Social influence and informational independence

Justin Sulik (justin.sulik@gmail.com)
Cognition, Values & Behavior, Ludwig Maximilian University of Munich

Bahador Bahrami
Department of Psychology, Ludwig Maximilian University of Munich; Royal Holloway, University of London

Ophelia Deroy
Faculty of Philosophy & Munich Center for Neuroscience, Ludwig Maximilian University of Munich

Abstract

We frequently use social information when making decisions. For instance, other people may know more about a problem than we do, so we might update our initial beliefs in light of their opinions. The epistemic value of these social cues depends in part on their informational independence. People should thus be sensitive to nonindependence in their weighting of social information. However, the current literature yields conflicting results. In one recent study, participants valued social information less when it was nonindependent; in another, participants were insensitive to nonindependence. We identify possible causes of this inconsistency, and present an experimental paradigm that aims to fill these gaps. Then, in a study (N=200) with pre-registered hypotheses and analyses, we find that participants were not sensitive to cue dependence. We highlight the relevance of this finding for the modern media context, where nonindependence of both traditional and social media sources can lead to the spread of bias or false belief.

Keywords: decision making; social information; cognitive bias; belief updating; independence;

When making decisions in the face of uncertainty, we often benefit from learning other people's beliefs (Boyd, Richerson, & Henrich, 2011; Farrell, 2011; Heyes, 2018; Mayo-Wilson, Zollman, & Danks, 2013; Morgan & Laland, 2012; Toyokawa, Whalen, & Laland, 2019; Tump, Wolf, Krause, & Kurvers, 2018). On the other hand, this social information can lead to groupthink (Kerr & Tindale, 2004), herding (Lorenz, Rauhut, Schweitzer, & Helbing, 2011), information cascades (Anderson & Holt, 1997), failure to explore sufficiently (Yahosseini, Reijula, Molleman, & Moussaïd, 2018; Zollman, 2010), or suppression of useful information (Gigone & Hastie, 1993; Stasser & Titus, 1985).

Optimal epistemic outcomes thus require a balance between the costs and benefits of social learning (Bernstein, Shore, & Lazer, 2018; Rendell et al., 2011; Toyokawa et al., 2019; Yahosseini & Moussaïd, 2019). One factor affecting the value of social learning — and thus the above balance — is informational independence (Hahn, von Sydow, & Merdes, 2019; Ladha, 1992). If you heard from ten people that they all predicted the same party would win an election, you would be less impressed by this apparent consensus if you found out that they had all just read it in the same newspaper than if they had all formed this belief independently.

The question, then, is whether people are sensitive to informational independence when incorporating social cues into their own beliefs. This is a pressing question, given the concentration of media outlets in the hands of a small group of moguls (Sweney, 2015, October 21), or the grave potential for the spread of misinformation online (Acerbi, 2019; Harvey et al., 2018; Pennycook et al., 2019).

Unfortunately, recent work yields conflicting results. Mercier and Miton (2019, MM) found that participants were sensitive to informational independence, while Yousif, Aboody, and Keil (2019, YAK) found that they were not. We aim to identify possible reasons for this inconsistency, and offer a new experimental paradigm that balances ecological validity with experimental control. Then, in an experimental study with pre-registered hypotheses and analyses, we test whether participants are sensitive to cue independence.

In both focal studies (MM, YAK) participants saw social information before making a judgment. Using 'source' to mean the ultimate or primary source of the information, most distant from the point of view of the participant, and 'cue' to mean the secondary source, nearest the participant, both studies contrasted a condition with informational dependence (with cues reflecting the same source) and a condition with informational independence (Fig. 1a).

The stimuli in MM were vignettes where social cues (friends Julia, Rob and Olivia) either reported their opinion of a restaurant (e.g., Julia: 'I don't think it's great.'), or reported their opinion as well as that a common source (Julia: 'I don't think it's great. Peter told me it was disappointing.'). MM found that participants weighted social information less when their social cues were dependent than when they were independent. The cues in YAK were newspapers discussing the future of the Japanese economy. The sources were economists quoted in the newspapers. In the dependent condition, all newspapers that were positive about the economy quoted the same economist, and all that were negative quoted another. In the independent condition, each newspaper quoted a separate economist. YAK found that participants' responses were not sensitive to cue independence.

