

The Implicit Honesty Premium:
Why Honest Advice Is More Persuasive than Highly Informed Advice

Uriel Haran

Ben-Gurion University of the Negev

Shaul Shalvi

University of Amsterdam

Article reference:

Haran, U., & Shalvi, S. (2020). The Implicit Honesty Premium: Why Honest Advice Is More Persuasive than Highly Informed Advice. *Journal of Experimental Psychology: General*, 149(4), 757-773. DOI: 10.1037/xge0000677.

© 2019, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. The data and materials for this article are available on the Open Science Framework (<https://osf.io/7hqkz>) and the Supplementary Online Material can be found at http://urielharan.com/eda_som. Thanks to Omer Lambez, Yael Levy and Asaf Mazar for help in data collection. This research was supported by the German-Israeli Foundation for Scientific Research and Development (grant I-2492-105.4/2017) and by the European Research Council (grant ERC-StG-637915). Correspondence concerning this article should be addressed to Uriel Haran at uharan@bgu.ac.il.

Abstract

Recipients of advice expect it to be both highly informed and honest. Suspecting either one of these attributes reduces the use of the advice. Does the degree of advice use depend on the reason for suspecting its accuracy? Five experiments tested the effect of the type of suspicion on advice taking. We find that recipients of advice discount it more severely when they suspect intentional bias than when they suspect unintentional error, for example, due to the advisor's insufficient knowledge. The effect persisted when we controlled for, and disclosed, the actual accuracy of the advice; it persisted when participants' own evaluations of the quality of the advice, as well as their desire to receive it, were equally high under both types of suspicion. Finally, we find the effect of suspicion on advice use stems from the different attributions of uncertainty associated with each type of suspicion. The results suggest people place an implicit premium on advisors' honesty, and demonstrate the importance of establishing reputation for advisors' success.

Keywords: advice, error, dishonesty, suspicion, uncertainty.

Advice is an important aid to decision-making. When people make a decision in an uncertain environment, they try to reduce uncertainty by seeking advice from others (Jonas, Schulz-Hardt, & Frey, 2005; Snizek & Van Swol, 2001). At the same time, advisees face new uncertainty regarding whether the advice is accurate or misleading. For example, the advice may be based on insufficient knowledge of the advisor and thus be prone to error. Alternatively, the advisor may have personal interests other than helping the advisee make a correct decision, which might motivate the advisor to provide biased advice. These factors can make the advice detrimental, rather than conducive, to the decision maker's ability to make a correct choice or accurate judgment.

Consider a car salesperson advising a client about which of two cars to buy. When asked about the cars' relative fuel efficiency, the salesperson highly recommends one car over the other, and informs the client that it gets very high gas mileage. To the extent that the advice is trusted, it helps reduce the client's experienced uncertainty. But the client may suspect the advice is not entirely accurate. The suspicion may arise for various reasons. For example, the client might think the salesperson does not really know the car's real gas mileage; therefore, the estimate is merely an uninformed guess. Alternatively, the client might suspect the salesperson has a special motive to sell this particular car, such as a higher commission he would earn for selling it, and would provide a high estimate regardless of whether it is accurate or inaccurate. In both cases, the question (i.e., How fuel efficient is the car?) and the advice (i.e., the gas-mileage figure stated by salesperson) are the same. Therefore, the accuracy or inaccuracy of the advice are identical, and the only difference is the reason for suspecting its accuracy (i.e., an error or intentional bias). Does the client's willingness to follow the salesperson's advice depend on the reason for the client's suspicion? This present research investigates this question.

We study the biasing effect of the reason to suspect the accuracy of advice. We measure people's willingness to follow advice when they suspect it suffers from a certain degree of unintentional error and when they suspect the same degree of inaccuracy, but for intentional reasons. Normatively, advice of the same quality should be used, or discounted, equally. However, we propose that different reasons to suspect the advice influence decision makers' sense of uncertainty, which in

turn determines their willingness to use it. As a result, advisees would discount potentially biased advice more severely than they would discount potentially erroneous advice, even when its actual quality, as well as the advisees' subjective evaluations of it, are similar.

Advice Use and Suspicion

Recipients' trust in advice depends on their beliefs about (a) the degree to which the advisor possesses genuine and accurate information, and (b) the advisor's perceived intention to share genuine information with them (Sperber et al., 2010). When recipients believe the advisor is sufficiently knowledgeable and motivated help them, they are expected to follow the advice. But suspicion of either one of these attributes makes people more vigilant and may reduce their propensity to incorporate the advice in their judgments. Although prior research has studied the effect of suspicion on the use of information, most of it focused on beliefs about the honesty of the source, and compared this type of suspicion to a state of uninterrupted trust. These studies suggest recipients of information generally assume the people who communicate the information are being honest, and thus take their statements at face value (Bond & DePaulo, 2006; McCornack & Parks, 1986; Zuckerman, DePaulo, & Rosenthal, 1981). On the other hand, people are averse to being exploited by others, or to "being duped" (ten Brinke, Vohs, & Carney, 2016; Vohs, Baumeister, & Chin, 2007). After experiencing intentional harm by others, people's trust levels typically decrease to a point that makes full recovery difficult, if not impossible (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Haselhuhn, Schweitzer, & Wood, 2010; Reeder & Brewer, 1979; Schweitzer, Hershey, & Bradlow, 2006). When suspicion about the source arises, it overrides the honesty assumption and leads recipients to discount the information (van Swol, 2009). Recipients may also expect deceivers to exert extra effort in masking the deception so as to convince their recipients that the information they communicate is accurate (Schul, Mayo, & Burnstein, 2008). As a result, they display greater preference for minimizing risk and relying on their own intuitions. All these findings suggest suspicion should decrease the willingness to take advice.

Whereas suspicion of intentional bias is expected to increase vigilance relative to no suspicion, what happens when the recipient trusts the honesty of the advisor but still suspects the advice might be misleading, for example, because the advisor may lack relevant knowledge? Should different reasons to suspect the advice have different effects on how recipients judge it? People may infer that dishonest intentions may increase the likelihood that the advice would be misleading, or they may believe lack of knowledge would sway the advice further away from the true answer. The reason for suspicion might influence expectations of its quality and accuracy. However, once recipients determine the expected accuracy of the advice, their use of it should be proportionate to the degree of its potential inaccuracy, regardless of why it may be inaccurate. Expecting a certain level of inaccuracy due to dishonesty and expecting the same level of inaccuracy due to the advisor's insufficient knowledge should harm the use of advice equally.

Very little research has been conducted on different reasons for suspicion, and that research has yielded mixed results. One study found the propensity to revise an opinion after realizing it was based on false information does not depend on the reason for the falsehood. Green and Donahue (2011) presented participants a magazine article and elicited their opinions about the article's topic. They then informed participants that certain important details in the article were false. One group learned the reporter was misinformed, whereas the other group was told the reporter "made up important facts." A subsequent questionnaire measured the shift in participants' opinions following this explanation, and found no differences between the groups. Other research finds people are sensitive to the way inaccurate information is produced when considering it. Schul, Mayo, Burnstein, and Yahalom (2007) showed participants a number of matchboxes, each containing either a blue or a yellow token. Participants predicted the color of the token in each box, with the help of a clue suggesting the correct color. Half of the participants received their clues from a computer, said to be programmed to randomly provide an incorrect clue one third of the time, and the other half received their clues from a person who was allowed to deceive them up to one third of the time. Participants in the deception condition chose according to the clues they received significantly less frequently than participants in the computer condition. The result suggests the reason for possible inaccuracy does

matter to recipients of information. However, because the reason for suspecting the information was determined both by the type of uncertainty associated with it (i.e., chance vs. deception) and by the source that communicated it (i.e., computer vs. person), the unique contribution of each factor is unclear. We attempt to resolve this issue by varying the type of uncertainty or reason for suspecting the information while keeping its source constant.

We study whether people treat the same advice differently when they suspect it is intentionally biased than when they view it as potentially erroneous, and if so, why. We address these questions by (a) varying the type of suspicion while keeping constant the advice itself, as well as its source, and (b) distinguishing between advice seeking (i.e., the willingness to receive advice or accept information from the advisor), advice taking (i.e., the motivation to follow vs. reject the advice), and advice use (i.e., the weight given to advice in the recipient's final estimate). In addition, we seek to uncover the underlying mechanism of the effect of suspicion by identifying the causal link between the reason for suspicion, the type of uncertainty, and advice use. We propose that although feeling uncertain generally increases the willingness to accept advice and revise one's opinion, random error and intentional bias are associated with different types of uncertainty, which may affect people's openness to advice.

Advice Use and Uncertainty

Uncertainty is an important condition for advice taking; the more uncertain people feel, the likelier they are to seek advice (Gino & Moore, 2007). When experiencing low knowledge, and therefore high uncertainty, people tend to give the advice higher weight in their final judgments (Yaniv, 2004). However, uncertainty can take multiple forms, and research on uncertainty examines not only its degree or severity, but also its type. The literature distinguishes between chance uncertainty and epistemic uncertainty. Chance uncertainty, also called aleatory uncertainty, is viewed as being caused by chance factors and randomly determined outcomes that cannot be known in advance. It is typically represented in relation to a class of possible outcomes and is naturally measured by relative frequency (Fox & Ulkumen, 2011). When people attribute uncertainty to chance,

they perceive it as more external to themselves and do not expect to have a great degree of control. They do not seek to eliminate the uncertainty as much as to manage it by determining the relative propensities of events (Howell & Burnett, 1978; Kahneman & Tversky, 1982). By contrast, epistemic uncertainty is attributed to intentional or controllable factors (Tannenbaum, Fox, & Ulkumen, 2017). Epistemically uncertain values are ones that are unknown to the judge but are “knowable” (or known to someone else). Epistemic uncertainty focuses on the extent to which an event is true or false, and can be reduced by searching for patterns or causality. Whereas assessment of purely aleatory uncertainty entails evaluation of the propensity of each possible outcome on a continuum, assessment of purely epistemic uncertainty generally entails evaluation of binary, true/false outcomes (Fox & Ulkumen, 2011).

Chance uncertainty encourages statistical thinking, which involves acknowledging one’s ignorance and a relatively high tolerance for error (Grove & Meehl, 1996). In a statistical mindset, people perceive error as inevitable and acceptable and the achievement of complete knowledge as impossible. Rather than perfect accuracy, the criterion for quality associated with statistical thinking is calibration, or a match between how often an event occurs and how often the forecaster predicts it will occur (Wallsten, Budescu, Erev, & Diederich, 1997). Epistemic uncertainty encourages clinical thinking, including attempts to understand and reduce the uncertainty of the information. The quality criterion for clinical thinkers is diagnosticity, or the match between the outcome of a given event and the forecaster’s prediction for that event (Wallsten et al., 1997). The clinical approach motivates attempts to understand the causal texture of the environment and seeks to achieve perfect predictability, although just like the statistical approach, it often fails to do so (Einhorn, 1986).

