Flexible use of confidence to guide advice requests

Nomi Carlebach & Nick Yeung

Department of Experimental Psychology, University of Oxford

Author Note

Correspondence concerning this paper should be addressed to Nomi Carlebach,
Department of Experimental Psychology, University of Oxford, Anna Watts Building, Oxford OX2

6GG, UK.

Email: naomi.carlebach@psy.ox.ac.uk

Abstract

A growing body of evidence indicates that subjective confidence plays an important role in guiding behaviour. Past studies have demonstrated a fixed relationship between confidence and behaviour, so that low confidence leads to one course of action, and high confidence to another. Here, we tested whether people display flexibility in their use of confidence, so that the mapping between confidence and behaviour is not necessarily fixed, but can instead vary depending on the specific context. This proposal was studied in three experiments, testing the hypothesis that the seemingly natural relationship between low confidence and advice requests varies according to whether people know, or do not know, the quality of the advice. Participants made an initial perceptual judgement, and then had to choose between re-sampling evidence or receiving advice from a virtual advisor, before making a final decision. The results from these experiments showed that, when no objective information about advisor reliability was available, people selected advice more often when their confidence was high rather than low, reflecting the use of confidence as a feedback proxy to learn about advisor quality. Importantly, participants were able to learn about the value of advice even in the absence of feedback and subsequently selected more advice from better advisors, indicating the value of using confidence in this way. In contrast, when participants had prior knowledge about the reliability of advisors, they requested advice more often when confidence was low, reflecting the use of confidence as a self-monitoring tool signalling that help should be solicited. These findings indicate that people use confidence in a way that is context-dependent and directed towards achieving their current goals.

Keywords: confidence, metacognition, decision making, advice, cognitive control

Introduction

People are faced with various choices almost on a momentary basis during an ordinary day. These choices range from the brief and inconsequential, such as what socks to wear before leaving the house in the morning, to those involving some deliberation, such as what films we want to watch, to those with lifelong impact, such as what career path to pursue. Often, many different considerations and uncertainties are involved in reaching a decision, making the best course of action unclear. This uncertainty is manifested in a universal experience of a varying sense of confidence that accompanies the decisions we reach. Typically, people can evaluate their own performance reliably, reporting levels of confidence that correlate with objective accuracy (Baranski & Petrusic, 1994; Koriat & Goldsmith, 1996), and detecting errors they have recently made (Rabbitt, 1966). Researchers have suggested that this ability to evaluate the quality of our decisions plays a crucial role in adaptive behavior, and is particularly important given that many decisions we reach are not accompanied by reliable and immediate external feedback (Meyniel, Sigman, & Mainen, 2015; Yeung & Summerfield, 2012, 2014).

Several recent studies that directly assessed the effects of subjective confidence on subsequent behavior have provided evidence supporting this proposal. In perceptual decision making, confidence has been found to modulate serial dependence, so that past decisions that were made with higher confidence more strongly bias current decisions (Samaha, Switzky, & Postle, 2019), and lead to more liberal response thresholds (Desender, Boldt, Verguts, & Donner, 2019). Furthermore, confidence has been shown to guide high-level strategic choices: People employ cognitive offloading strategies, such as setting external reminders, more often when unconfident that they will succeed without external help (Boldt & Gilbert, 2019; Gilbert, 2015; Hu, Luo, & Fleming, 2019); confidence guides decisions regarding whether to choose the option which is currently believed to have the best outcome, or explore new alternatives in the hope of finding an even better option (Boldt, Blundell, & De Martino, 2019); and people choose to engage

more with tasks in which they experience higher confidence in their performance (Carlebach & Yeung, 2020). Of particular relevance to the present study, two recent studies demonstrated that the amount of information people gather before committing to a choice depends on their subjective confidence, with people seeking out more information when they are initially less confident (Desender, Boldt, & Yeung, 2018; Desender, Murphy, Boldt, Verguts, & Yeung, 2019).

Taken together, these studies provide substantial evidence that confidence plays a functional role in a variety of adaptive behaviors. Importantly, these previous studies have all revealed a consistent mapping between confidence in a decision and subsequent behavior when performing a specific task, so that low confidence leads to one course of action and high confidence to another. For example, when performing an information sampling task, low confidence leads to gathering additional information, whereas high confidence leads to committing to a choice. However, it remains unclear whether this mapping is necessarily fixed, or could exhibit some degree of flexibility and context-sensitivity. Here, we propose that the relationship between confidence and subsequent behavior is not fixed. Rather, we suggest that this relationship varies depending on the specific way in which confidence is being used, such that people use confidence in a flexible fashion according to the immediate context and goal at hand. This framing is consistent with theories that view explicit, conscious representations – such as representations of decision confidence – as serving crucially to guide behavior in a flexible manner to support evaluation, planning, and voluntary action (Dehaene & Changeux, 2011).

To investigate this proposal, we studied the relationship between confidence and advice seeking. Previous social decision-making studies have shown that people utilize advice more when a task is difficult, and they are uncertain about their ability to make an accurate decision on their own (Gino & Moore, 2007). In these situations, confidence is used as a self-monitoring tool, signaling to a decision maker that they are unlikely to succeed and should seek the help of others.

We propose that this seemingly natural relationship between low confidence and advice requests is not fixed but rather context dependent. Specifically, we predicted that when people want to learn about advice quality, they will request more advice when their confidence is high, rather than low.

The importance of taking advisor quality into account when deciding whether to request or follow advice is intuitive. Accordingly, studies have shown that people are sensitive to the quality of advice based on the reputation of the advisor, relative to their own level of knowledge (Yaniv, 2004; Yaniv & Kleinberger, 2000). However, we often do not have any previous knowledge about the reliability of information provided by other people. In such situations, when feedback is available, people can integrate the reliability of advisors' responses over time, updating their trust in advisors through associative learning processes (Behrens, Hunt, Woolrich, & Rushworth, 2008). Importantly, it has recently been demonstrated that people can learn about advisor reliability even when feedback is not available, by comparing advice to their own beliefs when they are certain they know the correct answer (Pescetelli & Yeung, 2018): If we are certain that a given decision is correct, we can equally be certain that anyone who disagrees with us is wrong, and accordingly reduce our trust and reliance on their advice now and in the future. Conversely, if we reach a decision but with low confidence, we should still reduce our trust in advisors who disagree, but to a lesser extent. According to this agreement-in-confidence heuristic, if objective feedback is unavailable, our best opportunity to learn about the reliability of an advisor comes when we are ourselves sure of what we think about a particular decision. Based on this reasoning, we predicted that when people want to learn about advisor reliability in the absence of external feedback, they will strategically sample advice when confidence is high, breaking away from the typical relationship between low confidence and advice requests.

To illustrate this prediction, consider the situation of being unsure which of two movies to watch, and therefore seeking advice from online reviews—a straightforward example of

subjective confidence guiding information seeking. Suppose you find reviews for both movies written by the same film critic, who had a strong preference for one of the two movies. However, you have never heard of this critic before, and cannot be sure you should trust their opinion. Suppose, though, that on the website you notice that they have also reviewed a movie you saw last week and loved. Now you might be motivated to read the review, even though you are confident in your opinion about this movie, so as to learn about the reviewer: If you find that the critic loved the movie too, you might infer that the critic can be trusted, and therefore follow their recommendation for tonight's film choice.

The example above demonstrates the use of confidence both as a self-monitoring tool (indicating when you need advice), and as a proxy for feedback when objective feedback is lacking (as a yardstick against which to evaluate the reliability of advice). In the present study, we tested the prediction that people can utilize confidence in such a flexible way. We expected the relationship between confidence and advice requests to vary depending on the way in which confidence is currently being utilized. Specifically, we hypothesized that frequent advice requests when confidence is low reflect the use of confidence as a self-monitoring tool, signaling that help should be requested. In contrast, advice requests when confidence is high reflect an attempt to learn about advisor quality, by using confidence as an internal feedback proxy.

We tested our predictions in three experiments using a Judge-Advisor System (JAS) paradigm. In its most basic form, the JAS consists of a decision maker (the judge), and one or more advisors who communicate their advice to the judge (Bonaccio & Dalal, 2006; Sniezek & Buckley, 1989, 1995). In the present study, participants (i.e., acting as 'judge') performed a perceptual judgement task in which they reported which of two boxes contained more dots, and how confident they are in their response. Following each initial judgement, they chose between receiving advice from a virtual advisor or seeing the stimulus again. After receipt of the additional information chosen, they reported their final decision and confidence on which box contained

more dots. In Experiment 1, we tested the novel prediction that, when advisor reliability is initially unknown and no objective feedback is provided, participants will ask for advice when they themselves are confident in a decision (and can therefore learn if advisors are providing accurate advice). Experiments 2 and 3 manipulated the availability of explicit information about advisor reliability (rendering learning about advisors unnecessary) and of external feedback on task performance (so that the use of confidence as an internal feedback proxy is redundant). Our data reveal that the relationship between confidence and advice requests is not fixed, but instead varies depending on the availability of such information. We conclude that confidence is used strategically, both as a self-monitoring tool, and as a proxy for feedback, depending on the task context and current goal.