Neither schema in Fig. 1a offers a clean test of our question: whether people are sensitive to informational independence in social information. Firstly, the type of information presented by MM is not consistent across conditions (which could just as well be described as 'source present' vs. 'source absent'). A plausible implication is that the cues' opinions are based on hearsay in the dependent condition, but on direct experience in the independent condition. Thus, their results might be driven by epistemic modality rather than source in-

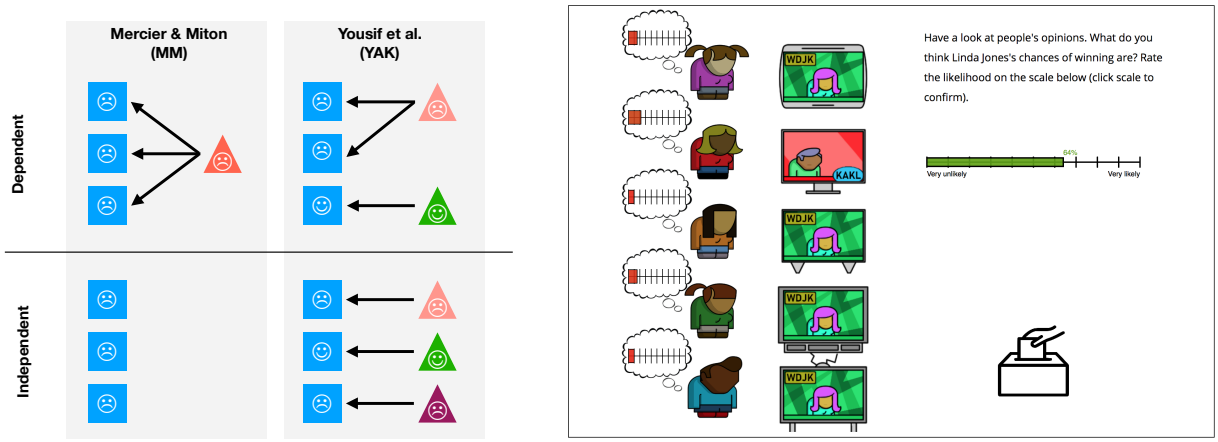

(a) (b)

Figure 1: (a) Schematic representations of the informational dependencies in the focal studies. Triangles represent sources and squares represent cues. Emoji symbolise positive or negative opinions about a restaurant (MM) or the future of the Japanese economy (YAK). Arrows represent the flow of information. (b) Our experimental interface. Prior to this point, the participant has been introduced to the town where these characters live, has been told about Linda Jones running for mayor, and has provided their initial estimate of Linda's chances of winning. Now, they are being shown the social information. Here, there is high informational dependency: most townspeople are watching the same channel. The townspeople's opinions are shown using the same response scale that the participant uses, varying from 'very unlikely' to 'very likely'. To highlight these responses, the scale turns redder for lower values and greener for higher values. It is clear that the people think Linda's chances are low. The participant is currently providing their final likelihood estimate (here, estimating Linda's chances at 64%). The ballot icon (lower right) reminds participants that they are predicting an election outcome.

dependence. The cues' opinions and their choice of source are perfectly correlated in YAK. It is not clear whether the newspapers are reporting neutrally what relevant experts have said, or whether the editors are taking a stance, and have intentionally chosen to quote experts that agree with their stance. This may imply different things about the source's expertise. Since assumptions about neutrality or source expertise may compete with source independence in evaluating social information, we manipulate whether participants are told that the cues intentionally chose their sources, or that the sources were randomly assigned.

Unlike MM and YAK, we make a firm distinction between cues and sources: our participants are told each cue's opinion, and are shown what source each cue got their information from, but are not told the sources' opinions. Thus, our schema is close to YAK (Fig. 1a), except that our sources' opinions are not revealed. This allows us to independently manipulate informational dependence and social opinions, while keeping epistemic modality constant. In our paradigm, participants are given social information about cartoon people who are watching the news on TV (Fig. 1b). The cartoons' opinions are revealed after the news broadcast, and participants can straightforwardly see what channel each person is watching (thus indicating informational dependencies), but the participant cannot hear the news anchors talking.

