UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Seeing Things Differently: The Role of Differing Perspectives in Advice-Taking

Permalink

https://escholarship.org/uc/item/7w8586gz

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors

Kumar, Aakriti
Tham, Ren-Hui Michelle
Steyvers, Mark

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Seeing Things Differently: The Role of Differing Perspectives in Advice-Taking

Aakriti Kumar (aakriti.kumar@kellogg.northwestern.edu)
Kellogg School of Management, Northwestern University, Evanston, IL 90202 USA

Ren-hui Michelle Tham (@uci.edu)
Department of Cognitive Sciences, University of California, Irvine, Irvine, CA 92697 USA

Mark Steyvers (mark.steyvers@uci.edu)
Department of Cognitive Sciences, University of California, Irvine, Irvine, CA 92697 USA

Abstract

Advice-taking plays a critical role in collaboration. Yet people tend to under-utilize advice, often to their own detriment. We investigate if people's utilization of advice improves when they know the advisor has access to different information compared to them. We examine how individuals integrate advice in an estimation task, where the advisee and the advisor have access to different perspectives of the same problem. We assess how individuals adjust their estimates when presented with estimates from a human advisor and an AI advisor, and when they are given information about the advisor's perspective. Our findings are consistent with egocentric discounting where individuals exhibit a general bias toward their own information. However, this discounting is lower for AI advisors compared to human advisors in our experiment. Our results also show that the advisor's estimate is taken more into account when the advisor has a more favorable viewpoint – for both human and AI advisors. This suggests a potential for optimizing advice-taking behavior by enhancing people's understanding of the advisor's viewpoint. This study furthers our understanding of advice-taking dynamics with human and AI advisors and the role of perspective-taking in decision-making processes.

Keywords: advice-taking; assisted decision-making; perspective-taking; egocentric discounting; human-ai collaboration

Introduction

Research on advice-taking reveals a common tendency: people undervalue suggestions from others compared to their own opinions (Bonaccio & Dalal, 2006) despite evidence that advice typically improves performance (Kämmer, Choshen-Hillel, Müller-Trede, Black, & Weibler, 2023). Concurrently, extensive research on overconfidence shows that people perceive their abilities as superior to those of their peers, exhibiting unwarranted confidence in themselves (Moore & Cain, 2007). Working with AI is no different. Research has shown that humans are susceptible to a variety of misjudgements and biases when seeking advice from machines (Goddard, Roudsari, & Wyatt, 2012; Dietvorst, Simmons, & Massey, 2015; Logg, Minson, & Moore, 2019). As humans interface more with AI assistants, it is important to understand how they incorporate AI advice in their decisions. This combination of undervaluing advice and overvaluing personal judgment presents a significant challenge in decision-making and advice utilization.

Several explanations have been proposed to explain why people may underutilize advice. Some suggest that people fail to appreciate the benefit of incorporating diverse perspectives (Larrick & Soll, 2006), while others attribute it to individual differences in agency (Schultze, Gerlach, & Rittich, 2018) or narcissism (Kausel, Culbertson, Leiva, Slaughter, & Jackson, 2015). A leading explanation for these tendencies in quantitative estimation tasks is egocentric discounting, a phenomenon where individuals underutilize external advice often due to factors such as asymmetric access to information (Yaniv & Kleinberger, 2000; Yaniv, 2004). It posits that individuals tend to favor their judgment over advice from others and struggle to accurately incorporate alternative viewpoints when evaluating advice. This is further confirmed by (Bailey, Leon, Ebner, Moustafa, & Weidemann, 2022) through a meta-analysis that shows people generally give more weight to their own estimates than to advice from others. Interestingly, some studies show that people are more likely to discount advice that differs from their initial beliefs (Yaniv & Milyavsky, 2007; Himmelstein & Budescu, 2023), while others show that advice that mirrors one's own beliefs is also discounted, being perceived as confirmation rather than new information (Allahverdyan & Galstyan, 2014; Himmelstein & Budescu, 2023). These findings suggest a complex interplay where the difference between the advice and the advisee's initial beliefs may influence the advice-taking process.

