Can Invalid Information be Ignored when it is Detected?

Adam T. Ramsey¹, Yanjun Liu², Jennifer S. Trueblood²

¹ Department of Psychology, Vanderbilt University

² Department of Psychological and Brain Sciences, Indiana University

Author Note

Correspondence may be addressed to Jennifer S. Trueblood, Department Psychological and Brain Sciences, Indiana University, 1101 E 10th St, Bloomington, IN 47405; Email: jstruebl@iu.edu

Abstract

With the rapid spread of information via social media, individuals are prone to misinformation exposure that they may utilize when forming beliefs. Over five experiments (total $N=815$), we investigated whether people could ignore quantitative information when they judged for themselves that it was misreported. Participants recruited online viewed sets of values sampled from Gaussian distributions to estimate the underlying means. They attempted to ignore invalid information, which were outlier values inserted into the value sequences. Results indicated participants were able to detect outliers. Nevertheless, participants' estimates were still biased in the direction of the outlier, even when they were most certain that they detected invalid information. The addition of visual warning cues and different task scenarios did not fully eliminate systematic over- and under-estimation. These findings suggest individuals may incorporate invalid information they meant to ignore when forming beliefs.

Keywords: misinformation, belief updating, outliers, estimation, judgment

Can Invalid Information be Ignored when it is Detected?**Statement of Relevance**

Data journalism has become a part of daily life. Take for example coverage of the COVID-19 pandemic, during which one may hear any number of arguments over “the numbers”. To form beliefs grounded in truth, it is important to accurately judge which data are legitimate, and avoid any that are invalid. In these experiments, we examined whether people are able to ignore invalid numerical information once they have detected it. We found that even when individuals were certain that a piece of information was invalid, they were not able to fully ignore it. This was true even when we provided warnings, or told people which data were invalid. These findings are important because they show how harmful simply encountering bad data can be. This highlights the importance of news and social media sites making concentrated efforts to fight misinformation before it can make it to the public; before its too late.

Introduction

With the rapid spread of information via social media, individuals are susceptible to seeing misinformation that could impact their beliefs. This is particularly true of infodemics involving massive amounts of information about a particular topic (e.g., COVID-19), including false and misleading information (Eysenbach, 2002; Greenspan and Loftus, 2021). Concerningly, false news often spreads more quickly and broadly online than true news (Vosoughi et al., 2018). The ease with which misinformation is disseminated online produces an environment where a few insistent voices sharing false information can sway much of the populace (Cook and Lewandowsky, 2016). Even those who do not intend to misinform may do so by sharing false news stories simply because they have seen the headline before (Effron and Raj, 2020).

It can be easy to think that identifying misinformation allows one to ignore it, but research into the continued influence effect (CIE; H. M. Johnson and Seifert, 1994) has shown that people will utilize information that has been retracted, even if they remember that the information is not legitimate. This effect is difficult to eliminate completely. Invalidated information continues to influence beliefs in circumstances where the retraction was much stronger than the false information (Ecker et al., 2011), and when people are forewarned about misinformation (Ecker et al., 2010).

Interventions to reduce the influence of retracted information include shifting one's focus towards evaluating accuracy when encoding false information (Pennycook et al., 2021). Similarly, reading a debunking message shortly after encountering false information diminishes, but does not eliminate, the CIE (Brashier et al., 2021; Wilkes and Leatherbarrow, 1988). Debunking messages that include truthful information which replaces the retracted information in one's mental model tend to be the most effective (Chan et al., 2017) and have been recommended for news sources, social media sites, and educators (Lewandowsky et al., 2012). However, these interventions apply to situations where an outside source has invalidated a piece of information. It is also important to consider whether individuals can ignore information that they have deemed invalid for themselves.

In the illusory truth paradigm, participants indicate the degree to which they believe different statements, some of which are false. Later, participants are presented with another list of statements, some of which they had previously encountered. Results show that people more strongly believe previously encountered statements, including false statements (Begg et al., 1992; Hasher et al., 1977). Additional repetitions of a statement increase belief of that statement even further (Hassan and Barber, 2021).

Illusory truth occurs even when one possesses knowledge that contradicts false information (Fazio et al., 2015). The increase in believability occurs for both plausible and implausible statements (e.g., “The Earth is a perfect square”; Fazio et al., 2019). Like the CIE, the illusory truth effect can be diminished, but rarely eliminated, by having participants focus on the accuracy of statements as they initially encounter them (Brashier et al., 2020).

Other research has also observed that ignoring information is not an easy task. 4-6 year old children found it difficult to ignore false information about a previous playdate (Schaaf et al., 2015). Adult jurors have difficulty ignoring inadmissible evidence when deliberating a verdict (London and Nunez, 2000). Even experienced judges who ruled evidence to be inadmissible struggle to ignore that evidence (Wistrich et al., 2004).

The previous literature has mainly investigated whether individuals utilize CIE or believe (illusory truth) false factoids, how misinformation spreads online (Effron and Raj, 2020; Lewandowsky et al., 2012; Vosoughi et al., 2018), and how to combat it (Brashier et al., 2021; Pennycook et al., 2021). The current paper extends this literature by investigating how individuals handle misinformation in the form of invalid reports of numerical quantities.

News sources are increasingly reporting data to consumers (Westlund and Hermida, 2021), with data journalists often presenting their findings as fact. However, they less often acknowledge limitations such as data collection practices or conflicts of interest in their work. This requires consumers to be vigilant and sample many disparate findings of varying quality in order to estimate the truth of a matter (e.g., Covid-19 vaccine efficacy rates, global temperature change, etc.). Therefore, it is important to understand how people process noisy numerical

information as they form beliefs of the underlying truth (Stubenvoll and Matthes, 2021), especially information they deem to be false.

In this paper we describe five experiments investigating whether people could ignore invalid numerical information when they detected it. We also examine manipulations aimed at helping people identify false reports (i.e., visual warning cues). Across our experiments, we show that it is difficult for people to disregard invalid numerical information even when this information is very easy to detect.

Method

In a series of five experiments, participants read imaginary scenarios in which they were presented sequences of reports including numerical information sampled from underlying Gaussian distributions. They attempted to ignore invalid reports while estimating the true mean of the underlying distribution. These five experiments were conducted sequentially, and the specific manipulations used in later experiments were inspired by findings of the previous ones.

