The role of decision confidence in advice-taking and trust formation

Niccolò Pescetelli^{1,2} and Nick Yeung¹

¹Department of Experimental Psychology, University of Oxford

²MIT Media Lab, Cambridge, MA

September 25, 2019

Abstract

In a world where ideas flow freely between people across multiple platforms, we often find ourselves relying on others’ information without an objective standard to judge whether those opinions are accurate. The present study tests an agreement-in-confidence hypothesis of advice perception, which holds that internal metacognitive evaluations of decision confidence play an important functional role—namely being a learning signal that allows to learn about the reliability of others in the absence of feedback— in the perception and use of social information, such as peers’ advice. We propose that confidence can

be used, computationally, to estimate advisors’ trustworthiness and advice reliability. Specifically, these processes are hypothesized to be particularly important in situations where objective feedback is absent or difficult to acquire. Here, we use a judge-advisor system paradigm to precisely manipulate the profiles of virtual advisors whose opinions are provided to participants performing a perceptual decision making task. We find that when advisors’ and participants’ judgments are independent, people are able to discriminate subtle advice features, like confidence calibration, whether or not objective feedback is available. However, when observers’ judgments (and judgment errors) are correlated—as is the case in many social contexts—predictable distortions can be observed between feedback and feedback-free scenarios. A simple model of advice reliability estimation, endowed with metacognitive insight, is able to explain key patterns of results observed in the human data. Finally, we use agent-based modeling to explore implications of these individual-level decision strategies for network-level patterns of trust and belief formation.

1 Introduction

Situations in which feedback on others’ performance is not immediately available or costly to acquire are abundant in our life. In many contexts, including in education and health, we rely on advice but may not have immediate feedback or other objective standards with which to judge the reliability of that

advice. Yet in these contexts we must learn to distinguish good from bad advice and consequently which advisors we should trust. How people do this, and how reliably they do so, are open questions.

Many influential models of trust rely on the trustee’s characteristics, like their ability or integrity (Mayer, Davis, & Schoorman, 1995). In the current work, we are interested in a specific aspect of trust that pertains the accuracy and reliability of the trustee. In some situations, the environment might provide contextual cues regarding such characteristics (e.g., individual, institutional or qualification-based reputation of the advisor). However, here we are interested in whether people can learn advice reliability through experience even in domains where feedback is sporadic, unreliable, costly or completely absent—and what mechanisms underlie trust formation in these contexts. Specifically, we investigate whether people can detect subtle informational differences among their social partners in the absence of any objective feedback and, if so, whether the trustor’s internal metacognitive states play a role in their estimation of advice reliability.

When available, external (reliable) feedback can guide learning through reinforcement learning mechanisms that have been thoroughly described in previous research (Behrens, Hunt, Woolrich, & Rushworth, 2008; Guggenmos, Wilbertz, Hebart, & Sterzer, 2016; Sutton & Barto, 1998). When applied to a social context, external feedback can be used to infer the reliability of external sources of advice by means of the same associative processes pre-

viously established for learning other, simpler, features of the environment (Behrens et al., 2008). But a similar kind of social inference is likely to occur also when such objective external feedback is not readily available. Our hypothesis is that the solution to this seemingly computationally intractable problem can be approximated by using two pieces of evidence that are always available when learning about the reliability of other social actors — namely, the advisor’s agreement rate with one’s own belief and one’s own internal decision confidence: If on a given decision we are certain we are correct, then we can equally be certain that anyone who disagrees with us is wrong, and accordingly down-weight their opinion in the future. If on the contrary we make a choice with less confidence, we should still down-weight disagreeing advice, but now to a lesser extent. In other words, we suggest that people can overcome the absence of objective feedback by exploiting the trial-by-trial covariation between their own internal decision confidence and the actual state of the environment (Pescetelli, Rees, & Bahrami, 2016). We call this heuristic the agreement-in-confidence heuristic. In contrast with other strategies based on classification variability (Weiss & Shanteau, 2003), which require repeated observations, this simple mechanism allows people to estimate an advisor’s reliability even in one-shot interactions. This strategy is consistent with previous studies in social psychology on advice taking, naïve realism and social judgment theory’s “latitude of acceptance”, which consistently have shown a tendency for people to discount disagreeing opinions, underweight advice as a factor of distance from one’s own opinion and consider

one’s own opinions as more objective than others’ (Ecken & Pibernik, 2016; Liberman, Minson, Bryan, & Ross, 2012; Minson, Liberman, & Ross, 2011; Schultze, Gerlach, & Rittich, 2018; Sherif, Sherif, & Nebergall, 1965; Soll & Larrick, 2009; Yaniv, 2004). However, it differs from these earlier findings by suggesting that these phenomena might be by-products of a best response strategy to situations where inference must be made without readily available objective feedback. The reasons why people fail to use advice is poorly understood (Rader, Larrick, & Soll, 2017). Our framework proposes a computationally viable mechanism, which we support with empirical evidence. We show that subjective metacognitive confidence can indeed interact with past agreement to provide a useful feedback approximation that can guide social perceptions of others’ reliability. Our proposal is also tangentially related to some accounts of the “false consensus” effect— whereby people who endorse an answer believe that others are also more likely to endorse it (Ross, Greene, & House, 1977)— that provide computational arguments for using one’s own belief as a significant sample of the population and to base predictions about others on such beliefs (Dawes, 1989).

Confidence can be formally defined as the (subjectively estimated) likelihood of being correct given the evidence and the decision made (Fleming & Daw, 2017). It has been shown to be an important factor that allows cognitive control (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Fleming & Dolan, 2012) and meta-learning (Flavell, 1979). Confidence has also been

shown to be involved in group decisions and information integration processes (Bonaccio & Dalal, 2006; Rader et al., 2017) and the importance of confidence in human interactions is well established. Much of this literature employs the judge-advisor system (JAS) paradigm (Bonaccio & Dalal, 2006; Rader et al., 2017; Sniezek & Buckley, 1989; Yaniv & Kleinberger, 2000), in which a participant (“judge”) is typically asked to give a first answer to a question and is then presented with the opinions of one or more advisors whose advice can then be used to update the initial judgment. Evidence suggests that confident eyewitness testimony is weighted more by members of a jury (Penrod & Cutler, 1995; Roediger III, Wixted, & Desoto, 2012), confident people are trusted more and are more influential within groups (Swol & Sniezek, 2005; Zarnoth & Sniezek, 1997), and confidence mediates the effect of power on advice taking (Tost, Gino, & Larrick, 2012). People have been shown to use confidence to arbitrate disagreement irrespective of true accuracy (Hertz, Romand-Monnier, Kyriakopoulou, & Bahrami, 2016; Mahmoodi et al., 2015; Penrod & Cutler, 1995). However being confident comes to a cost if confidence is not predictive of accuracy (i.e., if confidence is not *calibrated*). For instance people tend to distrust confident but unreliable testimony (Tenney, MacCoun, Spellman, & Hastie, 2007), at least when feedback is provided or easy to acquire (Sah, Moore, & Maccoun, 2013).

The hypothesis explored here is that the role of confidence in social scenarios goes beyond the simple weighting of diverse opinions, and can in addition

be used to approximate useful inferences about others. Two lines of reasoning converge to this conclusion. First, in a probabilistic framework, confidence can be interpreted as the subjective estimated likelihood of having made the correct decision, given the internal sensory states and the decision made (Aitchison, Bang, Bahrami, & Latham, 2015; Fleming & Daw, 2017; Meyniel, Sigman, & Mainen, 2015; Pouget, Drugowitsch, & Kepecs, 2016). Second, in many perceptual tasks used in group decision and advice-taking studies, accuracy correlates within-participant with confidence; that is, more confidently expressed judgments are, on average, more likely to be correct (Henmon, 1911; Koriat, 2012b). We suggest that these features make confidence signals extremely valuable when objective feedback is missing because they can be used as a proxy for feedback when the latter is absent. This proposal differs from a simple confidence heuristic described in the literature (Price & Stone, 2004) in that people should use their *own* confidence rather *advisor's* confidence as a proxy signal to advisor's accuracy. Previous studies had demonstrated that conveying information about confidence between participants is crucial for collective benefits in group performance (Bahrami et al., 2012b; Fusaroli et al., 2012; Koriati, 2012b). Bahrami et al. (2012a) demonstrated that in the absence of trial-by-trial feedback, participants were able to accrue collective benefit — that is, group performance exceeding that of the most accurate member's performance — if verbal interaction between participants was allowed. In contrast, providing objective feedback on members' trial accuracy but not allowing for verbal communication did not lead

to a collective benefit. This is an interesting finding considering that what should ultimately drive trust in a partner and the decision to follow their judgment is their accuracy rate.

We further speculate that judges' subjective confidence can be used (albeit with inherent noise) to infer the likelihood that a social advisor is correct in a similar way that external feedback is used to track advice reliability using simple reinforcement learning rules (Behrens et al., 2008). Although it might appear confusing to talk about reinforcement learning when no objective reward is present, our intuition is motivated by the fact that subjective confidence can be intrinsically rewarding and has been shown to provide a useful learning signal with neural markers similar to external rewards (Guggenmos & Sterzer, 2017; Guggenmos et al., 2016).

To the best of our knowledge, studies have so far investigated a person's confidence either as a signal to others (e.g., confident advisors are trusted more (Bahrami et al., 2010; Penrod & Cutler, 1995; Swol & Sniezek, 2005)), or as a signal for one self (e.g., confidence reflects a person's perceived likelihood of being correct (Aitchison et al., 2015; Botvinick et al., 2001; Fleming & Daw, 2017; Meyniel et al., 2015)). The present work's contribution is in exploring how people's own confidence also affects their learning about others and thus how influential received social information will be. The study investigates learning in feedback-free scenarios, which have received little research attention to date (see (Rouault, Dayan, & Fleming, 2019; Weiss &

Shanteau, 2003) for exceptions).

We investigate these questions using a perceptual decision paradigm, which have been used in group decision and JAS studies due to the high degree of control that experimenters have over the information provided to participants and the ability to define an optimal behavior (Bahrami et al., 2010; Bang et al., 2017; Sorkin, Hays, & West, 2001). Our findings show that people are indeed able to learn the reliability of an information source even when deprived of external feedback, and indicate that they do so by relying on internal signals of decision confidence. We further show that this simple mechanism works well when the advisor and advisee’s judgments are independent, but shows predictable downsides when information between them is correlated. Finally, we expand these empirical results on dyads to larger groups of simulated agents, to understand how these cognitive mechanisms may affect interactions in larger groups—ever more common in online debate spaces—under realistic assumptions. We find that these mechanisms can potentially alter communication network structures, depending on individuals’ information and bias.

1.1 Overview of the current research

To investigate the question of whether people use metacognitive information to judge the reliability of a source of advice when other cues are unavailable, we devised a computer-based judge-advisor system task in which participants performed a series of simple perceptual judgments where they first gave their

response and associated confidence, then were shown the opinion of different virtual advisors with predetermined informational profiles, and finally were asked to state their (perhaps revised) decision and confidence. Crucially we manipulated across participants the presence of objective trial-by-trial feedback. The group receiving this trial-by-trial feedback provides a baseline where we expected participants to accurately judge advisor reliability (Behrens et al., 2008). The behavior of interest is whether, in the absence of feedback, participants are still able to track the reliability of different advisors, and whether they use this reliability estimate to weight proffered advice appropriately. Two measures of trust were recorded to allow for dissociations between implicit and explicit behaviors. First, participants were explicitly asked to rate the reliability of advisors every second block of the experiment (see Supplementary Information for details). Second, the degree to which participants updated their judgments and associated confidence from pre- to post-advice provided another measure of judged reliability, with larger updates indicating greater advice influence (Bonaccio & Dalal, 2006). We call this measure “implicit” as participants were not explicitly prompted to report it. Although factors such as compliance and conformity might affect this measure without affecting the underlying beliefs Kelman (1958), these would be orthogonal to our experimental manipulation and average out over repeated observations.

The perceptual paradigm used allows direct manipulation of the evidence available to participants. At the same time, the use of “virtual” rather than

real human advisors (cf. Bahrami et al. (2010) and Zarnoth and Sniezek (1997)) allows us strict control over the informational content of the social sources. Bayes theorem can be used to compute the precise information carried by each advisor, as a normative framework within which to test our hypotheses and evaluate participants’ performance.

Experiment 1 was designed to test the value of using personal confidence information to learn subtle but important advisors’ characteristics, specifically advice accuracy and calibration. When feedback is readily available, both accuracy (Behrens et al., 2008) and calibration (Tenney et al., 2007) have been shown to be valued advice features affecting both advice trustworthiness and influence. By manipulating across participants the availability of trial-by-trial feedback, we show that people are capable of learning subtle cues in their advisors, and both with and without objective feedback they value accurate advice and calibrated advisors.

We use a simple Bayesian confidence-updating rule to generate formal predictions of different strategies for estimating advice reliability, when applied to the collected data. When objective feedback is available, the model can use the accuracy history of the advisor to estimate its reliability (*Accuracy* variant). When feedback is absent on the contrary, two alternative rules could in principle be used to overcome the absence of objective feedback, but only one is endowed with access to metacognitive information. Both model variants exploit the association existing between accuracy and agreement

among observers (Koriat, 2012b). Assuming judgments are independent and performance is above chance, agreement rate provides an indication of an advisor’s accuracy (e.g., two perfectly accurate observers will agree 100% of the time on the correct answer). Exploiting this feature, a first model variant (the *Confidence* variant) uses a nuanced strategy that weights agreement by subjective pre-advice confidence. According to this strategy, high confidence agreement is more informative than low confidence agreement. Again, this proposal differs from previous ones in suggesting that the advisee’s rather than advisor’s confidence will carry such information. We contrasted the *Confidence* variant with a simpler *Consensus* variant, which does not have access to pre-advice confidence and instead simply estimates the reliability of the advisors according to the rate of agreement with each advisor. We show that these simple algorithms capture patterns of the results seen in the human data.

In Experiment 1, advisors’ judgements were independent from the participant’s judgments. Experiments 2 and 3 extend Experiment 1 with two related aims. The first one is to distinguish between the *Consensus* and *Confidence* variants. The second is to explore inherent limitations in consensus and confidence-based estimates of reliability. Notice that agreement-based models have an inherent flaw in them: An agent using an agreement-based strategy will systematically overestimate the reliability of advisors who agree with the agent when the agent is wrong, or in other words when agent and advisor share the same biases in their judgments. Systematic deviations will

occur whenever two people are basing their judgments on the same, incomplete information, or making a decision with the same inherent biases (cf. (Tversky & Kahneman, 1974)). We expect these situations to be frequent in many everyday scenarios—e.g., online environments—where people can easily self-assort into information-insulated cliques (Sunstein, 2001). Key to both aims is manipulation of the coupling between agreement and accuracy by breaking the independence of initial judgment and advice that is present in Experiment 1. We predict that this will lead to predictable distortions in trust and influence, that differ from feedback-based objective reliability of advice, and that the patterns of trust and influence are best captured by a *Confidence* variant.

In this way, we combine experiments with model-informed analyses. The experiments track patterns of trust emerging in human participants in a task in which advice information and perceptual evidence is carefully manipulated. The models identify the patterns of trust and influence that should be observed if people adopted simple learning strategies based on accuracy estimated from feedback, simple agreement or confidence-weighted agreement. In a final section, we use agent-based modeling to explore the way in which the simple mechanisms of trust formation that we identify empirically might generalize to affect network formation in larger groups of interacting individuals. We show that, under realistic assumptions, the empirically identified mechanisms lead to divergent network structures, which are sensitive to the agents’ information and biases.

2 Experiment 1

Experiment 1 investigates whether people learn about subtle advice reliability cues—specifically, advice accuracy and calibration — in environments that lack objective feedback. Four virtual advisors with differing reliability were designed. They differed in accuracy (proportion correct) and calibration (how expressed confidence scales with probability correct (Fleming & Lau, 2014)). Participants repeatedly experienced each advisor to give them the opportunity to learn about advisor reliability. Learning was examined in terms of both explicit rating of trust in each advisor, and implicit expression of advice influence, to allow for possible dissociations between explicit and implicit trust. The crucial question of interest was whether these measures of trust and influence would be sensitive to differences in advisor accuracy and calibration even in the absence of objective feedback about the correct answer on each trial. By including a second group of participants, who did have access to trial-by-trial feedback, we could also assess the degree to which the putative agreement-in-confidence heuristic provided an adequate vs. imperfect proxy for objective feedback

2.1 Method

2.1.1 Participants

Volunteers ($N = 46$, females = 26, age = 23 ± 0.45) were recruited in exchange for monetary compensation or course credits. Half of the partici-

pants were assigned to the Feedback condition and the other half to the No-Feedback condition. The study was approved by local ethical committee. All participants gave informed consent prior to participation.

