"Fringe" beliefs aren't fringe

Louis Martí (I_Am_A_Robot@berkeley.edu)

Adam Conover

Celeste Kidd (celestekidd@berkeley.edu)
Department of Psychology, 2121 Berkeley Way
Berkeley, CA 94704 USA

Abstract

COVID-19 and the 2021 U.S. Capitol attacks have highlighted the potential dangers of pseudoscientific and conspiratorial belief adoption. Approaches to combating misinformed beliefs have tried to "pre-bunk" or "inoculate" people against misinformation adoption and have yielded only modest results. These approaches presume that some citizens may be more gullible than others and thus susceptible to multiple misinformed beliefs. We provide evidence of an alternative account: it's simply too hard for all people to be accurate in all domains of belief, but most individuals are trying. We collected data on a constellation of human beliefs across domains from more than 1,700 people on Amazon Mechanical Turk. We find misinformed beliefs to be broadly, but thinly, spread among the population. Further, we do not find that individuals who adopt one misinformed belief are more likely to engage in pseudoscientific or conspiratorial thinking across the board, in opposition to "slippery slope" notions of misinformation adoption.

Keywords: Pseudoscience; belief formation; misinformation; selection bias correction.

Introduction

Recent events surrounding QAnon and COVID-19 conspiracies have highlighted the potential dangers of misinformation (Romer & Jamieson, 2020; Woko, Siegel, & Hornik, 2020; Amarasingam & Argentino, 2020). In response, efforts to "pre-bunk" the conspiracies or "inoculate" the population against the spread of misinformation have arisen (Maertens, Roozenbeek, Basol, & van der Linden, 2020; Pennycook & Rand, 2020; Roozenbeek & van der Linden, 2019), with thus far only modest results (Banas & Rains, 2010). The potential robustness of this approach depends upon an untested assumption of human psychology: that some individuals are more gullible than others, and that they can be made less gullible through training.

Previous studies of conspiratorial and pseudoscientific belief focus on populations who hold misinformed beliefs in a single domain, for example flat earthers (Landrum, 2019), anti-vaxxers (Martinez-Berman, McCutcheon, & Huynh, 2020), climate change deniers (Uscinski, Douglas, & Lewandowsky, 2017), or incels (Young, 2019). However, examining beliefs in a single domain would necessarily make any variation appear as though some individuals are inherently more gullible than others. In fact, it is quite possible that everyone is trying their best to form beliefs that align with objective evidence in the world, and generally doing a decent but imperfect job across multiple knowledge domains.

We compare two hypotheses: one is that some individuals are fundamentally gullible and therefore are susceptible to multiple misinformed beliefs. The other is that it is simply hard to be accurate in all domains of belief, but most individuals are trying. In that case, we would expect misinformed beliefs to be broadly, but thinly, spread among the population.

To understand how beliefs arise and spread, you must look at constellations of beliefs. Here, we do just that. We collect a large set of judgements on a host of different types of beliefs—including conspiracies, pseudoscience, and other non-evidence-based beliefs[1]. We use this data to understand the prevalence overall of many misinformed beliefs, as well as whether belief in one tends to predict belief in others, as is widely espoused in "slippery slope" arguments (Wood, Douglas, & Sutton, 2012).

Our results demonstrate that rather than some portion of the population being gullible, most people hold one or more non-evidence-based beliefs. Our results suggest most of these beliefs do not predispose individuals to becoming more likely to adopt many other non-evidenced beliefs. The fact that misinformed beliefs are ubiquitous and generally not "gateway drugs" to other networks of misinformed beliefs has widespread implications for how we should structure efforts to combat misinformation in the world.

Methods

We recruited 2,036 participants using a custom built web interface on Amazon Mechanical Turk on November 24, 2020. We required participants to be from the U.S. and have at least a 95% approval rating from previous tasks. Responses were recorded on a secure server. After consenting to the experiment, participants were asked to type out a series of sentences, pledging to answer questions honestly. This was followed by a nine question demographics questionnaire (age, sex, race, ethnicity, state of residence, education, income, religion, and politics). Next, participants entered two practice trials where they rated the likelihood of statements ("Plants

[1]The misinformed beliefs we evaluated here included "conspiracies" and "pseudoscience", as well as other non-evidence-based beliefs, most but not all of which are commonly labeled "fringe beliefs". We recognize that these terms are not interchangeable, but investigate all to broadly understand how misinformed beliefs relate to one another in the population.

