Knowing whom to learn from: individual differences in metacognition and weighting of social information

Karsten Olsen^{1,2,*}

Andreas Roepstorff^{1,3}

Dan Bang^{1,4}

1 Interacting Minds Centre, Aarhus University

2 Department for Linguistics, Cognitive Science and Cognitive Semiotics

3 Social Anthropology, Aarhus University

4 Wellcome Centre for Human Neuroimaging, University College London

* Corresponding author

Abstract

Social learning enables us to acquire skills and knowledge more efficiently, provided that we learn from the right others. However, little is known about the cognitive factors that determine whom we decide to learn from and how much we benefit from such learning. Here we address this question using a perceptual task where participants had the opportunity to revise their responses (choice about a noisy stimulus and confidence in this choice being correct) in light of the responses made by two advisors of different reliability. We found consistent individual differences in the weights assigned to the social sources, and in the benefit obtained from the social sources – factors which we collectively refer to as Social Weighting Sensitivity (SWS). We also found that an individual's metacognition predicted SWS, with participants who were overconfident about their performance listening less to and benefiting less from social information. Finally, at the trial level, we found that participants adjusted the reliance on social information in light of the response and feedback history, indicating that they formed metacognitive estimates about the social sources. Our study highlights the importance of metacognition and SWS in human social learning.

Key words

Social learning; metacognition; decision-making; advice taking; confidence; individual differences

1 Introduction

Since the demonstration of Roger's Paradox [1], there has been great interest in the strategies that steer individuals towards reliable social models for learning through observation or interaction [2-4]. From what began as a thought experiment, the anthropologist Alan Rogers used a simple model to show, rather counter-intuitively, that in a changing environment, the fitness of social learners would, at equilibrium, be equal to the fitness of asocial learners. In other words, there was, mathematically, no selection for social learning. Given that social learning is ubiquitous amongst human and non-human animals, and generally thought to increase fitness [5], this result became known as Roger's paradox.

In an attempt to resolve Roger's paradox, subsequent research has sought to identify heuristics that direct social learning in an adaptive manner [6-8]. These heuristics are known as Social Learning Strategies (SLSs), and can be grouped according to the 'when', the 'what' and the 'whom' of social learning [4]. First, 'when' SLSs concern the conditions in which an individual should engage in social learning (e.g., copy others when hungry). Second, 'what' SLSs concern biases towards particular kinds of information (e.g., copy memorable information). Finally, 'whom' SLSs concern biases towards particular social models (e.g., copy successful individuals) [9]. Indeed, it pays off to be careful and selective when choosing whom to learn from [10, 11]. More generally, these strategies have been argued to be central to understanding the development of group-level coordination and cultural evolution [5, 12, 13].

Research on SLSs has been instrumental in extending social learning beyond undirected copying and demonstrating its adaptiveness. However, perhaps because of its cross-species perspective, the focus of this research has been on the behavioural consequences of SLSs rather than their underlying cognitive machinery. It has recently been proposed that, while most SLSs rely on domain-general processes of associative learning [14], some SLSs, unique to humans, specifically involve the ability to evaluate and report on properties of one's own cognitive processes [15]. This ability, known as (explicit) metacognition, facilitates novel ways of social learning [16]. For example, it enables individuals to recognize when they hold uncertain beliefs and therefore would benefit from obtaining information from others. It also enables individuals to share information to the benefit of others. For these reasons, metacognition has been proposed to be central to human group coordination and cultural development [15, 17, 18].

Research on metacognition and social cognition has, however, identified several factors which may impact on the adaptiveness of metacognitive SLSs. For example, in order to know when to turn to others, an individual must estimate the reliability of their own beliefs. However, people differ in their degree of insight [19, 20] and their general level of confidence [21-23]. This in turn means that, in order to harness social information, an individual must also estimate the reliability of others' beliefs independently of the confidence that they overtly express. While people do maintain such estimates [24, 25], they do not always use them optimally. For example, people show an 'equality bias', wrongly weighting others' beliefs as if they are as 'good' or as 'bad' as their own, regardless of objective differences in reliability [26]. In summary, in order to benefit from metacognitive SLSs, an individual must not only estimate the reliability of a social source but

also be able to use the social information in accordance with this estimate – an ability which we refer to as Social Weighting Sensitivity (SWS).

In this study, we investigate whether SWS is an individual trait that varies between individuals. Several studies have assessed individual differences in metacognition [19, 20, 27, 28], but we do not know whether there are individual differences in its social counterpart. If SWS is an individual trait, then we would expect SWS to be stable within individuals across different social contexts. For example, an individual's ability to harness information from unreliable and reliable social sources should be correlated. Further, even if involving distinct cognitive processes, we would expect a link between an individual's SWS and metacognition. For example, an individual's general level of confidence may constrain the benefit obtained from social information: if we have high confidence in our own beliefs, then we are unlikely to turn to others [29]. These aspects of social learning, which presumably are central to identifying reliable social models for learning through observation or interaction, have received little attention in the literature.