Although both types of uncertainty may increase advice taking, they might do so for different reasons. When experiencing chance uncertainty, people acknowledge their ignorance and have more trust in external sources of knowledge. When trusting a source, receivers of the information are willing to give up control over their judgment (Einhorn, 1986; Mayer, Davis, & Schoorman, 1995; Rousseau, Sitkin, Burt, & Camerer, 1998). The clinical mindset triggered by epistemic uncertainty seeks to increase understanding and maximize control (Einhorn, 1986; Spengler, 2013). Higher

epistemic uncertainty should increase the use of advice as long as its recipients believe it can minimize their uncertainty and lead them closer to the truth. When they do not trust the advice to be helpful, recipients might rely more on their own intuitions, or even refrain from making judgments altogether (Schul & Peri, 2015).

In light of these findings, we suggest suspecting intentional bias reduces the use of advice more than does the suspicion of random error of the same magnitude. Further, we suggest the difference in advice use occurs because intentional bias reduces the level of chance uncertainty associated with the advice. The present research advances our knowledge of perceived uncertainty, suspicion, and advice use in several ways. First, it provides a link between different types of uncertainty and advice taking. We test the role of suspicion and different uncertainty types in the exchange of information between advisor and advisee. By doing so, our work is the first to combine insights from the advice-taking literature, which has examined uncertainty as a unitary construct, and research on chance and epistemic uncertainty, which has mostly focused on the thought processes of independent judges. Second, we separate the source of information from the source of uncertainty associated with it. In previous studies, the two were often entangled, such that the source of information varied as part of the manipulation of uncertainty. For example, presenting information from a computer or a random sampling algorithm invoked chance uncertainty, whereas information from a dishonest person invoked epistemic uncertainty (Onkal, Goodwin, Thomson, Gonul, & Pollock, 2009; Schul et al., 2008, 2007). In the present work, we keep the source of information the same, namely, a person, in all of our studies. We also distinguish between the type of uncertainty and the type of suspicion, and measure their effects independently by varying one while keeping the other constant. By applying a moderation-of-process design, we seek to establish a causal link between suspicion, uncertainty type, and advice use. We first manipulate the reason for suspicion and measure its direct effect on advice use. Next, we measure the effect of suspicion on perceived chance and epistemic uncertainty, and test whether perceived uncertainty mediates the effect on advice use. Finally, we vary the level of chance uncertainty independently of suspicion, and measure the effects of these two factors on advice taking independently of each other.

Overview of studies

The present paper includes five experiments, in which we offered advice to decision makers and varied the reason for suspecting that advice. The designs of the experiments and their main findings are summarized in Table 1. Our research focuses on advice taking and information processing. To provide a direct estimate of people's evaluation of the information they receive, our experimental settings are intentionally kept free of ongoing relationships between advisors and advisees, because such relationships often involve the building and violation of trust by past interactions. Whereas studies of trust (e.g., Mayer, Davis, & Schoorman, 1995) examine evaluations of trustworthiness in people's character (e.g., factors that influence the trustworthiness of an advisor), we measure the assessment of information. Our experiments manipulate the specific type of suspicion in the information provided, and test how suspicion affects the way people process and use that information. In Experiments 1 and 4, we use a rather subtle suspicion manipulation by providing participants information about their advisors' knowledge and payoffs in a previous round. The advisor's incomplete knowledge prompts suspicion that subsequent advice would be erroneous, whereas previous incentives to mislead elicit suspicion that the advice may be intentionally biased. Experiments 2, 3, and 5 elicit suspicion more explicitly by directly telling participants their advisor would provide some erroneous or untruthful advice. Such blatant suspicion is common in word-of-mouth reviews and other client contexts. Using these various ways to manipulate suspicion helps us estimate the robustness of any observed effects.

The Experimental Paradigm

Our experiments employed the Judge-Advisor System (JAS; Gino & Moore, 2007; Sniezek & Buckley, 1995; Sniezek & Van Swol, 2001). This paradigm, the most commonly used method in advice research, elicits judgments and decisions from participants after providing them relevant advice. We measured advice taking and advice use in both choice and judgment tasks. In all experiments, participants made an initial estimate, received advice, and then had the opportunity to provide a revised, final estimate. Experiment 1 presented participants with two alternative answers to

a question, and their job was to choose the correct one. Advice taking was measured by the frequency with which participants revised their estimates to follow the advice. Experiments 2 and 5 administered numeric estimation tasks, and Experiments 3 and 4 elicited estimates of quantity. In these experiments, we calculated our dependent measure using a formula introduced by Harvey and Fischer (1997):

$$\text{Advice use} = (\text{final estimate} - \text{initial estimate}) / (\text{advice} - \text{initial estimate}).$$

The value yielded by this formula represents the degree to which a participant revised the estimate as a function of the difference between the advice and the initial estimate; higher values suggest greater influence of the advice. A value of zero means the decision maker was not swayed at all by the advice and kept the original estimate as it was. Negative values suggest the participant moved away from the advice. When the advice and the initial estimate are identical, the denominator's value is zero and the formula yields an undefined value, precluding the item from inclusion in the data analysis (see Bonaccio & Dalal, 2006). In our experiments, such instances occurred roughly 3% of the time.

We also collected additional measures, which were less central to our hypotheses, and included them in the Supplementary Online Material. We report how we determined our sample sizes, all data exclusions (if any), all manipulations, and all measures. All of the experiments were reviewed and approved by the Human Subjects Research Committee at Ben-Gurion University of the Negev.

Experiment 1

Experiment 1 tested whether suspicion of the advice, and the reason for suspecting it, affects recipients' propensity to accept the advice. We kept both the content and the source of the advice constant, and only varied recipients' suspicion of it. Participants completed two rounds of estimation tasks with the help of advice. The same advisor provided the advice in both rounds. After the first round, some participants received feedback on their advisor's performance. To make participants suspicious of subsequent advice, the feedback stated some of the advice they had received in the first

round may have been inaccurate. The feedback varied with regard to the reason for the possible inaccuracy. For one group, the inaccurate advice was attributed to missing knowledge, whereas for another group, it was the result of an incentive to mislead. A third group proceeded to the second round of feedback. We measured participants' advice taking both before and after the suspicion manipulation, which allowed us to control for differences in their individual tendencies to take advice. We predicted that (a) suspecting inaccuracy will reduce advice taking compared to not suspecting it; (b) suspecting intentional bias will lead to greater advice discounting than suspecting error.

Method

Two hundred forty residents of the US, Canada, and the UK ($M_{age} = 37.42$; 147 females), registered as workers on Amazon Mechanical Turk (MTurk), participated in the experiment in exchange for \$1 each and a chance to win a monetary bonus. We determined our sample size based on a power analysis of a medium-size effect ($f = 0.25$; $d = 0.5$) with 90% power, which suggested a minimum sample size of 207 participants.

The experiment consisted of 40 estimation tasks, divided into two rounds of 20 tasks. Each task presented a square matrix consisting of 100 cells, arranged in a 10×10 pattern. Each cell was in one of two colors (yellow and blue in round 1, green and orange in round 2).¹ The matrix was presented for one second. Participants were asked to choose the color that was represented in more cells in the matrix. They then received advice from an advisor who had more information about the matrices than they did and could either revise their choice or keep their initial one. Each correct final choice earned the participant a lottery ticket. At the end of data collection, we randomly selected five tickets and paid each of their holders a \$5 bonus.

The first round included 20 tasks. Participants received correct advice in 15 tasks and incorrect advice in the other five tasks. After completing this round, participants received feedback

¹ At the end of the study, we asked participants whether they suffered from any kind of color blindness that may have affected their performance. One participant reported having color blindness, but stated that he "was able to identify all colors in this study."

about the advice they received. We manipulated the feedback between three groups. The error group received the following text:

Advisor payment: For all 20 estimates, the advisor was paid (with a lottery ticket) each time you made a correct choice.

Advisor knowledge: For 10 of the 20 estimates, the advisor knew the correct answer before advising you. For the other 10 estimates, the advisor did not have any information and had to guess which color to advise you to choose. This means that on average, for those 10 estimates, the advisor made an error 5 times.

Participants in the intentional-bias group read:

Advisor knowledge: For all 20 estimates, the advisor knew the correct answer before advising you.

Advisor payment: For 10 of the 20 estimates, the advisor was paid (with a lottery ticket) each time you made a correct choice. For some of the other 10 estimates, the advisor was paid if you chose blue, even if blue was incorrect, and on the others the advisor was paid if you chose yellow, even if yellow was incorrect. This means that on average, for those 10 estimates, the advisor deceived you 5 times.

Participants in the control condition proceeded directly to round 2 without any feedback.

In round 2, participants completed another set of 20 tasks, with the same advisor as in round 1. Participants were informed that “as before, the advisor has more information about the matrices than you will have. But the advisor’s knowledge and payment are not necessarily the same as they were in the first part.” The advice in this round was correct in 14 tasks and incorrect in six; this fact was not disclosed to participants.

Measures. When the advice participants received was different from their initial choice, their options were to either take the advice by revising their choice or to ignore the advice. We measured advice use by the frequency at which participants revised their choices to follow the advice.

When the advice participants received matched their initial choice, their options were either to confirm their choice by not making a change or to exhibit active distrust by revising their final choice to be different from the advice. We measured active distrust by the frequency at which participants revised their choices to differ from the advice.

In addition, after the tasks, we elicited participants' evaluations of the advice they received in round 2. Participants estimated the number of tasks in round 2 in which they received correct advice and the degree to which they relied on the advice in this round (on a scale ranging from 0% = *not at all* to 100% = *total reliance*).

Results

Advice taking. We calculated advice-taking rates for rounds 1 and 2, and measured the effect of suspicion on advice taking using two mixed ANOVAs, with average advice taking in rounds 1 and 2 as repeated measures. One analysis measured the effect of suspicion by comparing the control group with the two treatment groups. The other tested the effect of the type of suspicion by comparing the error and intentional-bias groups. We found suspicion was detrimental to advice taking. As Figure 1 shows, the first analysis yielded a significant Round \times Suspicion interaction, $F(1, 238) = 5.22, p = .02, \eta^2 = .04$. As predicted, participants who underwent a suspicion manipulation were significantly less likely to take advice in round 2 than were participants in the control condition, who did not undergo the manipulation, $F(1,237) = 10.02, p = .002, \eta^2 = .04$. The second analysis compared the error group and the intentional-bias group. It also yielded a significant interaction between round and the reason for suspicion, $F(1, 158) = 7.45, p = .007, \eta^2 = .05$. Consistent with our prediction, participants who suspected intentional bias took the advice significantly less frequently than did those who suspected random error, $F(1,157) = 7.25, p = .008, \eta^2 = .04$.

