misinformation is important to understand, as Study 4 shows that advisees do choose to share information with advisors, even when they aren't sure this information is accurate. Further, when customer misinformation signals are paired with signals of trust, there is evidence to suggest that advisor deception increases even more. Indeed, this finding echoes the insight from Study 2 that acts of trust can be harmful for advisees. That is, since taxi drivers in Study 3 likely perceive the passenger's trust signal as a delegation of the route choice to the driver, the increase in driver dishonesty found in the Misinfo + Trust condition provides field evidence that acts of trust can be detrimental for advisees (consistent with Study 2). Overall, the findings from Studies 3 and 4 suggest that, unless customers are certain their information is accurate or that their advisors do not have COIs, it may be best for them not to share a guess of which product or service fits their needs.

Looking across the entire set of studies, the key takeaway is that simple signals and requests can have large consequences in COI situations. The results from this paper provide important insights for both customers and professionals. First, customers must understand how their communication will be perceived by experts. While customers may think that trust signals are an innocuous way to advance their relationships with professionals, in fact these signals could have serious consequences on customer outcomes. This is especially detrimental in more high-risk contexts such as healthcare and finance. On a more positive note, the research also provides a tool – simple statements of trust – that customers can use to improve outcomes in COI situations. Trust statements can be beneficial for customers as long as the statements do not imply further meaning (such as a lack of monitoring) to advisors.

Moreover, the research also raises an interesting point for marketing practitioners and firms, especially PSFs. Past work has found that, in contexts with professional service workers, customers rely heavily on social measures, such as reputation, when choosing a service provider (Rao et al. 2001). Thus, it is in a firm’s best interest, not only ethically but also monetarily, to keep its service professionals truthful. To help ensure professionals act honestly, managers can introduce trust requests into professional training programs. That is, advisors can literally ask customers to trust them, as Study 2 has demonstrated that doing so can greatly improve expert honesty, especially under decision delegation. This recommendation is not only virtually cost-free, but also easily implementable across a variety of fields and settings.

I conclude with a discussion of limitations and directions for future research. First, because the noise in Studies 1 and 2 is relatively large, some findings are a bit uncertain and I cannot make definitive claims. Future research could work to improve the clarity of findings, as I believe it is certainly a worthwhile task. Further, in these studies, I do not offer a traditional 2x2 experimental design, which means I do not fully explore all signaling contexts available. As there is much research to be done, my investigation was directed at what I perceived as especially interesting and important contexts, but future work could provide more “clean” experimental designs. Next, while I believe Study 3’s findings can generally be applied to different markets and geographical areas, it is likely the results would differ somewhat across cultures and locations. Indeed, an evaluation of how deception rates vary across locations could provide meaningful insights for firms looking to invest in, for example, Asian countries. Finally, while COI research is important and relevant specifically because COIs exist across various fields and contexts, this precise reason does make it difficult to generalize findings to all COI contexts. For instance, some COI situations are based in more long-term and high-risk contexts (e.g., healthcare, finance), while other contexts are more one-shot and low risk (e.g., taxi services, commissioned salespeople). In the present work, I focus primarily on the latter context, and while I believe the general insights can be extended to other types of contexts, I cannot claim that the results would hold in all COI situations. Therefore, future research that allows advisees and advisors to interact with identity disclosure and/or repeated interaction could generate fruitful insights.

References

Angelova V, Regner T (2013) Do voluntary payments to advisors improve the quality of financial advice? An experimental deception game. *Journal of Economic Behavior & Organization*. 93:205-218.

Arrow K (1974) The Limits of Organization. Norton.

Balafoutas L, Beck A, Kerschbamer R, Sutter M (2013) What Drives Taxi Drivers? A Field Experiment on Fraud in a Market for Credence Goods. *Review of Economic Studies*. 80(3):876-891.