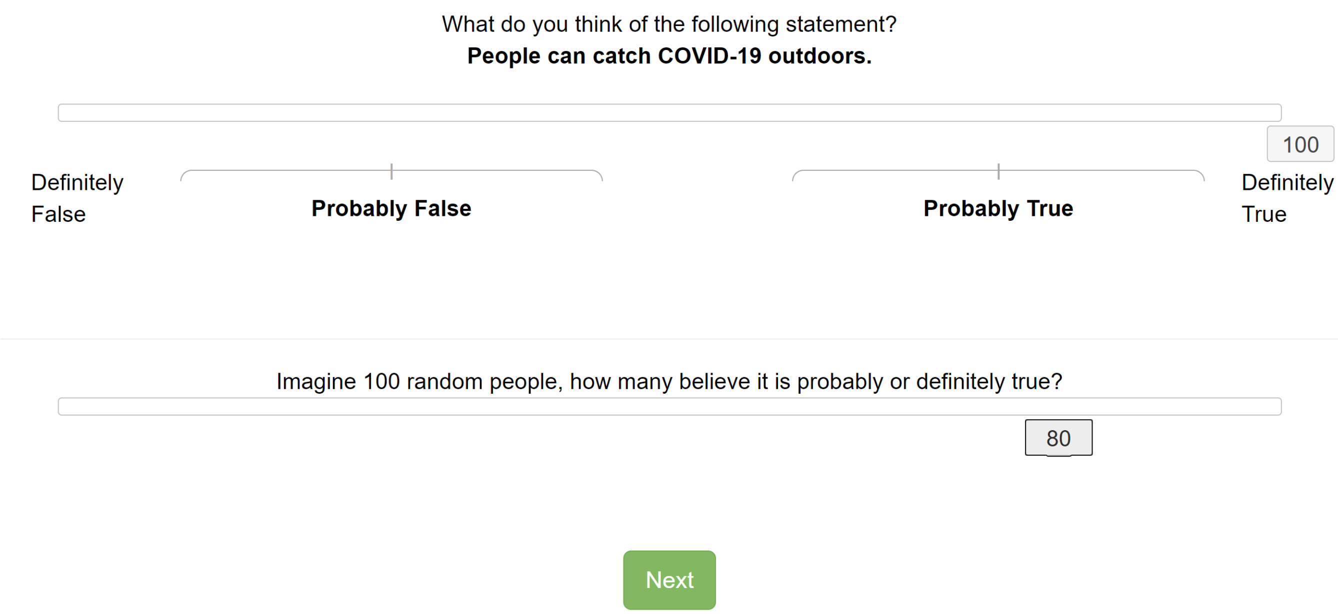

Figure 1: Participants saw 60 randomized trials as above.

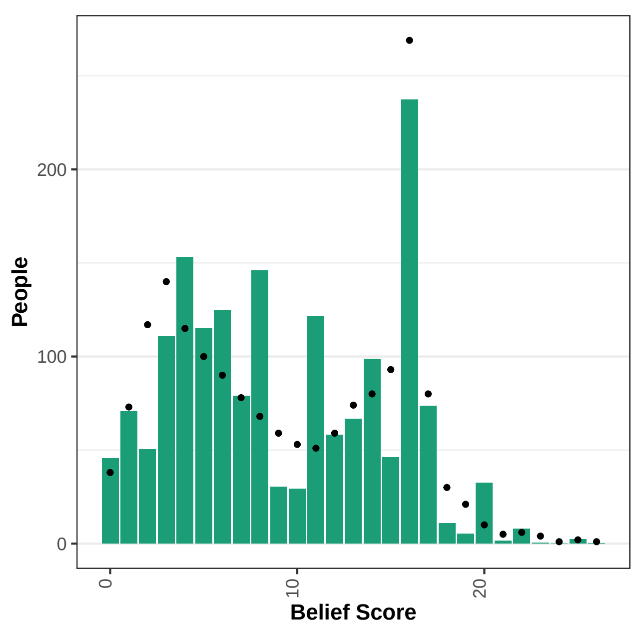

Figure 2: A histogram showing the number of people who believe x number of misinformed statements are more likely than not. The bars show data after sampling bias correction while the black dots show data before the correction. The maximum possible score is 30, one point for each statement. The median participant believes 9 out of 30 non-evidence-based statements. Despite being a conservative measure (we only tested a small minority of non-evidenced beliefs), we find that misinformed beliefs are widespread.