Active distrust. We calculated the rates at which participants revised their choices in the opposite direction of the advice, after initially agreeing with it, in rounds 1 and 2. We then subjected these rates to the same planned comparisons as advice taking. The analyses found that active distrust was rather infrequent, and it decreased from round 1 ($M = 0.03, SD = 0.07$) to round 2 ($M = 0.02, SD =$

0.06), $F(1, 238) = 6.64, p = .01, \eta^2 = .03$. Neither suspicion nor the reason for it had a significant effect on active distrust, $F_s \leq 1.08, p_s \geq .30, \eta^2 < .01$.

Discounting misleading advice vs. accurate advice. Reluctance to use advice following the suspicion manipulation is justified if suspicion helps advice recipients distinguish between accurate and misleading advice. Correct detection of misleading advice would result in its more severe discounting, without affecting the use of accurate advice. We tested this possibility by calculating separate advice-taking measures for correct advice and for incorrect advice. We subjected these measures to two mixed ANOVAs that measured the effects of suspicion and of the reason for suspicion. One analyses included task round and advice accuracy as repeated measures and suspicion as a between-subjects factor. The analysis yielded a significant difference between the use of correct advice and the use of incorrect advice, $F(1, 230) = 15.69, p < .001, \eta^2 = .06$, suggesting participants were able, to some extent, to identify when the advice they received was correct and when it was not. The comparison between the suspicion and control conditions yielded a significant Round \times Suspicion interaction, $F(1, 230) = 4.56, p = .03, \eta^2 = .02$ but no significant three-way interaction, $F < 1$. The second analysis was conducted among participants in the two suspicion conditions, and compared the use of advice by participants in the error and intentional-bias groups. This analysis also yielded a significant interaction between task round and reason for suspicion, $F(1, 151) = 5.82, p = .02, \eta^2 = .04$ but no three-way interaction, $F(1, 151) = 1.06, p = .30$. As Figure 2 shows, the suspicion manipulation, especially suspicion of intentional bias, reduced participants' willingness to use advice, but this was true for both correct and incorrect advice.

Post-task Advice-quality Evaluation. Participants underestimated the quality of the advice in round 2. They received correct advice in 14 of the 20 tasks, but estimated only 11.16 such instances, on average ($SD = 3.73$), $t(239) = 11.81, p < .001, d = 0.76$. A one-way ANOVA revealed a significant between-group effect, $F(2, 237) = 5.79, p = .003, \eta^2 = .05$. Planned contrasts found that suspicion of either error ($M = 10.29, SD = 3.63$) or intentional bias ($M = 10.96, SD = 3.89$) led to lower-quality assessments than those observed in the control condition ($M = 12.23, SD = 3.43$), $t(237) = 3.20, p = .002, d = 0.42$, but the type of suspicion had no effect on these assessments, $t(237) = 1.17,$

$p = .24$. This result suggests participants in the intentional-bias condition did not think they received advice inferior to that provided to the error group. Reported reliance on the advice displayed similar patterns, $F(2, 237) = 4.45, p = .01, \eta^2 = .04$. Planned contrasts reveal participants reported greater reliance on advice in the control condition ($M = 43.56\%, SD = 29.93$) than in the two suspicion conditions, $t(237) = 2.51, p = .01, d = 0.33$, whereas the difference between the error group ($M = 37.46\%, SD = 28.71$) and the intentional-bias group ($M = 30.19\%, SD = 26.39$) did not reach significance, $t(237) = 1.62, p = .11$.

Discussion

The findings of Experiment 1 supported both our predictions. First, suspicion of any kind decreased recipients' likelihood of revising their original choices after receiving advice. This finding suggests both intentional and unintentional failures to provide accurate advice decrease the persuasiveness of subsequent advice. Second, not all suspicion effects were created equal: Recipients were less likely to take advice when they suspected intentional bias than when they suspected error. Although their advisors' initial performance levels were identical, recipients who attributed prior imperfections to conflicting interests were more reluctant to take subsequent advice than those who attributed it to insufficient knowledge.

Advice-taking patterns suggest participants may have put a premium on honesty in their willingness to follow advice. The honesty premium, however, appears to be implicit: Participants in the intentional-bias condition did not think they were receiving lower-quality advice. Their willingness to follow that advice, however, was diminished nonetheless. Experiment 2 tested the generalizability of these findings in a different type of estimation setting and distinguished between advice seeking, advice taking, and advice use.

Experiment 2

Experiment 1 left some open questions. First, does the severe discounting of intentionally biased advice reflect *a priori* preference for honest (even if inaccurate) information? Participants who had been intentionally misled were more reluctant to follow advice than participants whose advisor

was uninformed, even though they were no less optimistic about the accuracy of the advice. If people simply prefer listening to honest advisors over dishonest ones, then we should observe a difference not only in the willingness to follow advice (i.e., advice taking), but also in the desire to receive it in the first place (i.e., advice seeking). In Experiment 2, we tested advice seeking separately from its use by participants, by measuring their willingness to pay (WTP) for advice before receiving it.

Another question that emerges from the results of Experiment 1 is whether the effect resulted from participants' reactions to being duped (Vohs et al., 2007). In Experiment 1, participants received advice after discovering the advisor had intentionally misled them. Their behavior may have been influenced by retributive motives, such that rejecting the advice was their way to punish or distance themselves from the advisor, rather than a calculated assessment of the advice's potential accuracy. Experiment 2 precluded these possible influences by notifying participants in advance of the expected quality of the advice, as well as the reason it may be inaccurate. Finally, does the quality of the advice matter? In Experiment 1, the rate of inaccurate advice was the same for all recipients. In Experiment 2, we varied advice quality by manipulating the number of pieces of accurate and inaccurate advice participants expected to receive.

Experiment 2 differed from the previous experiment in three additional attributes. First, rather than using categorical choice tasks, Experiment 2 elicited numerical estimates, which allowed us to test the generalizability of the effect of the reason for suspicion. We measured both advice taking, which we tested in Experiment 1, and advice use, defined as the degree of influence the advice had on its recipient's final judgment. Second, whereas Experiment 1 elicited suspicion indirectly, by making participants infer advisors' future performance from their past performance, in the present experiment, we directly informed participants of the reason the advice might be inaccurate. We also informed them explicitly about the expected outcome of the advice, namely, its expected accuracy rate.

Method

Two hundred eighty-five MTurk workers residing in the US, Canada, and the UK ($M_{age} = 37.42$; 147 females) participated in the experiment in exchange for \$0.50 each and a chance to win a

monetary bonus. We determined our sample size based on a power analysis of a medium-size effect ($f = 0.25$; $d = 0.5$) with 90% power, which suggested a sample of 231 participants. We increased the sample to account for the anticipated number of participants who would complete the tasks without the help of advice.

The experimental task differed from the one used in Experiment 1 in two major aspects. First, it included only one round of estimation tasks. Second, participants made numerical estimates, rather than categorical choices. In each task, they estimated the average of nine numbers arranged in a 3×3 matrix, which appeared on the screen for two seconds. After estimating the average, participants received advice from an advisor who had viewed the numbers for 10 seconds, and were given the opportunity to revise their estimates. An accurate final estimate (i.e., an estimate within 10 units of the true answer) earned the participant a lottery ticket. At the end of data collection, we randomly selected one ticket, and paid its owner a \$10 bonus.

Manipulations. Before beginning the tasks, participants learned their advisor would be randomly chosen from among the following four options: Advisor A had provided accurate advice for all six estimates; Advisor B had provided accurate advice for five estimates and inaccurate advice for one; Advisor C had provided accurate advice for four of the six estimates; and Advisor D had provided accurate advice for only three estimates. We specified that accurate advice was within 10 units of the correct answer, and that the inaccurate advice all missed the correct answer by more than 20 units. All four advisors were presented on the same page, so that they were easily comparable. We manipulated the reason to suspect the advice by varying the description of the inaccurate advice. The unintentional error group was told that in these cases, the advisor “was inaccurate, meaning that the advice missed the true answer by more than 20,” whereas the intentional-bias group learned the advisor “deliberately provided advice that missed the true answer by more than 20.”

Measures. We measured advice seeking by eliciting participants’ WTP for advice. Each advisor had a predetermined, undisclosed minimum price. Participants indicated the maximum price they would pay out of their monetary bonus to receive advice from each advisor. After the WTP

elicitation, every participant was randomly assigned one advisor. If the participant's WTP for this advisor's advice was equal to the advisor's minimum price or higher, the participant received advice throughout the experiment, for the price he or she was willing to pay. If WTP was lower than the minimum price, the participant completed the tasks without the help of advice.

To measure advice taking, we coded whether participants took the advice by revising their estimates in its direction, did not take the advice and stayed with their initial estimate, or displayed active distrust by adjusting their estimates away from the advice. We then calculated the relative frequency of each outcome for every participant. Advice use was measured based on Harvey and Fischer's (1997) advice-taking formula discussed above, which is sensitive to the degree to which participants revised their estimates following the advice.

Finally, after completing the tasks, participants evaluated the quality of the advice they had received. We elicited participants' estimates of the average error of the advice as well as the price they now thought would have been adequate to pay for it.

Results

Nineteen participants were unwilling to pay any amount for advice. Eight participants reported WTP more than their entire bonus payment, and six participants violated the dominance assumption, indicating higher WTP for advice of lower quality than for higher quality. We excluded these participants from all analyses, resulting in a sample of 252 participants.

Advice Seeking. We conducted a mixed ANOVA on participants' WTP, with advice quality as the within-subjects variable and reason for suspicion (error vs. intentional bias) as the between-subjects factor. As Figure 3 shows, we found no main effect of reason for suspicion, $F(1,250) = 1.31$, $p = .25$, or an interaction between the WTP for advice of each level of quality and the reason to suspect it, $F(3,750) = 1.31$, $p = .27$. This finding refutes the proposition of an *a priori* preference for advice prone to unintentional error over advice prone to intentional bias. It also suggests the type of suspicion may affect advice taking, but not advice seeking.

Advice Taking and Advice Use. Thirty participants reported lower WTP than the advisors' minimum price, and proceeded to complete the tasks without the help of advice. Each estimation task included numbers that were all in the same hundred series (e.g., between 200 and 299, between 600 and 699). Individual estimates that were outside the hundred series (e.g., an estimate of 168 for a group of numbers between 200 and 299) were considered invalid and were removed. We set an inclusion criterion of three valid estimates, such that only participants who provided a valid estimate in at least half the tasks were included in the data analysis. One participant provided only one valid estimate and was removed. The participant's removal did not affect the results.² We applied the same exclusion criterion in all subsequent experiments as well. The sample for the following analyses consisted of 221 participants.

