

Flattering Advice

Amanda Chen and David Hagmann

Department of Management, The Hong Kong University of Science and Technology

Author Note

Amanda Chen  <https://orcid.org/0000-0002-4428-627X>

David Hagmann  <https://orcid.org/0000-0002-2080-997X>

Correspondence concerning this article should be addressed to Amanda Chen, Email:
zchengj@connect.ust.hk

Abstract

Advice improves decisions when it transmits an unbiased, informed assessment, but the truth can be painful and fall short of the recipient's expectations. Across three preregistered and incentivized experiments involving real advisor-advisee interactions ($N = 3,982$), we show that advisors avoid delivering such disappointment at the expense of accuracy and their own earnings — even though their material incentives are aligned with the recipient's. Advisors provide inflated, flattering advice to advisees and do so to a greater extent when incentivized to be viewed as likable. This strategy works. Advisors who provide flattering advice are rated as more likable and no less trustworthy, even as their advice is less accurate (Study 1). When advisors observe an advisee's quiz performance along with the advisee's expressed expectations, knowledge of the advisee's expectations shifts the advice they give, although it carries no information the advisor does not already have (Study 2). Finally, because men expect to perform better than equally performing women, advisors who observe expectations (but not gender) more often tell men to enter a tougher but potentially more rewarding competition. The advice is followed, and it makes men worse off (Study 3). Advisors' efforts to avoid disappointment thus degrade the quality of advice and can generate systematic gender disparities in advice and outcomes even in the absence of any information about gender or intent to discriminate.

Keywords: advice, belief-based utility, interpersonal relationships, gender, discrimination

Flattering Advice

Introduction

Many of the most consequential decisions people make — whether to apply for a promotion, take on a high-visibility project, or enter a competitive market — are made with the help of advice from those who are better informed (Bonaccio & Dalal, 2006; Harvey & Fischer, 1997; Soll & Larrick, 2009). Advice is thus a central channel through which information is transmitted in organizations and markets (Gibbons et al., 1998; Stigler, 1961). When the economic incentives of advisor and advisee are aligned, informed advisors should simply reveal what they know: any distortion lowers the quality of the recipient's decision and the advisor's own payoff. The dominant paradigm in the literature on advice has built on this premise, treating advice as a channel for transmitting accurate information from an informed sender to a less-informed receiver, and generating robust insights into how people weight advice, when they discount it, and what makes advisors credible (Haran & Shalvi, 2020; Rader et al., 2017; See et al., 2011; Soll & Larrick, 2009). Yet this framing misses something fundamental about what advice is. When an advisor, for example, tells someone whether to apply for a promotion, they communicate not just a recommended course of action but an implicit evaluation of the person asking — one that both parties recognize as tailored to the recipient. Advice, in other words, is also a social signal of how the advisor sees the advisee. In this paper, we show that the social component of advice systematically distorts what gets communicated, reducing advice quality overall and introducing systematic differences between the advice received by men and women even when advisors are blind to the advisee's gender.

At its best, advice offers an accurate, informed assessment of the recipient's situation, helping them make better decisions than they would have made alone (Landis et al., 2022; Soll & Larrick, 2009). In practice, however, advisors often hold a less flattering view of the recipient than the recipient holds of themselves, inflated by self-serving interpretation of information. This gap raises a question that the informational framing cannot answer: should the advisor deliver an accurate but potentially painful assessment, or soften the recommendation to protect the

recipient's feelings and shield their beliefs? We show that advisors routinely choose the latter. We refer to such advice as flattering advice because it communicates a more optimistic view of the recipient than the sender in fact holds.

One reason to soften is concern for the recipient's hedonic welfare. Disappointment arises when information reveals that a desired outcome is less likely than anticipated, forcing recipients to update their beliefs about their future prospects (Gul, 1991; Kőszegi, 2010). Someone advised against applying for a promotion not only learns that they are less competitive than they had believed but must also revise their expectations about their career trajectory. These psychological costs are substantial enough that people go to considerable lengths to avoid information that threatens their beliefs, even when it is freely available (Golman et al., 2017; Ho et al., 2021; Loewenstein & Molnar, 2018). Bell (1985) shows that disappointment aversion implies a preference for breaking bad news gradually to soften its impact; advisors can go further and elect not to deliver bad news at all, shading their recommendations upward to spare the recipient. Beyond avoiding disappointment, advisors may also inflate their advice as a form of emotional support, encouraging recipients who are uncertain or anxious about their prospects (Dalal & Bonaccio, 2010; Feng & Magen, 2016).

Advisors also have self-interested reasons to flatter. Unlike the judge-advisor paradigm, which deliberately strips away social context to isolate informational dynamics (Harvey & Fischer, 1997; Snizek & Buckley, 1995), advice in organizational and personal settings is embedded in relationships. Advisors are keenly aware of how their recommendations reflect on them, and may soften their message to appear likable, avoid a difficult conversation, or preserve relational harmony (Chentsova-Dutton & Vaughn, 2012; Gneezy et al., 2017; Levine et al., 2020; Schaerer et al., 2018; Van Swol, 2009). Flattering advice may also shield the advisor from blame: people tend to attribute poor outcomes to the advisor even when the link between advice and outcome is uncertain (Palmeira et al., 2015; Weiner, 1985), but they punish deliberate misguidance far more harshly than unintended error, and an advisor who appears supportive is rarely suspected of ill intent (Haran & Shalvi, 2020; John et al., 2019). Communicators across consequential domains

routinely prioritize recipients' emotional welfare over accurate information delivery, and recipients reward them for doing so (Levine et al., 2018). Recipients penalize those who deliver unwelcome assessments and reward those whose recommendations match or exceed the self-image they are motivated to protect (Shalvi et al., 2019), appreciating well-meaning inflation even when they recognize it as such (Zhang & Epley, 2012). In our first experiment, advisors who flatter are indeed evaluated as more likable and warmer — and no less trustworthy. So even an advisor with no direct concern for the recipient's feelings has an incentive to flatter.

Which of these motives dominates will depend on the relationship and the stakes: advisors might flatter strangers out of politeness, encourage friends as a form of emotional support, or seek to motivate recipients with particularly ambitious goals. In this paper, we do not adjudicate between these motives. Instead, we document advisors' use of flattering advice, the interpersonal consequences of doing so, and the disparities it produces across groups of advisees.

Across all these motives, avoiding disappointment requires knowing what the recipient expects: it is not a matter of how good the news is, but of how it compares to the recipient's reference point (Bell, 1985; Kőszegi & Rabin, 2006, 2009). In the context of advice, recipients often supply that reference point themselves, because advice seekers frequently include their own thinking alongside their requests (Reif et al., 2024). An entrepreneur asking a mentor whether to pitch a top-tier investor may lead with how quickly early users have signed up and a student asking whether to apply to a competitive doctoral program may recount the praise her thesis received. Such information is usually offered as context, but it tells the advisor how the recipient expects to be evaluated. Expressed expectations also make an attractive basis for flattering advice because they are optimistic without being implausible: people systematically believe they performed better than they did (Moore & Healy, 2008), so advice anchored on their expectations is encouraging yet not dishonest — the advisor can flatter without deliberately misleading. An advisor who wants to avoid disappointing, or to reap the interpersonal rewards of flattering, will therefore calibrate their recommendation to the recipient's expressed expectations rather than to their own assessment. Importantly, they will do so even when those expectations carry no outcome-relevant information

that the advisor does not already have: they nonetheless convey information about the *recipient's beliefs*.