Previous research also indicates that people place greater weight on advice when they perceive the advisor to be an expert based either on the advisor's demonstrated accuracy in prior tasks or on the advisor's presumed capability based on status (Önkal, Gönül, Goodwin, Thomson, & Öz, 2017). Yaniv and Kleinberger (2000) also emphasize the dynamic nature of trust and the importance of advisor reputation on people's advice-taking behavior. However, these findings are grounded in scenarios where the advisee and the advisor do the same task and have access to identical information. In contrast, the key focus of our paper is how individuals integrate advice from others when the information available to the advisor differs from their own. Do people still exhibit egocentric discounting in such scenarios? Does this preference for one's own judgment over others' advice indicate a failure in perspective-taking? Or do individuals possess the capacity to discern between various forms of advice, selectively incorporating them based on the relevance of the advisor's information?

Through a meta-analysis, Bailey et al. (2022) found that the information about the *quality of advice* was the most significant predictor of advice-taking. In contrast, the key focus of our paper is to investigate if people can strategically alter their reliance on the advisor depending on the *quality of in-*

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society* ©2025 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

formation the advisor has. In particular, we investigate how visual perspective-taking influences advice integration in a visual quantity estimation task. Research on visual perspective taking has demonstrated that people are capable of accurately predicting the visual experience of other agents (Michelon & Zacks, 2006). However, it is unclear if this ability extends to advice integration, especially when the advisor has a superior vantage point to address a particular problem. Our research probes the interplay between egocentric discounting and perspective-taking in the context of advice integration. We examine how individuals integrate the advisor's estimate when individuals have knowledge about the quality of information available to the advisor.

In our experimental setup, an individual, referred to as the advisee, is tasked with estimating the number of objects in a jar, as shown in Figure 2. They are then provided with the estimate made by another individual, the advisor, regarding the same jar's contents. This setup mirrors a judge-advisor system (JAS) a widely used framework in decision-making research (Bailey et al., 2022; Himmelstein & Budescu, 2023; Önkal et al., 2017), where a judge initially makes a judgment and then has the chance to revise it after receiving advice. Unlike traditional JAS setups, our experiment presents the advisor and advisee with different views of the jar. These distinct viewpoints make the process of counting the objects more or less challenging.

The implications of our research extend beyond the scope of our experiments. They are relevant in various real-world contexts, including decision-making assisted by AI systems and collaborative problem-solving. The increasing prevalence of AI in decision-making (Steyvers & Kumar, 2022; Patel et al., 2019; Grgić-Hlača, Engel, & Gummadi, 2019) necessitates understanding the interaction between humans and AI advisors. This includes exploring different workflows, such as making AI advice always available or available only on demand to improve the efficacy of assisted decision-making. Advances in explainability methods also play a critical role in demystifying AI decisions for users (Ignatiev, 2020).

In the following sections, we seek to answer the following questions: Can the advisee, when informed about the advisor's view of the jar, accurately assess the difference in favorability between their view and the advisor's? When given the opportunity to revise their estimate, do individuals appropriately integrate the advisor's assessment into their estimation? How does people's weighting of advice differ across human and AI advisors? Furthermore, do individuals make optimal decisions based on who possesses the more advantageous view?

Overview of Experiments

The goal of our study is to assess advice-taking in situations where there is information asymmetry between the advisor and the advisee. To explore this theme, we conduct three experiments, with Experiments 1 and 2 serving as preliminary investigations that establish key background insights, and Experiment 3 as the primary focus of this paper.

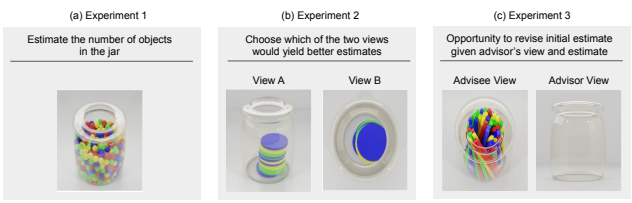

Figure 1: Illustration of the behavioral tasks in Experiments 1-3 to investigate advice-taking with information asymmetry between advisee and advisor.