We tested the effects of advice quality and reason for suspicion on participants' advice-taking outcomes (i.e., following the advice, not following the advice, and active distrust) using a multivariate ANOVA, and the effects on advice use (i.e., the weight participants gave the advice in their final estimates) with a univariate ANOVA. Consistent with our prediction, the analyses yielded a main effect of reason for suspicion on all advice-taking outcomes as well as on advice use. Consistent with the findings of Experiment 1, suspecting intentional bias was associated with lower willingness to take advice, $F(1,213) = 10.16, p = .002, \eta^2 = .05$, and with higher frequency of both remaining with one's initial estimate, $F(1,213) = 5.45, p = .02, \eta^2 = .03$, and of active distrust, $F(1,213) = 14.66, p < .001, \eta^2 = .06$. Suspecting intentional bias also reduced the weight participants gave the advice in their judgments, $F(1,213) = 9.97, p = .002, \eta^2 = .05$. Advice quality also affected all advice-taking outcomes, $F_s \geq 3.42, p \leq .02, \eta^2 = .05$, and advice use, $F(1, 213) = 13.85, p < .001, \eta^2 = .16$. The interaction between advice quality and suspicion produced no significant effects, $F_s < 1$.

Figure 4 presents the mean use of advice in each condition, grouped by three advice-quality categories: perfect advice (accurate in all estimates), high-quality advice (mostly accurate), and low-quality advice (as likely to be inaccurate as accurate). Simple effects tests reveal the effect was most

² Of the remaining 221 participants, 183 provided valid estimates in all tasks, 34 provided one invalid estimate, two provided two invalid estimates, and two provided three invalid estimates. Excluding all participants who provided one or more invalid estimates did not affect the pattern of the results.

pronounced when participants had reason to suspect the advice, but still expected it to be mostly accurate. The effect was significant among participants who received accurate advice in five of six tasks (advice taking: $t(58) = 2.19, p = .03, d = 0.58$; advice use: $t(58) = 2.22, p = .03, d = 0.57$) and in four of six tasks (advice taking: $t(49) = 1.94, p = .06, d = 0.55$; advice use: $t(49) = 2.29, p = .03, d = 0.65$). The suspicion manipulation did not affect the willingness to take or use low-quality advice, which was equally likely to be inaccurate as accurate, $t_s < 1$. The perfectly accurate advice was not subject to suspicion of either type, and, as expected, participants who received this advice were not significantly affected by the suspicion manipulation, (advice taking: $t(58) = 1.61, p = .11$; advice use: $t(58) = 1.18, p = .24$). A full report of the results of the advice-taking outcomes appears in the Supplementary Online Material.

Discounting misleading advice vs. accurate advice. As in Experiment 1, we tested whether suspecting intentional bias improved participants' ability to detect misleading advice, relative to participants who suspected random error. We calculated average advice-use rates for accurate and misleading advice, and conducted a mixed ANOVA with advisor quality and reason for suspicion as between-subjects factors and advice accuracy as a repeated measure. Note this analysis excludes participants who received accurate advice for all their estimates. On average, participants who suspected intentional bias were less receptive to advice than their counterparts who suspected error, both when the advice was misleading (intentional bias: $0.21 \leq M \leq 0.47, 0.35 \leq SD \leq 0.75$; error: $0.34 \leq M \leq 0.57, 0.30 \leq SD \leq 0.40$) and when it was accurate (intentional bias: $0.24 \leq M \leq 0.49, 0.26 \leq SD \leq 0.47$; error: $0.45 \leq M \leq 0.61, 0.33 \leq SD \leq 0.40$), $F(1, 150) = 5.36, p = .02, \eta^2 = .03$. We found no effect of advice accuracy, $F < 1$, suggesting participants gave similar weights to accurate advice in their estimates as they did to misleading advice. We also found no significant interaction between the reason for suspicion and advice accuracy, $F < 1$, or a three-way interaction with advisor quality, $F(2, 150) = 1.53, p = .22$. These results suggest that, as in Experiment 1, suspecting intentional bias did not improve participants' ability to detect misleading advice.

Post-task Advice-Quality Evaluation. Participants assessed the quality of the advice they received by estimating its average error and the price they now thought, having completed the task,

would be appropriate to pay for it.³ We subjected both measures to a two-way ANOVA. We found that whereas actual advice quality significantly affected both measures, $F_s \geq 12.87$, $p_s < .001$, the type of suspicion did not affect either measure, $F_s < 1$. These results are consistent with participants' WTP for advice, as well as with the findings of Experiment 1, and suggest the premium on honesty in advice taking and advice use is implicit, rather than a reflection of a conscious preference for honest, even if inaccurate, advice.

Discussion

Experiment 2 replicated the finding that the reason for suspicion affects the willingness to take advice. Participants who suspected intentional bias were less likely to take the advice than those who suspected it was prone to unintentional error of the same magnitude. Intentional bias also reduced the weight the advice received in recipients' final estimates. In the remaining studies, we only report analyses of advice use, which proved to be a more sensitive measure than advice taking.

The effect of the reason for suspicion persisted despite participants' advance knowledge of the actual quality of the advice. Whereas in Experiment 1, participants were less open to take advice after being intentionally misled, in Experiment 2, they exhibited the same effect even before any misleading took place. Discounting advice, then, seems to be more than a reaction to the feeling that one has been duped (Vohs et al., 2007). Additionally, the fact that participants were willing to pay as much for intentionally misleading advice as for erroneous advice suggests they were not consciously trading off advice quality for honesty. They exhibited rational behavior by being sensitive to the objective accuracy of the advice and unmoved by the reason for its inaccuracy. But the reason for inaccuracy still affected the persuasiveness of the advice. This finding suggests the type of suspicion affects the cognitive processing of the advice's content, if not the conscious desire to receive it.

Experiment 2 also identified a boundary for the effect. Participants were more receptive to advice suspected to be erroneous than to advice suspected to be intentionally misleading when they expected the advice to be accurate most of the time. When the quality of the advice was low, such that

³ We excluded one participant who reported a value of \$815 and three participants who estimated an average advice error of more than 100 units.

participants expected it to be inaccurate as frequently as they expected it to be close to the true answer, they discounted the advice at similar rates regardless of the reason for the frequent inaccuracy. Thus, the implicit honesty premium does not hold for any level of advice quality, but rather occurs when recipients trust the advisor to generally, though not always, be helpful.

Experiment 3

Experiments 1 and 2 found the persuasiveness of advice is reduced when recipients believe it is intentionally misleading than when they attribute its possible inaccuracy to unintentional factors. The effect seems to occur during the cognitive processing of the advice, even when recipients assess the quality of the advice equally, and show similar willingness to receive it. Experiment 3 investigated the cognitive aspects of advice use more closely by examining the relationship between suspicion and the type of uncertainty recipients experience, and how this relationship affects their use of advice.

We argue that suspecting random error and suspecting intentional bias are associated with different types of perceived uncertainty. Following the main conceptualizations of uncertainty in the literature, we distinguish between two types of uncertainty: epistemic and chance. When people attribute uncertainty to controllable factors, they are likely to apply a clinical mindset, which is characterized by attempts to reduce uncertainty (Einhorn, 1986; Tannenbaum et al., 2017). Conversely, the perception of uncertainty as the product of chance or of randomly determined outcomes that cannot be known in advance encourages statistical thinking. In this mindset, people acknowledge their ignorance and are more tolerant of error (Grove & Meehl, 1996). Whereas uncertainty of any kind may encourage reliance on advice, suspecting random error and suspecting intentional bias may trigger different levels of each type of uncertainty, which, in turn, may affect people's willingness to incorporate the advice in their judgments.

In Experiment 3, we manipulated the reason to suspect advice and measured participants' perceived chance and epistemic uncertainty, in addition to their rates of advice use. We predicted that recipients who suspect intentional bias in the advice would perceive lower chance uncertainty and higher epistemic uncertainty than those who attribute inaccuracies to unintentional error. Prior

research has found that when people are not knowledgeable or are not confident in their knowledge, they display higher advice use (Gino & Moore, 2007; Yaniv, 2004). These findings show an association between higher perceived and greater use of advice, but they are mute regarding the type of uncertainty participants felt. Therefore, we did not have a formal hypothesis about whether chance and epistemic uncertainty differ in their relationships with advice use. Nevertheless, we tested these relationships and whether they mediate the effect of the reason for suspicion on advice use. The hypotheses, experimental design, and planned analyses were preregistered online at <http://aspredicted.org/blind.php?x=4s8wu4>.

Method

Three hundred sixty participants (169 females, 190 males, one did not declare; $M_{age} = 35.15$) were invited via MTurk to participate in a study titled “How Much Money Is In the Bottle?” in exchange for \$0.50 and a chance to win a monetary bonus. A pilot study found a small effect on perceived uncertainty ($d = 0.35$); therefore, we determined our sample size of 350 based on a power analysis of an effect of this size for a two-group design, which suggested a sample of 346 participants, and adjusted for expected preregistered exclusions.

The experiment consisted of one round of six estimation tasks. In each task, participants saw a picture of a 0.5-liter water bottle partially filled with US coins in all denominations from \$0.01 to \$0.25. Participants’ job was to estimate the amount of money in each bottle. To provide a sense of the expected amounts, we presented participants with two pictures, one of a bottle that contained \$10.82 and one that contained \$19.05. In each task, participants were presented with a picture and estimated the amount of money that was in the bottle. They then received advice from an advisor who provided them advice for all six estimates. After receiving the advice, participants had the opportunity to make a revised final estimate. Each final estimate that was within \$1 of the correct amount awarded the participant a lottery ticket for a \$1 bonus. After the end of data collection, we held a draw among all lottery tickets won by participants and paid the owners of 10 drawn lottery tickets their bonus.

We manipulated the reason to suspect the advice by varying the advisor's knowledge and incentives. Before the task, half of the participants (the error group) received the following information:

Advisor payment: The advisor would get a bonus for every accurate final estimate you make (within \$1 of the exact correct amount).

Advisor knowledge: In 4 of the 6 rounds, the advisor received the correct amount, rounded down to the nearest whole number (for example, "\$15" for an amount of between 15.00 and 15.99). But in the other 2 rounds, the advisor did not have any information and had to guess. In these two rounds, the advisor guessed wrong, and gave you inaccurate advice that missed the correct answer by more than \$1.

The other half of participants (the intentional-bias group) received the following information:

Advisor knowledge: In every round, the advisor received the correct amount, rounded down to the nearest whole number (for example, "\$15" for an amount of between 15.00 and 15.99).

Advisor payment: In 4 of the 6 rounds, the advisor would get a bonus for every accurate final estimate you make (within \$1 of the exact correct amount). But in the other 2 rounds, the advisor would get a bonus if your answer misses the correct amount by more than \$1. In these two rounds, the advisor intentionally misled you, and gave you biased advice that missed the correct answer by more than \$1.

After the instructions, we administered a questionnaire measuring the level of chance and epistemic uncertainty participants perceived. We used Tannenbaum et al.'s (2017) EARS-4 measure, which is a reduced version of the 10-item EARS questionnaire introduced by Ulkumen, Fox, and Malle (2016), plus an additional item from the original questionnaire. Two items measured chance uncertainty ("The accuracy of the advice has an element of randomness;" "It feels like the accuracy of the advice is determined by chance factors"). Three other items measured epistemic uncertainty ("How accurate the advice will be has been determined in advance;" "The accuracy of the advice is knowable in advance, given enough information;" "The accuracy of the advice is something that well-informed people would agree on"). Participants answered each item on a 7-point scale (1 = *Not at all*,

7 = *Very much*). All five items appeared in a different, random order for each participant. After completing the questionnaire, participants proceeded to the estimation task. After the tasks, they evaluated the advisor's performance by estimating the average error in the advice they had received, and rated on a 1–5 scale the extent to which they used the advice, the level of influence the reason for the error had on their advice use, and their anger with the advisor.

