

Strategically Overconfident (to a Fault): How Self-Promotion Motivates Advisor Confidence

Alex B. Van Zant

Rutgers Business School

Article in press at *Journal of Applied Psychology*

Author Note: Alex B. Van Zant, Department of Management & Global Business, Rutgers Business School, Rutgers University—Newark and New Brunswick. Correspondence concerning this article should be addressed to: Alex Van Zant, 100 Rockefeller Rd, Piscataway Township, NJ 08854. Contact: alex.vanzant@rutgers.edu.

Abstract

Unlike judgments made in private, advice contexts invoke strategic social concerns that might increase overconfidence in advice. Many scholars have assumed that overconfident advice emerges as an adaptive response to advice seekers' preference for confident advice and failure to punish overconfidence. However, another possibility is that advisors robustly display overconfidence as a self-promotion tactic—even when it is punished by others. Across four experiments and a survey of advice professionals, the current research finds support for this account. First, it shows that advisors express more overconfidence than private decision-makers. This pattern held even after advice recipients punished advisors for their overconfidence. Second, it identifies the underlying motivations of advisors' overconfidence. Advisors' overconfidence was not driven by self-deception or a sincere desire to be helpful. Instead, it reflected *strategic* self-promotion. Relative to the overconfidence revealed by their private beliefs, advisors purposely increased their overconfidence while broadcasting judgments when (a) it was salient that others would assess their competence and (b) looking competent served their self-interest.

Keywords: overconfidence, confidence, advice giving, self-promotion, accountability

Strategically Overconfident (to a Fault): How Self-Promotion Motivates Advisor Confidence

Overconfident advice provokes many problems in applied settings. Financial advisors' overconfidence promotes excessive trading and risk exposure (Hackenthal, Haliassos, & Jappelli, 2012). Lawyers' over-optimistic forecasts in trial outcomes influence clients to pursue costly litigation (Goodman-Delahunty, Granhag, Hartwig, & Loftus, 2010; Wistrich & Rachlinski, 2012). Some have even connected state leaders' decisions to engage in costly wars to overconfident intelligence officials and military strategists (Johnson et al., 2006).

Managers and organizational decision-makers can benefit from curbing expert advisors' overconfidence. Organizations often elicit expert forecasts to inform financial decisions (Plous, 1995), hire consultants to inform strategic choices (Gino & Schweitzer, 2008), and solicit employees' domain expertise to enhance team performance (Zhang & Peterson, 2011).

Although many describe overconfidence as a pervasive cognitive bias shaped by how it is measured and assessed (Haran, Moore, & Morewedge, 2010; Juslin, Winman, & Hansson, 2007; Klayman, Soll, González-Vallejo, & Barlas, 1999), advice-giving is a social act. Unlike private decision-makers, advisors face strategic social concerns. Factors like advice seekers' preferences and the need to distill information for others can cause advisors' advice to diverge from their private judgments (Kray, 2000; Kray and Gonzalez, 1999; Jonas & Frey, 2003; Jonas, Schulz-Hardt, & Frey, 2005). However, research on advice confidence has tended to focus on how it impacts advice seekers (e.g., Price and Stone 2004; Sah, Moore, & MacCoun, 2013; Sniezek and Buckley 1995; Sniezek and Van Swol 2001; Vulllioud, Clément, Scott-Phillips, & Mercier, 2017; Yaniv & Foster, 1995). Thus, it is unclear how the strategic social dynamics of advice contexts shape advisors' confidence relative to private decision-makers not subjected to an audience. The current research explores why the dyadic interaction between advisors and advice seekers might

elicit *strategic* overconfidence where advisors purposely inflate their overconfidence above that observed in private decision-makers.

Many have argued that one reason for the prevalence of overconfident advice is that advice seekers fail to hold advisors accountable for their overconfidence (e.g., Radzevick & Moore, 2011; Tetlock, 2005). People tend to reward confident advisors (Price & Stone, 2004) and do not adequately punish overconfidence (Radzevick & Moore, 2011; Ronay, Ostrom, Lehmann-Willenbrock, Mayoral, & Rusch, 2019). Hence, the conclusion that advisors should adapt by displaying overconfidence to exploit its perks appears sensible. But would advisors inflate their confidence in light of signs that others are punishing overconfidence? If overconfident advice is a byproduct of adapting to advice seekers' preferences, advisors should also notice when an advice seeker is punishing their overconfidence and adapt by curbing their confidence. This logic informs the premise that accountability mechanisms enabling advice seekers to identify and penalize overconfidence can curb its pervasiveness (Meikle, Tenney, & Moore, 2016; Sah et al., 2013; Tetlock, 2005; Tetlock & Kim, 1987).

However, another possibility could be that advisors are so accustomed to inflating their confidence that they fail to adapt their approach when advice-seekers react negatively to their overconfidence. This pattern would suggest that punishing overconfidence might not necessarily be a cure-all for overconfident advice.

The current research tests whether advisors' overconfidence reflects an adaptive response to signals about advice seekers' preferences that declines when they punish overconfidence versus a robustly applied self-promotion tactic that occurs independently of feedback about advice seekers' preferences. It makes several theoretical, empirical, and practical contributions to the overconfidence and advice-giving literature. First, it more broadly explores how features of

the dyadic exchange between advice-givers and advice-seekers might shape advisors' overconfidence. Because prior overconfidence research has focused on either private judgments or publicly broadcasted advice without comparing the two (Meikle et al., 2016; Moore, Tenney, & Haran, 2015), it is unclear how the strategic social dynamics and reward systems of advice-giving contexts impact advisors' overconfidence compared to private decision-makers not subjected to these dynamics. The current research theoretically and empirically differentiates between several accounts for the prevalence of overconfident advice.

Second, the current research offers practical takeaways to help managers and other organizational decision-makers reduce their exposure to overconfident advice. This paper's findings have implications for whether punishing overconfidence is enough to reduce advisor overconfidence or whether other approaches that reduce advisors' motive to self-promote can be more effective. Finally, and more generally, the current research shows that audience concerns play a role in shaping overconfidence.

Defining Overconfidence in Advice

Scholars commonly define advice as providing judgments or suggestions to decision-makers solving a problem (Bonaccio & Dalal, 2006; Sniezek & Buckley, 1995). Overconfidence is a joint function of an advisor's confidence in a judgment and the judgment's accuracy. Thus, proving that an advisor is *overconfident* entails asserting that the advisor's confidence is excessively high compared to the advisor's accuracy (Meikle et al., 2016; Moore et al., 2015). Because drawing unbiased conclusions about an advisor's accuracy demands the presence of objective criteria (e.g., a prediction's outcome or claim's veracity),¹ this paper studies advisors'

¹ This is not to say advice seekers cannot adopt subjective criteria for judging an advisor's accuracy in some situations. However, one cannot logically make any claims about whether an advisor is *overconfident* in the absence of an objective accuracy criterion. For example, an advice seeker hoping for others to validate a preexisting belief might

overconfidence in contexts where advisors are providing others with judgments suited for comparison to an objective accuracy criterion.

This research examines overconfidence in the form of overprecision, defined as excessive certainty in the accuracy of one's judgment (Moore & Healy, 2008). Whereas an overconfident advisor assigns high confidence to an inaccurate judgment, a well-calibrated advisor assigns high confidence to an accurate judgment or low confidence to an inaccurate judgment. This paper identifies patterns in overconfidence driven by advisors' confidence above what is warranted by their accuracy.

Although overconfidence can take other forms,² overprecision is the form most directly tied to advisors' function of helping advice-seekers form accurate solutions to problems (Harvey & Fischer, 1997). Research on advice confidence has predominately studied how signals of advisors' accuracy like the probability that an advisor's advice is accurate or the margin of error around the advisor's judgment influence advice seekers' judgment (Bonaccio & Dalal, 2006). Overprecision is more robust than other forms of overconfidence (Moore et al., 2015). Therefore, this paper offers insight into how advice contexts shape the communication and spread of a notoriously pervasive form of overconfidence.

Two Potential Explanations for Overconfident Advice

The current research tests two accounts striving to explain what might cause advisors to inflate their overconfidence compared to private decision-makers. Whereas one assumes that advisors learn from social feedback about advice seekers' preferences and adapt their confidence displays accordingly, the other assumes that they adopt a robust approach of boosting their

perceive an advisor who confidently rejects the belief to be overconfident. Without an objective criterion to judge the advisor's accuracy, this perception could be biased.

² Overconfidence can also occur in the form of overestimating one's absolute or relative performance (Moore and Healy, 2008).

confidence separately from any feedback they receive about advice seekers' preferences.

Although the two accounts predict that advisors should display more overconfidence than private decision-makers when advice-seekers reward it (which is often the case), they make divergent predictions in contexts where advice-seekers punish overconfidence. Thus, studying contexts where advice-seekers punish overconfidence can help explain advisors' overconfidence.

However, the only studies to examine communicators' confidence in repeated interactions have focused on contexts where the costs of being exposed as overconfident do not override the benefits of displaying calibrated confidence (Hertz et al., 2017; Radzevick & Moore, 2011). This focus leaves it unclear which of these accounts might explain patterns in advisors' overconfidence. I advance hypotheses for what these two accounts would predict when the costs of overconfidence exceed any rewards of displaying high confidence.

Social Feedback Account

Perhaps the most popular argument for the prevalence of overconfident advice is that it is an adaptive response to social feedback about advice seekers' preferences. Because people prefer confident advisors (Price & Stone, 2004) and fail to adequately penalize overconfidence (Kennedy, Anderson, & Moore, 2013), advisors should learn that overconfidence is rewarded from their interactions with advice-seekers and exploit its perks (Radzevick & Moore, 2011; Ronay et al., 2019). As articulated by Radzevick and Moore (2011) on advice-seekers' susceptibility to overconfident advice: "...we see these wounds as largely self-inflicted. By rewarding advisors for expressing confidence while not adequately penalizing them for being wrong, customers in the market are essentially 'getting what they paid for'" (p. 99).

Yet if this *social feedback account* is valid, one would also expect that advisors should be cautious when interacting with an advice seeker who is punishing their overconfidence. One

reason overconfidence frequently goes unpunished is that people do not arm themselves with feedback about advisors' accuracy (Sah et al., 2013). However, when given clear feedback about advisors' accuracy, advice-seekers reward accurate (rather than confident) advisors and punish overconfident ones (Sah et al., 2013; Tenney, MacCoun, Spellman, & Hastie, 2007; Tenney, Spellman, & MacCoun, 2008; Vullioud et al., 2017). If the risks of displaying confidence outweigh its rewards, advisors' best approach should be displaying tempered confidence (cf. Johnson & Fowler, 2011). Punishing advisors' overconfidence may therefore be a simple and powerful prescription for obtaining well-calibrated advice. In light of evidence that their overconfidence is getting them into more trouble than it is worth, advisors should adjust by reducing their overconfidence.

Social Feedback Account Hypothesis 1: Advisors exhibit less overconfidence than private decision-makers *when given feedback that an advice seeker punishes overconfidence.*

It is also worth noting that when advice-seekers punish overconfidence, advisors should observe that they only penalize judgments expressed with high confidence when they turn out to be inaccurate. Advisors' overconfidence should therefore be sensitive to signals about their accuracy. In light of signs that advice seekers punish confidence when advisors attach it to inaccurate judgments, rational advisors acting in their self-interest should display less confidence in future judgments after learning that they are inaccurate. Because it reveals their poor calibration, accuracy feedback helps mitigate decision-makers' overconfidence (Minson & Umphres, in press; Moore et al., 2017; Moore & Schatz, in press). While this suggests that private decision-makers may temper their confidence after learning that their judgment is inaccurate, social feedback that an advice seeker is punishing their overconfidence should

provide advisors with an added motive to temper their confidence after an inaccurate judgment. Compared to private decision-makers, advisors could more quickly curb their overconfidence after getting negative feedback about their accuracy.

Social Feedback Account Hypothesis 2: Compared to private decision-makers, advisors should more strongly reduce their overconfidence after learning that they made an inaccurate judgment (versus making an accurate judgment) *when given feedback that an advice seeker punishes overconfidence.*

Strategic Self-Promotion Account

Another possibility is that advisors so robustly rely on overconfidence to serve their interests that they fail to adapt their approach in the presence of signals that an advice seeker punishes overconfident advisors. This *strategic self-promotion account* argues that advisors robustly employ a tactic of displaying confidence when motivated to project competence to an audience, despite any feedback they receive about an advice seeker's preferences. Actors engage in strategic self-promotion when they purposely attempt to project competence to others (Godfrey, Jones, & Lord, 1986; Jones & Pittman, 1982).³

Confident people often get ahead and seldom see their overconfidence punished by others (Anderson, Brion, Moore, & Kennedy, 2012; Kennedy et al., 2013). Observing this link, advisors may respond by deliberately inflating their confidence across many situations that motivate them to self-promote and fail to adapt their approach when faced with evidence that overconfidence hinders their self-promotion efforts. Although economic theory often assumes communicators

³ In a similar vein to how economists differentiate between behavior that is “strategic” versus “rational” or “optimal” (e.g., Bernheim, 1984; Pearce, 1984), I define behavior motivated by self-promotion to be strategic to the extent that it one enacts it because it is *perceived* to be effective at enhancing his or her competence in others’ eyes (cf. Goffman, 1959). Indeed, despite being strategic, self-promotion attempts can backfire when actors hold misguided beliefs about the efficacy of the tactics they employ (Bozeman & Kacmar, 1997; Schmitt & Buss, 1996).

employ tactics that maximize their chances of eliciting desired behaviors from others (DellaVigna 2009; Köszegi, 2014), communicators sometimes use suboptimal tactics (Daniels & Zlatev, 2019). At times, these tactics can be such deeply ingrained strategies that communicators fail to adjust when faced with evidence that they should adopt a different approach (Zlatev, Daniels, Kim, & Neale, 2017). Human behavior is variable and results in noisy feedback that constrains actors' ability to learn about the causal links between their actions and others' subsequent behavior (Ball, Bazerman, & Carroll, 1991; Brehmer, 1980; Tversky & Kahneman, 1986). In this vein, advisors might fail to learn from signals that others are punishing overconfidence. They may be so conditioned to display overconfidence as a means of strategic self-promotion that they continue to do so even when given evidence that overconfidence runs counter to their self-interest. Advisors are subject to social and economic concerns that might cause them to project competence (Jonas et al., 2005; Radzevick & Moore, 2011). Despite others' attempts to hold them accountable by punishing their overconfidence, advisors might display more overconfidence than private decision-makers.