2.1.2 Paradigm

The perceptual task used is a dot-count comparison task already described in Boldt and Yeung (2015), in which participants are asked to judge which of two boxes contains more dots (Figure 1A). The boxes are presented briefly, for 160 ms, to the left and right of a fixation cross, with dots randomly arranged in a 20 x 20 grid. One box contains more dots, specifically $ndots = 200 + d$, compared to the other with $ndots = 200 - d$. By manipulating the d parameter we can control the difficulty of the task. Difficulty was titrated to each individual sensitivity by applying a 2-down-1-up staircase procedure (Treutwein, 1995) to ensure similar overall accuracy across participants (nominal accuracy rate = 70.7%). The location of the box with more dots was predetermined in advance by pseudo-randomisation that ensured the number of left and right correct answers was balanced across the experiment.

Figure 1: (a): Experiment 1 paradigm. The task represents a computerised version of the classical judge-advisor system (Snizek & Buckley, 1989). In our version, the participant first makes a decision about which box contains most dots. The participant expresses his/her opinion on a semi-continuous confidence scale ranging from “100% sure left” to “100% sure right”. They are then presented with the opinion of a virtual advisor and are asked to update their original judgment with a new rating along the scale. (b): A simple model that estimates accuracy from pure agreement rate will underestimate the accuracy of any advisor unless the model itself is always correct in its judgments: Estimated accuracy = $a*b + (1-a)*(1-b)$, where a is the objective accuracy of the judge’s decisions and b is the advisor’s accuracy; i.e., their likelihood of agreeing on the correct answer plus their likelihood of agreeing on an incorrect answer.

After the brief visual presentation of the stimuli (160 ms) participants had unlimited time to enter their response and confidence judgment. They did so in a one-step decision by clicking with the mouse along a semi-continuous scale in 10 steps—ranging from “100% sure left” to “100% sure right”—and confirming their response pressing the spacebar. Text landmarks signalling 10% increases aided the interpretation of the scale. The middle point of

the scale (50% or total uncertainty) was removed and a gap appeared instead, meaning that participants had to commit to one interval (2-alternative forced-choice).

After confirming their response, one out of four different advisors appeared at the center of the screen as a head-shot picture. Pictures were selected from the NimStim database (Tottenham et al., 2009) and depicted four Caucasian, smiling female characters that were randomly assigned for each participant to one of the four accuracy/calibration conditions described below. Advice was provided as a pre-recorded female voice through active noise-cancelling headphones. Audio tracks were recorded in from native English speakers and their duration was altered with Audacity[®] so that all lasted for exactly two seconds. The association between advisor voice, face and advice profile was randomised across participants to avoid confounds due to appearance or voice characteristics. Advisors could express a binary level of confidence (low vs. high) and either agree or disagree with the participant’s judgment. Low confidence was expressed by the sentences “I think it was on the [LEFT/RIGHT]” and “It was on the [LEFT/RIGHT], I think”, with one of the two randomly assigned on every trial. Similarly, high confidence was expressed by the sentences “I’m sure it was on the [LEFT/RIGHT]!” and “It was on the [LEFT/RIGHT], I’m sure!”. The use of two inverted sentences for each confidence cue (“I’m sure” vs. “I think”) was to avoid over-repetition of a single sentence and to balance the differences in emphasis that the English language conveys when using the confidence cues (“I think” vs. “I’m sure”)

at the beginning or at the end of the sentence. The selection of LEFT or RIGHT depended on the advisor’s choice and accuracy as described below.

After the advice was given, participants had unlimited time to update their decision and confidence level using the same interface and input method as used in the pre-advice period. The question “What is your final decision?” appeared to prompt the update. The pre-advice confidence level remained on the scale as a shaded marker to remind participants of their initial opinion. Participants were allowed to stay in the same position along the scale, increase and decrease their confidence, or even change their minds (i.e., changing interval along the scale). Again they confirmed their final decision by pressing spacebar. In the Feedback condition only, after the final decision was confirmed, a high frequency error tone communicated whether the participant’s final decision was incorrect. In the No-Feedback condition, a new trial started immediately after participants had confirmed their final answer.

At the end of each block, a summary on the post-advice percentage accuracy of the participant (but not of the advisor) was provided to both groups. Notice that advisor presentation was balanced within blocks so the feedback participants received at the end of each block could not favor one advisor over the others. Participants performed 500 trials divided in 10 experimental blocks. Prior to these, two initial blocks with a fifth advisor served as practice and were removed from all the analyses. On each experimental block, each advisor appeared ten times. Ten randomly selected trials within each block were presented with a black silent silhouette and a post-advice decision

was not required (null trials). This was done to motivate participants to provide meaningful answers in their pre-advice answers on each trial and avoid pre-advice random guessing. After every two experimental blocks, participants answered a brief questionnaire about their explicit opinions about the four advisers. Four questions asked participants to directly rate on a scale from 1 (Not at all) to 50 (Extremely) how much they thought each adviser was accurate (Q1), confident (Q2), trustworthy (Q3) and influential on their own choices (Q4) (see SI for complete description).

2.1.3 Manipulation

We orthogonally manipulated the average accuracy of the four advisers and their confidence-to-accuracy calibration (Table 1). We kept the raw numbers of confident and unconfident advice trials equal across advisers (50% high confidence rate), to avoid people simply trusting the advisor who was the most confident *on average* when objective feedback is not available (Sah et al., 2013). We defined two accurate advisers by setting their accuracy to 80% and two inaccurate advisers by setting their accuracy to 60%. Then we set the calibrated advisers to be always accurate whenever confident, and the uncalibrated advisers' confidence to be entirely uninformative of accuracy. This led to the profiles shown in Table 1. Calibration is a term used in metacognitive studies indicating the strength of co-variation between confidence judgments and accuracy. Here we quantify calibration as Type 2 A_{ROC} (A''_{ROC}), a method which does not make assumptions about the gen-

erative model of confidence (Fleming & Lau, 2014). Uncalibrated advisors both had an A''_{ROC} of 0.5, meaning that confidence was totally uninformative in predicting the advisor’s trial-level accuracy. Due to the experimental design—in which overall accuracy of advisors was fixed at 60% or 80%, and calibrated advisors were always correct when high in confidence—calibrated advisors differed in their metacognitive sensitivity according to this metric (Accurate Calibrated advisor $A''_{ROC} = 0.72$; Inaccurate Calibrated advisor $A''_{ROC} = 0.84$).

We quantified the informational value of each advisor as the average information gain over possible advice encounters (see Supplementary Information). The mean absolute information gain so computed was lowest for the Inaccurate Uncalibrated advisor, intermediate for the Accurate Calibrated and Accurate but Uncalibrated advisors and the highest for the Calibrated but Inaccurate advisor (Table 1). This pattern can be intuitively understood by looking at Table 1: Although the Inaccurate Calibrated advisor’s accuracy rate is lower than the Accurate Calibrated advisor, outcomes can be better predicted by her judgment. In particular, her judgments correlate strongly positively with the correct answer when sure and strongly negatively when unsure. On the contrary, when the Accurate Calibrated advisor is unsure there is a much higher uncertainty about the final outcome. A related measure, the expected information gain, which takes into account the frequency of each event (see SI), confirmed a similar pattern of objective advice value.

Events count	Advisors			
	Accurate Calibrated	Accurate Uncalibrated	Inaccurate Calibrated	Inaccurate Uncalibrated
Incorrect Confident	0	1	0	2
Incorrect Unconfident	2	1	4	2
Correct Unconfident	3	4	1	3
Correct Confident	5	4	5	3
A''_{ROC}	.72	.5	0.84	.5
IG	0.29	0.26	0.38	0.08
IG_e	0.063	0.063	0.084	0.021

Table 1: Experiment 1 - Advisors profile. Values in the central section represent the number of times each event occurred over the course of a 50-trial block (10 null trials, with no advisor, are not shown). Type 2 area under the receiver operating characteristic curve A''_{ROC} represents the metacognitive sensitivity or calibration of each advisor. Information gain IG and expected information gain IG_e quantify the informativeness of the advice received.

2.1.4 Exclusion criteria

An exclusion criterion was set *a priori* for staircase convergence. Participants who showed progressively increasing thresholds (i.e., increasing d values across the experiment) were to be eliminated as this indicated that they were randomly guessing. None of the participants had to be removed when this criterion was applied to our sample. At the end of the experiment the average difficulty parameter d across participants (pooled data) was 9.6 ± 2.81 .

2.1.5 Three models of advisor reliability estimation

To identify the patterns of trust and influence that would result from adopting simple strategies for advice reliability estimation, we created models that used different types of information to assign value to the advice received. To mimic situations when feedback is provided on a trial-level basis, one model used objective feedback to track each advisor’s reliability (the *Accuracy* strategy). Two different strategies can be followed to estimate advisors’ reliability when feedback is absent. The first strategy is to use agreement rate with the advisors as an indicator of the advisor’s good performance. Assuming that a judge and advisor make independent decisions (below we explore cases where this assumption is violated), then agreement rate will vary as a simple linear function of an advisor’s accuracy, with a slope and maximum value that depends on the judge’s own accuracy (Figure 1B). Thus, judges can distinguish between better and worse advisors based on their agreement rates, but will systematically underestimate the accuracy of all advisors to the extent that their own task performance is imperfect.

However, people might solve the task using a more nuanced strategy that leverages the internal sense of confidence as it covaries with objective accuracy (Henmon, 1911; Koriat, 2012b; Roediger III et al., 2012). An effective strategy is to update trust more strongly from judgments when we are certain of our own answer than judgments where we are uncertain or guessing. Suppose, for example, a judge selects the LEFT interval and estimates the probability that they made the correct choice (i.e., their confidence) to be

80%. Suppose they now receive advice disagreeing with their judgment. Assuming perfect calibration, the judge can then assign a 20% probability (i.e., $1 - 0.80 = 0.20$) that, on that trial, the advisor is correct (i.e., it is 80% likely that they made an error); and the amount of learning that is possible about the advisors, scales monotonically with confidence. In contrast, a *Consensus* strategy, not using any subjective metacognitive information, would assign to a disagreement trial a 0% chance that the advisor is correct. Returning to Figure 1B, a judge learning via confidence-weighted agreement is not confined to estimating advisor reliability according to a fixed linear function that depends on their own accuracy, but rather can learn differentially across trials according to their confidence. An interesting and valuable feature of this approach is that it allows a judge to learn that an advisor is in fact more accurate than they are themselves (something that is impossible with simple agreement-based updating), because good advisors are characterized by agreeing most consistently with the judge in their confidently-made decisions (i.e., those that are most likely to be correct).

Three different model variants (see Supplementary Information for details) account for the Feedback condition (*Accuracy* model), the No-Feedback condition without metacognitive insight (*Consensus* model) and the No-Feedback condition with metacognitive insight (*Confidence* model) respectively. Trial-by-trial agreement, reported confidence and objective feedback from Experiment 1 were used to estimate trust and influence that these model variants would show with each advisor. As is shown below, the three model

variants make similar predictions when participants’ and advisors’ judgments are independent (as in Experiment 1), but generate different results when these initial judgments are correlated (as in Experiments 2 and 3).

2.2 Results

The first set of analyses were performed to show that participants are sensitive to all advice dimensions that were manipulated here, namely accuracy rate, calibration and advice confidence. Of particular interest is whether these main effects vary as a function of feedback availability (i.e., the between-participants manipulation). If the pattern of trust and influence changes as a function of feedback presence vs. absence, this would indicate that different advice dimensions are made more or less salient by the presence of feedback. These questions were investigated through the analysis of both explicit trust ratings and implicit influence measure. Given that in all analyses time was not a significant factor (e.g., when data were pooled across successive blocks to give a factor of time-on-task with 5 levels), the data from all blocks were collapsed together.

2.2.1 Trust ratings

A first analysis focused on participants’ explicit ratings of advisor reliability provided at the end of every second experimental block. Ratings were converted into a unitary measure of trust, after preprocessing steps involving normalisation, baseline correction and PCA as described in the Supplementary

Information. These steps aimed to remove confounds due to inter-individual differences in responding, to advisor appearance and to question wording. In both groups, questions one and four (regarding advisor accuracy and trustworthiness) scored highest on first PCA component. A mixed-design ANOVA was run to test whether Feedback (between-participants), advisor Accuracy (within-participants) and advisor Calibration (within-participants) affected explicit trust ratings (Table 1 in Appendix). Results showed significant main effects for advisor Accuracy ($F(1, 44) = 9.68, p = .003, \eta_G^2 = 0.079$) and advisor Calibration ($F(1, 44) = 12.32, p = .001, \eta_G^2 = 0.076$), but not for Feedback ($F < 1$). Participants trusted accurate over inaccurate advice, and calibrated over uncalibrated advice. No interaction term reached significance ($F(1, 44) < 1.9, p > .16$). Importantly, neither of the main within-participants effects interacted with Feedback, suggesting that participants were sensitive to the reliability of the advice and that this sensitivity did not vary consistently according to the presence or absence of feedback (Figure 2A).

Figure 2: Experiment 1 - Trust and influence patterns for humans and simulations. (a) Average trust ratings in the two feedback groups as a function of advisor accuracy and calibration. Error bars represent s.e.m. (b) The effect of advisor Accuracy, advice Confidence and Calibration, and Feedback group on influence measure. Error bars represent s.e.m. (c) Average estimated advisor reliability (θ) according to different variants of the Bayesian model computed from the data of Experiment 1.

We next ran planned 2-way ANOVAs separately for each feedback group, with within-participant factors of advisor Accuracy and Calibration, to assess whether the overall patterns described above were also reliable in each group. In the Feedback group this analysis revealed reliable main effects of both advisor Accuracy ($F(1, 22) = 8.26, p = .008, \eta_G^2 = .09$) and advisor Calibration ($F(1, 22) = 8.71, p = .007, \eta_G^2 = 0.12$), but no reliable interaction ($F(1, 22) = 1.24, p = .27, \eta_G^2 = .02$).

In the No Feedback group, although comparable numerical trends were apparent, neither main effect of Accuracy ($F(1, 22) = 3.42, p = .07, \eta_G^2 = 0.06$) and Calibration ($F(1, 22) = 4.03, p = .05, \eta_G^2 = 0.04$) reached statistical significance. The interaction term was not significant ($F < 1$).

2.2.2 Influence

Figure 2B plots averaged advisor influence, the implicit measure of learning about advice reliability. Influence was quantified as the signed difference between post and pre-advice confidence (Equation 3, Supplementary Information) and represents the shift in confidence observed after social information. Influence results were analysed using a mixed-design ANOVA that included the same factors of Feedback presence (between-participants), advisor Accuracy and Calibration (both within-participants) as in the trust analysis (Table 2 in Appendix). For the influence measure an additional within-participants factor was included according to the Confidence expressed by the advisor on a given trial. While both advisor's Accuracy and Calibration

are “trait” variables, advice Confidence is a trial-by-trial “state” variable (with values low vs. high). This analysis revealed significant main effects of advisor Accuracy ($F(1, 44) = 14.80, p < .001, \eta_G^2 = 0.02$) and Calibration ($F(1, 44) = 15.84, p < .001, \eta_G^2 = 0.01$), mirroring the observed effects of advisor reliability seen for explicitly expressed trust. A reliable main effect of advisor Confidence ($F(1, 44) = 55.82, p < .001, \eta_G^2 = .12$) indicated that more confidently expressed advice had greater influence. There was a significant interaction between Calibration and advisor Confidence ($F(1, 44) = 9.62, p = .003, \eta_G^2 = 0.004$) indicating that the effect of confidently (vs. unconfidently) expressed advice was greater for calibrated advisors than uncalibrated advisors. The analysis also revealed a significant 3-way interaction between advisor Accuracy, Calibration and Confidence ($F(1, 44) = 4.75, p = .03, \eta_G^2 = 8.5e - 04$), indicating that the two-way interaction between Calibration and Confidence was larger for inaccurate advisors than accurate ones. Importantly, there was no reliable main effect of Feedback ($F < 1$), nor any reliable interaction between Feedback and any other main effects (all $Fs(1, 44) < 1.1, ps > .29$) suggesting again that participants’ sensitivity to advisor reliability, here expressed in terms of the influence of their advice, did not depend significantly on the provision of trial-by-trial feedback.