Our primary task involves estimating the number of objects in a jar, a well-established paradigm for studying numerical estimation (Bevan, Maier, & Helson, 1963) and conformity (Jenness, 1932). To systematically vary information asymmetry, we present images of jars containing cylinders, disks, or spheres, each viewed from five different angles, as shown in Figure 2: 0° (side view), $22^\circ, 45^\circ, 66^\circ$ (intermediate viewing angles), and 90° (top view).These angles influence the visibility of objects, making some perspectives more favorable than others for estimation. For example, as shown in Figure 2, the top view (90°) allows for counting the cross-sections of cylinders but gives no information about the depth of spheres or disks. Conversely, the side view (0°) is favorable for counting disks but is not suitable for spheres or cylinders. Figure 1 illustrates the behavioral tasks across all three experiments.

In Experiment 1, we assess how object shape and jar viewing angle affect estimation accuracy. The estimates are made independently without any advisor. We use the estimates from Experiment 1 as advisor estimates in Experiment 3. In Experiment 2, we assess individuals' ability to identify the most favorable jar view for estimation to establish that individuals are capable of understanding the circumstances where another person would have better quality of information. Experiment 3 investigates how participants evaluate, process, and use the inputs provided by an advisor when they are presented with information regarding the advisor's viewpoint. In one condition, participants are presented with a human advisor's estimates (from Experiment 1). They are tasked with using this advice to formulate their own estimates, aiming for maximum accuracy. The experiment manipulates the perspective of the jar available to the advisee and advisor. In another condition, participants receive estimates from an AI advisor, along with information about the AI's viewpoint of the jar. We use a Wizard-of-Oz setup wherein the advice, although originating from a human, is portrayed as coming from an AI agent. Participants again have the option to use the AI advisor's estimate to produce an estimate that is as accurate as possible.

Methods

Stimuli

The images of jars used in the experiments were created using the Python interface of the Blender software (Blender, 2022). To capture different viewpoints of jars corresponding to the different 'views' accessible to the advisor and the advisee, we designed a standard jar using Blender and placed it in a virtual studio setup available on the platform. Next, we added five cameras to the scene: one positioned horizontally in front of the jar (0°), another looking down from the top (90°), and three placed in between (22°, 45°, 66°). We generated three distinct objects—spheres, cylinders, and disks—that were placed inside the jar. These objects were purposefully selected to create perspectives that are differentially favorable for counting each type of object. For instance, the side view is optimal for counting disks, whereas the top view works better for counting cylinders. The number of objects in each jar was randomly determined, falling within specified ranges of 50 and 1000 for spheres, 5 and 200 for cylinders, and 5 and 50 for disks. Finally, to virtually place the objects in the jar, we leveraged Blender's physics engine to simulate the dropping of objects into the jar. Five cameras captured snapshots of the jar once the objects were settled in their final positions in the jar. The final set of stimuli consisted of 150 images: 3 types of objects (spheres, cylinders, and disks) x 5 viewing angles x 10 jars with different numbers of objects. Each jar may be observed from five viewing angles.

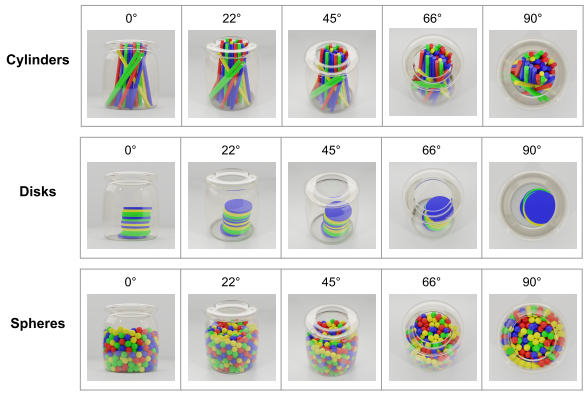

Figure 2: Illustration of the stimuli across viewpoints (columns) and types of objects (rows). For each type of object, the images show the jar with the same number of objects in the same configuration.

Procedure

In Experiment 1, 100 participants estimated the number of objects in 30 jar images. Participants provided their estimates for each jar in a text field, which accepted only numeric input. Figure 1 (a) shows a schematic of the experiment. The assignment of jar images to participants was counterbalanced, such that there were approximately 20 judgments per unique image. Participants were recruited from Prolific received $3.00 for their participation and $1.50 for good performance (i.e., if their performance is within the top 30% across all participants).