Strategic Self-Promotion Account Hypothesis 1: Advisors exhibit more overconfidence than private decision-makers *independently of feedback that an advice seeker punishes overconfidence.*

Suppose self-promotion drives advisors to continue inflating their confidence in light of signals that an advice seeker is punishing their overconfidence. In that case, feedback about their accuracy that exposes their confidence as inappropriately excessive (i.e., overconfidence) might also fail to change their approach. One viable possibility is that advisors respond to feedback about their accuracy in the same manner as private decision-makers. This pattern would suggest that, while feedback about their accuracy may help advisors calibrate their confidence like it

helps private decision-makers, the added social feedback they receive about advice-seekers punishing overconfidence does not lessen advisors' overconfidence in future judgments.

Two distinct patterns would even suggest that negative feedback about their accuracy might heighten advisors' inflated confidence. First, advisors could be less sensitive to feedback about their accuracy than private decision-makers and adopt an ongoing strategy of inflating their confidence irrespective of their accuracy. Whereas private decision-makers may reduce their confidence after learning about their inaccuracy, advisors may not reduce their confidence to the same degree. This pattern could increase the degree to which advisors inflate their overconfidence compared to private decision-makers after inaccurate judgments.

Strategic Self-Promotion Account Hypothesis 2a: Compared to private decision-makers, advisors should less strongly reduce their overconfidence after learning that they made an inaccurate judgment (versus making an accurate judgment) *when given feedback that an advice seeker punishes overconfidence.*

It could also be the case that advisors inflate their overconfidence after inaccurate judgments. Such a pattern would be evidence of compensatory self-promotion (e.g., Baumeister & Jones, 1978; Frey, 1978) where advisors overcompensate for their inaccuracy by bolstering their confidence in future predictions to signal that "last time was a fluke, but this time is different." Of all the strategic self-promotion patterns discussed above, this would result in the most significant divergence between advisors' and private decision-makers' overconfidence after inaccurate judgments.

Strategic Self-Promotion Account Hypothesis 2b: Compared to private decision-makers, advisors should increase their overconfidence after learning that they made an

inaccurate judgment (versus making an accurate judgment) *when given feedback that an advice seeker punishes overconfidence.*

Overview of Studies

One challenge of studying overconfidence in advice is that advisors' overconfidence can rarely be measured objectively with field data (Radzevick & Moore, 2011). A second challenge is that, even if one were to assess advisors' overconfidence from real-world discussions with advice seekers, advisors' private beliefs would be unobservable. To address these challenges, the studies in this paper adopt a similar approach to Radzevick and Moore (2011) by using experimental settings to study advisors' exchange with advice recipients. Like the expertise asymmetries inherent in most advice contexts, these studies involve dyadic interactions where focal participants endowed with relative expertise (advisors) convey judgments to a target (advice seeker) lacking their privileged information (e.g., Radzevick & Moore, 2011; Sah et al., 2013). Each study compares the overconfidence of advisors broadcasting their judgments to an advice seeker to private decision-makers.

Study 1 tests whether advisors' overconfidence exceeds that of private decision-makers in a repeated decision-making context where advisors receive noisy feedback that the penalty for overconfidence outweighs the reward for calibrated confidence. This paradigm permits a direct test of whether advisors' overconfidence reflects a process of learning from feedback about advice seekers' preferences and adapting in response (social feedback account) versus a robustly applied self-promotion tactic (strategic self-promotion account). Studies 2-5 focus on identifying why advisors' overconfidence diverges from private-decision makers and how it manifests. Study 2 more directly tests whether self-promotion motivates advisor overconfidence by testing whether advisors' overconfidence primarily diverges from private decision-makers when it is

highly salient that advice-seekers will assess their competence. In the process, it attempts to establish that advisors' overconfidence is deliberate and not a byproduct of self-deception (e.g., Chance, Norton, Gino, & Ariely, 2011). Studies 3 and 4 test whether advisors' overconfidence can be explained by self-interested strategic behavior instead of a sincere desire to help advice-seekers with informative judgments (cf. Yaniv & Foster, 1995, 1997). Finally, Study 5 tests for evidence of strategic self-promotion in the self-reported overconfidence of advice professionals.

Data and materials for all studies are accessible at

https://osf.io/gtypr/?view_only=f06e733ef13644e5b9a1de26c5157dc1. Planned sample sizes,

exclusion criteria, and analyses for Studies 1 and 5 are preregistered (Study 1:

https://osf.io/wp6c7/?view_only=bac75b85b76b4a05866d1aadb3001645, Study 5:

<http://aspredicted.org/blind.php?x=ni5ab6>). I report all data exclusions (if any),⁴ all

manipulations, and all measures. As recommended by Bliese and Wang (in press), I also report observed post-hoc power (denoted as $1-\beta$), which gives the cumulative probability of concluding that a specific effect is statistically significant (at $p < .05$) using an identical sample and statistical model. I analyzed the data for Studies 1 and 5 using Stata MP (version 16.1). For Studies 2-4, I used IBM SPSS Statistics (version 27).

Study 1: Does Overconfident Advice Persist When it is Punished?

Study 1 tests the social feedback and strategic self-promotion accounts in a setting where overconfidence is a liability to advisors. Participants made a series of predictions and indicated their confidence in each prediction. Some participants were randomly assigned to a control group of private decision-makers who did not broadcast their judgments to an advice seeker. In

⁴ In Studies 2-4, a small percentage of responses corresponded to IP addresses that appeared more than once in the dataset (3% of responses in Studies 2 and 4, 8% of Study 3 responses). Because I did not make an a priori decision to omit them from the dataset, I report results with these responses included; all reported effects hold in analyses that omit these responses. No duplicate IP addresses appeared in the dataset for Studies 1 and 5.

contrast, other participants were assigned to an advisor role condition designed to mimic the evolution of advisor-client relationships in applied settings.

Advisors broadcasted their confidence to a less-informed partner and were incentivized to have this partner elect to receive their advice; this type of incentive has been used in prior research to mimic the incentives inherent in many advice contexts where advisors are rewarded for convincing clients to “hire” them as an advisor (Radzevick & Moore, 2011). Advisors were punished for overconfidence: Their partner's probability of electing to receive their advice in the next round dropped directly with advisors' confidence following inaccurate predictions and was not impacted by advisors' confidence after accurate predictions. In the interest of testing the social feedback and strategic self-promotion accounts under conditions where the penalties for overconfidence are ecologically representative of the degree to which advice seekers are likely to punish overconfident advisors, advice seekers' decisions were simulated from the decisions of a separate sample of research participants who played the advice-seeker role.

Importantly, advisors received feedback about their accuracy and the hiring decisions of advice-seekers. If the social feedback account can explain advisors' overconfidence, then advisors should observe that displaying confidence is detrimental to their ability to self-promote in this context where an advice seeker punishes their overconfidence. This social feedback should cause their judgments to reflect *less* overconfidence than private decision-makers (Social Feedback Account Hypothesis 1), and they should temper their confidence more than private decision-makers after inaccurate judgments (Social Feedback Account Hypothesis 2).

However, if the strategic self-promotion hypothesis can account for advisors' overconfidence, they should not adapt to feedback that it is punished by an advice seeker. Consequently, they should display more overconfidence than private decision-makers (Strategic

Self-Promotion Account Hypothesis 1). Further, although advisors' overconfidence compared to private decision-makers might stay elevated to a similar level after learning about their (in)accuracy, it also remains possible that their relatively inflated overconfidence might increase even further after learning they made an inaccurate judgment (Strategic Self-Promotion Account Hypotheses 2a and 2b).

Method

Participants. Based on a preregistered data collection rule, I recruited 216 Amazon Mechanical Turk (MTurk) workers to complete the study ($M_{\text{Age}} = 34.39$ years, $SD = 10.34$, 36% female). In this and all subsequent studies, participants were not allowed to proceed to the main experiment until they successfully passed a comprehension check verifying their understanding of study procedures.

Procedure. Participants learned at the outset of the study that they would make a series of predictions about the value of stocks based on information about each stock's price at the beginning of each month over twelve months (selected at random from the S&P 500 and labeled stocks "A-J" to obscure their identity). For each stock, participants predicted whether the stock's value would be greater than or equal to a reference price at the start of the subsequent month ("Month 13"); the reference price was always equivalent to the mean of the stock's value over the first three months of the twelve-month data period. In the interest of providing advisors with a salient partner in a follow-up interaction and to reduce any potential suspicion about the validity of their partner's choices (which were simulated based on the decisions of previous participants), participants were then paired with another MTurk worker in a chatroom and prompted to exchange strategies for making accurate predictions during the study.

After exchanging strategies with a partner, participants were randomly assigned to an advisor or private decision-maker role. Regardless of their role, participants learned about several features of the study. First, in contrast to the twelve months of price data that they would use to make predictions about each stock, some other study participants would only be provided with access to the first three months of data for this period. The purpose of this instruction was to provide advisors and private decision-makers alike with the sense that they were in a position of relative expertise compared to some other participants. Because the mean of the three months of data that other participants could access determined the reference price, participants were aware that they had access to additional data to enhance their prediction accuracy.

Second, participants learned that they would record their level of confidence in their predictions by indicating the probability that a given prediction is accurate on a scale ranging from 50% to 100% ("There is a ___% chance that I accurately determined whether the stock's value will be above or below \$XX.XX at the start of Month 13; XX.XX was filled with the stock's reference price). Third, participants learned that they would receive information about each stock's value after making predictions and indicating their confidence for a given stock, but before proceeding to predict the next stock's value. Participants were guided through a sample prediction as they read about the procedure to reinforce their understanding.

Advisors. Participants assigned to the advisor role learned that their confidence for a given prediction would be broadcasted to their partner. Advisors were then informed that their partner—who was incentivized to be accurate—would have the opportunity to access their predictions in exchange for a small fee after seeing their confidence. This aspect of the study design simulates the dynamics of advice markets where advisors can publicly signal their confidence in their ability to provide sound advice but do not necessarily broadcast specific

information like formal predictions unless a client has enlisted their services (Radzevick & Moore, 2011). Advisors were also informed that irrespective of their partner's decision to access their prediction in the prior round, their partner would receive feedback about each prediction's accuracy between rounds. Thus, advisors were aware that they had an opportunity to build a reputation with their partner. To reinforce this instruction, participants read the following message after learning about their partner's decision of whether to access their prediction, but just before proceeding to predict the subsequent stock:

“The stock’s value ended up being [above / below] \$XX.XX at the start of Month 13. You predicted that the stock’s value will be [above / below] \$XX.XX. Your partner will receive this same information about the stock’s actual value in Month 13 and your prediction after making his or her own prediction for this stock, but before seeing your confidence on the next one.”

In line with the real-world incentives advisors typically earn for enlisting clients, advisors learned that they would earn one ticket for a \$50 bonus payment raffle each time their partner elected to access their prediction for a given stock.

Advisors also received feedback about whether their partner elected to "hire" them by accessing their prediction in a given round. They learned about their partner's decision before proceeding to the next prediction. Although the instructions did not specify who advisors' partner would be, they were written to elicit the pretense that their partner was the same person they interacted with earlier in a chatroom. In reality, advisors' "partner" was an algorithm that initially rewarded confidence but punished overconfidence after advisors' initial prediction (i.e., the time at which advice-seekers would have feedback about advisors' accuracy). To simulate partners' hiring decisions in an ecologically valid manner that captures variance in the actual choices of

advice-seekers (cf. Brunswik, 1955), I derived the algorithm's parameters from a model predicting the choices of a separate sample of advice-seekers based on their prior hiring decision, along with advisors' confidence, prior confidence, and prior prediction accuracy.⁵

Importantly, for advisors' initial prediction, each percentage point increase in their confidence resulted in a 5% increase in their odds of being hired. However, in subsequent rounds (where advisors and advice seekers alike had information about advisors' accuracy in prior rounds), advisors' confidence had no bearing on their odds of being hired unless they were inaccurate: For each inaccurate prediction, advisors' odds of being hired *decreased* by 5% with every percentage point increase in their confidence level. On balance, the algorithm rewarded tempered confidence more than high confidence.

In total, advisors completed ten rounds of making predictions, indicating their confidence, and receiving feedback about their accuracy and their partner's hiring decision. This number of rounds was selected based on prior research demonstrating that ten rounds provides enough feedback about others' behavior to elicit adjustments in communicators' influence tactics yet is also approximately the point at which communicators fail to show evidence of further learning in response to additional feedback (Zlatev et al., 2017). Critically, advisors were not told how many rounds of predictions they would complete or warned before the final round.

Private Decision Makers. Participants assigned to the private decision-maker role made predictions, indicated their confidence, and received feedback about their accuracy across the same ten rounds of predictions as advisors. However, they were informed that their predictions and confidence would be private and not shown to any other participants. As such, they were not

⁵ In a separate study conducted prior to Study 1, 101 advice recipients followed the procedure of advisors' partner as described to advisors. See the Supplemental Online Materials (SOM) for more detail about advice recipients' procedure (Study S1) and the process of developing an algorithm that could reasonably model their choices while ensuring that overconfidence was penalized on the balance.

paired with a partner and did not receive feedback about their partner's hiring decisions. Instead of being incentivized to be hired by a partner, private decision-makers learned that they would be entered into a \$50 bonus payment raffle for merely completing the study.