We addressed these questions using a social psychophysical task which required participants to combine their own perceptual judgements with those of others. On each trial, participants made a categorical decision about a random dot motion stimulus and reported their confidence in the decision being correct. They were then informed about the decision and confidence of a 'previous' participant (the advisor) on the corresponding trial and given the opportunity to revise their initial response in light of this information. The advisors were, in reality, computer agents designed to mimic unreliable and reliable participants. By means of this task design, we could test whether an individual's SWS was stable across social sources of varying reliability. In addition, we could test whether trait-level features of an individual's metacognition impact on the benefit that they obtain from social information. Finally, broadening the scope to the literature on advice taking [30], we could test how one's own as well as others' response and feedback history affect initial belief formation and subsequent belief revision.

2 Methods

2.1 Participants

Thirty-nine adults (27 female; mean age $\pm$ SD, 22.49 $\pm$ 2.10) with normal or corrected-to-normal vision participated in the study. All participants provided written informed consent.

2.2 Experimental task

Participants viewed a field of moving dots inside a circular aperture (Figure 1). On each update of the display, a fraction of the dots moved coherently, either to the left or to the right, whereas the remainder moved randomly. Once the display terminated, participants had to decide whether the net motion direction was left or right and indicate their confidence in this decision. Participants responded (decision and confidence) by sliding a marker along a scale with a fixed midpoint. The left-hand side corresponded to 'left' motion, and the right-hand side corresponded to 'right' motion. Each step further away from the centre indicated higher confidence (60% to 100% in steps of 10%). After having made a response, participants were informed about the response of another

participant (advisor) and given the opportunity to revise their initial response in the light of this new information. After having made their revised response, participants received feedback about the correct decision and their trial score (see below).

Figure 1. Schematic of experimental trial. Each trial began with the presentation of a fixation cross at the centre of a circular aperture. After a uniformly sampled delay (.5-1s), a random dot motion display was presented (1s). Participants were asked to maintain fixation during the stimulus presentation. The net direction of motion was either left or right. Once the random-dot display had terminated, participants were presented with a scale with a marker at its midpoint. They indicated their decision about the net direction of motion and their confidence in this decision by moving the marker along the scale. The left-hand side corresponded to left motion; the right-hand side corresponded to right motion. Each step further away from the centre indicated higher confidence (60% to 100%, by steps of 10%). After having made their response, participants were shown the response of another participant, the ‘advisor’. Participants were given the opportunity to revise their initial response in the light of this new information, before continuing to the next trial. Participants were paired with two advisors in total: a bad advisor (about 58% accuracy) and a good advisor (about 82% accuracy).

The fraction of coherently moving dots (coherence) had been calibrated in a separate session so as to achieve around 71% accuracy. More specifically, participants performed two blocks of the perceptual task without an advisor. In each block, the coherence was titrated using a 2-down-1-up procedure: after two correct decisions, coherence was decreased by .01; after one incorrect decision, coherence was increased by .01. In the first block, coherence was initialised at .2. In the second block, coherence was initialised at the average coherence across the first block. We used the average coherence in the second block for the main experiment.

Participants were told that the advisors were previous participants who had performed exactly the same task but, in reality, the advisors were simulated. Participants were paired with two advisors: a ‘bad’ advisor (about 58% accuracy) and a ‘good’ advisor (about 82% accuracy). The advisors used each confidence level the same proportion of times. An optimal agent would therefore assign more weight to the good advisor than the bad advisor for a given confidence level. Participants performed 2 blocks of 120 trials, one block for each advisor. The order of the advisors was counterbalanced across participants.

We submitted participants’ revised responses to a strictly proper scoring rule [31], with a the score (reward or loss) on trial t calculated as follows: $r = 100 \times (0.5 - (c_t - a_t)^2)$, where confidence c is $\{0.6, 0.7, 0.8, 0.9, 1\}$ and accuracy a is $\{0, 1\}$. Under this rule, it is in participants’ best interest to make as many correct decisions as possible and to indicate their confidence as accurately as possible. The rule thus encouraged participants to revise their response in light of advisor’s response. Participants’ scores ranged from -100 DKK (e.g., “100% LEFT” when motion direction was right) to 100 DKK (e.g., saying “100% LEFT” when motion direction was left). These earnings

were fictive, and all participants received a flat rate for participation. See Electronic Supplementary Materials for display settings, stimulus construction and simulation of advisors.