need water to grow." and "Birds lay eggs.") using a slider bar, and also guessed how many other people would find them to be probably or definitely true. Then, participants saw 60 more trials of the same format (see Figure 1). In order to test for differences between a broad-range of misinformed beliefs and those within a more narrow scope, each of these trials displayed one of 18 general statements (half non-evidence-based, half factual) or 12 COVID-19-relevant statements (seven non-evidence-based, five factual). Each statement was presented twice for a total of 60 trials as a way of assessing reliability. After every 15 trials, a free-response catch question was asked ("What is your favorite drink?", "What is your favorite movie?", "What is your favorite snack?", "What is your favorite aquatic animal?") to be used to filter careless or automated responses from polluting our participant pool.

Analysis

We used our data to estimate overall belief prevalence in the U.S. by correcting the sampling bias in our Amazon Mechanical Turk data. The Amazon Mechanical Turk's participant pool tends to be more white, male, young, and poor than the general U.S. population. (Moss, Rosenzweig, Robinson, & Litman, 2020)

To correct for this sampling bias, we employed an "iterative proportional fitting", or "raking" technique (Deming & Stephan, 1940). Raking applies a weight to each participant to offset sampling bias. For example, if the true proportion of males in the population is 50% but your sample is only 25% male, raking will apply a weight of 3 to each male and 1 to each female. If this process is performed for more than one variable, adjusting a participant's weight to match the true proportions for one variable may ruin the weight value for