Results

We preregistered the following criteria for participant inclusion in the analyses: First, we included only participants whose responses to the EARS questionnaire had a variance of more than zero. Ten participants responded with the same value to all items and were excluded. Second, we defined as invalid those estimates that yielded outlier advice-use scores, which made up 1.6% of the estimates in the dataset, and those estimates that did not have an advice-use score because the initial estimate was identical to the advice, which occurred 0.1% of the time. The preregistered inclusion criterion specified a minimum of three valid estimates for a participant to be included in the analyses. All participants achieved the criterion, and the final sample consisted of 350 participants.

Advice use and post-task evaluations. We tested the effect of the reason for suspecting the advice on participants' use of it using an independent samples t-test. As predicted, and consistent with the previous experiments, suspecting intentional bias was associated with lower advice use ($M = 0.46$, $SD = 0.34$) than was suspecting error ($M = 0.55$, $SD = 0.31$), $t(348) = 2.70$, $p = .007$, $d = .29$. As in previous experiments, the reason for suspicion did not affect participants' assessments of the advisor's performance, $t < 1$. Additionally, we found no effect on participants' anger toward the advisor, $t < 1$. According to participants' self-reports, intentional bias did not have a significantly greater effect on their behavior than did error due to insufficient knowledge, $t(348) = 1.45$, $p = .15$, but those who suspected error reported having used the advice to a greater extent ($M = 3.30$ on a 1-7 scale, $SD = 0.88$) than did those who suspected intentional bias ($M = 3.00$, $SD = 0.92$), $t(348) = 3.07$, $p = .002$, $d = 0.33$. Self-reported advice use correlated significantly with the actual weight given to the advice, $r = .547$, $p > .001$.

Discounting misleading advice vs. accurate advice. As in the first two experiments, we tested whether suspecting intentional bias improved participants' ability to detect misleading advice, relative to those who suspected random error. We calculated average advice-use rates for accurate and misleading advice, and conducted a mixed ANOVA with reason for suspicion as a between-subjects factor and advice accuracy as a repeated measure. Participants who suspected intentional bias were less receptive to advice than were their counterparts who suspected error, $F(1, 348) = 5.75, p = .02, \eta^2 = .02$. This was true both when the advice was misleading (intentional bias: $M = 0.51, SD = 0.44$; error: $M = 0.56, SD = 0.41$) and when it was accurate (intentional bias: $M = 0.43, SD = 0.34$; error: $M = 0.54, SD = 0.33$), which rendered the interaction effect non-significant, $F(1, 348) = 1.69, p = .19$. Interestingly, participants gave significantly more weight in their estimates to misleading advice than to accurate advice, $F(1, 348) = 5.39, p = .02, \eta^2 = .02$. These results provide no indication of an ability to correctly identify misleading advice, and suggest yet again that suspecting intentional bias had no positive effect on this ability.

Perceived uncertainty. A confirmatory factor analysis on the EARS items yielded two Promax-rotated factors that account for 67.61% of the variation. The two factors align with the items' categorization as measures of chance uncertainty (aleatory uncertainty, by Tannenbaum et al.'s (2017) terminology) and epistemic uncertainty. The factor loadings and reliability scores of the items are presented in Table 2.

We tested the effect of the reason for suspicion on the degree to which participants perceived the advice as characterized by chance and epistemic uncertainty. A mixed ANOVA, with suspicion condition as a between-subjects variable and the type of uncertainty as a repeated measure, found a significant interaction between the reason for suspicion and type of uncertainty, $F(1, 348) = 6.53, p = .01, \eta^2 = .02$. As predicted, suspecting error was associated with significantly higher ratings of chance uncertainty ($M = 4.86, SD = 1.26$) than suspecting intentional bias ($M = 4.49, SD = 1.49$), $t(348) = 2.48, p = .01, d = .27$. The difference in epistemic uncertainty, however, did not reach significance, contrary to our prediction (error: $M = 4.58, SD = 1.17$; intentional bias: $M = 4.73, SD = 1.26$), $t(348) = 1.22, p = .22$.

Next, we tested whether the two types of perceived uncertainty mediate the effect of suspicion on advice use by computing indirect effects of each of 5,000 bootstrapped samples (Hayes, 2013, model 4). Whereas both chance uncertainty ($\beta = .03, p = .03$) and epistemic uncertainty ($\beta = .04, p = .007$) were positively related to advice use, only the former mediated the effect of suspicion. The analysis found a significant indirect effect of intentional bias on advice use through perceived chance uncertainty (indirect effect = $-.01, SE = .006, 95\% CI [-.03 \text{ } -.0004]$), whereas its indirect effect via epistemic uncertainty was not significant (indirect effect = $.006, SE = .006, 95\% CI [-.004 \text{ } .02]$).

Discussion

Experiment 3 sheds light on the links between the reason for suspicion in advice, the type of uncertainty that characterizes it, and advice use. First, the experiment found both chance uncertainty and epistemic uncertainty predict use of advice. Second, the results provide partial support for the hypotheses regarding the relationship between the types of perceived uncertainty and advisees' suspicion. As predicted, suspecting intentional bias was associated with lower perceived chance uncertainty than suspecting error, and this difference led to lower use of advice. The relationship with epistemic uncertainty, however, was not as clear. Perceived epistemic uncertainty displayed a positive relationship with advice use but was not significantly affected by the reason for suspicion.

Therefore, although advice is more likely to influence one's judgment when one feels uncertain, the reason for suspecting the advice seems to influence one type of uncertainty more than the other. If chance uncertainty is responsible for the difference in advice taking under different suspicion contexts, then varying the level of this type of uncertainty could override the effect of suspicion and eliminate the difference in advice use. To test this proposition, the following two experiments manipulated characteristics of chance uncertainty independently of the reason for suspicion. Experiment 4 tested whether reducing perceived chance uncertainty by making the advice seem more deterministic reduces reliance on honest advice. In Experiment 5, we tested whether increasing chance uncertainty by introducing a random component to the advice can increase its use among recipients who suspect intentional bias.

Experiment 4

Experiment 4 tested whether controlling the level of chance uncertainty associated with the advice can overturn the effect of suspicion on advice use. We manipulated the chance uncertainty by varying the apparent randomness in the error of the advice. In the learning phase of the experiment, half of participants received advice that was systematically on one side of the correct answer, either consistently too high or consistently too low. Because chance uncertainty is associated with randomness, perceiving inaccuracies as non-random reduces the attribution of uncertainty to chance factors, regardless of the reason for the inaccuracy. The other participants completed the learning phase with advice that sometimes overestimated the true answer and sometimes underestimated it. Similarly to Experiment 1, we manipulated the reason for suspecting subsequent advice by varying the information about the advisor's knowledge and incentives in the learning phase. Chance uncertainty could therefore be reduced by either the suspicion of receiving intentionally biased advice or by the realization that the advice is consistently skewed in one direction, or both.

Our experimental design included four conditions. One condition produced high chance uncertainty by providing advice characterized by unintentional error that was randomly either above or below the true answer. The other three conditions were characterized by low chance uncertainty; in these conditions, the inaccuracies that caused suspicion were either intentional or systematic in their direction. We predicted that lower chance uncertainty, caused by advice that appeared systematically skewed in a certain direction, will lead to lower use of subsequent advice and will eliminate the difference between suspicion of error and suspicion of intentional bias.

Method

Two hundred MTurk workers residing in the US, Canada, the UK, and Australia ($M_{age} = 36.60$; 107 females) participated in the experiment in exchange for \$0.45 each and a chance to win a \$6 bonus. We determined our sample size based on a power analysis of a medium-size effect ($f = 0.25$; $d = 0.5$) with 90% power, which suggested a minimum sample size of 171 participants.

The experiment consisted of a learning phase that included 25 estimation tasks, and a test phase that included 10 similar tasks. Each task included a 20 cell \times 20 cell matrix, presented for three seconds. Every cell in the matrix was one of two colors (blue and orange in the learning phase, green and purple in the test phase). Participants estimated the number of cells of one color (blue in the learning phase, green in the test phase) that appeared in the matrix. After providing their initial estimates, participants received advice from an advisor and had the opportunity to revise their estimates.

We manipulated chance uncertainty via the variance of the apparent inaccuracy of the advice in the learning phase. The advice for all participants was of the same quality (mean absolute error = 14.93, $SD = 13.11$). Participants in the random-inaccuracy condition received advice that sometimes overestimated the true answer and sometimes underestimated it (underestimation of between 1 and 34 units, overestimation of between 2 and 34 units; $M = -0.13$, $SD = 20.27$). In the systematic-inaccuracy condition, the advice consistently appeared either too high or too low. Half of the participants in this condition received advice that overestimated the true answer in all tasks and the other half received advice that always underestimated the true answer. For all participants, the advice in this phase approximated the correct answer ($|error| < 10$) in nine tasks and inaccurate advice ($|error| \geq 25$) in six tasks. To ensure participants were aware of these differences, we gave them feedback at the end of the learning phase, which included a table presenting their final estimates, the correct answers, the advice they received, and whether it was “close enough” (i.e., within 10 units of the correct answer), “too high,” or “too low.”

The feedback at the end of the learning phase also served as our manipulation of the reason for suspicion. In the random-error condition, the feedback stated that the advisor’s incentive was to make the recipient’s estimates as accurate as possible. In nine tasks, the advisor saw the matrix for 20 seconds before providing the advice, but in the remaining six tasks, the advisor saw the matrix for only two seconds and had to guess the correct answer. The systematic-error group also read that in these six tasks, the advisor saw mainly [did not see enough] blue cells. In the intentional-bias conditions, the feedback stipulated that the advisor saw the matrix for 20 seconds in all tasks and was

incentivized to help the recipient make accurate estimates in nine tasks. The random-intentional-bias group read that in some of the other six tasks, the advisor was paid to make the recipient's advice as high as possible, and in others, as low as possible, whereas the systematic-intentional-bias group read that the advisor was incentivized to make the recipient's estimates of the number of blue cells as high [low] as possible.

Next, participants proceeded to the test phase of the experiment. This phase included 10 estimation tasks in the same format as those in the learning phase. Participants were informed they would receive advice from the same advisor who provided advice in the previous phase, and that the advisor again had more information about the matrices than they would have. We measured advice use with the same formula as in the other experiments. After completing the test phase, participants assessed the quality of the advice they received by estimating its average error.

