Collectively jumping to conclusions: social information amplifies the tendency to gather insufficient data

Justin Sulik

Cognition, Values & Behavior
Ludwig Maximilian University of Munich

Charles Efferson

Faculty of Business and Economics
University of Lausanne

Ryan McKay

Department of Psychology
Royal Holloway, University of London

False beliefs can spread within societies even when they are costly and when individuals share access to the same objective reality. Research on the cultural evolution of misbeliefs has demonstrated that a social context can explain *what* people think, but not whether it also explains *how* people think. We shift the focus from the diffusion of false beliefs to the diffusion of suboptimal belief-formation strategies, and identify a novel mechanism whereby misbeliefs arise and spread. We show that, when individual decision-makers have access to the data-gathering behaviour of others, the tendency to make decisions on the basis of insufficient evidence is amplified, increasing the rate of incorrect, costly decisions. We argue that this mechanism fills a gap in current explanations of problematic, widespread misbeliefs such as climate change denial.

Introduction

Humans are inveterate misbelievers (McKay & Dennett, 2009), prone to false beliefs about the self or the world. Misbeliefs may spread or be amplified in social groups (Efferson, McKay, & Fehr, 2020), often in the face of disconfirmatory evidence, and even when the beliefs are harmful or costly. Such beliefs — what we might call "collective delusions" — can have catastrophic social consequences. For example, pervasive suspicions that climate change is a hoax perpetrated by scientists seeking personal advantage, or that pharmaceutical companies hide data about the harmful side effects of vaccines (Lewandowsky, Gignac, & Oberauer, 2013;

Lewandowsky, Oberauer, & Gignac, 2013), decrease people's intentions to reduce their carbon footprint (Jolley & Douglas, 2014b) or to vaccinate (Jolley & Douglas, 2014a), potentially exacerbating environmental disasters, as well as public health emergencies such as the COVID-19 pandemic (Murphy et al., 2021).

To explain socially widespread misbeliefs, psychological research appeals to two main causal pathways (Fig. 1a). Though important, we will show that these pathways are insufficient to explain why collective delusions are robust to disconfirming evidence, or how they are maintained even when costly to the individuals holding those misbeliefs. We will argue for the importance of a further pathway (Fig. 1b), and then demonstrate experimentally that this pathway exists.

The first pathway in Fig. 1a is a matter of individual psychology, and the key idea is that some people have cognitive propensities that make them more prone to misbeliefs. For instance, a reliance on an intuitive rather than a reflective problem-solving style predicts science denial (Gervais, 2015; Lobato & Zimmerman, 2019), belief in the paranormal or conspiracies (Ståhl & van Prooijen, 2018; Wagner-Egger, Delouvée, Gauvrit, & Dieguez, 2018), and belief in fake news (Bronstein, Pennycook, Bear, Rand, & Cannon, 2019; Lazer et al., 2018). People who trust their intuitions also tend to make decisions with less evidence (Ross et al., 2016). This tendency to "jump to conclusions" (JTC) is thought to play a key role in the formation of delusional beliefs (Garety & Freeman, 2013).

(a)   (b)
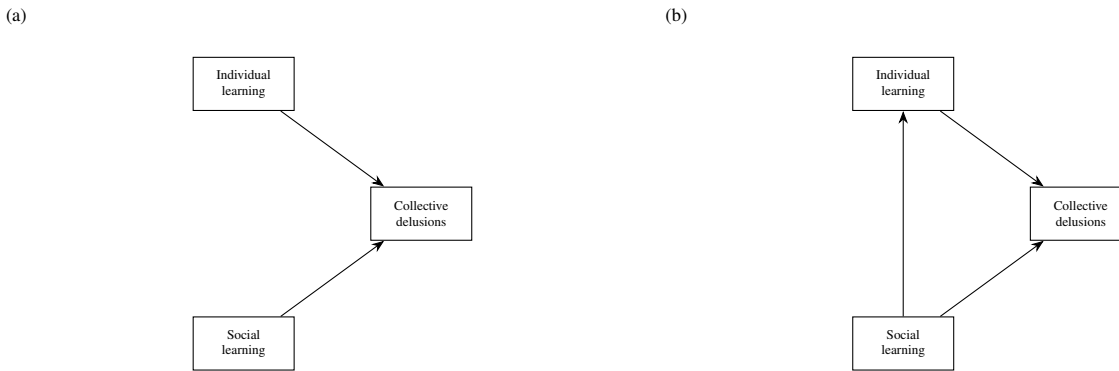

Figure 1. Causal models including individual factors, social factors, and their effects on false beliefs. (a) Most psychology research shows how individual and social learning both contribute, independently, to misbeliefs. (b) We aim to test, experimentally, whether a further pathway exists, whereby social information amplifies individual cognitive biases in a way that increase the risk of false beliefs. We note that further causal models are possible (e.g., a moderating effect of individual cognitive learning on the social pathway), but these do not concern us here since they do not involve a pathway from social information to individual cognition.

As stand-alone mechanisms, however, such biases do not provide a compelling account of collective delusions. In particular, a general tendency to adopt beliefs unreflectively or on limited evidence may lead individuals astray in random ways, but in itself will not support collective adherence to the same, specific false belief about the world.

Other types of cognitive bias, however, can boost the spread of beliefs with specific contents (Claidière & Sperber, 2007; Griffiths, Kalish, & Lewandowsky, 2008). For example, a particular belief (e.g., that vaccines or genetically modified foods are harmful) may spread and endure, in part, because it resonates with implicit cognitive proclivities and prejudices (Miton & Mercier, 2015), such as an essentialist bias that maps genetic modification onto tampering with an organism's essence, or a teleological bias that casts genetic modification as perverting the course of nature (Blancke, Van Breusegem, De Jaeger, Braeckman, & Van Montagu, 2015). Nevertheless, if certain types of content are simply intrinsically more appealing to human cognition in some universal sense, as per this account, then one would expect the resulting misbeliefs to be fairly ubiquitous, in which case something else is needed to explain why only some groups — but not others — arrive at collective misconstruals of reality.

Accordingly, the second pathway in Fig. 1a is social: people may espouse a false belief, not necessarily because of any particular deficiency in their belief-formation process, but because they are influenced by the beliefs of those around them (Asch, 1951; Cialdini & Goldstein, 2004; Sherif, 1935). In the examples of collective delusion discussed above, the relevant beliefs often cohere with the ideology of a social group whose membership is part of one's identity (Kahan, 2013; Rutjens, Heine, Sutton, & van Harreveld, 2018). For example, science denial correlates with both political and religious identity (Barone, Petto, & Campbell, 2014; Drummond & Fischhoff, 2017; Jylhä & Hellmer, 2020; Rutjens,

Heine, et al., 2018; Rutjens, Sutton, & van der Lee, 2018).

However, the evidence for a link between social identity and misbelief is typically correlational, rather than experimental, and thus cannot support claims about causation or explanation. What experimental evidence there is for the social pathway relies on somewhat artificial manipulations. For instance, in Sherif's autokinetic effect, in which the perception that a stationary point of light in an otherwise dark environment is moving can be influenced by the reports of others, participants could not rely on any empirical evidence to establish the truth of the matter. In contrast, collective delusions are held despite a wealth of empirical evidence attesting to the reality of climate change and vaccine safety. Similarly, the conformity experiments of Asch (1951, 1955) relied on deliberate deceit. However, such conformity is unstable without the continued influence of deceit (Jacobs & Campbell, 1961), whereas collective delusions seem to be sincerely held and stable.

Thus, although both social and individual pathways contribute to collective delusions, neither pathway on its own is adequate to explain key features of collective delusions. However, the problem runs even deeper: psychological research has tended to ignore the phenomenon of aggregation dynamics, whereby individual psychological processes scale up — in often unpredictable ways — to group-level phenomena. This aspect of social belief is more commonly studied in the field of Cultural Evolution.