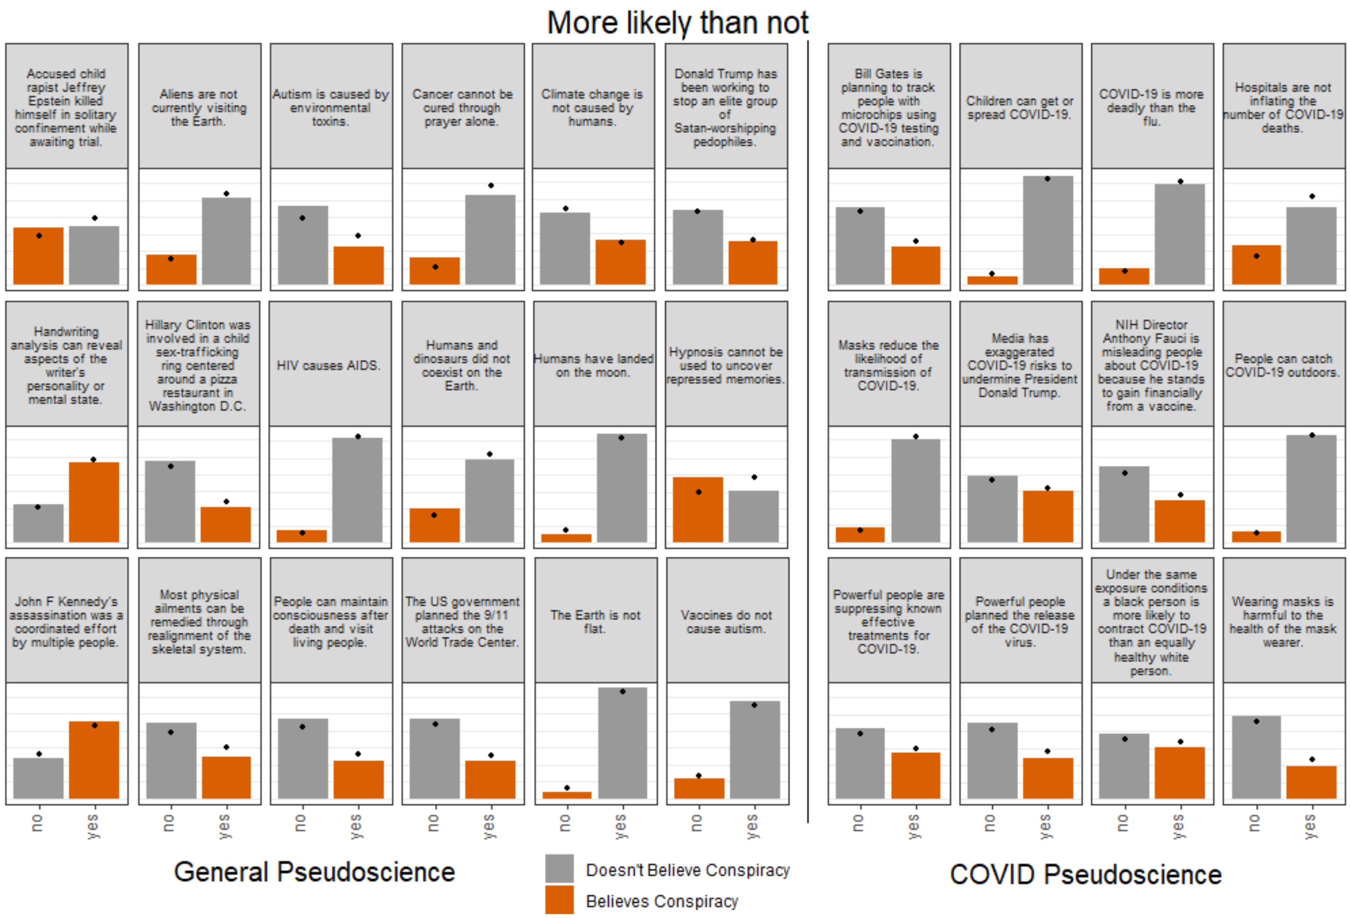

Figure 3: Proportions of people who believe each statement is more likely than not. Orange bars represent individuals who endorsed a misinformed statement while gray bars represent a rejection. Bars show data after sampling bias correction while the black dots show data before the correction. Misinformed beliefs are common and in certain cases, represent the majority view.

another variable. To correct for this, the algorithm is run for many iterations and only stops when all weighted proportions are within a set threshold (ε) of their true proportions. Our algorithm ran for 100 iterations with $\varepsilon = .000005$.

The true proportions for age, sex, race, ethnicity, state of residence, education, and income were calculated using the 2014–2018 American Community Survey Public Use Microdata Sample (PUMS) from the United States Census Bureau, after excluding all individuals under 18 (U.S. Census Bureau., 2020). The true proportions for religious and political affiliations were taken from the Pew Research Center 2014 Religious Landscape Study (Pew, 2014). We followed raking best practices (Battaglia, Izrael, Hoaglin, & Frankel, 2009) including combining certain demographic categories with extremely few entries in our data. The combined categories are "American Native or Other" for race, "Non-Mexican Hispanic/Latino" for ethnicity, "Master's/Doctorate/Professional degree" and "No High School Diploma or GED" for education, "Mainline Protestant", "Other Non-Christian", "Other Christian", and "Nothing in particular (religion not important)" for religion, and "Moderate" for politics.

Exclusion criteria

We applied a conservative exclusion criteria to our data which maintained 1,717 participants and excluded those who demonstrated inattention or who failed to demonstrate their humanity. We excluded 195 participants for not following the pledge-typing instructions at the onset of the experiment and 124 participants for not providing at least 3 out of 4 valid catch question responses.

For the remaining 1,717 participants, we examined their reliability by first labelling a repeated sentence as a bad trial if the participant gave likelihood scores that were greater than or equal to 20 points apart (out of 100). Using this criteria, the mean reliability for our participants is 88.7%, meaning 88.7% of all trial pairs were rated as less than 20 points apart. Note that even with a 20-point criteria (which is very conservative) our observed reliability remains very high. We do not exclude any participants with low reliability since it is possible that different likelihood scores on the same items is a reflection

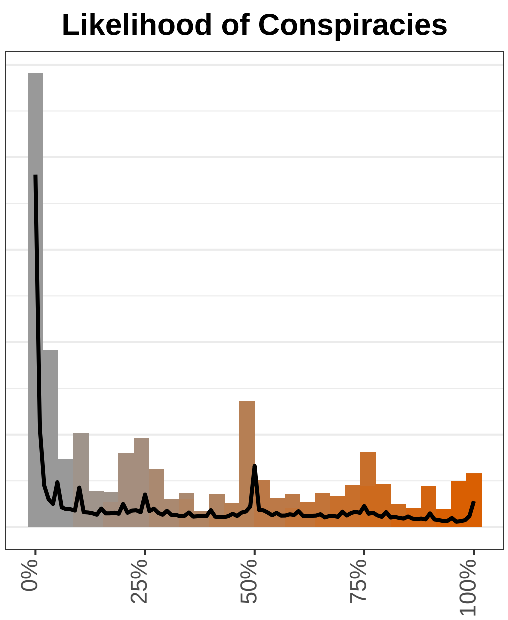

Figure 4: Weighted likelihood responses across all sentences binned by whether the sentences are non-evidence-based. The unweighted data is in the form of a line. The most common response by far is a total rejection of misinformed statements. When a misinformed statement is judged as possible, there is a large amount of variance in responses, indicating high uncertainty which implies these beliefs could be self-correcting over time. Very few responses indicate a non-evidence-based statement is 100% certain, signalling intellectual humility. Note that weighing the raw data raises prevalences relatively uniformly.

