Title
Beliefs are most swayed by social prevalence under uncertainty

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors
Orticio, Evan
Marti, Louis
Kidd, Celeste

Publication Date
2021

Beliefs are most swayed by social prevalence under uncertainty

Evan Orticio (orticio@berkeley.edu)

Louis Martí (lmarti13@gmail.com)

Celeste Kidd (celestekidd@berkeley.edu)

Department of Psychology, University of California, Berkeley

Abstract

We rely heavily on information from the social world to inform our real-world beliefs. How is this social information used, and when is it most influential? We assess the role of one kind of social information, the prevalence of a belief, in belief updating. Using real-world pseudoscientific and conspiratorial claims, we show that increases in people's estimates of the prevalence of a belief led to increases in their endorsement of said belief. Prevalence information elicited the strongest belief change when people were most uncertain of their initial belief, suggesting that people weigh social information rationally according to the strength of their initial evidence. We discuss the implications of our results in the context of the present misinformation epidemic.

Keywords: belief change; belief prevalence; informational social influence; misinformation

Introduction

Beliefs are essential in guiding our interactions with the world, yet the amount of data required to construct all of our beliefs "from scratch" is intractable given our limited attentional resources. What's more, most truths cannot be determined by direct reference to the physical world in the first place (Festinger, 1954). For most real-world beliefs, relevant evidence is inaccessible to the individual and must be mediated through other agents (Perfors & Navarro, 2019). Reliance on information sampled from the social world is thus central to belief formation. This is distinct from normative social influence, where one conforms their belief to an established social norm for the sake of maintaining positive social standing and group unity, regardless of potential conflicts with private information about the ground truth (Deutsch & Gerard, 1955). Instead, this social sampling, a form of informational social influence, is necessary for acquiring information about truths that are complex or otherwise unavailable to the learner.

People's beliefs have demonstrated sensitivity to social consensus information. Experimentally increasing perceptions of the scientific consensus about climate change led to increased belief in climate change as human-caused as well as increased support for public action (van der Linden, Leiserowitz, Feinberg & Maibach, 2015; van der Linden, Leiserowitz & Maibach, 2019). Appeals to expert consensus have been effectively used to intervene on beliefs in a variety of other domains, including GMOs (Kerr & Wilson, 2018) and vaccination (van der Linden, Clarke & Maibach, 2015). However, much of this work has relied on information from experts or authority figures, which may carry more weight than information about the beliefs of the general public. It is important to understand the role of public consensus in belief formation, as there are many cases in which expert opinion is ambivalent or distrusted; in the latter case, scientific consensus information can lose effectiveness in promoting belief change (Bolsen & Druckman, 2018). Previous work has discovered causal influences of public consensus information on e.g., stereotypes about race (Stangor, Sechrist & Jost, 2001) and obesity (Puhl, Schwartz & Brownell, 2005), but these findings conflate beliefs with attitudes and may not generalize to all kinds of beliefs. One notable exception is Lewandowsky, Cook, Fay & Gignac (2019), which found that beliefs formed about anthropogenic global warming on the basis of an internet blog were influenced by the extent to which reader comments endorsed or rejected its contents. While this finding is suggestive, it is impossible to parse the unique role of perceived consensus per se from the argumentative content of the reader comments. The precise causal role of one's perception of the general prevalence of a belief in their own belief change remains unclear.

It is not a given that information about the mere prevalence of a belief would influence people's endorsement of that belief, particularly in domains affected by widespread ignorance or polarization. Given the global rise in political polarization, and findings that lay perceptions of polarization tend to be even stronger than reality (Lees & Cikara, 2021), the possibility of bias in social evidence may be particularly salient, constraining its use in belief formation. Perfors, Navarro, and Shafto (2018) report a paradoxical finding that social presentation of stronger evidence did not always lead to stronger conclusions—because it prompted inferences about selection bias. In some cases, prevalence information can even lead to contrary updating. When presented with scientific consensus information about anthropogenic global warming, Americans who strongly supported free markets reduced their acceptance of the consensus and the reality of global warming further (Cook & Lewandowsky, 2016). People continually assess the reliability of new social information in light of their prior knowledge before integrating it to update their beliefs (or not).