Measuring Overconfidence. In this and all subsequent studies, I assessed the magnitude of advisors' overconfidence by regressing each participant's confidence on their accuracy and retaining the standardized residuals. The standardized residuals represent participants' confidence above and beyond what could have been predicted by their accuracy.⁶ To verify that any effects on participants' overconfidence reflected patterns in their confidence (rather than an artifact of their accuracy), I replicated all analyses of participants' residual overconfidence on their raw (i.e., unresidualized) confidence; all effects reported in this manuscript hold (see SOM).

Results

Main analyses. Following a preregistered data analysis plan, I analyzed participants' overconfidence across rounds using a multilevel linear regression model nesting rounds within participants using an advisor-specific random intercept.⁷ The results exclude three participants (one in the private decision-maker condition and two in the advisor condition) who expressed suspicion about whether their partner in either the chatroom (both roles) or main study (advisors only) was an actual person; all effects reported below hold when including these participants in analyses. Table 1 reports descriptive statistics and the correlations between variables. Table 2 shows the results of the models used for the main analyses.

⁶ Researchers have recommended this approach for measuring overconfidence because accuracy-based components like simple difference scores (i.e., subtracting accuracy from confidence) are primarily influenced by the component with higher variance (Edwards, 1994; Kennedy et al. 2013). Because overprecision is necessarily confounded with accuracy (Klayman et al. 1999; Olsson, 2014), accuracy can obscure any effects driven by confidence judgments if it is not partialled out.

⁷ I decided after completing data collection to focus the paper on overconfidence rather than confidence. As a result, I deviated from the preregistered plan to measure advisors' "confidence." All analyses of overconfidence reported in this manuscript replicate for participants' raw confidence that is unadjusted for accuracy (for analyses of participants' raw confidence in Study 1, see Table S3 of the SOM).

Participants exhibited a robust pattern of overconfidence, as their mean confidence of 81% exceeded their prediction accuracy of 53%, $t(215) = 24.85$, $p < .001$, $d = 1.69$, 95% CI = [1.48, 1.90]. This effect held for both advisors and private decision-makers (both $ps < .001$).

Relative to private decision-makers, advisors were more overconfident across predictions (Model 1: $p = .022$, $d = 0.32$, 95% CI = [0.05, 0.59], $1-\beta = 63\%$; Model 2: $p < .02$, $d = 0.33$, 95% CI = [0.05, 0.59], $1-\beta = 64\%$). An exploratory analysis failed to find evidence of this effect being qualified by round (Model 4: $p = .21$), indicating that advisors' relatively inflated overconfidence persisted to a similar degree across rounds of receiving feedback that overconfidence was punished. Despite exposure to repeated feedback that the risks of overconfidence outweighed the benefits of confidence, advisors displayed more overconfidence than private decision-makers in the tenth and final round of predictions ($M_{\text{Advisor}} = 0.15$, $SD = 0.09$ vs. $M_{\text{Decision-Maker}} = -0.19$, $SD = 0.10$), $t(211) = 2.57$, $p = .011$, $d = 0.35$, 95% CI = [0.08, 0.62], $1-\beta = 72\%$. Taken together, this is consistent with the strategic self-promotion account (Strategic Self-Promotion Account Hypothesis 1) and inconsistent with the social feedback account (Social Feedback Account Hypothesis 1).

Some evidence of an interaction between role and prior round accuracy also emerged (Model 3: $p = .085$, Model 4: $p = .062$). A simple slopes analysis of Model 3 revealed that private decision-makers' overconfidence was reduced after receiving negative feedback about their accuracy in the prior round, $B = -0.11$ ($SE = 0.05$), $z = 2.31$, $p = .021$, $d = 0.15$, 95% CI = [0.02, 0.28], $1-\beta = 64\%$. In contrast, advisors' overconfidence was insensitive to feedback about their accuracy, $B = 0.01$ ($SE = 0.05$), $z = 0.12$, $p = .90$, $d = 0.01$, 95% CI = [-0.12, 0.14]. These findings provide further evidence that diverges from the social feedback account. Witnessing advice-seekers punish their overconfidence did not increase advisors' responsiveness to feedback

that a prior judgment was inaccurate (Social Feedback Account Hypothesis 2). If anything, advisors were relatively insensitive to feedback about their accuracy (Strategic Self-Promotion Account Hypothesis 2a). Unlike private decision-makers, who reduced their overconfidence after learning that they were inaccurate, advisors' overconfidence persisted to a similar degree after receiving direct feedback about their poor calibration.

Post-hoc tests of the robustness of the strategic self-promotion account. Unlike private decision-makers, who only received feedback about their prediction accuracy, advisors also received feedback about a partner's hiring decision after each round. Advisors were hired in an average of 3.01 out of the nine prior rounds ($SD = 1.30$), entering the final round of predictions. At this hiring probability, a decrease of one standard deviation from advisors' mean confidence ($M = 82.10$, $SD = 13.79$) increased their probability of being hired in the subsequent round by 11%. A one standard deviation increase from their mean confidence at this hiring probability was associated with a 6% *decrease* in their subsequent hiring probability (see Table S5 of SOM for the impact of confidence on hiring outcomes for a prototypical advisor). However, the algorithm used to simulate advisors' hiring outcomes contained built-in variability that mimicked variance in advice-seekers' decisions. Thus, the degree to which an advice seeker punished overconfidence varied between advisors.

To test the robustness of the strategic self-promotion account, I examined how strong advisors' social feedback that an advice seeker punishes overconfidence needed to be for their overconfidence to align with private decision-makers.⁸ To capture the strength of advisors' feedback, I took the correlation between the parameter in the advice-seeker algorithm capturing an advisor's overconfidence (0 if accurate, confidence-50 if inaccurate) and the advisor's

⁸ I did not preregister these analyses. I thank the Associate Editor and two anonymous reviewers for raising concerns that inspired this analysis.

subsequent-round hiring outcomes (0 if not hired, 1 if hired) at the start of each round. I then multiplied the resulting correlation by negative one to arrive at *feedback strength* variable where positive values reflect a negative observed correlation between advisors' overconfidence in one round and their hiring likelihood in the subsequent round ($M = 0.25$, $SD = 0.50$).⁹

An analysis of advisors' overconfidence in Round 10—the point at which they would have received the most feedback possible about the extent to which an advice seeker penalized their overconfidence—revealed that they were more overconfident than private decision-makers across a wide range feedback strength values. As documented in Table 3, advisors were generally more overconfident than private decision-makers up to feedback strength values of 0.5. Additional analyses across all rounds of predictions also generally support the conclusion that advisors' inflated overconfidence persisted for feedback strength values in the 0.4-0.5 range (see SOM). These results, which I found additional support for in a separate set of analyses (see SOM), suggest that advisors were insensitive to feedback indicating that their advice recipient penalizes overconfidence. Thus, the strategic self-promotion account seems to hold even when advice-seekers apply a reasonably hefty penalty to overconfident advisors. It was not until isolating analyses to advisors exposed to correlations between prior overconfidence and hiring outcomes more strongly negative than approximately $r = -0.5$ that it becomes inconclusive whether they continued to exhibit more overconfidence than private decision-makers.

Discussion

The results of Study 1 support the strategic self-promotion account and oppose the social feedback account. Instead of adapting their approach, advisors' overconfidence persisted in the

⁹ Formal details about the measure's computation are in the SOM. I could not compute the feedback strength variable in any rounds for seven advisors whose accuracy, confidence, or hiring outcomes were invariant in Rounds 2-9. I did not include these advisors in analyses of feedback strength.

face of evidence that an advice seeker was on balance punishing their overconfidence (Social Feedback Account Hypothesis 1). Unlike private decision-makers, who reduced their overconfidence after inaccurate judgments, advisors' overconfidence was insensitive to feedback about their accuracy (Strategic Self-Promotion Account Hypothesis 2a). While not as conclusive as can be, this pattern raises the possibility that learning they are inaccurate might only exacerbate the degree to which advisors inflate their overconfidence compared to private decision-makers. At the very least, the pattern runs counter to the notion that social feedback about the penalty associated with overconfidence provides advisors with a powerful motive to temper their confidence after an inaccurate judgment (Social Feedback Account Hypothesis 2).

However, there are likely to be *some* circumstances where the social feedback account holds. I designed Study 1 with applicability to applied advice contexts in mind. Thus, one should be careful in concluding that punishing overconfidence *never* reduces advisors' overconfidence. The social feedback account might hold under conditions structured to maximize advisors' ability to learn about the penalties associated with overconfidence. Advisors received noisy feedback about their partner's preferences across only ten rounds of predictions. Although their overconfidence did not appear to be impacted by feedback, it could be possible that advisors require a much larger quantity of social feedback than provided in this study before they begin to show signs of tempering their overconfidence. If advisors received large quantities of feedback that an advice seeker *never* hired them after their overconfidence was exposed and *always* hired them after every well-calibrated prediction, they would stand a better chance of learning to temper their confidence.

Having acknowledged that advisors might learn to temper their overconfidence under conditions optimized for learning, it is worth questioning whether establishing these conditions is

realistic in many applied settings. This study's advisors received relatively instantaneous feedback about their accuracy and advice seekers' hiring decisions. Many biases go uncorrected in the real world because there is a considerable delay between actions and feedback about the consequences of those actions (Tversky & Kahneman, 1986). From the perspective of an advisor, it might often be months or years before learning about a prediction's accuracy—let alone observing the repercussions associated with the confidence one attached to a prediction. Therefore, even in contexts where advisors receive a large quantity of feedback from an advice seeker, delays in observing the consequences of their confidence displays are likely to hinder their ability to adapt in situations where the advice seeker punishes their overconfidence.

A second issue worthy of consideration is whether the social feedback account might hold if advice-seekers consistently apply stringent enough penalties for advisors to recognize that they are punishing overconfident advice. The post-hoc analyses of Study 1 indicated that advisors' inflated overconfidence persisted even among managers who had observed a correlation between overconfidence and hiring outcomes more extremely negative than $r = -0.4$. This correlation is stronger than the correlations observed in tightly controlled studies explicitly designed to test for evidence of advice-seekers punishing overconfidence¹⁰ and double the magnitude of the average published effect size in social cognition research ($r = 0.2$; Richard, Bond, & Stokes-Zootka, 2003). The sample of advisors who observed correlations between overconfidence and hiring outcomes more extremely negative than $r = -0.5$ was likely too small to adequately test for whether their inflated overconfidence persisted in the face of more substantial penalties ($N = 16$). However, the results indicate that the magnitude of the relationship between overconfidence and hiring outcomes advisors would need to observe before

¹⁰ For example, Sah et al. (2013) find correlations between advisor overconfidence and ratings of advisor credibility of $r = -.23$ (Study 1) and $r = -.28$ (Study 2).

learning to curb their overconfidence is likely substantially higher than the vast majority of causal relationships observed in real-world social interactions.¹¹

One final issue worth noting is that, unlike private decision-makers, advisors received an instruction between rounds of predictions (i.e., after learning about their accuracy in a given round, but before proceeding to the next round) reminding them that advice seekers would be able to see their predictions and the outcome of each prediction before observing their confidence in the subsequent round. The purpose of this instruction was to increase advisors' awareness that, irrespective of their partner's hiring decision, their partner could use information about their accuracy to hold them accountable in subsequent rounds of predictions.

Although this instruction intended to increase advisors' likelihood of recognizing that advice seekers could punish their overconfidence, it could have been written more clearly in retrospect. A more explicit statement about the link between overconfidence and hiring outcomes might have made it easier for advisors to notice that an advice-seeker punished their overconfidence (e.g., "your partner will have the opportunity to hold you accountable in the next round if you are overconfident"). However, one could also argue that such instructions would undermine the external validity of the current study. While advisors consider how their actions will impact their reputation (Sah & Loewenstein, 2014), they do not regularly see overt reminders that overconfidence could harm their reputation.

Study 2: Is Advisors' Overconfidence Deliberate or Self-Deceptive?

Study 1 provides evidence that, under parameters designed to resemble the degree to which actual advice-seekers punish overconfidence, advisors inflate their overconfidence above that of private decision-makers. This finding supports the strategic self-promotion account and

¹¹ Only 5% of all published effects in the field of social psychology exceed $r = 0.5$ (Richard et al., 2003).

runs counter to the social feedback account, which assumes that advisors' degree of overconfidence is a byproduct of social feedback about whether an advice seeker rewards or punishes overconfidence. Study 2 begins a series of studies that attempt to more precisely identify how advisors' strategic self-promotion manifests and differentiate the strategic self-promotion account from other similar accounts for the prevalence of overconfidence.

The first goal of Study 2 is to examine whether the strategic self-promotion account holds in the absence of feedback about advisors' accuracy and advice seekers' preferences. Although providing advisors immediate feedback was critical to testing for whether their overconfidence reflects a process of learning from social feedback in Study 1, most real-world environments do not allow for immediate feedback (Tversky & Kahneman, 1986). The mere presence of an opportunity to learn from social feedback can sometimes impact individuals' behavior by encouraging them to adopt a flawed hypothesis testing strategy where they engage in a specific behavior with the intent of confirming a hunch that it will yield desirable outcomes (Klayman & Ha, 1987). For example, because advisors in Study 1 were aware that they would be able to make adjustments after observing their advice seeker's decisions, they might have inflated their confidence to confirm a tentative hunch that it would be rewarded and continued to inflate their overconfidence after failing to notice that the advice seeker punished overconfidence. In the absence of social feedback about an advice seeker's preferences, advisors may be more reluctant to act on a hunch that an advice seeker will reward overconfidence. Because the strategic self-promotion account assumes that advisors' inflated overconfidence holds independently of any social feedback they might receive, demonstrating that it holds in its absence is crucial. If advisors' overconfidence reflects a robustly applied self-promotion tactic, they should continue

to exhibit more overconfidence than private decision-makers in the absence of feedback about advice seekers' preferences.