1. Experiment 1

Experiment 1 aimed to test our main hypothesis, that people use confidence flexibly depending on their immediate goal. To test this hypothesis, we employed a perceptual decision-making task in which participants had to report which of two boxes contained more dots and how confident they are in their decision. They then chose between seeing the same stimulus again, for a longer period, or receiving advice from a virtual advisor. We predicted a U-shaped relationship between confidence and advice requests, reflecting the use of confidence for two different purposes. Specifically, we predicted that people would seek advice more frequently when confidence was very low, to improve their accuracy by using the help of others, but critically also when confidence was very high, enabling them to learn about advisor quality. Key predictions and analyses were preregistered at <https://aspredicted.org/blind.php?x=n2pj9e>.

1.1. Methods

1.1.1. Participants

Twenty-six participants took part in the experiment (17 females, ages 18-35 years, $M = 23.5$, $SD = 5.5$). All participants had normal or corrected-to-normal vision. They provided written informed consent, and were compensated with course credit or payment for their participation. One participant whose data could not be divided into quantiles (as defined below) was removed from the data analyses, as described in the results section. All procedures were approved by the University of Oxford's Medical Sciences Interdivisional Research Ethics Committee.

1.1.2. Task and Procedure

Participants performed a series of trials in which they first reported which of two boxes contained more dots and how confident they were in their response, and then chose to sample one of two additional sources of information before giving a final response and confidence judgement. **Error! Reference source not found.** depicts an example trial in the main experimental blocks. Trials began with a fixation cross displayed at the center of the screen for 800 ms in total. The fixation cross flashed briefly 400 ms after it appeared, indicating the perceptual stimulus will be presented shortly after. The stimulus, two boxes containing different numbers of dots, appeared to the left and right of the fixation cross for 150 ms. Participants reported which of the two boxes contained more dots, and how confident they are in their response on a scale ranging from 50% (guess), to 100% (sure). After the initial response was registered, images representing the advisor and the dot stimulus appeared above and below the fixation cross. Participants chose between viewing the stimulus again for a longer period (making it easier to judge which box contained more dots) versus receiving advice from a virtual advisor, by selecting the appropriate image with the mouse cursor. If participants chose to view the stimulus again, they were presented with a fixation cross for 800 ms, followed by the same dot

stimulus for an extended period of 300 ms. If they chose advice, a

Fig 1. Experimental Design. Participants made a speeded judgement about which of two squares contained more dots, and reported their confidence in this decision. Following the initial judgement, participants chose between viewing the stimulus again for an extended period, or receiving advice from a virtual advisor by selecting the corresponding image. After receiving the additional information they selected, participants reported their final decision and confidence on which box contained more dots.

fixation cross appeared for 300 ms, after which the two boxes were presented to the right and left, but without the dot stimulus. For advice selections, they then saw an arrow pointing at one of the empty boxes (i.e. the advisor’s recommended response) displayed for 1500 ms. Advice trials were longer than review stimulus trials to counterbalance a bias displayed towards advice selections in a pilot study. Following the additional information (stimulus re-viewing or advice), participants reported their final response and confidence on which box contained more dots, with their initial response marked on the confidence scale for reference. The next trial then began (with no feedback provided to indicate the box that objectively contained more dots).

The experiment consisted of six blocks of 60 trials each. Participants encountered a new advisor during each block, represented by a picture of a Caucasian female with a neutral expression selected from the NimStim database (Tottenham et al., 2009). Advisor accuracy varied across the 6 blocks, at 50%, 60%, 70%, 80%, 90%, and 100% correct, with the order of advisor accuracy randomised across blocks for each participant (and, across participants, randomly assigned to the 6 different advisor faces). Participants were instructed that advisors would be of different quality, but received no information about their accuracies. To encourage

participants to choose according to their preference, the instructions indicated that ‘It is up to you to decide on each trial if receiving advice or viewing the stimulus again for a longer duration is more useful to you’. The locations in which the advisor and view-stimulus options were presented during information choice (i.e., above vs. below the fixation) were assigned randomly for each block, with each information type appearing in each location an equal number of times throughout the experiment.

To induce a wide range of confidence judgements (that would allow us to study the relationship between different confidence levels and information choice), each block contained an equal number of trials in three difficulty levels (easy, medium, hard), with trial order determined randomly. Trial difficulty was controlled by manipulating the difference in dot numbers between the boxes, with the numbers in each box determined by the equation $nDots = 200 \pm d * l$, where d represents the number of dots added and subtracted to the correct and incorrect boxes correspondingly, and $l = [3, 1, 1/3]$, for the easy, medium and hard difficulties correspondingly. To achieve similar accuracy levels across participants, the medium difficulty condition was staircased to ~70% accuracy using a 2-down 1-up procedure titrating the d parameter value throughout the experiment. The easy and difficult trials were defined relatively to the staircased medium difficulty by multiplying or dividing d by 3, as captured by the l parameter. The side of the box with more dots varied randomly across trials, and the locations of the dots within each box were sampled randomly on every trial. To encourage task engagement, participants received a bonus payment based on final-response accuracy during the main experimental blocks.

Before beginning the main experiment, participants viewed instructions explaining the task and performed three practice blocks. The first two blocks consisted only of the dot judgement task and contained 42 trials each. During these two training blocks, participants received feedback on their performance in the form of a tone after error trials, as well as feedback on their overall

accuracy at the end of each block (based on participants' final responses). The third practice block introduced the information choice component, with participants performing 36 trials of the full task. The practice advisor had an accuracy of 80%. As in the main experiment, participants did not receive information about advisor accuracy or feedback following each trial of this block, but received feedback on their overall final response accuracy at the end of the block.

After completing the main experiment, participants completed the Concern over Mistakes, Personal Standards and Doubts about Actions subscales from the Multi-dimensional perfectionism scale (Frost, Marten, Lahart, & Rosenblate, 1990). These surveys were included in this experiment, and in Experiment 2, to allow exploratory analysis of how perfectionism might relate to confidence and information choices. However, since no clear relationships were found, we do not report any analyses on the questionnaire data.

1.1.3. Stimuli and Apparatus

All stimuli were created and presented in Matlab using the Psychtoolbox3 extension (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997). Stimuli were presented on a 20-inch CRT screen with grey background (127, 127, 127), a refresh rate of 60 Hz and 1600 x 1200 resolution. Dot stimuli comprised two white rectangular frames in which white dots were presented in random locations in a 20 x 20 grid. Advisors were represented by a picture of a Caucasian female with a neutral expression selected from the NimStim database (Tottenham et al., 2009). The experiment took place in a dimly lit room, with participants sitting at a comfortable viewing distance from the screen.

1.1.4. Statistical analyses

Confidence Quantiles. To study the relationship between subjective confidence and behavior, trials were grouped according to confidence ratings. For each participant, N quantiles were computed on initial confidence ratings (with N depending on the specific analysis), and trials

grouped into bins according to trial confidence so that $Q(n - 1) < \text{conf}(t) \leq Q(n)$. If certain confidence ratings were used frequently, resulting in bins with shared cut off values or containing fewer than 20 trials, new quantiles were computed above and below these values. For example, if $Q(n) = Q(n - 1)$, $n - 1$ quantiles were calculated on $\text{conf} \leq Q(n)$, and the remaining quantiles were calculated on $\text{conf} > Q(n)$. If the last quantiles had shared cut off values, or the cutoff value was 100 (the highest confidence rating), $N - 1$ quantiles were calculated on $\text{conf} < Q(n)$, so that the highest confidence rating was binned in the last quantile.

1.2. Results

1.2.1. Accuracy and confidence judgements

To confirm that participants were complying with task instructions, our first analysis studied the differences in decision accuracy and confidence ratings between the initial and final responses. We expected that if, as instructed, participants used the additional information provided between their responses (i.e., receive advice or view stimulus again) to improve their decision accuracy, this would lead to an increase in accuracy in the final response. Consistent with this prediction, a paired t -test revealed that mean accuracy was higher on the final response, compared to the initial response ($M = 77\%$ vs. 74% , $SD = 3\%$ vs. 2% ; $t(24) = 6.14$, $p < .001$). This increase in accuracy was mirrored in participants' confidence ratings, with participants reporting higher confidence in their final responses, compared to their initial confidence ($M = 73$ vs. 70 , $SD = 8.6$ vs. 8.7 ; $t(24) = 7.06$, $p < .001$).

Next, we wanted to confirm that participants were sensitive to the difficulty manipulation, so that the varying difficulty levels induced a range of objective accuracy and subjective confidence scores. Given that our main goal in this experiment was to study the relationship between confidence and subsequent information choices, initial accuracy and confidence were of particular interest. However, for completeness, we report the effects of trial difficulty on both the

initial and final responses. Mean accuracy and confidence ratings were computed as a function

Fig 2. Results of Experiment 1: effects of task difficulty on accuracy and confidence. (A) Mean accuracy on the initial and final responses, divided according to trial difficulty. (B) Mean confidence ratings on the initial and final responses, divided according to trial difficulty. Error bars indicate standard errors of the mean (SEM).

of trial difficulty, for the initial and final responses separately (see Fig 2), and subjected to separate repeated-measures ANOVAs. Results showed that, as expected, accuracy varied according to trial difficulty for both the initial and final responses ($F(2, 48) = 625.25, p < .001, \eta^2 = .96$; $F(2, 48) = 443.73, p < .001, \eta^2 = .95$, for the initial and final responses, correspondingly). Within-subject linear contrasts were significant ($F(1, 24) = 1129.18, p < .001, \eta^2 = .98$; $F(1, 24) = 805.85, p < .001, \eta^2 = .97$), with higher accuracy rates on easier trials, for the initial ($M = 93\%, 71\%, 59\%, SD = 4\%, 2\%, 5\%$) and final responses ($M = 95\%, 76\%, 62\%, SD = 4\%, 4\%, 5\%$ for easy, medium and hard trials correspondingly). Most importantly, a similar pattern emerged for confidence judgements. Confidence differed monotonically according to task difficulty (initial response: $F(2, 48) = 133.77, p < .001, \eta^2 = .85$; final response: $F(2, 48) = 152.02, p < .001, \eta^2 = .86$), with significant within-subject linear contrasts ($F(1,24) = 154.84, p < .001, \eta^2 = .87$; $F(1,24) = 193.46, p < .001, \eta^2 = .89$), and higher confidence ratings for easier trials on the initial ($M = 74, 68, 67, SDs = 9$) and final responses ($M = 79, 71, 69, SDs = 9$). Thus, the difficulty manipulation had the expected effects, leading to varying accuracy levels and, importantly, a

range of confidence scores, with participants performing better, and reporting higher confidence, on easier trials.