Informational social influence as cue integration

The role of social information in belief formation can be understood within a multimodal cue integration framework, which originated in the physical perception literature (e.g., Knill & Richards, 1996; Ernst & Banks, 2002). In perceiving the world, observers combine signals from different senses to arrive at a unified percept. In forming a new belief, a learner similarly integrates private information, like domain-relevant knowledge, with social information, like the prevalence of a belief in a sampled population. Importantly, each cue is flexibly weighted according to its estimated reliability. Strong reliance on social information is often rational, especially under conditions of uncertainty. Within this framework, even the conformity to the majority in Asch's (1951) famous line matching experiments is an optimal choice, regardless of normative pressure. If the participant assumes that the confederates are unbiased and have low error rates, as is expected in simple perceptual tasks, the probability of the participant being correct given the confederates' convergence on a different answer is very low (Toelch & Dolan, 2015). Thus, an optimal Bayesian learner would give high weight to the social information.

The role of social prevalence in the spread of misinformation

As the Asch example illustrates, the optimal integration of social information can still lead to the formation of false beliefs. Can these principles similarly explain viral cases of online misinformation? Much of the existing literature on conspiracy and pseudoscientific beliefs has sought to explain them away by appealing to ancillary influences on belief formation, like emotionality (Vlasceanu, Goebel & Coman, 2020), political extremism (Van Prooijen, Krouwel & Pollet, 2015), inattentiveness (Pennycook & Rand, 2019), and bullshit receptivity (Pennycook & Rand, 2020). Similarly, social influences in belief formation have been described as heuristics which operate in opposition to "systematic processing" (Metzger, Flanagin & Medders, 2010; Kim, 2020). However, these arguments fail to acknowledge that many online platforms have created an environment that makes rational belief formation processes particularly susceptible to misinformation. False news has been found to spread "faster, farther, deeper, and more broadly" than true news online (Vosoughi et al., 2018), due in part to the formation of echo chambers. Given the high virality potential of online misinformation, it is likely that people encounter high social engagement metrics, like the number of likes and shares, accompanying false claims. These engagement metrics may serve as a cue to a belief's prevalence in the population, which may in turn impact the user's own impression of the belief's legitimacy.

The present study investigates this potential relationship between the inferred prevalence of a belief and its believability. In an online experiment, we modulate people's perceptions of the prevalence of various empirically unsupported beliefs to test whether this affects later endorsement of the beliefs. We predict that people will judge these beliefs as more likely when they believe the beliefs are more prevalent. Further, the weight of this prevalence cue should be rationally modulated according to the reliability of existing evidence. We use a set of items spanning a wide range of domains in order to capture beliefs of varying strength, enabling us to assess the relationship between initial certainty and belief change. We predict that prevalence information will have a greater influence on a belief when a participant is less certain of their initial belief.

Method

This experiment was preregistered prior to data collection at https://aspredicted.org/blind.php?x=qm6mt6.

Participants

Four hundred and three Americans fluent in English were recruited through Prolific (www.prolific.co) to participate in an 18-minute online experiment. Participants were compensated at a rate of $10/hour, with an opportunity to receive a bonus for good performance.

Procedure

Participants were shown a series of statements relating to real-world conspiratorial or pseudoscientific beliefs (e.g., "The earth is flat", "Wearing masks is harmful to the health of the mask wearer", "The US government planned the 9/11 attack on the World Trade Center", "Hillary Clinton was involved in a child sex-trafficking ring centered around a pizza restaurant in Washington DC"). On each trial, participants were asked to provide (1) a likelihood estimate of the belief ("How likely do you think it is that the statement is true?", on a 0-100 slider scale), and (2) a prevalence estimate ("How many people out of 100 do you think believe the statement is true?"). To ensure participants understood the likelihood scale, the experiment began with three practice trials whose answers were intended to be obvious. If the participant gave an unreasonable response to a practice trial, e.g., rated "Plants need water to grow" as less than 90% likely, they received an error message and were asked to change their response until they responded appropriately. We included 4 more trials of the same type in the main task to be used as attention checks (without feedback), for a total of 37 trials.