2.3 Data analysis

We performed main analyses by splitting the data into *initial* responses, submitted before seeing an advisor's response, and *revised* responses, submitted after seeing an advisor's response. We used repeated measures ANOVAs to test how *advisor reliability* (bad versus. good) and *response revision* (initial versus revised) affected *choice accuracy*, *mean confidence* and metacognition. We decomposed metacognition into *resolution* and *calibration* [32]. Resolution, a measure of metacognitive sensitivity, is the difference in mean confidence between correct and incorrect decisions – with higher values indicating higher resolution. Calibration, a measure of metacognitive bias, is the sum of differences between mean confidence and choice accuracy across confidence levels – with negative values indicating underconfidence and positive values indicating overconfidence. We quantified a participant's *susceptibility* to social advice using trial-by-trial multiple linear regression in which we predicted a participant's revised *signed* confidence (left, negative; right, positive) using an advisor's signed confidence, the participant's initial signed confidence and the participant's initial choice accuracy. We used the coefficient for the advisor's signed confidence as our measure of susceptibility. We investigated trial history effects on belief formation and belief revision by applying trial-by-trial multiple linear regressions to a participant's initial and revised confidence, respectively. We z-scored all confidence variables and contrast-coded choice accuracy (incorrect, -.5; correct, .5) in regression analyses. We tested significance by comparing the coefficients pooled across separate regression analyses for each participant against zero (one-sample *t*-test). We performed the individual difference analyses using robust multiple linear regression to minimise influence of outliers. Repeated measures ANOVAs were implemented in JASP; all other analyses were implemented in MATLAB.

3 Results

3.1 Choice accuracy and mean confidence

We first tested how advisor *reliability* and response *revision* affected choice accuracy and mean confidence (Figure 2A-B). Repeated measures ANOVAs showed positive main effects of reliability (accuracy, $F_{1,38} = 20.22$, $p < .001$; confidence, $F_{1,38} = 16.41$, $p < .001$) and revision (accuracy, $F_{1,38} = 69.51$, $p < .001$; confidence, $F_{1,38} = 78.45$, $p < .001$), and a positive interaction effect between reliability and revision (accuracy, $F_{1,38} = 98.52$, $p < .001$; confidence, $F_{1,38} = 21.54$, $p < .001$). Post-hoc paired *t*-tests showed that response revision increased choice accuracy (bad, $t(38) = 1.95$, $p = .058$; good, $t(38) = 9.94$, $p < .001$) and mean confidence (bad, $t(38) = 7.61$, $p < .001$; good, $t(38) = 8.59$, $p < .001$) for both advisors (Figure 2A-B). Further, despite initially comparable (Figure 2A-B; accuracy, $t(38) = 0.43$, $p = .671$; confidence, $t(38) = 1.96$, $p = .058$), choice accuracy and mean confidence were higher for the good compared to the bad advisor after response revision (Figure 2A-B; accuracy, $t(38) = 7.56$, $p < .001$; confidence, $t(38) = 5.21$, $p < .001$).

Figure 2. Effects of advisor reliability and response revision on behavioural responses. **A**, Choice accuracy (i.e. fraction of correct decisions). **B**, Mean confidence. **C**, Resolution defined as the difference in mean confidence for correct and wrong decisions. **D**, Calibration defined as the sum of differences between mean confidence and choice accuracy for each step on the confidence scale. **A-D**, Data are presented as group mean $\pm$ SEM. Each dot is a participant.

3.2 Metacognition

We next tested how advisor reliability and response revision affected metacognition – here decomposed into resolution and calibration (Figure 2C-D). In broad terms, resolution describes a participant's degree of insight into their choice accuracy (higher values indicates higher resolution), whereas calibration describes a participant's degree of underconfidence or overconfidence about their choice accuracy (negative, underconfidence; positive, overconfidence; zero, no bias).

A repeated measures ANOVA on resolution showed positive main effects of reliability ($F_{1,38} = 67.07$, $p < .001$) and revision ($F_{1,38} = 18.05$, $p < .001$), and a positive interaction effect between reliability and revision ($F_{1,38} = 25.72$, $p < .001$). Post-hoc paired t -tests showed that response revision increased resolution for both advisors (Figure 2C; bad, $t(38) = 4.00$, $p < .001$; good, $t(38) = 7.70$, $p < .001$). Further, despite initially comparable (Figure 2C; $t(38) = 0.96$, $p = .344$), resolution was higher for good compared to bad advisors after response revision (Figure 2C; $t(38) = 5.56$, $p < .001$).

A repeated measures ANOVA on calibration showed no main effect of reliability ($F_{1,38} = 0.36$, $p = .551$), but a positive main effect of revision ($F_{1,38} = 5.37$, $p < .001$), and a positive interaction effect between reliability and revision ($F_{1,38} = 18.36$, $p < .001$). Follow-up paired t -tests showed that calibration decreased for bad advisors but improved for good advisors after response revision (Figure 2D; bad, $t(38) = 2.75$, $p < .009$; good, $t(38) = -3.02$, $p = .004$). Further, despite initially comparable (Figure 2D; $t(38) = 1.23$, $p = .23$), calibration was better for good compared to bad advisors after response revision (Figure 2D; $t(38) = -4.87$, $p < .001$).

Figure 3. Decomposition of metacognitive measures. **A**, Resolution: mean confidence as a function of choice accuracy. **B**, Calibration: objective versus subjective performance as a function of response. The identity line corresponds to perfect calibration. Note that our confidence scale constrains the x-axis to the range 60%-100%. **A-B**, Data are presented as group mean $\pm$ SEM.