Balafoutas L, Kerschbamer R, and Sutter M (2015) Second-Degree Moral Hazard in a Real-World Credence Goods Market. *The Economic Journal.* 127(599):1-18.

Basu A, Lal R, Srinivasan V, Staelin R (1985) Salesforce Compensation Plans: An Agency Theoretic Perspective. *Marketing Science.* 4(4):267-291.

Beeler J, Hunton J (2002) Contingent Economic Rents: Insidious Threats to Auditor Independence*. Advances in Accounting Behavioral Research.* 5, 21-50.

Bekelman J, Li Y, Gross C (2003) Scope and Impact of Financial Conflicts of Interest in Biomedical Research: A Systematic Review. *JAMA.* 289(4):454-465.

Berg J, Dickhaut J, McCabe K (1995) Trust, Reciprocity, and Social History. *Games and Economic Behavior.* 10(1):122-142.

Bergen M, Dutta S, Walker O (1992) Agency Relationships in Marketing: A Review of the Implications and Applications of Agency and Related Theories. *Journal of Marketing.* 56(3):1-24.

Bhardwaj P (2001) Delegating pricing decisions. *Marketing Sci*. 20(2):143-169.

Bowman M, Pearle D (1988) Changes in Drug Prescribing Patterns Related to Commercial Company Funding of Continuing Medical Education. *The Journal of Continuing Education in the Health Professions.* 8(1):13-20.

Busse M, Israeli A, Zettelmeyer F (2017) Repairing the Damage: The Effect of Price Knowledge and Gender on Auto Repair Price Quotes. *Journal of Marketing Research.* 54(1):75-95.

Cain D, Loewenstein G, Moore D (2005) The Dirt on Coming Clean: Perverse Effects of Disclosing Conflicts of Interest. *Journal of Legal Studies.* 34(1):1-24.

Cain D, Loewenstein G, Moore D (2011) When Sunlight Fails to Disinfect: Understanding the Perverse Effects of Disclosing Conflicts of Interest. *Journal of Consumer Research.* 37(5):836-857.

Caudill T, Johnson M, Rich E, McKinney W (1996) Physicians, Pharmaceutical Sales Representatives, and the Cost of Prescribing. *Arch Fam Med.* 5(4):201-206.

Cegedim Strategic Data (2013) 2012 U.S. Pharmaceutical Company Promotion Spending. http://www.skainfo.com/health_care_market_reports/2012_promotional_spending.pdf

Church B, Kuang X (2009) Conflicts of Interest, Disclosure, and (Costly) Sanctions: Experimental Evidence. *The Journal of Legal Studies.* 38(2):505-532.

Crawford V, Sobel J (1982) Strategic Information Transmission. *Econometrica.* 50(6):1431-1451.

Currie J, Lin W, Zhang W (2011) Patient Knowledge and Antibiotic Abuse: Evidence from an Audit Study in China. *Journal of Health Economics.* 30(5):933-949.

Dana J, Loewenstein G, (2003) A Social Science Perspective on Gifts to Physicians from Industry. *JAMA.* 290(2):252-255.

Darby M, Karni E (1973) Free Competition and the Optimal Amount of Fraud. *Journal of Law and Economics.* 16(1):67-88.

Dong H, Bogg L, Wang K, Rehnberg C, Diwan V (1999) A description of outpatient drug use in rural China: evidence of differences due to insurance coverage. *International Journal of Health Planning and Management.* 14(1):41-56.

Dulleck U, Kerschbamer R, Sutter M (2011) The Economics of Credence Goods: An Experiment on the Role of Liability, Verifiability, Reputation, and Competition. *American Economic Review.* 101(2):526-555.

Eisenhardt K (1989) Agency Theory: An Assessment and Review. *The Academy of Management Review.* 14(1):57-74.

Evans A, Krueger J (2009) The Psychology (and Economics) of Trust. *Social and Personality Psychology Compass.* 3(6):1003-1017.