In a second block, these same trials were repeated in a shuffled order. However, before evaluating the likelihood and prevalence of each statement, participants were shown a sample of data indicating how many of 10 survey respondents believed the shown sentence. On half of the trials, randomly assigned, the sample for the 10 people approximately matched the participant's own estimate for the prevalence of the belief in block 1 (Control condition). On the other half of trials, the number of people in the sample endorsing the pseudoscientific or conspiratorial belief was 40% higher than the participant's initial prevalence estimate (Higher Prevalence condition). For example, if a participant estimated in block 1 that 19 of 100 people believe the given statement is true, then in block 2 they would be shown that 2 of the 10 survey respondents believe the statement (Control condition) or 6 of the 10 believe it (Higher Prevalence condition). Note that, for variation, some statements were related to an attested conspiracy or pseudoscientific belief but worded in their

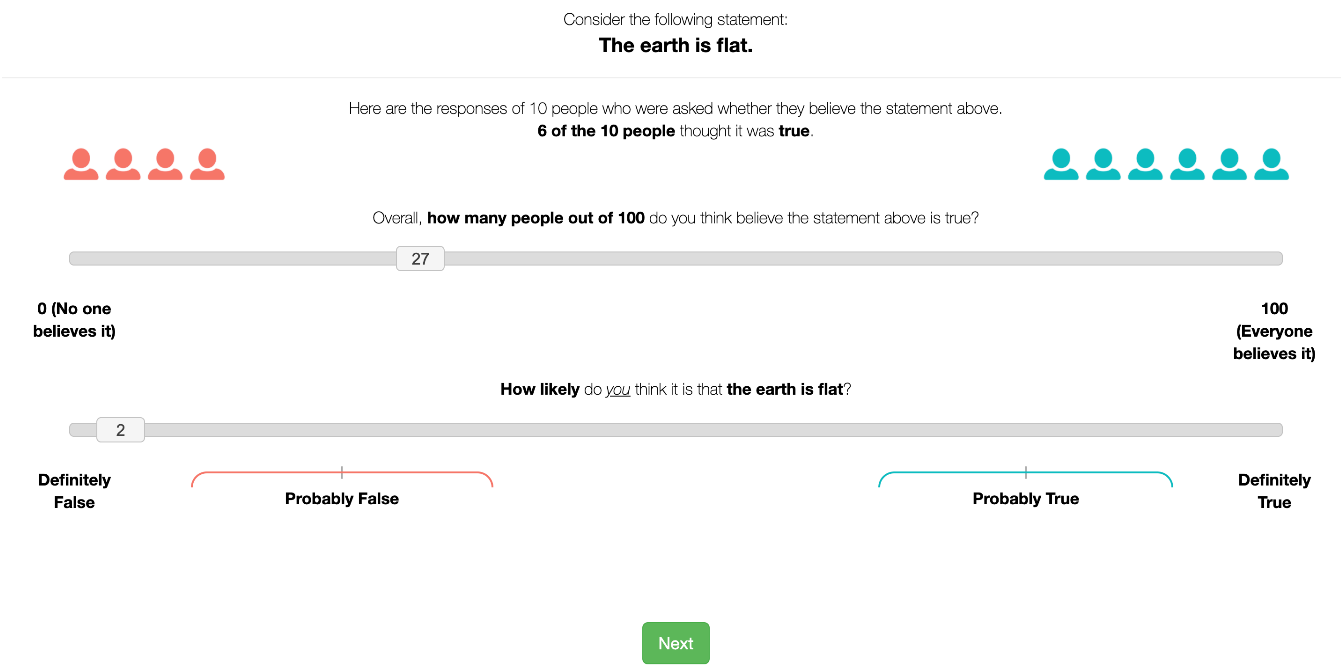

Figure 1: Example trial from block 2.

inverse (true) forms, e.g., "Humans have landed on the moon". In this case, the prevalence of the sample was made 40% lower, i.e., in the direction of the empirically unsupported belief, in the experimental condition. These items are reverse coded in our analyses, so we retain the Higher Prevalence label for simplicity. The prevalence data were also presented visuospatially with icons representing each of the 10 people (see Figure 1). The participants were then asked to give estimates for the prevalence and likelihood of the belief.

At the end of the experiment, participants were debriefed; they were shown a representative prevalence estimate of all of the beliefs based on a large sample of over 900 Americans (Martí, Conover & Kidd, 2021) and reminded that the prevalence of a belief does not necessarily indicate the ground truth.

Results

Exclusions

Following our preregistered exclusion plan, 18 participants were excluded for failing more than 1 of 4 attention checks, and an additional 19 participants were excluded for giving blank or unreasonable responses to one or more of 6 bot-catch questions (e.g., "What's your favorite frozen treat?") intermixed between experimental trials. These subject-level exclusions resulted in a final sample of 366.