We next deconstructed our measures of resolution and calibration into component parts in order to unpack the effects of advisor reliability and response revision on metacognition. First, by plotting mean confidence as a function of choice accuracy, advisor reliability and response revision, we see that the resolution effects were driven by a larger increase in confidence for correct decisions when paired with the good advisor (Figure 3A). This interpretation is supported by a repeated measures ANOVA on mean confidence with choice accuracy and response revision as within-subject factors and advisor reliability as between subject factor: the analysis shows a

positive interaction effect between accuracy, revision and reliability ($F_{1,76} = 23.43$, $p < .001$), in addition to positive main effects of accuracy ($F_{1,76} = 181.12$, $p < .001$) and revision ($F_{1,76} = 80.15$, $p < .001$), and positive interaction effects between accuracy and revision ($F_{1,76} = 75.21$, $p < .001$) and accuracy and reliability ($F_{1,76} = 5.85$, $p < .001$). Second, by plotting calibration curves as a function of advisor reliability and response revision, we see that the calibration effects were driven by a ‘correction’ of initial overconfidence (Figure 3B). This interpretation is supported by post-hoc paired t -tests comparing changes in calibration between bad and good advisors along the full calibration curve (60%, $t(35) = .78$, $p = .438$; 70%, $t(37) = 2.03$, $p = .049$; 80%, $t(37) = 3.19$, $p = .003$; 90%, $t(30) = 1.65$, $p = .110$; 100%, $t(27) = 1.65$, $p = .039$; degrees of freedom vary because of incomplete cell counts for some participants).

Figure 4. Individual variation in social traits. **A**, Susceptibility to advisors quantified using trial-by-trial regression model. **B**, Relationship between susceptibility to bad and good advisor. **C**, Benefit obtained from advisor quantified as difference in score for revised and initial responses. **D**, Relationship between benefits obtained from bad and good advisors. **A** and **C**, Data are presented as group mean $\pm$ SEM. Each dot is a participant. **B** and **D**, Black line is the best-fitting line from a robust regression. Each dot is a participant.

3.3 Social susceptibility and benefit

We next directly quantified participants' *susceptibility* to the advisors and the *benefit* that they obtained from the advisors. In line with the above results, good advisors had a higher influence on response revision than bad advisors (Figure 4A; $t(38) = 10.41$, $p < .001$, one-sample t -test on

social regression coefficient). However, the spread of individual data points indicated that there were individual differences in susceptibility. Indeed, the higher the influence of the good advisor, the higher the influence of the bad advisor (Figure 4B; $r(38) = .76$, $p < .001$, Pearson). There were also individual differences in benefit. Participants overall accrued more points when paired with good advisors (Figure 4C; $t(38) = 10.58$, $p < .001$, paired t -test), but the larger the score for the good advisor, the larger the score for the bad advisor (Figure 4D; $r(38) = .41$, $p = .009$, Pearson). Interestingly, there was a positive correlation between susceptibility and benefit, controlling for initial choice accuracy ($r(38) = .58$, $p < .001$, partial correlation).

3.4 Metacognitive predictors of social susceptibility and benefit

We next sought to identify factors which predicted susceptibility and benefit. To this end, we performed robust regressions in which resolution and calibration (calculated from the initial responses) competed to predict susceptibility or benefit. Calibration predicted both measures (Figure 5A-B; susceptibility, $t(35) = -2.40$, $p = .022$; benefit, $t(35) = -3.22$, $p = .003$). In particular, the higher the degree of overconfidence, the lower the influence of the advisors, and the smaller the benefit obtained from the advisors. In contrast, resolution predicted neither susceptibility nor benefit (Figure 5A-B; susceptibility, $t(35) = -0.46$, $p = .650$; benefit, $t(35) = -0.69$, $p = .492$). We controlled for initial choice accuracy in both analyses.

Figure 5. Relating individual variation in metacognition and social traits. **A**, Relationship between calibration (left) or resolution (right) and susceptibility. **B**, Relationship between calibration (left) or resolution (right) and benefit. **A-B**, Black line is the best-fitting line from a robust regression. Each dot is a participant.

Finally, our task allowed us to investigate how the trial history affected initial belief formation as well as belief revision in light of the advisor's response. To this end, we performed regression analyses in which we predicted participants' initial or revised confidence using their own (s) and the advisors' (p) confidence (C) and accuracy (A) on the previous trial (revised responses). We combined data across advisors, and only looked one trial back into time, in order to obtain reliable parameter estimates.