Second, Study 2 considers whether audience concerns impact advisors' overconfidence. Although a necessary precondition of self-promotion is that others can observe one's behavior, situations can vary in the extent to which people recognize that they can shape others' impressions (Leary & Kowalski, 1990). People primarily engage in self-promotion when they know their behavior can impact how others evaluate their attributes (Hewitt et al., 2003; Nezlek & Leary, 2002). If self-promotion drives advisors' overconfidence, their overconfidence should increase when they perceive that others will evaluate their competence.

A final goal of Study 2 is to test whether advisors' overconfidence reflects *deliberate* strategic behavior, as opposed to self-deception. Whereas strategic self-presentation represents a deliberate attempt to appear competent to others (Godfrey et al., 1986; Jones & Pittman, 1982), self-deception represents genuinely delusional beliefs that persist when accurate judgments are a primary goal (Chance et al., 2011; Schwardmann & Van der Weele 2019). Von Hippel and Trivers (2011) argue that self-deception is an adaptive behavior that enhances individuals' success at self-presentation; if one's goal is to convince others about his or her competence, then those efforts should be more successful if the person genuinely believes that he or she is highly competent. Consistent with this assertion, the mere desire for others to perceive oneself as competent has the potential to elicit genuine overconfidence that manifests in individuals' private judgment (Anderson et al., 2012; Schwardmann & Van der Weele, 2019). The potential for self-deception raises the possibility that advisors' overconfidence might not be a byproduct of strategic self-promotion at all. Instead, it could reflect delusional beliefs activated by advisors' motive to appear competent when they perceive that others will evaluate their competence.

Thus, Study 2 elicited participants' judgments twice: First, as advisors making public judgments and then as private decision-makers. This ordering was deliberate to maximize the chances of detecting an effect consistent with self-deception, as people are particularly prone to self-deception after claims about their competence are made public (Chance et al., 2011). Because private decision-makers in Study 1 did not have an incentive to honestly report their confidence (which is critical to asserting whether participants genuinely believe any inflated claims of confidence they made as advisors), Study 2 participants were incentivized to exhibit well-calibrated judgment while playing the private decision-maker role. If self-deception can account for advisors' overconfidence, then their subsequent overconfidence as private decision-makers should be influenced by their awareness that others scrutinized their competence while they played the advisor role.

However, if strategic self-promotion shapes advisors' overconfidence, then their public and private judgments should diverge. Although advisors should self-promote by deliberately inflating their confidence when they strongly perceive that others will evaluate their competence, these considerations should not influence their subsequent private judgments. To the extent that advisors' overconfidence reflects strategic self-promotion, they should inflate their advice confidence compared to the confidence they later display as private decision-makers when it is highly salient that others will evaluate their competence.

Method

Participants. I posted 360 assignments to MTurk and stopped data collection once they were all submitted, which resulted in a total of 362 MTurk workers completing the experiment ($M_{\text{Age}} = 34.02$ years, $SD = 11.16$, 40% female).

Procedure. Study 2 followed a 2 (role: advisor, private decision-maker) X 2 (confidence frame: evaluative, non-evaluative) mixed design with confidence frame manipulated between-subjects and role varying within subjects. The study's beginning closely followed the procedure for advisors in Study 1, but with three exceptions. First, advisors did not receive feedback about their prediction accuracy or the advice seeker's decision to accept their advice. Instead, the advice seeker was described as a participant who would complete the study at a later point in time.¹² The purpose of this modification was to assess advisors' use of strategic overconfidence in the absence of social feedback. However, this does not mean that participants did not have the opportunity to establish a reputation with the advice seeker: As with Study 1 advisors, participants were informed that advice seekers would receive feedback about each stock's actual price. As a second modification, participants evaluated only half of the stocks presented to Study 1 participants (randomized order). Because participants played the role of both advisor and private decision-maker, five of the stocks from Study 1 were presented to participants in the interest of keeping the study length more manageable.

Third, I manipulated confidence frame. To do this, I capitalized on prior research demonstrating that the language used to produce confidence judgments impacts the extent to which communicators' confidence elicits internal attributions about their competence versus external attributions about the decision context (Kahneman & Tversky, 1982; Løhre & Teigen, 2016; Ülkümen, Fox, & Malle, 2016). Relative to statements that refer externally to an event's likelihood ("there is a ___% chance of [event] occurring"), people associate self-referent confidence statements like the one used in Study 1 ("there is a ___% chance that I accurately determined whether...") with communicators' degree of knowledge or skill. I reasoned that the

¹² In order to determine advisors' payoffs, I later recruited a separate sample of participants to play the advice recipient role.

framing of advisors' confidence judgments would influence their perception that others will evaluate their competence; a pretest confirmed that the manipulation succeeded in shaping the extent to which participants perceived that others would evaluate their competence (see SOM, Study S2A for details about the pretest). Participants randomly assigned to the evaluative frame condition sent advice-seekers a message indicating how confident they were in the accuracy of their judgment using a frame similar to the one in Study 1 ("There is a ___% chance that I can accurately tell you whether the stock's value will be above or below \$XX.XX at the start of Month 13"). In contrast, those in the non-evaluative frame condition sent a message indicating the likelihood of an outcome that happened to match their prediction ("There is a ___% chance that the stock's value will be [above / below] \$XX.XX at the start of Month 13").¹³

After playing the advisor role, all participants proceeded to play the private decision-maker role; they were assured that others would not see their confidence or predictions for these judgments. Participants then proceeded to make predictions and confidence judgments for the same stocks they evaluated as advisors. Participants' confidence judgments were then elicited using a frame consistent with the one they saw as advisors. Although participants knew they did not have an audience evaluating their competence while playing the private decision-maker role (and thus should not be concerned with others evaluating their competence), I kept the confidence frame consistent across roles.

Instead of being rewarded based on their ability to be hired by an advice seeker (as they were in the advisor role), participants received "accuracy points" determined by a modified Brier score (Brier, 1950); Brier scores incentivize the honest reporting of confidence by rewarding accurate guesses expressed with high confidence and inaccurate guesses expressed with

¹³ The portion containing [above / below] appeared as a blank line in participants' message so that they knew their message would not reveal their actual prediction unless the advice recipient elected to pay a small fee.

tempered confidence. The more accuracy points participants earned, the greater their chances of winning a separate raffle for a \$50 bonus payment. To help participants understand the scoring, they saw a table demonstrating the number of points they would earn at various confidence levels, conditional on their prediction accuracy. To reinforce their understanding of their incentives, they read the following: "you earn more raffle tickets for a \$50 bonus payment when you honestly report your true confidence in the accuracy of your predictions."

Results

As in Study 1, participants exhibited a robust pattern of overconfidence. Compared to their mean confidence of 76%, their predictions were accurate 54% of the time, $t(361) = 19.11$, $p < .001$, $d = 1.00$, 95% CI = [0.88, 1.13]; this pattern held across conditions and roles (all $ps < .001$).

Because participants made their confidence judgments in the absence of feedback about their accuracy and advisors' hiring decisions (and could therefore not be impacted by these factors), I collapsed participants' overconfidence across all five rounds of predictions. A 2 X 2 mixed ANOVA ($\eta^2 = .04$) revealed two main effects on participants' overconfidence (see Figure 1 for conditional means and standard errors). First, participants displayed more overconfidence as advisors than the overconfidence reflected in their judgments as private decision-makers, $F(1, 360) = 21.94$, $p < .001$, $\eta_p^2 = .06$, $1-\beta > 99\%$. Second, they were more overconfident in the evaluative confidence frame condition than in the non-evaluative frame condition, $F(1, 360) = 10.33$, $p = .001$, $\eta_p^2 = .03$, $1-\beta = 89\%$.

More importantly, a Role X Confidence Frame interaction emerged, $F(1, 360) = 8.15$, $p = .005$, $\eta_p^2 = .02$, $1-\beta = 81\%$. Consistent with the strategic self-promotion account, the confidence frame manipulation impacted participants' overconfidence when they played the advisor role,

$t(360) = 4.37, p < .001, d = 0.46, 95\% \text{ CI} = [0.25, 0.67], 1-\beta = 99\%$. Relative to the non-evaluative frame condition, participants in the evaluative frame condition were more overconfident when making judgments as advisors. In contrast, the confidence frame did not impact participants' overconfidence when they were later incentivized to honestly report their confidence as private decision-makers, $t(360) = 1.38, p = .17, d = 0.14, 95\% \text{ CI} = [-0.06, 0.35]$. This resulted in a pattern where participants' overconfidence as advisors exceeded their overconfidence as private decision-makers in the evaluative confidence frame condition, $t(180) = 4.67, p < .001, d = 0.35, 95\% \text{ CI} = [0.20, 0.50], 1-\beta > 99\%$. However, despite making the same predictions using the same confidence frame, their overconfidence as advisors and private decision-makers did not differ in the non-evaluative frame condition, $t(180) = 1.55, p = .12, d = 0.12, 95\% \text{ CI} = [-0.03, 0.26]$. Thus, although participants were overconfident on the whole, they inflated their overconfidence as advisors relative to their private beliefs—but only to the extent that they strongly perceived others would evaluate their competence. Further, because participants' greater overconfidence in the evaluative confidence frame condition did not persist in the private judgments that immediately followed their broadcasted advice, participants did not engage in self-deception.

Discussion

Study 2 provides more direct evidence for the strategic self-promotion account and sheds further light on why advisors' overconfidence is greater than that of private decision-makers. Because advisors' overconfidence increased as it became more apparent that advice seekers would evaluate their competence, this indicates that their overconfidence increased to the extent they perceived that high confidence could allow them to project competence. Critically, this effect only held for the judgments they broadcasted to advice-seekers and did not impact their

judgments as private decision-makers. This finding suggests that advisors' overconfidence was deliberate and not a byproduct of delusional beliefs resulting from self-deception. Notably, this occurred despite the presence of conditions that are conducive to finding evidence of self-deception. People are particularly motivated to engage in self-deception about their abilities after engaging in ethically questionable behavior that makes them appear competent to others (Chance et al., 2011). Because participants in the evaluative confidence frame condition made judgments as private decision-makers *after* inflating their confidence as advisors, they should have been motivated to respond in a manner consistent with those estimates as private decision-makers. However, they actively lowered their confidence as private decision-makers, suggesting that they were aware of their deceptive behavior.

Although I pretested the confidence frame manipulation and found that it impacted advisors' overconfidence in the main study, I did not directly assess whether advisors consciously considered how their confidence would impact advice seekers' perception of their competence. If advisors' elevated overconfidence in the evaluative confidence frame condition was indeed strategic, then they should have perceived that their expressed confidence would impact their others' perception of their competence. I tested this in a replication of the advisor role portion of the study. In addition to replicating the effect of confidence frame on advisors' overconfidence, the study also confirmed that advisors more strongly associated their confidence with advice seekers' perception of their competence in the evaluative frame condition than in the non-evaluative frame condition (see SOM, Study S2B).

The Impact of Self-Promotion Incentives on Advisor Overconfidence

Studies 1 and 2 support the strategic self-promotion account. This account relies on the premise that advisors strategically inflate their confidence due to a (sometimes mistaken) belief

that overconfidence is in their self-interest. However, an alternative explanation for the findings in Studies 1 and 2 is that advisors' strategic overconfidence was instead driven by a sincere desire to help advice seekers.

The hypothesis that people consciously convey overconfidence to provide judgments others find informative was first advanced by Yaniv and Foster (1995, 1997) to explain the prevalence of overconfidence. Although the authors did not empirically test this assertion by comparing the overconfidence of individuals expressing confidence in front of an audience to those making judgments in private, it remains a popular account for overconfidence's robustness (Moore et al., 2015). This account holds that people deliberately display overconfidence because others perceive confident judgments to be informative. When they believe others will evaluate their competence, advisors may be particularly likely to realize that they can tailor their confidence to increase their advice's perceived informativeness. Given that communicators sometimes engage in strategic deception that helps others even when they have no incentive to do so (Erat & Gneezy, 2012), it could be the case that a desire to be helpful could inspire advisors to purposely inflate their overconfidence even when it does not serve their self-interest.

Because Studies 1 and 2 simulated the real-world incentives of advisors whose livelihood depends on their ability to attract clients, it is unclear whether these incentives triggered strategic self-promotion or whether advisors would have kept inflating their confidence in their absence. While scholars have decried the widespread incentives in advice markets that reward overconfidence (Kahneman, 2011; Radzevick & Moore 2011; Tetlock, 2005), there has yet to be an empirical test of what causal role these incentives play in driving advisors' confidence displays. If advisors' strategic overconfidence is motivated by self-promotion, they should curb their bent toward inflating their confidence (compared to their judgments as private decision-

makers) in the absence of incentives to engage in self-promotion. However, if their overconfidence is instead motivated by a desire to be helpful, their inflated confidence should persist to a similar degree in the absence of self-promotion incentives.

Strategic Self-Promotion Account Hypothesis 3: Compared to private decision-makers, advisors should more strongly increase their overconfidence in the presence of incentives to engage in self-promotion than in the absence of self-promotion incentives.

Studies 3-5: Testing the Effect of Self-Promotion Incentives

Studies 3-5 test the effect of self-promotion incentives on advisors' overconfidence. In so doing, they attempt to differentiate the strategic self-promotion account from an alternative account where advisors' inflated overconfidence reflects motives to be helpful. Study 3 tests the effect of a hiring incentive on advisors' overconfidence. To expand on Study 3, Study 4 tests the effect of a more general self-promotion incentive on a different form of overconfidence (excessively narrow confidence interval estimates). Finally, Study 5 attempts to find evidence of strategic self-promotion by comparing the self-reported overconfidence of advice professionals in situations where self-promotion serves their self-interest to those where it does not.

Study 3: The Effect of Hiring Incentives

Study 3 attempts to replicate the finding in Study 2 that advisors inflate their overconfidence when it is highly salient that others will evaluate their competence while exploring whether the presence of a hiring incentive moderates this pattern. The strategic self-promotion account predicts that advisors' tendency to inflate their confidence (relative to their beliefs as private decision-makers) should reduce in the absence of an incentive to be hired.