In addition to these preregistered exclusions, we also excluded trials in which participants initially endorsed the empirically unsupported belief with a likelihood rating of over 60%, because the 40% exaggeration in the Higher Prevalence condition could not be applied to these. These trials only constituted 13.1% of our data, and their removal did not significantly affect the interpretation of any analyses.

Participants revise their estimates of prevalence in light of prevalence data

If our manipulation worked as intended, participants' estimates of the prevalence of these beliefs should have increased after seeing the Higher Prevalence data, but remained the same after seeing the Control prevalence data. On average, participants' prevalence estimates increased by 21.0% in the Higher Prevalence condition and decreased marginally by 1.7% in the Control condition. A t-test indicates that this difference is statistically significant ($t(6905) = 65.75$, $p < 0.001$, $d = 1.35$), confirming the effectiveness of the manipulation.

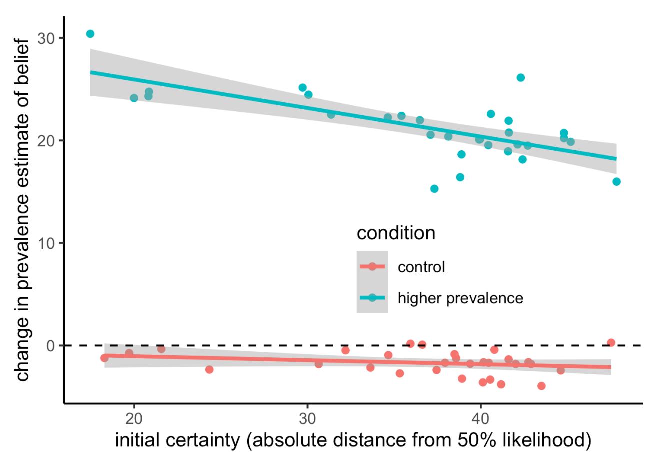

Figure 2: Mean change in participants' estimates of belief prevalence after viewing new prevalence data for each item, by condition. Colored lines represent the fit from linear regression for each condition, with 95% confidence intervals.

Changes in prevalence estimates depend on initial certainty of belief

First, we asked whether the amount of change in prevalence estimates was predicted by participants' initial certainty in their belief. As a metric of certainty, we used participants' initial likelihood estimates for each belief and calculated their distance from 50% on the likelihood scale (where 50% indicates being equally distant from "definitely false" and "definitely true" on the judgment scale). Thus, higher values of certainty indicate the most extreme beliefs. Given that participants tended to rate most of these empirically unsupported beliefs as rather unlikely, certainty ratings were high overall (mean = 37.8/50, median = 46, SD = 16.0).

Figure 2 shows the mean change in prevalence estimates for each belief in the Control vs. Higher Prevalence conditions as a function of initial certainty of belief. We ran a linear mixed-effects model predicting change in prevalence estimates using standardized certainty and condition (Control vs. Higher Prevalence), with random intercepts for item and participant. The model revealed a significant main effect of condition ($\beta = 22.8$, $p < 0.001$) and a significant condition by certainty interaction ($\beta = -2.19$, $p < 0.001$). Within the Higher Prevalence condition, participants updated their estimates of a belief's prevalence more in response to new prevalence data when they were less certain of the belief itself. This model accounted for 50.0% of the variance in the change in prevalence estimates (conditional R2, Nakagawa et al., 2017).

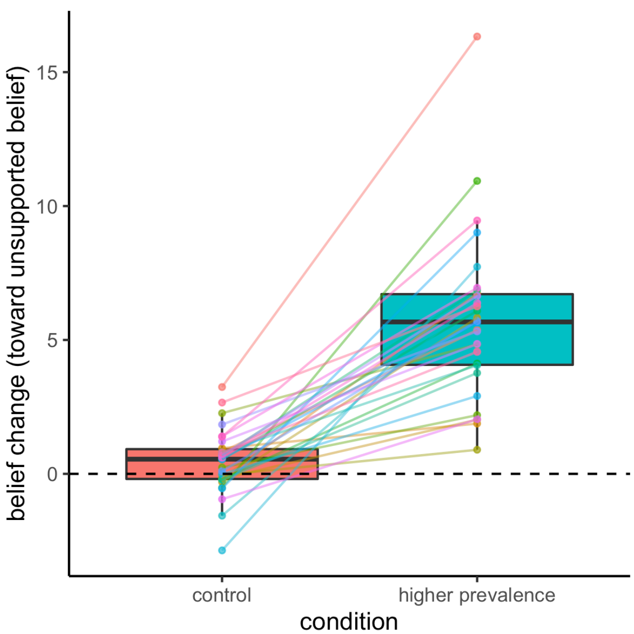

Figure 3: Boxplot of belief change after seeing a sample of data that either matched participants' expectations (Control) or indicated a higher prevalence of belief (Higher Prevalence). Points represent the mean belief change per item, with lines showing the effect of prevalence condition for each item.