Results

We conducted an exclusion procedure similar to that used in the previous experiments. Because Experiment 4 involved estimates of quantities between 0 and 400, we defined invalid estimates as ones that missed the correct answer by 200 units or more, and removed them from the dataset. We applied the same inclusion criterion used in Experiments 2 and 3, which required valid estimates in at least half the tasks for a participant to be included in the data analysis. All participants in Experiment 4 achieved this criterion.

Advice Use. The average rate of advice use in the test phase was 0.43 ($SD = 0.29$). We tested our prediction using a planned contrast, which compared the random-error condition, in which chance uncertainty was high, with the three conditions characterized by low chance uncertainty, due to either intentional bias or systematic inaccuracy in the learning phase. Figure 5 shows that, as predicted, participants who experienced high chance uncertainty in the learning phase used subsequent advice significantly more than the other groups, whose advice was characterized by a lower level of chance uncertainty, $t(196) = 3.36$, $p = .001$, $d = 0.48$. As predicted, among participants who suspected unintentional error, observing systematic inaccuracy in the learning phase reduced advice use in the

test phase, $t(96) = 2.79$, $p = .006$, $d = 0.57$, whereas the manipulation did not affect the use of advice suspected to be intentionally misleading, $t < 1$. A two-way ANOVA on advice use in the test phase, using advice use in the learning phase as a covariate, found a significant Suspicion \times Systematic-Inaccuracy interaction, $F(1,195) = 4.89$, $p = .03$, $\eta^2 = .02$.

Post-task Advice-Quality Evaluation. Consistent with the previous experiments, epistemic uncertainty did not significantly affect participants' estimates of the average error in the advice they received, $t(196) = 1.59$, $p = .11$. Although the reason for suspicion had a marginally significant main effect across the entire sample, $F(1,196) = 2.98$, $p = .09$, $\eta^2 = .09$, it did not significantly affect error estimates in either the random or systematic inaccuracy condition, $ts \leq 1.33$, $ps \geq .19$. All other main, interaction, and simple effects were non-significant.

Discussion

Complementing the findings of Experiment 3, Experiment 4 tested the role of chance uncertainty in the effect of suspicion on advice use. In Experiment 4, we varied the level of chance uncertainty independently of the reason for suspecting the advice. Whereas unintentional error increases the attribution of inaccuracy to chance, intent to mislead reduces the perceived randomness associated with the advice, and, accordingly, the attribution of uncertainty to random factors. Because chance uncertainty was found to mediate the association between the reason for suspicion and advice use, we predicted that the degree to which advice quality is determined by random factors will affect participants' use of the advice.

In Experiment 4, we varied the level of perceived randomness associated with the inaccuracy of the advice in both suspicion conditions, and found support for our prediction. When participants observed systematic inaccuracy, the attribution of the inaccuracy to either intentional or unintentional factors did not affect their subsequent use of advice. When we eliminated differences in the type of perceived uncertainty, the effect of the type of suspicion on advice use disappeared. These results establish a causal link between the type of suspicion and advice use through perceived chance uncertainty, using a moderation-of-process design (Spencer, Zanna, & Fong, 2005). Moderation-of-process designs provide experimental rather than correlational evidence for the causal role of process

variables by directly manipulating the proposed mechanism (Adam, Shirako, & Maddux, 2010). When perceived uncertainty was not controlled, the type of suspicion affected advice use; directly manipulating uncertainty, while keeping the reason for suspicion constant, had a significant effect on advice use and eliminated the effect of suspicion. This finding demonstrates the causal relationship between the reason for suspicion and advice use, via the level of chance uncertainty in the advice. If, as the present experiment found, reducing apparent randomness lowers advice use, then increasing the perceived randomness associated with the advice might also increase its use by recipients, even when they suspect the advice is prone to intentional manipulation. Experiment 5 tested this proposition.

Experiment 5

Experiment 5 complements Experiment 4 by testing the causal link between the type of suspicion and advice use through perceived uncertainty. Whereas in Experiment 4 we reduced chance uncertainty by making the error in the advice appear less arbitrary, in Experiment 5, we increased chance uncertainty by adding a random component to the advice-giving process. In one condition, participants received advice from the same source for all estimates, whereas in another condition, they received advice from one of two different sources who randomly alternated between estimates. The accuracy of the advice was identical for all participants, and in each condition, one group attributed inaccurate advice to unintentional error, whereas the other group perceived it as intentionally biased. We predicted greater advice use in settings characterized by high chance uncertainty than in those where chance uncertainty is low. To test this prediction, we compared the three conditions in which chance uncertainty was high, due to either suspected error or randomly alternating sources, with the group that experienced low chance uncertainty by suspecting intentional bias in advice provided by a single source.

Method

Two hundred eleven MTurk workers residing in the US ($M_{age} = 36.33$; 107 females) participated in the experiment in exchange for \$0.50 each and a chance to win a \$10 bonus. We

determined our sample size based on a power analysis of a medium-size effect ($f = 0.25$; $d = 0.5$) with 90% power, which suggested a minimum sample size of 171 participants.

We employed the same estimation task and suspicion manipulation used in the high-quality-advice condition of Experiment 2: All participants made six numerical estimates with the help of advice. Four pieces of advice were accurate and two were inaccurate, for either intentional or unintentional reasons. We also varied the perceived randomness of the advice process, by manipulating the source that provided the advice. In the single-source condition, one advisor provided the advice for all six estimates; the advice was accurate in four estimates and inaccurate in two. In the randomly alternating sources condition, participants learned the advice was obtained from two advisors. Each advisor provided advice for three estimates, accurate advice for two estimates, and inaccurate advice for one. Participants in this group were informed that the order of the estimates and of the advisors would be random, and they would not be told who gave the advice in each task. After completing the task, participants assessed the quality of the advice by estimating its average error.

Results

We used the same process for coding and removing invalid estimates and the same criterion of providing valid estimates in at least half of the tasks for inclusion in the analyses that we used in Experiments 2, 3, and 4. Six participants provided either zero or one valid estimate and were removed from the data set. All other participants met the inclusion criterion, bringing the final sample to 205 participants.

Average advice use across all conditions was 0.41 ($SD = 0.21$).⁴ We conducted a planned comparison between the three conditions characterized by high chance uncertainty (i.e., the two error groups and the intentionally misleading randomly alternating sources groups) and the one in which chance uncertainty was low (i.e., single, intentionally misleading source). As predicted, we found significantly lower use of advice among the latter group, where chance uncertainty was lowest, than in

⁴ We observed two outlying individual estimates, which displayed advice-use levels more than 8.5 standard deviations higher than their respective means, and omitted them from the analyses. Including these two estimates increases the mean advice-use score in the randomly alternating sources error condition from 0.46 to 0.49 ($SD = 0.28$), but does not change the direction or significance of any of the effects.

the groups that experienced higher chance uncertainty, $t(203) = 2.95, p = .004, d = 0.41$. Simple effects tests find that randomly alternating source selection significantly increased the use of advice suspected to be intentionally biased (single source: $M = 0.34, SD = 0.16$; randomly alternating sources: $M = 0.42, SD = 0.22$), $t(98) = 2.16, p = .03, d = 0.44$. However, for recipients who suspected random error, alternating between sources made no difference (single source: $M = 0.42, SD = 0.21$; randomly alternating sources: $M = 0.46, SD = 0.25$), $t < 1$. As in previous experiments, chance uncertainty had no effect on assessments of advice quality, $t < 1$, suggesting the effect is not due to overt preferences between different types of advice.

Discussion

Experiment 5 demonstrated that when suspecting intentional bias in advice, increasing chance uncertainty makes people more persuaded by it. In both the single-source and randomly alternating sources conditions, participants received the same advice and half of them suspected intentional bias. The only difference was that chance uncertainty was emphasized in one condition by randomly alternating between two sources of advice, rather than using the same source throughout the task. When chance uncertainty was higher, so was the degree to which participants followed the advice. Consistent with the results of Experiment 4, these findings support our prediction that increasing the uncertainty associated with apparent randomness increases the use of advice, and this difference drives the effect of the reason for suspicion on advice use.

The increased use of intentionally biased advice when its sources randomly alternated refutes the possibility that the effect of suspicion is due to a general aversion to dishonesty. In both the single-source and randomly alternating sources conditions, recipients experienced the same type of suspicion. Both groups received the same advice; therefore, the random-selection process improved neither the advice's expected nor actual quality. The only change was in the degree of apparent randomness and chance uncertainty. Increasing this type of uncertainty surrounding the advice eliminated its excessive discounting by recipients who suspected intentional bias.

General Discussion

Seeking advice is a common human activity. People seldom make decisions without first acquiring some information from an external source. Although obtaining perfectly accurate advice would be ideal, advice takers rarely know for certain how accurate the advice they receive would be. Instead, they must estimate the level of accuracy of the advice before determining whether to follow it, and if so, to what extent. Such estimation is based on beliefs about the advice's proneness to random error and to deliberate manipulation. Five experiments demonstrated how these beliefs shape recipients' use of advice. When they suspected a certain degree of inaccuracy in the advice, attributing it to unintentional reasons allowed recipients to accept the advice more often and give it more weight in their judgments than when they suspected intentional bias. This effect was unrelated to the quality of the advice. Controlling for actual quality, participants showed equal willingness to seek advice that was prone to error and advice that was prone to intentional bias. They also perceived the potentially biased advice as equally valuable and as likely to be accurate. Nevertheless, suspecting intentional bias made the advice less persuasive than advice suspected to be unintentionally erroneous. The studies varied in the amount of information participants had about the advisors' knowledge and incentives as well as in the timing of this knowledge. In Experiment 2, 3, and 5, participants were forewarned that they should expect to receive some erroneous or intentionally biased advice, whereas in Experiments 1 and 4, they learned they had received some erroneous or intentionally biased advice in a previous round of estimations. All these elicitations of suspicion resulted in the same pattern of results.

The findings of Experiments 3, 4, and 5 shed light on the mechanism that underlies the effect of suspicion on the use of advice. We reasoned that intentional bias is associated with epistemic uncertainty, whereas error is associated with chance uncertainty, and predicted these associations would mediate the effect of suspicion on advice use. The results of Experiment 3 provided partial support for our prediction. Whereas both types of uncertainty were associated with greater use of advice, only uncertainty that was attributed to chance mediated the effect of suspicion on advice use. Following this finding, we manipulated the level of perceived chance uncertainty in Experiments 4

and 5 and tested the use of advice by participants who suspected either error or intentional bias. Consistent with our reasoning, unintentional error that appeared non-random led participants to discount the advice as severely as advice prone to intentional manipulation. By contrast, increasing the perceived randomness associated with the advice significantly increased advice use under suspicion of intentional bias. These results help establish a causal link between suspicion and advice use through perceived chance uncertainty, in a moderation-of-process design (Spencer et al., 2005).