Method

Participants. I posted 216 assignments to MTurk and stopped data collection once they were all submitted, resulting in 220 MTurk workers completing the experiment ($M_{\text{Age}} = 33.09$ years, $SD = 11.01$, 39% female).

Procedure. In the interest of providing conditions that elicit strategic overconfidence from advisors, participants followed the same procedure as those in the evaluative confidence frame condition of Study 2. Participants first played the advisor role before playing the private decision-maker role. However, while playing the advisor role, they were randomly assigned to either receive a financial incentive to be "hired" by their advisor or not. Those in the hiring incentive condition received the same incentive as advisors in Studies 1 and 2. Participants assigned to the no incentive condition were also entered into a raffle for a \$50 bonus payment but told that the number of times their partner hires them has no bearing on their chances of winning. After playing the advisor role, participants played the private decision-maker role. They received the same incentive as in Study 2 to honestly report their confidence. Regardless of which incentive condition they were assigned to as advisors, private decision-makers did not have a hiring incentive.

Results

As in Studies 1 and 2, participants exhibited a robust pattern of overconfidence. Compared to their mean confidence of 76%, their predictions were accurate 59% of the time, $t(219) = 10.67$, $p < .001$, $d = 0.72$, 95% CI = [0.57, 0.87]; this pattern held across conditions and roles (all $ps < .001$).

A 2 (role: advisor, private decision-maker) X 2 (advisor financial incentive: yes, no) mixed ANOVA ($\eta^2 = .06$) revealed two main effects on participants' overconfidence (see Figure 2 for conditional means and standard errors). First, participants displayed more overconfidence

as advisors than the overconfidence reflected in their judgments as private decision-makers, $F(1, 218) = 9.38, p = .002, \eta_p^2 = .04, 1-\beta = 86\%$. Second, they were more overconfident in the advisor financial incentive condition than in the no financial incentive condition, $F(1, 218) = 9.07, p = .003, \eta_p^2 = .04, 1-\beta = 85\%$.

More importantly, a Role X Advisor Financial Incentive interaction emerged, $F(1, 218) = 27.61, p < .001, \eta_p^2 = .11, 1-\beta > 99\%$. Consistent with participants' behavior being motivated by self-interest, the financial incentive to be hired impacted participants' overconfidence when they played the advisor role, $t(218) = 5.17, p < .001, d = 0.70, 95\% \text{ CI} = [0.42, 1.00], 1-\beta > 99\%$. Relative to the no advisor financial incentive condition, participants in the advisor hiring incentive condition were more overconfident as advisors. However, exposure to a financial incentive while playing the advisor role did not influence participants' subsequent overconfidence when they were later incentivized to honestly report their confidence as private decision-makers, $t(218) = 0.48, p = .63, d = 0.07, 95\% \text{ CI} = [-0.20, 0.33]$. Replicating Study 2, participants' overconfidence as advisors exceeded their overconfidence as private decision-makers in the advisor financial incentive condition, $t(110) = 5.18, p < .001, d = 0.49, 95\% \text{ CI} = [0.29, 0.69], 1-\beta > 99\%$. However, their overconfidence as advisors was slightly *lower* than their overconfidence as private decision-makers in the no advisor financial incentive condition, $t(108) = 1.85, p = .067, d = 0.18, 95\% \text{ CI} = [-0.01, 0.37]$. This suggests that advisors' overconfidence was not motivated by a desire to be helpful. Instead, consistent with the strategic self-promotion account, advisors only inflated their overconfidence above the overconfidence they displayed as private decision-makers when appearing competent served their self-interest.

Discussion

Study 3 provides further evidence consistent with the strategic self-promotion account in a context where advisors strongly perceived that others were evaluating their competence. As in Study 2, advisors deliberately inflated their confidence in the presence of a financial incentive to have an advice seeker elect to receive their advice. In the absence of the incentive, however, they no longer inflated their confidence. Further, because exposure to the incentive as an advisor did not impact participants' judgment as private decision-makers, this suggests that it did not trigger self-deception. In addition to corroborating the findings of Study 2, this provides additional evidence that advisors consciously engaged in self-promotion without biasing their judgment.

More importantly, Study 3 provides direct evidence that advisors' strategic overconfidence was motivated by a desire to self-promote in situations where it serves their self-interest. Although other scholars have suggested that one mechanism for the prevalence of overconfidence is that people deliberately communicate it in order to provide information that others find helpful and informative (Yaniv & Foster, 1995, 1997), this hypothesis has not been empirically tested by comparing the judgments of private-decision makers to communicators broadcasting judgments to an audience (Moore et al., 2015). Study 3 finds evidence that this hypothesis cannot account entirely for advisors' strategic overconfidence. Because the removal of a financial incentive to self-promote curbed (and even slightly reversed) advisors' strategic overconfidence, advisors' overconfidence was, in part, motivated by self-interest. Thus, their overconfidence was not strictly motivated by a desire to be helpful.

It is noteworthy that participants in Studies 2 and 3 always played the private decision-maker role after playing the advisor role and were incentivized to calibrate their confidence accurately. Although these aspects of the study design were critical to test for evidence of self-deception adequately, it raises the possibility that idiosyncrasies of these design features could

have driven advisors' strategic overconfidence in the presence of an evaluative confidence frame and a hiring incentive. To rule out this possibility, I collected data from a sample of undergraduate business students assigned to the advisor financial incentive condition in Study 3 (see SOM, Study S3). However, I modified the condition so that participants' estimates as private decision-makers were not tied to accuracy incentives and always elicited before playing the advisor role. As in Studies 2 and 3, advisors' overconfidence exceeded their overconfidence as private decision-makers. This finding suggests advisors' strategic overconfidence in situations where they strongly perceive that others will evaluate their competence and have an incentive to self-promote is unlikely to be accounted for by the ordering of judgments or the presence of an accuracy incentive. Study 4 addresses this issue further by counterbalancing the order of judgments and not providing participants with an accuracy incentive.

Study 4: The Effect of Self-Promotion Incentives on Confidence Interval Estimates

Study 4 aims to generalize beyond situations where an advice seeker must pay for advice. Aside from advising clients, advisors might also advise their employer and colleagues in situations where they have relevant domain expertise for solving a problem at hand. While they may not earn a commission for providing sound advice, advisors' ability to project competence in these situations can impact their chances of achieving economically valuable outcomes like promotions and performance bonuses. Therefore, Study 4 attempts to extend beyond the idiosyncrasies of advisor-client interactions by conceptually replicating many features of advice-giving contexts. However, rather than testing the effect of an incentive to be hired by a client, the study tests the strategic self-promotion account by assessing whether a more general incentive to be perceived as competent by an audience causes communicators to strategically inflate their confidence relative to their judgments as private-decision makers.

Additionally, Study 4 examines whether advisors' strategic overconfidence persists with an alternative signal of their confidence: the margin of error surrounding their judgment. The most robust way of eliciting overconfidence is to ask people to provide a confidence interval around their numerical estimate of some quantity such that the confidence interval contains the quantity's actual value a prespecified percentage of the time (Alpert & Raiffa, 1982; Klayman et al., 1999); the narrower an individual's confidence interval, the more confident the individual is that their estimate is close to the truth. To rule out potential confounds introduced by the ordering in which participants played each role and attaching incentives for good calibration to the private decision-maker role (i.e., Studies 2 and 3), Study 4 counterbalances the order in which participants play the communicator and private decision-maker roles and does not provide private decision-makers with an incentive to report their confidence honestly.

Method

Participants. I posted 360 assignments to MTurk and stopped data collection once they were all submitted, which resulted in a total of 360 MTurk workers completing the experiment ($M_{\text{Age}} = 36.92$ years, $SD = 11.42$, 51% female).

Procedure. Study 4 followed a 2 (role: communicator, private decision-maker) X 2 (self-promotion incentive: yes, no) mixed design with self-promotion incentive manipulated between subjects and judgment publicity manipulated within subjects. Participants completed a weight-guessing task adapted from prior research, where they viewed five different full-body photos of different individuals and provided a 90% confidence interval for each pictured individual's weight (Gino & Moore, 2007; Moore & Klein, 2008). To do this, they gave both a "lower bound" and "upper bound" (in their choice of pounds or kilograms) such that there was only a 5% probability that the pictured individual could weigh less (lower bound) or more (upper

bound) than that number (see Moore, Carter, & Yang, 2015). Participants played the role of both communicator and private decision-maker in a randomized, counterbalanced order.

Communicator role. Before playing the communicator role, participants learned they had a 10% chance of having a subset of their judgments displayed to a future evaluator (described as another MTurk worker) who would be judging their expertise at weight guessing. Evaluators were described as only having access to participants' confidence interval estimates and a face-only version of each photo when evaluating their expertise (in contrast to participants' full-body photos). This aspect of the procedure maintained the appearance of an expertise asymmetry between potential evaluators and participants that mimics the expertise asymmetries inherent in most advisor-client interactions (Radzevick & Moore, 2011; Sah et al., 2013),

Participants were randomly assigned to self-promotion incentive or no incentive conditions. In the self-promotion incentive condition, participants learned that their judgments as communicators would impact their final payment. Specifically, they were told that, if they were selected to have their public judgments evaluated, they would earn "expertise points" such that the more expertise others perceived them to possess at weight guessing, the higher their expected bonus payment (maximum: \$10). To reinforce the manipulation, participants read: "your expected bonus payment is the highest when others think you have expertise at weight guessing on the basis of your lower bound and upper bound estimates."

Participants read the same language in the no incentive condition clarifying that their public confidence interval estimates would potentially be evaluated by someone else to assess their expertise. However, they also learned that this assessment would have no bearing on their final payment: "While these upcoming lower bound and upper bound estimates may be shown to somebody else, they will have no bearing on your final payment."

Private decision-maker role. Before playing the private decision-maker role, participants learned that their confidence interval estimates would not be shown to anybody else and would have no bearing on their final payment. They then made judgments about the same set of photos as they did in the communicator role. Unlike Studies 2 and 3, participants did not have an incentive to make accurate judgments.

Measuring overconfidence. As in prior studies, I measured overconfidence by regressing each participant's confidence on their accuracy and retaining the standardized residuals. Because high confidence corresponds to smaller (i.e., narrower) confidence intervals, I multiplied the resulting residuals by negative one so that higher values correspond to more overconfidence. Following the lead of Yaniv and Foster (1995, 1997), I computed a measure of accuracy unconfounded with confidence interval width by taking the midpoint of participants' confidence intervals and recording the absolute difference between this value and the actual weight of each pictured individual.¹⁴

Results

Participants exhibited a robust pattern of overconfidence, as their 90% confidence intervals only contained the correct weight 61% of the time, $t(359) = 20.85$, $p < .001$, $d = 1.10$, 95% CI = [0.26, 0.32]. This pattern held across conditions and roles (all $ps < .001$).

¹⁴ One can also assess accuracy as a binary measure indicating whether a confidence interval contained the correct weight, but this measure is confounded with confidence interval width. For example, consider two individuals who have identical beliefs about an individual's most likely weight. If Guesser A provides a wider confidence interval around that point estimate than Guesser B, Guesser A will necessarily have a greater likelihood of providing a confidence interval that contains the correct estimate. However, this apparent accuracy improvement is not a reflection of Guesser A holding a more accurate belief about the pictured individual's weight, but rather a reflection of Guesser A's wider confidence interval. Measuring accuracy as a difference between the midpoint of the confidence interval (a proxy of one's perception of an individual's most likely weight) avoids this dependency on confidence interval width. That said, the general pattern of results identified in this study holds in analyses of an alternative measure of overconfidence computed from the standardized residuals derived from regressing confidence on a binary measure of accuracy (see SOM).

I conducted a 2 X 2 mixed ANOVA ($\eta^2 = .02$) on participants' overconfidence (see Table 4 for conditional means and standard deviations). A main effect emerged where the self-promotion incentive increased participants' overconfidence, $F(1, 358) = 7.99, p = .005, \eta_p^2 = .02, 1-\beta = 80\%$. I did not find evidence of a role main effect, $F(1, 358) = 1.43, p = .23, \eta_p^2 < .01$.

More importantly, a Role X Self-Promotion Incentive interaction emerged, $F(1, 358) = 5.97, p = .015, \eta_p^2 = .01, 1-\beta = 68\%$. Consistent with advisors in Study 3, the self-promotion incentive impacted participants' overconfidence when they played the communicator role, $t(358) = 3.72, p < .001, d = 0.39, 95\% \text{ CI} = [0.18, 0.60], 1-\beta = 96\%$. Relative to the no incentive condition, participants in the self-promotion incentive condition were more overconfident when making judgments as communicators. Although the self-promotion incentive elicited a similar effect on participants' overconfidence as private decision-makers, the effect was smaller in magnitude, $t(358) = 1.77, p = .078, d = 0.19, 95\% \text{ CI} = [-0.02, 0.39]$. This difference in the extent to which the self-promotion incentive increased participants' overconfidence as communicators versus their overconfidence as private decision-makers resulted in a pattern where communicators' overconfidence exceeded private decision-makers' overconfidence in the self-promotion incentive condition, $t(181) = 2.03, p = .044, d = 0.15, 95\% \text{ CI} = [0.004, 0.30], 1-\beta = 52\%$. However, their overconfidence as communicators did not differ from their overconfidence as private decision-makers in the no incentive condition, $t(177) = 1.46, p = .15, d = 0.11, 95\% \text{ CI} = [-0.04, 0.26]$. This supports the strategic self-promotion account.

Discussion

By demonstrating that more general self-promotion incentives cause communicators to convey more overconfidence than reflected by their private beliefs, the results of Study 4 suggest that the findings in prior studies extend beyond advisor-client interactions. Although

communicators were not playing an advisor role per se, they were broadcasting their confidence to an audience with inferior domain expertise attempting to evaluate their competence. Because an incentive to project expertise elicited strategic overconfidence, this suggests that expert advisors are likely motivated to display overconfidence in a variety of situations where appearing competent serves their self-interest.