The design and analysis plans for Experiments 3, 4, and 5 were pre-registered at AsPredicted and can be accessed at https://aspredicted.org/NNB_QGM, https://aspredicted.org/OAB_ATF, and https://aspredicted.org/MTN_FZN. Example experimental code (for Experiment 5) can be accessed at: https://osf.io/9ybnx/?view_only=10c1c405a3884898b6806e919a0b3a98.

Participants

For each experiment, we intended to have approximately 50 participants per condition. In total, we recruited 106 adult participants for Experiment 1, 107 adult participants for Experiment 2, 382 adult participants for Experiment 3, 134 adult participants for Experiment 4, and 274 adult participants for Experiment 5. The sample size was determined prior to starting recruitment and was based on previous studies of the CIE and illusory truth effect (Ecker et al., 2011; Hassan and Barber, 2021).

We recruited participants from Amazon Mechanical Turk (MTurk) using the CloudResearch platform, and targeted those who typically submitted good work as indicated by their over 95% approval rating for all human intelligence tasks they had completed. The study was approved by the Institutional Review Board at the first author's university. Each participant received \$0.50 for every 15 minutes of participation. This payment rate was determined by the median earning rate of Mturk workers (Hara et al., 2018). In Experiments 1 and 2, participants received \$1.00 upon completion of the approximately 30-minute task. In Experiment 3-5, participants received \$1.50 upon completion of the approximately 45-minute task. Data was analyzed only after all data had been collected.

Due to the nature of the experimental manipulation, participants needed to view at least three reports (i.e., stimuli) in a trial to possibly encounter the invalid report of interest to our hypotheses (see Procedure section). Therefore, participants with a median number of stimuli viewed per trial of less than three were excluded from the analyses. This criterion also served to protect against participants who merely clicked through the experiment as fast as possible without dutifully completing the task. After excluding participants due to low median stimuli viewed per trial, we retained 85 participants (45 women, 39 men, 1 non-binary; age: $M = 39.88$, $SD = 13.31$) in Experiment 1, 83 participants (49 women, 31 men, 1 non-binary and 2 unknown; age: $M = 41.37$, $SD = 13.70$) in Experiment 2, 314 participants (197 women, 112 men, 1 non-binary, 1 genderqueer, and 1 transmasculine participant, and 2 unknown; age: $M = 43.75$, $SD = 14.25$) in Experiment 3, 103 participants (76 women, 27 men; age: $M = 42.51$, $SD = 13.47$) in Experiment 4, and 230 participants (156 women, 72 men, 1 non-binary and 1 unknown; age: $M = 40.94$, $SD = 14.42$) in Experiment 5 for data analyses.

Materials

Each trial consisted of a sequence of screens, each showing the results from a hypothetical medical study (e.g., Figure 1A) involving a certain number of patients (e.g., Figure 1B). Participants sampled information (i.e., results of medical tests such as 7 out of 20 participants had

negative side effects) until they felt ready to make a judgment about the true underlying mean. In Experiments 1-3, the number of patients involved in each medical test was 20. In Experiments 4-5, we used 35 patients to investigate whether the pattern of effects would be observed with different stimuli values. To construct these sequences, a list of 51 fictional reports (i.e., stimuli) were created for each trial. The number of reports available in a given trial was chosen to be 51 somewhat arbitrarily, as it would be a large enough number of reports that participants would most likely cease their information search before exhausting the reports (which is critical to investigating information sampling behavior). Participants could sample as many as up to 51 fictional reports in each trial before they made an estimation (Figure 1C). This self-paced task allowed participants to decide how many stimuli they viewed on each trial.

The key within-subject manipulation was the presence of invalid reports (i.e., outlier test results) on some trials. In each experiment, there were 52 trials consisting of three main types: 16 Control trials without outliers, 16 Test trials with outliers, and 20 Catch trials.

In Experiments 1-3, the 16 Control trials were constructed by sampling from a Gaussian distribution with mean of 8 and a standard deviation of 2 (i.e., $N(8, 2)$). In Experiments 4-5, the control trials were sampled from $N(20, 2)$ to investigate whether the pattern of effects would be observed with different numerical values. No reports indicating a value further than 2 SD from the mean were included in the control trial lists. For each trial, we sampled 51 values from its corresponding Gaussian distribution and rounded these values to their closest integers to construct 51 possible reports for participants to view.

The Test trials were designed to investigate the effect of invalid reports on information seeking and one's final estimates. There were two sets of 8 test trials (one set with low outliers and one with high outliers) created by following the same procedure for generating the control trials, and then randomly inserting an outlier as report three, four, or five in the sequence. Placing the outlier early in the sequence of reports ensured that participants were likely to see it when it appeared. Low outliers were 1, 2, and 3 in Experiments 1-3, and 13, 14, and 15 in Experiments 4-5. High outliers were 13, 14, and 15 in Experiments 1-3, and 25, 26, and 27 in Experiments 4-5.

Figure 1

This figure depicts an example of the different screens participants viewed on each trial in Experiment 1. Panel A: An example of the orienting story participants read at the start of each trial. The medical condition and fictional drug names were determined randomly for each trial. Panel B: An example of the stimuli participants viewed on each trial. Each stimulus represents one fictional researcher's report, and participants sampled reports until they felt comfortable estimating the underlying true prevalence rate. Panel C: The response screen where participants typed their estimates after terminating their information search. Panel D: The outlier detection response screen. One group in Experiment 1 and all participants in the following experiments indicated how likely it was that they encountered invalid information at the end of each trial.

The presence of these outliers should not affect people's estimation of the underlying true mean of the distributions if ignored. In other words, the mean value of the test trials was equivalent to that of the control trials, after ignoring the invalid information present in outliers.

The Catch trials were designed to ensure that there was sufficient variety in the stimuli. In each experiment, there were 20 catch trials made up of 4 sets of 5 catch trials. In Experiments 1-3, the 4 sets were sampled from normal distributions with means of either 5 or 11 and SDs of either 1 or 2. In Experiments 4-5, the 4 sets were from normal distributions with means of either 17 or 23 and SDs of either 1 or 2. These additional trials served as a distraction to help prevent participants from learning the structure underlying the control and test stimuli.