In addition to these schematic issues, there are differences between the focal studies that may well be confounds. For instance, the null results in YAK might reflect the cognitive effort required for reading long, technical texts (or reflect concomitant inattentiveness). In our paradigm (Fig. 1b), text is kept to a minimum, and information about social cues and their (in-)dependence is presented in an accessible visual manner. Thus, if we find no sensitivity to informational independence, we can rule out cognitive effort as the cause.

MM measure beliefs (how good the restaurant is), whereas YAK measure confidence in beliefs (how confident they are in the belief that the economy will improve). To illustrate the distinction: you can believe there is a 50% chance that a fair coin will come up heads, but have 100% confidence that this is correct. Consequently, we collect both belief estimates and confidence. If confidence is less influenced by social information than the beliefs themselves, that might explain the contrasting results in those studies.

Both focal studies only measure participants' final beliefs, after seeing social information. However, our core concern is social weighting: how much participants adjust towards social information. Thus, in line with Molleman, Kurvers, and van den Bos (2019), we employ a multi-stage approach, where participants provide an initial belief, then receive social information, then provide a final belief. This affords explicit measurement of how much people adjust their beliefs in light of social information, and thus how much weight they have accorded to the social information (Molleman et al., 2019).

Finally, there are stable individual differences in people's

tendency to make use of social information (Molleman et al., 2019; Toelch, Bruce, Newson, Richerson, & Reader, 2014), though Olsen, Roepstorff, and Bang (2019) identify some contexts in which this tendency can be more flexible. The question, then, is whether informational independence is such a context. If not, these stable tendencies might explain the insensitivity observed by YAK. To test this, we incorporate a measure of participants' baseline social information use.

In sum, recent evidence concerning people's sensitivity to informational dependency in social cues is inconclusive. We have identified several possible explanations of this inconsistency, and have outlined how our paradigm addresses these gaps. We pre-registered four hypotheses (https://osf.io/9pmqy). (H1) If people's evaluation of social cues is sensitive to informational dependency, participants will adjust their estimates more when sources of social information are diverse than when they are homogeneous; (H2) Participants will be sensitive to whether sources are intentionally selected vs. randomly assigned; (H3) The effect of informational dependency will be larger for people's likelihood ratings than for their confidence ratings; (H4) By including a measure of baseline social information use as a covariate in a regression, we will be better able to detect an effect of condition.

Methods

Participants

We recruited 200 participants (mean age 37; 103 male, 92 female, 5 self-described or skipped the question on gender) via Amazon's Mechanical Turk platform. Our sample size calculation was pre-registered (https://osf.io/9pmqy). Participation was managed by Turkprime (Litman, Robinson, & Abberbock, 2017), and was limited to those with IP addresses in the USA, with approval ratings higher than 95%. Participants provided informed consent, and the study received ethical approval from the University of London. Participants were paid $2.20 (mean study duration: 12.4 minutes).

Procedure

A demonstration of the experiment is available at `https://guarded-coast-18242.herokuapp.com/demo`. Participants undertook a simple English test, motivated by worries about data quality on Mechanical Turk, driven by users with poor English (Kennedy, Clifford, Burleigh, Waggoner, & Jewell, 2018). We excluded data from 3 participants who scored lower than 80% on this test (a pre-registered criterion).

Participants then undertook 5 trials of the BEAST task (Molleman et al., 2019), to provide an index of their baseline tendency to use social information. On each trial, 50–130 small images of animals were presented for 6 seconds. Participants provided an estimate of how many animals they saw. They were then shown the estimate of another person (from a pilot study). Participants then provided a final estimate. Their baseline social information use is the difference between their initial and final estimate, as a proportion of the difference between their initial estimate and the social information, averaged across trials.

Next, participants practiced using our response scales to rate whether or not a scenario is morally acceptable (e.g., spanking children) or likely (e.g., a Democrat becoming the next governor of their state), and to rate confidence. They did four practice trials. Two of these had an intersubjectively agreed answer (whether genocide is bad, and whether Ronald McDonald is likely to become next president of the USA). We pre-registered an exclusion criterion for these trials, and excluded 27 participants (2 of whom had failed the English test anyway). This exclusion does not alter our conclusions.

Then, in our experimental task, participants were introduced to three fictional American towns. For each town, they were given a moral scenario (e.g., 'The people of Greenville have been thinking about whether cannabis should be decriminalized. In general, how do you view decriminalization of cannabis?') and a prediction scenario (e.g., 'Bob Smith is one of two candidates running for mayor of Greenville. His platform involves increasing the town's education budget, building more bicycle paths, and instituting equal pay. How likely do you think Bob Smith is to win?').