Although we observed no reliable effects of feedback, we ran planned follow-up ANOVAs for each feedback group separately to assess whether the overall patterns described above were also reliable in each group. In

the Feedback group, this analysis revealed significant effects on influence for advisor Accuracy ($F(1, 22) = 8.23, p = .008, \eta_G^2 = 0.02$) and Calibration ($F(1, 22) = 6.35, p = .01, \eta_G^2 = .03$), and advice Confidence ($F(1, 22) = 32.92, p < .001, \eta_G^2 = 0.18$). A two-way interaction was found between Calibration and Confidence ($F(1, 22) = 7.37, p = .01, \eta_G^2 = .008$) but not between Accuracy and Calibration nor Accuracy and Confidence (both $F < 1$). As shown in Figure 2B, the interaction term reflects the fact that the average influence difference between calibrated and uncalibrated advisors was mainly shown for highly confident advice compared to uncertain advice. Finally a reliable three-way interaction between Accuracy, Calibration and Confidence ($F(1, 22) = 8.32, p = .008, \eta_G^2 = .003$) suggests that such interaction was stronger for inaccurate advisors than accurate ones (Figure 2B).

A corresponding ANOVA on influence measures in the No-Feedback group revealed main effects of advisor Accuracy ($F(1, 22) = 6.87, p = .01, \eta_G^2 = .02$) and Calibration ($F(1, 22) = 9.48, p = .005, \eta_G^2 = .01$), and advice Confidence ($F(1, 22) = 22.95, p < .001, \eta_G^2 = .07$). The results show that participants were more influenced by Accurate and Calibrated advisors, and by Confident advice compared to uncertain advice, paralleling the effects observed when feedback was provided. None of the two-way interaction terms reached significance (Accuracy rate x Calibration: $F(1, 22) = 1.35, p > .25, \eta_G^2 = .001$; Accuracy rate x advice Confidence: $F(1, 22) = 2.90, p = .1, \eta_G^2 = .002$; Calibration x advice Confidence: $F(1, 22) = 2.59, p = .12, \eta_G^2 = .001$), although numerically there were greater effects of confidence for calibrated advisors

(as observed in the Feedback condition). The three-way interaction between Accuracy, Calibration and Confidence was not significant ($F < 1$).

2.2.3 Model results

The three model variants were run based on the actual series of each participants decisions and (where appropriate) their associated confidence, and the advice they received and (where appropriate) its accuracy, in Experiment 1. We looked at how the different variants of the model fared in evaluating different advisors' profiles based on this input. The aim of the following analyses was to verify how the pattern of final model's trust (θ) in each advisor differed when different pieces of information were used to compute it. The model is not intended as a mechanistic description of participants' behaviour. Instead, it aims to show that a simple reliability estimation rule will develop differentiated patterns of trust across advisors depending on the type of information it relies upon. For this analysis, data from the Feedback and No-Feedback groups were pooled together to increase statistical power, as the presence of feedback did not affect the variables that model's variants were based on, namely advisors' accuracy, agreement rates and participant's pre-advice confidence ratings respectively.

The model's estimates of advisor reliability (θ parameter) can be used to predict nominal confidence changes using Bayes rule (see Supplementary Information). Confidence changes predicted by the model could have been used to plot the model's implicit trust in the advisors in a manner similar to

that already shown with the behavioral data. However, contrary to θ values, confidence changes are an implicit measure of trust and they are affected not only by the value assigned to an advisor but also by the pre-advice confidence of that trial. Thus, although the two measures are related, θ represents a direct measure of the model’s belief about advisors’ reliability (a trial-by-trial measure that is not available in the human data).

Each model variant’s final θ values—i.e., the model’s advisor reliability estimates on the last trial of the experiment—were analyzed using a 2x2 repeated measures ANOVA with factors of advisor Accuracy (high vs low) and Calibration (high vs. low). Scaling factors were applied to agreement to take into account the fact that in this experiment advisors provided a confidence judgment with their advice. The *Accuracy* variant is, in this experiment, fully pre-determined by the advisors’ set accuracy rates and thus no statistical analysis were run due to the absence of variability across participants. We plot however its trust value θ for visual reference as it shows that when the model is provided with information about the objective performance of the participant (and thus of the advisors), it is able to distinguish advisors both in terms of their Accuracy rate and their confidence Calibration. The difference between calibrated and uncalibrated advisors is larger for inaccurate than accurate advisors, due to the weights used to convey advice confidence (Figure 2C). In the absence of externally given objective feedback, both the *Consensus* variant—which es-

timates advisors’ reliability by assuming that advisors are correct whenever they agree with the participants own first decision, and wrong otherwise—and the *Confidence* variant—which uses agreement as a proxy for feedback like the *Consensus* variant, but scales them by pre-advice confidence—show greater trust for Accurate ($F(1, 45) > 165.84, p < .001, \eta_G^2 = .44$) and Calibrated ($F(1, 45) = 32.70, p < .001, \eta_G^2 = .13$) advisors compared to inaccurate or uncalibrated ones. Neither variant showed a significant interaction between the two factors ($F(1, 45) < 2.65, p > .11, \eta_G^2 = .01$). Notice that advisors were not constrained to agree with the participant a pre-determined amount of times, thus explaining the variability observed across participants according to the specific sequence of decisions they made and advice they received.

Taken together, these modeling results show that simple computations of advisor reliability perform well at this task even when trial-level feedback is absent, effectively capturing key patterns of trust observed in the human data across feedback conditions. For the task used in Experiment 1, the three variants do not make contradictory predictions on which advisors should be trusted. In particular the two No-Feedback variants (that base trust on simple agreement, or agreement weighted by confidence) cannot be disentangled using the data collected from this experiment, with both showing sensitivity to both accuracy and calibration of an advisor.

2.3 Discussion

In this experiment we tested whether participants would distinguish among advisors with differing profiles, and whether patterns of trust and influence would vary according to the availability of objective feedback. Across conditions, participants expressed greater trust in, and were more influenced by, advice characterised by high confidence, high accuracy and high calibration. Although the first two dimensions are not a surprise (Snizek & Van Swol, 2001), the fact that participants also value advice calibration is less well documented. Tenney et al. (2007) showed, in mock jury decisions, that confident but inaccurate witnesses are discredited more than equally inaccurate but less confident ones. Sah et al. (2013) further confirmed that when feedback is readily available, confident but incorrect advice backfires on an advisor’s reputation. However, when feedback was not available or costly to acquire the same authors found evidence for a confidence heuristic whereby confident advisors were more influential regardless of accuracy (Price & Stone, 2004).

The current study extends this previous work in two respects. First, the behavioral as well as the simulation findings highlight the value in using information regarding agreement and confidence in social learning: Although the patterns of trust and influence were not identical across Feedback and No-Feedback conditions, no consistent (statistically reliable) differences were observed, and overall the results indicate that people can learn useful distinguishing characteristics (accuracy and calibration) even in the absence of objective feedback, and even when average advisor confidence is kept con-

stant (Price & Stone, 2004; Sah et al., 2013). Second, we replicate in a highly controlled perceptual task setting the finding that judges value advisor’s confidence calibration. By allowing repeated interactions with the same advisors and by fixing advisors’ average expressed confidence, the current paradigm was able to quantify advice calibration and show that people value it over and above the overall level of confidence with which advice is expressed. Although the large number of interactions in this study (500 trials) is unlikely to occur in certain daily-life scenarios, the mechanisms described are likely to operate also at smaller scales. Furthermore, these numbers are well within the range of many situations people encounter with colleagues or family members. The results thus should be interpreted as highlighting some fundamental social mechanisms occurring in realistic social contexts.

Simulation results based on straightforward implementations of learning strategies—estimating advice reliability according to objective feedback, agreement, or agreement weighted by confidence—were consistent with the empirical results in showing similar patterns of trust across advisors regardless of the presence or absence of objective feedback. Thus, calculating agreement rates and weighting agreement by subjective confidence can be effective heuristics for estimating advisor reliability when objective feedback is absent. The reason for this is that both agreement and subjective confidence covary with accuracy. Thus, both are useful cues, and in the present data the two are difficult to disentangle. To disambiguate predictions of the two

hypotheses, and to explore crucial limitations in the use of agreement and confidence in estimating advisor reliability, the following two experiments included advice that was not independent of participants' initial decisions. Instead, advice varied systematically to allow orthogonal manipulation of (1) advisors' agreement and accuracy rates and (2) advisors' agreement rates and subjective confidence. While Experiment 1 showed the normative value of using such heuristics when advice is independent from a judge's initial decision, Experiments 2 and 3 explored the impact of these heuristics when advice is not independent.

3 Experiment 2

As shown by the *Consensus* model, the frequency of agreement can be used as a cue to the reliability of advice if observations from individual observers are independent. This observation echoes the Wisdom of Crowds (WoC) effect (Galton, 1907; Surowiecki, 2004), whereby aggregation of individual opinions results in better-than-expert performance. Wisdom of crowds is often cited as the archetype of collective intelligence (Krause, Ruxton, & Krause, 2010). The classic explanation of the effect is that individual judgments are the sum of a signal component and random noise. Averaging individual judgments has the effect of enhancing signal by averaging out noise, with a similar mechanism as repeated measures in statistics (Armstrong, 2001).

However, people's judgments are rarely independent and distorted only

by random noise (Koriat, 2012b; Krause et al., 2010): We watch the same news, browse the same websites and are affected by the same cognitive biases as others we interact with. Thus, there may be dependencies between judgments that people are unlikely account for (Yaniv, Choshen-Hillel, & Milyavsky, 2009). One source of dependence between opinions is the use of similar cognitive heuristics (Tversky & Kahneman, 1974). Other sources can arise from using similar strategies to sample information from the environment (Vandormael, Herce Castañón, Balaguer, Li, & Summerfield, 2017), being exposed to similar signals (Kao & Couzin, 2014; Kao, Miller, Torney, Hartnett, & Couzin, 2014) or belonging to the same network clique (Jamieson & Cappella, 2008; Jasny, Waggle, & Fisher, 2015; Sunstein, 2001). People tend not to take into account advisors’ judgments interdependence (Yaniv et al., 2009). As a result, crowds are known to be susceptible to large error cascades (Le Bon, 1895; Mackay, 1841), economic bubbles (De Martino, ODoherty, Ray, Bossaerts, & Camerer, 2013), polarisation (Myers & Lamm, 1976) and *groupthink* (Janis, 1972; Turner & Pratkanis, 1998). An open question is whether these examples of error magnification are due to a lack of judgment independence. Dependencies in error sources may lead to greater consensus among observers but to a decorrelation from actual accuracy—e.g., by agreeing on errors (Koriat, 2012b). In the context of advice, the focus of the present research, opinions and suggestions might be framed in relation to the person we are advising whether accidentally or deliberately. This framing will influence the extent to which opinions agree or disagree with each

other, and thus the information carried by other people’s consensus. As a first exploration of biased advice scenarios, we thus investigated whether in these contexts people trust accurate advisors or advisors who tend to agree.

Specifically, advisor accuracy was manipulated so that two advisors were on average highly accurate (around 80% accuracy) and two advisors were on average relatively inaccurate (around 60% accuracy). Orthogonally, advisor agreement rate with the participant’s judgment was manipulated so as to create two advisors who agreed with the participant frequently (around 80% of trials) and two advisors who tended to have a lower agreement rate with the participant (around 60%). This created an advisor who is highly accurate and tends to agree with the participant, an advisor who is highly accurate but tends to disagree with the participant, an advisor who is not very accurate but still tends to agree with the participant, and an advisor who is neither accurate or agreeing. The two intermediate profiles are an important test case where the two dimensions of accuracy and agreement work against each other. The agreeing but inaccurate advisor represent someone who shares biases with the participant and so makes similar (correlated) mistakes. The accurate but disagreeing advisor instead represents an advisor who uses different information so is accurate and tends to be correct when the participant makes mistakes and vice versa. Knowing what advice characteristics are made salient by the presence of feedback can shed light on the mechanisms of trust formation and advice influence.

3.1 Methods

3.1.1 Participants

The experiment included 46 participants equally divided between the two feedback groups (37 females in total, 18 of whom were in the Feedback group, mean age: 21.63 ± 3.02). Participants were recruited through Oxford University’s online recruitment platform and local news advertisement. Participants were compensated either with money or course credits. All participants provided informed consent before the experiment. All procedures were approved by the local ethical committee.

3.1.2 Paradigm

The overall design was very similar to Experiment 1, with advice provided and its impact measured in the context of the same dot-count perceptual decision task. Participants performed ten blocks of 44 trials each. In this experiment, advice was presented in the form of a binary judgment, expressed by the sentences “I think it was on the [LEFT/RIGHT]” and “It was on the [LEFT/RIGHT], I think”, depending on advisor’s opinion and random selection of one of the two sentences across trials (i.e., there was no across-trial manipulation of advisor confidence). For participants in the Feedback group only, after the final choice was confirmed, a high frequency beep sound (duration 140 ms) was played whenever the participant’s final choice was incorrect.

At the end of every second block, a computer-based questionnaire was presented to participants prompting them to evaluate the four advisors on a 50-point scale (1=“Not at all”, 50=“Extremely”). The questions were the same as those used in Experiment 1 apart from question two (question about advisor’s confidence) that was replaced with a question about advisor likeability, to account for the advisors not expressing confidence along with their judgments in this experiment.

In each block, the presentation of the four advisors was randomly shuffled across trials, with each advisor appearing exactly ten times. Four additional trials were randomly inserted into the block and served as null trials, where a silent silhouette was displayed instead, to encourage participants to register meaningful initial confidence judgments as well as final (post-advice) judgments. The difficulty of the first-order task was again titrated to a target level of 70.7% accuracy using a 2-down-1-up staircase method. Prior to commencing the 10 experimental blocks, participants completed two practice blocks, with ten practice trials each, with a fifth advisor. Practice trials were removed from all the analyses.

3.1.3 Manipulation

To disentangle agreement rate from the accuracy of the advisors, the probability of agreement conditional on the participant’s choice accuracy was manipulated. Through the staircase procedure it was expected that all participants would converge to an accuracy level of about 70%. This enabled us

to manipulate advisor accuracy and agreement rate by pre-determining the probability of agreement after a participant’s error or a participant’s correct response. Both accuracy and agreement were manipulated to have two levels (high=80% and low=60%). This gave rise to the four advisor profiles defined in Table 2. Probabilities are expressed as a fraction over the number of participants’ expected correct (7) and incorrect (3) judgments, over the number of encounters with one advisor during one block (10).

	Advisors			
	High Accuracy High Agreement	High Acc. Low Agr.	Low Acc. High Agr.	Low Acc. Low Agr.
$p(Agr Cor_s)$	6.5/7	5.5/7	5.5/7	4.5/7
$p(Agr Inc_s)$	1.5/3	0.5/3	2.5/3	1.5/3
Expected Acc. rate	80%	80%	60%	60%
Expected Agr. rate	80%	60%	80%	60%
IG	0.28	0.27	0.03	0.06
IG_e	0.09	0.13	0.01	0.03

Table 2: Experiment 2 advisors’ profiles. Expected accuracy and agreement rates of different advisors are disentangled by manipulating the probability of the advice agreeing with the participant, conditional on the participant’s accuracy. In the table, probabilities are expressed as a fraction of the number of participants’ expected correct (7) and incorrect (3) judgments, during the number of encounters with one advisor (10) in a single experimental block. Information gain and expected information gain— IG and IG_e respectively—indicate average informational value of the advice, computed as information gain and expected information gain respectively (see Supplementary Information).

The dissociation between accuracy and agreement was possible by making the *accurate but disagreeing* advisor disagree more often on trials when the participant had made an incorrect initial decision (meaning a correct response for the advisor). Similarly, an *inaccurate but agreeing* advisor was created by making the advisor more likely to agree when the participant’s initial judgment was incorrect. For the two advisors for whom accuracy and agreement work in the same direction—i.e., *accurate and agreeing* vs. *inaccurate and disagreeing*—it can be expected that participants will favor the former over the latter. However when accuracy and agreement pattern work in opposite directions—as is the case for the agreeing but inaccurate and the disagreeing but accurate advisors—it can be expected that the availability of external feedback will favor the accurate advisor in the Feedback group but not in the No-Feedback group.

Similar to Experiment 1, we used conditional probabilities and the participants’ expected accuracy to compute the informational value of each advisor (Table 2). Advisors’ mean absolute information gain IG and expected information gain IG_e were computed as in the previous experiment. Contrary to Experiment 1, however, advisors did not express different levels of confidence. This created only two possible situations on each trial, either agreement or disagreement. As can be seen in Table 2, information gain was highest for the accurate advisors and the lowest for the *inaccurate but agreeing* advisor. This advisor is in fact counter-intuitively anti-predictive of the correct answer: In contrast with the other advisors, this advisor’s probability of agreement is

higher when the participant’s initial judgment is incorrect than when it is correct. Thus, an optimal observer would decrease their confidence when in agreement and increase it when in disagreement with this advisor. For all other advisors, agreement should increase confidence and disagreement should decrease it.