Study 4 also extends on the prior studies by demonstrating that the strategic self-promotion account extends to advisors' judgments about the margin of error surrounding their judgment. Although Yaniv and Foster (1995, 1997) have suggested that one reason for the excessive narrowness of individuals' confidence intervals is their desire to provide estimates that others find informative, they only examined private decision-makers' confidence intervals. This study provided an opportunity to directly compare the width of participants' confidence intervals as communicators to their intervals as private decision-makers. Communicating with a potential audience alone did not cause participants to express narrower confidence intervals. Instead, consistent with Study 3 and the strategic self-promotion account, their publicly broadcasted confidence intervals were only narrower than their private judgments when self-promotion served their self-interest.

Study 5: Strategic Overconfidence Among Advice Professionals

Collectively, Studies 1-4 exhibit a robust pattern where advisors (and communicators more generally) engage in a pattern of strategic self-promotion by deliberately expressing overconfidence when they believe self-promoting their expertise will serve their self-interest. The studies attempted to simulate dynamics in real-world advice contexts, but it is unclear whether their findings generalize to advice professionals' behavior. Study 5 tests whether advice

professionals' self-reports of attempting to deliberately project overconfidence increases in situations where projecting expertise serves their self-interest.

Method

Participants. Based on a preregistered data collection rule, 100 advice professionals completed the study ($M_{\text{Age}} = 41.34$ years, $SD = 11.65$, 60% female, $M_{\text{Industry Experience}} = 13.37$ years, $SD = 9.17$). I recruited participants through a Prime Panel managed by CloudResearch.¹⁵ Before being invited to complete the study, I prescreened potential participants based on their employment in occupations that are typically associated with advising others (accounting, consulting, finance, human resources, IT, legal, marketing, sales) and their passage of a customized screener verifying that advising others is a primary function of their job. After following a preregistered procedure to eliminate participants who provided incomplete or nonsensical responses, I included 94 participants in the final analyses; the results hold in analyses of all 100 participants.

Procedure. Participants wrote about two situations where somebody solicited their advice at work (randomized and counterbalanced order) and assured the confidentiality of their written responses. In one of the situations, they had an incentive to engage in self-promotion. In the other, they did not. Participants read the following prompt [self-promotion incentive condition / no self-promotion incentive condition]:

Think of a time at work when somebody solicited your advice about something you were uncertain about and you felt that your ability to convince the person about your expertise was [critical to / not going to impact your chances of] retaining a client's business,

¹⁵ Prime Panels is a compilation of online research panels that provides access to a larger, more diverse pool of participants than Amazon Mechanical Turk.

keeping your job, earning a commission, earning a performance bonus, or earning a promotion.

Participants described each situation and then proceeded to indicate whether, when interacting with others in the situation described, they (a) conveyed more confidence than they knew was justifiable at the time (deliberate overconfidence), (b) an appropriate level of confidence (appropriate confidence), or (c) less confidence than they knew was justifiable at the time (deliberate tempered confidence).

Participants' responses were subjected to the Bayesian Truth Serum algorithm (Prelec, 2004), which incentivizes honest self-reports by comparing participants' responses to the mean of their estimates for how prevalent each response is among sampled participants. This scoring system rewards “surprisingly common” responses and penalizes “surprisingly uncommon” responses. The system was described to participants as follows: “the important property of the formula is that it rewards truthful answers”; participants were then provided with a link to a paper describing the algorithm. Consistent with prior research (John, Loewenstein, & Prelec, 2012), I donated to participants' charity of choice and informed them that the size of the donation made on their behalf would increase with their truthfulness. Research has demonstrated that these procedures compare favorably to other honesty-inducing techniques at increasing truth-telling in situations involving admissions of socially undesirable behavior (John et al., 2012; Weaver & Prelec, 2013).

Results

Following a preregistered data analysis plan, I examined participants' confidence using a multilevel ordered logistic regression model that nested responses within participants using a participant-specific random intercept (Pseudo $R^2 = .05$, $SD_{\text{Random Intercept}} < 0.01$, $ICC < .01$).

Consistent with the prior studies, participants conveyed more confidence in the presence of a self-promotion incentive, $B = 0.83$ ($SE = 0.30$), $z = 2.75$, $p = .006$, $OR = 2.30$, 95% CI = [1.27, 4.17], $1-\beta = 79\%$.

As documented in Table 5, a follow-up multilevel logistic regression on participants' deliberate overconfidence (1 = deliberate overconfidence, 0 = appropriate confidence or deliberate tempered confidence) revealed a main effect of the self-promotion incentive, $z = 3.26$, $p = .001$, $OR = 3.24$, 95% CI = [1.60, 6.56], $1-\beta = 90\%$. Relative to the no incentive condition (16%), participants were more likely to indicate that they attempted to convey more confidence than they knew was justifiable in situations involving a self-promotion incentive (37%). In contrast, a separate follow-up logistic regression on participants' deliberate tempered confidence (1 = deliberate tempered confidence, 0 = appropriate confidence or deliberate overconfidence) did not find evidence of the self-promotion incentive impacting participants' likelihood of self-reporting that they attempted to convey less confidence than they knew was justifiable, $z = 0.44$, $p = .66$, $OR = 0.83$, 95% CI = [0.35, 1.95]. Thus, the self-promotion incentive increased participants' confidence by raising their likelihood of deliberately attempting to project overconfidence but not by reducing their likelihood of deliberately attempting to project tempered confidence.

Discussion

Study 5 finds support for the strategic self-promotion account in the self-reports of advice professionals. Consistent with the prior studies, advisors were more likely to admit that they attempted to strategically display overconfidence in situations where self-promoting their expertise to others served their self-interest than in situations where it did not.

General Discussion

The current research tests whether cogently learning and adapting to advice-seekers' preferences can account for overconfident advice or whether it is instead a byproduct of a robustly applied self-promotion tactic. Scholars have suggested that advice seekers' inability to punish overconfidence adequately is partly to blame for the prevalence of overconfident advice (Radzevick & Moore, 2011; Ronay et al., 2019; Tetlock, 2005). However, the current research results suggest that merely punishing overconfidence is unlikely to temper advisors' strategic overconfidence. Even when paired with an advice seeker who penalized confidence on balance, advisors were more overconfident than private decision-makers; this inflated overconfidence persisted in the face of feedback about advice seekers' decisions. Instead, advisors' overconfidence appeared to be exacerbated by a pattern of strategic self-promotion. They deliberately inflated their confidence to the extent that it was (a) highly salient others would evaluate their expertise, and (b) projecting competence served their self-interest.

Importantly, overconfidence persisted even in the absence of motives to engage in strategic self-promotion. Private decision-makers not broadcasting their judgments also exhibited overconfidence, irrespective of whether they had an incentive to calibrate their confidence. This pattern is consistent with previous findings that overconfidence persists even in the presence of incentives for accurate calibration (Williams & Gilovich, 2008; Lebreton et al., 2018) and the acquisition of expertise that should enhance calibration (Liersch, & Yaniv, 2008; Moore et al., 2017). Because overconfidence is such a pervasive cognitive bias, strategic self-promotion is merely a factor that can contribute to it in advice contexts; it by no means offers a complete explanation for advisors' overconfidence.

However, because prior overconfidence research has almost exclusively focused on either the perspective of private decision-makers or advisors communicating publicly, without differentiating between advisors' public communications and private judgments (Meikle et al., 2016; Moore et al., 2015), the current research offers new insights into how the strategic social considerations inherent in many advice contexts might exacerbate overconfidence. In so doing, it informs several theoretical accounts for the social motivations underlying overconfidence. Further, in identifying the causal impact of several key features of advice contexts on advisor overconfidence, the current research offers several practical implications for managers and other types of organizational decision-makers hoping to reduce their exposure to overconfident advice.

Theoretical Implications

This research makes several theoretical contributions to the literature on overconfidence and advice-giving. First, it offers one explanation for why advice-giving contexts might exacerbate overconfidence. In contrast to the private decision contexts typically studied by overconfidence researchers (e.g., Haran et al., 2010; Juslin et al., 2007; Klayman et al., 1999; Moore et al., 2015), advice contexts present advisors with strategic social considerations that alter their recommendations (Kray, 2000; Kray & Gonzalez, 1999) and information search (Jonas & Frey, 2003; Jonas et al., 2005). The current research documents yet another manner in which advice contexts shape individuals' judgment: They can motivate them to become more overconfident when giving advice. A fundamental aspect of many real-world advice contexts is that they present advisors with motives to project competence to an audience (Radzevick & Moore, 2011; Tetlock, 2005). As demonstrated by the current research, these motives can result in advisors strategically employing overconfidence as a means of self-promoting their expertise. Although some have suggested that advisors' overconfidence emerges as a byproduct of advice-

seekers' failure to punish overconfidence (Radzevick & Moore, 2011; Ronay et al., 2019; Tetlock, 2005), the current findings indicate that this account might oversimplify the reasons for the prevalence of overconfident advice. While accountability can reduce private decision-makers' overconfidence (Tetlock & Kim, 1987), merely punishing overconfidence does not appear to override advisors' tendency to display it when motivated to self-promote.

Another contribution of this research is in documenting precisely *how* self-promotion motivates advisor overconfidence. Although social motives to appear competent can sometimes delude individuals into becoming genuinely overconfident (Anderson et al., 2012; Schwarzmann & Van der Weele, 2019), advisors exhibited *strategic* overconfidence by inflating their publicly communicated confidence above the overconfidence reflected by their private beliefs. In combination with advice professionals' self-reports, this indicates that they were deliberately overconfident while engaging in self-promotion. These findings add to prior studies documenting that advisors' communicated beliefs sometimes diverge from those they hold privately (Kray 2000; Kray & Gonzalez, 1999; Jonas & Frey, 2003; Jonas et al., 2005).

Third, the current research points to advisors' self-interest as a driver of their strategic overconfidence. One popular explanation for the prevalence of overconfidence is that people are deliberately overconfident to be informative to others (Yaniv & Foster, 1995, 1997). However, because scholars have largely neglected to consider how communicators' confidence varies in private versus in front of an audience, the current research offers a rare direct test of this hypothesis. While advisors did inflate their overconfidence relative to their private beliefs, this only occurred when projecting competence to an audience was in their self-interest. Thus, self-promotion, rather than a genuine desire to be informative, motivated advisor overconfidence.

Finally, this research offers more general insights into when social motives can exacerbate overconfidence. Scholars often assume that people deliberately exploit the benefits of confidence when communicating to an audience (e.g., Meikle et al., 2016; Van Zant & Moore, 2013), but direct supporting evidence has proven to be surprisingly elusive (Moore et al., 2015). Because previous failed attempts at detecting this effect are mostly unpublished and anecdotal at present, ascertaining precisely what methodological features might differentiate the current studies from prior attempts is a challenge. However, the current research suggests that the effect is particularly likely to emerge in contexts where communicators perceive that others will evaluate their competence and self-promotion serves their self-interest.

Practical Implications

The current research also offers insights for managers and decision-makers who want to reduce their exposure to overconfident advice. One insight is that attempting to curb advisors' overconfidence solely by punishing them for their overconfidence is unlikely to be sufficient in reducing one's exposure to overconfident advice. Strategic overconfidence is such a robustly applied tactic that advisors continued to rely on it even in the presence of feedback that their overconfidence was undermining their economic success. This finding is similar to prior studies documenting that people often fail to learn from real-world feedback about how their behavior impacts others' decisions (Ball et al., 1991; Zlatev et al., 2017).

Because advisors appear to use confidence as an instrument to achieve self-presentational goals aligned with their self-interest, it might be tempting to conclude that advice seekers would be best avoiding all information about advisors' confidence. However, this may not necessarily be advisable. One reason is that advisor confidence tends to correlate with the advisor's judgmental accuracy and objective level of expertise (McKenzie et al., 2008; Moore et al., 2017;

Sniezek & Van Swol, 2001). Thus, although advisors might disingenuously exaggerate their confidence when it serves their self-interest, their confidence is still likely to provide some diagnostic information about the quality of their recommendations that can help advice seekers. Further, the notion that advice seekers can avoid exposure to information about advisors' confidence altogether is probably wishful thinking. In addition to expressing their confidence through written communication, people also convey confidence through their nonverbal behavior (Tenney, Meikle, Hunsaker, Moore, & Anderson, 2019; Van Zant & Berger, 2020). Whether they like it or not, advice seekers are likely to be exposed to subtle verbal and nonverbal indicators of an advisor's confidence during interactions with the advisor.

The current research does, however, offer some promising solutions for curbing advisors' likelihood of strategically displaying overconfidence in the first place. One is that advice seekers hoping to receive honest confidence estimates should avoid advisors primarily compensated based on their ability to recruit clients, take on new accounts, or solicit hourly consultation fees. Instead, they might opt for advisors who are salaried or earn bonuses based on client satisfaction. Likewise, organizational leaders might promote an organizational culture where employees do not constantly perceive their prospects in the organization as hinging on their response to a superior's question.

However, incentives to self-promote are often unavoidable. What can managers and advice seekers do in these situations? It appears to be critical that they avoid projecting the impression that the intent behind any questions they ask is to scrutinize advisors' competence. A simple way to avoid this impression is by framing questions about advisors' confidence in a manner that refers to the likelihood of events instead of advisors' judgmental accuracy. Whereas advisors are relatively unlikely to perceive that their competence is under scrutiny when others

ask questions framed around event likelihoods (e.g., "How likely is [outcome] to occur?"), it becomes highly salient that others are evaluating their competence when they ask questions framed around their judgmental accuracy (e.g., "How confident are you that your prediction about [outcome] is accurate?").

Limitations and Future Directions

The current research also raises several important theoretical and methodological issues. First, it does not truly address how people come to associate overconfidence with strategic self-promotion in the first place. Although advisors in the current studies did not show evidence of adapting to advice seekers' preferences, a lifetime of repeated exposure to audiences who reward overconfidence might have already ingrained the idea that people reward overconfidence in many situations (Johnson & Fowler, 2011). Communicators might struggle to adapt their behavior in situations where overconfidence backfires because they are so strongly socialized to display it as a self-promotion tactic. Another possibility could be that communicators' strategic overconfidence emerges as a byproduct of perspective-taking. If advisors tend to take advice seekers' perspective, they should recognize that they would personally reward a confident advisor and act on this inference by inflating their own confidence. Research that considers how people develop lay intuitions about others' likely response to overconfidence could shed light on why people display it when self-promoting.