Fehr E, Gachter S, Kirchsteiger G (1997) Reciprocity as a contract enforcement device: Experimental evidence. *Econometrica.* 65(4):833-860.

Fehr E, Klein A, Schmidt K (2007) Fairness and contract design. *Econometrica*. 75(1):121-154.

Fehr E, Schmidt K (1999) A Theory of Fairness, Competition, and Cooperation. *QJE.* 114(3):817-868.

Fischbacher U (2007) Z-tree: Zurich toolbox for readymade economic experiments. *Experiment. Econom*. 10(2):171-178.

Greenwood R, Li S, Prakash R, Deephouse D (2005) Reputation, Diversification, and Organization Explanations of Performance in Professional Service Firms. *Organization Science.* 16(6):661-673.

Ham S, Wu J, Lim N, Koch I (2019) Conflict of Interest in Third-Party Reviews: An Experimental Study. Working Paper.

Hampson L, Agrawal M, Joffe S, Gross C, Verter J, Emanuel E (2006) Patients' Views on Financial Conflicts of Interest in Cancer Research Trials. *N Engl J Med*. 355(22):2330-2337.

Herbig P, Milewicz J (1993) The relationship of reputation and credibility to brand success. *Journal of Consumer Marketing*. 10(3):18-24.

Ho T-H, Lim N, Camerer C (2006) Modeling the Psychology of Consumer and Firm Behavior with Behavioral Economics. *JMR.* 43(3):307-331.

Ho T-H, Weigelt K (2005) Trust Building Among Strangers. *Management Science.* 51(4):519-530.

Ismayilov H, Potters J (2013) Disclosing advisor's interests neither hurts nor helps. *Journal of Economic Behavior & Organization.* 93: 314-320.

Jensen M, Meckling W (1976) Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure. *Journal of Financial Economics.* 3(4):305-360.

Joseph K (2001) On the optimality of delegating pricing authority to the sales force. *J. Marketing.* 65(1): 62-70.

Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. *Econometrica.* 47(2):263-291.

Ke B, Yu Y (2006) The Effect of Issuing Biased Earnings Forecasts on Analysts' Access to Management and Survival. *Journal of Accounting Research.* 44(5):965-999.

Kerschbamer R, Neururer D, Sutter M (2016) Insurance Coverage of Customers Induces Dishonesty of Sellers in Markets for Credence Goods. *PNAS.* 113(27):7454-7458.

Lal R (1990) Improving Channel Coordination Through Franchising. *Marketing Science.* 9(4):299-318.

Libby R, Hunton J, Tan H, Seybert N (2008) Relationship Incentives and the Optimistic/Pessimistic Pattern in Analysts' Forecasts. *Journal of Accounting Research.* 46(1):173-198.

Lim N, Ham S (2014) Relationship Organization and Price Delegation: An Experimental Study. *Management Science.* 60(3):586-605.

Loewenstein G, Cain D, Sah S (2011) The Limits of Transparency: Pitfalls and Potential of Disclosing Conflicts of Interest. *American Economic Review.* 101(3):423-28.

Lu F (2014) Insurance Coverage and Agency Problems in Doctor Prescriptions: Evidence from a Field Experiment in China. *Journal of Developmental Economics.* 106:156-167.

McCabe K, Rigdon M, Smith M (2003) Positive reciprocity and intentions in trust games. *Journal of Economic Behavior & Organization.* 52:267-275.

Mehran H, Stulz R (2007) The Economics of Conflicts of Interest in Financial Institutions. *Journal of Financial Economics.* 85: 267-296.

Michaely R, Womack K (1999) Conflict of Interest and the Credibility of Underwriter Analyst Recommendations. *The Review of Financial Studies.* 12(4):653-686.

Mimra W, Rasch A, Waibel C (2016) Second Opinions in Markets for Expert Services: Experimental Evidence. *Journal of Economic Behavior & Organization.* 131:106-125.