Once advice is calibrated to expressed expectations, any systematic group difference in expressed confidence translates into a systematic group difference in advice. Research consistently documents that men receive more aspirational advice than women, who are more likely to receive cautious, relational, or generally encouraging guidance (Bear et al., 2017; Gallen & Wasserman, 2024; Kanze et al., 2018). These differences have typically been attributed to the expectations advisors hold about men and women, to beliefs that women's feelings should be protected from negative information (Jampol & Zayas, 2021; Moss-Racusin et al., 2012), or to unconscious bias (Banaji & Greenwald, 1995; Nosek et al., 2009) — though interventions designed to reduce such bias have been largely unsuccessful (Chang et al., 2019; Paluck & Green, 2009). All of these accounts require the advisor to observe or infer the recipient's gender. We propose a mechanism that is not inherently discriminatory but produces discrimination through group differences in expectations, driven by advisors' tendency to avoid disappointing recipients. This is similar to recent work on cognitive errors related to base rates as a source of discrimination against numerically smaller groups, which in many domains would lead to gender discrimination (Hagmann et al., 2026).¹

Men are on average more confident in their performance than equally performing women, expressing higher expectations even when their objective ability is identical (Exley & Nielsen, 2024; Lundeberg et al., 1994; Möbius et al., 2022). When seeking advice, men are more likely to highlight their strengths, recent successes, and ambitious goals; women with the same performance are more likely to express themselves modestly, setting a lower reference point for the advisor. Advisors who calibrate to these expressed expectations will therefore give men more ambitious advice than equally performing women, even when they have no information about gender and hold no beliefs about what men and women are capable of. The confidence gap

¹ Although beyond the scope of this paper, there may also be strategic considerations when there is a conflict of interest and advisors try to take advantage of (perceived) naivete (Bhattacharya et al., 2024).

becomes an advice gap through the intention to be kind rather than the intention to discriminate. Moreover, the dynamic may be self-reinforcing: more ambitious advice encourages men to pursue more challenging and better compensated opportunities (Coffman et al., 2024; Niederle & Vesterlund, 2007), which further bolsters their confidence, while more cautious advice steers women toward safer paths, deepening rather than correcting the original gap.

We present results from three preregistered and incentivized experiments involving real advisor-advisee interactions ($N = 3,982$). Each study consists of three asynchronous stages: a group of participants (advisees) performs an initial task, a second group (advisors) provides advice, and the original participants return to receive the advice and make an incentivized decision. Throughout, the advisor's earnings are tied to the quality of the advisee's decision, so distorting advice is costly to the advisor. In Study 1, we show that advisors use flattering advice as a tool to appear more likable, and that this strategy is effective: advisees rate flattering advisors as more likable and no less trustworthy, even as the advice they receive is less accurate. In Study 2, we use a randomly assigned default to nudge advisees to express high or low expectations and show that advisors respond to these expectations even though they carry no information: the advisor directly observes the performance that determines the advisee's payoff. Advisors who observe expressed expectations are more likely to tell those nudged to have high expectations to aim higher, even though both groups performed equally well. Study 3 then shows the consequence of this mechanism for group disparities. Because men express greater confidence than equally performing women in the context of our experiment, advisors who observe stated expectations recommend more ambitious options to men than to women with identical objective performance, even though advisors have no information about gender. Notably, this advice gap leads to worse recommendations for the very group it appears to favor: while most of those who do well in a competitive task are men, men are also disproportionately likely to be advised to enter competitions that they would lose.

This paper makes three contributions. First, we speak to the literature on advisors' motives in advice-giving. Prior work has assumed that advisors either act in the advisee's best interest and

provide accurate advice, or distort their recommendations because of a conflict of interest or a desire to impose their own views (Haran & Shalvi, 2020; Mackinger et al., 2017; Milyavsky & Gvili, 2024; Sah & Loewenstein, 2015; Van Swol, 2009). We show that distortion arises even when material incentives are fully aligned: advisors shade their recommendations out of concern for recipients' feelings at a cost to accuracy and to their own earnings. Second, we advance the literature on belief-based utility by showing that concern for belief utility operates not only on the demand side of information but on the supply side as well. Advisors internalize the psychological cost their advice imposes on recipients and tailor their recommendations to preserve the recipient's self-perceptions (Golman et al., 2017; Loewenstein & Molnar, 2018). Third, we identify a novel mechanism behind gender disparities in advice. Men are on average more confident than equally performing women (Exley & Nielsen, 2024; Lundeberg et al., 1994; Möbius et al., 2022), and because advisors calibrate their recommendations to implied expectations, this confidence gap becomes an advice gap — one that emerges not from discriminatory intent or beliefs about gender, but from the well-intentioned desire to flatter rather than disappoint. Because the mechanism operates without the advisor ever observing gender, it would survive interventions that target advisors' beliefs about groups.

Experiments

The three preregistered studies share a three-stage design in which participants are paired anonymously in two roles. Following the direction in which advice flows, we refer to the participant who gives advice as the *Sender* and the participant who receives it as the *Receiver*. (Receivers also report their own guesses and expectations, but only the Sender transmits advice.) In Stage 1, we measure a Receiver's performance and their expectations of how well they did. In Stage 2, we show the measured performance to a Sender and ask them to advise the Receiver on how to maximize their chances of earning a bonus; in Studies 2 and 3, half of the Senders additionally observe the Receiver's stated expectations. In Stage 3, Receivers return, view the advice they received, and make their final, incentivized decisions. This structure captures a core feature of organizational advice-giving, where managers or mentors often know both an

employee's track record and their confidence or aspirations, and could adjust their advice to avoid disappointing them.

Two features of the design rule out standard explanations for distorted advice. First, material incentives are aligned: the Sender's bonus depends on the decision the Receiver makes, so the advice that maximizes the Receiver's expected earnings also maximizes the Sender's. Under standard assumptions, a Sender should therefore simply reveal their best assessment. Second, when Senders observe the Receiver's stated expectations (Studies 2 and 3), those expectations carry no decision-relevant information: the Receiver's bonus is based solely on past performance, which the Sender directly and perfectly observes but the Receiver can only estimate. Even if two Receivers with the same score differed in latent ability — one having been lucky and the other unlucky on specific questions — that difference would be irrelevant, because the bonus turns on the realized score alone. Nor can expectation-consistent advice be justified as a motivational boost, since the performance being evaluated is past and fixed.

We nonetheless propose that Senders recognize the psychological toll of disappointing news, incurring a hedonic cost when they anticipate distressing the Receiver. This cost, we hypothesize, is sufficient to distort advice upward, away from accuracy and toward confirming the Receiver's prior beliefs. Stated expectations play a central role in this account not because they are informative but because they reveal the reference point against which the Receiver will evaluate the advice — and thus which recommendations would disappoint.

Our experiments are in other respects a conservative test of this hypothesis. Senders remain fully anonymous throughout: they have no ongoing relationship with the Receiver to protect, cannot be identified or held accountable for poor advice, and interact through a medium that attenuates the emotional richness of face-to-face communication. Interpersonal considerations would likely be stronger when advisor and advisee have an existing relationship and when advice is delivered face-to-face. Moreover, unlike in our experiments, advisors in the field usually do not suffer any direct costs when their flattering advice leads to a poor outcome for

the recipient, particularly when it is not clear what would have happened under a counterfactual.² If the desire to avoid disappointment shapes advice even in our setting, it is likely more pronounced in the organizational contexts we seek to understand.