In addition to the within-subject manipulation of trial types, each experiment employed a set of between-subject manipulations (Table 1). In Experiment 1, we investigated how Response Types interacted with the effect of Trial Types. One group of participants provided only their estimates of medical cases (Estimate-Only; Figure 1C). The other group provided these estimates, but also indicated the likelihood of observing a outlier in their sampled reports using a 5-level Likert scale on each trial (Estimate-Detect; Figure 1D).

Table 1

Summary of experimental design. Bold text denotes the key manipulations in each experiment.

Experiment	Within-Subject Variables	Between-Subject Variables			
	Trial Type	Response Type	Warning Cue	Response Order	Task Scenario
Exp 1	Control	Estimate-Only	No Cue	Estimate First	Side Effects
	Low-Outlier	Estimate-Detect			
	High-Outlier				
Exp 2	Control	Estimate-Detect	70% Cue (trial-level)	Estimate First	Side Effects
	Low-Outlier		100% Cue (trial-level)		
	High-Outlier				
Exp 3	Control	Estimate-Detect	No Cue	Estimate First	Side Effects
	Low-Outlier		70% Cue (trial-level)	Detect First	
	High-Outlier		100% Cue (trial-level)		
Exp 4	Control	Estimate-Detect	No Cue	Estimate First	Side Effects
	Low-Outlier				Health Improvements
	High-Outlier				
Exp 5	Control	Estimate-Detect	No Cue	Estimate First	Side Effects
	Low-Outlier		100% Cue (report-level)		Health Improvements
	High-Outlier				

All experiments also involved catch trials that are not listed on the table.

In Experiments 2 and 3, we used visual warning cues to alert participants of upcoming invalid information on trials. The warning cue was a red and white-colored, triangular warning sign with an exclamation point inside it (see Figure 2A). This design and color scheme was chosen to ensure participants did not miss the cue. Participants were informed about the nature of the cues at the start of experiment. The warning cue appeared in the space above the story that started each trial (Figure 2A), and disappeared when the story was no longer on the screen. There were two levels of reliability for these trial-level warning cues: 70% and 100% reliable. In the 100% reliable cue group, the warning cue showed up on every test trial (i.e., the trials with outliers). In the 70% reliable group, the cue had a 70% chance of appearing on outlier trials and a 30% chance of appearing on non-outlier (i.e., control and catch) trials. In Experiment 2, we

assigned participants randomly to 70% and 100% Cue groups. In Experiment 3, we included a third group which did not see any warning cues to serve as control (No-Cue) in addition to the two Cue groups.

Figure 2

Panel A: An example of the orienting story that participants read in Experiments 2 and 3. The cue at the top of the screen warned that at least one outlier would appear in the reports for that trial. For some participants, this cue was 100% reliable. For other participants, the cue indicated there was a 70% chance that they would see an outlier in the upcoming reports. Participants were randomly assigned to a cue type before beginning the first trial. Panel B: The outlier cue employed in Experiment 5. In this experiment, one group of participants were shown the cue above any report that had been fabricated. The other group was not shown any cue during the task.

In Experiment 3, we also examined the effect of response order on participants' ability to ignore invalid test results. Half of participants were required to report an estimate before reporting the likelihood of outlier detection (i.e., the Estimate-Detect group); whereas the other half of participants were asked to first report the likelihood of detecting an outlier and then submitted their numerical estimates (i.e., the Detect-Estimate group).

In Experiment 4, we examined whether the task scenario affected participants' ability to ignore outliers. Half of participants were required to estimate numerical evidence in the context of negative side effects (i.e., the Side-Effects group), as in Experiments 1-3. The other half of participants estimated the numerical evidence in the context of positive effects. That is, they were required to estimate the proportion of patients that would show improved health for each medication (i.e., the Health-Improvements group). For the Health-Improvements group, any language referencing side effects was modified to describe health improvements. This included

the instructions, the short story at the start of each trial, and the response screens. Only the language was manipulated; the program that governed the presentation of numerical stimuli was identical for both groups.

In Experiment 5, we investigated the interaction between the task scenario and warning cues on the effect of invalid test results on numerical estimates. We employed a stronger report-level cuing manipulation than the trial-level cues employed in Experiments 2 and 3. In this experiment, the report-level cue appeared above the invalid report (Figure 2B), and disappeared when the report left the screen. In other words, rather than warning participants about the possible appearance of invalid information in upcoming reports as in Experiments 2 and 3, the warning cues in Experiment 5 were 100% accurate and flagged each invalid report for participants.

Procedures

In each experiment, participants were randomly assigned to one of the between-subject groups before the experimental session began. There were two between-subject groups in Experiments 1, 2, and 4, six between-subject groups in Experiment 3, and four between-subject groups in Experiment 5 (see Table 1).

Participants were instructed that they would be viewing a number of fictional scenarios in which a medical lab had developed a new drug to treat an ailment. In Experiments 1-3, the lab's researchers were each said to have administered the drug to 20 different patients to see how many developed negative side effects. In Experiments 4-5, the lab's researchers were each said to have administered the drug to 35 different patients to see how many developed negative side effects (Side-Effects group) or how many showed positive health improvements (Health-Improvements group). Participants were then shown an example of the reports they would be viewing. Participants were informed that their task was to view the researchers' reports to determine the true underlying prevalence rate of the effect for each medication.

While reading the instructions for the task, participants were also made aware of two key points. First, participants were reminded that it is normal for lab members' results to differ from

one another since they are different people testing different samples of patients. They were also informed that they may encounter reports from lab members who misreported their results. Participants were told that misreported results would be “much higher or much lower than the other reports for a given medication.” Participants were specifically instructed to ignore any report that they believed was misreported in this way.

For participants assigned to no-cue groups, after reading the instructions, a comprehension question assessed whether the participant understood the instructions, and how to determine which reports were invalid. For participants assigned to warning-cue groups, they were also informed about the nature of the cues, including the cue’s reliability. A comprehension check was added to ensure participants understood the meaning of the particular warning cue. Any participant that did not pass these checks were removed from the experiment and were not permitted to begin the task.