We first analysed how the trial history affected initial belief formation (Figure 6A). We found that participants' current confidence was modulated by an interaction between their report and feedback history ($C_{s,t-1}$, $t(38) = 1.97$, $p = .057$; $A_{s,t-1}$, $t(38) = 1.28$, $p = .208$; $C_{s,t-1} \times A_{s,t-1}$, $t(38) = 3.70$, $p < .001$). Post-hoc visualisation of these effects showed that, while participants' previous confidence had no effect on their current confidence when the previous decision was incorrect, it biased current confidence when the previous decision was correct (Figure 6B). There were no effects of the advisor's report and/or feedback history ($C_{p,t-1}$, $t(38) = -1.74$, $p = .090$; $A_{s,t-1}$, $t(38) = -0.09$, $p = .931$; $C_{s,t-1} \times A_{s,t-1}$, $t(38) = 1.27$, $p = .210$).

We next analysed how the trial history affected belief revision (Figure 6C). We found a positive relationship between, on one hand, participants' revised confidence and, on the other hand, their initial and previous confidence ($C_{s,t}$, $t(38) = 17.80$, $p < .001$; $C_{s,t-1}$, $t(38) = 2.53$, $p = .016$; $C_{s,t} \times C_{s,t-1}$, $t(38) = -0.64$, $p = .527$). There were no effects of their feedback history ($A_{s,t-1}$, $t(38) = 1.72$, $p = .093$; $A_{s,t-1} \times C_{s,t-1}$, $t(38) = 0.53$, $p = .601$; $C_{s,t} \times A_{s,t-1}$, $t(38) = -0.64$, $p = .527$; $C_{s,t} \times C_{s,t-1} \times A_{s,t-1}$, $t(38) = -0.96$, $p = .344$). We also found that participants' revised confidence was modulated by an interaction between the advisors' report and feedback history ($C_{p,t}$, $t(38) = 11.68$, $p < .001$; $A_{p,t-1}$, $t(38) = 0.17$, $p = .862$; $C_{p,t-1}$, $t(38) = -0.92$, $p = .361$; $A_{p,t-1} \times C_{p,t-1}$, $t(38) = 0.28$, $p = .779$; $C_{p,t} \times C_{p,t-1}$, $t(38) = -1.33$, $p = .193$; $C_{p,t} \times A_{s,t-1}$, $t(38) = 3.76$, $p < .001$; $C_{p,t} \times C_{p,t-1} \times A_{s,t-1}$, $t(38) = 2.27$, $p = .029$). Post-hoc visualisation of these effects revealed an intricate pattern (Figure 6D). Specifically, when the advisor's previous decision was incorrect, the less confident they were on the previous trial, the higher the impact of their current confidence on the participant's current confidence (left panel). However, when the advisor's previous decision was correct, the impact of their current confidence on the participant's current confidence was unaffected by their previous confidence (right panel).

4 Discussion

Empirical and theoretical work on social learning has revealed strategies that adaptively guide copying and learning behaviour in humans and non-human animals. These Social Learning Strategies (SLSs) have been proposed to be key to the evolution of group coordination and culture [3, 33]. However, it has been hypothesised that humans are unique in that they also target their social learning by asking 'who knows?', and that these metacognitive SLSs, which involve the ability to explicitly evaluate and report on properties of one's cognitive processes, hold the key to understanding the sophisticated forms of group coordination and culture found in humans [15]. Extending on these ideas, we highlight that, once a social learner has estimated the reliability of a social source of information, they must then also weight the social information in accordance with this estimate. This can be a challenge, however, as individuals in social contexts may be subject to biases and not always able to use such estimates optimally [26]. Here we build on these hypotheses by relating metacognition to social learning in human participants.

Participants performed a social psychophysical task which required them to combine their own perceptual judgements with those of others. On each trial, after having indicated a choice and the confidence in this choice being correct, participants were informed about the response of another participant (the advisor) and then given the opportunity to revise their initial response. Participants were told that the advisor was a previous participant who had performed exactly the same task,

but, in reality, the advisor was a computer agent. Participants were paired with two advisors of varying reliability; one advisor had a lower choice accuracy than the participant, whereas the other advisor had a higher choice accuracy than the participant. In order to be successful, participants must not only evaluate the reliability of their own beliefs, that is, have good metacognition. They must also estimate the reliability of the social source and use the social information in accordance with this estimate – an ability which we refer to as Social Weighting Sensitivity (SWS). By means of this design, we could test the predictions that SWS is an individual trait, i.e. stable within an individual across advisors, and that an individual's metacognition constrains their ability to benefit from social information.

In line with our predictions, we found that SWS was stable within an individual across advisors: the higher the influence of the good advisor on response revision, the higher the influence of the bad advisor, and the larger the benefit obtained from seeing the good advisor's response, the larger the benefit from seeing the bad advisor's response, and vice versa. Moreover, we found that these SWS measures, susceptibility and benefit, were related to an individual's metacognition. We quantified these using initial responses which, critically, were made before seeing an advisor's response. More specifically, we found that, while an individual's degree of insight into choice accuracy (i.e. resolution) did not correlate with susceptibility or benefit, their degree of confidence bias relative to choice accuracy (i.e. calibration) correlated with both measures, with higher confidence bias – an indicator of self-efficacy – associated with lower susceptibility and benefit. In other words, the more overconfident an individual is, the less likely they are to consider and thereby benefit from social information. More generally, these results show that the tendency for some people to insist on their own opinion when faced with beneficial but conflicting advice, a bias known as *egocentric advice discounting* [34, 35], varies between individuals and is intimately linked to an individual's self-efficacy.