The studies differ in what we manipulate and in the question each answers. Study 1 varies the Sender's incentive (a bonus tied either to the accuracy of the Receiver's decision or to how likable the Receiver rates them) to test whether Senders use flattering advice as an interpersonal strategy and whether that strategy pays off. Study 2 isolates the causal effect of expressed expectations on advice: a randomly assigned default nudges Receivers to express high or low expectations, so that expressed expectations vary while performance does not. Study 3 turns to naturally occurring expectations: Receivers estimate their score on a mathematics quiz, and men estimate higher scores than women with identical performance. Because Senders observe these estimates but never gender, any gender gap in advice can arise only through the expectations Receivers express. Together, the studies trace flattering advice from its interpersonal returns (Study 1) to the mechanism that produces it (Study 2) to its distributional consequences (Study 3).

Open Science Statement

We report all manipulations, measures, and data exclusions in our experiments, and none of the experiments involved deception. The preregistration reports, screenshots of all experimental materials, and the analysis code to replicate all statistical analyses and figures are available on the Open Science Framework.³

Study 1: Flattering Advice and Its Interpersonal Returns

We begin by examining whether Senders recognize that flattering advice carries hedonic value for Receivers and anticipate that this will reflect positively on them. Attractiveness is a useful domain for this test because it is ego-relevant enough that negative feedback can hurt, and prior work shows that people routinely avoid giving such feedback honestly face-to-face (Gneezy

² In some cases, however, repeated interaction might also increase incentives for honest feedback. Someone who is known to persistently give overoptimistic advice may be viewed as less trustworthy in the long run.

³ https://osf.io/8r3d4/?view_only=5ad7bafcd16b4d4ba08bb28b0e2bd02d

et al., 2017). That work documents feedback avoidance driven by self-protection from blame, an effect that is mitigated when feedback is anonymous; here we test whether flattery can also be wielded strategically to leave a positive impression. Specifically, we examine if Senders believe that flattering advice will improve how they are perceived, and if so, whether this perception is true. We do this by manipulating the incentives for Senders, who receive a bonus either based on how they are evaluated by a Receiver, or a bonus for the Receiver's accuracy.

The study takes place in three stages. First, Receivers upload selfies and are grouped with nine other Receivers of the same sex. We then recruit Senders of the opposite sex to rank the members of a group from most to least attractive and to provide advice. Specifically, Senders advise the Receiver they ranked as the 7th most attractive on what rank that Receiver should bet they were ranked by the larger group of raters. Senders were randomly assigned to two treatments, receiving a bonus payment either if the Receiver guessed their rank accurately or if the Receiver evaluated the Sender as likable. We hypothesize that Senders incentivized to be liked will recommend betting on a lower and thus more attractive rank.

Methods

Stage 1. We recruited 300 Receivers from Prolific and, after asking demographic questions, invited them to upload selfies to be rated on attractiveness by other participants. We obtained selfies from 100 men and 107 women who adhered to our instructions, which required that the photo show the Receiver's face and include no other person. In line with our preregistration, we selected the first 100 photos for each gender to arrive at a gender-balanced sample of 200 Receivers ($M_{\text{Age}} = 39.37$ years). Receivers were informed that their selfie would be randomly grouped with those of nine other Receivers of their sex and ranked on attractiveness by a group of new Prolific participants of the opposite sex. They then made an unincentivized guess of their rank in this group of ten.

Stage 2. Next, we recruited Senders from Prolific; after the exclusions described below, the final sample consists of 472 Senders ($M_{\text{Age}} = 41.03$ years; 49.79% Female). We first collected demographic information, then matched each Sender to a group of ten Receivers of the opposite

sex. Senders ranked the ten selfies from most to least attractive, so that a lower rank corresponds to greater attractiveness, by entering a rank into a textbox next to each picture. Because of a limitation with the survey software, we could not validate that each rank was given only once, and we removed 115 Senders who failed to follow instructions and did not provide a complete ranking.⁴

After submitting their rankings, Senders saw the photo of the Receiver they had ranked as the 7th most attractive. We reminded them of the rank they had just given to that Receiver and informed them that this Receiver would be invited back and could earn a \$1 bonus for correctly guessing their rank in the group. That rank was determined by the aggregate ratings of all Senders who ranked the group, and because Receivers never observed the other nine people in their group, they would depend on the Sender's recommendation along with their own assessment.

The key manipulation was the Sender's incentive for giving advice. We randomly assigned Senders to one of two schemes. In the "Accuracy" treatment, they received a bonus identical to the Receiver's: \$1 if the Receiver guessed their rank correctly. In the "Likability" treatment, we informed them that the Receiver would rate how likable they considered the Sender on a 5-point Likert scale, with each scale point translating into a bonus of 20 cents. Senders then selected the rank, from 1 to 10, that they recommended the Receiver bet on.

Stage 3. Finally, we invited back the Receivers who had been ranked as the 7th most attractive by at least one Sender and who had therefore received advice. Following our preregistration, we kept the follow-up survey open for seven days. In total, 144 Receivers ($M_{\text{Age}} = 40.1$ years, 75 men, 69 women) completed the follow-up survey. We reminded them of the selfie they uploaded in Stage 1 and informed them that a group of ten selfies, including theirs, had been ranked by other Prolific participants. They then saw the advice from a randomly selected Sender and made their own guess of their rank, with a \$1 incentive for accuracy. Finally, they rated the Sender's likability, warmth, friendliness, good-naturedness, trustworthiness, and sincerity on


⁴ Moreover, we could not collect data from Senders who did not rank any Receiver as the 7th most attractive, as that image would be shown on subsequent pages.

5-point Likert scales adapted from Fiske et al. (2007).

Results

Figure 1

Advice given in Study 1. Senders advise the Receiver they ranked 7th most attractive out of ten (dashed box). Under the accuracy incentive, the modal advice is the Sender's own ranking; under the likability incentive, advice shifts toward more attractive ranks. The inset summarizes advice relative to the Sender's own ranking of 7.



We begin by examining the prior beliefs of the Receivers who uploaded selfies. On average, men guessed that they ranked 5.76 in their group of ten and women guessed that they ranked 6.37 ($t(205) = 2.15$, $p = .033$), a first glimpse of the gender gap in self-expectations that Study 3 exploits. Notably, Receivers' self-perceptions correlated strongly with the aggregate ratings of the Senders (0.43 , $t(193) = 6.70$, $p < .001$). However, there was substantial heterogeneity in perceptions of attractiveness. Of the 200 Receivers, 149 were ranked as the 7th most attractive by at least one Sender. On average, men in this subset estimated that they ranked 5.94th and women estimated that they ranked 6.59th ($t(143) = 1.94$, $p = .054$).

Next, we turned our attention to the advice, shown in Figure 1. As predicted, Senders in

the Likability condition recommended a more attractive rank (5.38) than those in the Accuracy condition (6.19), indicating that Senders inferred flattering someone with pleasant advice would make themselves appear more likable ($t(470) = 5.44, p < .001$). Notably, Senders in the Accuracy condition also gave flattering advice, recommending a more attractive rank than the rank 7 they had themselves assigned ($t(234) = -8.82, p < .001$). Thus, even when incentivized for accuracy, Senders offered flattering advice.⁵ At the same time, Senders did not simply tell Receivers that they were the most attractive person in the group; they may have inferred that flattering advice needs to be somewhat realistic to be believable. We return to this in the General Discussion.