1.2.2. Confidence-Accuracy calibration

To further study the relationship between subjective confidence and objective accuracy, we grouped trials according to confidence ratings, pooling across all trial difficulties. For each participant, trials were grouped into six bins according to initial-confidence sextiles, and mean accuracy rates computed for each bin (see Fig 3A). A repeated-measures ANOVA showed that mean accuracy varied reliably as a function of confidence ($F(5,120) = 76.02, p < .001, \eta^2 = .76$), with accuracy increasing monotonically with the level of confidence, as captured by a significant linear trend ($F(1,24) = 286.99, p < .001, \eta^2 = .92; M = 61\%, 66\%, 71\%, 78\%, 83\%, 91\%$, and $SD = 8\%, 5\%, 7\%, 7\%, 5\%, 6\%$). Thus, participants' subjective reports of confidence were well-calibrated to objective accuracy. This pattern was seen across experiments but is reported here only, for the sake of brevity.

Fig 3. Results of Experiment 1: (A) Mean accuracy according to reported confidence. For each participant, trials were divided into six bins according confidence ratings. (B) Mean proportion of advice selections according to adviser accuracy. Error bars indicate standard errors of the mean (SEM).

1.2.3. *Learning about advisor quality*

Across blocks, participants encountered advisors whose accuracy varied from 50% to 100%, but no explicit trial-by-trial feedback was provided against which to evaluate the quality of this advice. Participants nevertheless learned about advisor quality when deciding whether to ask for advice: As demonstrated in Fig 3B, the proportion of advice requested throughout the experiment depended on advisor accuracy. A repeated-measures ANOVA on the proportion of advice requests as a function of advisor accuracy showed a significant effect ($F(5, 120) = 2.80, p = .02, \eta^2 = .818$), with a significant within-subject linear contrast ($F(1,24) = 6.87, p = .015, \eta^2 = .71$), indicating that advice requests increased monotonically as a function of advisor quality. These results provide further evidence that people can learn about the quality of advice even when objective feedback is lacking (Pescetelli & Yeung, 2020). As described above, our hypothesis was that this learning would be driven by participants seeking advice when they were already sure about the correct answer, so that they could evaluate advisors against these firmly held beliefs, rather than only seeking advice when low in confidence as seen in previous research.

1.2.4. *Confidence effects on information choices*

Overall, participants chose to receive advice on .39 of the trials on average, with large individual differences in preferences for advice or viewing the stimulus again (proportion advice choice: range = 0–.91, $SD = .26$). Both advice and viewing the stimulus again increased participants' accuracy on their final response compared to their initial response (from 74% to 77% and 72% to 74%, respectively), indicating that both types of information were of value ($t(23) = 3.97, p < .001$ and $t(24) = 4.23, p < .001$, for advice and view-stimulus trials, respectively; note that one participant never requested advice).

To answer our main question—whether people use confidence strategically depending on their immediate goal—we studied the relationship between confidence and information choices. For each participant, we divided trials into six bins according to initial-confidence sextiles, and

then computed the proportion of trials in which advice was selected for each bin. We predicted a U-shaped relationship between confidence and advice selections, whereby people would choose to receive advice, rather than view the stimulus again, more often both when they were initially low in confidence (reflecting an attempt to use advice in order to improve accuracy), and when they were initially high in confidence (reflecting an attempt to learn about advisor quality). As expected, a repeated-measures ANOVA revealed a significant effect of confidence on information choices ($F(5, 120) = 5.34, p = .012, \eta^2 = .18$), and a significant quadratic trend ($F(1, 24) = 5.82, p = .024, \eta^2 = .195$). Surprisingly, however, we found that people requested advice more often when confidence was high, but not low—the significant quadratic trend reflected a levelling off, rather than increase, in advice choices at the lowest levels of confidence (see Fig. 4A). Pairwise comparisons of confidence sextile bins with Bonferroni correction revealed significant differences in the frequency of advice selections between trials with the highest confidence ratings (Q6) and medium confidence trials (Q5 and Q4, $ps < .05$), but no significant differences for the lowest confidence trials (Q1). Thus, while these results support the hypothesis that people use confidence as a feedback proxy to learn about advisor reliability, they do not support our initial prediction of a U-shaped relationship between confidence and advice requests.

Since confidence varied reliably as a function of trial difficulty, an alternative interpretation of the above results is that information choices were driven by trial difficulty, rather than confidence. Indeed, because the difficulty manipulation was introduced precisely with the intention of inducing variability in confidence ratings, trials of different difficulty levels were not equally represented in confidence bins, making it difficult to separate these effects. For example, the highest confidence bin comprised on average 36, 14 and 9 trials from the easy, medium and hard conditions, respectively. Importantly, despite these differences in confidence distributions, there was also a considerable overlap in the confidence ratings between the different difficulties.

Fig. 4. Results of Experiment 1: (A) Mean proportion of advice selections according to confidence bins across the entire experiment. (B) Normalised initial confidence ratings, divided according to trial difficulty. Confidence ratings were first normalized within each participant, across all trials regardless of difficulty. Then, normalised confidence ratings were aggregated across all participants and grouped according to trial difficulty. (C) Mean proportion of advice selections according to confidence bins and trial difficulty. For each participant, trials within each difficulty were divided into four bins according to confidence ratings. (D) Mean proportion of advice selections according to confidence and time in block. For each participant, trials were divided into two time periods (first half vs. second half of block), and then into six bins according to confidence ratings, within each time period. Error bars represent SEMs.

Fig. 4B shows the distributions of initial confidence ratings for each difficulty level plotted across all participants, after normalizing confidence within each participant (across all trials regardless of difficulty). The visible overlap in confidence ratings between the different difficulty levels, renders an alternative explanation, whereby behavior is driven by objective difficulty rather than confidence, unlikely.

To test directly whether confidence affects information seeking above and beyond trial difficulty, we performed a two-way repeated-measures ANOVA on the proportion of advice requests, with factors of trial difficulty and confidence. For each participant, trials were divided first according to task difficulty, and then based on confidence ratings. Confidence quartiles were computed for each participant within each difficulty, and trials divided into bins according to confidence (here we used four bins, rather than six, due to the smaller number of trials available within each difficulty level). We found that both trial difficulty ($F(2, 48) = 5.96, p = .019, \eta_p^2 = .20$) and confidence ($F(3, 72) = 4.11, p = .041, \eta_p^2 = .15$) reliably predicted information choices, and the interaction between them was not significant ($F(6, 144) = 1.25, p = 0.28, \eta_p^2 = .05$). As shown in Fig. 4C, participants requested advice more frequently on easy trials, compared to medium and hard trials ($ps = .02$). In addition, participants asked for advice more often when they were highly confident ($p = .049, .044, .010$ for Q4 compared to Q1, Q2 and Q3 correspondingly). Notably, an increase in advice requests for the highest confidence bin is visible within all three difficulty conditions. These results indicate that confidence influences information choices above and beyond trial difficulty.

1.2.5. Evidence for use of confidence for learning

We predicted this high rate of advice requests at the highest levels of confidence based on the hypothesis that participants were using their own judgements as a yardstick against which to judge (and learn about) advisor quality. As a more direct test of this idea, we examined how the relationship between confidence and information choices varied during blocks, as participants presumably learned about the quality of advice offered by the current advisor. Based on our learning hypothesis, we expected more advice selections on high-confidence trials at the beginning of each block, when participants did not have any information about the current advisor's accuracy. As the block progressed, we expected less advice choices when confidence

was high, since participants should have gained knowledge about advisor quality and would no longer needed to use high confidence trials to learn about advice.

Consistent with these predictions, a two-way repeated-measures ANOVA on the proportion of advice choices, with time in block (first half, second half) and confidence (Q1-Q6) as factors, revealed significant main effects for time ($F(1, 24) = 19.95, p < .001, \eta_p^2 = .45$) and confidence ($F(5, 120) = 5.02, p = .016, \eta_p^2 = .17$), and, importantly, a reliable interaction effect between these factors ($F(5, 120) = 3.95, p = .01, \eta_p^2 = .14$): As shown in Fig. 4D, advice choices varied more strongly as a function of confidence in the first half of blocks than the second. Follow-up one-way ANOVAs on advice choices for each block half separately revealed that confidence reliably predicted advice requests during the first time period ($F(5, 120) = 7.04, p = .003, \eta^2 = .23$) but not during the second time period ($F(5, 24) = 1.93, p = .15, \eta^2 = .07$).