The same two measures of interest as in Experiment 1—explicit trust and implicit influence—were calculated. Similarly, the three model variants were run on the new behavioral dataset. For pairwise comparisons, both p-values and Bayes factors are reported. Bayes factors are reported using the notation suggested in Dienes and McIatchie (2017). The priors on the effect sizes to compute Bayes factors are informed by results of Experiment 1 (while no priors were available for Experiment 1).

3.1.4 Exclusion criteria

The first two experimental blocks were removed from the analysis to allow the staircase procedure to fully adapt to each individual’s threshold. This was necessary given that our manipulation was heavily dependent on the expected accuracy rate of the participants. A further exclusion criterion was set to exclude all participants whose threshold never converged, which suggests a random response strategy. None of the participants had to be removed on the basis of this criterion. The perceptual task difficulty d (dot difference between boxes) after staircasing was 9.93 ± 2.96 (pooled data).

3.2 Results

A first set of analyses explored the effect of our manipulation on the trust and influence measures as a function of whether objective trial-level feedback was provided. In particular we are interested in whether any of the manipulated within-participants factors (advice agreement rate and accuracy) interacted with feedback presence.

3.2.1 Trust ratings

A mixed-design ANOVA was run on explicit trust ratings with feedback group as a between-participants factor and advisor accuracy (low vs. high) and agreement rate (low vs. high) as within-participants factors (Table 1 in Appendix). This analysis revealed significant main effects for both accuracy ($F(1, 44) = 8.36, p = .005, \eta_G^2 = .06$) and agreement rates ($F(1, 44) = 22.52, p < .001, \eta_G^2 = .1$), but not for feedback ($F < 1$). A significant interaction was found between feedback and accuracy ($F(1, 44) = 8.41, p = .005, \eta_G^2 = .06$) but not between feedback and agreement rate ($F(1, 44) = 1.88, p = .17, \eta_G^2 = .01$) or between agreement and accuracy ($F < 1$), nor a significant three-way interaction ($F < 1$). The results suggest that in their explicit ratings of trust, participants reported on average to prefer advice coming from accurate sources and sources that tended to agree with their own judgments. Importantly, however, the effect of accuracy was modulated by the presence of feedback, with effects of advisor accuracy apparent only in the Feedback condition (Figure 3A).

Figure 3: Experiment 2 - Trust ratings and influence for human participants and simulations. (a) Average trust ratings in the two feedback groups, divided by advisor accuracy rate and agreement rate. Error bars represent s.e.m. (b) Influence that advice had on participants' opinions in the two feedback groups and divided by advisor accuracy rate and agreement rate. Error bars represent s.e.m. (c) Model's θ for the three variants, representing the trust of the model in each advisor. Error bars represent s.e.m.

We then looked at trust ratings in the two feedback conditions individually. A repeated measures ANOVA on trust ratings from the Feedback group revealed significant effects of both advisor's Accuracy ($F(1, 22) = 21.36, p < .001, \eta_G^2 = .20$) and Agreement rate ($F(1, 22) = 5.62, p = .02, \eta_G^2 = .06$), but no significant interaction ($F < 1$). The effect of accuracy was much stronger than the one observed for agreement as indicated by the generalised eta squared values (Bakeman, 2005). Nonetheless, it is notable that agreement had a significant effect on trust even when controlling for accuracy.

A corresponding ANOVA on trust rating data from the No-Feedback group found that Agreement rate ($F(1, 22) = 18.91, p < .001, \eta_G^2 = .18$) but not Accuracy ($F < 1$) affected participants' explicit reliability judgments. No significant interaction was found between the two factors ($F < 1$). These findings suggest that when feedback was not directly available to estimate partners' reliability, participants expressed greater trust in agreeing advisors over disagreeing ones, but not in accurate advisors over inaccurate ones.

3.2.2 Influence

The next analysis focused on influence, as an implicit measure of trust. A mixed-design ANOVA on the influence measure was run with feedback as a between-participants factor and advisor Accuracy (low vs. high) and Agreement rate (low vs. high) as within-participants factors (Table 2 in Appendix). This analysis revealed significant main effects of advisor Accuracy ($F(1, 44) = 14.79, p < .001, \eta_G^2 = .04$) and Agreement rate ($F(1, 44) =$

13.91, $p < .001$, $\eta_G^2 = .03$), and a non-significant interaction between these factors with a numerical trend for participants to be more influenced when feedback was not available ($F(1, 44) = 3.19$, $p = .08$, $\eta_G^2 = .04$) (Figure 3B). Importantly, no significant interactions were found between Accuracy and Feedback ($F < 1$), between Agreement and Feedback ($F(1, 44) = 2.01$, $p = .16$, $\eta_G^2 = .005$) and between Accuracy and Agreement rate ($F < 1$). No significant three-way interaction was found either ($F < 1$). The results show that participants were more influenced by accurate advisors and by advisors characterized by high agreement rates with their own judgments. Importantly, however, the effects did not interact reliably with feedback, suggesting that similar effects were found in the two condition groups independently of feedback availability. We then looked at individual pattern of results within each experimental group.

A two-way ANOVA on influence measures in the Feedback group showed significant effects for both advisor's Accuracy ($F(1, 22) = 12.71$, $p = .001$, $\eta_G^2 = .06$) and Agreement rate ($F(1, 22) = 5.81$, $p = .02$, $\eta_G^2 = .01$). No significant interaction was observed between the two factors ($F < 1$). As expected, when looking at opinion change, participants' opinions were more affected by accurate advisors over inaccurate ones. Interestingly, mirroring effects observed in explicitly rated trust, participants were affected by advisors who tended to agree with their own judgment, although feedback indicated equivalent accuracy rates for pairs of advisors characterized by different agreement rates.

When looking at the No-Feedback condition, a similar pattern emerged. A 2x2 repeated measures ANOVA on influence showed significant main effects for both advisor’s Agreement rates ($F(1, 22) = 8.60, p = .007, \eta_G^2 = .06$) and Accuracy ($F(1, 22) = 4.09, p = .05, \eta_G^2 = .02$), although the effect size of Agreement was the greater of the two. No significant interaction was observed ($F < 1$). Thus, in the No-Feedback group, the results suggest a dissociation between what people reported in the questionnaires, with higher trust reported for agreeing advisors regardless of accuracy, and people’s performance in the task, where influence showed effects of both accuracy and agreement rate.

3.2.3 Model results

The models described above were applied to data from Experiment 2, to understand whether the *Consensus* and *Confidence* model variants behaved differently in scenarios where advice accuracy and advice agreement rate are dissociated. Each model variant was used to provide an estimate of trust in each advisor based on each participant’s expressed confidence, experienced advice and, in the case of the *Accuracy* model, advice objective accuracy. In this experiment, advisors did not express a confidence judgment about their opinions. Thus all model variants could be simplified by not taking into account advice confidence (see Supplementary Information for details).

Given that we are not trying to fit the model to the human data, but only observing how trust patterns change as a function of the information

fed to the model, data from the two feedback groups were pooled together. Trial-by-trial pre-advice confidence and advice were input to each of the three model variants and resulting θ -values for each advisor were compared (Figure 3C). Both the *Accuracy* and the *Confidence* models' θ values show a significant effect of Accuracy ($F(1, 45) > 8.85, p < .005, \eta_G^2 > .05$), while the *Consensus* model only showed a non significant marginal effect ($F(1, 45) = 3.22, p = .07, \eta_G^2 = .02$). Both the *Confidence* and *Consensus* models show a significant effect of Agreement ($F(1, 45) > 434.7, p < .001, \eta_G^2 > .71$), but no reliable interaction between the two factors ($F(1, 45) < 2.04, p > .15, \eta_G^2 < .007$).

Not surprisingly, the results suggest that when provided with objective feedback on trial-by-trial performance, a simplified model of reliability estimation (*Accuracy* variant) was able to dissociate advisors based on their accuracy, but did not differentiate based on advisors' agreement profile. On the contrary, human data suggest that agreement influences trust even when trial-by-trial feedback is provided. More surprisingly, a model without access to objective accuracy but endowed with metacognitive insight (*Confidence* variant) was also able to discriminate between equally agreeing but differently accurate partners. As shown in Table 2, the accurate agreeing advisor tends to agree more often than the inaccurate agreeing advisor when the participant is objectively correct (6.5 times out of 7 against 5.5 times out of 7) and less often when the participant is objectively wrong (1.5 times out of 3 against 2.5 times out of 3). Trials when participants' initial judgment

is correct are usually associated with greater confidence ratings (Fleming & Lau, 2014; Henmon, 1911; Koriatic, 2012b), thus a strategy of reliability estimation relying on confidence can exploit this covariation to detect differences in accuracy, notwithstanding equal agreement rates.

3.3 Discussion

Experiment 2 orthogonally manipulated the agreement rate and accuracy of advice presented to participants. Results show different behaviors according to whether or not participants (and the model) were provided with trial-by-trial feedback. Human participants, but not the model (*Accuracy* variant), showed an effect of agreement even when given objective feedback on the advisors performance. When feedback was removed, participants' explicit trust reports showed an effect of agreement but not of accuracy, while their implicit confidence changes showed both effects. Similarly, a *Confidence* model distinguished both dimensions of advice, namely agreement and accuracy, by scaling agreement with internal metacognitive signals. On the contrary, a *Consensus* model showed a strong preference for agreeing advisors and showed only limited evidence of being able to distinguish accurate from inaccurate advisors.

Unsurprisingly, advisor accuracy affected reported trust and influence measure in the Feedback group (Behrens et al., 2008). The finding that these participants also showed an agreement effect is, however, more surprising.

The finding seems to suggest that agreement is perceived as a strong indicator of the reliability of advice, independent of the actual underlying accuracy. A possible explanation for this effect is that participants show a confirmation bias (Nickerson, 1998) by which confirming evidence (in this case an agreement trial) is weighted more than disproving evidence (a disagreement trial). The *Accuracy* model created to describe trust formation in cases when feedback is provided by the experimenter is thus not enough to describe participants' behavior as it fails to show a mediating effect of agreement. A combination of accuracy and internal confidence could perhaps explain the mechanism giving rise to participants' actual behavior.

When feedback was unavailable, both participants' implicit (influence) and explicit (trust) behavior is dominated by the agreement rate of the advisors rather than by their accuracy. This results are in agreement with the established findings in social psychology — e.g., phenomena of naïve realism and latitude of acceptance — that people tend to discard disagreeing opinions and tend to see subjective views are more objective than others' (Liberman et al., 2012; Minson et al., 2011; Sherif et al., 1965; Shultz, Katz, & Lepper, 2001).

One possibility for the diminished effect of advisor accuracy is that the preprocessing performed on original questionnaire ratings through PCA to aggregate different questions together (see Supplementary Information) caused a loss of the information originally communicated by participants. This hy-

pothesis was discarded after obtaining similar results when including only the first trust question (Q1), which directly asked participants to rate how “accurate” each advisor was. The dominance of agreement over accuracy might have resulted in a halo-dumping effect (Clark & Lawless, 1994) whereby, when prompted to discriminate among advisors, participants used only the most accessible dimension. Nevertheless, an effect of advisor accuracy was still observed in implicit confidence change behavior—and, as predicted by the *Confidence* model—suggesting that some proximal cues to accuracy were available to them. It is unlikely that this cue was the agreement of the advice, as this dimension was orthogonal to accuracy by design. A more plausible hypothesis is that participants accumulated over time some internal uncertainty quantity (e.g. confidence) for each advisor separately.

Collectively, these results demonstrate that agreement is a crucial cue to reliability, such that its effects are seen even when objective feedback is provided. The fact that agreement was such a strong predictor of advice impact on participants’ opinions made it difficult to isolate the effect of internal confidence on advice evaluation. Experiment 3 was designed to prevent this issue and to match all advisors in terms of agreement and accuracy rates, while at the same time varying the amount of shared information and bias between the advisors and the participants. It provides a direct test of the hypothesis that confidence judgments are used to make useful inferences about others.

4 Experiment 3

Experiment 3 aimed to provide stronger evidence that subjective confidence contributes to the formation of judgments about advice reliability. In this experiment, we directly manipulated the probability that advice would agree with the participant’s initial judgment, conditional on their initial confidence in that judgment. Three advisors were created by making advisors’ judgments more or less biased towards the participant’s confidence judgments in trials when the participant’s initial judgment is correct: (1) an unbiased advisor who tended to agree with the participant’s choice about 70% of the time, independent of participant’s confidence; (2) a bias-sharing advisor who tended to agree with the participant’s initial choice particularly when the participant is confident; (3) an “anti-bias” advisor who tended to agree with the participant particularly when the participant is unsure. The labeling of the advisors in terms of bias does not reflect an actual bias manipulation, but rather reflects our aim to capture the patterns of agreement that would result from shared biases with the advisors. The manipulation is a special case of the general idea of judgment dependence among observers that was described above. This design is intended to capture a property that may hold in many real-world situations, where the opinions and choices of individuals are not independent, but instead are correlated (e.g., reflecting their shared reliance on common sources of information). Crucially, the effects of agreement-by-confidence were studied while matching the overall agreement

rate and accuracy of the different advisors.

4.1 Methods

4.1.1 Participants

50 participants were tested and divided in the two experimental groups. Due to participants failing to attend or complete sessions, numbers across groups were unbalanced with 24 participants in the No-Feedback group and 26 in the Feedback group. Participants were recruited through local advertisement and the Oxford University online recruitment platform. Prior the start of the experiment all participants signed a consent form approved by the local ethical committee. Volunteers were compensated for their time via monetary compensation or course credits.

4.1.2 Paradigm

The experiment consisted of 12 experimental blocks of 35 trials each, with each of the three advisors seen on 10 trials and with 5 null trials. The perceptual task was the same as for Experiments 1 and 2. Display parameters, input modality and stages within a trial remained unaltered. A different confidence rating scale was used in this experiment, because the 10-point scale used in Experiments 1 and 2 would not allow us to distinguish fine gradations in confidence that are needed in this design. Thus, in this experiment, participants rated their confidence on a 100-point scale (50 points per

interval, left vs. right). All participants' stimuli were titrated to converge to 70.7% decision accuracy. Two blocks of 25 trials served as the practice blocks and used a fourth practice advisor. Similar to Experiments 1 and 2, trust questionnaires were presented every two block and questionnaire data went through the same preprocessing pipeline.

4.1.3 Manipulation

The advice profiles of the three advisors were manipulated so that all advisors showed on average the same level of accuracy and agreement rate with the participants' judgments. The pattern of agreement was manipulated, however, such that the three different advisors' likelihood of agreement with the participant varied according to the participant's initial confidence, as described in Table 3 and illustrated in Figure 4.

The distribution of the participant's pre-advice confidence judgments was divided into three confidence bins: the low, middle and high confidence bin, which included 30%, 40% and 30% of the trials, respectively. On trials where the participant's pre-advice judgment was incorrect all advisors had a probability of agreement of 30% independent of the participant's confidence level. On trials where the participant's judgment was correct, however, the three advisors had different agreement patterns. An unbiased advisor had a probability of agreement of 70% independent of the participant's confidence. A bias-sharing advisor was defined as agreeing around 80% of the

time when the participant was highly confident and 60% of the time when s/he was uncertain. An anti-bias advisor was designed to agree 60% of the time when the participant was highly confident and 80% of the time when s/he was uncertain. Importantly, however all advisors had an equal chance of agreement when the participant’s pre-advice confidence fell in the middle bin (70%). This ensured that all advisors were matched across all trials in terms of average agreement rate ($0.7 * 0.7 + 0.3 * 0.3 = 0.58$) and accuracy ($0.7 * 0.7 + 0.3 * 0.7 = 0.7$).

By limiting analyses to trials within the intermediate confidence bin, we could compare advisors on trials that were matched for confidence and a priori likelihood of advice agreement. For example, asymmetric confidence changes in agreement and disagreement and larger advice influence in uncertain trials might contribute to advisor-specific distortions in the influence measure, if all trials were pooled together. The confidence reference distribution used to assign trials to the low, middle and high confidence bins was first set up on the basis of each participants confidence ratings in the first two practice blocks. The reference distribution was updated after each block to reflect the distribution of confidence judgments provided during the previous two blocks. This precaution was taken to take into account possible shifts of confidence during the course of the experiment.

	Advisors		
	Bias-sharing	Unbiased	Anti-bias
$p(\text{Agree} \text{Correct}^s, \text{Confidence}_{\text{low}}^s)$	60%	70%	80%
$p(\text{Agree} \text{Correct}^s, \text{Confidence}_{\text{mid}}^s)$	70%	70%	70%
$p(\text{Agree} \text{Correct}^s, \text{Confidence}_{\text{high}}^s)$	80%	70%	60%
$p(\text{Agree} \text{Incorrect}^s)$	30%	30%	30%
$AUC(IG_e)$	14.63	14.74	15.78

Table 3: Experiment 3 advisors’ profiles. Agreement probability of different advisors is manipulated conditional on the participant’s pre-advice confidence and accuracy. This manipulation allowed to create three different advisors who were matched in terms of agreement rate and accuracy, but who differed in terms of information as quantified by $AUC(IG_e)$, namely the area under the expected information gain curve (see Supplementary Information for details).