Indeed, the Likability incentive came at a cost to accuracy: the gap between the advised rank and the Receiver's actual final rank was significantly larger in the Likability condition than in the Accuracy condition ($t(458) = -3.97, p < .001$). That said, even under likability incentives, Senders still provided somewhat accurate advice on average, with the mean gap being only 1.

Table 1

Receivers who receive advice implying greater attractiveness, that is, a lower rank, perceive the Sender as more likable (Column 1) and warmer (Column 2). Column 3 shows that Senders are rated as more trustworthy when they advise lower ranks, but this relationship is only directional.

	(1)	(2)	(3)
Advised Rank	-0.103*	-0.128**	-0.050
	(0.046)	(0.044)	(0.043)
Constant	3.752***	3.885***	3.385***
	(0.278)	(0.267)	(0.258)
N	144	144	144

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001


Finally, we examined whether flattering advice led to more positive evaluations of Senders or was instead dismissed as insincere. Following our preregistration, we averaged the ratings on

⁵ This shading could be due to concerns of avoiding disappointment. However, it could also be that Senders are uncertain about the rankings they have given and make a recommendation that combines their own belief with a uniform prior. Therefore, our analyses focus on the difference between the two conditions.

likability, warmth, friendliness, and good-naturedness to create a scale of likability ($\alpha = 0.93$); and we created a scale of trustworthiness by averaging the ratings of trustworthiness and sincerity ($\alpha = 0.87$).

Figure 2

Evaluations of the Sender by advised rank in Study 1. Small points show individual ratings; the large white-rimmed points mark the mean rating at each advised rank. Receivers rate Senders as more likable and warmer when the advice implies a more attractive (lower) rank.




As seen in Figure 2, Senders who suggested that the Receiver was more attractive were indeed rated as more likable ($b = -0.10$, 95% CI $[-0.19, -0.01]$, $t(142) = -2.23$, $p = .027$) and warm ($b = -0.13$, 95% CI $[-0.22, -0.04]$, $t(142) = -2.91$, $p = .004$; Columns 1 and 2 of Table 1). Figure 3 traces the same pattern against the Receiver's own initial guess as the reference point: evaluations fall steeply when the advice implies a less attractive rank than the Receiver had guessed, whereas flattery beyond the guess earns little additional credit. Interestingly, these benefits did not come at the cost of sincerity: Senders who recommended a more favorable rank were viewed as no less trustworthy (Column 3 of Table 1).

In exploratory analyses, we also compared evaluations across the two incentive treatments,

Figure 3

Mean evaluations of the Sender in Study 1, with 95% confidence intervals, by how far the advised rank fell below or above the Receiver's own guess; positive values mean the advice implied the Receiver was more attractive than they had guessed, and the tails are pooled at three or more ranks. Curves are loess fits to the individual ratings. Evaluations drop once advice falls short of the Receiver's own guess, while flattery beyond the guess earns little extra credit.



a comparison we did not preregister and were not powered to detect. We found no difference in likability and warmth across the two treatments ($t(142) = 1.13$, $p = .262$, and $t(142) = 0.60$, $p = .550$, respectively). Similarly, among the advice shown to Receivers who returned, the discrepancy between the advised rank and the correct rank did not differ across the two treatments ($t(142) = -0.36$, $p = .723$).

Discussion

When advice communicates ego-relevant information (here, the Receiver's attractiveness), Senders tailor their recommendations to their incentives. Specifically, when they are rewarded for being likable, they recommend that the Receiver bet on a more favorable rank than when they are incentivized for accuracy. People understand that the advice they give can influence how others perceive them and use flattering advice as a tool to boost their likability. Importantly, Receivers

do not discount flattering advice and instead evaluate people who advise them to bet on a more attractive rank as warmer and more likable. These gains to interpersonal perceptions do not come at the cost of trustworthiness, even when the advice is inflated.

Study 2: Advice Responds to Uninformative Expectations

We next examine whether Senders take into account the expectations of Receivers when providing a recommendation. Receivers complete a quiz consisting of ten questions that draw on ego-relevant domains. They then express a non-binding preference to compete against a group of high performers or a group of low performers on the same task, based on their past performance. We nudge Receivers toward one of the two groups by selecting it as the default. Senders observe Receivers' score on the quiz and, in one treatment, also which group they had selected. We predicted that Receivers who were nudged toward the high-performer group would be more likely to receive advice to compete against that group when Senders observed their non-binding choice.

Methods

Competitor Recruitment. In a preliminary stage, we recruited 50 participants from Prolific ($M_{\text{Age}} = 35.54$ years; 64% Female) and gave them five minutes to complete a quiz consisting of ten items. The quiz included word puzzles, identifying emotions from photos, and selecting the best responses for hypothetical scenarios. These questions were adapted from surveys that measure problem-solving, emotional intelligence, and communication skills, which a pretest confirmed people consider important in modern society, giving participants a stake in doing well. Participants received a bonus of five cents for each correctly answered question. We then ranked them by their score and labeled the 20 highest scorers the High Performer Group and the 20 lowest scorers the Low Performer Group.

Stage 1. We recruited 201 Receivers ($M_{\text{Age}} = 41.17$ years; 51.74% Female). They completed the same 10-item quiz, also earning five cents for each correct answer, and then expressed a non-binding preference for competing against either the High or the Low Performer Group. We informed them that they would be invited back at a later date to make a binding decision and could earn a bonus if their score on this quiz was equal to or higher than that of a

randomly selected member of their comparison group. Our treatment randomized which of the two groups was selected by default; Receivers were free to select the other group. The survey concluded with basic demographic questions.

Stage 2. For the focal part of our experiment, we then recruited 1,000 Senders ($M_{\text{Age}} = 38.95$ years; 59.1% Female). We informed them of the ability quiz that the Receivers had completed, how the Low and High Performer Groups were constructed, and the choice and incentives that Receivers faced.

Senders were randomly assigned to one of two treatments. In the “Performance” treatment, they observed only the Receiver’s score. In the “Performance + Expectation” treatment, they additionally observed the Receiver’s non-binding choice of which group to compete against. Because the outcome depended only on the quiz score, which Senders observed, this choice carried no decision-relevant information. Senders then advised ten Receivers on which group each should compete against. They could earn the same bonus as one of the Receivers they had advised who returned to make a decision, and the survey concluded with basic demographic questions.

Stage 3. Finally, we invited Receivers back for the follow-up survey. Following our preregistration, we kept the survey open for seven days. In total, 176 Receivers returned ($M_{\text{Age}} = 41.9$ years; 51.14% Female). The short survey reminded them of the task they completed in the previous survey, informed them that Senders had observed their real score and given them advice on which group to compete against, and reminded them of their initial choice of comparison group. They also learned that if they chose to compete against the High Performer Group and win, they would earn a bonus of 50 cents, whereas if they chose to compete against the Low Performer Group and win, they would earn a bonus of 20 cents. However, they were not informed of their true score or the scores of the groups they could compete against. Receivers then observed the advice from a randomly selected Sender and made their decision.