Participants revise their beliefs in line with new prevalence information

We hypothesized that showing participants samples of people that more often espoused an unsupported belief than the participant expected would increase their own endorsement of said belief. Figure 3 shows the mean amount of belief change per item according to prevalence condition. A positive value of belief change corresponded to a stronger endorsement of the unsupported belief. As predicted, participants' ratings of the likelihood of these beliefs increased by 5.44% in the Higher Prevalence condition and remained relatively stable (increased by 0.46%) in the Control condition ($t(9047) = 17.82$, $p < 0.001$, $d = 0.365$), indicating that exposure to samples with higher belief prevalence influences the believability of the belief.

Belief change is commensurate with change in prevalence estimate

If social prevalence is treated as an independent source of information that is integrated rationally with prior beliefs, then greater changes in one's estimation of the prevalence of a belief should result in greater changes to the belief itself. Figure 4 shows the relationship between change in prevalence estimate and resulting belief change within the Higher Prevalence condition. We ran a linear mixed-effects model predicting belief change with condition and change in prevalence estimate as fixed effects and random intercepts per participant and item. This model revealed significant main effects of condition ($\beta = 1.29$, $p < 0.001$) and change in prevalence estimate ($\beta = 0.16$, $p < 0.001$), suggesting that larger increases in prevalence estimates of a belief led to larger increases in personal belief endorsement. There was no significant interaction ($p = 0.81$). The model accounted for 19.0% of the variance in belief change.

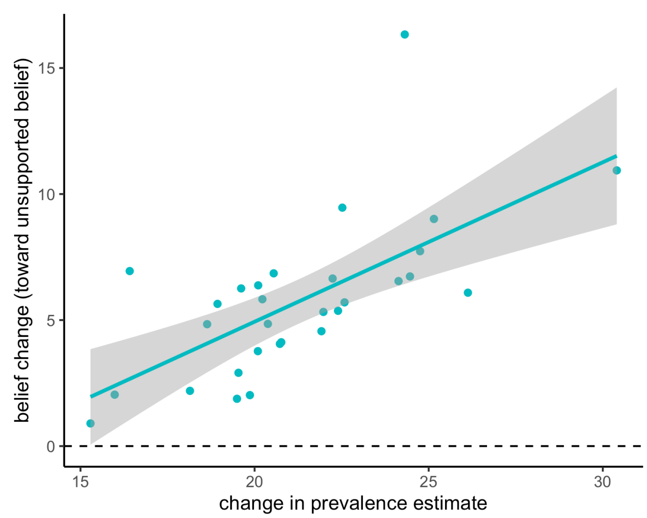

Figure 4: Mean belief change vs. mean change in prevalence estimate for each item in the Higher Prevalence condition. Line represents the fit from linear regression with 95% confidence intervals.

Belief change is dependent on initial certainty

We predicted that rational belief updating should also depend on the initial certainty of the belief. While novel evidence about social prevalence should bear significant weight under conditions of uncertainty, high certainty beliefs should be relatively resistant to updating regardless of social prevalence. Figure 5 illustrates how initial certainty about a belief predicts how that belief changed after exposure to both expected (control) and Higher Prevalence information.

To assess the relationship between initial certainty and belief change, we fit a linear mixed-effects model with standardized certainty and prevalence condition as fixed effects and random intercepts per participant and item. The model revealed main effects of both certainty ($\beta = 2.42$, $p < .001$) and prevalence condition ($\beta = 4.93$, $p < .001$), as well as a significant interaction in the predicted direction ($\beta = -2.84$, $p < .001$). The model accounted for 18.1% of the variance in belief change. In the Control condition, higher certainty predicted more belief change, likely because high certainty corresponded to the lower end of the likelihood scale where any movement had to be upward. However, in the Higher Prevalence condition where belief change was motivated by data, lower levels of certainty predicted higher belief change as hypothesized.