Cultural Evolution posits similar dual pathways, distinguishing individual learning — where individuals gather their own information about the environment — from social learning — where they learn from other individuals (Boyd & Richerson, 1985; Laland, 2004; Morgan, Rendell, Ehn, Hoppitt, & Laland, 2011). In particular, cultural evolutionary theory suggests that when individuals conform to those around them, groups may converge either on an equilibrium in which

everyone has an accurate view of the world and behaves accordingly, or on some other equilibrium in which everyone has an inaccurate view and behaves accordingly (Boyd & Richerson, 1985; Efferson, Lalive, Richerson, McElreath, & Lubell, 2008; Efferson et al., 2020; Morgan & Laland, 2012). Thus, a combination of individual and social learning seems to explain why — sometimes, at least — a group can collectively form the same false belief.

Although early experimental work seemed to support this theoretical claim (Anderson & Holt, 1997), the misbeliefs formed by these processes were later shown to be unstable (Goeree, Palfrey, Rogers, & McKelvey, 2007). In particular, if even just a few people in the group engage in individual learning, gathering even a small amount of information about the actual state of the world, this has profound knock-on effects on aggregate social dynamics by destabilizing outcomes in which most people in the group are wrong, thus pushing the group toward an accurate understanding of their environment (Efferson et al., 2020; Goeree et al., 2007; Perreault, Moya, & Boyd, 2012).

Thus, the aggregate dynamics of individual and social learning do not explain how collective delusions are stable over time when there is even a trickle of accurate information about the environment. Indeed, one of the most puzzling features of collective delusions is how robust they are in the face of disconfirming evidence (e.g., that climate change is real and vaccines are beneficial). Thus, a core aspect of this phenomenon crying out for explanation is how collective delusions can be resistant to the steady flow of empirical evidence, since the findings from cultural evolution show that they are susceptible to just that. Clearly, then, something is missing from the schema in Fig. 1a.

There are thus challenges on multiple fronts for the current state of the art in experimental psychology: to explain how collective delusions persist even though a small trickle of evidence is able to destabilize the effect of social learning, and to explain how the same delusion can be widespread within one group but not another, if driven by individual biases. This is our point of departure in the present study. Here, we test whether another causal pathway exists, whereby social information affects the learning biases themselves (Fig. 1b), in addition to the beliefs those biases occasion. Can social context explain not just what people learn, but also how they learn? This question entails a shift in focus, from the diffusion of false beliefs to the diffusion of suboptimal learning strategies, and implicates a novel catalyst for the emergence of destructive, socially shared misbeliefs.

There is qualitative evidence (Metz, Weisberg, & Weisberg, 2018) and evidence from computational modeling (Druckman & McGrath, 2019) suggesting that this pathway is plausible. One specific mechanism suggested by such research is that people vary in their estimation of the value of relevant evidence, and this variation may depend on social factors. However, there is no experimental evidence demonstrating the existence of this pathway. We aim to show experimentally that social information can amplify an individual learning bias, such that people grow more likely to undervalue empirical evidence.

We do this by adapting a probabilistic reasoning task — the Beads Task (Huq, Garety, & Hemsley, 1988) — which is known to reveal biased data-gathering behavior in clinical and healthy populations: people frequently jump to conclusions, making decisions with insufficient evidence (Furl & Averbeck, 2011; Garety & Freeman, 2013; McLean, Mattiske, & Balzan, 2017; Ross, McKay, Coltheart, & Langdon, 2015; Van der Leer, Hartig, Goldmanis, & McKay, 2015). The task involves randomly drawing colored beads from an urn, in order to discover the color of the majority of beads in the urn. The task affords participants the opportunity to gather their own data about the objective state of the world (the majority color in the urn). However, as participants can decide how much data they would like to gather (specifically, how many beads they would like to draw), differences in data-gathering strategies may lead to participants holding different beliefs about the state of the world. If participants form false beliefs about the majority color in the urn because they do not gather enough evidence, it is not because the evidence is unavailable (as per Sherif, 1936) or deceptive (as per Asch, 1951).

To explore whether a social context can amplify this individual learning bias, we embed the Beads Task in a social transmission-chain procedure (Fig. 2a). In such a procedure, participants are randomly assigned to different chains or groups. Within each group, an initial participant (assigned to position 1 within their group) undertakes the Beads Task. Subsequently, a second participant (assigned to position 2 within their group) is given information about the first participant's data-gathering strategy before also undertaking the Beads Task. Thereafter, the participant in position 3 is given information about the data-gathering decisions of participants in positions 1 and 2, and so on. Fig. 2b illustrates how this might appear to a participant in position 11 within a particular chain. All participants within a chain are drawing from the same (virtual) urn, and thus face the same state of the world.

To evaluate the effect of social information, we contrasted the above procedure (which describes the social condition) with an asocial condition. In the asocial condition, participants were told what position they were in within their group (purely to make the task as similar as possible to the social condition), but were not given information about the decisions of previous participants. Thus, the presence or absence of social information is the only difference between conditions.

We tested whether individual tendencies to jump to conclusions on limited evidence would increase when those indi-

Figure 2. (*a*) A schematic illustration of the transmission-chain procedure. 2000 participants are divided into 100 groups, with participants in each group being arranged in a chain of 20 positions or generations. In the social condition (odd-numbered chains, green), participants are given information about the data-gathering decisions of previous participants who had done the task. In the asocial condition (even-numbered chains, red) participants are given minimal information: they are just told their position within the group. (*b*) A screenshot of the experiment interface for a participant in the social condition in position 11 of their chain. The number of beads seen by participants in each of the previous 10 positions are shown with black dots. In this example, the focal participant has already drawn three beads (which happen to have been green, pink, pink, recorded in the final column), so they have already spent $0.03 of their $0.15 endowment. They are reminded of the costs and bonuses and asked to decide whether they want to pay to see another bead, or whether they have seen enough. The asocial control condition would appear similar, except that the previous positions would lack any black dots indicating the number of beads seen by each previous participant. The horizontal blue line is merely a convenience to help participants compare their decisions with those of previous positions.

viduals can observe the evidence-gathering behavior of others in their chain (the social condition), relative to those who cannot (the asocial control). We stress that individuals in the social condition receive veridical information about others' data-gathering decisions (Fig. 2b), whereas classical conformity experiments involving misbelief (e.g. Asch, 1951, 1955) employ deceptive confederates who deliberately misinform participants. Our experiment is thus aimed at explaining misbelief in the absence of deliberate deception. Participants in both conditions pay a cost (from an endowment) each time they wish to gain some information by drawing a bead, but they win a bonus if they correctly guess the color of the majority of beads in the urn. This combination of incentives allows us to infer how much participants value the information, and to classify participants' decisions as optimal or sub-optimal (Van der Leer et al., 2015).

To avoid a potential confound, participants in the social condition were only told how many beads had been drawn by previous participants (cf. Fig. 2b). They were not told what color those beads were, nor what each participant thought the color of the majority of the beads in the urn was. They were also not told how profitable the decisions of the previous participants were. Otherwise, participants in the social condition would gain additional information about the urn, thus simplifying their task. As it is, the social information

we provided is strictly irrelevant to each person's decision, because beads are drawn randomly, and this experiment is thus not merely about probing the social pathway in Fig. 1a in more depth.