In Experiment 2, 200 participants completed a total of 30 trials to predict which of two views of the same jar would lead to the most accurate estimate. The order of the 30 unique jar images was randomized. The first view of the jar, referred to as 'View A', was chosen out of the 5 views. The second view, or 'View B', was chosen from the 4 remaining views. The combination of views across A and B was counterbalanced such that there were approximately 10 judgments per unique view combination. Figure 1 (b) illustrates the experiment. Participants from Prolific received $2.00 for their participation and $1.00 for good performance.

Experiment 3 involved a two-step estimation task. First, participants independently estimated the number of objects in a jar. Next, they viewed their advisor's perspective (without the objects in it) of the same jar along with the advisor's estimate (Figure 1(c)) before providing a final, revised estimate. Advisors were randomly selected from Experiment 1, with each yoked to two Experiment 3 participants. Although viewing the same jar, participants' and advisors' perspectives could vary and these view pairings were counterbalanced (approx. 8 estimates per combination). Participants (N=275, recruited via Prolific) were informed their advisor was either human (n=200) or AI (n=75). They received $3.00 for participation, plus a $1.50 bonus for performance in the top 30%. Figure 1 (c) illustrates the experiment.

Analysis

Mean Estimation Error To evaluate participants' estimation error for different images, we computed the Mean Absolute Error (MAE) for each viewing angle and shape combination. If X is the true number of objects in a jar, and X_i is person i's estimate of the number of objects in that jar, the MAE for that trial is $|X - X_i|$. Across the three experiments, we used a logarithmic transformation to standardize participants' estimates.

Weight of Advice Weight of Advice (WoA) is measured by how much the advisee's estimate shifts towards the advice given, proportionally to the difference between their original belief and the advice (Hogarth & Einhorn, 1992; Soll & Larrick, 2009; Bailey et al., 2022; Himmelstein, 2022). WOA typically has a trimodal distribution (Soll & Larrick, 2009), with spikes at 0 (advice declined) and 1 (advice adopted). WoA (w) can be understood through a simple mathematical relationship involving the initial estimate of the advisee (c), the estimate provided by the advisor (p), and the revised estimate of the advisee after receiving the advice (r). The revised estimate of the advisee is a weighted average of the advisee's initial estimate and the advisor's estimate: $w = \frac{r-c}{p-c}$.

WoA measures how much the advisee's estimate shifts towards the advisor's estimate relative to their initial estimate.

If the advisee ignores the advisor's advice, WoA is 0, and if the advisee fully adopts the advisee's estimate, WoA is 1. We utilize WoA as a metric to infer participants' reliance on the advisor's estimate. To handle negative and extreme WoA values, any empirical WoA less than 0 was clipped to 0, and any value greater than 1 was clipped to 1.

Bayesian Regression Models The key question in this study is whether participants take into account the advisor's view relative to their own view when weighting the advisor's estimate. We investigate this question by examining three regression models that aim to predict the Weight of Advice (WoA) based on various types of information about the advisee and advisor perspectives. Bayesian linear regression models were used to compare different models of how individuals determine the weight of advice based on their own and the advisor's viewpoint. All models encoded the independent variables as categorical variables.

The *Relative Favorability* model allows for the possibility that the WoA is dependent on the relative favorability of the advisor and advisee views. The model represents relative favorability based on the preference probability derived from Experiment 2. Relative favorability quantifies the perceived informational advantage of the advisor's viewing angle compared to the advisee's. To operationalize this, we used the preference probabilities derived from Experiment 2, where participants chose which of two views of the same jar would likely yield a more accurate estimate. We divided preference probabilities into three categories: less than .33 (advisee view is preferred over advisor view), between .33 and .66 (no clear preference for advisor or advisee view), and greater than .66 (advisor view is preferred over advisee view). The goal of this model is to determine how people's preferences for the advisor's point of view influenced the weight they assigned to the advisor's estimate.

The second and third models both function as null models. The *Advisee View* model assumes that only the advisee view predicts the WoA and ignores the advisor view (e.g., the advisee may rely more on the advisor if their view is less favorable regardless of whether the advisor has a favorable view). Similarly, in the *Advisor View* model, the only determinant of WoA is the advisor's point of view and ignores the advisee's point of view. We used the BayesFactor package (Morey, Rouder, & Jamil, 2018) in R to implement the models and to calculate the Bayes factor, a metric for assessing the strength of evidence in Bayesian Statistics. Bayes Factors allow us to quantify the evidence supporting each model relative to others. Each of the three regression models was applied separately to each object shape.