For each scenario, they provided an initial estimate of their belief and confidence. They were shown 5 cartoon townspeople (the social cues). They were told that the townspeople get all their information from TV (the sources). Informational independence was manipulated by varying the number of channels shown (e.g., only two channels in Fig. 1b). The TVs displayed animated news anchors. The townspeople's beliefs concerning the scenario were shown with the same scale that the participant had used. We checked that participants were attending to the social information (e.g., asking them to click on the townsperson with the strongest belief). Participants provided a final estimate of their belief and confidence.

To reduce the chance that participants became aware of the key manipulation, four of the six trials were distractors (all three moral scenarios, and one of the prediction scenarios), where the townspeople's responses might agree or disagree with the participant's initial response, and where the number of TV channels shown varied between two and four. In the experimental trials (two prediction scenarios), the townspeople were shown disagreeing with the participant's initial estimate, so that there was ample scope for the participant to adjust their opinion in the direction of the social information. As a within-subjects condition, one of these experimental trials had high informational independence (4 channels shown, so two townspeople were watching the same channel) and one had low informational independence (2 channels shown, so 4 townspeople were watching the same channel). As a between-subjects condition, participants were either told that the channels were randomly selected, or that the townspeople had chosen their most trusted channel. Finally, participants provided basic demographic information.

Results

We begin with pre-registered analyses (we depart from the pre-registration only in using mixed-effects regressions in place of fixed-effects regressions at the request of a reviewer; this has no effect on our conclusions). Our core question is whether participants' weighting of social information is sensitive to informational diversity. In particular, if MM are right, we would expect that participants weight social information less when there is low informational independence than when there is high informational independence (H1). To measure people's weighting of social information, we calculated their belief adjustment in a similar way to the aforementioned BEAST task (Molleman et al., 2019): the difference between final and initial estimate, as a proportion of the difference between the initial estimate and the mean of the social cues. We modelled the effect of informational independence on social weighting with a linear regression. The effect of informational diversity was not significant (linear regression $b = -0.01$ bootstrapped 95% CIs $[-0.06, 0.04]$, $SE = 0.03$, $t = -0.29$, $p = .78$), and the difference between conditions was small (Fig. 2a).

We wondered if the different results in the focal studies might be due to participants making inferences about source quality (e.g., that a newspaper had quoted a particular economist because of their expertise in the matter). If so, we hypothesized that participants' social weighting would be sensitive to source choice — whether the sources were intentionally chosen or randomly assigned (H2). We modelled the effect of source choice on social weighting with a linear regression. The effect of source choice was not significant (linear regression $b = 0.01$ bootstrapped 95% CIs $[-0.09, 0.11]$, $SE = 0.05$, $t = 0.22$, $p = .83$), and again, the difference between conditions was small (Fig. 2b).

If the null results in YAK are driven by their use of confidence ratings (vs. belief ratings in MM), then adjustment should be lower for confidence ratings than the other estimates (H3). Since there is no social information on confidence (unlike for the belief estimates for H1 & H2), we cannot use a proportional measure of adjustment. Instead, we calculated the absolute change in rating for each scale (belief estimate and confidence) and modeled the effect of scale type and informational diversity on absolute change. There was a significant main effect of scale type ($b = 0.18$ bootstrapped 95% CIs $[0.15, 0.21]$, $SE = 0.02$, $t = 10.79$, $p < .001$). Participants adjusted their belief estimates significantly more than their confidence ratings (Fig. 2c). However, there was no significant main effect of source diversity ($b = 0.01$ bootstrapped 95% CIs $[-0.01, 0.04]$, $SE = 0.02$, $t = 0.53$, $p = .59$) and no significant interaction ($b = -0.02$ bootstrapped 95% CIs $[-0.07, 0.02]$, $SE = 0.023$, $t = -0.8$, $p = .42$). Thus, the insensitivity to informational diversity is not explained by the different response types.