As the statistical decisions involved in performing the Beads Task can be rather complex, we propose a series of related hypotheses. Each hypothesis models participants' behavior as a function of inputs of varying complexity. For instance, a parsimonious view might be that participants are sensitive to the number of beads drawn and nothing else, from which we derive H1 below. Participants may also be sensitive to the quality of that information (yielding H3 below). The most complex model — but the only one able to distinguish optimal from suboptimal behavior — is if participants are additionally sensitive to the costs incurred in gathering that evidence and to the likelihood of winning the bonus, as a function of further draws (yielding H2). Finally, we can remain agnostic as to the belief-formation processes themselves, and simply measure how likely it is that participants are to form true or false beliefs about the color of the majority of beads (H4).

H1 Participants will gather less evidence (draw fewer beads) in the social condition than in the asocial condition ('JTC in a relative sense').

H2 Participants will be more likely to jump to conclusions (making objectively premature decisions) in the social condition than in the asocial condition ('JTC in an absolute sense').

H3 Participants in the social condition will make their decisions about the majority color of the urn at a point of lower objective certainty than will participants in the asocial condition.

H4 Participants will be more likely to form inaccurate beliefs about the state of the world in the social condition, relative to the asocial control.

We pre-registered our data-collection procedures, experiment design, analysis plan, and the above four hypotheses (https://osf.io/db3ws).

Methods

Participants and Incentives

We used the Turkprime platform (L. Litman, Robinson, & Abberbock, 2017) to recruit 2000 participants (1072 female, 909 male, the remainder choosing to self-describe their gender, or not to respond to this question; age: $M = 36.6$, $SD = 13.8$) via Amazon's Mechanical Turk service (MTurk). Given that no previous data on the social effect under investigation here exist, rather than undertaking a power analysis, we sought to ensure that we had more participants per condition (here, 1000 per condition) than the largest-yet study using the Beads Task had (558 overall, in Ross et al., 2016).

Participation was limited to those with IP addresses registered in the USA, who had MTurk approval ratings of at least 95%, and who had completed at least 100 tasks on MTurk. Participants were paid a base rate of $0.50. Additionally, they were given an endowment of $0.15, and were awarded a bonus of $0.92 if they guessed the correct urn. They used the $0.15 endowment to pay $0.01 each time they wanted to draw a bead and see its color. They could thus afford to draw up to 15 beads, and retained any unspent endowment. This incentive scheme was designed so that, on average, the optimal number of beads to draw (assuming risk neutrality) would fall in the middle of the range, to avoid floor or ceiling effects (for the relevant calculations, see https://osf.io/wfbuq/).

Before participating, participants provided informed consent. The consent form included information about the incentives, and informed participants that there would be comprehension tests. The study received ethical approval through the Royal Holloway, University of London self-certification process.

Materials

We coded the experiment using the jsPsych JavaScript Library (De Leeuw, 2015). To avoid participants misunderstanding the nature of the task (Balzan, Delfabbro, Galletly, & Woodward, 2012), we provided detailed written instructions and incorporated extensive comprehension checks. Both the instructions and the trial itself were accompanied by animated visual illustrations, so that participants could see the beads being drawn and replaced in the urn, rather than just being told textually that this was the case. The experiment script is available via https://osf.io/9f2ks/, as is a link to a demonstration of the task.

Procedure

Participants were randomly assigned to one of 100 groups of 20 people each. Fifty of the groups were in the social condition, in which each participant was shown the number of beads seen by each of the people in the group who had previously drawn beads. Fifty of the groups were in the asocial condition, where participants saw the number of people in their group who had previously drawn beads, but not how many beads each of these people had drawn Fig. 2b. For the first person in each group the social and asocial groups were identical, because there were no prior persons in the group whose behavior could be displayed.

Participants were given a demonstration of the task, comprising both textual instructions and animations that visually illustrated the procedure.

First, participants were shown two sets of 100 beads each. Each bead was either pink or green, and the two sets' colors were in complementary ratios (i.e., one set had 60 green and 40 pink beads; the other 40 green and 60 pink). Each set of beads was shuffled and dropped into one of two urns, to obscure the beads' colors. The urns were shuffled so that participants did not know which urn had which majority color. Then one urn was removed, while the other remained on the screen. Participants were told that 15 beads would be drawn with replacement from the remaining urn, and that their task was to decide how many of these beads they wished to see before choosing what color the majority of the beads in the urn was. For details of the randomization procedure (including an inconsequential coding error) see https://osf.io/9f2ks/.

After a demonstration of the task, participants' comprehension was tested. For instance, they were asked whether beads were drawn from one urn or both (answer: one), and whether the ratio of beads remained constant over time (answer: true). They could proceed only when they scored 100% on the comprehension test, and otherwise looped through the demonstration and instructions again, as many times as were necessary to score 100%. 72% of participants passed this comprehension test on the first attempt.

Second, the incentives were explained, including the

$0.15 endowment (from which they paid a $0.01 cost to see each bead) and the $0.92 bonus for correctly guessing the urn's majority color at the end of the trial.

Third, the social nature of the experiment was explained to participants in all positions except the first in each chain: either each participant would simply see the number of preceding decision makers in the chain (asocial control), or each participant would additionally see the number of beads that each of the preceding decision makers had chosen to see (social condition).

Finally, their comprehension of the incentive structure and social aspect was tested. As previously, they could proceed only when they scored 100% on the comprehension test. Otherwise the relevant demonstrations were repeated. 80% of participants passed this comprehension test on the first attempt.

When the trial began, participants were shown what position they were in, and (with the exception of participants in the first position) were told which condition they were in. To avoid participants drawing fewer beads just to end the task early (Van der Leer et al., 2015), all participants had to *draw* 15 beads but they only paid a cost if they chose to *see* the color of a drawn bead. When they decided they had seen enough beads, all subsequent draws (if any) were greyed out and did not incur a cost. Each time a participant paid to draw a bead, its color was recorded on screen (to avoid a dependence on memory) and the bead was replaced in the urn (Fig. 2b). After 15 beads had been drawn, participants were asked what they thought the majority color in the urn was. They were told whether they were correct and how much additional money they had earned.

Finally, participants answered some demographic questions, including age, gender, and whether they were a native speaker of English (1906 indicated they were). They were also able to leave comments if they wished.

Results

Pre-registered analyses

Analyses were pre-registered at https://osf.io/db3ws and carried out in R (R Core Team, 2018). The full analysis script is available at https://osf.io/9f2ks. This includes technical details for the analyses described below, such as random effects structures. In general, we report the results of (generalized) linear mixed-effects models built with package `lme4` (Bates, Maechler, & Bolker, 2011), including 95% bootstrapped confidence intervals for fixed-effect coefficients (generated by function `bootMer` from the same package), where possible. We use package `lmerTest` (Kuznetsova, Brockhoff, & Christensen, 2017) to provide p-values and package `glmmTMB` (Brooks et al., 2017) for more complex (e.g., beta or zero-inflated) regression models.

We hypothesized that participants in the social condition would show a *relative* JTC bias (H1), indexed by taking fewer draws to decision (DTD) than those in the asocial control. Overall, the mean DTD was 6.55 (bootstrapped 95% CIs [6.317, 6.782]). DTD in the social condition was significantly lower than the asocial condition (linear mixed-effects regression, $B = -0.69$ [$-1.262, -0.158$], $SE = 0.289$, $t = -2.388$, $p = 0.019$, Fig. 3a). Thus, social information amplified JTC in a relative sense. There was no main effect of position ($B = -0.019$ [$-0.057, 0.019$], $SE = 0.02$, $t = -0.946$, $p = 0.344$) and no interaction between position and condition ($B = 0.005$ [$-0.083, 0.097$], $SE = 0.045$, $t = 0.108$, $p = 0.914$). Thus, the social amplification effect did not detectably increase over time. We discuss possible reasons for this in the more detailed analysis (which includes R scripts) at https://osf.io/9f2ks.