1.3. Discussion

The results of Experiment 1 provided some support of our main hypothesis, but also produced some surprising results. As predicted, we found that participants were able to learn about the accuracy of their advisors even without objective feedback, and that they sought advice frequently when they were most confident in their own initial decisions, particularly early in each block as they were learning about the quality of advice on offer. These findings are consistent with our hypothesis that confidence is used strategically depending on peoples' immediate goal: Advice-seeking is not solely associated with situations of high uncertainty, as explored in previous research, but can also be seen in situations of high confidence, when an important goal is to learn about the current task context. Surprisingly, however, we did not find that participants were more likely to request advice when they were unconfident in their initial response, such that we did not observe the U-shaped relationship we had originally predicted between confidence and advice seeking: Instead, participants' preferences for receiving advice vs. re-viewing the stimulus were relatively stable from medium to low levels of confidence. The explanation for this

pattern is unclear, and as we will see below, the effect varied across experiments, with some evidence that participants are more likely to request advice when confidence is low in Experiments 2 and 3.

Experiment 2 built on these findings to compare conditions in which participants received prior information about advice quality, rendering learning unnecessary, to conditions in which they did not receive prior information, and thus had to learn about the quality of advice. We predicted that participants would show evidence of using confidence in a strategic fashion, such that the relationship between confidence and advice-seeking would vary across these conditions.

2. Experiment 2

Experiment 2 tested whether confidence is used flexibly depending on the availability of external information, which we varied in a between-participants design. Using a similar task to Experiment 1, we manipulated the availability of explicit prior information about advice quality: In the "Explicit information available" condition, participants performed the same task as in Experiment 1 but, importantly, were notified of the current advisor's average accuracy at the start of each block. The "Explicit information absent" condition was a replication of Experiment 1, so that explicit information about advisor accuracy was not provided (and instead advisor accuracy had to be learnt). We expected the relationship between confidence and information choices to change depending on information availability. Specifically, we predicted that people will use trials on which they are highly confident to learn about advisor quality only when explicit information about advice was absent. In contrast, when explicit information about advisor quality was available, and learning about advisors unnecessary, we did not expect to see a higher frequency of advice selections when confidence was high. We did not have a specific prediction regarding the relationship between confidence and advice seeking in this case: It could be that people might request more advice when confidence is low in order to improve their immediate accuracy, but it seemed equally plausible that they might not use confidence to guide advice

requests in this condition, and that advice requests would depend only on advisor quality, which was known. Key predictions and analyses were preregistered at <http://aspredicted.org/blind.php?x=ay2ng3>.

2.1. Methods

2.1.1. Participants

Fifty-one participants took part in Experiment 2. One participant whose accuracy on the initial response was more than two standard deviations below the mean accuracy was removed, resulting in a final sample of 50, with 25 participants in each condition (35 females, ages 18-33 years, $M = 21.7$, $SD = 4.2$). All participants had normal or corrected-to-normal vision. They provided written informed consent, and were compensated with course credit or payment for their participation. All procedures were approved by the University of Oxford's Medical Sciences Interdivisional Research Ethics Committee.

2.1.2. Task and Procedure

Experiment 2 used a mixed design, with explicit information about advisor accuracy (absent, available) varying between participants. For each participant, advisor accuracy varied randomly between blocks across 5 levels, from 50% - 90%, and within each block participants performed the perceptual task across three levels of task difficulty. The perceptual task and information choice were identical to Experiment 1. Thus, the 25 participants assigned to the information absent condition provided a replication of Experiment 1. The remaining 25 participants were assigned to the information available condition in which, at the beginning of every block, a message on the screen informed them of the advisor's accuracy in the current block (e.g., "Advisor accuracy on this block: 70%"). This experiment did not include a block with 100% accurate advice, to avoid participants in the information available condition always requesting advice from the 100% accurate advisor. The end of block feedback differed slightly from

Experiment 1, with participants receiving feedback on their average initial-response accuracy rate, in addition to average final-response accuracy.

2.1.3. Stimuli and Apparatus

The stimuli and apparatus were identical to Experiment 1, except for advisor stimuli: In Experiment 2, advisors were represented by a picture of a smiling Caucasian female selected from the Chicago Face database (Ma, Correll, & Wittenbrink, 2015).

2.2. Results

2.2.1. Accuracy and confidence judgements

Our first analysis assessed whether participants in both experimental conditions were performing the first and second order tasks in compliance with instructions. We performed a two-way mixed measures ANOVA on mean accuracy, with factors for explicit information (absent, available) and response (initial, final). The analysis showed a reliable effect for response ($F(1, 48) = 108.27, p < .001, \eta_p^2 = 0.69$), as well as a reliable interaction between response and external information ($F(1, 48) = 6.38, p < .02, \eta_p^2 = 0.12$). Follow-up tests showed that, as expected, mean accuracy was higher on the final response compared to the initial response, for participants in both experimental conditions. The difference was larger for participants who received information about advisor accuracy, likely because they could better use advice to improve their accuracy ($M = 78\%$ vs. 75% , $SD = 2\%$ vs. 3% , $t(24) = 5.90, p < .001$ and $M = 79\%$ vs. 74% , $SDs = 3\%$, $t(24) = 8.68, p < .001$, for participants without explicit information, and with information, respectively). A similar ANOVA on confidence revealed a significant effect for response ($F(1, 48) = 51.08, p < .001, \eta_p^2 = .52$), but not for external information or the interaction between response and external information ($Fs < 1$). Participants in both experimental groups reported higher confidence in their final responses compared to the initial responses ($M = 73$ vs. 71 , and $M = 74$ vs. 71), with this difference slightly larger numerically for the group with explicit

Fig. 5. Results of Experiment 2: effects of response (initial vs. final) and task difficulty on accuracy and confidence according to the experimental condition. Participants who did not receive explicit information about advisor accuracies are depicted in blue (A, B), and participants who did receive explicit information about advisors are depicted in orange (C, D). (A, C) Mean accuracy on the initial and final responses divided according to trial difficulty. (B, D) Mean confidence ratings on the initial and final responses, divided according to trial difficulty. Error bars indicate standard errors of the mean (SEM).

information. Thus, participants used additional information to improve their accuracy, and reliably tracked this improvement in their confidence reports. Also replicating the results of Experiment 1, mean accuracy in both experimental conditions varied reliably as a function of task difficulty, for both the initial ($F(2, 48) = 424.41, p < .001, \eta_p^2 = .946, F(2, 48) = 542.34, p < .001, \eta_p^2 = .96$) and the final responses ($F(2, 48) = 463.54, p < .001, \eta_p^2 = .95, F(2, 48) = 240.79, p < .001, \eta_p^2 = .91$), for participants without external information, and with information, respectively. As shown in Figure 5, similar patterns were found for confidence, but are not reported for the sake of brevity.

Fig. 6. Results of Experiment 2: (A) Mean proportion of advice selections according to advisor accuracy and explicit information availability. Participants who did not receive explicit information about advisor accuracies are depicted in blue, and participants who did receive explicit information about advisors are depicted in orange. Error bars indicate SEMs. (B) Ten trial running average showing the proportion of advice requests according to advisor accuracy, across blocks and participants when explicit information about advisors was not available, and (C) when explicit information about advisor accuracy was available.

2.2.2. Effects of advisor quality

A two-way mixed measures ANOVA on the proportion of advice requests with factors for advisor accuracy and external information availability showed reliable effects for advisor accuracy ($F(4, 192) = 42.46, p < .001, \eta_p^2 = .47$) and for the interaction effect between advisor accuracy and information availability ($F(4, 192) = 12.46, p < .001, \eta_p^2 = .21$). Critically, participants in both groups showed sensitivity to advice quality in their choices to receive advice vs. re-view the stimulus. Thus, replicating the results of Experiment 1, the proportion of advice requests was reliably predicted by advisor accuracy even when external information was not available ($F(4, 96) = 8.38, p < .001, \eta^2 = .26$; see Fig. 6A), with a reliable linear contrast ($F(1, 24) = 21.73, p < .001, \eta^2 = .475$). Not surprisingly, we found a stronger effect of advisor accuracy on the proportion of advice selections for participants who were told about advisor accuracy before each block ($F(4, 96) = 36.52, p < .001, \eta^2 = .60$), again with a significant linear trend ($F(1, 24) = 69.73, p < .001, \eta^2 = .74$).

Fig. 6B and 6C show a ten-trial running average of the proportion of advice requests according to advisor accuracy and information condition, and demonstrate the differences in learning according to information availability. When explicit information was not available, participants learned about advisor accuracy throughout the block, asking for less advice from low quality advisors as the block progressed. In contrast, when explicit information about advisor accuracy was available, the proportion of advice requests varied as a function of advisor quality stably throughout the block. Thus, patterns of advisor choice reflected advisor quality for both groups of participants, but only the group without external information about advisor quality showed evidence of learning. Of interest, then, was whether confidence influenced advisor choice differently across groups, reflecting the different needs and goals of the two groups of participants.

2.2.3. Effects of confidence on information choices depend on explicit information availability

Our results from Experiment 1 showed that participants selected more advice when confidence was high, presumably in order to learn about advice quality. In the current experiment, we predicted that participants would use confidence differently depending on the availability of information about advisor accuracy. Specifically, we predicted that when no explicit information about advice quality was available (i.e., replicating the conditions of Experiment 1), participants would again request more advice when confidence was high, allowing them to learn about the advisors. In contrast, we did not expect to see this pattern for participants in the explicit information condition, since learning about advisor accuracy was unnecessary.