Finally, since there are individual differences in people's tendency to use social information, we wondered if including a measure of this baseline tendency as a covariate in our regressions might change the above conclusions (H4). We modeled the effects of source diversity and the BEAST measure on social weighting (outcome measure calculated as for H1, H2) with a linear regression. There was no significant main effect of the BEAST measure ($b = 0$, bootstrapped 95% CIs $[-0.37, 0.39]$, $SE = 0.2$, $t = -0.02$, $p = .986$). Including the BEAST measure in the regression did not reveal a main effect of source diversity ($b = 0.04$, bootstrapped 95% CIs $[-0.11, 0.18]$, $SE = 0.07$, $t = 0.54$, $p = .58$). There was no significant interaction ($b = -0.16$, bootstrapped 95% CIs $[-0.57, 0.30]$, $SE = 0.21$, $t = -0.78$, $p = .43$, Fig. 2d).

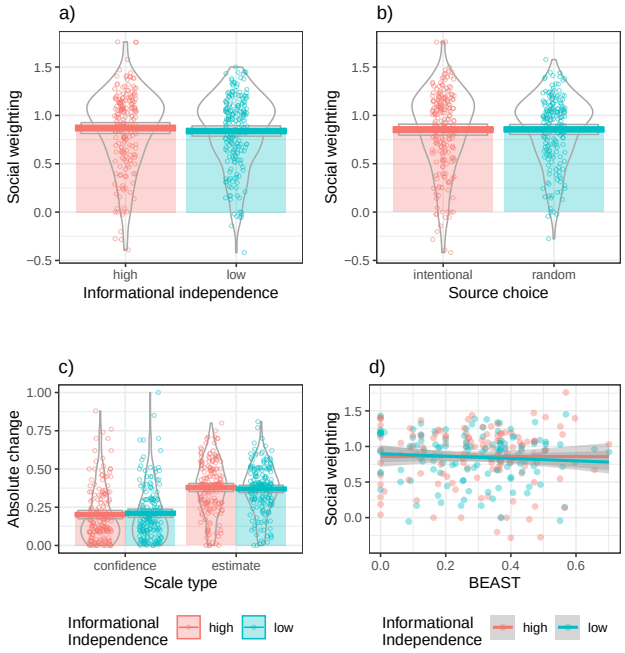

Figure 2: Plots of pre-registered hypotheses (a) Participants' social weighting (proportional adjustment) was not sensitive to informational dependence; (b) Participants' social weighting was not sensitive to how the sources were selected; (c) Participants adjusted the confidence scales less than the belief estimate scales; (d) Participants' baseline social weighting tendency (the BEAST measure) did not predict their social weighting in our experimental tasks.

We shift now to exploratory analyses. Our conclusions above support the claim by YAK that people are not generally sensitive to informational diversity (in their terms, do not distinguish between a true and a false consensus). However, what reason do we have to suppose that these results are not driven by inattentive or random responding? Here, we examine signatures of data quality, checking that this null result reflects people's sincere responses.

People adjusted towards the social cues overall (mean proportional change > 0, Fig. 2a). If they were disregarding the social cues or answering randomly, this would be un-

likely. Further, we included practice trials with intersubjectively correct answers (see Procedure). 87.5% of participants responded correctly to these questions, and the remaining 13.5% (a failure rate not at all remarkable for online studies) were dropped from analysis.

We compare experimental trials (election predictions, where social cues disagreed with the participant's initial estimate) and distractor trials (either moral issues, or where social cues agreed with the participant). We expect participants to be more confident in their moral beliefs than in their election predictions. Participants reported higher initial confidence in their morality judgments (mean confidence = 0.85) than in their election predictions (mean confidence = 0.61, bootstrapped 95% CIs for difference in means: [0.20, 0.27]). We expect moral beliefs to be less susceptible to change. Participants adjusted their beliefs for prediction trials (mean change = 0.29) more than for morality trials (mean change = 0.16, bootstrapped 95% CIs for difference in means: [0.11, 0.15]). We expect that participants would adjust less for agreement trials than for disagreement ones. The change in belief was lower for agreement trials (mean change = 0.09) than for disagreement ones (mean change = 0.36, bootstrapped 95% CIs for difference in means: [0.24, 0.28]).