To evaluate *absolute* JTC (H2) we calculated an expected-value matrix representing the probable payout for all possible states, where each state is defined by the following two values.

- Δ: the absolute difference in the number of beads of each color (so if 3 green and 3 pink beads are drawn, $\Delta = 0$, and if 3 green and 5 pink beads or 5 green and 3 pink beads are drawn, $\Delta = 2$).

- t: time — counted as number of draws — since the start of the trial, such that before any beads have been drawn $t = 0$ and after 15 beads have been drawn, $t = 15$.

The expected value in each state depends on the cost paid to see t beads and the probability of winning the bonus as a function of Δ (if we assume participants pick the majority color if $\Delta > 0$ and pick randomly if $\Delta = 0$). Based on these expected values, if a rational decision is just a decision that would maximise the agent's expected payoff, then a corresponding decision matrix can be calculated for each possible state, where the decision is represented as 0 if it is rational in that state to pay to see another bead, and 1 if it is rational to stop paying.

Our full calculations for this decision matrix are available at https://osf.io/wfbuq/, but we provide a brief and informal overview of a single such decision here, for the sake of illustration. The full decision matrix simply generalizes this kind of decision.

Imagine that a participant has drawn 14 beads. If 7 beads are pink and 7 are green, then at that point, there is a 50% chance of guessing correctly, since there is no evidence either way ($\Delta = 0$, in the terminology of our analyses). Thus, stopping at 14 beads in this scenario would yield a 50% chance of winning the 92 cent bonus, yielding an expected value of $0.5 * 92 = 46$ cents. However, if the participant were to draw one more bead, they would be in one of two possible states. They could either end up in a state where they have 7 pink

and 8 green or one with 8 pink and 7 green beads. In both cases, $\Delta = 1$. When $\Delta = 1$, there is a 60% chance of guessing correctly. If they draw one more bead, the expected value of the bonus will be $0.6 * 92 = 55.2$ cents. It would cost them an additional cent if they wish to proceed after the 14th bead, but this cent would yield an increase of 9.2 cents in expected winnings, so it is worth drawing the final bead. However, if a participant did stop after 14 beads, they would be classed as Jumping to Conclusions.

Overall, a high proportion of participants jumped to conclusions (mean = 0.67, bootstrapped 95% CIs [0.652, 0.693]). If participants were aiming at the optimal point and if they erred randomly (i.e., were equally likely to undershoot as to overshoot), we might expect fewer than 50% of them to undershoot (given that some participants might not err at all). Since the bootstrapped CIs don't contain the value 0.5, the data reflect an overall tendency to jump to conclusions, consistent with findings in the literature (Furl & Averbeck, 2011; Van der Leer et al., 2015).

We hypothesized that social information would amplify an absolute JTC bias, and that participants would thus be more likely to jump to conclusions in the social condition than in the asocial control (H2). As predicted, participants were significantly more likely to jump to conclusions in the social condition relative to the asocial control (binomial mixed-effects regression $B = 0.299$ [0.049, 0.562], $SE = 0.129$, $z = 2.315$, $p = 0.021$, Fig. 3b). Thus, the social context amplified the JTC bias. Again, there was no effect of the participant's position within their chain ($B = 0.012 [-0.005, 0.031]$, $SE = 0.009$, $z = 1.309$, $p = 0.19$), and no interaction between condition and position ($B = -0.021$ [$-0.058, 0.014$], $SE = 0.017$, $z = -1.203$, $p = 0.229$).

For H3, we hypothesized that participants would make their decision at a point of lower objective certainty in the social condition than in the asocial control. 'Objective certainty' here means the posterior probability that an ideal Bayesian agent would have after observing a given realized sequence of beads (as a function of Δ). Overall, participants decided with mean objective certainty of 0.67 (bootstrapped 95% CIs $[0.669, 0.68]$). Participants in the social condition decided at a significantly lower objective certainty (beta mixed-effects regression, $B = -0.08$ [$-0.136, -0.025$], $SE = 0.028$, $z = -2.826$, $p = 0.005$). There was no effect of position ($B = -0.003$ [$-0.007, 0.001$], $SE = 0.002$, $z = -1.336$, $p = 0.181$) and no significant interaction between condition and position ($B = 0.002$ [$-0.007, 0.011$], $SE = 0.004$, $z = 0.452$, $p = 0.652$).

Finally, we hypothesized that participants in the social condition would be less accurate (H4). Accuracy here is a binary variable with value 1 if the participant guessed the correct urn, and 0 otherwise. Overall, mean accuracy was 0.75 (bootstrapped 95% CIs $[0.73, 0.768]$). Participants in the social condition were significantly less accurate (bino-

mial mixed-effects regression, $B = -0.267$ [$-0.47, -0.059$], $SE = 0.104$, $z = -2.579$, $p = 0.01$). Again, there was no effect of position ($B = -0.014$ [$-0.031, 0.005$], $SE = 0.009$, $z = -1.522$, $p = 0.128$) and no interaction between condition and position ($B = 0.029$ [$-0.006, 0.065$], $SE = 0.018$, $z = 1.623$, $p = 0.105$).

We thus found a significant effect of condition for all four pre-registered hypotheses, but no effect of position and no interaction between condition and position. A social context has thus amplified the Jumping to Conclusions bias, increasing the likelihood that people would make costly decisions on the basis of insufficient evidence.

Exploratory analyses

On average participants earned \$1.27 (bootstrapped 95% CIs $[1.258, 1.291]$)[1]. Participants in the social condition earned significantly less (linear mixed-effects regression, $B = -0.039$ [$-0.072, -0.006$], $SE = 0.017$, $t = -2.299$, $p = 0.022$). Thus, a social context has amplified the JTC bias enough to cause costly misbeliefs.

The range of DTDs — the socially transmitted information — is distributed as in Fig. 4a, which shows a number of striking features. Most obviously, the most common responses were on the extremes: participants frequently drew 0 beads or the maximum 15. This behavior is consistent with reliance on some form of heuristic, rather than calculating (or even roughly estimating) probabilities, costs or benefits. Based on participants' comments (left optionally at the end of the study), we describe two potential such heuristics.

The first plausible heuristic is to gain the maximum amount of information, exemplified by comments such as 'I definitely felt it was worth .15 to get the maximum amount of information,' or 'I was going to guess pink after 10 beads but it wasn't worth the 5 cents to risk being wrong.' The second heuristic is the (mistaken) intuition that the information provided by the beads is worthless — in the sense that winning the bonus would remain a 50:50 chance no matter how many beads are drawn — whereas the endowment, if unspent, is a sure thing. This is exemplified by comments such as 'I decided not to see any beads because 60/40 was close to a coinflip so I figured the .15 was "a bird in the hand"', or, 'I figured my odds of guessing would be virtually no better by spending any endowment to see the beads. Even if I did spend the entire endowment my odds of guessing the majority correctly wouldn't have been much different than just guessing.'