Our task also allowed us to investigate how the advisors affected participants' metacognition. Interestingly, we found that the advisors affected both resolution and calibration as assessed by comparing features of participants' initial and revised responses. More specifically, both advisors improved resolution, but the improvement was larger for the good advisor, and was due to a selective increase in confidence for correct decisions. In contrast, the advisors had opposite effects on calibration: the good advisor improved calibration, but the bad advisor impaired calibration. Follow-up analysis showed that the good advisor lead to a decrease in overconfidence for easy decisions, whereas the bad advisor lead to an increase in overconfidence for difficult decisions. This effect is, however, likely to reflect a mixture of changes in both choice accuracy and reported confidence (e.g., a reduction in overconfidence may result from an increase in choice accuracy). Taken together, these results indicate that people engage in metacognitive learning at multiple levels. More generally, these results fit with the hypothesis that confidence is subject to social influences over and above task performance [36].

Finally, our task allowed us to investigate belief formation (initial response) and belief revision (revised response) in social situations at the trial level. We found that belief formation was shaped by participants' own but not the advisors' trial history. Interestingly, participants were only biased by their past successes and not their past mistakes: if the previous decision was correct, then the

higher the confidence in that decision, the higher the confidence in the current decision; in contrast, if the previous decision was incorrect, then there was no serial effect. Turning to belief revision, we found that participants were affected by an interaction between the advisors' past confidence and past accuracy: the advisors' influence was larger when they reported high confidence after (a) reporting low confidence in a decision that turned out to be incorrect or (b) having made a correct decision. Taken together, these results show that people are sensitive not only to others' task performance but also their metacognitive ability and adjust the degree to which they rely on them accordingly. More broadly, these results highlight the importance of metacognition and SWS in social learning contexts. For example, tracking metacognitive information allows an individual to assess the 'evidence' underpinning others' behaviour and in turn use this estimate to optimise the degree to which they rely on others as a model for social learning.

Future research should seek to uncover the cognitive and neural basis of SWS. For example, there is emerging evidence that metacognition is a domain-general ability, with features of an individual's metacognition correlating across task domains [21, 22, 37]. However, it is not known whether the same holds for its social counterpart. Furthermore, several studies have related individual differences in metacognition to neural anatomy [19, 27] and neural activity [38] – with lateral prefrontal cortex as a candidate neural substrate for metacognition. However, the neural basis of SWS is unknown. Other avenues for future research is to establish how SWS relates to other cognitive factors, such as developmental change [39, 40], and the between-group variation in social learning observed cross-culturally [41]. In order to address these open questions, researchers may benefit from a close integration of the cognitive sciences and evolutionary biology [2, 42]. In conclusion, we believe that understanding the cognitive and neural mechanisms underlying metacognition and social learning and their interaction is an important stepping stone toward explaining the sophisticated forms of group coordination and culture found in humans.

Ethics

The study was approved by the Research Ethics Committee of Central Region Denmark.

Data, code and materials

Anonymised behavioural data and code supporting main analyses are available on GitHub (<https://github.com/danbang/article-who-knows/>).

Competing interests

The authors report no competing interests.

Author contributions

KO, DB and AR conceived the study and designed the experiment. KO performed data collection. KO and DB analysed and interpreted the data. KO drafted the manuscript. KO and DB wrote the manuscript. AR critically revised the manuscript. All authors gave final approval for publication and agree to be held accountable for the work performed therein.

Acknowledgements

This project was supported by seed funding from the Interacting Minds Centre at Aarhus University. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z).