To test these predictions, each participant's trials were divided into six bins according to confidence, and the proportion of advice requests computed. We performed a two-way mixed ANOVA on the proportion of advice requests (which again varied substantially across participants: range = 0-.88) with factors for explicit information (available, absent) and

Fig. 7. Results of Experiment 2. (A) Mean proportion of advice choices according to explicit information availability and confidence. Participants who did not receive explicit information about advisor accuracies are depicted in blue, and participants who did receive explicit information about advisors are depicted in orange. (B, C) Mean proportion of advice choices according to information availability, confidence and time in block. For each participant, trials were divided into two time periods (first half vs. second half of block), and then into six bins according to confidence ratings, within each time period. (B) Participants who did not receive explicit information about advisor accuracy. (C) Participants who did receive explicit information about advisor accuracy. Error bars represent SEMs.

confidence (Q1-Q6, see Figure 7A). As predicted, this analysis revealed a significant interaction between confidence and information availability ($F(5, 240) = 6.73, p < .001, \eta_p^2 = .12$), as well as a significant main effect for confidence ($F(5, 240) = 3.54, p = .02, \eta_p^2 = .07$), but not for information availability ($F < 1$). Given the reliable interaction effect, indicating that confidence influences information sampling differently depending on the availability of external information about advice, we performed separate follow-up repeated measure ANOVAs, to study the relationship between confidence and information sampling within each condition.

We found that when explicit information was not available, confidence predicted information choices reliably ($F(5, 120) = 7.51, p < .001, \eta^2 = .24$), with a reliable quadratic trend found in the data ($F(1, 24) = 22.57, p < .001, \eta^2 = .49$). Importantly, consistent with the results of Experiment 1, the proportion of advice requests was highest for the highest confidence bin (Q6), with Bonferroni-corrected pairwise comparisons revealing reliable differences in the frequency of advice requests between high confidence trials (Q6) and medium confidence trials

(Q2, Q3, Q4 and Q5, $ps < .05$). Interestingly, the data revealed a numerical increase in advice requests for lowest confidence (Q1) compared to Q2 ($t(24) = 2.75$), but this difference was not reliable when correcting for multiple comparisons ($p = .17$). For participants who were provided with information about advisor quality before each block, confidence reliably predicted advice requests ($F(5, 120) = 3.05, p = .046, \eta^2 = .11$), with a linear trend approaching significance ($F(1, 24) = 4.17, p = .052$). Crucially, for these participants the proportion of advice requests decreased as a function of initial decision confidence, with no evidence of increased advice seeking at the highest levels of initial confidence (though note that pairwise comparisons did not reveal any reliable differences after Bonferroni corrections). Additional analyses ensuring that advice requests were driven by confidence, above and beyond trial difficulty, can be found in the Supplemental Information.

2.2.4. Effects of time in block depend on need to learn

In Experiment 1, we demonstrated that the effects of confidence on information sampling depend on time, with people selecting advice when confidence is high more often only during the first half of each block, but not the second. We suggested that these differences are a result of learning – once participants have learned how good advisors are, they do not request advice when confidence is high. If this is the case, we should not see a change in the relationship between confidence and advice as a function of time under conditions where learning is redundant. To test this prediction, for each participant trials were divided according to time in the block and confidence ratings. We performed two-way repeated-measures ANOVAs on the proportion of advice requests, as a function of time in block (first vs. second half), and confidence bin (Q1-Q6), for each condition separately.

Replicating our results from Experiment 1, we found that when external information about advisors' accuracy was not available, advice requests were reliably predicted by confidence ($F(5, 120) = 7.49, p < .001, \eta_p^2 = .24$), and time ($F(1, 24) = 6.85, p = .015, \eta_p^2 = .22$). Importantly,

the interaction term between time and confidence was also significant ($F(5, 120) = 3.72, p = .014, \eta_p^2 = .13$). Confidence reliably predicted advice requests during the first time period ($F(5, 120) = 8.64, p < .001, \eta^2 = .26$), but not the second ($F(5, 120) = 2.63, p = .061, \eta^2 = 0.10$). When external information was available, we found that advice requests were reliably predicted by confidence ($F(5, 120) = 3.29, p = .04, \eta_p^2 = .12$), but, crucially, there were no significant effects of time ($F(1, 24) = 1.95, p = .5, \eta_p^2 = .08$), or the interaction between time and confidence ($F < 1$, see Fig. 7B and 7C).

2.3. Discussion

Experiment 2 investigated whether, when choosing between advice and re-sampling the evidence, people use their confidence differently depending on the availability of explicit information about advisor reliability. Replicating the results from Experiment 1, we found that when explicit information about advisor quality was not available, participants requested advice more frequently when confidence was high. In contrast, when external information about advisor quality was available, participants tended to request advice more often when confidence was low. Our results further demonstrated that the relationship between confidence and information choices varied across blocks only when explicit information about advisors was not available. Importantly, participants requested advice when confidence was high only during the first half of each block. These results support our interpretation that people use their confidence as an internal feedback cue, allowing them to learn how good advice is by comparing the advice they received with their own beliefs when they are highly confident. When participants do not want to learn about advice quality, either because explicit information is available or because they have already learned earlier in the block, they do not use confidence as a feedback proxy for evaluating advisors.

Taken together, Experiments 1 and 2 provide evidence for our two key hypotheses: that confidence guides adaptive information seeking in a flexible manner, and that confidence can

function as an internal feedback cue, used to evaluate the quality of advice when explicit information about advisor accuracy is not available. In Experiment 3, we aimed to provide further support for these claims by investigating how people use confidence when explicit information about advisor accuracy is not available but external feedback is provided. If, as we hypothesize, confidence was previously used as an internal feedback proxy, we would expect that its effects would change when external feedback, that renders this proxy redundant, is provided.

3. Experiment 3

Experiment 3 extended our investigation of the flexible use of confidence to conditions in which participants had to learn about advice quality, but did not need their confidence to achieve this learning. Participants performed a perceptual decision task, similar to that used in Experiments 1 and 2. To evaluate the influence of external feedback on the use of confidence, half of the participants received trial-by-trial feedback, while the other half of the participants did not. We expected that only participants who did not receive external feedback would request advice more frequently on high confidence trials, using confidence as a feedback proxy in order to learn about advisor quality. We did not expect this relationship between confidence and advice requests when participants received external feedback, since feedback would provide information about advice accuracy, regardless of subjective confidence on a given trial. Key predictions and analyses were preregistered at <http://aspredicted.org/blind.php?x=8um4u6>.

3.1. Methods

3.1.1. Participants

Eighty-eight participants were tested in order to achieve a final sample of 40 participants in each experimental group (32 females, ages 19-42 years, $M = 27.9$, $SD = 5.9$). A larger sample was recruited compared to Experiments 1 and 2, because we expected online effects might be weaker. Four participants were excluded because their data could not be divided into appropriate

confidence bins, and four additional participants excluded because their accuracy was more than 2 standard deviations below the average accuracy rate. Participants were recruited via Prolific Academic (<https://www.prolific.co>) and were compensated for participation. They received online instructions and agreed to take part before beginning the experiment. All procedures were approved by the Oxford University Medical Sciences Interdivisional Research Ethics Committee.

3.1.2. Stimuli and Apparatus

Experiment 3 was an online study, run on Prolific. All stimuli were created and presented using Javascript and CSS. Stimulus sizes were defined in pixels, and so varied according to the specific screen used by each participant.

3.1.3. Task and Procedure

Participants performed a perceptual judgement task, similar to the task used for Experiments 1 and 2. Several changes, described below, were made to the procedure to accommodate the change to an online platform, but which did not alter the substance of the task. Each trial began with a black fixation cross at the center of a grey screen, with two white framed boxes appearing to the left and right of the fixation. After 500 ms the fixation flashed, indicating the dot stimuli will appear briefly. The dots appeared after an additional 300 ms, and remained on the screen for 150 ms. Participants had to report which box contained more dots and how confident they are in their decision, in a single response given on a scale ranging from 100% left to 100% right. If participants selected 50%, the message "please choose one side or the other" appeared in red under the confidence scale, and they had to update their response and select one side before they could proceed. After recording their initial response, participants chose between viewing the stimulus again or receiving advice from a virtual advisor. Two boxes representing the information types appeared above and below the center of the screen, and participants selected which information to receive by pressing the corresponding box with the left mouse button. If

participants chose to receive advice, the box representing the stimulus choice disappeared, and a text box emerged from the advisor box, pointing in the direction of the advisor's response, with the text "I think it was on the RIGHT / LEFT" appearing within. After 1000 ms the confidence scale appeared below the advice, with the advice remaining on the screen until participants submitted their final response. If participants chose to view the stimulus again, they were presented once more with a fixation for 800 ms (flashing after 500 ms), followed by the same dot stimulus for an extended period of 300 ms. The confidence scale appeared immediately after the dot stimulus disappeared.