---

[1]In order to ensure a more acceptable rate of payment (e.g., \$14.00 per hour in California in 2021), we offer the following two suggestions that we have adopted since running this (our first study on MTurk with bonuses depending in part on random outcomes). First, budget on the assumption that all participants will win the maximal bonus. Second, ensure the base rate (without bonuses) is above the acceptable minimum.
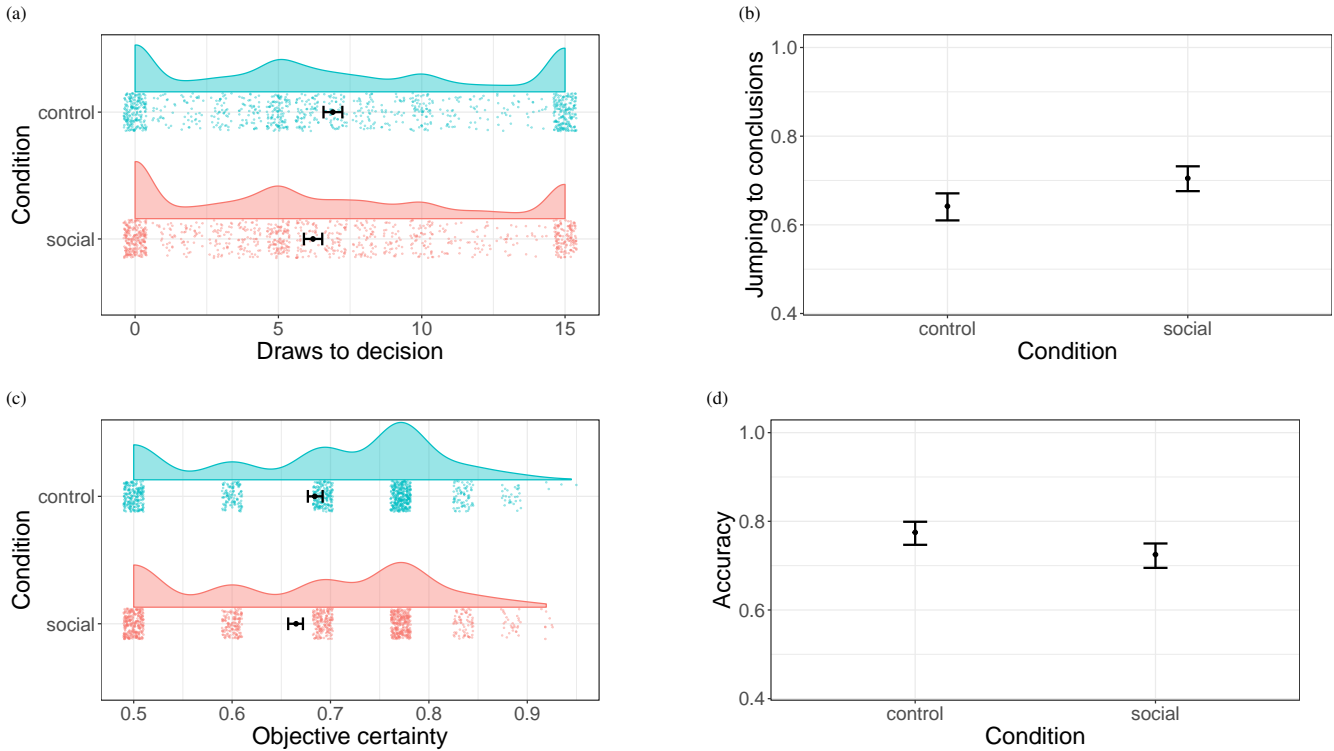
(a)

(b)

(c)

(d)

Figure 3. Plots of the outcome variables for each pre-registered hypothesis. Where the outcome variable is numeric (a, c), the plot shows the individual datapoints (jittered to avoid overlap) and kernel density estimates ('raincloud' plots, Allen, Poggiali, Whitaker, Marshall, & Kievit, 2018). Where the outcome variable is binary (b, d), the plot shows group means along with 95% confidence intervals.
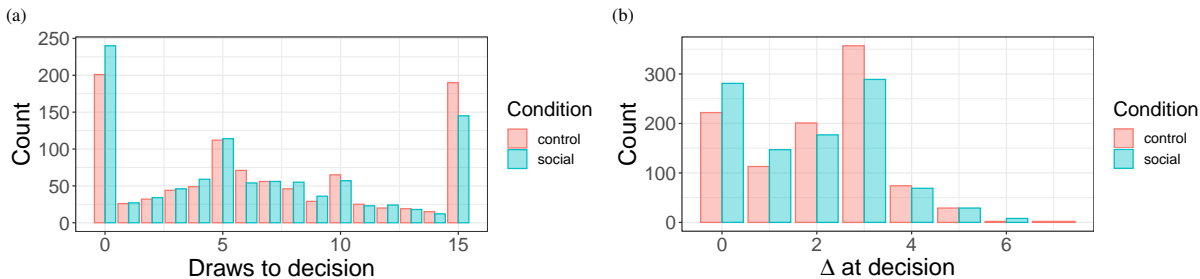
(a)

(b)

Figure 4. Histograms plotting (a) the distribution of draws to decision (DTD) and (b) Δ at decision (Δ_{AD}), colored by condition.

Another striking feature of Fig. 4a is that the counts for most DTDs are very similar between conditions, except for the extremes of the distribution. Thus, the effect of condition seems to be driven largely by participants in the social condition being more likely to decide without seeing any beads, and less likely to draw the maximum.

For ease of modelling, we first reframe the issue in terms of the quantity of diagnostic evidence seen, indexed by Δ at decision (or Δ_{AD}), rather than number of beads seen. If two participants saw 5 beads before picking an urn, but participant A saw 1 green and 4 pink and participant B saw 3 green and 2 pink, then the Δ_{AD} for participant A is 3 and for participant B is 1. In this case, participant A has a better chance of guessing the right urn even though they had the same DTD.

Two of the outcome variables already considered (JTC and objective confidence) are a function of Δ. In a sense, then, Δ is psychologically simpler than those outcome variables, which require additional calculation. If participants are relying on simpler heuristics, as suggested above, then simpler input variables such as this may be informative.

The variable Δ_{AD} is distributed as shown in Fig. 4b. We model this count variable with a Poisson regression. Additionally, we include a zero-inflation component to model the spike where $\Delta_{AD} = 0$. Table 1 shows two versions of the zero-inflation component: one with only an intercept (Model 1) and one with a term for condition (Model 2). Model 1 tests whether there are significantly more zeroes overall in this distribution than expected under a Poisson model, and

Model 2 tests whether there are significantly more zeroes in the social condition in particular. When the zero-inflation component excludes condition (Model 1), the Poisson component has a significant negative effect of condition: Δ_{AD} was lower in the social condition, so participants in that condition saw significantly less diagnostic evidence before deciding. When the zero-inflation component includes condition (Model 2), there is a significant positive effect of condition, but condition becomes non-significant in the Poisson component. Thus, participants in the social condition were significantly more likely to decide without any diagnostic evidence (where $\Delta_{AD} = 0$), and once this factor is accounted for, the two conditions are otherwise not significantly different. The Akaike Information Criterion (AIC) for Model 1 is 6816.84 and for Model 2, AIC = 6812.69. With a difference in AIC of approximately 4, there is positive evidence for Model 2 relative to Model 1. Thus, we conclude that social information increased the likelihood that participants decided without any diagnostic evidence.

Finally, we note that social transmission effects tend to vary widely across chains. This reflects the kind of path-dependence associated with frequency-dependent social learning strategies like conformity (Boyd & Richerson, 1985), though the domain of interest here concerns how to learn, rather than the behavior itself. For illustrative purposes, we plot the cumulative rate of accuracy in each chain in Fig. 5.

These by-group patterns show that no chain in the asocial control was anything other than overwhelmingly accurate, whereas this tendency was reduced in the social condition. We will highlight several features of the specific paths in Fig. 5 that speak to this claim. In the asocial control, 3 chains maintained 100% accuracy past position 10, whereas only 1 chain in the social condition did. Fourteen chains in the asocial control ended up with a cumulative accuracy $\geq$ 0.8 compared to 5 in the social condition. Only 2 chains in the asocial condition ended with cumulative accuracy $\leq$ 0.6 whereas 10 chains in the social condition did. Indeed, only in the social condition was any chain more wrong than right by the end (with cumulative accuracy $<$ 0.5 in position 20, in the chain highlighted in black). Such a chain represents a form of "collective delusion", where most people in the chain make costly, incorrect decisions.