about their uncertainty, not about a lack of attention. [2]

Results

Fringe beliefs are not fringe

We calculated a belief score for each participant representing the number of non-evidence-based statements that they believe are more likely than not.[3] While any given misinformed belief may be uncommon, when examined in aggregate, the vast majority of people believe in at least several misinformed statements (see Figure 2). In our sample, 98% of participants believe at least one statement, and 52% believe at least **nine**. It is important to note that this measure is conservative, as we only tested a small minority of all non-evidence-based beliefs. Interestingly, unsubstantiated COVID-19 beliefs are substantially less common than all other misinformed beliefs, possibly because they have had less time to spread throughout the population, or perhaps due to public health messaging (see Figure 3). The median General Pseudoscience score is 6 (out of a possible 18) while the median COVID Pseudoscience score is 3 (out of a possible 12).

Fringe beliefs are often weakly held

While these beliefs are widespread, they are not strongly held. Figure 4 shows the weighted aggregate of all responses across all sentences. Among all misinformed statements that participants believe are more likely than not, there still exists a high degree of uncertainty. Very few non-evidence-based statements are rated as 100% true, but the same level of uncertainty is not present when participants rule out a misinformed belief (an overwhelming number of responses rate them as 0% likely).

Only some beliefs are possible gateways

As Figure 5 shows, only some beliefs provide evidence for the slippery slope argument. Each boxplot partitions participants depending on whether or not they believe a particular misinformed statement. If a slippery slope existed, participants who believed any given misinformed statement would have a significantly higher conspiracy score than the participants who did not believe it. In other words, the boxes (which represent 95% confidence intervals) within a statement would not overlap and the conspiracy box would be higher. Instead, we find that this is only the case for 13 out of the 30 statements. As supported by a linear regression, if you believe that aliens are currently visiting the Earth, you are no more likely to believe other misinformed beliefs ($\beta = -0.0588$, $SE = 0.0048$, $t = -12.14$, $p < 0.001$). On the other hand, if you believe autism is caused by environmental toxins, you *are* more likely to believe other misinformed beliefs, as a linear regression confirms ($\beta = 0.1377$, $SE = 0.0036$, $t = 38.70$,

[2] Applying reasonable reliability exclusion criteria ends in very few participants being excluded and does not change our overall results.

[3] This score was reverse coded for the 14 statements that were factual.