At the beginning of each trial, participants read a brief story to orient them to the scenario. For example, “Imagine that a lab has developed a pill named Corfenib for treatment of sinus infections. While testing the prevalence of negative side effects, different lab members reported the following results:”. Imaginary medication names were created for this experiment to ensure participants could not utilize existing information about real medications to inform their estimates. The medication and condition names were randomly selected each trial, as the specific names were not of interest to the current hypotheses.

After reading the imaginary scenario, the participant was shown the first fictional researcher’s report. Each report was presented as a number such as “8 out of 20”. Each report remained on the screen for two seconds, after which participants pressed a key to either see another report, or provide their estimate. After seeing the first report, the participant was then shown a screen that allowed them to press the right arrow on the keyboard to see another report, or the spacebar to progress to a screen where they could indicate their estimate and then move on to the next trial. In this way, participants could sample as many reports per trial as they wanted until they felt comfortable estimating the true side effect prevalence rate (or health improvements

prevalence rate in Experiments 4 and 5) of the medication.

When participants were ready to indicate their estimate, they pressed the spacebar and typed in their answer on a response screen. Participants in Estimate-Detect groups were also asked to rate how likely it was that they encountered an invalid report on that trial on a Likert scale (1 = “Not very likely” to 5 = “Very likely”). Participants repeated this sequential sampling of the reports, choosing their stopping point, and estimating the true prevalence rate of negative side effects for each trial. In total, each participant completed 52 trials, within each of which they could view up to 51 reports. Participants were given the option of taking a short break after completing half the trials.

Results

All analyses described in this paper were conducted using jamovi computer software (Jamovi project, 2021; R Core Team, 2020). All mixed effects models were fit using the GAMLj jamovi module (Gallucci, M., 2019; Ripley, B., Venables, W., Bates, D. M., Hornik, K., Gebhardt, A., & Firth, D., 2018). In addition to reporting estimated coefficients for each regression analysis in the main paper, we report the statistical results of each main effect and interaction term from the omnibus likelihood ratio test in the Supplementary Material, which compares the full model against restricted models where one effect of interest is excluded. In addition, for each linear mixed model, we report the conditional R^2 (P. C. Johnson, 2014), which reflects the variance explained by the mixed model over the total variance of the dependent variable. For each ordinal logistic regression model, we reported McFadden’s pseudo R^2 (R^2_{McF} ; McFadden et al., 1973) to reflect the variance explained by the model relative to the total variance of the dependent variable. Note that the Catch trials were not included in these analyses since they served only as distractors and were not of interest to the hypotheses of this experiment. Additionally, the results of paired t-tests (Table 1 in the Supplementary Material) suggest that the accuracy estimated from the check trials was similar to those estimated from the control trials.

All data and analyses have been made publicly available via the Open Science Framework

and can be accessed at https://osf.io/9ybnx/?view_only=10c1c405a3884898b6806e919a0b3a98.

Can people detect invalid information?

Number of Reports Viewed

As shown in Table 2, the observed (i.e., true) mean of the valid reports viewed by participants was similar across different trial types and consistent with experimental manipulations. In addition, we observed that the number of reports viewed by participants was slightly higher in test trials (i.e., low-outlier and high-outlier trials) than in control trials (i.e., trials that included only valid reports) in all five experiments.

Table 2

Descriptive statistics for Experiments 1-5. Control trials only contained valid reports. Low-outlier (high-outlier) trials included one low (high) invalid report.

Experiment		Observed Mean of Valid Reports			Number of Reports Viewed			Outlier Likelihood Ratings			Estimation Error Scores		
		Control	Low Outlier	High Outlier	Control	Low Outlier	High Outlier	Control	Low Outlier	High Outlier	Control	Low Outlier	High Outlier
Exp1	M	8.00	8.08	8.03	9.07	9.93	9.75	3.05	4.12	3.99	0.03	-0.23	0.44
	SD	0.18	0.33	0.31	4.54	4.6	4.59	0.81	0.78	0.76	0.38	0.59	0.62
Exp2	M	8.01	8.03	8.06	9.49	9.45	10.20	2.86	4.29	4.10	-0.03	-0.24	0.36
	SD	0.18	0.28	0.27	4.12	3.64	4.08	0.84	0.62	0.69	0.41	0.61	0.59
Exp3	M	8.01	7.99	7.99	9.66	10.00	10.30	2.86	4.18	3.90	-0.03	-0.13	0.40
	SD	0.18	0.35	0.32	5.05	5.3	5.19	0.75	0.73	0.78	0.42	0.91	0.85
Exp4	M	20.00	20.00	20.00	9.58	10.10	9.97	2.53	3.60	3.48	-0.16	-0.52	0.37
	SD	0.20	0.27	0.34	4.68	4.93	4.71	0.71	0.81	0.86	0.43	0.88	1.06
Exp5	M	20.00	20.00	20.00	10.30	10.70	10.70	2.36	4.01	3.98	-0.22	-0.48	-0.08
	SD	0.19	0.30	0.28	6.48	6.68	6.28	0.89	0.88	0.88	0.90	1.05	1.16

To examine if the presence of outliers affected the number of reports participants sampled, we fit a linear mixed effects regression model to predict the number of reports participants viewed based on Trial Type (Control, Low Outlier, High Outlier) and corresponding between-subject manipulations (Table 1) in each experiment, as well as the interaction terms. We also included by-subject random intercepts¹. The results of the estimated coefficients are summarized in Table 3, and the results of omnibus tests are summarized in the Supplementary Material (SM: Table 2).

Consistent with our observations, the model revealed a significant increase in the number of reports participants viewed when they encountered either low outliers or high outliers in test trials as compared to when they encounter no invalid reports in control trials. There was no

¹ By-subject random slopes were not included due to model convergence issues.

significant effect of warning cues on the number of reports participants sampled in Experiments 2 and 3. In Experiment 5, when warning cues flagged the exact reports containing invalid information, participants viewed significantly fewer reports in test trials ($M = 10.4$) than in control trials ($M = 10.73$). Other factors, including Response Type (Estimate-Only vs. Estimate-Detect) and Task Scenario (Side Effects vs. Health Improvements), did not significantly impact the number of reports participants viewed.