Another question unanswered by the current research concerns whether advisors continue to rely on overconfidence as a strategic self-promotion tactic in the presence of opportunities to provide a rationale for their confidence to an audience. While many advice contexts do not provide advisors with the opportunity to provide detailed justifications for their confidence until they are "hired" by a client (Radzevick & Moore, 2011), it could be possible that advisors might

not display inauthentic confidence when provided an opportunity to rationalize their uncertainty. I considered this possibility in one study but continued to find evidence that self-promotion incentives cause advisors to inflate their overconfidence in a context where they have an opportunity to provide their audience with written explanations for their confidence (see Study S4 of the SOM). However, it could be possible in some circumstances that advisors might combine tempered confidence with a compelling rationale to provide a unique and contrarian perspective.

Finally, it is worth exploring the conditions under which punishing overconfidence might cause advisors to curb their confidence. The current research fails to find evidence of advisors adapting to feedback that others punish overconfidence, but experimental research paradigms that systematically vary the degree to which others penalize overconfidence under conditions optimized for learning might more precisely identify how strongly advice-seekers would need to penalize overconfidence before advisors learn to temper their confidence.

Relatedly, if advisors were to receive feedback more explicitly linking their overconfidence to advice-seekers' decisions, they should be more likely to recognize when advice seekers are punishing overconfidence. In applied advice contexts, this would require that advice recipients directly confront advisors with the rationale behind their decisions to reject advice. However, advice-seekers often experience discomfort with the prospect of directly confronting a distrusted advisor (Sah, Loewenstein, & Cain, 2013). Instead, they often opt to reject suspect advice discreetly. Thus, it could be possible that many advice-seekers would be reluctant to directly confront advisors about their overconfidence. Instead, they might opt to reject their advisor under different pretenses that allow the advisor to save face (e.g., "I cannot afford to pay you right now"). A lack of direct feedback would pose challenges to advisors'

ability to learn that advice-seekers are punishing their overconfidence. Future research on how social feedback can be optimized to improve advisors' ability to recognize when others are punishing overconfidence should consider the content of the feedback, along with interventions that enhance advice seekers' likelihood of choosing to provide direct feedback to advisors.

Conclusion

Whether playing the role of a consumer seeking professional advice about an important life decision, an employee seeking a coworker's opinion, or even a manager soliciting employee input while devising a course of action, we all stand to benefit from well-calibrated advisors. While it may be tempting to think holding people accountable by punishing their overconfidence should force them to rethink their instincts to self-promote through displays of confidence, recent events call this conclusion into question. Donald Trump's persistence in confidently making many dubious claims despite facing intense media scrutiny, losing a presidential reelection bid, and repeatedly having his mistaken claims exposed by non-partisan websites like PolitiFact is one example that comes to mind. The current findings are similarly pessimistic about whether punishing overconfidence is effective at changing advisors' behavior. With social media providing more avenues for advice professionals to self-promote their expertise than ever before, punishing their overconfident claims are unlikely to override their tendency to be strategically overconfident. Instead, advice-seekers would be better off putting advisors' concerns about being evaluated at ease and more generally rethinking how they choose to reward advisors.

References

- Alpert, M., & Raiffa, H. (1982). A progress report on the training of probability assessors. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and Biases* (pp. 294-305). Cambridge: Cambridge University Press.
- Anderson, C., Brion, S., Moore, D.A., & Kennedy, J.A. (2012). A status-enhancement account of overconfidence. *Journal of Personality and Social Psychology, 103*(4), 718-735.
- Ball, S.B., Bazerman, M.H., & Carroll, J.S. (1991). An evaluation of learning in the bilateral winner's curse. *Organizational Behavior and Human Decision Processes, 48*(1), 1-22.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica, 52*(4), 1007-1028.
- Baumeister, R.F., & Jones, E.E. (1978). When self-presentation is constrained by the target's knowledge: Consistency and compensation. *Journal of Personality and Social Psychology, 36*(6), 608-618.
- Bliese, P. D., & Wang, M. (in press). Results provide information about cumulative probabilities of finding significance: Let's report this information. *Journal of Management*.
- Brier, G.W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review, 78*(1), 1-3.
- 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.
- Bozeman, D. P., & Kacmar, K. M. (1997). A cybernetic model of impression management processes in organizations. *Organizational Behavior and Human Decision Processes, 69*(1), 9-30.
- Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica, 45*(1), 223-241.

- Chance, Z., Norton, M.I., Gino, F., & Ariely, D. (2011). Temporal view of the costs and benefits of self-deception. *Proceedings of the National Academy of Sciences*, *108*(Supplement 3), 15655-15659.
- Daniels, D.P., & Zlatev, J.J. (2019). Choice architects reveal a bias toward positivity and certainty. *Organizational Behavior and Human Decision Processes*, *151*(1), 132-149.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, *47*(2), 315-372.
- Edwards, J.R. (1994). Regression analysis as an alternative to difference scores. *Journal of Management*, *20*(3), 683-689.
- Erat, S., & Gneezy, U. (2012). White lies. *Management Science*, *58*(4), 723-733.
- Frey, D. (1978). Reactions to success and failure in public and in private conditions. *Journal of Experimental Social Psychology*, *14*(2), 172-179.
- Gino, F., & Moore, D.A. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making*, *20*(1), 21-35.
- Gino, F., & Schweitzer, M. E. (2008). Blinded by anger or feeling the love: How emotions influence advice taking. *Journal of Applied Psychology*, *93*(5), 1165-1173.
- Godfrey, D.K., Jones, E.E., & Lord, C.G. (1986). Self-promotion is not ingratiating. *Journal of Personality and Social Psychology*, *50*(1), 106-115.
- Goodman-Delahunty, J., Granhag, P.A., Hartwig, M., & Loftus, E.F. (2010). Insightful or wishful: Lawyers' ability to predict case outcomes. *Psychology, Public Policy, and Law*, *16*(2), 133-157.
- Goffman, E. (1959). *The presentation of self in everyday life*. Garden City, NY: Doubleday/Anchor Books.

Hackethal, A., Haliassos, M., & Jappelli, T. (2012). Financial advisors: A case of babysitters?

Journal of Banking & Finance 36(2):509-524.

Haran, U., Moore, D.A., & Morewedge, C.K. (2010). A simple remedy for overprecision in

judgment. *Judgment and Decision Making* 5(7):467-476.

Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and

sharing responsibility. *Organizational Behavior and Human Decision Processes*, 70(2),

117-133.

Hertz, U., Palminteri, S., Brunetti, S., Olesen, C., Frith, C.D., & Bahrami, B. (2017). Neural

computations underpinning the strategic management of influence in advice giving.

Nature Communications, 8(1), 2191.

Hewitt P.L., Flett, G.L., Sherry, S.B., Habke, M., Parkin, M., Lam, R.W., ... Stein, M.B. (2003).

The interpersonal expression of perfection: Perfectionistic self-presentation and

psychological distress. *Journal of Personality and Social Psychology*, 84(6), 1303-1325.

John, L.K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable

research practices with incentives for truth telling. *Psychological Science*, 23(5), 524-

532.

Johnson, D.D.P., & Fowler, J.H. (2011). The evolution of overconfidence. *Nature*, 477(7364),

317-320.

Johnson, D.D.P., McDermott, R., Barrett, E.S., Cowden, J., Wrangham, R., McIntyre, M.H., &

Rosen, S.P. (2006). Overconfidence in wargames: Experimental evidence on

expectations, aggression, gender and testosterone. *Proceedings of the Royal Society B-*

Biological Sciences, 273(1600), 2513-2520.

- Jonas, E., & Frey, D. (2003). Information search and presentation in advisor–client interactions. *Organizational Behavior and Human Decision Processes*, 91(2), 154-168.
- 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.
- Jones, E.E., & Pittman, T.S. (1982). Toward a general theory of strategic self-presentation. In J. Suls (Ed.), *Psychological Perspectives on the Self*. (pp. 231-262). Hillsdale, NJ: Erlbaum.
- Juslin, P., Winman, A., Hansson, P. (2007). The *naïve* intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, 114(3), 678-703.
- Kahneman, D. (2011, October). Don't blink! The hazards of confidence. *New York Times*.
- Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition*, 11(2), 143-157.
- Kennedy, J.A., Anderson, C., & Moore, D.A. (2013). When overconfidence is revealed to others: Testing the status-enhancement theory of overconfidence. *Organizational Behavior and Human Decision Processes*, 122(2), 266-279.
- Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2), 211-228.
- Klayman, J., Soll, J.B., González-Vallejo, C., & Barlas, S. (1999). Overconfidence: It depends on how, what, and whom you ask. *Organizational Behavior and Human Decision Processes*, 79(3), 216-247.
- Köszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature*, 52(4), 1075-1118.
- Kray, L.J. (2000). Contingent weighting in self-other decision making. *Organizational Behavior and Human Decision Processes*, 83(1), 82-106.

Kray, L.J., & Gonzalez, R. (1999). Differential weighting in choice versus advice: I'll do this, you do that. *Journal of Behavioral Decision Making*, 12(3), 207-218.

Leary, M.R., & Kowalski, R.M. (1990). Impression management: A literature review and two-component model. *Psychological Bulletin*, 107(1), 34-47.

Lebreton, M., Langdon, S., Slieker, M. J., Nooitgedacht, J. S., Goudriaan, A. E., Denys, D., . . .

Luigjes, J. (2018). Two sides of the same coin: Monetary incentives concurrently improve and bias confidence judgments. *Science Advances*, 4(5), eaaq0668.

Løhre, E., & Teigen, K.H. (2016). There is a 60% probability, but I am 70% certain:

Communicative consequences of external and internal expressions of uncertainty.

Thinking and Reasoning, 22(4), 369-396.

McKenzie, C.R.M., Liersch, M.J., & Yaniv, I. (2008). Overconfidence in interval estimates:

What does expertise buy you? *Organizational Behavior and Human Decision Processes*, 107(2), 179-191.

Meikle, N.L., Tenney, E.R., & Moore, D.A. (2016). Overconfidence at work: Does

overconfidence survive the checks and balances of organizational life? *Research in Organizational Behavior*, 36(1), 121-134.

Minson, J. A., & Umphres, C. (in press). Confidence in context: Perceived accuracy of

quantitative estimates decreases with repeated trials. *Psychological Science*.

Moore, D.A., Carter, A.B., & Yang, H.H.J. (2015). Wide of the mark: Evidence on the

underlying causes of overprecision in judgment. *Organizational Behavior and Human Decision Processes*, 131(1), 110-120.

Moore, D.A., & Healy, P.J. (2008). The trouble with overconfidence. *Psychological Review*,

115(2), 502-517.

- Moore, D.A., & Klein, W.M.P. (2008). Use of absolute and comparative performance feedback in absolute and comparative judgments and decisions. *Organizational Behavior and Human Decision Processes*, 107(1), 60-74.
- Moore, D. A., & Schatz, D. (in press). Overprecision increases subsequent surprise. *PLoS One*.
- Moore, D.A., Swift, S.A., Minster, A., Mellers, B., Ungar, L., Tetlock P., ... Tenney, E.R. (2017). Confidence calibration in a multiyear geopolitical forecasting competition. *Management Science*, 63(11), 3552-3565.
- Moore, D.A., Tenney, E.R., & Haran, U. (2015). Overprecision in judgment. In G. Keren & G. Wu (Eds.), *Wiley-Blackwell Handbook of Judgment and Decision Making* (pp. 182-209). West Sussex, UK: John Wiley & Sons, Ltd.
- Nezlek, J.B., & Leary, M.R. (2002). Individual differences in self-presentational motives in daily social interaction. *Personality and Social Psychology Bulletin*, 28(2), 211-223.
- Olsson, H. (2014). Measuring overconfidence: Methodological problems and statistical artifacts. *Journal of Business Research*, 67(8), 1766-1770.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4), 1029-1050.
- Plous, S. (1995). A comparison of strategies for reducing interval overconfidence in group judgments. *Journal of Applied Psychology*, 80(4), 443-454.
- Prelec, D. (2004). A Bayesian Truth Serum for subjective data. *Science*, 306(5695), 462-466.
- Price, P.C., & Stone, E.R. (2004). Intuitive evaluation of likelihood judgment producers: Evidence for a confidence heuristic. *Journal of Behavioral Decision Making*, 17(1), 39-57.

- 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.
- Richard, F. D., Bond, C. F., & Stokes-Zoota, J. J. (2003). One hundred years of social psychology quantitatively described. *Review of General Psychology*, *7*(4), 331-363.
- Ronay, R., Oostrom, J.K., Lehmann-Willenbrock, N., Mayoral, S., & Rusch, H. (2019). Playing the trump card: Why we select overconfident leaders and why it matters. *Leadership Quarterly*, *30*(6), 10136.
- Sah, S., & Loewenstein, G. (2014). Nothing to declare: Mandatory and voluntary disclosure leads advisors to avoid conflicts of interest. *Psychological Science*, *25*(2), 575-584.
- Sah, S., Loewenstein, G., & Cain, D. M. (2013). The burden of disclosure: Increased compliance with distrusted advice. *Journal of Personality and Social Psychology*, *104*(2), 289-304.
- Sah, S., Moore, D.A., & MacCoun, R.J. (2013). Cheap talk and credibility: The consequences of confidence and accuracy on advisor credibility and persuasiveness. *Organizational Behavior and Human Decision Processes*, *121*(2), 246-255.
- Schmitt, D. P., & Buss, D. M. (1996). Strategic self-promotion and competitor derogation: Sex and context effects on the perceived effectiveness of mate attraction tactics. *Journal of Personality and Social Psychology*, *70*(6), 1185-1204.
- Schwardmann, P., & van der Weele, J. (2019). Deception and self-deception. *Nature Human Behavior*, *3*(10), 1055-1061.
- Sniezek, J.A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, *62*(2), 159-174.
- 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.