Participants were divided into the two experimental groups randomly (feedback available vs. feedback absent). For both groups, after participants submitted their final response, the box frames appeared on the screen again for 300 ms. Participants in the feedback available group received feedback after every trial, in the form of the dot stimulus appearing again in the correct box for 300 ms (the incorrect box remained empty). Participants in the feedback absent group viewed the empty boxes, without the dot stimulus, for the same duration. The next trial began after a 300 ms inter-trial interval. At the end of every block, all participants received feedback on their average accuracy on the block, and could take a break. Each participant performed five blocks of 60 trials, with advisor accuracy ranged from 60% - 100% accurate, randomly ordered across blocks. As in the previous experiments, trials were in three varying levels of difficulty, with the medium difficulty staircased using a 2-down 1-up procedure.

Before beginning the main experimental blocks, participants received online instructions and training. During the training, participants performed two blocks of 54 trials consisting of only the dot judgement task, with trials ending after the initial response (skipping the information choice and final response sections). All participants received feedback during these blocks, so they could learn to perform the task properly. After two training blocks, participants viewed additional instructions explaining the entire task and introducing the information choice. Then,

participants performed a 12-trial practice block of the entire task – judging which box had more dots, choosing between viewing the information again or receiving advice, and reporting a final response. Feedback in this block was contingent on the experimental condition of each participant. At the end of this block, participants were presented with a brief reminder of how to perform the task, and asked to make sure they are in a quiet and comfortable environment. To encourage task engagement, they were informed they could earn additional bonus payment depending on their accuracy levels.

3.2. Results

3.2.1. Accuracy and confidence judgements

As in the previous experiments, accuracy improved between the initial and final decision, for participants in both experimental conditions ($M = 81\%$ vs. 76% , $SD = 4\%$ vs. 3% , for participants in the feedback absent condition, and $M = 83\%$ vs. 76% , $SD = 4\%$ vs. 3% , for the feedback available condition). A two-way mixed measures ANOVA on mean accuracy, with factors for feedback availability (absent, available) and response (initial, final) showed a reliable effect for response ($F(1, 78) = 259.94$, $p < .001$, $\eta^2 = 0.77$), but not for feedback availability ($F < 1$) or the interaction between them ($F(1, 78) = 1.57$, $p = .21$, $\eta^2 = 0.02$). The increase in accuracy was reflected in confidence judgements, with participants reporting higher confidence in their final judgements ($M = 78$ vs. 72 , $SD = 9$ vs. 10 for the feedback absent condition, and $M = 81$ vs. 74 , $SDs = 10$, for the feedback available condition). A similar ANOVA on confidence revealed a reliable effect for response ($F(1, 78) = 140.01$, $p < .001$, $\eta^2 = 0.77$), but not for feedback availability ($F(1,78) = 1.07$, $p = .30$, $\eta^2 = 0$) or the interaction term ($F < 1$). Thus, participants used the additional information provided to improve their accuracy, and updated their confidence accordingly.

Fig. 8. Results of Experiment 3: effects of task difficulty on accuracy and confidence according to the experimental condition. Participants who did not receive trial-by-trial feedback are depicted in blue (A-B), and participants who did receive trial-by-trial feedback are depicted in orange (C-D). (A,C) Mean accuracy on the initial and final responses divided according to trial difficulty. (B,D) Mean confidence ratings on the initial and final responses, divided according to trial difficulty. Error bars indicate standard errors of the mean (SEM).

As shown in Fig. 8, the difficulty manipulation once again had the expected effects on accuracy and confidence, in both experimental conditions. Repeated-measures ANOVAs revealed that difficulty reliably predicted accuracy when feedback was not provided, for the initial ($F(2, 78) = 796.82, p < .001, \eta^2 = .95$) and final responses ($F(2, 78) = 570.44, p < .001, \eta^2 = .94$). Accuracy increased monotonically the easier the task, for both the initial and final responses, as demonstrated by significant linear contrasts ($F(1,39) = 1471, p < .001, \eta^2 = .97$ and $F(1,39) = 974.99, p < .001, \eta^2 = .96$, for the initial and final responses respectively). Similar effects were found when feedback was provided, with difficulty reliably predicting accuracy, ($F(2, 78) = 906.75, p < .001, \eta^2 = .96$, and $F(2, 78) = 335.39, p < .001, \eta^2 = .90$), and accuracy

increasing monotonically on easier trials ($F(1, 39) = 1508, p < .001, \eta^2 = .97$, and $F(1, 39) = 514.43, p < .001, \eta^2 = .93$), for both initial and final responses, correspondingly.

Importantly, the difficulty manipulation also led to differences in confidence judgements. In the feedback absent condition, confidence varied reliably as a function of task difficulty, for the initial ($F(2, 78) = 88.52, p < .001, \eta^2 = .69$) and final responses ($F(2, 78) = 99.20, p < .001, \eta^2 = 0.72$), with significant linear trends ($F(1, 39) = 92.02, p < .001, \eta^2 = .70$, and $F(1, 39) = 105.06, p < .001, \eta^2 = .73$, for the initial and final responses correspondingly). Similarly, when feedback was provided, difficulty predicted confidence reliably on the initial ($F(2, 78) = 100.85, p < .001, \eta^2 = .72$), and final responses ($F(2, 78) = 119.02, p < .001, \eta^2 = .75$), with significant linear contrasts showing that confidence increased on easier trials, for both the initial ($F(1, 39) = 101.99, p < .001, \eta^2 = .72$), and final responses ($F(1, 39) = 125.54, p < .001, \eta^2 = .76$). These results demonstrate that online participants performed the first and second order tasks in compliance with instructions, and that the difficulty manipulation led to a range of confidence judgements as intended.

3.2.2. *Effects of advisor quality*

Participants in both groups showed sensitivity to advisor accuracy in terms of their proportion of advice requests (see Fig. 9A). A two-way mixed measures ANOVA on the proportion of advice requests with factors for advisor accuracy and feedback availability showed reliable effects for advisor ($F(4, 312) = 10.08, p < .001, \eta_p^2 = .11$) and for the interaction effect between advisor accuracy and feedback availability ($F(4, 312) = 4.93, p < .002, \eta_p^2 = .06$). Follow-up ANOVAs were performed to examine the influence of advisor within each group separately. For participants in the no-feedback condition, the influence of advisor accuracy on advice requests approached significance ($F(4, 156) = 2.47, p = .058, \eta_p^2 = .06$) with a significant linear trend ($F(1, 39) = 5.77, p = .02, \eta_p^2 = .13$), indicating that participants requested more advice from

Fig. 9. Results of Experiment 3. Average proportion of advice selections according to: (A) Feedback availability and advisor accuracy. (B) Feedback availability and confidence bin. (C, D) Feedback availability, confidence and time in block. For each participant, trials were divided into two time periods (first half vs. second half of block), and then into six bins according to confidence ratings, within each time period. (C) Participants who did not receive trial-by-trial feedback, and (D) Participants who did receive trial-by-trial feedback. Error bars indicate standard errors of the mean (SEM).

better advisors, even when feedback was not available. When feedback was provided, the proportion of advice requests was reliably predicted by advisor quality ($F(4, 156) = 11.25, p < .001, \eta_p^2 = .22$). Participants requested more advice from better advisors, as demonstrated by a reliable linear contrast ($F(1, 39) = 31.21, p < .001, \eta_p^2 = .44$).

3.2.3. *Effects of confidence on information choices depend on feedback availability*

As in the previous experiments, participants displayed large individual differences in their preferences for selecting advice or viewing the stimulus again. On average, participants chose to receive advice on .30 of the trials when feedback was not available (range = 0-.82, $SD = .25$), and .37 of trials when feedback was provided (range = 0-.88, $SD = .28$). Figure 9B shows the relationship between confidence and advice requests within each feedback condition. As predicted, a two-way mixed ANOVA on the proportion of advice selections with factors for feedback (absent, available) and confidence (Q1-Q6) revealed a significant interaction between these two factors ($F(5, 390) = 3.36, p = .036, \eta_p^2 = .04$), demonstrating that confidence influenced advice-seeking differently according to the availability of feedback. No significant main effects were found for feedback ($F(1, 78) = 1.56, p = .22, \eta_p^2 = 0$), or confidence ($F(5, 390) = 2.20, p = .11, \eta_p^2 = .03$).

To study the relationship between confidence and advice requests within each feedback condition, follow-up one-way ANOVAs were performed for each group separately. For participants who did not receive trial-by-trial feedback, we found that confidence reliably predicted advice requests ($F(5, 195) = 3.48, p = .035, \eta^2 = .08$). As shown in Fig. 9B, participants requested more advice when confidence was lowest (Q1), and highest (Q6), compared to medium confidence trials. This was captured by a significant quadratic trend ($F(1, 39) = 19.03, p < .001, \eta^2 = .33$), but pairwise comparisons did not show significant results after Bonferroni corrections. Importantly, for participants receiving explicit trial-by-trial feedback, confidence did not reliably predict advice requests ($F(5, 195) = 1.82, p = .5, \eta^2 = .04$).