References

- [1] Rogers, A.R. 1988 Does biology constrain culture? *American Anthropologist* **90**, 819-831.
- [2] Kendal, R.L., Boogert, N.J., Rendell, L., Laland, K.N., Webster, M. & Jones, P.L. 2018 Social learning strategies: bridge-building between fields. *Trends in Cognitive Sciences*.
- [3] Rendell, L., Fogarty, L., Hoppitt, W.J., Morgan, T.J., Webster, M.M. & Laland, K.N. 2011 Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends in Cognitive Sciences* **15**, 68-76.
- [4] Hoppitt, W. & Laland, K.N. 2013 *Social learning: an introduction to mechanisms, methods, and models*, Princeton University Press.
- [5] Boyd, R. & Richerson, P.J. 1985 *Culture and the evolutionary process*, University of Chicago Press.
- [6] Heyes, C.M. 1994 Social learning in animals: categories and mechanisms. *Biological Reviews* **69**, 207-231.
- [7] Kline, M.A. 2015 How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain Sciences* **38**.
- [8] Laland, K.N. 2004 Social learning strategies. *Animal Learning & Behavior* **32**, 4-14.
- [9] Muthukrishna, M., Morgan, T.J. & Henrich, J. 2016 The when and who of social learning and conformist transmission. *Evolution and Human Behavior* **37**, 10-20.
- [10] Wisdom, T.N., Song, X. & Goldstone, R.L. 2013 Social learning strategies in networked groups. *Cognitive Science* **37**, 1383-1425.
- [11] Rendell, L., Boyd, R., Enquist, M., Feldman, M.W., Fogarty, L. & Laland, K.N. 2011 How copying affects the amount, evenness and persistence of cultural knowledge: insights from the social learning strategies tournament. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* **366**, 1118-1128.
- [12] Chudek, M. & Henrich, J. 2011 Culture–gene coevolution, norm-psychology and the emergence of human prosociality. *Trends in Cognitive Sciences* **15**, 218-226.
- [13] Powell, A., Shennan, S. & Thomas, M.G. 2009 Late Pleistocene demography and the appearance of modern human behavior. *Science* **324**, 1298-1301.
- [14] Heyes, C. & Pearce, J.M. 2015 Not-so-social learning strategies. *Proceedings of the Royal Society of London B: Biological Sciences* **282**, 20141709.
- [15] Heyes, C. 2016 Who knows? Metacognitive social learning strategies. *Trends in Cognitive Sciences* **20**, 204-213.
- [16] Dunstone, J. & Caldwell, C.A.J.P.C. 2018 Cumulative culture and explicit metacognition: a review of theories, evidence and key predictions. *Palgrave Communications* **4**, 145.
- [17] Frith, C.D. 2014 How the brain creates culture. *Nova Acta Leopoldina* **120**, 3-26.
- [18] Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C. & Frith, C.D. 2014 Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences* **18**, 186-193.
- [19] Fleming, S.M., Weil, R.S., Nagy, Z., Dolan, R.J. & Rees, G. 2010 Relating introspective accuracy to individual differences in brain structure. *Science* **329**, 1541-1543.
- [20] Song, C., Kanai, R., Fleming, S.M., Weil, R.S., Schwarzkopf, D.S. & Rees, G. 2011 Relating inter-individual differences in metacognitive performance on different perceptual tasks. *Consciousness and Cognition* **20**, 1787-1792.
- [21] Ais, J., Zylberberg, A., Barttfeld, P. & Sigman, M. 2016 Individual consistency in the accuracy and distribution of confidence judgments. *Cognition* **146**, 377-386.
- [22] Rouault, M., McWilliams, A., Allen, M. & Fleming, S. 2018 Human metacognition across domains: insights from individual differences and neuroimaging. *Personality Neuroscience* **1**.
- [23] Aitchison, L., Bang, D., Bahrami, B. & Latham, P.E. 2015 Doubly Bayesian analysis of confidence in perceptual decision-making. *PLoS Computational Biology* **11**, e1004519.

- [24] Bang, D., Fusaroli, R., Tylén, K., Olsen, K., Latham, P.E., Lau, J.Y., Roepstorff, A., Rees, G., Frith, C.D. & Bahrami, B. 2014 Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and cognition* **26**, 13-23.
- [25] Behrens, T.E., Hunt, L.T., Woolrich, M.W. & Rushworth, M.F. 2008 Associative learning of social value. *Nature* **456**, 245-249.
- [26] Mahmoodi, A., Bang, D., Olsen, K., Zhao, Y.A., Shi, Z., Broberg, K., Safavi, S., Han, S., Ahmadabadi, M.N. & Frith, C.D. 2015 Equality bias impairs collective decision-making across cultures. *Proceedings of the National Academy of Sciences* **112**, 3835-3840.
- [27] Allen, M., Glen, J.C., Mullensiefen, D., Schwarzkopf, D.S., Fardo, F., Frank, D., Callaghan, M.F. & Rees, G. 2017 Metacognitive ability correlates with hippocampal and prefrontal microstructure. *Neuroimage* **149**, 415-423.
- [28] Navajas, J., Hindocha, C., Foda, H., Keramati, M., Latham, P.E. & Bahrami, B. 2017 The idiosyncratic nature of confidence. *Nature Human Behaviour* **1**, 810-818.
- [29] Morgan, T., Rendell, L., Ehn, M., Hoppitt, W. & Laland, K. 2012 The evolutionary basis of human social learning. *Proceedings of the Royal Society of London B: Biological Sciences* **279**, 653-662.
- [30] Bonaccio, S. & Dalal, R.S. 2006 Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes* **101**, 127-151.
- [31] Brier, G.W. 1950 Verification of forecasts expressed in terms of probability. *Monthly Weather Review* **78**, 1-3.
- [32] Fleming, S.M. & Lau, H.C. 2014 How to measure metacognition. *Frontiers in Human Neuroscience* **8**, 443.
- [33] Richerson, P.J. & Boyd, R. 2005 *Not by genes alone*, University of Chicago Press.
- [34] Yaniv, I. 2004 The benefit of additional opinions. *Current directions in psychological science* **13**, 75-78.
- [35] Yaniv, I. & Kleinberger, E. 2000 Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational behavior and human decision processes* **83**, 260-281.
- [36] Bang, D., Aitchison, L., Moran, R., Castanon, S.H., Rafiee, B., Mahmoodi, A., Lau, J.Y., Latham, P.E., Bahrami, B. & Summerfield, C.J.N.H.B. 2017 Confidence matching in group decision-making. *Nature Human Behaviour* **1**, 0117.
- [37] Morales, J., Lau, H. & Fleming, S.M. 2018 Domain-general and domain-specific patterns of activity supporting metacognition in human prefrontal cortex. *Journal of Neuroscience* **38**, 3534-3546.
- [38] Fleming, S.M., Huijgen, J. & Dolan, R. 2012 Prefrontal contributions to metacognition in perceptual decision making. *Journal of Neuroscience* **32**, 6117-6125.
- [39] Paulus, M., Tsalas, N., Proust, J. & Sodian, B. 2014 Metacognitive monitoring of oneself and others: Developmental changes during childhood and adolescence. *Journal of Experimental Child Psychology* **122**, 153-165.
- [40] Weil, L.G., Fleming, S.M., Dumontheil, I., Kilford, E.J., Weil, R.S., Rees, G., Dolan, R.J. & Blakemore, S.-J. 2013 The development of metacognitive ability in adolescence. *Consciousness and Cognition* **22**, 264-271.
- [41] Mesoudi, A., Chang, L., Dall, S.R. & Thornton, A. 2016 The evolution of individual and cultural variation in social learning. *Trends in Ecology & Evolution* **31**, 215-225.
- [42] Heyes, C. 2018 Enquire within: cultural evolution and cognitive science. *Philosophical Transactions of the Royal Society B: Biological Sciences* **373**, 20170051.