Our findings contribute to the knowledge on how individuals react to uncertainty. Previous research examined related issues, such as differences between uncertainty and ambiguity (Ellsberg, 1961), uncertainty about the content of the information versus its conveyer, and the effects of chance versus epistemic uncertainty (Schul et al., 2007). Our experiments kept both the source of the information and the content of the advice identical for all participants. Recipients of advice who suspected intentional bias did not believe the advice was less accurate than did those who suspected error, but nevertheless allowed the advice a lower degree of influence on their judgment. The finding suggests the degree to which individuals follow advice is determined not only by the factors that influence its perceived accuracy, but also by other factors that affect the advice seeker's mindset during the estimation process.

Practical Implications

Our findings provide applicable insight to both takers and givers of advice. The objective of the advice taker is to make the most effective use of advice and, ultimately, make accurate judgments and decisions. We found that whereas the expected accuracy of advice determined its influence on recipients, advice that was prone to intentional bias was consistently less persuasive than advice that was prone to random error, regardless of its expected accuracy. Advisees would benefit from relying more on their expectations of the overall quality of the advice and from ignoring irrelevant factors that can potentially bias their judgments.

Advisors and consultants can also learn a valuable lesson from our findings. Advisors want to provide accurate advice, but to be successful, their advice must first be persuasive enough to be

followed (Radzevick & Moore, 2011). Our findings suggest perceived honesty affects adherence to advice more than perceived expertise and knowledge. Participants followed advice associated with perfect honesty but imperfect knowledge more than they followed advice characterized by dubious honesty, even when the source's expertise was not questioned. Advisors and consultants should therefore pay special attention to their credibility as honest practitioners in the eyes of their clients, to ensure their advice is followed.

Our findings also contribute to the study of conflicts of interest. Advisors' potential conflicts of interest increase the perceived likelihood of bias in the information they share with their clients (Chugh, Bazerman, & Banaji, 2005; Pierce, 2012). Our studies highlight the peril in such a situation, even if the advisor effectively acts solely in the best interest of the recipient. It has been suggested that disclosing conflicts of interest to recipients could help maintain and enhance their trust in their advisors. The research on conflicts of interest has challenged this idea. Cain, Loewenstein, and Moore (2005) found that disclosing conflicts of interest reduces the pressure on advisors to reconcile their obligations toward their clients with their personal interests, and may cause them to provide more biased advice than they otherwise would. Our findings suggest such disclosure could also reduce recipients' reliance on the advice. Although it might protect clients from dishonest conduct by the advisor, it will likely not help the relationship between the parties.

Directions for Future Research

Our studies open the door for further research. For example, whereas our experiments kept the tests of suspicion free of confounding effects of communication and ongoing advisor relationships, such factors can influence the behavior of both advice givers and advice takers (Dana & Cain, 2015; Schwartz, Luce, & Ariely, 2011). Examining the role of suspicion on advice in more complex relationships may provide additional valuable insight. Another path for future research concerns the attributions of error and bias. Random error and intentional bias can have different reasons and motives. Advice may be unintentionally wrong because the advisor had insufficient knowledge, because the advisor was negligent, or because of unforeseen changes in the environment.

Similarly, an advisor might consciously mislead an advisee because of malicious motivations or pressure from superiors and commitment to the interests of the workplace (Moore & Gino, 2013). In Experiments 1, 3, and 4, we provided a reason for each type of inaccuracy; the results were consistent with those of the studies that did not include such disclosure. However, a more complex examination of higher-order assumptions made by advisees about their advisors may provide a more nuanced account of the effect and its sources.

Conclusion

At a time of abundant information and broad access to various sources of knowledge, receiving advice has become quite easy. Determining the quality of advice, however, remains as difficult as ever. Advice seekers wish to receive the most accurate and diagnostic information available. To that end, they must determine whether the information they receive is reliable, both in its imperviousness to random error and in the honesty of its source. Our findings suggest suspecting intentional bias reduces adherence to advice relative to suspecting error, independent of how accurate the advice is, or even how accurate the recipient believes it is. Taking such information into consideration may be beneficial to both givers and recipients of advice, and help maintain effective advisor-advisee relationships.

Context

The idea for this paper originated in a discussion between the authors about why people should either believe or disregard advice they receive from strangers. When we reviewed the literature on suspicion, we were surprised to find it has mostly focused on lie detection, aversion to deception, and other forms of suspecting intentional bias. Reasons for unintentional error, which recipients of advice often assume, have been largely neglected. We set out to investigate whether the type of suspicion matters, combining our background in the study of dishonesty (e.g., Shalvi, Dana, Handgraaf, & De Dreu, 2011; Shalvi, Eldar, & Bereby-Meyer, 2012; Weisel & Shalvi, 2015) and uncertainty in information processing (e.g., Haran & Moore, 2014; Haran, Ritov, & Mellers, 2013;

Moore, Tenney, & Haran, 2016). We hope our research can help people identify and remove an obstacle that impedes them from optimizing performance and maintaining fruitful social interactions.

References

- Adam, H., Shirako, A., & Maddux, W. W. (2010). Cultural variance in the interpersonal effects of anger in negotiations. *Psychological Science, 21*(6), 882–889. doi:10.1177/0956797610370755
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5*(4), 323–370. doi:10.1037//1089-2680.5.4.323
- Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes, 101*(2), 127–151. doi:10.1016/j.obhdp.2006.07.001
- Bond, C. F., & Depaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review, 10*(3), 214–234.
- Cain, D. M., Loewenstein, G., & Moore, D. A. (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies, 34*(1), 1–25. doi:10.1086/426699
- Chugh, D., Bazerman, M. H., & Banaji, M. R. (2005). Bounded ethicality as a psychological barrier to recognizing conflicts of interest. In D. A. Moore, D. M. Cain, G. Loewenstein, & M. H. Bazerman (Eds.), *Conflicts of Interest: Challenges and Solutions in Business, Law, Medicine, and Public Policy* (pp. 74–95). Cambridge, MA: Cambridge University Press.
- Dana, J., & Cain, D. M. (2015). Advice versus choice. *Current Opinion in Psychology, 6*, 173–176. doi:10.1016/j.copsyc.2015.08.019
- Einhorn, H. J. (1986). Accepting error to make less error. *Journal of Personality Assessment, 50*(3), 387–395. doi:10.1207/s15327752jpa5003_8
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics, 75*(4), 643–669.
- Fox, C. R., & Ulkumen, G. (2011). Distinguishing two dimensions of uncertainty. In W. Brun, G. Keren, G. Kirkeboen, & H. Montgomery (Eds.), *Perspectives on Thinking, Judging, and Decision Making* (pp. 21–35). Oslo: Universitetsforlaget.
- Gino, F., & Moore, D. A. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making, 35*, 21–35. doi:10.1002/bdm
- Green, M. C., & Donahue, J. K. (2011). Persistence of belief change in the face of deception: The effect of factual stories revealed to be false. *Media Psychology, 14*, 312–331. doi:10.1080/15213269.2011.598050
- Grove, W. M., & Meehl, P. E. (1996). Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: The clinical–statistical controversy. *Psychology, Public Policy, and Law, 2*, 293–323.
- Haran, U., & Moore, D. A. (2014). A better way to forecast. *California Management Review, 57*(1), 5–15.
- Haran, U., Ritov, I., & Mellers, B. A. (2013). The role of actively open-minded thinking in information acquisition, accuracy, and calibration. *Judgment and Decision Making, 8*(3), 188–201.
- Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes, 70*(2), 117–133. doi:10.1006/obhd.1997.2697

- Haselhuhn, M. P., Schweitzer, M. E., & Wood, A. M. (2010). How implicit beliefs influence trust recovery. *Psychological Science, 21*(5), 645–648. doi:10.1177/0956797610367752
- Howell, W. C., & Burnett, S. A. (1978). Uncertainty measurement: A cognitive taxonomy. *Organizational Behavior and Human Performance, 22*(1), 45–68. doi:10.1016/0030-5073(78)90004-1
- Jonas, E., Schulz-Hardt, S., & Frey, D. (2005). Giving advice or making decisions in someone else's place: The influence of impression, defense, and accuracy motivation on the search for new information. *Personality and Social Psychology Bulletin, 31*(7), 977–990. doi:10.1177/0146167204274095
- Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition, 11*, 143–157.
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review, 20*(3), 709–734. doi:10.5465/AMR.1995.9508080335
- McCornack, S. A., & Parks, M. R. (1986). Deception detection and relationship development: The other side of trust. *Annals of the International Communication Association, 9*(1), 377–389. doi:10.1080/23808985.1986.11678616
- Moore, C., & Gino, F. (2013). Ethically adrift: How others pull our moral compass from true North, and how we can fix it. *Research in Organizational Behavior, 33*, 53–77. doi:10.1016/j.riob.2013.08.001
- Moore, D. A., Tenney, E. R., & Haran, U. (2016). Overprecision in judgment. In G. Wu & G. Keren (Eds.), *Handbook of Judgment and Decision Making* (pp. 182–209). Hoboken: Wiley-Blackwell.
- Onkal, D., Goodwin, P., Thomson, M., Gonul, S., & Pollock, A. (2009). The relative influence of advice from human experts and statistical methods on forecast adjustments. *The Journal of Behavioral Decision Making, 22*, 390–409. doi:10.1002/bdm.637
- Pierce, L. (2012). Organizational structure and the limits of knowledge sharing: Incentive conflict and agency in car leasing. *Management Science, 58*(6), 1106–1121. doi:10.1287/mnsc.1110.1472
- Radzevick, J. R., & Moore, D. A. (2011). Competing to be certain (but wrong): Market dynamics and excessive confidence in judgment. *Management Science, 57*(1), 93–106. doi:10.1287/mnsc.1100.1255
- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review, 86*(1), 61–79. doi:10.1037/0033-295X.86.1.61
- Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. (1998). Not so different after all: A cross-discipline view of trust. *Academy of Management Review, 23*(3), 393–404.
- Schul, Y., Mayo, R., & Burnstein, E. (2008). The value of distrust. *Journal of Experimental Social Psychology, 44*(5), 1293–1302. doi:10.1016/j.jesp.2008.05.003
- Schul, Y., Mayo, R., Burnstein, E., & Yahalom, N. (2007). How people cope with uncertainty due to chance or deception. *Journal of Experimental Social Psychology, 43*(1), 91–103. doi:10.1016/j.jesp.2006.02.015
- Schul, Y., & Peri, N. (2015). Influences of distrust (and trust) on decision making. *Social Cognition, 33*(5), 414–435.
- Schwartz, J., Luce, M. F., & Ariely, D. (2011). Are consumers too trusting? The effects of relationships with expert advisers. *Journal of Marketing Research, 48*, S163–S174. doi:10.1509/jmkr.48.SPL.S163