3.2.4. *Confidence and advice-seeking over time*

Next, we examined whether confidence was used differently, depending on how much experience participants had with an advisor. Though we did not find a significant effect of confidence on information choices when feedback was available, we report the analysis for both

feedback conditions, for completeness. Within each experimental condition, trials were divided according to time in the block (first half vs. second half), and then according to confidence bins (Q1-Q6), as shown in Fig. 9C and 9D. Two-way repeated measures ANOVAs on the proportion of advice requests, as a function of time-in-block and confidence bin were performed for each condition separately. When feedback was available, we did not find any reliable effects for time-in-block ($F(1, 39) = 1.46, p = .23, \eta_p^2 = .04$), confidence ($F(5, 195) = 1.91, p = .09, \eta_p^2 = .05$), or their interaction between these factors ($F < 1$). More surprisingly, results reveal that the interaction between time and confidence was not significant when feedback was absent ($F < 1$). The effects of time ($F(1,39) = 1.26, p = .27, \eta_p^2 = .03$) and confidence were also not reliable, but confidence approached significance ($F(5, 195) = 2.85, p = .06, \eta_p^2 = .07$). As can be seen in Figure 9C, there is a trend towards participants requesting more advice on high-confidence trials, regardless of the time spent in the block. These results are consistent with the weak learning effect found in this experiment, suggesting that participants continued to use confidence to try and learn about advice quality throughout the entire block.

3.3. Discussion

Experiment 3 evaluated whether the availability of feedback modulates how people use confidence when deciding whether to request advice or re-sample stimulus information. Our aim was to extend our findings from Experiment 2 to conditions in which participants had to learn about advisor accuracy (i.e., explicit information about advice quality was not available), but did not have to rely on their confidence for learning. The results demonstrate that participants used confidence to guide information choices only when external feedback was unavailable. These results lend further support to our suggestion that, when external information is not available, confidence is used as a feedback proxy, allowing assessment of advisor accuracy.

The present results revealed weaker learning effects when external feedback was unavailable, compared with Experiments 1 and 2. Perhaps correspondingly, we also found that

in this experiment participants' use of confidence was not reliably modulated by time. Instead, they continued to request advice more often when confidence was high throughout the entire block. This discrepancy might be attributed to Experiment 3 being performed on-line, such that it drew on a different (perhaps less engaged) participant sample than the in-lab studies in Experiments 1 and 2.

4. *General Discussion*

Confidence reflects a subjective estimate of the likelihood that a decision reached was a good one, and as such can play an important role in guiding adaptive behavior. Recent studies have demonstrated the involvement of confidence in a variety of behaviors, such as learning (Guggenmos, Wilbertz, Hebart, & Sterzer, 2016; Hainguerlot, Vergnaud, & de Gardelle, 2018) and information seeking (Desender et al., 2018; Desender, Murphy, et al., 2019). However, these previous studies have typically revealed consistent mappings between confidence and behavior, whereby low confidence leads to one course of action, and high confidence to another. In the present study, we tested the hypothesis that the relationship between confidence and behavior will vary strategically according to the specific goal and context, based on the notion that explicit representations of confidence and uncertainty serve crucially to guide behavior in a flexible, goal-directed manner (Dehaene & Changeux, 2011; Nelson & Narens, 1990).

This hypothesis was tested in three experiments that examined the relationship between confidence and advice seeking. In our JAS paradigm, participants had to choose between receiving advice and re-sampling the evidence after each initial decision. We found that, in most cases, participants' information choices were predicted by their previous confidence judgements, and, critically, that this relationship varied depending on context. Specifically, the relationship between advice seeking and confidence depended on the availability of information about advisor reliability. When participants did not have any external information about advisor accuracy, they selected advice more often when confidence was high. In contrast, when information about

advisor quality was available, participants tended to request advice most often when confidence was low. These patterns of behavior appeared in our data consistently, regardless of the type of information available about advisor quality: whether it was provided explicitly, learnable via external feedback, or learnt through exposure earlier in the block. Thus, whereas low-confidence advice requests reflect the use of confidence as a self-monitoring system—with low confidence acting as a cue that participants are unlikely to succeed in the task—high-confidence advice requests reflect the use of confidence as a feedback proxy, allowing participants to learn about advisor reliability by comparing advice to the expected answer when they are certain they know the correct answer.

Supporting the suggestion that confidence is used as a feedback proxy, we found evidence in all three experiments demonstrating that participants were able to learn about the quality of advice, and requested more advice from better advisors, even in the absence of external feedback. These results are in line with a recent study, which showed that confidence can be used to learn about the reliability of advisors when feedback is not available (Pescetelli & Yeung, 2020). Together, these results extend findings from previous studies, which have demonstrated that confidence can be used as a learning signal during perceptual tasks (Guggenmos et al., 2016), to show that confidence can also be valuable as a feedback signal when learning about social information.

The present findings complement previous research into advice taking using the Judge Advisor System. Like many other actions, the decision whether to seek advice can be viewed in terms of the costs and benefits associated with each line of action (Van Swol, Paik, & Prah, 2018). Often, soliciting the help of external advisors requires the investment of time and money, particularly when advice is given by professionals or experts. Typically, the main benefit of advice is to improve performance and gains (for example when requesting advice from a banker about investments). Interestingly, there have been suggestions that advice might provide benefits

in addition to improved performance, such as joint responsibility (El Zein, Bahrami, & Hertwig, 2019; Harvey & Fischer, 1997). Here we demonstrate that advice seeking can also provide knowledge about advisor reliability. By requesting advice when cost is low (e.g. when we already know the answer), people can learn about the reliability and accuracy of others, and build up their trust in an advisor. This information can be important when deciding whether to invest in advice in the future.

Our finding that confidence is used to learn about the reliability of others resonates with the conclusions of recent studies investigating the formation of global estimates of self-performance (SPEs; Boldt, Schiffer, Waszak, & Yeung, 2019; Rouault, Dayan, & Fleming, 2019). In these studies, the researchers demonstrated that local confidence (confidence in a single decision) is aggregated and used for building a global self-performance estimate over time. These estimates affect neural preparation for stimulus processing (Boldt, Schiffer, et al., 2019), and, importantly, allow people to decide which tasks to pursue next, based on their beliefs of how likely they are to succeed (Carlebach & Yeung, 2020; Rouault et al., 2019). Here, we demonstrate that local confidence is used not only to form a global estimate of self-performance, but also global estimates about the performance of others. This ability to monitor and evaluate one's own performance, and compare it to the performance of others, is particularly important in social contexts, when decisions have to be made as a group, or when deciding whether to seek or use other peoples' advice. Thus, our results suggest that confidence acts as a signal that allows not only the comparison and prioritisation of different tasks (Carlebach & Yeung, 2020; de Gardelle, Le Corre, & Mamassian, 2016; de Gardelle & Mamassian, 2014), but also the comparison of different sources of information. Future studies might extend these findings to consider different types and sources of information beyond advisors.

Previous studies exploring advice taking have demonstrated that people use advice more often when a task is difficult (Gino & Moore, 2007), and they are unconfident (Gino, Brooks, &

Schweitzer, 2012). Based on these results, we initially predicted a U-shaped relationship between confidence and advice requests when external information about advisor accuracy was not available, a condition which we tested in all three experiments. Notably, participants consistently requested advice more often when confidence was high, but did not show a consistent effect for low confidence across the three experiments. While the data from Experiments 2 and 3 reveal a trend towards frequent advice requests when confidence was low, this pattern was not present in Experiment 1. Several methodological differences between the experiments might have caused these inconsistencies in the results, such as the testing methodology (on-line vs. in-lab testing), different advisor accuracies, and differences in the end-of-block feedback. Importantly, however, though we did not find strong evidence for this pattern of behavior when people did not have information about advisor accuracy, our results did show that, when participants believed advice was accurate, they requested advice more often when confidence was low. These results were most evident in Experiment 2, when participants had explicit information about advisor accuracies. By demonstrating that people do not always request advice more often when their confidence is low, as we initially expected, but rather adjust their behavior based on the specific context and goal, these results provide a degree of support to our hypothesis that confidence is used flexibly.

Seeking advice when confidence is low can be viewed as a form of cognitive offloading – a use of external resources in order to reduce cognitive demand (Risko & Gilbert, 2016). Typically, studies investigating cognitive offloading have studied this in the context of memory, looking into the use of external devices to set reminders (Gilbert, 2015), or save information that might otherwise be forgotten (Hu et al., 2019). These previous studies have demonstrated that offloading strategies can improve memory performance and, most relevant to the present study, that people use metacognitive evaluation cues when determining whether they should rely on such external devices (Hu et al., 2019). Our study extends these findings to demonstrate that

people rely on metacognitive evaluations of their performance to decide when to offload to external sources also in tasks which are not memory related. In other words, much like people would decide whether to write a shopping list based on how likely they believe they are to remember all the items otherwise, they also decide whether they should ask for advice in perceptual tasks based on how likely they think they are to succeed on their own. When people think a task is hard, they reduce the demand on themselves by using the help of others.

Critically, though the present study investigated the use of confidence for two specific purposes—as a cue that a task is hard and external assistance should be used, and as a feedback signal for learning—our main interest was not in these specific uses. Rather, the key idea which we address is peoples' ability to use confidence strategically, as demonstrated by the flexibility in the relationship between confidence and behavior. The flexible use of confidence in our study is perhaps most evident when considering the results of all three experiments together. Results from Experiments 1 and 2 showed that, when no external information about advice was available, the relationship between confidence and advice requests changed as time progressed, with people requesting less advice when confidence was high once they had learned about the quality of advice. In Experiment 3, where participants did not learn about advice quality as effectively, participants continued requesting advice when confidence was high throughout the entire experiment. Taken together, these results show that participants changed their use of confidence strategically, in a way that advanced their current goals. These results are somewhat surprising given previous evidence that the confidence signal is integrated in the brain automatically when making decisions (Lebreton, Abitbol, Daunizeau, & Pessiglione, 2015). Thus, it might be expected that consistent associations are learned between confidence and behavior. For example, a learned association might be that when confidence is low external help is sought to improve performance, while when confidence is high advice is discarded. Our results demonstrate that the relationship between confidence and behavior is more flexible than would be expected based on

such associative mechanisms. An interesting question that remains open is how people are able to learn different uses of confidence and apply them to situations appropriately, particularly when situations differ only in peoples' internal state or level of knowledge, as in the experiments we describe here.