Preliminary Findings

Estimation Error Decreases as the Viewing Angle of the Jar Increases

Figure 3 shows participants' mean estimation error for each shape and viewing angle in Experiment 1. Notably, for cylinders, estimation error decreases as the viewing angle of the jar increases. The opposite pattern is observed for disks, where estimation error increases with the viewing angle. For spheres, the results reveal a non-monotonic and U-shaped trend with the intermediate angles having a lower estimation error than 0 and 90 degrees. In comparison to disks and cylinders, estimation error and variance for spheres are considerably higher. This higher error in estimation can be attributed to the larger number of spheres in the stimuli (Chesney, Bjalkebring, & Peters, 2015). Our findings indicate that participants exhibit reasonably good performance in the estimation task, with higher errors for less favorable viewpoints and lower errors for more favorable viewpoints.

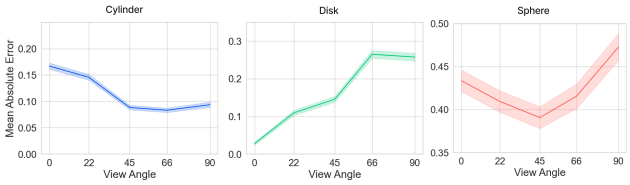

Figure 3: Participants' mean absolute errors for cylinders, disks, and spheres across various viewing angles. The colored bands represent the standard error of the mean.

Participants Favor Viewing Angles that Lead to Lower Estimation Errors

To quantify the relative preference of View B over View A, we calculate the preference probability for View B. Suppose the number of people who chose View A is X, and the number of people who chose View B is Y.The preference probability for View B is then evaluated as $\frac{Y}{X+Y}$. This metric helps us understand how likely participants are to prefer View B over View A.

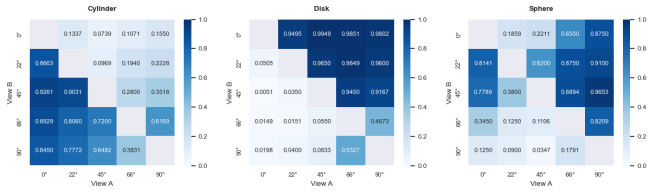

Figure 4: Preference probability for View B over View A for each shape and viewing angle combination.

Figure 4 illustrates participants' view preferences from Experiment 2. Darker shades of blue indicate a stronger preference for View B over View A. Notably, for cylinders, darker cells are primarily clustered below the diagonal and as the angle of View A increases, the cells in each row transition from a darker to a lighter shade. This result shows that participants not only favor views with larger angles but also exhibit a stronger preference when the difference between the

two viewing angles is larger. Conversely, for disks, the dark blue cells are concentrated above the diagonal. The shade of this area is darker and more homogeneous compared to the lower triangular area of the cylinder. This shows that participants have a strong preference for lower angles, and the difference between the angles of View B and View A does not strongly impact the magnitude of this preference. Finally, for the spheres, cells tend to be darker when View B is at a more intermediate angle compared to View A. In addition, the top view (90°) is less preferred over any other viewpoint, presumably because participants intuit that the top view leads to ambiguities in the overall count. Overall, the results from Experiment 2 indicate that participants have a strong an intuitive understanding that some viewpoints lead to lower quality information. Their preferences for certain perspectives (Figure 4) align well with views that were empirically shown to result in lower estimation errors (Figure 3).

	Human Advisor		AI Advisor	
Shape	Advisee View	Advisor View	Advisee View	Advisor View
Disk	1.7×10^{36}	3.5×10^{21}	7.4×10^{14}	6.1×10^{18}
Sphere	6.2×10^{2}	7.4×10^{3}	6.2	1.5×10^{3}
Cylinder	8.2	4×10^{4}	2×10^{1}	2×10^{1}

Table 1: Bayesian Linear Regression Analysis on Weight of Advice. The table compares the Relative Favorability model with the two baseline models (Advisee and Advisor View models) across human and AI advisors, and three object types: Disks, Spheres, and Cylinders. The values reflect Bayes Factors of the Relatively Favorability model against the Advisee and Advisor View models.