Although average accuracy and agreement rates were matched across advisors, the information that participants could gain by receiving advice was not. In the present experiment, computing advisor’s information gain from prior knowledge of their conditional probability was not possible as this was dependent on each participant’s individual confidence distribution. However, according to a Bayesian normative model, greater changes along the confidence scale should take place when subjective initial confidence is low, and lower shifts should be observed when initial confidence is high. The anti-bias advisor is in this respect in the good position of giving supporting evidence (agreement) when it is needed the most (low subjective initial confidence). The bias-sharing advisor on the contrary tends to agree when their judg-

ments are less effective (high subjective initial confidence). The intuition that the anti-bias advisor was the most informative was confirmed using a numerical simulation based on the nominal confidence updates of an ideal Bayesian observer performing the task (see Supplementary Information and Figure S1).

Figure 4: Experimental manipulation of Experiment 3. The probability of advisor’s agreement conditional on participant’s accuracy and confidence was manipulated so that the three advisors differed in their pattern of agreement (i.e., bias) despite being equal on average agreement rate and accuracy rate. Continuous lines represent expected agreement rates, dashed lines represent empirical data pooled across the two feedback groups.

Following Experiment 2, we expected different patterns of results to emerge from feedback and feedback-free conditions. Specifically, based on

the *Confidence* model, we predict that when feedback is not available, trust reports and influence measure will favor the bias-sharing advisor over the other advisors. This follows from the fact that both participants and the *Confidence* model will experience more high-confidence agreements with the bias-sharing advisor compared to the non-bias sharing one. This will happen even if advisors are matched in terms of accuracy and agreement, because participants and the *Confidence* model will learn the association between internal metacognitive states (like confidence) with external contingencies (like agreements/disagreements). Neither the *Consensus* or the *Accuracy* models are expected to distinguish among advisors. When objective feedback is provided, we might expect that participants will not distinguish between advisors given that they are all equally reliable ($\sim 70\%$ accuracy).

The same measures of interest adopted in the first two experiment—namely the explicit trust measure and the implicit advice influence measure—were adopted here to compare how different advisors’ profiles were perceived by the participants. Again, the three model variants were applied to the new data.

4.1.4 Exclusion criteria

An exclusion criterion based on staircase convergence was set so to exclude all participants who showed random guessing. Application of this criterion resulted in the exclusion of one participant from the Feedback group and one participant from the No-Feedback group, leaving a total of 25 and 23

participants in these groups, respectively. Average difficulty parameter d was 9.98 ± 2.82 (pooled data).

4.2 Results

The following analyses were performed to evaluate whether advisor types differed according to the measures of explicit trust and implicit opinion change. Furthermore we were interested in knowing whether this effect was modulated by the presence of feedback. Feedback acted as a between-participants factor while advisor type acted as a within-participants factor in the following mixed-design ANOVAs. For within-participants variables, degrees of freedom were corrected for violations of sphericity according to the Greenhouse-Geisser procedure, with epsilon values reported as appropriate.

4.2.1 Trust ratings

A mixed-design ANOVA was run on explicitly rated trust scores from the end-of-block questionnaires, with Feedback as a between-participants factor and Advisor Type as within-participants factor (Table 1 in Appendix). This analysis revealed no significant main effect of Advisor Type ($F(2, 92) = 1.66, p = .19, \eta_G^2 = .03, \epsilon = 0.99$) or Feedback ($F < 1$), but a significant interaction between these factors ($F(1, 92) = 6.64, p = .002, \eta_G^2 = .12, \epsilon = 0.99$). The results suggest that although different advisors were similarly trusted on average, the effect was modulated by the presence of feedback. Figure 5A shows how the effect of advisor was influenced by the presence of feedback.

Figure 5: Experiment 3 - trust ratings and influence of human participants and simulations. (a) The effect of advisor type on explicitly reported trust, separately for the Feedback and No-Feedback condition. Error bars represent s.e.m. (b) The effect of advisor type on the influence measure, divided by the two feedback conditions. Error bars represent s.e.m. (c) Average θ -values, representing the model's trust in different advisors, differ depending on what pieces of information the model has access to. Error bars represent s.e.m.

Separate within-groups analyses were carried out to understand differences among advisors present in each experimental group. For the group of participants receiving trial-by-trial feedback, a one-way repeated measures ANOVA (Figure 5A) revealed a significant effect of Advisor Type ($F(2, 48) = 4.90, p = .01, \eta_G^2 = .16$), with rated trust being highest for the anti-bias advisor, intermediate for the unbiased advisor, and lowest for the bias-sharing advisor. Follow-up pairwise comparisons indicated that the bias-sharing advisor was significantly less trusted than the anti-bias advisor ($t(24) = 3.09, p = .004, d = 1.07, B_{H[0,1]} = 47.34$), with no reliable differences observed between the anti-bias and the unbiased advisors ($t(24) = 1.60, p = .12, d = 0.51, B_{H[0,1]} = 1.67$) or between the unbiased and the bias-sharing advisors ($t(24) = 1.56, p = .13, d = 0.56, B_{H[0,1]} = 1.83$). This pattern, with a high level of trust in the anti-bias advisor and lower in the bias-sharing advisor was not predicted (a priori, or by the models as described below), but is in accordance with the pattern of information gain of each advisor (Table 3), suggesting that participants' reported trust was sensitive to the information carried by each advisor, rather than their simple accuracy. Similar results are reported below for implicit trust, suggesting this pattern is robust to different trust elicitation methods.

For trust ratings from the No-Feedback group, a repeated measures one-way ANOVA also revealed a significant difference among Advisors ($F(2, 44) = 3.56, p = .03, \eta_G^2 = .13$), but the pattern observed was very different. As predicted, rated trust was highest for the bias-sharing advisor when feedback was

absent, numerically so compared with the anti-bias advisor ($t(22) = 1.48, p = .07, d = 0.53$, one-tail, $B_{H[0,1]} = 1.74$) and significantly so compared with the unbiased advisor ($t(22) = 2.81, p = .005, d = 0.98$, one-tail, $B_{H[0,1]} = 22.42$). The anti-bias and the unbiased advisors did not significantly differ from each other ($p > .1, d = 0.41, B_{H[0,1]} = 1.13$), but rated trust was, unexpectedly, numerically higher for the anti-bias observer than the unbiased advisor.

4.2.2 Influence

Similar patterns of trust were apparent as measured implicitly via advisor influence (confidence change following advice). A mixed-design ANOVA on measured influence with Feedback as a between-participants factor and Advisor Type as a within-participants factor (Table 2 in Appendix), revealed no significant main effect of Feedback ($F(1, 46) = 1.36, p = .24, \eta_G^2 = .01$) nor Advisor ($F(2, 92) = 1.12, p = .33, \eta_G^2 = .009, \epsilon = 0.77$), but a significant interaction between the two ($F(2, 92) = 4.80, p = .01, \eta_G^2 = .03, \epsilon = 0.77$), indicating that the presence of feedback modulated the influence that different advisors had on participants' opinions (Figure 5B).

When looking at influence in the Feedback condition only, a one-way ANOVA showed a marginally significant effect of Advisor Type ($F(2, 48) = 2.88, p = .06, \eta_G^2 = .05$). Planned comparisons (two-tailed t-tests) showed that the bias-sharing advisor was less influential than both the anti-bias advisor ($t(24) = 1.98, p = .05, d = 0.54, B_{H[0,5]} = 3.71$) and the unbiased advisor ($t(24) = 2.26, p = .03, d = 0.44, B_{H[0,5]} = 6.56$). The anti-bias ad-

visor was numerically more influential than the anti-bias advisor, but this difference was not reliable ($t(24) = 1.10, p = .28, d = 0.26, B_{H[0,5]} = 1.47$).

A one-way ANOVA on influence measured in the No-Feedback condition revealed a significant effect of Advisor ($F(2, 44) = 3.25, p = .04, \eta_G^2 = .03$). Planned comparisons showed that the bias-sharing advisor was significantly more influential than the unbiased advisor ($t(22) = 2.63, p = .007, d = 0.46$, one-tail, $B_{H[0,5]} = 13.62$) and numerically more influential than the anti-bias advisor ($t(22) = 1.46, p = .07, d = 0.29$, one-tail, $B_{H[0,5]} = 2.06$). No significant difference was found between the unbiased and the anti-bias advisors ($p > .1, d = 0.16, B_{H[0,5]} = 1.07$), but the direction of the difference was the same as for explicitly rated trust, with the anti-bias advisor somewhat more influential on participants' decisions and confidence than the unbiased advisor. Thus, these results seem to suggest that the bias-sharing advisor was more influential than the other two when trial-by-trial feedback was not available.

In summary, as in Experiment 2, we find that participants' perception of advice reliability varies systematically according to whether trial-by-trial feedback is present or absent when advice is correlated with their own initial decisions. In particular, trust in advisors follows their relative informativeness when feedback is provided, but when feedback is absent participants tend to trust in and be influenced by advisors who share their judgment biases (i.e., in advisors who agree with their confidently held judgments).

4.3 Model results

The following simulations and analyses were performed to explore the differing patterns of trust across advisors that are expected from simple models of estimating advisor reliability. Differences in θ -values observed among the three model variants show how same experiences of stimuli and advice can give rise to different trust judgments depending on the information an observer has access to. Simulations are also useful in the context of this experiment, because they allow us to check that our designs were controlled as intended (e.g., for agreement rates across advisors) even though we had less precise control over conditions because counterbalancing depended on an evolving estimate of participants confidence distributions. Figure 5C shows the pattern of results that the three model variants produce.

Both when the model has access to trial-by-trial feedback (*Accuracy* variant), and when it only has access to past agreement (*Consensus* variant), no significant effect of Advisor is observed ($F(2, 94) < 1.70, p > .18, \eta_G^2 < .02$), nor is there a difference between the bias-sharing and the anti-bias advisor. These patterns are expected because the three advisors were matched for accuracy and agreement rates by design in this experiment. On the contrary, a *Confidence* variant—which uses metacognitive information and past agreement (but lacked access to trial-level feedback)—showed a significant effect of Advisor ($F(2, 94) = 7.95, p < .001, \eta_G^2 = .10$). Specifically, simulated trust was higher for the bias-sharing advisor than the anti-bias advisor ($t(47) = 3.54, p = .001, d = .74, B_{[0,.05]} = 114.64$), and higher for the unbiased advisor

than the anti-bias advisor ($t(47) = 2.99, p = .004, d = .57, B_{[0,.05]} = 18.16$). Simulated trust was higher for the unbiased than the bias-sharing advisor, but this difference was not reliable ($t(47) = 1.21, p = .22, d = .25, B_{[0,.05]} = 0.41$). These findings indicate that by accessing metacognitive signals (as provided in the model by participants' confidence ratings) the model was able to discriminate among different advisors. This model correctly predicts greatest levels of trust in a bias-sharing advisor, but also predicts lowest levels of trust in anti-bias advisors, whereas our experimental participants expressed (numerically) lowest levels of trust in unbiased advisors.

Post-advice confidence correlations. As described above, our main analyses for each experiment focused on qualitative predictions arising from different strategies for inferring advisor reliability. Collectively, the results are consistent with the hypothesis that people use their internal sense of confidence in making these inferences—showing sensitivity to advisor accuracy in the absence of objective feedback, even when advisors are matched for agreement rate, and developing differing patterns of trust when advisor agreement rates vary with their own expressed decision confidence. Our final analysis of the empirical data attempted a more quantitative comparison of model predictions, specifically focusing on the whether the *Consensus* or *Confidence* models better predicted post-advice confidence ratings across trials for the participants in the No Feedback conditions of Experiments 1-3.

For this analysis we used Bayes rule to infer the trial-by-trial post-advice

confidence ratings that each variant would express given a participant’s expressed pre-advice confidence and advisor agreement (see Supplementary Information). The within-participant correlation between participants’ post-advice confidence and model’s post-advice confidence was computed for each experiment, for No-Feedback groups only. Second-order statistics were performed to test, across experiments, which variant was more strongly correlated with the human data. A 2x2 ANOVA on correlation coefficients with Model (*Consensus* vs. *Confidence*) and Experiment as factors showed that the *Confidence* variant was significantly more correlated with participants’ responses than the *Consensus* variant ($F(1, 22) = 8.18, p = 0.009, \eta_G^2 = 0.0049$). No significant effect of experiment nor interaction between the two were found ($F(2, 44) < 1.3, p > .25$), suggesting that, across experiments, the *Confidence* model’s post-advice confidence more strongly co-varied with participants’ true post-advice responses. As a check for the soundness of this model comparison method, the same 2x2 ANOVA was run on the correlation coefficients between participants’ post-advice confidence and the model’s post-advice confidence predictions, after randomly shuffling trials within each participant. This operation should ensure that any advantage of the *Confidence* variant over the *Consensus* variant is not due to unspecific factors (like being overall more conservative in updating confidence), but rather to trial-level variability. After reshuffling, the *Consensus* and *Confidence* variants were not significantly different from each other ($F(1, 22) = 1.95, p = 0.17, \eta_G^2 = 9.54e - 04$), corroborating our conclusions.

4.4 Experiment discussion

Experiment 3 showed again that systematic differences in trust emerge as a function of feedback when advisors' judgements are non-independent from the advisee, and further demonstrated a strong influence of our agreement-in-confidence manipulation on trust. Here, the presence of feedback partly reversed the pattern of trust ratings and influence measure that was observed when trial-by-trial feedback was unavailable: When objective feedback was provided, people trusted more and were more influenced by an advisor who more frequently agreed with them when they themselves were unsure (vs. less frequently when they were sure), and showed less trust in an advisor who tended to agree with them in decisions in which they were already sure. On the contrary, when objective trial-by-trial feedback was removed, this pattern partly reversed such that trust and influence were greatest for the advisor who tended to agree with participants' confidently made judgments. Participants expressed less trust in an advisor who tended to disagree with their confidently made judgments and even more so in an unbiased advisor whose tendency to agree or disagree did not vary with the participants initial decision confidence.

Participants in the Feedback group showed an effect of advisor type even though the feedback indicated that all advisors were equal in terms of both agreement rate and accuracy. These findings contrast somewhat with the results of Experiment 2, where ratings of trust and measures of advice influ-

ence were dominated by agreement rates when these varied markedly across advisors. Evidently, though agreement appears to be a critical determinant of trust (with or without feedback), participants remain sensitive to other dimensions of advice. Here, they appeared sensitive to the informational value (“information gain”, see Supplementary Information) of the advice and not only to its reliability *per se*: They trusted and were more influenced by advisors whose judgments were less redundant with their own, in particular showing greatest trust in the advisor whose advice was most likely to improve their performance (by reinforcing, through agreement, initial judgments that were correct but most likely to waver due to low initial confidence).

As predicted, participants in the No Feedback group showed greatest trust in advisors who tended to agree with their own confidently expressed judgments. However, more surprisingly, participants did not show least trust in the anti-bias advisor who agreed with them more frequently when their initial judgment was made with low confidence. If anything, trust was lower in the unbiased advisor. The difference was not reliable and hence must be interpreted with caution, but we note a possible link to an existing proposal about how expertise is assessed in the absence of feedback. Thus, according to Weiss and Shanteau (2003), scarcely differentiated judgments across different observations tend to be indicative of lower expertise. In the present experiment, the unbiased advisor was characterized by this lack of differentiation, in that her advice agreement rate was constant as a function of participants’ initial decision confidence, in contrast to the two other advi-

sors. Nonetheless, regardless of the explanation of this unexpected result, the findings of Experiment 3 demonstrate participants’ sensitivity to advice in relation to their own confidence in their judgments, and learn differentiated patterns of trust accordingly.

5 Agent-based modeling: network-level impacts of individuals’ trust strategies

The experimental and simulation results described above demonstrate that the absence of trial-level objective feedback does not preclude agents from inferring the trustworthiness of other agents. Simple heuristics based on agreement and decision confidence can be good proxies for objective feedback. However, our findings also highlight that systematic deviations between feedback and feedback-free scenarios—and, correspondingly, systematic biases in trust formation—can be observed when judges’ and advisors’ judgments are correlated (e.g., due to shared biases or common sources of information). Thus far we have focused on the situation of an individual decision maker learning about the quality of advice from different sources. To extend this work, we explored how these effects might play out in more complex, multi-actor situations using agent-based modeling.