Moore D, Loewenstein G, Tanlu L, Bazerman M (2003) Auditor Independence, Conflict of Interest, and the Unconscious Intrusion of Bias. Working Paper.

Morris T, Empson L (1998) Organisation and expertise: an exploration of knowledge bases and the management of accounting and consulting firms. *Accounting Organizations and Society.* 23(5):609-624.

O'Brien P, McNichols M, Lin J-W (2005) Analyst Impartiality and Investment Banking Relationships. *Journal of Accounting Research.* 43(4):623-650.

Özer Ö, Subramanian U, Wang Y (2018) Information Sharing, Advice Provision, or Delegation: What Leads to Higher Trust and Trustworthiness? *Management Science*. 64(1):474-493.

Özer Ö, Zheng Y, Chen K-Y (2011) Trust in Forecast Information Sharing. *Management Sci.* 7(6):1111-1137.

Özer Ö, Zheng Y, Ren Y (2014) Trust, trustworthiness, and information sharing in supply chains bridging China and the United States. *Management Sci*. 60(10):2435-2460.

Perlis R, Perlis C, Wu Y, Hwang C, Joseph M, Nierenberg A (2005) Industry sponsorship and financial conflict of interest in the reporting of clinical trials in psychiatry. *Am J Psychiatry.* 162(10):1957-1960.

Pesendorfer W, Wolinsky A (2003) Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice. *Review of Economic Studies.* 70(2):417-437.

Rao H, Greve H, Davis G (2001) Fool's gold: Social profit in the initiation and abandonment of coverage by Wall Street analysts. *Admin. Sci. Quart.* 46:502-526.

Rousseau D, Sitkin S, Burt R, Camerer C (1998) Not So Different After All: A Cross-Discipline View of Trust. *Academy of Management Review.* 23(3):393-404.

Sah S, Loewenstein G (2013) Nothing to Declare: Mandatory and Voluntary Disclosure Leads Advisors to Avoid Conflicts of Interest. *Psychological Science.* 25(2):575-584.

Sah S, Loewenstein G (2015) Conflicted advice and second opinions: Benefits, but unintended consequences. *Organizational Behavior and Human Decision Processes.* 130:89-107.

Sah S, Loewenstein G, Cain D (2013) The burden of disclosure: Increased compliance with distrusted advice. *Journal of Personality and Social Psychology.* 104(2):289-304.

Schneider H (2012) Agency Problems and Reputation in Expert Services: Evidence from Auto Repair. *The Journal of Industrial Economics.* 60(3):406-433.

Shapiro C (1982) Consumer Information, Product quality and Seller Reputation. *The Bell Journal of Economics*. 98:659-79.

Silverman G, Loewenstein G, Anderson B, Ubel P, Zinberg S, Schulkin J (2010) Failure to discount for conflict of interest when evaluating medical literature: a randomized trial of physicians. *J Med Ethics.* 36(5):265-270.

Spence M (1973) Job Market Signaling. *Quarterly Journal of Economics*. 87(3):355-374.

Spence M (2002) Signaling in Retrospect and the Information Structure of Markets. *American Economic Review*. 9(3):434-459.

Sülzle K, Wambach A (2005) Insurance in a Market for Credence Goods. *Journal of Risk & Insurance.* 72(1):159-176.

Thaler R (1980) Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior and Organization.* 1(1): 39-60.

Wazana A (2000) Physicians and the pharmaceutical industry: is a gift ever just a gift? *JAMA.* 283(3):373-380.

Wilks T (2002) Predecisional Distortion of Evidence as a Consequence of Real-Time Audit Review. *The Accounting Review.* 77(1):51-71.