Results

We begin by examining the choices and performance of the Receivers. As expected, the default manipulation changed their initial choices: in the “Low Default” treatment, only 22% of

Receivers preferred to compete against the High Performer Group, compared to 65% in the “High Default” treatment ($\chi^2(1, n = 201) = 36.62, p < .001$).

Because our default treatment took place after Receivers completed the ability quiz, we did not expect a difference in performance across the Low and High Default treatments. Indeed, the two groups scored no different from one another (5.10 and 5.18 for the Low Default and High Default treatments, respectively, $t(199) = 0.31, p = .756$).

Table 2

Advice to compete against the High Performer Group in Study 2. Displaying the non-binding choice of Receivers for whom the Low Performer Group was selected by default makes it less likely that they are advised to compete against the High Performer Group (Column 1). Column 2 restricts the analysis to Senders who observe Receivers’ initial non-binding choice and controls for Receivers’ score on the quiz. Standard errors are clustered at the Sender level.

	Advice to Compete Against High Performers	
	All Advice	Expectations Shown
High Default	0.011 (0.013)	0.025* (0.011)
Expectation Shown	−0.041* (0.016)	
Expectation x High Default	0.037+ (0.019)	
Score		0.168*** (0.004)
Constant	0.389*** (0.012)	−0.503*** (0.020)
N	10 000	5000

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We then examined whether Senders took the Receivers’ expectations into account. Because Receivers who scored higher were also more likely to pick the High Performer Group, we follow our preregistration and use the random assignment of the default as an instrument for higher expressed expectations. We estimate linear probability models for the decision of the Senders, equal to 1 if they recommend competing against the High Performer Group and 0 if they

recommend the Low Performer Group. Because each Sender made ten recommendations, we cluster standard errors at the Sender level throughout.


Column 1 of Table 2 shows our main specification with the experimental assignment for Receivers (High vs. Low Default), the assignment for Senders (Expectation Shown vs. Hidden), and their interaction. Displaying the expectations of Receivers in the Low Default treatment reduced recommendations to compete against the High Performer Group by 4 percentage points, or about 11% relative to when expectations were hidden. For Receivers in the High Default treatment — who were nudged to express high expectations, though they did not necessarily hold them — this decrease was marginally smaller, and the advice they received did not differ based on whether their expectations were shown or hidden. Column 2 of Table 2 restricts the sample to Senders in the Expectation Shown condition and controls for Receivers' true quiz score. We observe a significant main effect of the default treatment, where a high default increased the likelihood of receiving flattering advice. Since only 65% of Receivers in the High Default treatment actually retained their default choice, this represents a conservative test of the effect of high expectations on the likelihood of receiving flattering advice. We show these results graphically in the left panel of Figure 4; the right panel previews the analogous pattern for men and women in Study 3.

Discussion

When a Sender knows an outcome that a Receiver can only estimate, learning this estimate should not change the advice that is given. However, in line with our argument that Senders take into account the belief utility of the Receiver and prefer to avoid disappointing them, we find that providing information about expectations makes a difference. Specifically, Receivers who expressed a preference for stronger competitors were more likely to receive flattering advice, encouraging them to pursue that option. This occurs even though Senders share the potential costs of recommending a more challenging competition.

Figure 4

Advice in Studies 2 and 3. Showing the Receiver's expectations to Senders lowers advice to compete for the group with lower expressed expectations — Receivers nudged by a low default (Study 2, left) and women (Study 3, right) — while advice to the high-expectations group is unchanged. Error bars are 95% confidence intervals.



Study 3: Avoiding Disappointment Leads to Gender Bias

In Study 3, we examine whether the desire to avoid disappointment can generate gender disparities in advice, even in the absence of any gender information. We ask Receivers to estimate their score on a mathematics test and anticipate that men will guess higher than women. Moreover, we expect that Senders take into account the expected score when recommending whether someone should compete against a group of High or Low performers. As a result, we hypothesize that, when expectations are known to Senders, men will be advised to compete against High Performers more than women, even when Senders do not know gender and expectations are uninformative.

Methods

Competitor Recruitment. We began by recruiting a sample of 50 participants from Prolific to complete a 10-question multiple choice mathematics quiz. The questions were taken

from a paper version of the ASVAB standardized exam, such that answers are not available online. Participants had five minutes to answer the quiz and were paid 10 cents for each correctly answered question. As in Study 2, we defined the top 20 scorers as “High Performers” and the bottom 20 scorers as the “Low Performers.” On average, participants answered 4.42 questions correctly, High Performers scored between 5 and 10, and Low Performers scored between 0 and 3. To anchor the expectations of Receivers in our main experiment, we simulated 1,000 groups of five participants from this sample, with the 5th percentile of groups scoring an average of 2.6 and the 95th percentile scoring an average of 6.4. We reported these averages to Receivers in the Low Expectations and High Expectations treatments, respectively.

Stage 1. We recruited 1,002 Receivers for Stage 1 of our main experiment. To arrive at a gender-balanced sample, we dropped the last two men to complete the survey, ending up with a sample of 500 men and 500 women ($M_{\text{Age}} = 42.16$ years). Receivers completed the same 10-item mathematics quiz as the preliminary sample and were informed that their performance would affect their bonus earnings in a follow-up stage to be conducted a few days later.

After completing the quiz, Receivers learned the average score of a group of five participants from the preliminary survey. We randomly assigned them to learn about the 5th percentile of groups, which scored 2.6, in the “Low Expectations” treatment, or about the 95th percentile of groups, which scored 6.4, in the “High Expectations” treatment. Receivers then made an unincentivized guess of how many questions they had answered correctly, and the survey concluded with basic demographic questions.

Stage 2. We then recruited 1,000 Senders for Stage 2. We began by informing them of the mathematics quiz that the preliminary sample and the Receivers had completed, along with the average score in the preliminary sample. Senders had to recommend whether a Receiver should compete against the Low Performers or the High Performers, the terms we used in the survey. We anticipated that being told to compete against High Performers is more flattering and hence that being told to compete against the Low Performers would be disappointing to someone who had expected to do well. Receivers would earn a bonus if their score was equal to or higher than that

of a randomly selected member from their chosen group: 50 cents if they outperformed a High Performer and 30 cents if they outperformed a Low Performer.

Senders were randomly assigned to one of two treatments. In the “Performance” treatment, they only observed the score of the Receiver on the mathematics quiz. In the “Performance + Expectation” treatment, they observed the score as well as the Receiver’s guess of how many questions they had answered correctly. As in Study 2, the guess carried no information beyond the score the Sender already observed. Moreover, in *neither* treatment did they receive any demographic information about the Receivers. Senders gave recommendations to ten Receivers, who unbeknownst to them were five men and five women matched to have identical performance on the quiz.⁶ If their advice was shown to a Receiver who returned for the follow-up survey, the Sender would receive the identical bonus as that Receiver. The survey then concluded with basic demographic questions.

Stage 3. Finally, we invited the Receivers back for the follow-up survey. Following our preregistration, we kept the survey open for seven days. In total, 949 Receivers (479 men, 470 women) returned. The brief survey reminded them of the task they completed in Stage 1, informed them that Senders had observed their real score and given them advice on which group to compete against, and reminded them of how many questions they had guessed they answered correctly. Importantly, they were not informed of their true score or the scores of the groups they could compete against. Receivers then observed the advice from a randomly selected Sender who had been paired with them and made their decision.