Using an agreement (or agreement-in-confidence) heuristic can be an effective strategy to infer partners’ reliability when more objective cues are not available or costly to acquire. However, predictable distortions are expected

when independence is not preserved among judges. If patterns of trust vary depending on feedback availability, we might expect different macro-level patterns of trust in networks where feedback is available rather than difficult to obtain. Furthermore, the presence of bi-directional information channels between agents (as opposed to judge-advisor systems) may lead to clustering of people who share the same biases, increasing the polarization of their views (i.e., confidence magnitude) as well as producing phenomena of *groupthink* (Janis, 1972; Turner & Pratkanis, 1998). These effects parallel important observations about human social networks. Informational echo-chambers—whereby individuals with similar characteristics tend to form clusters that are impenetrable to external information—have been often described in social sciences (Jasny et al., 2015; Sunstein, 2001). When exposed to competing opinions on social media, people belonging to the same cluster may prefer, consume and share within-cluster information more than between-cluster information (Bessi, 2016; Del Vicario et al., 2016), leading to a spurious consensus effect (Yaniv et al., 2009). Similarly, people in a group are known to prefer discussion about shared information than private information (Schulz-Hardt, Brodbeck, Mojzisch, Kerschreiter, & Frey, 2006; Stasser & Titus, 2003). If so, the present ideas might provide a normative account of the origin of these effects that otherwise seem suboptimal: According to our findings, preference for agreeing opinions (particularly those that agree with our confidently held views) is an adaptive strategy for estimating advice reliability in the absence of an objective standard, but can lead to these systematic biases in trust

when opinions and judgments are non-independent.

To explore this hypothesis, we used agent-based modeling to explore how the manipulation of feedback availability affects network structure and information flow. Agent-based models are useful tools to study phenomena of emergence from simple interacting parts. Manipulating agents’ access to metacognitive signals can thus shed light on how the cognitive structure of agents can affect the structure of a network (Epstein, 2013). Giving agents choice over what information to sample—for example, if reliability estimates were used to guide advisor choice (Aiello et al., 2012; Denrell, 2005)—might further remove agents from being exposed to useful but uncorrelated information (cf. our “anti-bias” advisor in Experiment 3).

In the following section, we test how macro-scale patterns of trust and beliefs are affected by the presence or absence of feedback in networks of agents employing the cognitive strategies described so far in this work. We expect that (a) When judgments in the population are independent, agreement-based heuristics are useful in reliably approximating true underlying expertise; (b) when judgments are correlated—e.g., due to the presence of shared biases in the population—agents sharing similar biases will tend to segregate due to the use of a confidence-agreement heuristic in feedback-free (but not feedback-rich) scenarios; (c) The use of trust to discount advice will be beneficial to agents’ accuracy when judgments are independent but not when they are correlated.

5.1 Model description

An agent-based model was programmed using NetLogo (Wilensky, 1999) and is available at https://github.com/chri4354/trust_formation_without_feedback. It simulates development of trust among agents as they make decisions, and revise these decisions on the basis of learning about the opinions of others, in a series of simple binary choices: Is a stimulus from category A or B, as a generalised case of the specific dot discrimination task used in our experiments. Of interest was how trust among agents was affected by the relative quality of their decisions, the availability of feedback and, crucially, the degree to which agents shared biases in their decision processes. Our models formalised these biases as differences in prior (base rate) expectations about the likelihood of A or B being the correct answer, which via Bayesian belief updating will bias the interpretation of incoming information in the decision process (as an analogue of what might be expected to happen in everyday decisions as a function of our political leanings, musical preferences, tastes in wine, etc.).

The model was initialised as a fully connected directed network of N agents. A directed edge from agent i to agent j represents the trust $\theta_{i,j}$ that i has in j 's opinions. We simulate agents on a lattice network performing repeated binary A/B decisions, receiving advice from other agents, inferring their reliability and updating their own initial decisions. We let the simulation run for a 1000 steps. A signal s with strength S was drawn from a uniform distribution between $-\frac{S}{2}$ and $+\frac{S}{2}$. This represents the decision quan-

tity to estimate (e.g., difference in dots or true state of the world). The task of each agent was to determine if s was positive (event A) or negative (event B). Each agent estimated the posterior probability of A given the perceptual information generated by s as follows:

$$p'(A) = p(A|E_p) = \frac{p(A)E_p}{p(A)E_p + p(\bar{A})\bar{E}_p} \quad (1)$$

$$E_p = L(s + \mathcal{N}(0, \sigma)) \quad (2)$$

where $p(A)$ is the prior probability of observing A s before seeing any stimulus, L is a logistic sigmoid mapping from sensory evidence to probability; E_p is the perceptual evidence resulting from such mapping and $\mathcal{N}$ is independent individual perceptual Gaussian noise with mean 0 and standard deviation σ . Bars represent complement probability. Each agent's perceptual noise (and thus accuracy) was manipulated by varying the noise parameter σ . Each agent's bias was manipulated by varying the value $p(A)$. Agents' confidence was represented as the distance from the uncertainty point 0.50:

$$C = .50 + |p'(A) - .50| \quad (3)$$

Trust, represented by the network's edges, was initialized to 0.50 for every agent and updated after social interaction. After making a judgment, agents selected one other agent to interact with either at random (random sampling) or proportionally to their trust (biased sampling). Agents then updated their

initial judgment $p'(A)$ as follows:

$$\hat{p}(A) = p(A|E_s) = \frac{p'(A)E_s}{p'(A)E_s + p'(\bar{A})\bar{E}_s} \quad (4)$$

where E_s represents social evidence and is obtained from the advisor's judgment either by taking the advisor's raw judgment $p'(A)$ (without advice discounting) or by discounting the advisor's judgment proportionally to the agent's trust in the advisor (with advice discounting). Advice discounting consisted in a linear regression toward the uncertainty point 0.50. (equation S20). After updating their judgments, each agent i updated its trust judgments (i.e., outward edges θ_i) based on the available information about other agents. If feedback is available, the agent updates its current trust in agent j by virtue of a delta rule in the form:

$$\theta_{i,j}^{t+1} = \theta_{i,j}^t + \alpha(F_j - \theta_{i,j}^t) \quad (5)$$

where F_j is the accuracy of agent j and α is a learning rate set to 0.1. If feedback is not available on the contrary, the agent replaces F with $\hat{F}$, or the *estimated* partner's accuracy. $\hat{F}$ was calculated using the agreement or agreement in confidence heuristics described above. In our simulations, we assessed the effect of feedback availability as it varied parametrically, from being available after every decision (i.e., p-feedback = 1.0), available after only some decisions (i.e., $0 < \text{p-feedback} < 1$), or never available (i.e., p-feedback = 0), rather than the simpler case of feedback presence/absence

that we studied experimentally above.

The emergence of trust patterns when using agreement-based heuristics can be expected to track true accuracy when judgments are independent but generate clustering of populations when judgment correlations emerge within such populations. We defined a network’s clustering coefficient as the ratio between average trust toward agents who initially share the same bias (in-group trust) and total average trust: $\bar{\theta}_{in-group}/(\bar{\theta}_{in-group} + \bar{\theta}_{out-group})$. A ratio of 0.5 represents no preference (i.e., no difference in trust) toward agents sharing the same bias, while a ratio greater than 0.5 represents a preference toward agents sharing the same initial biases. We test how network clustering is shaped by the presence or absence of objective feedback and show that bias-specific segregation arises only in the absence of feedback.

Finally, once bias-specific segregation is established, we ask whether such clustering remains stable. In particular, after 500 iterations we allow agents to dynamically change their original bias as a function of experience. For example, it is known that the bias observed in people performing binary judgments is influenced by their recent history of decisions (Akaishi, Umeda, Nagase, & Sakai, 2014; Zylberberg, Wolpert, & Shadlen, 2018). In the present context, if an agent systematically reports “A” but receives negative feedback, they should reduce their bias by decreasing their prior probability $p(A)$. Similarly, when feedback is absent, an agent who systematically reports “A” but finds themselves, after interacting with other agents, believing that B s are more frequent than expected, should reduce their bias towards A s. Con-

versely, bias should get stronger if the social contexts reinforces it (although see (Bail et al., 2018)). We modelled bias update with a delta rule:

$$p(A)^{t+1} = p(A)^t + \alpha(I^t - p(A)^t) \quad (6)$$

where I^t is an indicator variable that represents the final belief in the event A . When objective feedback is available, I takes the value of 1 if an event A occurred and 0 otherwise. When feedback is not available, I is set to the discrete or continuous final subjective belief in the event A . In the following section, we show the results obtained when a discrete final belief is used in the absence of feedback:

$$I = \begin{cases} 1, & \text{if } \hat{p}(A) \geq 0.5 \\ 0, & \text{otherwise} \end{cases}$$

Similar results were obtained setting I to the continuous belief $\hat{p}(A)$.

5.2 Results

We first tested the hypothesis that agreement-based heuristics are adaptive in situations of independent judgments because they reliably track others' expertise without the need to rely on external forms of feedback. To this end, we set all agents' initial bias $p(A)$ to 0.50 and the probability of feedback to 0. We then created two sub-populations of agents that orthogonally varied in their overall judgment accuracy (modeled as the degree of perceptual noise,

with decreasing accuracy as a function of increasing noise). We then calculated the average trust toward a target sub-population (here, Population 2). The results show that average trust in Population 2 agents correctly tracks their underlying perceptual noise: Trust in Population 2 is highest when their judgments are more accurate (i.e., when their perceptual noise is low: upper pixels in Figure 6A) and decreases as their perceptual noise increases. Interestingly, the effect is not linear and we observe an interaction between the accuracy of Population 1 and the accuracy of Population 2 (Figure 6A, and Figure S2), mirroring the simple numerical analysis shown in Figure 1B: In the absence of feedback, poor performers in Population 1 fail to distinguish the accuracy of others, evident as a narrower range of trust values in rightward pixels in Figure 6A [cf. (Kruger & Dunning, 1999)]. Nevertheless, overall, we see that trust in Population 2 systematically tracks its level of performance, even when agents lack any objective standard against which to judge their own decisions or those of others.

We explored these simulation results further to identify the benefits of learning differentiated patterns of trust according to confidence-weighted agreement. To this end, Figure 6B plots the impact on Population 1 agent accuracy of learning the opinion of Population 2 agents, as a function of their respective decision accuracies (noise levels), specifically in the case of agents who do not discount advice according to their level of trust in advisors. As shown in Figure 6B, Population 1 agents generally improve the accuracy of their decisions by receiving advice from Population 2 members (most pixels

in Figure 6B are light-colored), but this benefit becomes an accuracy cost if Population 2 agents are much worse at the task than Population 1 agents (dark pixels in the lower left of Figure 6B; cf. Bahrami et al., 2010). However, this cost of receiving bad advice is greatly mitigated when advice is discounted according to trust learnt through confidence-weighted agreement (Figure 6C, light-colored pixels in the lower left corner). Thus, learning differentiated patterns of trust enables agents to benefit from advice when it is useful, and downweight or ignore it otherwise.

Figure 6: A: average trust shown in agents belonging to population 2. Trust in these agents increases as the noise affecting them decreases. Interestingly, trust formation is affected by perceptual noise of the observed agent(s) as well as perceptual noise of the observing agent(s) [cf.(Kruger & Dunning, 1999)]. B: impact of advice for Population 1. Pixels’ color represent the difference between post and pre-advice accuracy. C: the use of a trust-based advice discounting strategy increases post-interaction accuracy of accurate agents’ but not inaccurate ones.

Our second hypothesis was that when the true state of the world is difficult to discern (e.g., when objective feedback is absent, signal strength is weak or perceptual noise is large), bias-specific segregation can arise. As described above, we quantify this segregation in terms of a *clustering* metric, which indicates the degree to which trust in in-group members exceeds trust

in out-group members (indicated by clustering > 0.5).

In this simulation, we varied the bias $p(A)$ shown by two subpopulations of equal size and equal average accuracy (equivalent average perceptual noise). For each individual, we randomly draw $p(A)$ from a uniform distribution in the range $[0, 0.50]$ (Population 1) and $[0.50, 1]$ (Population 2): That is, the two populations differ in their base rate estimates of the relative likelihood of the two outcomes (A and B), with Population 1 biased towards judging events as “B”s and Population 2 biased towards judging events as “A”s. On each iteration, agents interacted with another randomly selected agent and updated their trust according to equation 5. Figure 7 shows the average clustering shown after 1000 iterations as a function of signal strength (from low in the upper panels to high in the lower panels), perceptual noise (variation along the x-axis of each panel), feedback probability (y-axis), and whether agents update their bias depending on recent experience (left column: without bias update; vs. right column: with bias update).

Figure 7: Clustering as a function of signal strength, probability of feedback and noise. Each pixel in each panel indicates the average clustering value (over 20 simulations), computed as described in the main text, after 1000 iterations of decision, advice and update in a network of agents with fixed stimulus strength, perceptual noise and feedback probability parameters. Right and left columns distinguish simulations with and without bias update respectively. Rows from top to bottom represent increasing signal strengths.

Average clustering decreased as the probability of receiving objective feedback increased, evident in Figure 7 as darker pixels in the lower parts of each panel. When objective feedback is consistently presented, trust is determined by an advisor’s objective accuracy rather than their agreement with the agents’ own judgments on each trial. Trust therefore converges to the same value for all agents, regardless of group, reflecting the fact that underlying accuracy (dependent on the level of perceptual noise) was equal on

average across the entire population. A similar principle explains why clustering also decreased as signal strength increased—evident in Figure 7 as darker pixels on average in the lower panels of the figure than in the upper panels: When signal strength is high, agents’ decisions are dominated by external evidence rather than their prior expectations (as per Equation 1 above), thus agreement levels are high with both in-group and out-group members, and agents learn similar levels of trust in others across the whole population. On the other hand, clustering tends to decrease as perceptual noise increases (i.e., as agents’ average accuracy decreases), evident in Figure 7 as darker pixels to the right-hand side of each panel. This observation reflects the fact that as agents’ accuracy tends towards chance levels, their likelihood of agreement with others will correspondingly tend towards chance (regardless of others’ accuracy) (Figure S4), and they therefore lose the ability to distinguish advisor accuracy in the absence of an objective standard (cf. Figure 1B), and learn low levels of trust in both in-group and out-group members (Figure S5).

On the contrary, agents in each sub-population tended to trust in-group members more than outgroup members as a function of decreasing feedback availability and signal strength, with this dependence disrupted under conditions of increasing perceptual noise. When signal strength is low, agents’ decisions are more strongly influenced by their prior expectations (as they should be, according to Bayes theorem). As a consequence, when feedback is infrequent or absent, agents’ trust is dominated by agreement over objective

accuracy, and clustering with like-minded (bias-sharing) in-group members results, a function of increased agreement rates among members of homogeneous populations. This finding suggests that a simple delta rule of trust update can perform very differently depending on what source of information is used to supervise learning. The circular nature of using one’s own belief to estimate others’ accuracy produces higher clustering and segregation in scenarios where feedback is infrequent or absent, and decisions are often ambiguous. Meanwhile, our results are also in accordance with recent studies in ecology and network science, which suggest that the presence of noise is beneficial to group decision making to the extent that it reduces the risk of getting stuck on local optima (Couzin et al., 2011; Kao & Couzin, 2014; Shirado & Christakis, 2017): Here, high levels of noise disrupt the formation of clustering according to pre-existing biases (albeit by reducing trust in both in-group and out-group members alike) (Figure S5).

Our third and final hypothesis was that segregation between in-group and out-group members, once established, can resist modification and in turn affect the persistence of shared beliefs. In this simulation, agents first made decisions with fixed biases (i.e., a fixed $p(A)$ value) and learnt trust in other agents across multiple iterations, as in the previous simulation. As before, the two populations drew their initial biases from random distributions favoring either A or B as a prior ($p(A) > .5$ or $p(A) < .5$, respectively), with feedback availability and perceptual noise level fixed in each simulation run but varied across runs to explore the impact of these parameters on network

behavior. Signal strength was fixed at 1, where a wider range of behaviors was previously observed (Figure 7B). However, crucially in this simulation, after 500 iterations we allowed agents to update their bias, whereby their experience of the relative prevalence of A and B outcomes (as determined by their belief-influenced interpretation of sensory evidence, opinions of partners, and objective feedback where available) leads to modification of their estimates of $p(A)$. Of interest was the way in which the biases of the two groups evolved across interactions as a function of feedback availability, perceptual noise, and information sampling strategy.

The results are shown in Figure 8, which plots the results of separate simulations for populations of agents that select partners either at random (columns A and B) or according to their current levels of trust in other agents (columns C and D). Columns A and C plot bias (i.e., mean $p(A)$ value) across 1000 iterations of the decision task for the two groups in the population—one initially biased toward A decisions (red lines) and the other initially biased toward B decisions (green lines). Columns B and D plot the distribution of biases across members of each group after the final (1000th) iteration.