- Schweitzer, M. E., Hershey, J. C., & Bradlow, E. T. (2006). Promises and lies: Restoring violated trust. *Organizational Behavior and Human Decision Processes*, *101*(1), 1–19. doi:10.1016/j.obhdp.2006.05.005
- Shalvi, S., Dana, J., Handgraaf, M. J. J., & De Dreu, C. K. W. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes*, *115*(2), 181–190. doi:10.1016/j.obhdp.2011.02.001
- Shalvi, S., Eldar, O., & Bereby-Meyer, Y. (2012). Honesty requires time (and lack of justifications). *Psychological Science*, *23*(10), 1264–70. doi:10.1177/0956797612443835
- Sniezek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*. doi:10.1006/obhd.1995.1040
- Sniezek, J. A., & Van Swol, L. M. (2001). Trust, confidence, and expertise in a judge-advisor system. *Organizational Behavior and Human Decision Processes*, *84*(2), 288–307. doi:10.1006/obhd.2000.2926
- Spencer, S. J., Zanna, M. P., & Fong, G. T. (2005). Establishing a causal chain: Why experiments are often more effective than mediational analyses in examining psychological processes. *Journal of Personality and Social Psychology*, *89*(6), 845–51. doi:10.1037/0022-3514.89.6.845
- Sperber, D., Clement, F., Heintz, C., Mascaro, O., Mercier, H., Origgin, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, *25*(4), 359–393.
- Tannenbaum, D., Fox, C. R., & Ulkumen, G. (2017). Judgment extremity and accuracy under epistemic versus aleatory uncertainty. *Management Science*, *63*(2), 497–518. doi:10.1287/mnsc.2015.2344
- ten Brinke, L., Vohs, K. D., & Carney, D. R. (2016). Can ordinary people detect deception after all? *Trends in Cognitive Sciences*, *20*(8), 579–588. doi:10.1016/j.tics.2016.05.012
- Ulkumen, G., Fox, C. R., & Malle, B. F. (2016). Two dimensions of subjective uncertainty: Clues from natural language. *Journal of Experimental Psychology: General*, *145*(10), 1280–1297. doi:10.1037/xge0000202
- van Swol, L. M. (2009). The effects of confidence and advisor motives on advice utilization. *Communication Research*, *36*(6), 857–873. doi:10.1177/0093650209346803
- Vohs, K. D., Baumeister, R. F., & Chin, J. (2007). Feeling duped: Emotional, motivational, and cognitive aspects of being exploited by others. *Review of General Psychology*, *11*(2), 127–141. doi:10.1037/1089-2680.11.2.127
- Wallsten, T. S., Budescu, D. V., Erev, I., & Diederich, A. (1997). Evaluating and Combining Subjective Probability Estimates. *Journal of Behavioral Decision Making*, *10*, 243–268.
- Weisel, O., & Shalvi, S. (2015). The collaborative roots of corruption. *Proceedings of the National Academy of Sciences*, *112*(34), 10651–10656. doi:10.1073/pnas.1423035112
- Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, *93*(1), 1–13.
- Zuckerman, M., DePaulo, B. M., & Rosenthal, R. (1981). Verbal and nonverbal communication of deception. *Advances in Experimental Social Psychology*, *14*, 1–59.

Tables

Table 1. *Summary of all experiments' designs and findings.*

Experiment	Estimate type	Advice quality	Method of eliciting suspicion	Main findings
1	Categorical	High (15/20 accurate)	Disclosure of advisor's knowledge and incentives in previous round.	Suspecting intentional bias reduced advice taking relative to suspecting error.
2	Numeric	Perfect (6/6 accurate) High (5/6 accurate) High (4/6 accurate) Low (3/6 accurate)	Disclosure of the reason for expected inaccuracy.	Suspecting intentional bias reduced use of high-quality advice relative to suspecting error. Suspicion type did not affect willingness to pay for advice.
3	Quantity	High (4/6 accurate)	Disclosure of advisor's knowledge and incentives.	Suspecting intentional bias reduced use of advice relative to suspecting error. Suspecting error was associated with higher perceived chance uncertainty than suspecting intentional bias. Perceived chance uncertainty mediated the effect of the reason for suspicion on advice use.
4	Quantity	High (9/15 accurate)	Disclosure of advisor's knowledge and incentives in previous round. Disclosure of the degree and direction of inaccuracies in previous round.	Suspecting systematic inaccuracy, both intentional and unintentional, reduced advice use, relative to suspecting random error.
5	Numeric	High (4/6 accurate)	Disclosure of the reason for expected inaccuracy. Number of sources of advice: one source vs. two randomly-alternating sources.	Introducing randomness to the advice process increased advice use and attenuated the effect of suspicion.

Table 2. *Factor loadings and reliability scores of measures of chance and epistemic uncertainty in Experiment 3.*

	Factor 1	Factor 2
<u>Chance uncertainty</u>	$\alpha = .674$	
The accuracy of the advice has an element of randomness.	.870	
It feels like the accuracy of the advice is determined by chance factors.	.863	
<u>Epistemic uncertainty</u>		$\alpha = .687$
How accurate the advice will be has been determined in advance.		.756
The accuracy of the advice is knowable in advance, given enough information.		.818
The accuracy of the advice is something that well-informed people would agree on.		.780

Figures




Figure 1. Average rates of advice taking by task round and reason for suspicion setting in Experiment 1. Error bars represent ± 1 SEM.

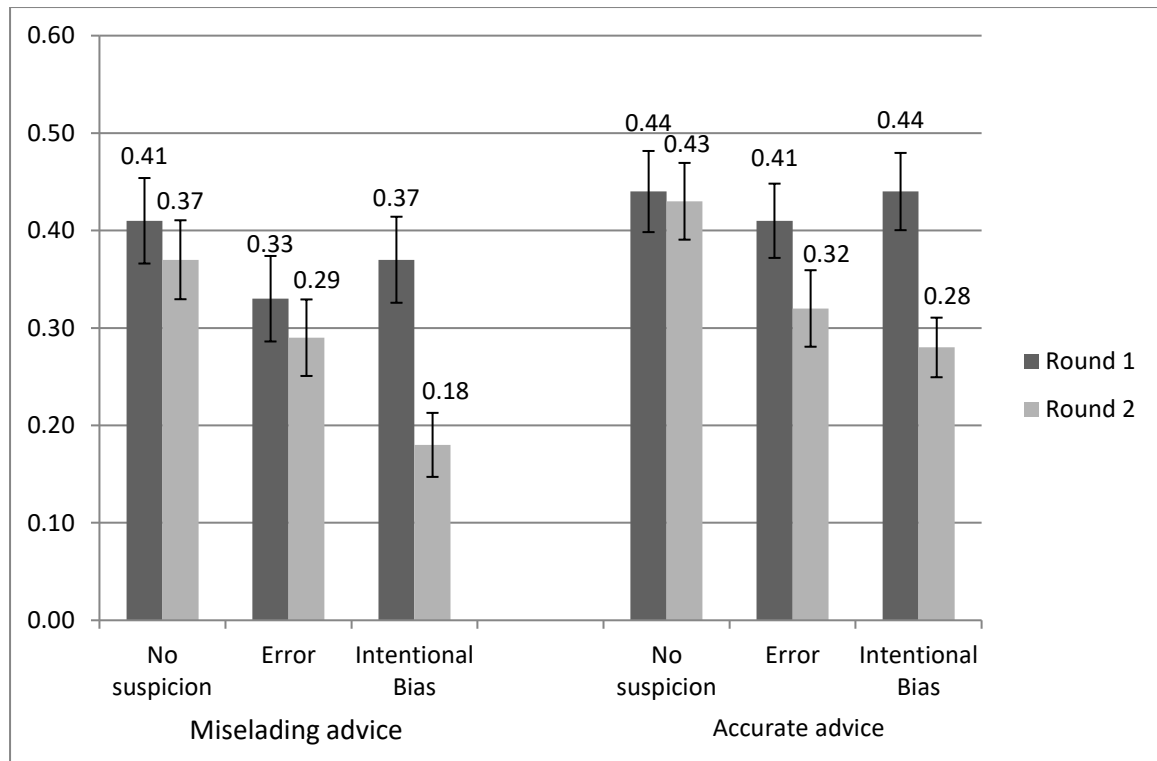


Figure 2. Average rates of advice taking by advice accuracy, task round, and reason for suspicion setting in Experiment 1. Error bars represent ± 1 SEM.




Figure 3. Mean willingness to pay for advice by advice quality and reason for suspicion in Experiment 2. Error bars represent ± 1 SEM.




Figure 4. Average rates of advice use by advice quality and reason for suspicion in Experiment 2. Error bars represent ± 1 SEM.

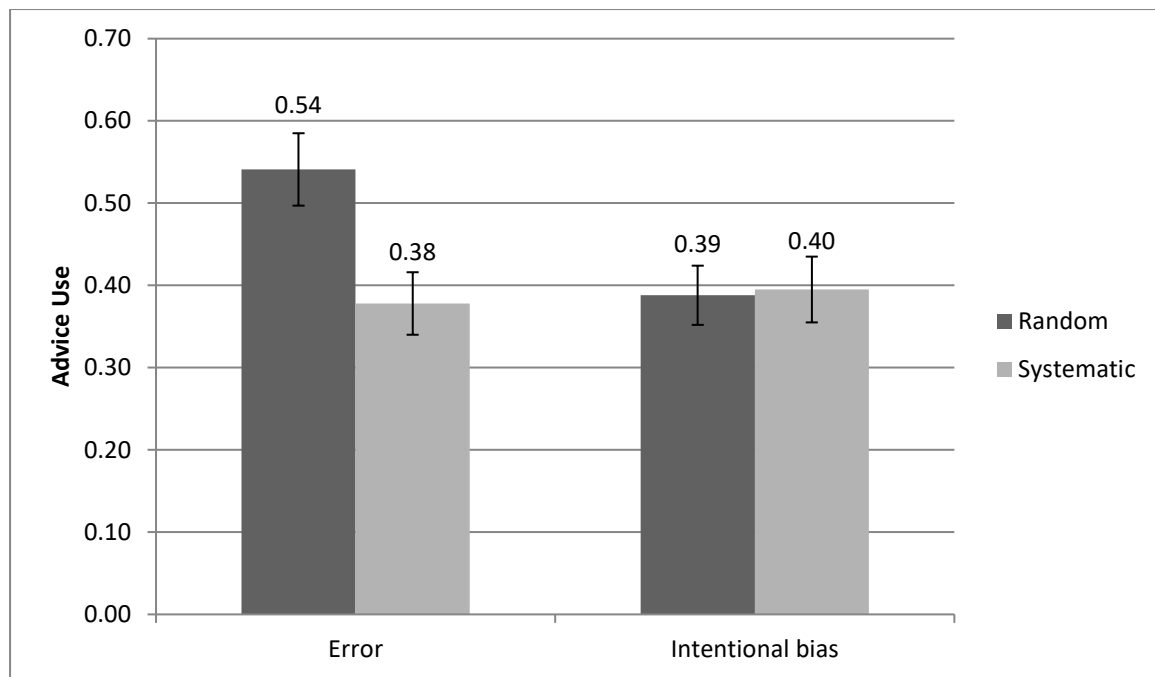


Figure 5. Average rates of advice use by reason for suspicion and direction of prior inaccuracy in Experiment 4. Error bars represent ± 1 SEM.