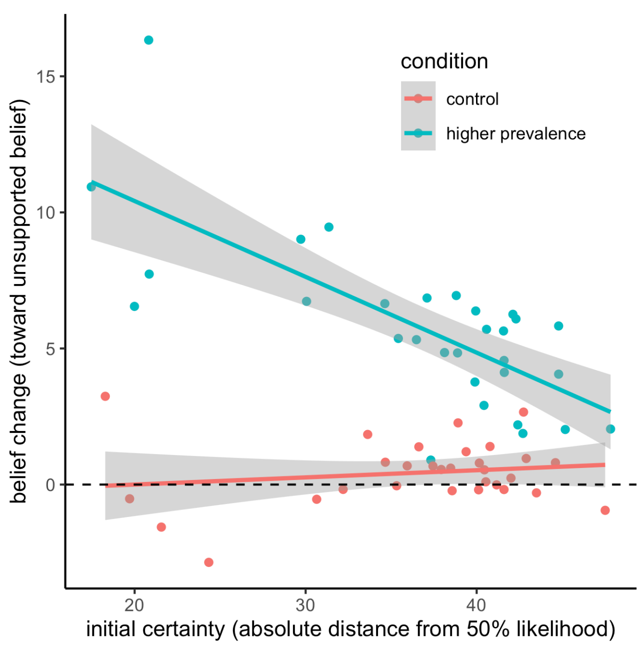

Figure 5: Mean belief change vs. mean initial certainty for each item, by condition. Lines represent the fit from linear regression for each condition, with 95% confidence intervals.

Discussion

In this experiment, we demonstrated that increasing people's perceptions of the general prevalence of a belief, divorced from any direct evidence, can cause them to endorse that belief more strongly. This prevalence-based belief change was significant for 27 of the 30 beliefs tested (see Figure 3), despite the fact that all items pertained to uncommon, empirically unsupported beliefs in a wide variety of domains. Participants were presented with new prevalence data from an anonymous 10-person sample that conflicted with their prior about a belief's prevalence by a uniform amount. Participants' initial certainty governed the extent to which they updated their estimates of each belief's prevalence in response to this new data. Larger changes in prevalence estimates, as well as lower initial certainty in beliefs, led to higher ultimate belief change. Taken together, these findings suggest that prevalence information serves as an independent cue that people rationally integrate with existing evidence to form a belief.

This work builds upon evolutionary models of social learning, clarifying the role of certainty in belief formation. Existing models have proposed a copy-when-uncertain social learning strategy, but precise interpretations vary. Learners are predicted to copy the behavior of others when they have no relevant private information or when their private information is unreliable or outdated (Kendal, Cooland & Laland, 2009). There is empirical evidence to support these models (e.g., Toelch, Bruce, Newson, Richerson & Reader, 2014), but they tend to assume a binary choice between copying social information or maintaining private information. Instead, we find evidence that social prevalence information is rationally integrated with prior beliefs. The reliability of new social evidence is evaluated against the initial certainty of the belief, and both sources of information are combined to form a final belief. This framework allows for social information to have a graded effect on belief change according to its relative reliability, and can accommodate cases in which such information is largely dismissed.

The malleability of participants' beliefs in line with the minimal social data from our experiment is particularly striking given that items were generally high-certainty and low-probability. The median initial likelihood rating for items in our study was a mere 7%. Social prevalence information may have an even stronger effect on other beliefs about which people are more ambivalent or generally lack information. In addition, participants were likely aware of the fact that all the items were related to pseudoscientific or conspiratorial claims, which may have made them more wary of the items and increased their motivation to remain consistent in their reported beliefs between blocks. The social prevalence manipulation may induce even stronger belief updating in more neutral contexts.

Demand characteristics are a concern for our experiment because participants were asked to re-rate their beliefs directly after new prevalence data was presented. However, there are several reasons to doubt that demand characteristics played a significant role in producing our pattern of results. First, we implemented a blocked design to limit participants' access to their initial prevalence estimates of the beliefs. Participants may not have been aware of the extent to which their prevalence estimates changed in the second block, reducing demand characteristics. Second, since the prevalence data in the Higher Prevalence condition differed from participants' original estimates by a fixed amount, demand characteristics should be consistent across trials. However, participants did not update their beliefs equally across all trials within this condition, but rather did so flexibly in accordance with their initial certainty about the beliefs. Finally, data shown in the Higher Prevalence condition, although significantly higher than participants' initial expectations, were usually still ambivalent. The modal distribution of data presented in this condition was 6 out of 10 people endorsing the belief. Such cases of mixed evidence likely bring weaker demand characteristics than data with higher consensus.