Results

Participants exhibited mean absolute errors of .13, .16, and .38 in independently estimating cylinders, disks, and spheres in the jar in the human advisor condition. In the AI advisor condition, participants had mean absolute errors of .12, .15, and .38 when estimating the number of cylinders, disks, and spheres respectively. These results are similar to the performance of participants in Experiment 1, where the mean absolute errors were .12, .16, and .43 for the same shapes. Notably, when participants had access to an advisor's estimate, in the human advisor condition, their errors reduced to .10, .11, and .36 and in the AI advisor condition, their errors reduced to .09, .10, and .36 for cylinders, disks, and spheres. The average weight given to the human advisor's advice was .38, .39, and .39, while the average weight given to the AI advisor's advice was .49, .46, and .49 for each of these shapes.

Relative Favorability of Views Modulates Weight of Advice

Each of the three regression models was applied separately to each object shape in Experiment 3. The results of this analysis are presented in Table 1. The Bayes Factors represent the relative evidence of the Relative Favorability model over the two null models (Advisor and Advisee View). For both the Disks and Spheres, the results provide substantial evidence (BF $\geq$ 100) for the Relative Favorability model, suggesting participants weigh advice based on the Relative Favorability of different views rather than solely their own or the advisor's view. For Cylinders, the Bayes factor of 8.2 for the comparison between Relative Favorability and Advisee View suggests more modest evidence in favor of Relative Favorability. In contrast, The Bayes factor of 4×10^4 for Relative Favorability' against Advisor View indicates very strong evidence for Relative Favorability. In summary, our findings reveal that participants tend to base their advice-taking decision on the relative favorability of views instead of solely relying on their view or the advisor's view of the jar.

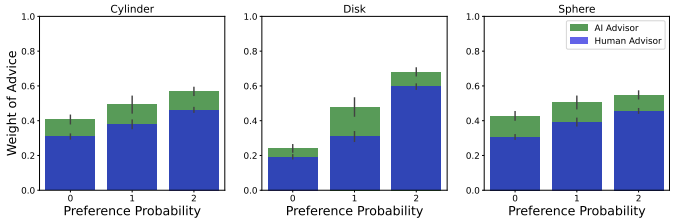

Figure 5: People's observed weight of advice across different preference probabilities for both human and AI advisors. People's WoA increased as the advisor's view became more favorable compared to their own. Error bars show standard error of the mean.

Difference between Observed and Optimal Weight of Advice

We also investigate to what extent participants are optimal in their weighting of advice. To excel in the advice-taking task, an ideal participant would need to adjust their weight of advice to minimize their estimation error. The optimal WoA would also have to take into account the combination of advisor and advisee views. To formalize the analysis of optimal WoA, we introduce some notation. For each trial i in a collection of n advisor-advisee view combination trials, let t_i be the true number of objects in the jar, p_i be the advisor's estimate, c_i be the advisee's initial estimate, and w be the weight of advice. The participant's objective is to minimize: $\min(t_i - (w_i p_i + (1-w_i)c_i))^2$; which gives us a closed-form expression for the optimal weight of advice for trial i, $w_i^{opt} == \frac{t_i - c_i}{p_i - c_i}$.

Figure 6 illustrates the comparison between participants' observed and optimal weight of advice for different preference probabilities for both human and AI advisors. Similar to empirical WoA, to handle negative and extreme optimal WoA values, any optimal WoA less than 0 was clipped to 0, and any value greater than 1 was clipped to 1. Our findings reveal two distinct, consistent patterns across all shapes. First, relative favorability of the views matters. As the advisor's view becomes more favorable relative to the advisee's view of the jar, observed WOA increases. This is evident in the up-

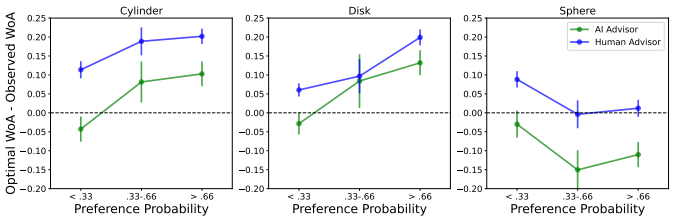

Figure 6: Deviation of people's observed weight of advice from the optimal weight of advice across different preference probabilities for both human and AI advisors. We see: 1) increase in observed WOA as the advisor's view became more favorable, and 2) consistent underutilization of advice (egocentric discounting) especially for the human advisor.

ward trend of WOA values across preference bins. Second, we see a notable discrepancy across all preference probabilities: the observed WOA consistently falls short of the optimal WoA. This indicates that individuals tend to assign less weight to the advisor's estimate than would be considered optimal. This reflects a bias where participants are inclined to favor their estimates over those of the advisor, even though it negatively impacts their performance. However, this discounting is lower for the AI advisor compared to the human advisor.