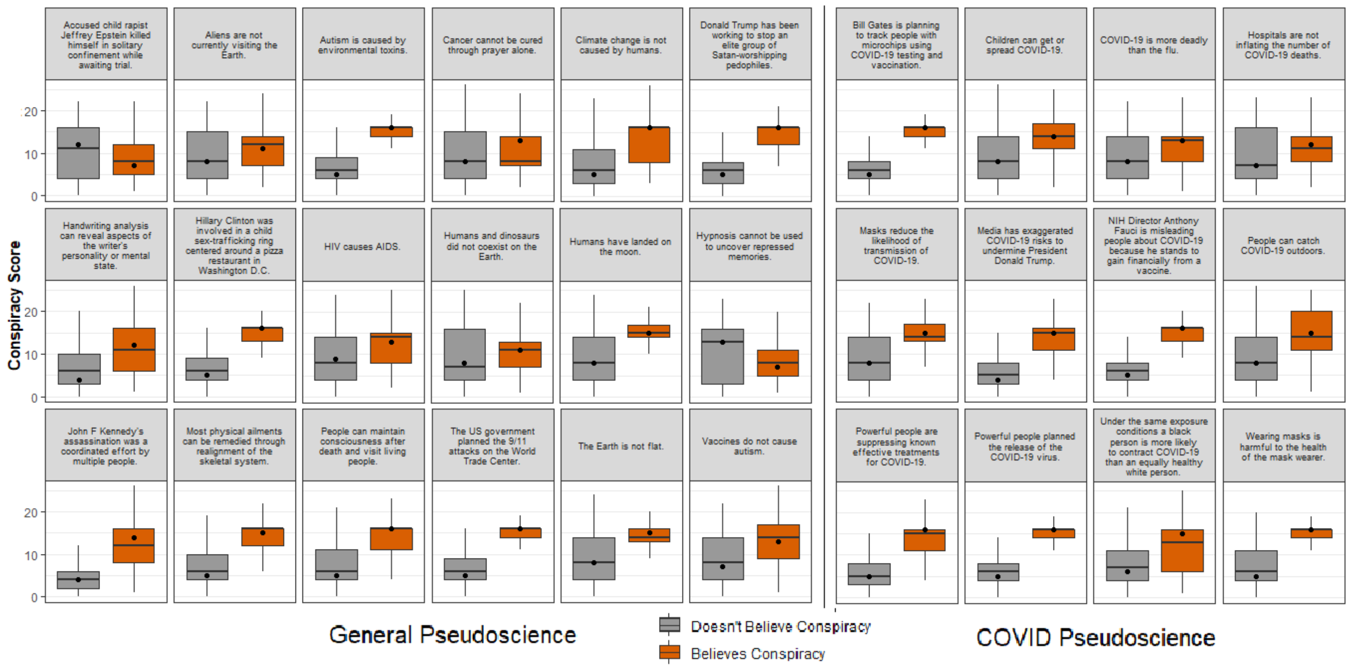

Figure 5: Conspiracy scores (total number of misinformed statements believed) binned by whether participants believe a particular statement. Boxplots are weighted data, black dots are the unweighted medians. Orange boxes represent individuals who believe a non-evidence-based statement while gray boxes represent non-believers. Boxes show the median 50% quantiles of conspiracy scores. Non-overlapping boxes within a statement indicate 95% confidence that the true medians differ. Most box pairs overlap indicating an overall lack of support for the slippery slope theory.

$p < 0.001$). It is important to note that this is still not a confirmation of a slippery slope in these cases, as causality would need to be determined. Comparing COVID-19 scores in the same manner results in 7 out of 12 significant differences, indicating that perhaps the slope is a bit slipperier when dealing with misinformed beliefs which are closely related.

In aggregate, participants are very good at predicting beliefs

As Figure 6 shows, the predictions participants made about the prevalence of these beliefs were very accurate. With each data point representing a different statement, all of them are either just above or just below the line of perfect prediction $y = x$. A generalized logistic regression confirms this ($\beta = 1.6201$, $SE = 0.0344$, $z = 47.12$, $p < 0.001$).

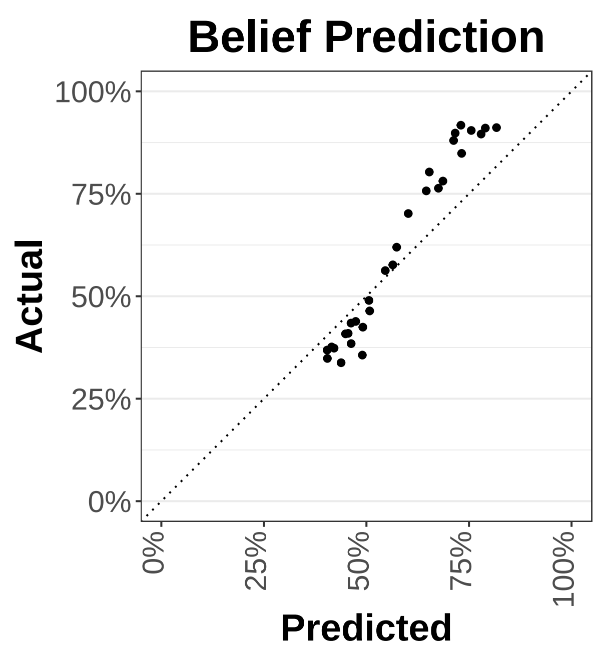

Figure 6: Mean predictions and the actual percentage of participants who believe each statement is more likely than not. Perfect predictions would lie along the $y = x$ line; participants are very good at predicting the beliefs of others in aggregate.

Conclusions and Discussion

Summary

Although any given misinformed belief may be uncommon, in aggregate these beliefs are extremely common. The median number of non-evidence-based beliefs held by individuals is 9, and 98% of people believe at least one. These beliefs are almost always accompanied by a degree of uncertainty however.

We find only modest evidence to support the slippery slope theory. Most endorsed beliefs do not show a significantly higher conspiracy score. Only when looking at closely related beliefs such as beliefs related to COVID-19, does belief in one statement tend to predict belief in the others, signaling that individuals do not view all theories equally.