Outlier Likelihood Ratings

We also observed that the reported likelihood of seeing an outlier was higher in test trials than in control trials at the mean level (Table 2), suggesting that participants might be able to detect outliers when they were present. An ordinal logistic regression of outlier likelihood ratings was conducted for each experiment to examine this hypothesis. The model predicted the outlier likelihood ratings in the Estimate-Detect condition using Trial Type and the key between-subject manipulations (Table 1) in each experiment as predictors. The model also included interaction terms and by-subject random intercepts. The results of the estimated coefficients are summarized in Table 4, and the respective results of omnibus tests are summarized in the Supplementary Material (SM: Table 3).

In all five experiments, the model supported that likelihood ratings were significantly predicted by trial type (see SM: Table 3). When participants encountered an outlier higher than the underlying mean (i.e., in High-Outlier test trials), they reported higher likelihoods of detecting the outlier than when they did not encounter any outlier (i.e., in control trials). Similarly, when they encountered an outlier lower than the underlying mean (i.e., in Low-Outlier test trials), participants also reported higher likelihood ratings compared to control trials (Table 4). This indicates that participants were able to detect the presence of outliers in test trials.

The interaction between the presence of warning cues and trial type also significantly predicted participants' outlier likelihood ratings (SM: Table 3). In Experiment 2, the increase in likelihood ratings between test trials (i.e., Low-outlier and High-Outlier trials) and control trials

was significantly larger in the 100% cue group than in the 70% cue group, reflecting that participants were more certain about encountering an outlier in trials when the cues were more reliable. Regression results for Experiment 3 also suggested a significant increase in likelihood ratings when participants were provided with cues of higher reliability, replicating the findings of Experiment 2. In Experiment 5, when warning cues flagged the exact report containing invalid information, the increase in likelihood ratings between test and control trials was even larger (Table 4).

No consistent evidence was found to support the effects of other factors, including response orders (Experiment 3) and task scenarios (Experiments 4 and 5) on changes in likelihood ratings across different trial types.

Conclusions

Our results showed that when encountering outliers in trials, participants sampled more reports before making an estimate and reported a higher likelihood of seeing an outlier. In addition, the presence of warning cues facilitated participants' outlier detection. These findings suggest that participants were able to detect outliers.

Can people ignore invalid information?

To assess participants' performance at estimating prevalence rates of side effects (or health improvements in Experiments 4 and 5), we converted the estimates to an error score by subtracting the true mean of all valid reports a participant viewed on a trial, that is excluding outlier values, from participants' reported mean (Figure 3, 4, and 5). This error score reflects how far a reported estimate was from the ground truth. A positive error value indicates an overestimation of ground truth, and a negative error value indicates an underestimation of ground truth.

To assess the effect of outlier presence on participants' estimates, we were especially interested in examining the changes in error scores when comparing between test trials and control trials. Deviation in error scores between test trials (i.e., including either low outliers or

high outliers) and control trials reflects the biasing effects of outliers on numerical estimation. Note that with this approach we define an estimation bias due to an outlier through the comparison of outlier trials to control trials and not ground truth.

Biasing Effects in Each Experiment

As demonstrated in Table 2, the observed average error scores differed across trial types in all five experiments. Compared to the error scores observed in control trials, the error scores in High-Outlier trials were higher, and the error scores in Low-Outlier trials were lower. A linear mixed effects regression model was constructed to statistically test this observation. For each experiment, the model predicted error scores as a linear function of predictors including Trial Type (No outlier, Low Outlier, High Outlier) and the corresponding between-subject groups in that experiment (Table 1). The model also included interaction terms and by-subject random intercepts. The results of the estimated coefficients are summarized in Table 5, and the respective results of omnibus tests are summarized in the Supplementary Material (SM: Table 4).

In all five experiments, the model indicated a significant effect of Trial Type on error scores (SM: Table 4). Moreover, the direction of change in error scores between test and control trials was positively associated with the type of outliers present in the test trials (Table 5). When participants encountered an outlier lower than the underlying mean in Low-Outlier trials, their error scores were significantly lower than in control trials. In contrast, when participants encountered an outlier higher than the underlying mean in High-Outlier trials, their error scores were significantly higher than in control trials. These results support the significant biasing effects of the presence of outliers on participants' numerical estimation.

The model in Experiment 1 also revealed a significant main effect of Response Type on error scores. When participants were explicitly required to report outlier likelihood ratings along with their estimates (i.e., Estimate-Detection group), their error scores were lower ($M = -0.061$) compared to when they only reported estimates (i.e., Estimate-Only group, $M = 0.147$). However, the interaction between Response Type and Trial Type was not significant, suggesting that

Figure 3

Mean error in each of the experimental conditions of Experiment 1. For each trial, error was computed by first determining the mean of all report values a participant saw in that trial, minus any outliers. This mean was then subtracted from the estimate the participant provided. Thus error is the difference between the participant's estimate and the mean of all valid reports they viewed on a trial, with positive values indicating an overestimation and negative values indicating underestimation. Overall, estimates were biased in the direction of outliers when they were present. Error bars represent the $\pm$ standard errors of the means.

Response Type might not markedly affect the degree to which outliers influence participants' estimates (Figure 3).

For Experiments 2 and 3, the model revealed no significant main effect of the trial-level Warning Cue nor its interaction with Trial Type on error scores (Figure 4; also see Table 4 in SM). In Experiment 2, the results suggested that the reliability of warning cues did not significantly affect participants' error scores, irrespective of the type of trial. In Experiment 3, the model results additionally suggested that the presence or absence of these cues did not significantly affect participants' error scores. These results indicate that cues of upcoming invalid reports did not effectively reduce the influence of outliers on participants' numerical estimates.

Figure 4

Mean errors for Experiments 2 and 3. Panel A: Participants' estimates in Experiment 2 were biased in the direction of an outlier when it was present. This bias was present in both the 70%- and 100%-Cue groups. Panel B: Results from a pre-registered replication of Experiment 2 (i.e., Experiment 3), including a group which was shown no warning cues. Results from Experiment 2 were replicated, and the No-Cue group's estimates were also biased in the direction of outliers. Note that the presence of warning cues in Experiments 2 and 3 was at the trial level. Error bars represent the $\pm$ standard error of the means.