Study 1 Appendix

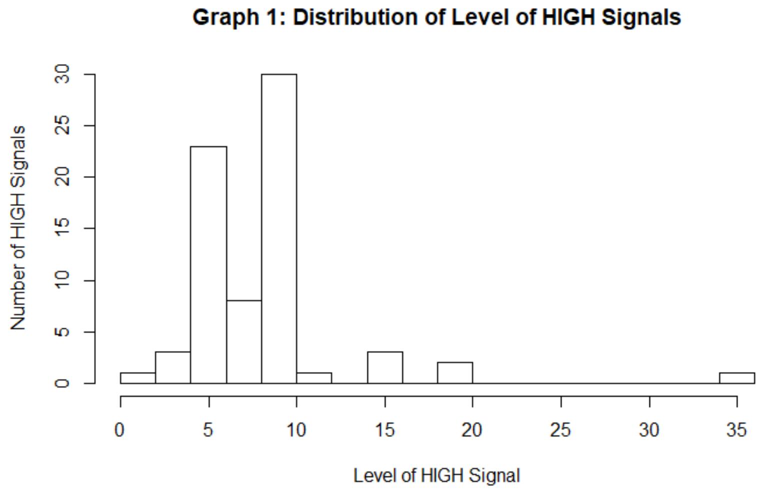

Study 3 Appendix

Table A1: OLS Regressions Showing ECP Longer than PIE

	DV: Ride Duration (Mins)		
	Control	Misinfo	Misinfo + Trust
	(1)	(2)	(3)
ECP	-0.209	3.912***	2.802***
	(1.898)	(0.691)	(1.045)
Rainy	1.036**	2.327**	1.167
	(0.468)	(0.981)	(1.252)
Traffic jam	1.432	-1.612***	3.579***
	(0.929)	(0.506)	(0.752)
Driver 45+	1.227***	-0.683	0.199
	(0.451)	(0.558)	(0.962)
Comfort	-0.362	-0.891	-1.402*
	(0.568)	(0.560)	(0.782)
10:00-11:29	-0.012	0.400	-0.822
	(0.434)	(0.598)	(1.243)
11:30-12:29	0.370	0.147	-1.409**
	(1.139)	(0.972)	(0.694)
12:30-14:29	-0.295	2.413***	0.679
	(0.397)	(0.354)	(1.350)
14:30-16:30	1.855**	2.790***	1.781
	(0.847)	(0.518)	(1.239)
Female RA	-0.732*	0.223	1.233*
	(0.432)	(0.387)	(0.699)
Vietnamese RA	-0.350	0.411	-1.178***
	(0.286)	(0.675)	(0.346)
Constant	21.481***	21.807***	23.195***
	(0.767)	(0.750)	(0.626)
Observations	80	108	98
R^2	0.268	0.485	0.298
Adjusted R^2	0.150	0.426	0.208
F Statistic	2.263**	8.210***	3.316***
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A2: OLS Regressions Showing ECP More Expensive than PIE

	DV: Ride Fare (S)		
	Control	Misinfo	Misinfo + Trust
	(1)	(2)	(3)
ECP	2.597	1.571***	2.755***
	(1.699)	(0.519)	(0.881)
Rainy	-0.774	0.291	-0.875
	(1.214)	(0.309)	(1.084)
Traffic jam	0.794	-0.149	2.236***
	(0.813)	(0.439)	(0.689)
Driver 45+	-0.907**	0.428	-0.214
	(0.390)	(0.453)	(0.714)
Comfort	-0.562	-0.484*	-0.066
	(0.484)	(0.256)	(0.558)
10:00-11:29	1.362***	0.801	0.401
	(0.478)	(0.636)	(0.515)
11:30-12:29	0.083	0.383	-0.517
	(0.456)	(0.637)	(0.620)
12:30-14:29	0.314	0.123	2.706**
	(0.694)	(0.713)	(1.085)
14:30-16:30	0.739	0.079	1.798*
	(0.619)	(0.350)	(1.011)
Female RA	-0.115	0.433	0.470
	(0.271)	(0.494)	(0.544)
Vietnamese RA	-0.575**	-0.184	0.667
	(0.264)	(0.487)	(0.869)
Constant	19.134***	18.096***	17.433***
	(0.332)	(0.297)	(0.833)
Observations	80	108	98
R^2	0.194	0.132	0.310
Adjusted R^2	0.064	0.032	0.022
F Statistic	1.492	1.323	3.516***