As shown in Figure 8, when feedback is always available (saturated red and green lines), decision biases rapidly diminish such that agents from both populations converge on accurate estimates of the relative likelihood of A vs. B outcomes (i.e., unbiased estimates that vary narrowly around $p(A) = 0.5$). This convergence occurs irrespective of the level of perceptual noise and the way agents select their partners, reflecting the obvious value of hav-

ing an objective standard against which to evaluate decisions. However, as feedback becomes less and less available (indicated in Figure 8 as reducing saturation of line colors), population biases are more persistent in a manner that exhibits sensitivity to both information selection strategy and levels of perceptual noise. Consider first the simulations of populations with random sampling of partners (columns A and B). When agents randomly sample their advisors, they tend to converge on accurate estimates of $p(A)$ (i.e., ~ 0.5). This occurs even when feedback is not available (as indicated by low saturation lines) and independently of noise level (as indicated by different rows). The reason is that when agents sample their advisors at random, agents are equally likely to receive advice that the correct answer is A vs. B (or in other words advice that is characterized by their same bias vs. opposite bias). However, when agents weight advice in proportion to trust, they are more strongly influenced by other agents who share their initial bias (and therefore in whom they learn more trust based on patterns of agreement, as explored in the previous simulation) (Figure S4). Thus, social information can stabilise initial group biases insofar that agents rely on this social information as feedback probability decreases (populations with lighter colored lines exhibit systematically larger biases in Figure 8 and S6). The lower panels of Figure S6 A-B show that these population biases are disrupted when agents have higher levels of perceptual noise: This effect is partly due to noise reducing the influence of prior expectations on individual agents own decisions, but is also due to noise disrupting trust in other agents who share

an initial bias (because noise will tend to reduce agreement rates with these individuals (Figures S4-S5)).

As shown in the right hand panels of Figure 8, the stabilising effect of social information is also apparent when agents select partners to receive advice from as a function of their trust in those agents (columns C and D). Under these circumstances, stable group biases can persist even when feedback probability is as high as 0.8 (in the simulation when noise is absent, Figure 8C, top panel), even if those biases are relatively small. This stabilisation is a natural consequence of preferentially seeking advice from agents who share a bias toward believing $p(A)$ is systematically greater than (red group) or lower than (green group) its true value. Perhaps even more strikingly, the simulations also reveal cases where social information leads to extreme biases when objective feedback is rarely available: In these cases (light colored lines, e.g., $p(\text{feedback}) < .4$ in the top panel of Figure 8C), group biases actually become more extreme over time. This phenomenon emerges from a positive feedback loop whereby agents who share an initial bias (of, say, believing that $p(A) > 0.5$) will tend to ask each other for advice more often, thereby compounding each others biased beliefs, and thereby increasing their tendency to seek advice from like-minded agents, thereby compounding their bias, and so on across iterations. The two groups therefore diverge to separate stable networks that exhibit extreme biases (of $p(A) \sim 0$ and $p(A) \sim 1$) and a strong tendency only to trust and seek advice from in-group members.

In this way, sub-optimal behaviors at the network-level—such as self-

sustaining biases and even extremes of polarization—can emerge from decision and learning strategies that are normatively sensible at the level of individual agents: Using decision confidence as a proxy for objective accuracy can allow agents to learn about the reliability of social information sources when feedback is infrequent or even absent. Moreover, using social information can in turn allow agents to learn about the world under these conditions of diminished feedback. However, these learning strategies exhibit important limitations when agents initial decisions are not independent but rather exhibit patterns of shared vs. distinct biases across individuals within a population.

Figure 8: Bias as a function of time, feedback probability and agents’ perceptual noise. Rows represent increasing levels of noise. Lines alpha values (saturation) represent probability of objective feedback. Columns A and B show agents’ bias evolution over time and final bias distribution when agents choose their partners at random. Columns C and D show the same graphs when agents choose their partner proportionally to their trust.

Overall, the results of the agent-based modelling show the macro-scale effects of feedback availability, use of different reliability estimation strategies (feedback vs. agreement-based), and partner-selection strategies. Contrary to judge-advisor systems—where information flows from the advisor to the judge—here we allow agents to share information bidirectionally and use this information to estimate each other’s reliability. Once again, agreement-based strategies provide an effective way to learn the reliability of social partners when feedback is not available and judgments are independent. However,

they lead to predictable patterns of clustering and trust when agents' agreement rates are inflated by their shared biases or access to correlated information. Moreover, bias updating and partner-selection strategies also affect the structure of the network by changing the segregation of individuals sharing the same bias, leading to stabilisation of networks with systematically biased or even extremely polarized beliefs about the world (in our simulations, the relative likelihood of A vs. B being the correct decision).

6 General Discussion

The present research explored how people learn the reliability of social information sources in contexts where objective external feedback is not readily available. In these scenarios, we hypothesize that metacognitive confidence provides a useful proxy for objective feedback. This is due to the fact that in many tasks, confidence provides a finely calibrated estimate of underlying accuracy (Fleming, 2016; Henmon, 1911; Koriati, 2012a; Pescetelli et al., 2016) as it represents the subjectively estimated likelihood of a correct decision being made (Aitchison et al., 2015; Fleming & Daw, 2017). Thus, people can apply this internal confidence estimate to received advice in order to estimate its reliability (as $p(\text{correct})$ in the case of agreement, and $1-p(\text{correct})$ in the case of disagreement). Experiment 1 empirical data and simulations show the power of this agreement-in-confidence heuristic in identifying subtle differences in advice quality when advice provides new independent information:

Participants (and simple models) learned to distinguish advisors of differing accuracy and confidence calibration even in the absence of objective feedback, and did so to a similar extent to participants (and simple models) who had access to objective feedback after every decision. Thus, we propose that people are able to follow internal signals (e.g., decision confidence) in a comparable manner to external signals—such as rewards or feedback—to learn about the reliability of advice and advisors according to associative learning principles (Behrens et al., 2008; Sutton & Barto, 1998).

These findings extend our understanding of the uses of metacognitive signals. The simulation results show that a model endowed with access to its own metacognitive signals can discriminate between advisors better than models that only have access to objective feedback or agreement rate. Thus, metacognitive signals like confidence judgments are useful not only for internal uncertainty monitoring (Yeung & Summerfield, 2012), information seeking (Desender, Boldt, & Yeung, 2018), cognitive control (Botvinick et al., 2001) and social coordination (Bahrami et al., 2010; Shea et al., 2014), but are also useful in evaluating the reliability of external information sources. The present findings demonstrate the potential bi-directionality of this inference process: Confidence is not only the end-product of information flow from the external stimulus to a perceptual inference, but it can feed back to help making inferences about external events. Confidence represents a probabilistic estimation that once formed can be used to infer the state of variables that did not directly generate it. Evidence seems to support this

notion, as confidence signals have been shown to parallel reward prediction error signals in feedback-free situations (Guggenmos et al., 2016; Zylberberg et al., 2018).

Our findings also extend previous work on the role of confidence in group decision making (Bahrami et al., 2010; Koriat, 2012b; Sorkin et al., 2001). Previous studies have documented the way that the confidence accompanying advice provides a useful and impactful signal of its reliability (Bahrami et al., 2010; Swol & Sniezek, 2005). Here, we explore the converse situation of how an advisee’s internal sense of confidence can modulate perceptions of advice. Thus, whereas previous studies have shown how confidence judgments from two individuals can be used to combine optimally the individual estimates of external stimulus attributes, our experiments and simulations indicate that an observer’s internal confidence about the state of the external stimulus can also be used to estimate the reliability of a social partner above and beyond what can be inferred from objective external cues. We used a traditional Judge-Advisor System paradigm (Bonaccio & Dalal, 2006; Sniezek & Buckley, 1989; Yaniv & Kleinberger, 2000) and used a perceptual decision-making paradigm to allow precise control over important variables such as strength of evidence available for the initial decision, the information carried by advice, and the manipulation of objective feedback to enable social decision-making to be understood in terms of current theories of confidence as a probabilistic computation (Pouget et al., 2016).

Experiments 2 and 3 explored crucial limitations of the agreement-in-confidence heuristic, which leads to systematic biases in trust and influence when advice is not independent of an individual’s own opinions. Much of the work in psychology in the last decades has focused on heuristics and biases affecting human judgments (Gigerenzer, 2008; Tversky & Kahneman, 1974). Given our limited cognitive capacities, we must often adapt to use approximations or “short-cuts” to find good-enough solutions to otherwise intractable problems. Heuristics are computationally cheap solutions that work well in environments characterized by the statistical properties of the environments they evolved in (Tversky & Kahneman, 1983). However, they can lead to systematic errors when these assumptions are violated. In the current work, the agreement-in-confidence heuristic provides a solution to the intractable problem of estimating the reliability of information in the absence of an objective standard. This solution works well when agreement correlates with accuracy, as is the case where initial judgment and advice are independent (Experiment 1). Thus, in Experiment 1, trust and influence patterns were little affected by the presence or absence of an objective standard. However, the process goes astray when the independence between judgment and advice is broken, as in Experiments 2 and 3 where advice was contingent on participants’ initial decision and expressed confidence. Adverse consequences emerged in terms of distinct patterns of trust and influence with and without an objective standard. The result is explained by the fact that when feedback is removed participants are only left with their own initial probabilistic

estimate. Our findings are in agreement with studies showing that people often tend to ignore the correlation structure of advisors and heavily rely on agreement among information sources (Yaniv et al., 2009). [Any attempt to use this estimate to attribute feedback to their advisors is cursed by the error characterizing their own original judgment. As an example, imagine you are certain that you are drinking a very expensive wine at a blind wine tasting event. You will likely judge anybody who says otherwise as incompetent in wine matters. In the (not so unlikely) event that you are mistaken, you could keep thinking that a potentially brilliant sommelier is only a novice, with potentially negative consequences for you in terms of reputation and future wine choices.] SCRATCH THIS OFF?

Our results were generally in line with the idea that we trust advisors according to their objective accuracy when feedback is present and agreeing-in-confidence advisors when it is absent. However, there were some notable (and surprising) exceptions. First, advisor agreement rate had an effect over and above accuracy even when feedback was available in Experiment 2. This finding suggests that multiple strategies might idiosyncratically converge to an estimate of advice reliability, with people relying on sometimes-useful cues (like agreement) even if they are strictly redundant in the current context. Second, people seem to value information over pure accuracy, as shown by participants in the Feedback group of Experiment 3, who preferred an advisor who agreed with them more frequently when they themselves were

correct but low in confidence (the anti-bias advisor) over equally accurate but less informative advisors. This result is in agreement with literature on judge-advisor systems showing people’s preference for advisors with unshared information (Van Swol & Ludutsky, 2007). Finally, a result with a less clear explanation is that participants in Experiment 3 did not prefer the neutral to anti-bias advisor when feedback was absent. If anything, they tended to prefer the latter. We speculate that this unexpected result can be explained in light of the fact that, in domains where feedback is removed, expertise can still be estimated by comparing the classification variability observed for similar stimuli with the classification variability observed across different stimuli (Weiss & Shanteau, 2003). One possibility is that participants noticed the lack of variability in agreement rates (across different trial difficulties as experienced subjectively and expressed in terms of decision confidence) of the neutral advisor and consequently discounted this advisor.

7 Conclusions

The current work aimed to show that confidence is a valuable attribute of someone’s judgment in social decision-making. It helps others discriminate when one is more likely to be correct (and thus value their contribution) but also helps a decision-maker to make consistent judgments about others irre-

spective of feedback availability. However, this potentially adaptive solution to an intractable problem of learning in the absence of feedback can backfire when judgments from different observers are not independent, leading to systematic biases in trust and influence.

References

- Aiello, L. M., Barrat, A., Schifanella, R., Cattuto, C., Markines, B., & Menczer, F. (2012). Friendship prediction and homophily in social media. *ACM Transaction on the Web*, 6(2), 9. Retrieved from <https://hal.archives-ouvertes.fr/hal-00718085>
- Aitchison, L., Bang, D., Bahrami, B., & Latham, P. (2015). Doubly Bayesian Analysis of Confidence in Perceptual Decision-Making. *PLoS Computational Biology*, 11(10), 1.
- Akaishi, R., Umeda, K., Nagase, A., & Sakai, K. (2014, 1). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron*, 81(1), 195–206. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/24333055> doi: 10.1016/j.neuron.2013.10.018
- Armstrong, J. S. (2001). *Principles of forecasting: a handbook for researchers and practitioners*. Norwell, MA: Kluwer Academic.
- Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. (2012a, 2). Together, slowly but surely: the role of social interaction and feedback on the build-up of ben-

- efit in collective decision-making. *Journal of experimental psychology. Human perception and performance*, 38(1), 3–8. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3268462&tool=pmcentrez&rendertype=abstract> doi: 10.1037/a0025708
- Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. (2012b, 4). What failure in collective decision-making tells us about metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594), 1350–1365. Retrieved from <http://rstb.royalsocietypublishing.org/cgi/doi/10.1098/rstb.2011.0420> doi: 10.1098/rstb.2011.0420
- Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010, 8). Optimally interacting minds. *Science (New York, N.Y.)*, 329(5995), 1081–5. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/20798320> doi: 10.1126/science.1185718
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., ... Volfovsky, A. (2018, 9). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37), 9216–9221. Retrieved from <http://www.pnas.org/lookup/doi/10.1073/pnas.1804840115> doi: 10.1073/pnas.1804840115
- Bakeman, R. (2005). Recommended effect size statistics for repeated measures designs. *Behavior Research Methods*, 37(3), 379–384.

- Bang, D., Aitchison, L., Moran, R., Castanon, S. H., Rafee, B., Mahmoodi, A., ... Summerfield, C. (2017). Confidence matching in group decision-making. *Nature Human Behaviour*, 1(0117), 1–7. doi: 10.1038/s41562-017-0117
- Behrens, T. E. J., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. S. (2008, 11). Associative learning of social value. *Nature*, 456(7219), 245–9. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2605577&tool=pmcentrez&rendertype=abstract> doi: 10.1038/nature07538
- Bessi, A. (2016, 12). Personality traits and echo chambers on facebook. *Computers in Human Behavior*, 65, 319–324. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0747563216305817> doi: 10.1016/j.chb.2016.08.016
- Boldt, A., & Yeung, N. (2015, 2). Shared Neural Markers of Decision Confidence and Error Detection. *Journal of Neuroscience*, 35(8), 3478–3484. Retrieved from <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.0797-14.2015> doi: 10.1523/JNEUROSCI.0797-14.2015
- Bonaccio, S., & Dalal, R. S. (2006, 11). 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. Retrieved from <http://linkinghub>

- .elsevier.com/retrieve/pii/S0749597806000719 doi: 10.1016/j.obhdp.2006.07.001
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624–652. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.108.3.624> doi: 10.1037/0033-295X.108.3.624
- Clark, C. C., & Lawless, H. T. (1994). Limiting response alternatives in time-intensity scaling: an examination of the halo-dumping effect. *Chemical Senses*, 19(6), 583–594. Retrieved from <https://ezproxy-prd.bodleian.ox.ac.uk:5876/chemse/article/19/6/583/article> doi: 10.1093/chemse/19.6.583
- Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., ... Leonard, N. E. (2011). Uninformed Individuals Promote Democratic Consensus in Animal Groups. *Science (New York, N.Y.)*, 334(6062), 1578–1580. doi: 10.1126/science.1210280
- Dawes, R. M. (1989, 1). Statistical criteria for establishing a truly false consensus effect. *Journal of Experimental Social Psychology*, 25(1), 1–17. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/002210318990036X> doi: 10.1016/0022-1031(89)90036-X
- Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2016, 12). Echo Chambers: Emotional Contagion and Group Polarization on Facebook. *Scientific Reports*, 6(1), 37825.