On the other hand, in Experiment 5, when 100% reliable cues for misinformation were presented simultaneously with invalid reports (i.e., report-level cues), error scores on test trials were closer to those in control trials (Figure 5). The model results confirmed a significant interaction between Warning Cue and Trial Type in Experiment 5 (SM: Table 4). In the No-Cue group, participants' error scores were significantly biased in the direction of the outliers they encountered, replicating what we observed in previous experiments. In contrast, when participants received clear warnings on outliers they encountered (i.e., report-level 100% Cue group), the deviation in error scores between test trials (either Low-Outlier trials or High-Outlier trials) and control trials was significantly reduced (Table 5). These results indicate that instantaneous alerts for the presence of misinformation might reduce the biasing effects of that information. In the 100% cue group, the results of paired t-tests further revealed that the mean

difference in error scores between Low-Outlier and control trials ($M = -0.193$) was significant, $t(121) = -4.20, p < 0.001$, Cohen's $d = -0.380$, but not between High-Outlier and control trials ($M = 0.077$), $t(121) = 1.180, p = 0.239$, Cohen's $d = 0.107$. These results suggest that providing clear and reliable alerts for invalid information might reduce biasing effects, at least in some cases (e.g., when high outliers are present).

The model results for Experiment 5 also revealed a significant main effect of Task Scenario on error scores, but the interaction between Task Scenario and Trial Type was not significant (SM: Table 4). On average, error scores were higher in the Side-Effects group ($M = -0.125 \pm 1.30$) than in the Health-Improvements group ($M = -0.380 \pm 1.64$). This indicates that when evaluating numerical evidence, participants tended to be more conservative in the context of the positive scenario (e.g., Health-Improvements) than in the context of the negative scenario (e.g., Side-Effects). We observed a similar pattern of error scores across different task scenarios in Experiment 4, but these patterns were not statistically significant. Finally, no evidence was found to support the effect of Response Order (i.e., Estimate-Detect vs. Detect-Estimate; also see results in Experiment 3).

Biasing Effects Across Experiments

The results of within-experiment analyses consistently revealed the biasing effects of outliers on numerical estimation in all experiments, despite different between-subject manipulations across these experiments. However, variations in certain between-subject factors, such as warning cues and task scenarios, appeared to affect estimation accuracy and the biasing effects of outliers in some experiments.

Since the design of these experiments were highly similar, we conducted a mixed effects linear regression analysis on data combined from all five experiments to assess the influence of these between-subject factors on the biasing effects of outliers. The mixed linear regression model was constructed to include Trial Type, Warning Cue, Task Scenario, as well as the interaction terms, as predictors for error scores. We note that only the Estimate-Detection group

from Experiment 1 was included, as this response type was used in all of the other experiments (Table 1). The results of the estimated coefficients are summarized in Table 6 ($n = 764$), and the respective results of omnibus tests are summarized in the Supplementary Material (SM: Table 5).

In line with our findings from the within-experiment analyses, the results using the combined data confirmed the biasing effects of outliers on participants' estimates, where error scores were positively associated with the type of outliers participants encountered. Moreover, the model confirmed the significant interaction between Warning Cue and Trial Type (SM: Table 5), showing that 100% reliable warning cues at the report level significantly reduced these biasing effects (Table 6). Additionally, the model confirmed a significant effect of Task Scenario on estimates, suggesting that participants underestimated prevalence rates to a greater extent in the positive scenario (i.e., Health Improvements) than in the negative scenario (i.e., Side Effects).

Biasing Effects When Outliers Were Detected

Finally, we examined whether participants were able to ignore invalid information when they were highly certain about its presence in a trial. To this end, we focused on a subset of trials in each experiment, where an outlier was shown (i.e., test trials) and participants were maximally certain that they encountered the outlier (i.e., likelihood ratings = 5). Table 7 shows that error scores in Low-outlier trials still deviated negatively from the ground truth, while in High-Outlier trials error scores deviated positively.

To validate these observations, we fit a linear mixed effects regression model to predict error scores when participants gave likelihood ratings of 5, for each experiment separately. The predictors included Trial Type (i.e., high and low outlier trials) and the key between-subject manipulations (Table 1), as well as interaction terms and by-subject random intercepts. For these analyses, we only included high and low outlier trials and not control trials as our objective is to see if estimates differ when outliers are present and detected. The results of the estimated coefficients are summarized in Table 8, and the respective results of omnibus tests are summarized in the Supplementary Material (SM: Table 6).

The models revealed a significant effect of Trial Type on error scores in each experiment (SM: Table 6). The presence of high outliers generally resulted in higher estimates than the presence of low outliers (Table 8). This indicates that participants were still biased by outliers, despite being aware of their presence. The presence of report-level warning cues reduced the influence of high outliers (Experiment 5). In addition, the model revealed a main effect of task scenario on estimates (see Experiments 4 and 5 in Table 6 in SM), suggesting that error scores were significantly lower in the Health-Improvements scenario than in Side-Effects scenario. No consistent evidence was found to support significant interactions between Trial Type and the between-subject factors.

Conclusions

Our results show that participants were unable to ignore outliers when evaluating numerical evidence, even when they detected outliers. The use of reliable warning cues regarding the presence of outliers at the report-level reduced the biasing effects, but the reduction was not sufficient to fully eliminate these effects.

General Discussion

Across five experiments, participants' numerical estimates were biased in the direction of outliers. Moreover, the biasing effects persisted across two different ranges of stimuli values. Overall, participants' information seeking slightly increased when they saw an outlier, but this did not affect the bias. These findings suggest that people may not be able to fully ignore invalid numerical information once they have encountered it.

This is consistent with the continued influence effect (CIE) and the illusory truth literature, showing that it is difficult for individuals to ignore or disbelieve invalid information that they have encountered (Hasher et al., 1977; H. M. Johnson and Seifert, 1994; Wistrich et al., 2004). Our work expanded the traditional CIE paradigm by having participants decide for themselves which pieces of information (numerical reports) were invalid. Even when participants actively detected invalid information and made their final response shortly thereafter (approx.

10-25sec), they were still influenced by it.

In addition, consistent with previous research on misinformation, it appears that numerical misinformation is difficult to ignore even with interventions. In the current study, receiving warnings for upcoming invalid information increased participants' ability to detect outliers, but did not eliminate the influence of invalid reports on their estimates. Furthermore, when outliers were clearly flagged, the biasing effects were reduced but still existed. These findings are also consistent with existing evidence in Brashier et al. (2021), suggesting that timing of belief revision is critical.