Note: *p<0.1; **p<0.05; ***p<0.01

Table A3: Logit Regressions Showing Overtreatment Effect

	DV: Driver Took ECP = 1	
	Base = Control (1)	Base = Misinfo (2)
Control	-	-1.580*** (0.486)
Misinfo	1.580*** (0.486)	-
Misinfo + Trust	2.375*** (0.764)	0.796* (0.421)
Rainy	0.254 (0.659)	0.254 (0.659)
Traffic jam	1.052 (0.717)	1.052 (0.717)
Driver 45+	-0.526** (0.247)	-0.526** (0.247)
Comfort	0.030 (0.310)	0.030 (0.310)
10:00-11:29	-0.148 (0.439)	-0.148 (0.439)
11:30-12:29	-0.837 (0.849)	-0.837 (0.849)
12:30-14:29	0.219 (0.419)	0.219 (0.419)
14:30-16:30	-0.549 (0.582)	-0.549 (0.582)
Female RA	-0.023 (0.239)	-0.023 (0.239)
Vietnamese RA	0.003 (0.187)	0.003 (0.187)
Constant	-2.867*** (0.604)	-1.287*** (0.391)
Observations	286	286
R^2	0.175	0.175
chi^2 (df = 12)	31.552***	31.552***

Note: *p<0.1; **p<0.05; ***p<0.01

Table A4: OLS Regression Analyzing Overtreatment with Misinfo Condition as Baseline

	DV: Driver Took ECP = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
Control	-0.130***	-0.128***	-0.130***	-0.118***	-0.117***	-0.119***
	(0.026)	(0.025)	(0.025)	(0.024)	(0.025)	(0.026)
Misinfo + Trust	0.109*	0.110*	0.126**	0.133**	0.133**	0.126*
	(0.059)	(0.059)	(0.061)	(0.061)	(0.062)	(0.067)
Rainy		0.039	0.041	0.027	0.026	0.030
		(0.093)	(0.094)	(0.094)	(0.095)	(0.099)
Traffic jam			0.149	0.147	0.150	0.151
			(0.111)	(0.111)	(0.107)	(0.111)
Driver 45+				-0.085**	-0.086**	-0.077**
				(0.035)	(0.036)	(0.035)
Comfort					0.013	0.003
					(0.046)	(0.044)
10:00-11:29						-0.021
						(0.065)
11:30-12:29						-0.100
						(0.097)
12:30-14:29						0.036
						(0.063)
14:30-16:30						-0.070
						(0.071)
Female RA	-0.007	-0.006	-0.020	-0.021	-0.021	-0.016
	(0.028)	(0.027)	(0.025)	(0.026)	(0.026)	(0.028)
Vietnamese RA	-0.008	-0.007	-0.019	-0.018	-0.018	0.0001
	(0.025)	(0.025)	(0.025)	(0.027)	(0.027)	(0.026)
Constant	0.174***	0.167***	0.158***	0.214***	0.207***	0.228***
	(0.022)	(0.021)	(0.024)	(0.039)	(0.038)	(0.055)
Observations	286	286	286	286	286	286
R^2	0.063	0.064	0.077	0.088	0.088	0.101
Adjusted R^2	0.049	0.047	0.057	0.065	0.062	0.061
F-Statistic	4.70*** (4, 281)	3.821*** (5, 280)	3.880*** (6, 279)	3.811*** (7, 278)	3.335*** (8, 277)	2.55*** (12, 273)

Note: *p<0.1; **p<0.05; ***p<0.01

Study 4 Appendix

Figure A1: Extensive Form of Game in Study 4