Discussion

Collaboration often involves seeking and incorporating advice from others. However, research on advice-taking consistently reveals that people tend to under-utilize the advice they receive (Kämmer et al., 2023; Bonaccio & Dalal, 2006; Bailey et al., 2022). Our research takes a novel approach by investigating how individuals process and integrate advice when they are aware that their advisor has access to different information. We show that people are able to evaluate the relevance and utility of the information available to the advisor for a specific task, enabling them to make informed adjustments in their reliance on advice.

We examined participants' performance in a simple, yet challenging visual counting task. This task, inspired by everyday experiences, involves estimating the number of objects in a jar from various perspectives. Through two preliminary studies, we learned that participants were better at the estimation task when shown some views of the jar compared to others. This is intuitive – some perspectives of the jar obstruct critical information. For example, a top view of the jar with spheres obscured the depth of the objects, impacting the accuracy of people's estimates. We also confirmed that people's intuitions were correct - people preferred views of the jar which led to lower estimation errors. Finally, in Experiment 3, we presented another person's estimate and view of the jar (tagged as human or AI advice) as additional information that participants could use to improve their estimates. Our analysis revealed that participants adjusted their initial estimates to incorporate the advisor's estimate, with the degree of adjustment or weight of advice increasing when the advisor had a more favorable view of the jar. This suggests participants have a nuanced understanding of how different perspectives can influence the accuracy of estimates.

Consistent with prior work, our study reveals persistent egocentric discounting even when individuals acknowledge an advisor's different perspective, highlighting an opportunity for improvement. Yaniv and Choshen-Hillel (2012) showed that while predicting an advisor's view can initially decrease this bias, its effects may fade, and they suggested perspective-taking as a solution. Kämmer et al. (2023) also identified information asymmetry regarding the advisor's rationale as a factor. Our research supports these ideas, indicating that clarifying the foundation of an advisor's viewpoint (their visual perspective) can lessen, but not fully alleviate, egocentric discounting.

Our results also align with a growing body of literature on human-AI collaboration, which suggests that people often interact with AI systems in ways that differ from their interactions with humans (Logg et al., 2019; Steyvers & Kumar, 2022; Kelly, Kumar, Smyth, & Steyvers, 2023). A key finding from our study is the reduced egocentric discounting observed when advice came from an AI advisor compared to a human. Future research should directly investigate the mechanisms underlying this differential weighting of advice. For example, measuring participants' trust in and perceptions of the AI advisor (e.g., perceived objectivity, competence) could help to disentangle the relative contributions of trust, perceived rationality, and social factors. Incorporating dialogue, and manipulating the justification provided by the AI (e.g., 'I have a clear view' vs. 'I am using a sophisticated algorithm') could also shed light on the relative importance of transparency versus perceived algorithmic expertise.

Understanding the dynamics of advice-taking and egocentric discounting can significantly improve decision-making processes, ultimately enhancing the quality of collective judgments and facilitating more effective collaboration. While our work contributes to a deeper understanding of these dynamics, several aspects remain unexplored. One key limitation of our study is the absence of feedback from participants during the experiment. In real-life scenarios, feedback plays a crucial role in refining one's understanding of both the advisor and the task, allowing for adjustments in mental models based on actual outcomes. Additionally, our study does not examine the interaction between the overall proficiency of the advisor at the task and their access to differential information. One important empirical extension is to look at how people account for advisor reputation when incorporating advice, especially in scenarios with varying levels of information availability.

References

Allahverdyan, A. E., & Galstyan, A. (2014). Opinion dynamics with confirmation bias. *PloS one*, *9*(7), e99557.