It is also important to note that participants in our experiment were not simply blindly updating their beliefs to match the prevalence information that they were shown. Recall that the prevalence data shown in the Control condition matched participants' initial estimates of the broader prevalence of a belief, and not their personal ratings of its likelihood. Participants were aware of this distinction; their initial prevalence estimates differed from their own initial likelihood estimates by a mean of 10.9% across conditions. Thus, before encountering the prevalence manipulation, participants demonstrated an implicit understanding that their belief may not be representative of that of the broader population. Further, the prevalence information in the Control condition primed this discrepancy by providing additional evidence matching their initial prevalence estimate. The fact that participants only modulated their beliefs in the Higher Prevalence condition suggests that the key factor was not simply a difference between their belief and the new prevalence data, but an error signal in their estimation of the prevalence per se. This comports with recent evidence that prediction error linearly predicts belief change (Vlasceanu, Morais & Coman, 2021), and further highlights social prevalence as an independent factor affecting human belief.

One limitation of the present study is that the nature of the experiment may have biased participants to assume the prevalence data they were shown was particularly reliable or representative of the general population. The association of the experiment with a research university may have increased the authoritativeness of the manipulated prevalence data, eliciting a stronger effect than may occur in more naturalistic contexts. However, there is some evidence indicating that this is of little concern. The critical manipulation in our experiment showed a sample with a belief prevalence 40% higher than the participant indicated that they expected, yet participants' final prevalence estimates changed by a mean of only 21%. Thus, participants did not uncritically trust the prevalence data, but rather integrated it with their existing belief about the prevalence of each claim. Still, future work should explore whether the method of presentation of prevalence data affects its influence on people's beliefs.

Implications for countering false beliefs online

Our findings have important implications for interventions against the spread of misinformation online. The existing interventions used by prominent social media sites typically consist of warning labels to flag false or misleading information on their posts. Although there is some evidence that this kind of intervention may be modestly effective (e.g., Clayton et al., 2020; Walter & Murphy, 2018), other studies find that it doesn't help at all (Grady, 2019; Pennycook, Cannon & Rand, 2018). In any case, posts flagged with warning labels typically still have visible social engagement metrics. Given our finding that inferring a conspiratorial belief is prevalent in a population increases one's likelihood of believing it themselves, engagement metrics are a dangerous cue that may work against the effectiveness of warning labels. Indeed, people tend to rate news from low-credibility sources as higher quality when engagement metrics are present (Chung, 2017). Further, high engagement metrics elicited more sharing and less fact-checking from participants in a simulated social media feed (Avram et al., 2020). This is of particular concern for viral misinformation posts, which can bear engagement numbers in the hundreds or thousands, as opposed to our experimental manipulation with data on 10 people. Hiding social engagement metrics entirely for posts relaying false or misleading information may therefore help reduce false belief. Future work should test interventions along these lines in an effort to counter the online misinformation crisis.

Acknowledgments

We thank Holly Palmeri, Emily Sumner, Carolyn Baer, and Steve Piantadosi for helpful comments on this work. We also thank the Berkeley Center for New Media, the Jacobs Foundation, the Hellman Scholars Fund, and a Georgia Lee Fellowship for funding in support of this work.

References

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. *Groups, leadership and men: Research in human relations,* 177.

Avram, M., Micallef, N., Patil, S., & Menczer, F. (2020). Exposure to social engagement metrics increases vulnerability to misinformation. *arXiv preprint arXiv:2005.04682.*

Bolsen, T., & Druckman, J. N. (2018). Do partisanship and politicization undermine the impact of a scientific consensus message about climate change?. *Group Processes & Intergroup Relations, 21*(3), 389-402.

Chung, M. (2017). Not just numbers: The role of social media metrics in online news evaluations. *Computers in Human Behavior, 75*, 949-957.

Clayton, K., Blair, S., Busam, J. A., Forstner, S., Glance, J., Green, G., ... & Nyhan, B. (2020). Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior, 42*(4), 1073-1095.