In conclusion, the results of the current investigation show that people use their confidence both as a self-monitoring tool, and as a feedback proxy allowing them to learn about the reliability of others. Our findings fit in with the concept of subjective confidence specifically, and metacognitive cues in general, acting as a tool that can be used to monitor behavior and adjust performance accordingly (Nelson & Narens, 1990). As such, our subjective sense of confidence is a powerful tool that we carry with us constantly and use flexibly according to our current goals and needs.

Acknowledgements

We thank Niccolo Pescetelli for his help with Matlab scripts and Matt Jaquiere for his help with Javascript. Nomi Carlebach was supported by a DPhil studentship from the Department of Experimental Psychology, University of Oxford, and grants from the British Friends of the Hebrew University, Funds for Women Graduates, B'nai B'rith Leo Baeck (London) Lodge and the Anglo-Israel Association.

Appendix A. Supplementary material

Supplemental Information to this manuscript can be found in the attached file.

References

- Baranski, J. V., & Petrusic, W. M. (1994). The calibration and resolution of confidence in perceptual judgments. *Perception & Psychophysics*, 55(4), 412–428. <https://doi.org/10.3758/BF03205299>
- Behrens, T. E. J., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. S. (2008). Associative learning of social value. *Nature*, 456(7219), 245–249. <https://doi.org/10.1038/nature07538>
- Boldt, A., Blundell, C., & De Martino, B. (2019). Confidence modulates exploration and exploitation in value-based learning. *Neuroscience of Consciousness*, 2019(1), 1–12. <https://doi.org/10.1093/nc/niz004>
- Boldt, A., & Gilbert, S. J. (2019). Confidence guides spontaneous cognitive offloading. *Cognitive Research: Principles and Implications*, 4(1). <https://doi.org/10.1186/s41235-019-0195-y>
- Boldt, A., Schiffer, A.-M., Waszak, F., & Yeung, N. (2019). Confidence Predictions Affect Performance Confidence and Neural Preparation in Perceptual Decision Making. *Scientific Reports*, 9(1), 4031. <https://doi.org/10.1038/s41598-019-40681-9>
- 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. <https://doi.org/10.1016/j.obhdp.2006.07.001>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9176952>
- Carlebach, N., & Yeung, N. (2020). Confidence Acts as an Internal Cost-Benefit Factor When Choosing Between Tasks. *Journal of Experimental Psychology : Human Perception and Performance*.
- de Gardelle, V., Le Corre, F., & Mamassian, P. (2016). Confidence as a Common Currency between Vision and Audition. *PLOS ONE*, 11(1), e0147901. <https://doi.org/10.1371/journal.pone.0147901>
- de Gardelle, V., & Mamassian, P. (2014). Does Confidence Use a Common Currency Across Two Visual Tasks? *Psychological Science*, 25(6), 1286–1288. <https://doi.org/10.1177/0956797614528956>
- Dehaene, S., & Changeux, J. P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70(2), 200–227. <https://doi.org/10.1016/j.neuron.2011.03.018>
- Desender, K., Boldt, A., Verguts, T., & Donner, T. H. (2019). Confidence predicts speed-accuracy tradeoff for subsequent decisions. *ELife*, 8, 1–25. <https://doi.org/10.7554/eLife.43499>
- Desender, K., Boldt, A., & Yeung, N. (2018). Subjective Confidence Predicts Information Seeking in Decision Making. *Psychological Science*. <https://doi.org/10.1177/0956797617744771>
- Desender, K., Murphy, P., Boldt, A., Verguts, T., & Yeung, N. (2019). A postdecisional neural marker of confidence predicts information-seeking in decision-making. *Journal of Neuroscience*, 39(17), 3309–3319. <https://doi.org/10.1523/JNEUROSCI.2620-18.2019>
- El Zein, M., Bahrami, B., & Hertwig, R. (2019). Shared responsibility in collective decisions. *Nature Human Behaviour*, 3(6), 554–559. <https://doi.org/10.1038/s41562-019-0596-4>

- Frost, R. O., Marten, P., Lahart, C., & Rosenblate, R. (1990). The dimensions of perfectionism. *Cognitive Therapy and Research*, 14(5), 449–468. <https://doi.org/10.1007/BF01172967>
- Gilbert, S. J. (2015). Strategic use of reminders: Influence of both domain-general and task-specific metacognitive confidence, independent of objective memory ability. *Consciousness and Cognition*, 33, 245–260. <https://doi.org/10.1016/j.concog.2015.01.006>
- Gino, F., Brooks, A. W., & Schweitzer, M. E. (2012). Anxiety, advice, and the ability to discern: Feeling anxious motivates individuals to seek and use advice. *Journal of Personality and Social Psychology*, 102(3), 497–512. <https://doi.org/10.1037/a0026413>
- Gino, F., & Moore, D. A. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making*, 20(1), 21–35. <https://doi.org/10.1002/bdm.539>
- Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *eLife*, 5, 1–19. <https://doi.org/10.7554/eLife.13388>
- Hainguerlot, M., Vergnaud, J.-C., & de Gardelle, V. (2018). Metacognitive ability predicts learning cue-stimulus associations in the absence of external feedback. *Scientific Reports*, 8(1), 5602. <https://doi.org/10.1038/s41598-018-23936-9>
- Harvey, N., & Fischer, I. (1997). Taking Advice Accepting Help, Improving Judgment. *Organizational Behavior and Human Decision Processes*, 70(2), 117–133. Retrieved from <https://pdfs.semanticscholar.org/2987/5f69c4e39bee1c415c6e9fbae7b1f02bc19b.pdf>
- Hu, X., Luo, L., & Fleming, S. M. (2019). A role for metamemory in cognitive offloading. *Cognition*, 193(August 2018). <https://doi.org/10.1016/j.cognition.2019.104012>
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. *Perception*, 36(14), 1–16.
- Koriat, A., & Goldsmith, M. (1996). Monitoring and control processes in the strategic regulation of memory accuracy. *Psychological Review*, 103(3), 490–517. <https://doi.org/10.1037/0033-295X.103.3.490>
- Lebreton, M., Abitbol, R., Daunizeau, J., & Pessiglione, M. (2015). Automatic integration of confidence in the brain valuation signal. *Nature Neuroscience*, 18(8), 1159–1167. <https://doi.org/10.1038/nn.4064>
- Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, 47(4), 1122–1135. <https://doi.org/10.3758/s13428-014-0532-5>
- Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, 88(1), 78–92. <https://doi.org/10.1016/j.neuron.2015.09.039>
- Nelson, T. O., & Narens, L. (1990). Metamemory: A Theoretical Framework and New Findings. *The Psychology of Learning and Motivation*, 26, 125–173. [https://doi.org/10.1016/S0079-7421\(08\)60053-5](https://doi.org/10.1016/S0079-7421(08)60053-5)
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9176953>
- Pescetelli, N., & Yeung, N. (2018). The role of decision confidence in advice-taking and trust formation. *ArXiv Preprint ArXiv:1809.10453*, (October). Retrieved from <http://arxiv.org/abs/1809.10453>

- Pescetelli, N., & Yeung, N. (2020). The role of decision confidence in advice-taking and trust formation. *Journal of Experimental Psychology: General*.
- Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. *Experimental Psychology*, 71(2), 264–272. <https://doi.org/10.1037/h0022853>
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive Offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Rouault, M., Dayan, P., & Fleming, S. M. (2019). Forming global estimates of self-performance from local confidence. *Nature Communications*, 10(1), 1–11. <https://doi.org/10.1038/s41467-019-09075-3>
- Samaha, J., Switzky, M., & Postle, B. R. (2019). Confidence boosts serial dependence in orientation estimation. *Journal of Vision*, 19(4), 1–13. <https://doi.org/10.1167/19.4.25>
- Snizek, J. A., & Buckley, T. (1989). Social influence in the advisor-judge relationship. *Annual Meeting of the Judgment and Decision Making Society, Atlanta, Georgia*.
- Snizek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, Vol. 62, pp. 159–174. <https://doi.org/10.1006/obhd.1995.1040>
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., ... Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, 168(3), 242–249. <https://doi.org/10.1016/j.psychres.2008.05.006>
- Van Swol, L. M., Paik, J. E., & Prael, A. (2018). *Advice Recipients* (Vol. 1; E. L. MacGeorge & L. M. Van Swol, Eds.). <https://doi.org/10.1093/oxfordhb/9780190630188.013.2>
- Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, 93(1), 1–13. <https://doi.org/10.1016/j.obhdp.2003.08.002>
- 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. <https://doi.org/10.1006/obhd.2000.2909>
- Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 367(1594), 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>
- Yeung, N., & Summerfield, C. (2014). Shared mechanisms for confidence judgements and error detection in human decision making. In S. M. Fleming & C. D. Frith (Eds.), *The Cognitive Neuroscience of Metacognition* (Vol. 9783642451, pp. 147–167). https://doi.org/10.1007/978-3-642-45190-4_7