Discussion

We used a social-transmission experiment to study whether a social context would amplify individual data-gathering biases. In a task that involved drawing colored beads from an urn to decide the color of the majority of beads in the urn, participants in a social condition were shown how much evidence had been gathered by each previous participant in their chain. Compared to participants in an asocial control condition, social condition participants drew fewer
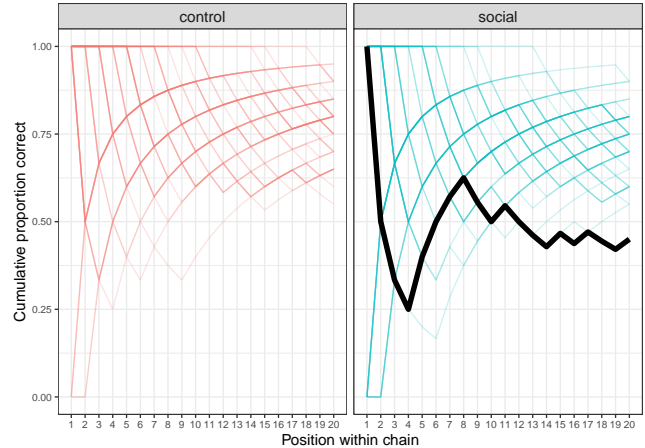

Figure 5. The cumulative proportion of participants guessing correctly in each chain. For the sake of illustration, one chain in the social condition is highlighted in black; the remaining chains are colored by condition, and are lightened so that the overlap of more than one chain can be indicated with deeper reds or blues. In the highlighted chain, the participant in position 1 guessed correctly (so the cumulative proportion correct at position 1 is 1). The participant in position 2 in the same chain guessed incorrectly, so the cumulative value in position 2 is $\frac{1+0}{2}$, or 0.5. The participant in position 3 likewise guessed incorrectly, so the cumulative value at that point is $\frac{1+0+0}{3}$. By position 20, this illustrative chain was overall more incorrect than correct.

beads, were more likely to jump to conclusions and to arrive at false beliefs, and earned less money.

Importantly, it was not the false beliefs themselves that were transmitted socially, as social condition participants were not given information about the decisions of their forebears regarding the state of the world (the majority color of beads in the urn), nor were they privy to the nature of the evidence previous participants had observed (the sequence of beads). Rather, what was transmitted was information about the belief-formation processes of others. What we have shown, with a carefully incentivized paradigm, is that people can socially acquire a disregard for evidence *per se*. Thus, not only is there a causal pathway from social information to individual cognitive bias (Fig. 1b), but this specific bias — the Jumping to Conclusions bias — directly impacts individual learning by causing people to gather less evidence about the state of the world.

We have argued that individual factors alone cannot explain why misbeliefs are socially widespread, while social factors cannot explain why misbeliefs are stable, given that even a trickle of information about the world — gathered via individual learning — can destabilize social traditions of misbelief (Goeree et al., 2007). In this light, our result represents an advance towards understanding collective delusions, because we have shown not only how social and individual factors can interact, but also how social information can

Table 1

Parameters for mixed-effects zero-inflated Poisson regressions, modeling the effect of condition on Δ at decision (Δ_{AD}) with a random intercept for chain.

Component	Parameter	Model 1				Model 2			
		B	SE	z	p	B	SE	z	p
Poisson	Intercept	0.889	0.024	36.62	<0.001[***]	0.875	0.025	34.58	<0.001[***]
	Condition	-0.074	0.035	-2.11	0.035[*]	-0.043	0.037	-1.16	0.246
Zero-infl.	Intercept	-1.578	0.081	-19.58	<0.001[***]	-1.78	0.124	-14.31	<0.001[***]
	Condition	-	-	-	-	0.398	0.162	2.45	0.0142[*]

make participants significantly more likely to refrain from gathering evidence about the world.

We have thus identified a novel mechanism that helps explain costly social misbelief. Moreover, our work also represents a methodological advance. Whereas classical conformity experiments (e.g., Asch, 1955) typically require "exogenous" sources of misbelief, such as deceptive confederates who deliberately misinform participants, the mechanism we identify is "endogenous" in that misbeliefs arise naturally without artificial influence, due to participants' cognitive biases. Further, as participants took part online (as compared with Asch's study, where participants gave responses in full view of confederate 'participants'), this result is unlikely to be explained by social pressures to comply (cf. Levy, 1960), though it may still represent a general bias towards social learning (Perreault et al., 2012).

Another strand of psychological research into conformity (the autokinetic effect, Sherif, 1936) involves cases where participants are artificially deprived of the ability to weigh empirical evidence directly and objectively. In contrast, our paradigm makes objective evidence available to all participants. Our incentives are designed to make it always optimal to draw at least a few beads, yet many participants chose not to gather any evidence at all, and social information boosted this tendency. The results thus speak to the puzzle of why groups of humans can believe *demonstrably* false things, by showing that a social context can amplify a tendency to undersample potentially disconfirming evidence.

Apart from these theoretical and methodological contributions, our results have three main implications for the future study of misbeliefs in the real world, where the objective value of information is not always clear, and where other people's information gathering habits are not so explicitly available.

First, these results can be construed as showing that a social context can shift the balance between reliance on external evidence and reliance on introspection, gut feelings, or subjective intuitions. Multiple participants expressed the false intuition that the beads were informationally worthless, believing the 60:40 ratio to be close enough to 50:50 that it was roughly at chance. Thus, they privileged their intuition over the readily available evidence. As mentioned earlier, an intuitive cognitive style — where people assent to the first,

intuitive explanation that occurs to them, rather than probing further — predicts belief in fake news (Bronstein et al., 2019; Pennycook & Rand, 2020) and disbelief in science (Gervais, 2015; Lobato & Zimmerman, 2019). Our results suggest the intriguing possibility that the effect of intuitive cognitive style can be amplified by particular social contexts. It would be interesting to see how far this possibility extends. For instance, people who endorse evolution are more likely to justify their beliefs with reference to evidence, whereas people who deny evolution are more likely to justify their beliefs based on what they feel to be 'true in their heart' (Metz et al., 2018). Naturally, this example is confounded with religious factors, and yet that does not preclude the possibility that part of the explanation may be a socially learned disregard for empirical evidence, as religious belief is itself correlated with intuitive cognitive style (Pennycook, Ross, Koehler, & Fugelsang, 2016).

Second, our results offer a concrete way to understand social misbeliefs in light of what is known about clinical and sub-clinical delusions. Measures of delusion-proneness have been shown to predict belief in fake news (Bronstein, Pennycook, Bear, Cannon, & Rand, 2018) or the paranormal (Ståhl & van Prooijen, 2018). Further, delusion-prone people are typically less able to integrate incoming evidence to update their false beliefs, showing a Bias Against Disconfirmatory Evidence (BADE, McLean et al., 2017), which make speak to the inflexibility of collective delusions. What our approach adds, by using the Beads Task — which is adapted from the delusions literature (Huq et al., 1988) — and studying the Jumping to Conclusions bias — which is thought to play a role in the formation of delusions (Garety & Freeman, 2013) — is a specific mechanism for bridging the social and individual pathways in Fig. 1, as we experimentally controlled the kind of social information shared in the social condition.