Finally, we check for ideological consistency. People approving of tougher immigration policy were more likely to disapprove of tougher gun-licence background checks (linear regression $b = -0.55$, $SE = 0.08$, $t = -6.99$, $p < .001$); were less likely to think that a mayoral candidate running on a liberal platform would win ($b = -0.11$, $SE = 0.04$, $t = -3.08$, $p = .002$); but were more likely to think that a candidate with a conservative platform would win ($b = 0.12$, $SE = 0.05$, $t = 2.71$, $p = .007$). Thus, their beliefs appear consistent. Overall, this pattern of responses suggests that participants responded sincerely and honestly.

Discussion

If people integrate social information into their individual beliefs, then the weighting of the information should optimally be sensitive to its informational independence (Hahn et al., 2019; Ladha, 1992). However, recent empirical work on this has yielded contradictory results: (Yousif et al., 2019, YAK) found that people were not sensitive to independence; (Mercier & Miton, 2019, MM) that they were. We have developed an experimental paradigm that aims to overcome some gaps in these focal studies, and to explore whether certain differences between them might explain this contradiction.

Our data support the finding by YAK: our participants were not sensitive to informational independence in updating their beliefs. None of the potential factors we identified (whether intentionally choosing a source implies anything about source quality; whether rating beliefs rather than confidence in beliefs makes a difference; and whether people's baseline tendency to weight social information explains their insensitivity to our main manipulation) did anything to alter this conclusion. We presented several exploratory analyses suggesting that our participants responded sincerely and honestly, in which case this result seems to be a genuine failure to incorporate the relevant information into their decision making.

Why, then, did we find no significant difference, unlike MM? We suggested that MM's design implies a difference in epistemic modality (direct experience vs. hearsay). Their focus is on 'evolutionary valid cues', ones that our brains evolved to attend to and to process. They argue that hearsay is such a cue. Without disagreeing with this specific claim, we note that this represents a confound for our particular question, since the kind of information presented is not consistent across their conditions (Fig. 1a). In more recent work (Altay, Claidière, & Mercier, 2020), Mercier and colleagues argue that the 'friend of a friend' attribution is a important factor in social transmission. In that case, our choice of stimuli — people watching TV, rather than people reporting their friends' beliefs — might explain our results.

Why should it matter if people are insensitive to nonindependence? Nonindependence plays a role in social phenomena such as denial of climate change (Harvey et al., 2018), but it is just one of many factors, and the interaction between such factors is doubtless complex (Hahn, Harris, & Corner, 2016). To have any hope of understanding such complex phenomena 'in the wild', it is important to also try understand the individual factors, such as sensitivity to informational independence, in controlled conditions. For instance, an agent-based simulation shows that throttling social information, by increasing independence, boosts group epistemic outcomes — a benefit of 'transient diversity' (Zollman, 2010). However, if human biases mean that we fail to exploit the benefits of such diversity, then advice gained from normative approaches (such as computational models) might not translate into concrete benefits.

Turning to the broader context, is it surprising that people were insensitive to statistical nonindependence? Kahneman and Tversky (1972) showed that people do not typically incorporate sample size into their judgments of probability. Insensitivity to the nonindependence of a sample – much like the size of the sample – may just mean that human cognition is prone to process statistical information suboptimally.

Finally, our main motivation in including the BEAST measure (Molleman et al., 2019) was to provide a baseline measure of individuals' tendency to use social information. It is somewhat surprising that this measure did not correlate with social information weighting in our experimental task, especially since Toelch et al. (2014) report a moderate correlation ($r = .29$) in individuals' social information across two quite different tasks. We note briefly that, unlike Molleman et al., we did not remind participants of their own initial estimates while they were making their final estimate, but otherwise leave this problem for future research.

Acknowledgements

We gratefully acknowledge support by the NOMIS Foundation through the project "Diversity in Social Environments".

References

Acerbi, A. (2019). Cognitive attraction and online misinformation. *Palgrave Communications*, *5*(1), 15.

Altay, S., Claidière, N., & Mercier, H. (2020, January 25). It happened to a friend of a friend: Inaccurate source reporting in rumor diffusion. *PsyArXiv*. doi: https://doi.org/10.31234/osf.io/5czka

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American economic review*, 847–862.

Bernstein, E., Shore, J., & Lazer, D. (2018). How intermittent breaks in interaction improve collective intelligence. *PNAS*, *115*(35), 8734–8739.

Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *PNAS*, *108*(Supplement 2), 10918–10925.

Farrell, S. (2011). Social influence benefits the wisdom of individuals in the crowd. *PNAS*, *108*(36), E625–E625.