Results

We begin by examining the performance and expectations of the Receivers. As intended, the expectations treatment affected how well they thought they performed: Receivers in the “Low Expectations” treatment guessed a score of 3.89, compared to 4.42 in the “High Expectations” treatment ($t(998) = 3.78, p < .001$). The manipulation was successful, albeit small in magnitude.

Because our treatment took place after Receivers completed the mathematics quiz, we

⁶ We made this decision to account for the possibility of gender differences in performance.

would not expect a difference in performance across the Low and High Expectations treatments. Indeed, the two groups scored no different from one another (4.79 and 4.55 for the Low Expectations and High Expectations treatments, respectively, $t(998) = -1.64$, $p = .102$). Furthermore, we found a gender difference in performance: men scored 5.02 on average, while women scored 4.33 ($t(998) = -4.73$, $p < .001$). Aligned with our prediction, men expressed greater confidence in their performance than equally performing women, suggesting a confidence gap ($b = 0.77$, 95% CI [0.56, 0.99], $t(997) = 7.15$, $p < .001$). Neither the difference in raw performance nor the difference in beliefs about performance, however, affects the interpretation of our findings, which rely on an interaction between the Receiver's gender and an experimental treatment assigned to Senders.

Table 3


Column 1 displays advice to compete against the High Performer Group based on whether Receivers were primed with low or high expectations, among Senders who observed expectations. Column 2 displays advice to compete against the High Performer Group based on the Receiver's gender and whether expectations were shown to the Sender. Column 3 displays the expected bonus of the advice received based on the Receiver's gender and whether expectations were shown. All standard errors are clustered at the Sender level.

	(1)	(2)	(3)
High Expectation	0.012 (0.011)		
Performance	0.128*** (0.003)		
Expectation Shown		-0.034* (0.013)	0.000 (0.003)
Receiver Male		0.002 (0.005)	0.000 (0.001)
Expectation x Male		0.032*** (0.009)	-0.005** (0.002)
Constant	-0.263*** (0.015)	0.339*** (0.010)	0.258*** (0.002)
N	5000	10 000	10 000

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 5

Advice by quiz score in Study 3. Top: probability of being advised to compete against the High Performer Group for men and women whose expectations were shown to Senders, with 95% confidence intervals; the grey step line is the hidden-expectations baseline. Bottom: the male-female gap in that probability at each score, with the hidden-expectations gap shown as a placebo. The gap opens only in the middle of the score range, where expectations are most informative about what advice would disappoint.



Next, we examined whether Senders took the Receivers' expectations into account. Column 1 of Table 3 reports a linear probability model of advising to compete against the High Performer Group among Senders in the "Performance + Expectation" treatment, controlling for the true performance of the Receiver. Because each Sender made ten recommendations, we cluster standard errors at the Sender level throughout. Contrary to our expectations, we did not find a significant effect of the expectations treatment, possibly because the induced difference in expectations was too small.

However, recall that men were more confident in their performance than equally performing women. Our theory thus predicts that showing expectations should lead to more flattering advice for men than for women — that is, more recommendations to compete against the High Performer Group. Column 2 of Table 3 reports a linear probability model with advice to compete against the High Performer Group as the outcome, and the Receiver's gender, whether expectations were shown to the Sender, and their interaction as predictors. As predicted, we find a significant interaction effect ($b = 0.03$, $SE = 0.01$, $p < .001$): men were more likely than women to be advised to compete against the High Performers when expectations were shown than when they were hidden. We show this result graphically in Figure 5. As the top panel makes clear, this effect was driven by advice given to Receivers who scored in the middle of the possible range; the bottom panel plots the male-female gap in advice directly and shows that it vanishes when expectations were hidden. Thus, when it was clear that someone should compete against Low or High Performers, Senders were not deferring to expectations.

To examine advice quality, we computed the expected bonus earnings for a Receiver who followed the recommendation. For each Receiver, we matched their score against all 20 members of the relevant group, calculated the proportion of members they would outperform, and multiplied this proportion by the respective bonus amounts of 50 cents for the High Performer Group and 30 cents for the Low Performer Group. To see whether including expectations leads to worse advice for men, Column 3 of Table 3 reports a regression of this expected bonus on the experimental treatment of the Sender, the gender of the Receiver, and their interaction. Displaying

expectations led men to be advised to compete against the High Performer Group more often, and this advice turned out to be costly: men received worse advice than women when expectations were displayed, but not in their absence.⁷

To determine whether flattering advice is truly costly, however, we need to examine the outcomes of the Receivers, who could recognize flattering advice for what it is and ignore it. In line with our prediction, Receivers who were primed with high expectations earned less when their expectations were shown to the Senders, although this result is only directional (Column 1 of Table 4). Similarly, as shown in Column 2 of Table 4, men earned less when their expectations were shown. These findings suggest that flattering advice is not without consequences. Moreover, because men were more likely to follow such advice (Column 4 of Table 4), they competed against the High Performer Group more often as a result of flattering advice.

Discussion

Our findings suggest that attempts to avoid disappointment can be a novel source of gender differences in advice. Women underestimate their mathematics test scores more than men do. When these expectations are shown to Senders, Senders are more likely to tell men than women to compete against high performers. Notably, this turns out to be poor advice: men whose Senders are aware of their expectations receive worse advice, and men with the same scores as their female counterparts end up earning less, although the earnings result is only directional. The disparity thus operates through a chain in which men receive more favorable advice, are more likely to follow it, and ultimately face worse outcomes.

General Discussion

While advice is typically studied as helpful input that facilitates better decision making, we argue that advice also conveys an implicit evaluation of the advisee by the advisor. Because

⁷ This analysis was not preregistered, and we note here that the reduction in expected earnings is small. However, it is interesting that expectations have a negative effect for men, who underestimate their performance on average. One possibility is that Senders suggest the High Performer Group more often than is optimal. This is consistent with our finding from Study 1, in which Senders offered flattering advice even when incentivized only for accuracy.

Table 4

Column 1 displays Receivers' actual bonus based on their performance and whether they were primed with high expectations, among Receivers whose expectations were shown to the Sender. Column 2 displays the actual bonus based on performance and the Receiver's gender in the same subsample. Columns 3 and 4 display the likelihood of following the advice based on the Receiver's gender, whether expectations were shown to the Sender, and whether the advice was to compete against the High Performer Group; Column 3 includes only main effects, while Column 4 adds the interactions.

	(1)	(2)	(3)	(4)
Performance	0.038*** (0.003)	0.039*** (0.003)		
High Expectation	-0.020 (0.013)			
Receiver Male		-0.003 (0.013)	-0.046+ (0.026)	-0.144*** (0.041)
Expectation x Male				0.111* (0.053)
Expectation Shown			-0.003 (0.026)	-0.057 (0.040)
Advice: High Performer			-0.138*** (0.028)	-0.206*** (0.051)
Advice: High Performer x Male				0.130* (0.056)
Expectation x Advice: High Performer				-0.010 (0.056)
Constant	0.087*** (0.016)	0.077*** (0.015)	0.858*** (0.024)	0.906*** (0.030)
N	481	481	949	949

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

both parties recognize this evaluative dimension, advice-giving is not a purely informational act but also a social one, and advisors are keenly aware that their recommendations signal how they see the person asking. This creates interpersonal concerns that cause advisors to shade their recommendations toward a more optimistic evaluation than objective performance warrants. Across three preregistered experiments, we document that advisors do not always provide the

most accurate advice available to them. Instead, they provide flattering advice: recommendations that incorporate the advisee's expressed expectations even when objective performance is fully observable and expectations carry no additional information. They do so consistently even when their material incentives are aligned with the advisee's and even when they remain fully anonymous — conditions under which standard accounts of communication predict truthful revelation.