Cook, J., & Lewandowsky, S. (2016). Rational irrationality: Modeling climate change belief polarization using Bayesian networks. *Topics in Cognitive Science, 8*(1), 160-179.

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology, 51*(3), 629.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429-433.

Festinger, L. (1954). A theory of social comparison processes. *Human relations, 7*(2), 117-140.

Grady, R. (2019). *Examining, correcting, and failing to correct politically biased judgments and memories in real-world contexts.* Unpublished doctoral dissertation, University of California, Irvine.

Kendal, R. L., Coolen, I., & Laland, K. N. (2009). Adaptive Trade-offs in the Use of Social and Personal Information. In *Cognitive ecology II* (pp. 249-271). University of Chicago Press.

Kerr, J. R., & Wilson, M. S. (2018). Changes in perceived scientific consensus shift beliefs about climate change and GM food safety. *PloS one, 13*(7), e0200295.

Kim, J. (2020). The Meaning of Numbers: Effect of Social Media Engagement Metrics in Risk Communication. *Communication Studies*, 1-19.

Knill, D. C., & Richards, W. (Eds.). (1996). *Perception as Bayesian inference.* Cambridge University Press.

Lees, J., & Cikara, M. (2021). Understanding and combating misperceived polarization. *Philosophical Transactions of the Royal Society B, 376*(1822), 20200143.

Lewandowsky, S., Cook, J., Fay, N., & Gignac, G. E. (2019). Science by social media: Attitudes towards climate change are mediated by perceived social consensus. *Memory & cognition, 47*(8), 1445-1456.

Martí, L. Conover, A., & Kidd, C. (2021). *Prevalence of pseudoscientific beliefs*. Unpublished manuscript.

Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010). Social and heuristic approaches to credibility evaluation online. *Journal of communication, 60*(3), 413-439.

Nakagawa, S., Johnson, P. C., & Schielzeth, H. (2017). The coefficient of determination R2 and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of the Royal Society Interface, 14*(134), 20170213.

Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of experimental psychology: general, 147*(12), 1865.

Pennycook, G., & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition, 188*, 39-50.

Pennycook, G., & Rand, D. G. (2020). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of personality, 88*(2), 185-200.

Perfors, A., & Navarro, D. J. (2019). Why do echo chambers form? The role of trust, population heterogeneity, and objective truth. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society.*

Perfors, A., Navarro, D. J., & Shafto, P. (2018). Stronger evidence isn't always better: A role for social inference in evidence selection and interpretation. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society.*

Puhl, R. M., Schwartz, M. B., & Brownell, K. D. (2005). Impact of perceived consensus on stereotypes about obese people: a new approach for reducing bias. *Health Psychology, 24*(5), 517.

Stangor, C., Sechrist, G. B., & Jost, J. T. (2001). Changing racial beliefs by providing consensus information. *Personality and Social Psychology Bulletin, 27*(4), 486-496.

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 B: Biological Sciences, 281*(1776), 20132864.

Toelch, U., & Dolan, R. J. (2015). Informational and normative influences in conformity from a neurocomputational perspective. *Trends in cognitive sciences, 19*(10), 579-589.

van der Linden, S. L., Clarke, C. E., & Maibach, E. W. (2015). Highlighting consensus among medical scientists increases public support for vaccines: evidence from a randomized experiment. *BMC public health, 15*(1), 1-5.

van der Linden, S. L., Leiserowitz, A. A., Feinberg, G. D., & Maibach, E. W. (2015). The scientific consensus on climate change as a gateway belief: Experimental evidence. *PloS one, 10*(2), e0118489.

van der Linden, S., Leiserowitz, A., & Maibach, E. (2019). The gateway belief model: A large-scale replication. *Journal of Environmental Psychology, 62*, 49-58.

Van Prooijen, J. W., Krouwel, A. P., & Pollet, T. V. (2015). Political extremism predicts belief in conspiracy theories. *Social Psychological and Personality Science, 6*(5), 570-578.

Vlasceanu, M., Goebel, J., & Coman, A. (2020). The Emotion-Induced Belief Amplification Effect. In *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society.*

Vlasceanu, M., Morais, M. J., & Coman, A. (2021). The effect of prediction error on belief update across the political spectrum. *Psychological Science.*

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science, 359*(6380), 1146-1151.

Walter, N., & Murphy, S. T. (2018). How to unring the bell: A meta-analytic approach to correction of misinformation. *Communication Monographs, 85*(3), 423-441.