Gigone, D., & Hastie, R. (1993). The common knowledge effect: Information sharing and group judgment. *Journal of Personality and Social Psychology*, *65*(5), 959–974.

Hahn, U., Harris, A. J., & Corner, A. (2016). Public reception of climate science: Coherence, reliability, and independence. *Topics in Cognitive Science*, *8*(1), 180–195.

Hahn, U., von Sydow, M., & Merdes, C. (2019). How communication can make voters choose less well. *Topics in Cognitive Science*, *11*(1), 194–206.

Harvey, J. A., Van Den Berg, D., Ellers, J., Kampen, R., Crowther, T. W., Roessingh, P., ... Mann, M. E. (2018). Internet blogs, polar bears, and climate-change denial by proxy. *BioScience*, *68*(4), 281–287.

Heyes, C. (2018). Enquire within: cultural evolution and cognitive science. *Phil. Trans. R. Soc. B*, *373*(1743), 20170051.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, *3*(3), 430–454.

Kennedy, R., Clifford, S., Burleigh, T., Waggoner, P., & Jewell, R. (2018). The shape of and solutions to the MTurk quality crisis. *Available at SSRN*. doi: http://dx.doi.org/10.2139/ssrn.3272468

Kerr, N. L., & Tindale, R. S. (2004). Group performance and decision making. *Annu. Rev. Psychol.*, *55*, 623–655.

Ladha, K. K. (1992). The Condorcet jury theorem, free speech, and correlated votes. *American Journal of Political Science*, 617–634.

Litman, L., Robinson, J., & Abberbock, T. (2017). Turkprime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*(2), 433–442.

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *PNAS*, *108*(22), 9020–9025.

Mayo-Wilson, C., Zollman, K., & Danks, D. (2013). Wisdom of crowds versus groupthink: learning in groups and in isolation. *International Journal of Game Theory*, *42*(3), 695–723.

Mercier, H., & Miton, H. (2019). Utilizing simple cues to informational dependency. *Evolution and Human Behavior*, *40*(3), 301–314.

Molleman, L., Kurvers, R., & van den Bos, W. (2019). Unleashing the BEAST: a brief measure of human social information use. *Evolution and Human Behavior*, *40*(5), 492–499.

Morgan, T. J. H., & Laland, K. N. (2012). The biological bases of conformity. *Frontiers in neuroscience*, *6*, 87.

Olsen, K., Roepstorff, A., & Bang, D. (2019, March 24). Knowing whom to learn from: individual differences in metacognition and weighting of social information. *PsyArXiv*. doi: https://doi.org/10.31234/osf.io/jqheu

Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G. (2019). Understanding and reducing the spread of misinformation online. *PsyArXiv*. doi: https://doi.org/10.31234/osf.io/3n9u8

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*(2), 68–76.

Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, *48*(6), 1467.

Sweney, M. (2015, October 21). *UK media plurality threatened by dominant group of large firms – report.* The Guardian. Retrieved from `https://www.theguardian.com/media/2015/oct/21/uk-media-plurality-threatened-by-dominant-group-of-large-firms-report`

Toelch, U., Bruce, M. J., Newson, L., Richerson, P. J., & Reader, S. M. (2014). Individual consistency and flexibility in human social information use. *Proceedings of the Royal Society of London B*, *281*(1776), 20132864.

Toyokawa, W., Whalen, A., & Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, *3*(2), 183–192.

Tump, A. N., Wolf, M., Krause, J., & Kurvers, R. H. (2018). Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information. *Journal of the Royal Society Interface*, *15*(142), 20180155.

Yahosseini, K. S., & Moussaïd, M. (2019). The social dynamics of collective problem-solving. *bioRxiv*. doi: https://doi.org/10.1101/771014

Yahosseini, K. S., Reijula, S., Molleman, L., & Moussaïd, M. (2018). Social information can undermine individual performance in exploration-exploitation tasks. In C. Kalish, M. Rau, J. Zhu, & T. Rogers (Eds.), *Cogsci 2018* (pp. 2473–2478).

Yousif, S. R., Aboody, R., & Keil, F. C. (2019). The illusion

of consensus: A failure to distinguish between true and false consensus. *Psychological Science*, *30*(8), 1195–1204.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, *72*(1), 17–35.