- Tenney, E.R., MacCoun, R.J., Spellman, B.A., & Hastie, R. (2007) Calibration trumps confidence as a basis for witness credibility. *Psychological Science, 18*(1), 46-50.
- Tenney, E.R., Meikle, N.L., Hunsaker, D., Moore, D.A., & Anderson, C. (2019). Is overconfidence a social liability? The effect of verbal versus nonverbal expressions of confidence. *Journal of Personality and Social Psychology, 116*(3), 396-415.
- Tenney, E.R., Spellman, B.A., MacCoun, R.J. (2008). The benefits of knowing what you know (and what you don't): How calibration affects credibility. *Journal of Experimental Social Psychology, 44*(5), 1368-1375.
- Tetlock, P.E. (2005). *Expert political judgment: How good is it? How can we know?*. Princeton, NJ: Princeton University Press.
- Tetlock, P.E., & Kim, J.I. (1987). Accountability and judgment processes in a personality prediction task. *Journal of Personality and Social Psychology, 52*(4), 700-709.
- Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business, 59*(4), 251-278.
- Ülkümen, G., Fox, C.R., & Malle, B. (2016). Two dimensions of subjective uncertainty: Clues from natural language. *Journal of Experimental Psychology: General, 145*(10), 1280-1297.
- Van Zant, A.B., & Berger, J. (2020). How the voice persuades. *Journal of Personality and Social Psychology, 118*(4), 661-682.
- Van Zant, A. B., & Moore, D. A. (2013). Avoiding the pitfalls of overconfidence while benefitting from the advantages of confidence. *California Management Review, 55*(2), 5-23.

- Von Hippel, W., & Trivers, R. (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences*, 34(1), 1-56.
- Vulllioud, C., Clément, F., Scott-Phillips, T., & Mercier, H. (2017). Confidence as an expression of commitment: Why misplaced expressions of confidence backfire. *Evolution and Human Behavior*, 38(1), 9-17.
- Weaver, R., & Prelec, D. (2013). Creating truth-telling incentives with the Bayesian Truth Serum. *Journal of Marketing Research*, 50(3), 289-302.
- Williams, E. F., & Gilovich, T. (2008). Do people really believe they are above average? *Journal of Experimental Social Psychology*, 44(4), 1121-1128.
- Wistrich, A.J., & Rachlinski, J.J. (2012). How lawyers' intuitions prolong litigation. *Southern California Law Review*, 86(1), 571-636.
- Yaniv, I., & Foster, D.P. (1995). Graininess of judgment under uncertainty: An accuracy-informativeness trade-off. *Journal of Experimental Psychology: General*, 124(4), 424-432.
- Yaniv I., & Foster, D.P. (1997). Precision and accuracy of judgmental estimation. *Journal of Behavioral Decision Making*, 10(1), 21-32.
- Zhang, Z., & Peterson, S. J. (2011). Advice networks in teams: The role of transformational leadership and members' core self-evaluations. *Journal of Applied Psychology*, 96(5), 1004-1017.
- Zlatev, J.J., Daniels, D.P., Kim, H., & Neale, M.A. (2017). Default neglect in attempts at social influence. *Proceedings of the National Academy of Sciences*, 114(52), 13643-13648.

Zlatev, J.J., Daniels, D.P., Kim, H., & Neale, M.A. (2018). Reply to Jung et al.: Default neglect persists over time and across contexts. *Proceedings of the National Academy of Sciences*, *115*(35), E8107-E8108.

Table 1

Study 1: Means, Standard Deviations, and Correlations Between Variables

Variable	<i>M</i>	<i>SD</i>	1	2	3	4
1. Advisor Role	0.50	0.50				
2. Round	5.50	2.87	.00			
3. Prior Round Confidence	80.56	14.32	.12***	.04 [†]		
4. Prior Round Accuracy	0.55	0.50	.02	-.12***	.02	
5. Overconfidence	0.00	1.00	.12***	.04 [†]	.55***	.03

Note. Advisor Role = 0 for private decision-makers, 1 for advisors; Prior Accuracy = 0 for inaccuracy in prior round, 1 for accuracy in prior round.

[†] $p < .10$. *** $p < .001$.

Table 2

Study 1: Regressions Predicting Overconfidence

Variable	Model 1		Model 2		Model 3		Model 4	
	<i>B</i> (<i>SE</i>) [<i>B</i> _{95% CI}]	<i>z</i>	<i>B</i> (<i>SE</i>) [<i>B</i> _{95% CI}]	<i>z</i>	<i>B</i> (<i>SE</i>) [<i>B</i> _{95% CI}]	<i>z</i>	<i>B</i> (<i>SE</i>) [<i>B</i> _{95% CI}]	<i>z</i>
Advisor Role (0 = no, 1 = yes)	0.23 (0.10) [0.03, 0.44]	2.30*	0.23 (0.10) [0.04, 0.42]	2.33*	0.29 (0.10) [0.09, 0.49]	2.79**	0.39 (0.13) [0.13, 0.65]	2.98**
Round	0.01 (0.01) [0.002, 0.02]	2.40*	0.02 (0.01) [0.005, 0.03]	2.69**	0.02 (0.01) [0.004, 0.03]	2.63**	0.02 (0.01) [0.01, 0.04]	2.76**
Prior Round Confidence			0.01 (0.002) [0.002, 0.01]	3.44***	0.01 (0.002) [0.002, 0.01]	3.47***	0.01 (0.002) [0.002, 0.01]	3.41***
Prior Round Accuracy (0 = inaccurate, 1 = accurate)			0.05 (0.03) [-0.01, 0.12]	1.55	0.11 (0.05) [0.02, 0.20]	2.31*	0.11 (0.05) [0.02, 0.21]	2.39*
Advisor Role X Prior Round Accuracy					-0.11 (0.07) [-0.25, 0.02]	1.72†	-0.13 (0.06) [-0.26, 0.01]	1.87†
Advisor Role X Round							-0.02 (0.01) [-0.04, 0.01]	1.26
Intercept	-0.19 (0.08) [-0.35, -0.04]	2.50*	-0.68 (0.15) [-0.98, -0.40]	4.69***	-0.72 (0.15) [-1.01, -0.43]	4.87***	-0.76 (0.15) [-1.06, -0.47]	5.01***
Pseudo <i>R</i> ²	.16		.25		.25		.25	
<i>SD</i> _{Random Intercept}	0.71		0.67		0.67		0.67	
<i>SD</i> _{Residuals}	0.70		0.70		0.70		0.70	
<i>ICC</i>	.51		.48		.48		.48	

Note. Parameter estimates are from multilevel regression models predicting advisor overconfidence in a given round.

† $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$.

Table 3

Study 1: The Effect of Role on Overconfidence at Increasing Levels of Feedback Strength

Feedback Strength Threshold	N_{Advisors}	$M_{\text{Advisor Overconfidence}} (SD)$	$t(df), p\text{-value}$	$d, 95\% CI$
0.0	78	0.16 (0.95)	$t(183) = 2.31, p = .021$	0.34, [0.05, 0.64]
0.1	71	0.18 (0.95)	$t(176) = 2.40, p = .017$	0.37, [0.07, 0.67]
0.2	63	0.25 (0.91)	$t(168) = 2.77, p = .006$	0.44, [0.13, 0.75]
0.3	50	0.36 (0.89)	$t(155) = 3.23, p = .002$	0.55, [0.21, 0.89]
0.4	31	0.28 (0.93)	$t(136) = 2.28, p = .024$	0.47, [0.06, 0.87]
0.5	16	0.09 (1.01)	$t(121) = 2.28, p = .024$	0.27, [-0.26, 0.80]

Note. All values represent the results of t -tests comparing a subsample of advisors' Round 10 overconfidence to all private decision-makers' Round 10 overconfidence ($M = -0.19, SD = 1.04$). The Feedback Strength Threshold column corresponds to the minimum value of feedback strength that an advisor had to exceed entering Round 10 to be included in analyses. N_{Advisors} = the number of advisors who exceeded a given Feedback Strength Threshold.

Table 4

Study 4: Overconfidence by role and self-promotion incentive

	<u>Role</u>	
	Communicator	Private Decision-Maker
Self-Promotion Incentive	0.17 ^a (0.64)	0.07 ^b (0.87)
No Self-Promotion Incentive	-0.14 ^c (0.95)	-0.10 ^{bc} (0.97)

Note. Numbers represent conditional means (standard deviations in parentheses). Cells with different superscripts differ at $p < .05$.

Table 5

Study 5: Regressions Predicting Advisors' Deliberate Overconfidence and Tempered Confidence

Variable	Deliberate Overconfidence		Deliberate Tempered Confidence	
	<i>B</i>	<i>SE</i>	<i>B</i>	<i>SE</i>
Self-Promotion Incentive	1.17**	0.36	-0.19	0.44
Intercept	-1.74***	0.29	-1.83***	0.30
Pseudo R^2		.05		<.01
$SD_{\text{Random Intercept}}$		0.01		<0.01
<i>ICC</i>		<.01		<.01

Parameter estimates are from multilevel logistic regression models predicting deliberate overconfidence and deliberate tempered confidence. Self-Promotion Incentive = 0 for no self-promotion incentive, 1 for self-promotion incentive.

** $p < .01$. *** $p < .001$.

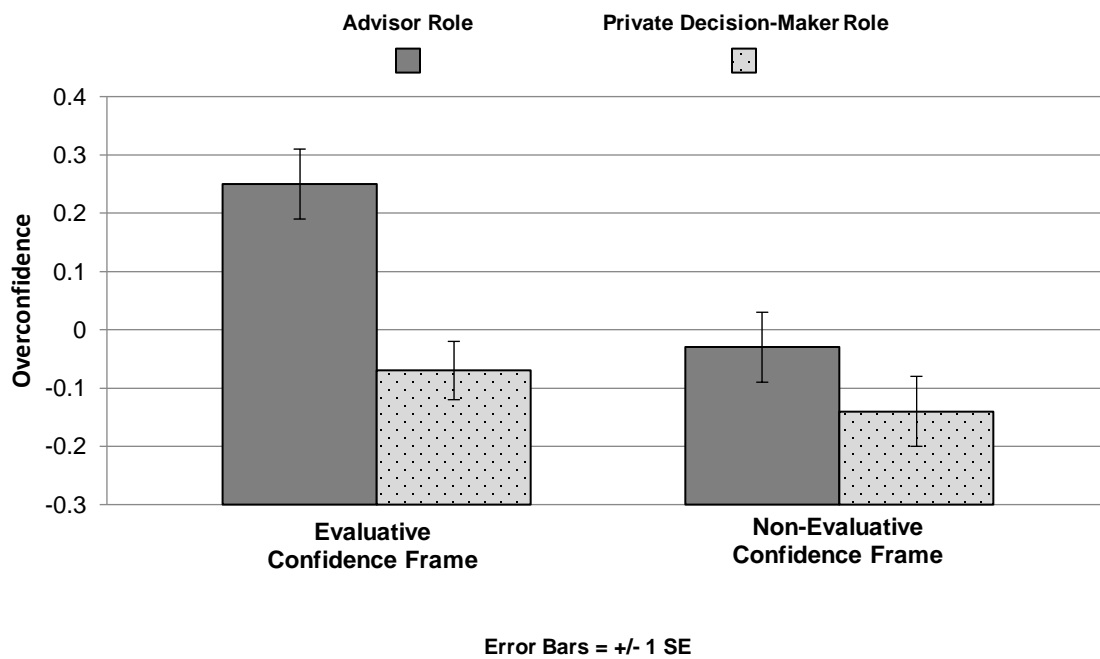
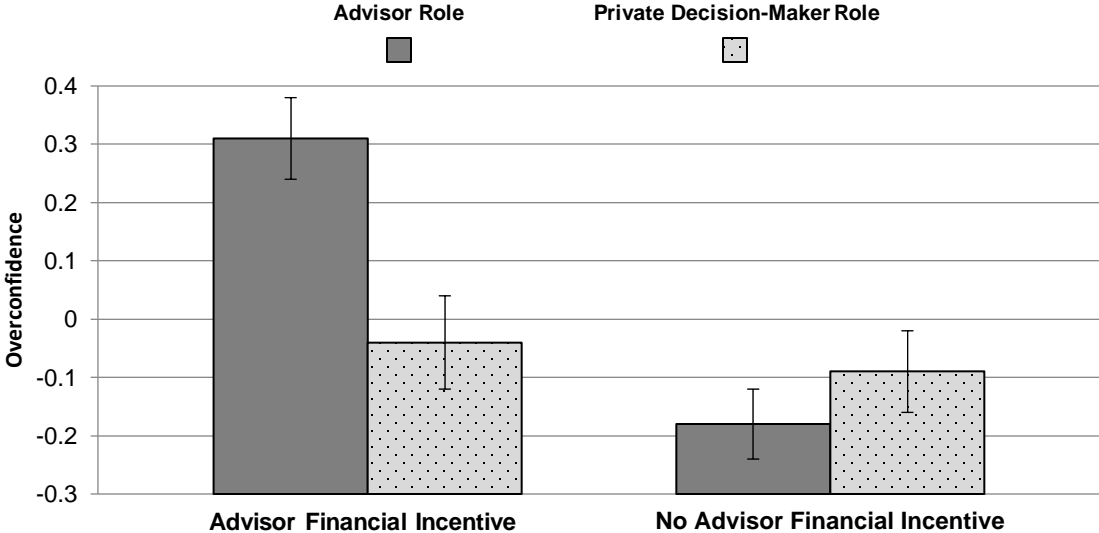


Figure 1. Study 2: Overconfidence by role and confidence frame.



Error Bars = +/- 1 SE

Figure 2. Study 3: Overconfidence by role and advisor financial incentive.