Study 1 shows that advisors use flattering advice as a deliberate interpersonal strategy, and that the strategy works: advisors who flatter are rated as more likable and no less trustworthy. Nor is the tendency merely a response to explicit likability incentives — even advisors rewarded purely for accuracy inflate their recommendations. Study 2 shows that advisors calibrate their advice to the advisee's expressed expectations even when those expectations carry no information the advisor does not already possess. Study 3 traces the consequence for group disparities: because men express greater confidence than equally performing women, advisors who observe stated expectations — but never gender — recommend more ambitious options to men than to women with identical performance, and this seemingly favorable advice ultimately leaves men worse off.

Our findings extend the advice literature by taking the advisor's perspective. The judge-advisor paradigm has been productive in examining how recipients weight or discount the recommendations they receive (Harvey & Fischer, 1997; Snizek & Buckley, 1995; Soll & Larrick, 2009), but its design abstracts away the social context in which advice is exchanged, and less is known about what shapes the advice givers provide. People rarely seek advice as blank slates: they reveal their expectations and aspirations, and in doing so signal how they hope to be evaluated. We show that advisors respond to these signals at the expense of accuracy. Prior work has attributed departures from accurate advice to conflicts of interest (Sah et al., 2013; Sah & Loewenstein, 2015); we find that they emerge even when incentives are fully aligned. And whereas research on the tension between benevolence and honesty has examined how observers judge communicators who face it (Levine et al., 2020; Levine & Schweitzer, 2014), we show how advisors themselves resolve it: not by withholding unfavorable information, but by actively

distorting what they convey.

Our findings also extend the literature on belief-based utility from the demand side of information to the supply side. People derive utility from holding favorable beliefs about themselves and avoid information that threatens those beliefs, even at material cost (Golman et al., 2017; Ho et al., 2021; Loewenstein & Molnar, 2018). We show that informed parties internalize this calculus on others' behalf: advisors anticipate the psychological cost that a disappointing recommendation would impose on the recipient and adjust what they transmit accordingly, at a material cost to themselves. The hedonic distortions documented on the receiving end of information have a mirror image on the sending end.

Finally, we identify a mechanism behind gender disparities in advice that operates without the advisor ever observing gender. Prior accounts attribute the more ambitious advice men receive to stereotypical beliefs about men's and women's capabilities (Bear et al., 2017; Kanze et al., 2018; Moss-Racusin et al., 2012) or to paternalistic motives to shield women from negative evaluations (Gallen & Wasserman, 2024; Jampol & Zayas, 2021) — explanations that require the advisor to perceive or infer the advisee's gender. In our experiments, the gap emerges instead from the interaction between group differences in expressed confidence (Exley & Nielsen, 2024; Lundeberg et al., 1994; Möbius et al., 2022) and advisors' deference to those expressions. Because the gap is jointly produced by both parties, interventions that target advisor attitudes alone (Chang et al., 2019; Paluck & Green, 2009) will leave the mechanism untouched. The dynamic may also be self-reinforcing: ambitious advice emboldens men to pursue challenging opportunities and express still higher expectations, while cautious advice steers women toward safer paths. What the literature treats as a stable group difference may thus be partly the product of a cycle in which the confidence gap and the advice gap feed each other, with well-intentioned advisors unknowingly contributing to a disparity they were trying to avoid.

For organizations, the challenge is that flattering advice is difficult to detect precisely because it feels supportive. Mentors who ground their recommendations in expressed confidence rather than demonstrated performance may strengthen their relationships with advisees while

systematically steering talent toward paths that match stated confidence rather than ability. Our results suggest that the resulting advice gap will not be closed by bias training, which targets beliefs about gender that this mechanism does not require. More promising are approaches that target the mechanism directly: advisory protocols that anchor recommendations to objective performance rather than stated expectations, and calibration exercises that help advice seekers convey their performance accurately. The goal is not to discourage advisees from sharing their expectations, but to equip advisors to recognize how those expectations can distort their recommendations.

Several features of our design suggest boundary conditions worth exploring. Our experiments involved anonymous online dyads, which attenuate the emotional richness of face-to-face advice (Kiesler et al., 1985) and remove any accountability for poor recommendations; the effects we document are likely larger where relationships are ongoing and communication is direct. Relationship closeness itself, however, could cut either way: close advisors are more invested in how their recommendations are received (Feng & Magen, 2016) and are expected to be benevolent rather than bluntly honest (Levine et al., 2020), yet they may also trust the relationship to withstand an uncomfortable truth; whether closeness amplifies or dampens flattery is an open empirical question. The inflation we observe is, moreover, modest. Advisors continued to weight objective performance heavily, flattering enough to encourage but not enough to be incredible. This balance should shift with outcome uncertainty and accountability: where results are hard to predict and poor advice cannot be traced to its source, flattery becomes cheaper. The effects we document in a controlled setting are therefore likely a lower bound on what occurs in the field.

Advice has long been regarded as transmitting accurate information from those who know more to those who know less. That framing misses something fundamental about what advice is. Because advice also conveys an evaluation of the person asking, it is inherently a social act, and advisors are pulled toward recommendations that flatter rather than disappoint. This pull persists when material incentives favor honesty, degrades the quality of advice, and, when it meets the

confidence gap between men and women, produces a gender gap in advice that no one intends.

The kindness is real, but so are its costs.

References

- Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology*, 68(2), 181–198.
<https://doi.org/10.1037/0022-3514.68.2.181>
- Bear, J. B., Cushenbery, L., London, M., & Sherman, G. D. (2017). Performance feedback, power retention, and the gender gap in leadership. *The Leadership Quarterly*, 28(6), 721–740.
- Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations Research*, 33(1), 1–27.
- Bhattacharya, U., Kumar, A., Visaria, S., & Zhao, J. (2024). Do Women Receive Worse Financial Advice? *The Journal of Finance*, 79(5), 3261–3307. <https://doi.org/10.1111/jofi.13366>
- 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.
- Chang, E. H., Milkman, K. L., Gromet, D. M., Rebele, R. W., Massey, C., Duckworth, A. L., & Grant, A. M. (2019). The mixed effects of online diversity training. *Proceedings of the National Academy of Sciences*, 116(16), 7778–7783.
- Chentsova-Dutton, Y. E., & Vaughn, A. (2012). Let me tell you what to do: Cultural differences in advice-giving. *Journal of Cross-Cultural Psychology*, 43(5), 687–703.
- Coffman, K. B., Collis, M. R., & Kulkarni, L. (2024). Whether to apply. *Management Science*, 70(7), 4649–4669.
- Dalal, R. S., & Bonaccio, S. (2010). What types of advice do decision-makers prefer? *Organizational Behavior and Human Decision Processes*, 112(1), 11–23.
- Exley, C. L., & Nielsen, K. (2024). The gender gap in confidence: Expected but not accounted for. *American Economic Review*, 114(3), 851–885.
- Feng, B., & Magen, E. (2016). Relationship closeness predicts unsolicited advice giving in supportive interactions. *Journal of Social and Personal Relationships*, 33(6), 751–767.
- Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth

- and competence. *Trends in Cognitive Sciences*, 11(2), 77–83.
- Gallen, Y., & Wasserman, M. (2024). *Informed Choices: Gender Gaps in Career Advice*.
- Gibbons, R. V., Landry, F. J., Blouch, D. L., Jones, D. L., Williams, F. K., Lucey, C. R., & Kroenke, K. (1998). A comparison of physicians' and patients' attitudes toward pharmaceutical industry gifts. *Journal of General Internal Medicine*, 13(3), 151–154.
- Gneezy, U., Gravert, C., Saccardo, S., & Tausch, F. (2017). A must lie situation - avoiding giving negative feedback. *Games and Economic Behavior*, 102, 445–454.
<https://doi.org/10.1016/j.geb.2017.01.008>
- Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1), 96–135. <https://doi.org/10.1257/jel.20151245>
- Gul, F. (1991). A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, 667–686.
- Hagmann, D., Sajons, G. B., & Tinsley, C. H. (2026). Base rate neglect as a source of inaccurate statistical discrimination. *Management Science*, 72(6), 4974–4999.
<https://doi.org/10.1287/mnsc.2023.00603>
- Haran, U., & Shalvi, S. (2020). The implicit honesty premium: Why honest advice is more persuasive than highly informed advice. *Journal of Experimental Psychology: General*, 149(4), 757.
- Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, 70(2), 117–133.
- Ho, E. H., Hagmann, D., & Loewenstein, G. (2021). Measuring information preferences. *Management Science*, 67(1), 126–145.
- Jampol, L., & Zayas, V. (2021). Gendered white lies: Women are given inflated performance feedback compared with men. *Personality and Social Psychology Bulletin*, 47(1), 57–69.
- John, L. K., Jeong, M., Gino, F., & Huang, L. (2019). The self-presentational consequences of upholding one's stance in spite of the evidence. *Organizational Behavior and Human Decision Processes*, 154, 1–14.

- Kanze, D., Huang, L., Conley, M. A., & Higgins, E. T. (2018). We ask men to win and women not to lose: Closing the gender gap in startup funding. *Academy of Management Journal*, 61(2), 586–614.
- Kiesler, S., Sproull, L., & Eccles, J. S. (1985). Pool halls, chips, and war games: Women in the culture of computing. *Psychology of Women Quarterly*, 9(4), 451–462.
- Kőszegi, B. (2010). Utility from anticipation and personal equilibrium. *Economic Theory*, 44(3), 415–444.
- Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4), 1133–1165.
- Kőszegi, B., & Rabin, M. (2009). Reference-dependent consumption plans. *American Economic Review*, 99(3), 909–936.
- Landis, B., Fisher, C. M., & Menges, J. I. (2022). How employees react to unsolicited and solicited advice in the workplace: Implications for using advice, learning, and performance. *Journal of Applied Psychology*, 107(3), 408.
- Levine, E. E., Hart, J., Moore, K., Rubin, E., Yadav, K., & Halpern, S. (2018). The surprising costs of silence: Asymmetric preferences for prosocial lies of commission and omission. *Journal of Personality and Social Psychology*, 114(1), 29.
- Levine, E. E., Roberts, A. R., & Cohen, T. R. (2020). Difficult conversations: Navigating the tension between honesty and benevolence. *Current Opinion in Psychology*, 31, 38–43.
- Levine, E. E., & Schweitzer, M. E. (2014). Are liars ethical? On the tension between benevolence and honesty. *Journal of Experimental Social Psychology*, 53, 107–117.
- Loewenstein, G., & Molnar, A. (2018). The renaissance of belief-based utility in economics. *Nature Human Behaviour*, 2(3), 166–167.
- Lundeberg, M. A., Fox, P. W., & Punčochař, J. (1994). Highly confident but wrong: Gender differences and similarities in confidence judgments. *Journal of Educational Psychology*, 86(1), 114.
- Mackinger, B., Jonas, E., & Mühlberger, C. (2017). When advisors' true intentions are in

- question. How do bank customers cope with uncertainty in financial consultancies? *Frontiers in Psychology*, 8, 1112.
- Milyavsky, M., & Gvili, Y. (2024). Advice taking vs. Combining opinions: Framing social information as advice increases source's perceived helping intentions, trust, and influence. *Organizational Behavior and Human Decision Processes*, 183, 104328.
- Möbius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science*, 68(11), 7793–7817.
- Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502.
- Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. (2012). Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41), 16474–16479.
- Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *The Quarterly Journal of Economics*, 122(3), 1067–1101.
- Nosek, B. A., Smyth, F. L., Sriram, N., Lindner, N. M., Devos, T., Ayala, A., Bar-Anan, Y., Bergh, R., Cai, H., Gonsalkorale, K., et al. (2009). National differences in gender–science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences*, 106(26), 10593–10597.
- Palmeira, M., Spassova, G., & Keh, H. T. (2015). Other-serving bias in advice-taking: When advisors receive more credit than blame. *Organizational Behavior and Human Decision Processes*, 130, 13–25.
- Paluck, E. L., & Green, D. P. (2009). Prejudice reduction: What works? A review and assessment of research and practice. *Annual Review of Psychology*, 60, 339–367.
- Rader, C. A., Larrick, R. P., & Soll, J. B. (2017). Advice as a form of social influence: Informational motives and the consequences for accuracy. *Social and Personality Psychology Compass*, 11(8), e12329.
- Reif, J. A., Larrick, R. P., & Soll, J. B. (2024). The inclusion of anchors when seeking advice:

- Causes and consequences. *Organizational Behavior and Human Decision Processes*, 185, 104378.
- Sah, S., & Loewenstein, G. (2015). Conflicted advice and second opinions: Benefits, but unintended consequences. *Organizational Behavior and Human Decision Processes*, 130, 89–107. <https://doi.org/10.1016/j.obhdp.2015.06.005>
- 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.
- Schaerer, M., Tost, L. P., Huang, L., Gino, F., & Larrick, R. (2018). Advice giving: A subtle pathway to power. *Personality and Social Psychology Bulletin*, 44(5), 746–761.
- See, K. E., Morrison, E. W., Rothman, N. B., & Soll, J. B. (2011). The detrimental effects of power on confidence, advice taking, and accuracy. *Organizational Behavior and Human Decision Processes*, 116(2), 272–285.
- Shalvi, S., Soraperra, I., Weele, J. J. van der, & Villeval, M. C. (2019). *Shooting the messenger? Supply and demand in markets for willful ignorance*.
- Snizek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, 62(2), 159–174.
- Soll, J. B., & Larrick, R. P. (2009). Strategies for revising judgment: How (and how well) people use others' opinions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3), 780.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213–225.
- Van Swol, L. M. (2009). The effects of confidence and advisor motives on advice utilization. *Communication Research*, 36(6), 857–873.
- Weiner, B. (1985). An attributional theory of achievement motivation and emotion. *Psychological Review*, 92(4), 548.
- Zhang, Y., & Epley, N. (2012). Exaggerated, mispredicted, and misplaced: When “it’s the thought

that counts” in gift exchanges. *Journal of Experimental Psychology: General*, 141(4), 667.