Implications

We have demonstrated that misinformed beliefs are widespread throughout the population, but often weakly held. This pattern is consistent with the idea that people are trying their best to understand what is true in the world but making errors. Since uncertainty is associated with higher degrees of interest and curiosity, we would expect these particular misinformed beliefs to be more likely to update and self-correct as further evidence is sampled and integrated. Thus, these kinds of misinformed beliefs—those which are held with higher degrees of uncertainty—are less problematic in context because we would not expect them to be stubbornly held.

Further evidence in support of this idea is our finding that misinformed COVID-19 beliefs were less common than misinformed beliefs at large. If this measurable difference can be directly linked to the efficacy of public policy education efforts, it will provide clear evidence that effective interventions to combat misinformation should provide people with clear, specific evidence. We suspect that, given the patterns we observed here, these approaches may be more effective than generally trying to train people against dubious belief adoption at large.

Social prevalence encourages misinformed beliefs

The "illusory truth effect" says that people will tend to believe something more strongly the more they encounter it (Hasher, Goldstein, & Toppino, 1977). The effect starts with as little as two exposures, has been found in children as young as 5 (Fazio & Sherry, 2020), and is just as strong regardless of an individual's cognitive ability, need for cognitive closure, or cognitive style (De Keersmaecker et al., 2020). Compounding this effect is evidence that individuals overestimate the amount of information they use to form a belief (Klein & O'Brien, 2018). These facts—in tandem with their relationship to relatively sudden recent changes in how people gather their information from the world—may help explain the prevalence we observe here.

We note that the most commonly endorsed misinformed beliefs in our study are, by definition, *not* fringe beliefs. Certain statements, such as the belief that handwriting analysis can reveal an individual's personality, are endorsed not only by certain authorities and popular culture, but by the majority of people in the world (70% in our sample). Since we know that social prevalence is one important cue people use in order to infer what is true (Orticio, Marti, & Kidd, 2021), the more people already believe this kind of non-evidence-based belief, the more we'd expect will believe it in the future. This isn't irrational given the limitations of human access to truth. Since no one has direct access to truth, the best a person can do engage in inference from sparse, often indirect sources of evidence, like the opinions of others.

Misinformed beliefs in informational ecosystems

The changing way in which people sample evidence from the world in order to infer truth is important to consider in light of our results.

People increasingly rely on online sources for information, and almost all recommend content based on algorithms that maximize engagement (clicks, likes, comments, or hang time on the visual image). Maximizing engagement likely results in promoting material that is *less likely* to be true. For example, recent work suggests interestingness-if-true is a strong predictor of news sharing, not the user's assessment of its likelihood of actually being true (Altay, de Araujo, & Mercier, 2020). This is important because we know that estimates of social prevalence increase the likelihood of the adoption of misinformation as belief (Orticio et al., 2021).

The engagement-based reward systems used by platforms like YouTube, TikTok, and Twitter to incentivize content creation also likely add bias to the information pool. Creators are rewarded with views, likes, or money proportional to how many sets of eyes their content attracts and the duration they keep them glued—not the veracity of what they post. These pressures likely incentivize the creation of sensationalized, conspiratorial, and fringe content due to its novelty and subsequent interest.

If our information ecosystem is polluted, it is particularly problematic in light of human fallibility. People have built-in mechanisms designed to help them sparsely sample information in the world to draw quick inferences and act (Kidd & Hayden, 2015; Wade & Kidd, 2019). People seek out information to reduce their uncertainty, but move from uncertain to relatively sure on the basis of heuristics like feedback (Martí, Mollica, Piantadosi, & Kidd, 2018). Once certain, people tend to stick stubbornly with their established beliefs, and it is difficult to prompt them to revise. These cognitive mechanisms are useful in preventing people from wasting time from material that is already understood.

However, this aspect of human psychology may be problematic in environments with misleading feedback signals and which offer several points of related feedback quickly, as is true when people seek answers online.

Future Directions

Future research should examine the complexities of these beliefs that our analysis ignores. Examining higher-order beliefs could better uncover relationships between beliefs. An individual who holds the general belief that governments hide the truth from the public might be more likely to believe specific beliefs such as 9/11 was an inside job and that Epstein was murdered.