Third, we offer a concrete prediction about collective delusions, on the basis of our results. We have argued that the devaluation of evidence is a mechanism whereby social information can affect individual learning in such a way as to stabilize misbeliefs. If that is the case, then when two social groups take opposing sides in any disagreement over empirically verifiable matters, we hypothesize that one side will only be able to maintain their position long-term if their group identity includes epistemic commitments, such

as scorn of scientists as sources of empirical evidence, scorn of journalists as sources of accurate reporting, and warnings against doubting received wisdom (for instance, see James 1:6 '... the one who doubts is like a wave of the sea that is driven and tossed by the wind'). Religion is an obvious case where social identity carries epistemic injunctions, but we are arguing that this could be a much more general phenomenon. Naturally, the supposition that epistemic standards can be part of, say, political ideology is hardly novel (cf. Orwell's *Nineteen Eighty-Four*), but we are proposing a specific mechanism for why this may be a general phenomenon.

Our paradigm also includes several limitations. First, it involves mathematical reasoning, and if people are mathematically uncertain, then they may be more disposed to follow social information (Cross, Brown, Morgan, & Laland, 2017). However, this is not incompatible with our claims: if people are uncertain about whether the beads are worth seeing, but become more confident in drawing zero beads when they see that others have done so, then this is merely one way in which social information has the effect we are claiming. Further, this is one reason we posed four related hypotheses. The simplest of these merely involves counting the number of beads observed, which is simple enough for primary school mathematics. In any case, we modeled optimal behavior in a way that incorporates probability and expected value, but this was to support the analytic aim of comparing human performance with an objective, optimal benchmark, rather than to claim that many (or any) people actually carry out such calculations. This research was originally motivated by gaps in current explanations of collective delusions such as climate change denial, and lay people are probably at least as uncertain about the physics of climate change as they are about the mathematics of probability in our task. Thus, uncertainty is inherently a part of people's decision making in the phenomenon to be explained, rather than an unfortunate accident in our empirical study of it.

Second, although the social information about bead draws was strictly irrelevant, people may have assumed that this information was relevant — indeed, the assumption of relevance is foundational in human communication (Sperber & Wilson, 1995) — and perhaps we should thus have probed participants about their assumptions here. Indeed, this assumption has been shown to explain apparently irrational behavior (Sperber, Cara, & Girotto, 1995). However, appeal to this principle alone cannot explain why social information biased people towards drawing the minimum number of beads (rather than the maximum, which was the other popular strategy). Thus, it does not undermine our account. Relatedly, although our design intentionally makes social information useless, we acknowledge both that it is frequently useful — indeed, it is hard to see how our species would have evolved to be such keen social learners, otherwise — and that there are adaptive signals of its usefulness. For instance, it is

often beneficial to copy prestigious individuals (Henrich & Gil-White, 2001), and displays of pride can serve as signals of such prestige (Martens & Tracy, 2013). Nonetheless, we wish to highlight that such matters concern deciding *whom* to copy when copying, whereas our particular focus here has been on how social information can discourage people from engaging in individual learning, to detrimental effect.

Third, although we included an extensive system of comprehension checks, we did not check every aspect of participants' understanding. It may have been worth checking that participants had properly read the paragraph in the instructions from the social condition that said they would only see how many beads had been drawn by previous participants (in case they were confused and thought that the social information displayed actual black beads seen) or that the bead colors in the experimental task would not be the same as in the demonstration.

Finally, it is somewhat surprising that we observed no increase in the rate of Jumping to Conclusions as the transmission chain grew in length. We explore the data behind this in some detail in the expanded analysis (https://osf.io/9f2ks), but our best current explanation is based on the observation that people vary in their base-level sensitivity to social information (Mesoudi, Chang, Dall, & Thornton, 2016; Molleman, Kurvers, & van den Bos, 2019). In our data, only a subset of participants drew zero beads, and only a subset of *those* participants can have done so because of the experimental manipulation. Thus, as any effect of time would be observable only for a subset of a subset of participants, and as this subset of a subset of the population also vary in their base-level sensitivity to social information (Molleman et al., 2019), then it is unsurprising if it is hard to detect a signal of how much social information is needed for this sub-subset of people to be swayed.

We conclude with two questions for future research. First, what, if anything, is cognitively distinctive about the participants who were susceptible to the faulty intuition that the evidence was worthless, or who tended to imitate this heuristic in others? So far, we have mentioned JTC, a reliance on intuition rather than evidence, and heterogeneous social-information use, but another possibility might be that such people are particularly intolerant of ambiguity. A 60:40 ratio of bead colors involves more ambiguity than, say, a 75:25 ratio. Higher tolerance of ambiguity is associated with curiosity about new information (J. A. Litman, 2010), whereas intolerance of ambiguity is associated with conservatism (Jost, Glaser, Kruglanski, & Sulloway, 2003), itself a predictor of some forms of science denial (Rutjens, Heine, et al., 2018). A second question is, what patterns in social information are responsible for our main effect? Our paradigm involves presenting participants in the social condition with naturalistic inputs, in the sense that these were veridical responses of other participants at earlier time points. This was motivated

by the desire to avoid deception (cf. Asch's studies). However, this naturalistic data may have obscured any effect of particular kinds of social information (for instance, the number of people who previously drew zero beads) and a different approach will be needed to probe such effects.

References

Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., & Kievit, R. (2018, August). Raincloud plots: a multi-platform tool for robust data visualization. *PeerJ Preprints*, *6*, e27137v1. Retrieved from https://doi.org/10.7287/peerj.preprints.27137v1 doi: 10.7287/peerj.preprints.27137v1

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, *87*(5), 847–862.

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership and men* (pp. 177–190). Pittsburg, PA: Carnegie Press.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, *193*(5), 31–35.

Balzan, R., Delfabbro, P. H., Galletly, C. A., & Woodward, T. S. (2012). Over-adjustment or miscomprehension? A re-examination of the jumping to conclusions bias. *Australian & New Zealand Journal of Psychiatry*, *46*(6), 532–540.

Barone, L. M., Petto, A. J., & Campbell, B. C. (2014). Predictors of evolution acceptance in a museum population. *Evolution: Education and Outreach*, *7*(23), 1–22.

Bates, D., Maechler, M., & Bolker, B. (2011). lme4: Linear mixed-effects models using s4 classes [Computer software manual]. Retrieved from http://CRAN.R-project.org/package=lme4 (R package version 0.999375-42)

Blancke, S., Van Breusegem, F., De Jaeger, G., Braeckman, J., & Van Montagu, M. (2015). Fatal attraction: the intuitive appeal of GMO opposition. *Trends in Plant Science*, *20*(7), 414–418.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.

Bronstein, M. V., Pennycook, G., Bear, A., Cannon, T., & Rand, D. (2018). Do similar cognitive mechanisms encourage delusion-like ideation and belief in fake news? *Schizophrenia Bulletin*, *44*(Suppl 1), S147.

Bronstein, M. V., Pennycook, G., Bear, A., Rand, D., & Cannon, T. D. (2019). Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking. *Journal of Applied Research in Memory and Cognition*, *8*, 108–117.

Brooks, M. E., Kristensen, K., van Benthem, K. J., Magnusson, A., Berg, C. W., Nielsen, A., . . . Bolker, B. M. (2017). glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *The R Journal*, *9*(2), 378–400. Retrieved from https://journal.r-project.org/archive/2017/RJ-2017-066/index.html

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, *55*, 591–621.

Claidière, N., & Sperber, D. (2007). The role of attraction in cultural evolution. *Journal of Cognition and Culture*, *7*(1-2), 89–111.

Cross, C. P., Brown, G. R., Morgan, T. J., & Laland, K. N. (2017). Sex differences in confidence influence patterns of conformity. *British Journal of Psychology*, *108*(4), 655–667.

De Leeuw, J. R. (2015). jsPsych: A javascript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, *47*(1), 1–12.

Druckman, J. N., & McGrath, M. C. (2019). The evidence for motivated reasoning in climate change preference formation. *Nature Climate Change*, *9*(2), 111-119.