Bailey, P. E., Leon, T., Ebner, N. C., Moustafa, A. A., & Weidemann, G. (2022). A meta-analysis of the weight of advice in decision-making. *Current Psychology*, 1–26.

Bevan, W., Maier, R. A., & Helson, H. (1963). The influence of context upon the estimation of number. *The American Journal of Psychology*, *76*(3), 464–469.

Blender. (2022). Blender (version 3.1.0) [computer software].

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.

Chesney, D., Bjalkebring, P., & Peters, E. (2015). How to estimate how well people estimate: Evaluating measures of individual differences in the approximate number system. *Attention, Perception, & Psychophysics*, *77*, 2781–2802.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, *144*(1), 114.

Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: a systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, *19*(1), 121–127.

Grgić-Hlača, N., Engel, C., & Gummadi, K. P. (2019). Human decision making with machine assistance: An experiment on bailing and jailing. *Proceedings of the ACM on Human-Computer Interaction*, *3*(CSCW), 1–25.

Himmelstein, M. (2022). Decline, adopt or compromise? a dual hurdle model for advice utilization. *Journal of Mathematical Psychology*, *110*, 102695.

Himmelstein, M., & Budescu, D. V. (2023). Preference for human or algorithmic forecasting advice does not predict if and how it is used. *Journal of Behavioral Decision Making*, *36*(1), e2285.

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, *24*(1), 1–55.

Ignatiev, A. (2020). Towards trustable explainable ai. In *Ijcai* (pp. 5154–5158).

Jenness, A. (1932). The role of discussion in changing opinion regarding a matter of fact. *The Journal of Abnormal and Social Psychology*, *27*(3), 279.

Kämmer, J. E., Choshen-Hillel, S., Müller-Trede, J., Black, S. L., & Weibler, J. (2023). A systematic review of empirical studies on advice-based decisions in behavioral and organizational research. *Decision*.

Kausel, E. E., Culbertson, S. S., Leiva, P. I., Slaughter, J. E., & Jackson, A. T. (2015). Too arrogant for their own good? why and when narcissists dismiss advice. *Organizational Behavior and Human Decision Processes*, *131*, 33–50.

Kelly, M., Kumar, A., Smyth, P., & Steyvers, M. (2023). Capturing humans' mental models of ai: An item response theory approach. In *Proceedings of the 2023 acm conference on fairness, accountability, and transparency* (pp. 1723–1734).

Larrick, R. P., & Soll, J. B. (2006). Intuitions about combining opinions: Misappreciation of the averaging principle. *Management science*, *52*(1), 111–127.

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90–103.

Michelon, P., & Zacks, J. M. (2006). Two kinds of visual perspective taking. *Perception & psychophysics*, *68*, 327–337.

Moore, D. A., & Cain, D. M. (2007). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes*, *103*(2), 197–213.

Morey, R. D., Rouder, J. N., & Jamil, T. (2018). *Bayesfactor: Computation of bayes factors for common designs. r package version 0.9. 12-4.2.*

Önkal, D., Gönül, M. S., Goodwin, P., Thomson, M., & Öz, E. (2017). Evaluating expert advice in forecasting: Users' reactions to presumed vs. experienced credibility. *International Journal of Forecasting*, *33*(1), 280–297.

Patel, B. N., Rosenberg, L., Willcox, G., Baltaxe, D., Lyons, M., Irvin, J., ... others (2019). Human–machine partnership with artificial intelligence for chest radiograph diagnosis. *NPJ Digital Medicine*, *2*(1), 1–10.

Schultze, T., Gerlach, T. M., & Rittich, J. C. (2018). Some people heed advice less than others: Agency (but not communion) predicts advice taking. *Journal of Behavioral Decision Making*, *31*(3), 430–445.

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.

Steyvers, M., & Kumar, A. (2022). Three challenges for ai-assisted decision-making.

Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational behavior and human decision processes*, *93*(1), 1–13.

Yaniv, I., & Choshen-Hillel, S. (2012). When guessing what another person would say is better than giving your own opinion: Using perspective-taking to improve advice-taking. *Journal of Experimental Social Psychology*, *48*(5), 1022–1028.

Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational behavior and human decision processes*, *83*(2), 260–281.

Yaniv, I., & Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgments. *Organizational Behavior and Human Decision Processes*, *103*(1), 104–120.