One explanation of these findings could relate to memory retrieval errors. According to exemplar theories of memory, individuals in our task store instances in memory representing each stimulus they encounter. When participants are asked to make judgments, they sample values from memory and combine them to generate an estimate (André et al., 2021). In this account, invalid information would be encoded, but would have a potentially lower (but non-zero) activation weight than other values encountered on a trial. When participants generate their estimates, the outlier values could be sampled and incorporated into the final judgment, producing estimates biased towards outliers.

It should be noted that the current experiments did not ascertain the exact moment participants detected an outlier. Participants only indicated the likelihood rating of outlier detection at the end of the trial. Thus, we could not determine in which trials participants detected the outlier upon seeing it, which might determine whether they encoded the value. Similarly, we cannot say in which trials participants encoded the outlier as valid, and later determined it to be invalid after seeing more values in the set. Future work could probe participants' detection after each stimulus to investigate this issue.

Another possible explanation of our findings is that our outlier values were not surprising enough. Filipowicz et al. (2018) found that participants who would usually show increased belief updating to surprising stimuli instead showed no updating when new information was extremely surprising. They suggested that their participants may have judged information to be too

surprising to be legitimate and ignored it. Our experiments used outliers that were 2.5 to 3.5 SD from the mean in order to mimic deceptive statistics that one may encounter in their daily life. Future work could employ more extreme outliers to investigate the boundary conditions for incorporating outliers into one's beliefs.

It should also be noted that our experiments utilized a narrow range of values with a small variance, presented in a format that might not reflect data seen in the real world. For example, people often encounter data in the form of percentages, such as in demographic information cited by politicians (Prévost and Beaud, 2015), or they encounter large numbers such as mortality predictions for the COVID-19 pandemic (Allyn, 2020). They may also encounter numerical data that are not normally distributed or with large volatility, such as stock prices. Future work could employ different numerical formats, magnitudes, and variances to investigate whether the biasing effects of invalid information would generalize to these numerical variations.

The current experiments add to the literature concerning how we deal with misinformation. Much of the current CIE and illusory truth work examines how we share and spread misinformation, and how we may curtail that spread (Chan et al., 2017; Lewandowsky et al., 2012; Pennycook et al., 2021). Few studies have examined how people handle numerical misinformation (but see Stubenvoll and Matthes, 2021), and our work contributes to the literature by illuminating that people have difficulty ignoring misinformation in the form of invalid numerical reports, even when they do not believe the report is legitimate. The influence of bad numerical information on our beliefs suggests a vulnerability to sloppy data reporting or outright false statistical information in individuals' voting, consumer, and health decisions.

Author Contributions

A. Ramsey and J. Trueblood developed the study concept and protocol designs. A. Ramsey created the stimulus presentations. A. Ramsey and Y. Liu analyzed the data. All authors interpreted the data. A. Ramsey drafted the manuscript, and all authors provided critical revisions. All authors approved the final version of the manuscript for submission.

Open Practices

All data and analyses have been made publicly available via the Open Science Framework and can be accessed at https://osf.io/9ybnx/?view_only=10c1c405a3884898b6806e919a0b3a98. The design and analysis plans for studies 3, 4, and 5 were pre-registered at AsPredicted and can be accessed at https://aspredicted.org/NNB_QGM, https://aspredicted.org/OAB_ATF, and https://aspredicted.org/MTN_FZN.

Acknowledgements

This work was also supported by NSF grant 1846764.

Table 3

Estimated coefficients from the mixed effects linear regression models for number of reports participants viewed as predicted by key manipulations in each experiment, after accounting for by-subject random intercepts.