Drummond, C., & Fischhoff, B. (2017). Individuals with greater science literacy and education have more polarized beliefs on controversial science topics. *Proceedings of the National Academy of Sciences*, *114*(36), 9587–9592.

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., & Lubell, M. (2008). Conformists and mavericks: the empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior*, *29*(1), 56–64.

Efferson, C., McKay, R., & Fehr, E. (2020). The evolution of distorted beliefs versus mistaken choices under asymmetric error costs. *Evolutionary Human Sciences*, *2*, E27.

Furl, N., & Averbeck, B. B. (2011). Parietal cortex and insula relate to evidence seeking relevant to reward-related decisions. *Journal of Neuroscience*, *31*(48), 17572–17582.

Garety, P., & Freeman, D. (2013). The past and future of delusions research: from the inexplicable to the treatable. *The British Journal of Psychiatry*, *203*(5), 327–333.

Gervais, W. M. (2015). Override the controversy: Analytic thinking predicts endorsement of evolution. *Cognition*, *142*, 312–321.

Goeree, J. K., Palfrey, T. R., Rogers, B. W., & McKelvey, R. D. (2007). Self-correcting information cascades. *The Review of Economic Studies*, *74*(3), 733–762.

Griffiths, T. L., Kalish, M. L., & Lewandowsky, S. (2008). Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philosophical*

Transactions of the Royal Society B: Biological Sciences, *363*(1509), 3503–3514.

Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, *22*(3), 165–196.

Huq, S., Garety, P. A., & Hemsley, D. R. (1988). Probabilistic judgements in deluded and non-deluded subjects. *The Quarterly Journal of Experimental Psychology*, *40*(4), 801–812.

Jacobs, R. C., & Campbell, D. T. (1961). The perpetuation of an arbitrary tradition through several generations of a laboratory microculture. *The Journal of Abnormal and Social Psychology*, *62*(3), 649.

Jolley, D., & Douglas, K. M. (2014a). The effects of anti-vaccine conspiracy theories on vaccination intentions. *PloS One*, *9*(2), e89177.

Jolley, D., & Douglas, K. M. (2014b). The social consequences of conspiracism: Exposure to conspiracy theories decreases intentions to engage in politics and to reduce one's carbon footprint. *British Journal of Psychology*, *105*(1), 35–56.

Jost, J. T., Glaser, J., Kruglanski, A. W., & Sulloway, F. J. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin*, *129*(3), 339–375.

Jylhä, K. M., & Hellmer, K. (2020). Right-wing populism and climate change denial: The roles of exclusionary and anti-egalitarian preferences, conservative ideology, and antiestablishment attitudes. *Analyses of Social Issues and Public Policy*.

Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection: An experimental study. *Judgement and Decision Making*, *8*, 407–424.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13), 1–26. doi: 10.18637/jss.v082.i13

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, *32*(1), 4–14.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., . . . others (2018). The science of fake news. *Science*, *359*(6380), 1094–1096.

Levy, L. (1960). Studies in conformity behavior: a methodological note. *The Journal of Psychology*, *50*(1), 39–41.

Lewandowsky, S., Gignac, G. E., & Oberauer, K. (2013). The role of conspiracist ideation and worldviews in predicting rejection of science. *PloS One*, *8*(10), e75637.

Lewandowsky, S., Oberauer, K., & Gignac, G. E. (2013). NASA faked the moon landing—Therefore, (climate) science is a hoax: An anatomy of the motivated rejection of science. *Psychological Science*, *24*(5), 622–633.

Litman, J. A. (2010). Relationships between measures of i- and d-type curiosity, ambiguity tolerance, and need for closure: An initial test of the wanting-liking model of information-seeking. *Personality and Individual Differences*, *48*(4), 397–402.

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.

Lobato, E. J., & Zimmerman, C. (2019). Examining how people reason about controversial scientific topics. *Thinking & Reasoning*, *25*(2), 231–255.

Martens, J. P., & Tracy, J. L. (2013). The emotional origins of a social learning bias: Does the pride expression cue copying? *Social Psychological and Personality Science*, *4*(4), 492–499.

McKay, R. T., & Dennett, D. C. (2009). The evolution of misbelief. *Behavioral and Brain Sciences*, *32*(6), 493–561.

McLean, B. F., Mattiske, J. K., & Balzan, R. P. (2017). Association of the jumping to conclusions and evidence integration biases with delusions in psychosis: a detailed meta-analysis. *Schizophrenia Bulletin*, *43*(2), 344–354.

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*(3), 215–225.

Metz, S. E., Weisberg, D. S., & Weisberg, M. (2018). Nonscientific criteria for belief sustain counter-scientific beliefs. *Cognitive Science*, *42*(5), 1–27.

Miton, H., & Mercier, H. (2015). Cognitive obstacles to pro-vaccination beliefs. *Trends in Cognitive Sciences*, *19*(11), 633–636.

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.

Morgan, T. J. H., Rendell, L., Ehn, M., Hoppitt, W., & Laland, K. (2011). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*, *279*(1729), 653–662.

Murphy, J., Vallières, F., Bentall, R., Shevlin, M., McBride, O., Hartman, T., . . . Hyland, P. (2021). Psychological characteristics associated with COVID-19 vaccine hesitancy and resistance in Ireland and the United Kingdom. *Nature Communications*, *12*, 29.

Pennycook, G., & Rand, D. G. (2020). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Person-*

ality, *88*(2), 185–200.

Pennycook, G., Ross, R. M., Koehler, D. J., & Fugelsang, J. A. (2016). Atheists and agnostics are more reflective than religious believers: Four empirical studies and a meta-analysis. *PloS One*, *11*(4), e0153039.

Perreault, C., Moya, C., & Boyd, R. (2012). A bayesian approach to the evolution of social learning. *Evolution and Human Behavior*, *33*(5), 449–459.

R Core Team. (2018). R: A language and environment for statistical computing [Computer software manual]. Retrieved from http://www.R-project.org/

Ross, R. M., McKay, R., Coltheart, M., & Langdon, R. (2015). Jumping to conclusions about the beads task? A meta-analysis of delusional ideation and data-gathering. *Schizophrenia Bulletin*, *41*(5), 1183–1191.

Ross, R. M., Pennycook, G., McKay, R., Gervais, W. M., Langdon, R., & Coltheart, M. (2016). Analytic cognitive style, not delusional ideation, predicts data gathering in a large beads task study. *Cognitive Neuropsychiatry*, *21*(4), 300–314.

Rutjens, B. T., Heine, S. J., Sutton, R. M., & van Harreveld, F. (2018). Attitudes towards science. *Advances in Experimental Social Psychology*, *57*, 125–165.

Rutjens, B. T., Sutton, R. M., & van der Lee, R. (2018). Not all skepticism is equal: Exploring the ideological antecedents of science acceptance and rejection. *Personality and Social Psychology Bulletin*, *44*(3), 384–405.

Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology (Columbia University)*, *27*(187).

Sherif, M. (1936). *The psychology of social norms*. New York: Harper.

Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, *57*(1), 31–95.

Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Malden, MA: Blackwell Publishing.

Ståhl, T., & van Prooijen, J.-W. (2018). Epistemic rationality: Skepticism toward unfounded beliefs requires sufficient cognitive ability and motivation to be rational. *Personality and Individual Differences*, *122*, 155–163.

Van der Leer, L., Hartig, B., Goldmanis, M., & McKay, R. (2015). Delusion proneness and 'jumping to conclusions': Relative and absolute effects. *Psychological Medicine*, *45*(6), 1253–1262.

Wagner-Egger, P., Delouvée, S., Gauvrit, N., & Dieguez, S. (2018). Creationism and conspiracism share a common teleological bias. *Current Biology*, *28*(16), R867–R868.