Electronic Supplementary Materials

Knowing whom to learn from: individual differences in metacognition and weighting of social information

Karsten Olsen*, Andreas Roepstorff & Dan Bang

*Corresponding author: karsolsen@gmail.com

Methods

Display parameters and response mode

Each participant viewed an LCD display (Fujitsu Siemens AMILO SL 3220W, 22"; resolution, 1280×1020; background luminance, 62.5 Cd/m²) at a distance of about 57 cm. Stimulus presentation was controlled by the Cogent Toolbox (www.vislab.ucl.ac.uk/cogent.php) for MATLAB. Participants responded using the arrow keys on a QWERTY keyboard (left arrow key, move left; right arrow key, move right; down arrow key, confirm).

Stimulus construction

The circular aperture (8 degrees in diameter) was centred on the fixation cross (0.2 degrees in diameter). The motion stimulus was made up of three independent sets of dots (each dot was 0.12 degrees in diameter) shown in consecutive frames. Each set of dots was shown for one frame (about 16ms) and then replotted again three frames later (about 50ms). Some dots were displaced in the direction of motion at a speed of 2 degrees s⁻¹ while the rest of the dots were displaced at random locations within the aperture. Coherence refers to the fraction of dots displayed in the motion direction. The use of three independent sets of dots means that the positions of dots in frame one, for example, are only correlated with the positions of dots in frame four, seven, eleven and so forth. The dot density was fixed at 16 dots degrees⁻² s⁻¹. The number of dots in each frame was thus: $16 * \pi * 8^2 * (1/60) \approx 53$. To help participants keep fixation, a black circular mask (0.7 degrees in diameter) centred on the fixation cross was shown during the presentation of the motion stimulus. The two motion directions (left, 180 degrees; right, 0 degrees) were counterbalanced across the trials.

Advisor simulation

We used a signal detection model to simulate the two advisors. On each trial, the advisor receives noisy sensory evidence, x , sampled from a Gaussian distribution, $x \in N(d\theta, \sigma^2)$. The mean of the distribution is given by a motion direction, $d \in \{-1, \text{left}; 1, \text{right}\}$, and a motion coherence, θ , which was a fixed value determined in the calibration session. The standard deviation of the distribution, σ , specifies the level of sensory noise and was chosen so as to achieve levels of accuracy that were lower (bad advisor: $z - .12$) or higher (good advisor: $z + .12$) than the mean accuracy (z) of the participant in the calibration session. We thresholded the vector of sensory evidence for each advisor so as to achieve a uniform distribution across responses (i.e. "left, 100%", "left, 90%" and so forth were each selected on a twelfth of the trials). This approach has been used previously to simulate realistic agents in perceptual decision tasks [1].

References

- [1] Bang, D., Aitchison, L., Moran, R., Castanon, S.H., Rafiee, B., Mahmoodi, A., Lau, J.Y., Latham, P.E., Bahrami, B. & Summerfield, C.J.N.H.B. 2017 Confidence matching in group decision-making. *Nature Human Behaviour* **1**, 0117.