Retrieved from <http://www.nature.com/articles/srep37825> doi: 10.1038/srep37825

- De Martino, B., O'Doherty, J. P., Ray, D., Bossaerts, P., & Camerer, C. (2013, 9). In the Mind of the Market: Theory of Mind Biases Value Computation during Financial Bubbles. *Neuron*, 79(6), 1222–1231. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0896627313005680> doi: 10.1016/j.neuron.2013.07.003
- Denrell, J. (2005). Why Most People Disapprove of Me: Experience Sampling in Impression Formation. *Psychological Review*, 112(4), 951–978. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.112.4.951> doi: 10.1037/0033-295X.112.4.951
- Desender, K., Boldt, A., & Yeung, N. (2018, 5). Subjective Confidence Predicts Information Seeking in Decision Making. *Psychological Science*, 29(5), 761–778. Retrieved from <http://journals.sagepub.com/doi/10.1177/0956797617744771> doi: 10.1177/0956797617744771
- Dienes, Z., & Mclatchie, N. (2017, 3). Four reasons to prefer Bayesian analyses over significance testing. *Psychonomic Bulletin & Review*, 1–12. Retrieved from <http://link.springer.com/10.3758/s13423-017-1266-z> doi: 10.3758/s13423-017-1266-z
- Ecken, P., & Pibernik, R. (2016, 7). Hit or Miss: What Leads Experts to Take Advice for Long-Term Judgments? *Management Science*, 62(7), 2002–2021. Retrieved from <http://pubsonline.informs.org/doi/10.1287/mnsc.2015.2219> doi: 10.1287/mnsc.2015.2219

- Epstein, J. M. (2013). *Agent_Zero: Toward Neurocognitive Foundations for Generative Social Science*. Princeton and Oxford: Princeton University Press.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*(10), 906–911. Retrieved from <http://content.apa.org/journals/amp/34/10/906> doi: 10.1037/0003-066X.34.10.906
- Fleming, S. M. (2016). Changing our minds about changes of mind. *eLife*, *5*. doi: 10.7554/eLife.14790
- Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision performance: A general Bayesian framework for metacognitive computation. *Psychological review*, *124*(1), 1–59. doi: 10.1002/sml.))
- Fleming, S. M., & Dolan, R. J. (2012, 4). The neural basis of metacognitive ability. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1594), 1338–1349. Retrieved from <http://rstb.royalsocietypublishing.org/cgi/doi/10.1098/rstb.2011.0417> doi: 10.1098/rstb.2011.0417
- Fleming, S. M., & Lau, H. C. (2014, 7). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*. Retrieved from http://www.frontiersin.org/Human_Neuroscience/10.3389/fnhum.2014.00443/abstract doi: 10.3389/fnhum.2014.00443
- Fusaroli, R., Bahrami, B., Olsen, K., Roepstorff, A., Rees, G., Frith, C., & Tylén, K. (2012, 7). Coming to Terms: Quantifying the Benefits of

- Linguistic Coordination. *Psychological science, online Jul*, 1–9. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/22810169> doi: 10.1177/0956797612436816
- Galton, F. (1907, 3). Vox Populi. *Nature*, 75(1949), 450–451. Retrieved from <http://www.nature.com/doifinder/10.1038/075450a0> doi: 10.1038/075450a0
- Gigerenzer, G. (2008). Why Heuristics Work? *Perspectives on Psychological Science*, 3(1), 20–29.
- Guggenmos, M., & Sterzer, P. (2017). A confidence-based reinforcement learning model for perceptual learning. *bioRxiv*.
- Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016, 3). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *eLife*, 5. Retrieved from <http://elifesciences.org/lookup/doi/10.7554/eLife.13388> doi: 10.7554/eLife.13388
- Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review*, 18(3), 186–201. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/h0074579> doi: 10.1037/h0074579
- Hertz, U., Romand-Monnier, M., Kyriakopoulou, K., & Bahrami, B. (2016). Social influence protects collective decision making from equality bias. *Journal of Experimental Psychology: Human Perception and Performance*, 42(2), 164–172. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/xhp0000145> doi: 10.1037/xhp0000145

- Jamieson, K. H., & Cappella, J. N. (2008). *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. Oxford: Oxford University Press. Retrieved from https://books.google.co.uk/books?id=1390a4M0sAgC&redir_esc=y
- Janis, I. L. (1972). *Victims of groupthink*. Boston: Houghton Mifflin.
- Jasny, L., Waggle, J., & Fisher, D. R. (2015, 5). An empirical examination of echo chambers in US climate policy networks. *Nature Climate Change*, 5(8), 782–786. Retrieved from <http://www.nature.com/doifinder/10.1038/nclimate2666> doi: 10.1038/nclimate2666
- Kao, A. B., & Couzin, I. D. (2014, 4). Decision accuracy in complex environments is often maximized by small group sizes. *Proceedings of the Royal Society B: Biological Sciences*, 281(1784), 20133305–20133305. Retrieved from <http://rspb.royalsocietypublishing.org/cgi/doi/10.1098/rspb.2013.3305> doi: 10.1098/rspb.2013.3305
- Kao, A. B., Miller, N., Torney, C., Hartnett, A., & Couzin, I. D. (2014, 8). Collective Learning and Optimal Consensus Decisions in Social Animal Groups. *PLoS Computational Biology*, 10(8), e1003762. Retrieved from <http://dx.plos.org/10.1371/journal.pcbi.1003762> doi: 10.1371/journal.pcbi.1003762
- Kelman, H. (1958). Compliance, identification, and internalization: Three processes of attitude change. *Journal of Conflict Resolution*, 2(1), 51–60.
- Koriat, A. (2012a, 1). The self-consistency model of subjective confi-

- dence. *Psychological review*, 119(1), 80–113. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/22022833> doi: 10.1037/a0025648
- Koriat, A. (2012b, 4). When are two heads better than one and why? *Science (New York, N.Y.)*, 336(6079), 360–2. Retrieved from <http://www.sciencemag.org/cgi/doi/10.1126/science.1216549><http://www.ncbi.nlm.nih.gov/pubmed/22517862> doi: 10.1126/science.1216549
- Krause, J., Ruxton, G. D., & Krause, S. (2010, 1). Swarm intelligence in animals and humans. *Trends in Ecology & Evolution*, 25(1), 28–34. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0169534709002298> doi: 10.1016/j.tree.2009.06.016
- Kruger, J., & Dunning, D. (1999). Unskilled and Unaware of It: How Difficulties in Recognizing One’s Own Incompetence Lead to Inflated Self-Assessments. *Journal of personality and social psychology*, 77(6), 1121–1134.
- Le Bon, G. (1895). *La psychologie des foules* (Félix Alca ed.). Paris: Ancienne Librairie Germer Bailliere et Cie. Retrieved from <http://socserv.mcmaster.ca/econ/ugcm/3113/lebon/Crowds.pdf>
- Liberman, V., Minson, J. A., Bryan, C. J., & Ross, L. (2012, 3). Naïve realism and capturing the wisdom of dyads. *Journal of Experimental Social Psychology*, 48(2), 507–512. Retrieved from <https://linkinghub.elsevier.com/retrieve/pii/S0022103111002599> doi: 10.1016/j.jesp.2011.10.016

- Mackay, C. (1841). *Extraordinary Popular Delusions and the Madness of Crowds* (Wordsworth ed.). Ware, UK: Wordsworth Edition Limited.
- Mahmoodi, A., Bang, D., Olsen, K., Zhao, Y. A., Shi, Z., Broberg, K., ... Bahrami, B. (2015). Equality bias impairs collective decision-making across cultures. *Proceedings of the National Academy of Sciences*, 201421692. Retrieved from <http://www.pnas.org/lookup/doi/10.1073/pnas.1421692112> doi: 10.1073/pnas.1421692112
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995, 7). An Integrative Model of Organizational Trust. *The Academy of Management Review*, 20(3), 709. Retrieved from <http://www.jstor.org/stable/258792?origin=crossref> doi: 10.2307/258792
- Meyniel, F., Sigman, M., & Mainen, Z. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, 88(1), 78–92. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0896627315008284> doi: 10.1016/j.neuron.2015.09.039
- Minson, J. A., Liberman, V., & Ross, L. (2011, 10). Two to Tango. *Personality and Social Psychology Bulletin*, 37(10), 1325–1338. Retrieved from <http://journals.sagepub.com/doi/10.1177/0146167211410436> doi: 10.1177/0146167211410436
- Myers, D. G., & Lamm, H. (1976). The group polarization phenomenon. *Psychological Bulletin*, 83(4), 602–627. doi: 10.1037/0033-2909.83.4.602
- Nickerson, R. (1998). Confirmation bias: a ubiquitous phenomenon in many

- guises. *Review of general psychology*, 2, 175–220.
- Penrod, S., & Cutler, B. (1995). Witness confidence and witness accuracy: Assessing their forensic relation. *Psychology, Public Policy, & Law*, 1, 817–845.
- Pescetelli, N., Rees, G., & Bahrami, B. (2016). The perceptual and social components of metacognition. *Journal of Experimental Psychology: General*, 145(8), 949–965. doi: 10.1037/xge0000180
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016, 2). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19(3), 366–374. Retrieved from <http://www.nature.com/doifinder/10.1038/nn.4240> doi: 10.1038/nn.4240
- Price, P. C., & Stone, E. R. (2004, 1). Intuitive evaluation of likelihood judgment producers: evidence for a confidence heuristic. *Journal of Behavioral Decision Making*, 17(1), 39–57. Retrieved from <http://doi.wiley.com/10.1002/bdm.460> doi: 10.1002/bdm.460
- Rader, C. A., Larrick, R. P., & Soll, J. B. (2017, 8). Advice as a form of social influence: Informational motives and the consequences for accuracy. *Social and Personality Psychology Compass*, 11(8), e12329. Retrieved from <http://doi.wiley.com/10.1111/spc3.12329> doi: 10.1111/spc3.12329
- Roediger III, H. L., Wixted, J. H., & Desoto, K. A. (2012, 7). The Curious Complexity between Confidence and Accuracy in Reports from Memory. In *Memory and law*

- (pp. 84–117). Oxford University Press. Retrieved from <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199920754.001.0001/acprof-9780199920754-chapter-4> doi: 10.1093/acprof:oso/9780199920754.003.0004
- Ross, L., Greene, D., & House, P. (1977, 5). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13(3), 279–301. Retrieved from <https://linkinghub.elsevier.com/retrieve/pii/002210317790049X> doi: 10.1016/0022-1031(77)90049-X
- Rouault, M., Dayan, P., & Fleming, S. M. (2019, 12). Forming global estimates of self-performance from local confidence. *Nature Communications*, 10(1), 1141. Retrieved from <http://www.nature.com/articles/s41467-019-09075-3> doi: 10.1038/s41467-019-09075-3
- Sah, S., Moore, D. A., & Maccoun, R. J. (2013). Cheap talk and credibility: The consequences of confidence and accuracy on advisor credibility and persuasiveness. *Organizational Behavior and Human Decision Processes*, 121, 246–255. doi: 10.1016/j.obhdp.2013.02.001
- Schultze, T., Gerlach, T. M., & Rittich, J. C. (2018, 7). Some People Heed Advice Less than Others: Agency (but Not Communion) Predicts Advice Taking. *Journal of Behavioral Decision Making*, 31(3), 430–445. Retrieved from <http://doi.wiley.com/10.1002/bdm.2065> doi: 10.1002/bdm.2065
- Schulz-Hardt, S., Brodbeck, F. C., Mojzisch, A., Kerschreiter, R., & Frey,

- D. (2006, 12). Group decision making in hidden profile situations: dissent as a facilitator for decision quality. *Journal of personality and social psychology*, 91(6), 1080–93. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17144766> doi: 10.1037/0022-3514.91.6.1080
- Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., & Frith, C. D. (2014, 2). Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences*. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S1364661314000230> doi: 10.1016/j.tics.2014.01.006
- Sherif, C., Sherif, M., & Nebergall, R. (1965). *Attitude and attitude change*. Philadelphia: W.B. Saunders Company.
- Shirado, H., & Christakis, N. A. (2017, 5). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, 545(7654), 370–374. Retrieved from <http://www.nature.com/doifinder/10.1038/nature22332> doi: 10.1038/nature22332
- Shultz, T. R., Katz, J. A., & Lepper, M. R. (2001). Clinging to beliefs: A constraint-satisfaction model. In *Proceedings of the twenty-third annual conference of the cognitive science society* (pp. 928–933). Mahwah, NJ: Erlbaum.
- Snizek, J. A., & Buckley, T. (1989). Social influence in the advisor-judge relationship. In *Annual meeting of the judgment and decision making societ.* Atlanta, Georgia.
- Snizek, J. A., & Van Swol, L. M. (2001, 3). Trust, Confidence, and Expertise in a Judge-Advisor System. *Organizational behavior and human*

- decision processes*, 84(2), 288–307. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11277673> doi: 10.1006/obhd.2000.2926
- Soll, J. B., & Larrick, R. P. (2009, 5). Strategies for revising judgment: how (and how well) people use others' opinions. *Journal of experimental psychology: Learning, memory, and cognition*, 35(3), 780–805. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19379049> doi: 10.1037/a0015145
- Sorkin, R. D., Hays, C. J., & West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review*, 108(1), 183–203. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.108.1.183> doi: 10.1037/0033-295X.108.1.183
- Stasser, G., & Titus, W. (2003). Hidden Profiles : A Brief History. *Psychological Inquiry*, 14(3), 304–313.
- Sunstein, C. R. (2001). *Republic.com*. Princeton University Press.
- Surowiecki, J. (2004). *The Wisdom of Crowds. Why the Many are Smarter than the Few*. (Abacus ed.). London: Little, Brown Book Group.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: an introduction*. Cambridge, MA: MIT Press.
- Swol, L. M., & Sniezek, J. A. (2005). Factors affecting the acceptance of expert advice. *British journal of social psychology*, 44(3), 443–461.
- Tenney, E. R., MacCoun, R. J., Spellman, B. a., & Hastie, R. (2007, 1). Calibration trumps confidence as a basis for witness credibility. *Psychological Science*, 18(1), 46–50. Retrieved from <http://www.ncbi.nlm>

- .nih.gov/pubmed/17362377 doi: 10.1111/j.1467-9280.2007.01847.x
- Tost, L. P., Gino, F., & Larrick, R. P. (2012, 1). Power, competitiveness, and advice taking: Why the powerful don't listen. *Organizational Behavior and Human Decision Processes*, 117(1), 53–65. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0749597811001233> doi: 10.1016/j.obhdp.2011.10.001
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., ... Nelson, C. (2009, 8). The Nim-Stim set of facial expressions: judgments from untrained research participants. *Psychiatry research*, 168(3), 242–9. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3474329&tool=pmcentrez&rendertype=abstract> doi: 10.1016/j.psychres.2008.05.006
- Treutwein, B. (1995). Adaptive Psychophysical Procedures. *Vision Research*, 35(17), 2503–2522.
- Turner, M. E., & Pratkanis, A. R. (1998, 2). Twenty-Five Years of Groupthink Theory and Research: Lessons from the Evaluation of a Theory. *Organizational Behavior and Human Decision Processes*, 73(2-3), 105–115. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S074959789892756X> doi: 10.1006/obhd.1998.2756
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. Retrieved from <http://www.jstor.org/stable/1738360Copy>

- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315. Retrieved from <http://content.apa.org/journals/rev/90/4/293> doi: 10.1037/0033-295X.90.4.293
- Vandormael, H., Herce Castañón, S., Balaguer, J., Li, V., & Summerfield, C. (2017, 3). Robust sampling of decision information during perceptual choice. *Proceedings of the National Academy of Sciences*, 114(10), 2771–2776. Retrieved from <http://www.pnas.org/lookup/doi/10.1073/pnas.1613950114> doi: 10.1073/pnas.1613950114
- Van Swol, L. M., & Ludutsky, C. L. (2007, 6). Tell Me Something I Don't Know. *Communication Research*, 34(3), 297–312. Retrieved from <http://journals.sagepub.com/doi/10.1177/0093650207300430> doi: 10.1177/0093650207300430
- Weiss, D. J., & Shanteau, J. (2003, 1). Empirical assessment of expertise. *Human factors*, 45(1), 104–16. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12916584>
- Wilensky, U. (1999). *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.
- Yaniv, I. (2004, 1). Receiving other peoples advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, 93(1), 1–13. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0749597803001018> doi: 10.1016/j.obhdp.2003.08.002
- Yaniv, I., Choshen-Hillel, S., & Milyavsky, M. (2009). Spurious consensus

and opinion revision: Why might people be more confident in their less accurate judgments? *Journal of experimental psychology. Learning, memory, and cognition*, 35(2), 558–563. doi: <http://dx.doi.org/10.1037/a0014589>

Yaniv, I., & Kleinberger, E. (2000, 11). Advice Taking in Decision Making: Egocentric Discounting and Reputation Formation. *Organizational behavior and human decision processes*, 83(2), 260–281. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11056071> doi: 10.1006/obhd.2000.2909

Yeung, N., & Summerfield, C. (2012, 5). Metacognition in human decision-making: confidence and error monitoring. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 367(1594), 1310–21. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3318764&tool=pmcentrez&rendertype=abstract> doi: 10.1098/rstb.2011.0416

Zarnoth, P., & Snizek, J. A. (1997, 7). The Social Influence of Confidence in Group Decision Making. *Journal of Experimental Social Psychology*, 33(4), 345–366. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0022103197913263> doi: 10.1006/jesp.1997.1326

Zylberberg, A., Wolpert, D. M., & Shadlen, M. N. (2018, 9). Counterfactual Reasoning Underlies the Learning of Priors in Decision Making. *Neuron*, 99(5), 1083–1097. Retrieved from <https://linkinghub>

.elsevier.com/retrieve/pii/S0896627318306330 doi: 10.1016/
j.neuron.2018.07.035