Our work in progress examines demographic predictors of misinformed beliefs using the rich demographic information collected from participants. We are also developing novel methods for examining our data for clusters of beliefs, with particular attention to COVID and QAnon clusters. We will also investigate avenues for quantifying potential harms for each belief and examine whether harmful beliefs are more or less common.

Acknowledgements

We would like to thank members of the Kidd Lab for providing valuable feedback, especially Sarah Stolp, David O'Shaughnessy, and Holly Palmeri. This work was supported by awards from DARPA (Machine Common Sense TA1, BAA number HR001119S0005), Berkeley Center for New Media, the Jacobs Foundation, and the Georgia Lee Fellowship through the Hellman Scholars Fund.

References

Altay, S., de Araujo, E., & Mercier, H. (2020). "if this account is true, it is most enormously wonderful": Interestingness-if-true and the sharing of true and false news.

Amarasingam, A., & Argentino, M.-A. (2020). The qanon conspiracy theory: A security threat in the making. *CTC Sentinel*, *13*(7), 37–44.

Banas, J. A., & Rains, S. A. (2010). A meta-analysis of research on inoculation theory. *Communication Monographs*, *77*(3), 281–311.

Battaglia, M. P., Izrael, D., Hoaglin, D. C., & Frankel, M. R. (2009). Practical considerations in raking survey data. *Survey Practice*, *2*(5), 1–10.

De Keersmaecker, J., Dunning, D., Pennycook, G., Rand, D. G., Sanchez, C., Unkelbach, C., & Roets, A. (2020). Investigating the robustness of the illusory truth effect across individual differences in cognitive ability, need for cognitive closure, and cognitive style. *Personality and Social Psychology Bulletin*, *46*(2), 204–215.

Deming, W. E., & Stephan, F. F. (1940). On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, *11*(4), 427–444.

Fazio, L. K., & Sherry, C. L. (2020). The effect of repetition on truth judgments across development. *Psychological Science*, *31*(9), 1150–1160.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of verbal learning and verbal behavior*, *16*(1), 107–112.

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, *88*(3), 449–460.

Klein, N., & O'Brien, E. (2018). People use less information than they think to make up their minds. *Proceedings of the National Academy of Sciences*, *115*(52), 13222–13227.

Landrum, A. (2019). Youtube as a primary propagator of flat earth philosophy. In *2019 annual meeting.*

Maertens, R., Roozenbeek, J., Basol, M., & van der Linden, S. (2020). Long-term effectiveness of inoculation against misinformation: Three longitudinal experiments. *Journal of Experimental Psychology: Applied*.

Martí, L., Mollica, F., Piantadosi, S., & Kidd, C. (2018). Certainty is primarily determined by past performance during concept learning. *Open Mind*, *2*(2), 47–60.

Martinez-Berman, L., McCutcheon, L., & Huynh, H. P. (2020). Is the worship of celebrities associated with resistance to vaccinations? relationships between celebrity admiration, anti-vaccination attitudes, and beliefs in conspiracy. *Psychology, health & medicine*, 1–10.

Moss, A. J., Rosenzweig, C., Robinson, J., & Litman, L. (2020). Demographic stability on mechanical turk despite covid-19. *Trends in cognitive sciences*, *24*(9), 678–680.

Orticio, E., Marti, L., & Kidd, C. (2021). Beliefs are most swayed by social prevalence under uncertainty. In *Cogsci.*

Pennycook, G., & Rand, D. (2020). The cognitive science of fake news.

Pew. (2014). Religious landscape study. *Pew Research Center*.

Romer, D., & Jamieson, K. H. (2020). Conspiracy theories as barriers to controlling the spread of covid-19 in the us. *Social Science & Medicine*, *263*, 113356.

Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, *5*(1), 1–10.

U.S. Census Bureau. (2020, January). *2014-2018 american community survey 5-year public use microdata samples.*

Uscinski, J. E., Douglas, K., & Lewandowsky, S. (2017). Climate change conspiracy theories. In *Oxford research encyclopedia of climate science.*

Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic bulletin & review*, *26*(4), 1377–1387.

Woko, C., Siegel, L., & Hornik, R. (2020). Explaining the association of race and covid-19 vaccination intentions: the role of behavioral beliefs and trust in covid-19 information sources.

Wood, M. J., Douglas, K. M., & Sutton, R. M. (2012). Dead and alive: Beliefs in contradictory conspiracy theories. *Social psychological and personality science*, *3*(6), 767–773.

Young, O. (2019). What role has social media played in violence perpetrated by incels?