Experiment	Term	β	SE	t	df	p-value	95% CI
Exp 1	Intercept: Estimate-only, Control	9.598	0.500	19.215	84	< 0.001	[8.619, 10.577]
	Low Outlier	0.810	0.210	3.864	2475	< 0.001	[0.399, 1.220]
	High Outlier	0.646	0.209	3.094	2475	0.002	[0.237, 1.056]
	Estimate-Detect	0.143	0.999	0.143	84	0.887	[-1.815, 2.101]
	Low Outlier * Estimate-Detect	-0.597	0.419	-1.423	2475	0.155	[-1.418, 0.225]
	High Outlier * Estimate-Detect	-0.318	0.418	-0.761	2475	0.447	[-1.137, 0.501]
	Conditional R^2	0.531					
Exp 2	Intercept: 70% Cue (trial-level), Control	9.735	0.424	22.948	82	< 0.001	[8.904, 10.567]
	Low Outlier	-0.055	0.180	-0.305	2490	0.760	[-0.407, 0.297]
	High Outlier	0.641	0.180	3.564	2490	< 0.001	[0.288, 0.993]
	100 % Cue	-1.162	0.848	-1.370	81.600	0.174	[-2.825, 0.501]
	Low Outlier * 100 % Cue (trial-level)	0.553	0.359	1.539	2490	0.124	[-0.151, 1.257]
	High Outlier * 100 % Cue (trial-level)	-0.253	0.360	-0.704	2490	0.482	[-0.958, 0.452]
	Conditional R^2	0.521					
Exp 3	Intercept: No Cue, Control, Estimate-First	9.950	0.288	34.600	310	< 0.001	[9.386, 9.386]
	Low Outlier	0.317	0.100	3.178	9190	0.001	[0.122, 0.122]
	High Outlier	0.575	0.100	5.768	9189	< 0.001	[0.380, 0.380]
	70 % Cue (trial-level)	-0.507	0.700	-0.725	310	0.469	[-1.878, -1.878]
	100 % Cue (trial-level)	-0.008	0.707	-0.011	310	0.991	[-1.394, -1.394]
	Detect-First	0.006	0.575	0.010	310	0.992	[-1.122, -1.122]
	Low Outlier * 70% Cue (trial-level)	0.335	0.244	1.377	9190	0.169	[-0.142, -0.142]
	High Outlier * 70% Cue (trial-level)	-0.199	0.243	-0.816	9189	0.414	[-0.675, -0.675]
	Low Outlier * 100% Cue (trial-level)	0.291	0.246	1.186	9190	0.236	[-0.190, -0.190]
	High Outlier * 100% Cue (trial-level)	0.280	0.245	1.141	9189	0.254	[-0.201, -0.201]
	Low Outlier * Detect-First	0.010	0.200	0.051	9190	0.959	[-0.381, -0.381]
	High Outlier * Detect-First	-0.078	0.199	-0.391	9189	0.696	[-0.469, -0.469]
	70 % Cue (trial-level) * Detect-First	0.396	1.399	0.283	310	0.778	[-2.347, -2.347]
	100 % Cue * Detect-First	1.155	1.414	0.817	310	0.415	[-1.617, -1.617]
	Low Outlier * 70% Cue (trial-level) * Detect-First	-0.069	0.487	-0.142	9190	0.887	[-1.023, -1.023]
	High Outlier * 70% Cue (trial-level) * Detect-First	0.456	0.487	0.938	9189	0.348	[-0.497, -0.497]
	Low Outlier * 100% Cue (trial-level) * Detect-First	0.704	0.491	1.433	9190	0.152	[-0.259, -0.259]
	High Outlier * 100% Cue (trial-level) * Detect-First	1.191	0.491	2.427	9189	0.015	[0.229, 0.229]
	Conditional R^2	0.625					
Exp 4	Intercept: Control, Side-Effects	9.911	0.465	21.310	102	< 0.001	[8.999, 10.822]
	Low Outlier	0.582	0.163	3.569	3045	< 0.001	[0.262, 0.901]
	High Outlier	0.386	0.164	2.359	3045	0.018	[0.065, 0.707]
	Health Improvement	0.313	0.930	0.336	102	0.737	[-1.510, 2.136]
	Low Outlier * Health-Improvements	0.442	0.326	1.357	3045	0.175	[-0.197, 1.081]
	High Outlier * Health-Improvements	0.315	0.327	0.961	3045	0.337	[-0.327, 0.956]
	Conditional R^2	0.614					
Exp 5	Intercept: Control, No-Cue, Side-Effects	10.501	0.419	25.074	227	< 0.001	[9.681, 11.322]
	Low Outlier	0.448	0.125	3.578	6796	< 0.001	[0.202, 0.693]
	High Outlier	0.354	0.125	2.842	6796	0.004	[0.110, 0.599]
	100 % Cue (report-level)	-0.320	0.838	-0.382	227	0.703	[-1.962, 1.322]
	Health Improvement	0.156	0.838	0.186	227	0.853	[-1.486, 1.797]
	Low Outlier * 100% Cue (report-level)	-0.708	0.250	-2.826	6796	0.005	[-1.198, -0.217]
	High Outlier * 100% Cue (report-level)	-0.569	0.249	-2.283	6796	0.022	[-1.058, -0.080]
	Low Outlier * Health-Improvements	-0.073	0.250	-0.291	6796	0.771	[-0.563, 0.418]
	High Outlier * Health-Improvements	-0.716	0.249	-2.871	6796	0.004	[-1.205, -0.227]
	100 % Cue (report-level) * Health-Improvements	2.959	1.675	1.766	227	0.079	[-0.324, 6.243]
	Low Outlier * 100% Cue (report-level) * Health-Improvements	-0.765	0.501	-1.529	6796	0.126	[-1.746, 0.216]
	High Outlier * 100% Cue (report-level) * Health-Improvements	-0.012	0.499	-0.024	6796	0.981	[-0.989, 0.965]
	Conditional R^2	0.693					

Table 4

Estimated coefficients from the ordinal logistic regression models for outlier likelihood ratings as a function of key manipulations in each experiment.

Experiment	Term	Log of Odds Ratio	SE	Z	p-value	95% CI
Exp1	Low Outlier - Control	1.590	0.148	10.760	< 0.001	[1.300, 1.880]
	High Outlier - Control	1.290	0.143	9.050	< 0.001	[1.010, 1.570]
	R^2_{McF}	0.05				
Exp2	Low Outlier - Control	1.897	0.136	13.980	< 0.001	[1.632, 2.164]
	High Outlier - Control	1.405	0.130	10.800	< 0.001	[1.151, 1.660]
	100% Cue – 70% Cue (trial-level)	-0.272	0.098	-2.780	0.005	[-0.464, -0.081]
	Low Outlier * 100% Cue (trial-level)	0.876	0.183	4.790	< 0.001	[0.518, 1.235]
	High Outlier * 100% Cue (trial-level)	1.044	0.180	5.810	< 0.001	[0.692, 1.397]
	R^2_{McF}	0.104				
Exp3	Low Outlier - Control	1.558	0.109	14.337	< 0.001	[1.346, 1.772]
	High Outlier - Control	1.070	0.107	10.038	< 0.001	[0.861, 1.279]
	70% Cue (trial-level) – No Cue	-0.004	0.081	-0.050	0.960	[-0.163, 0.155]
	100 % Cue (trial-level) – No Cue	-0.318	0.087	-3.649	< 0.001	[-0.489, -0.147]
	Detect First – Estimate First	-0.043	0.087	-0.497	0.620	[-0.214, 0.128]
	Low Outlier * 70% Cue (trial-level)	0.224	0.154	1.459	0.145	[-0.077, 0.525]
	High Outlier * 70% Cue (trial-level)	0.379	0.152	2.498	0.012	[0.082, 0.677]
	Low Outlier * 100% Cue (trial-level)	1.581	0.173	9.139	< 0.001	[1.243, 1.921]
	High Outlier * 100% Cue (trial-level)	1.302	0.163	7.973	< 0.001	[0.982, 1.623]
	Low Outlier * Detect-First	0.373	0.164	2.279	0.023	[0.052, 0.695]
	High Outlier * Detect-First	0.069	0.161	0.428	0.669	[-0.247, 0.385]
	70 % Cue (trial-level) * Detect-First	0.180	0.123	1.471	0.141	[-0.060, 0.421]
	100 % Cue (trial-level) * Detect-First	-0.166	0.125	-1.329	0.184	[-0.410, 0.079]
	Low Outlier * 70% Cue (trial-level) * Detect-First	-0.248	0.230	-1.079	0.281	[-0.699, 0.203]
	High Outlier * 70% Cue (trial-level) * Detect-First	-0.101	0.226	-0.447	0.655	[-0.544, 0.342]
	Low Outlier * 100% Cue (trial-level) * Detect-First	-0.228	0.243	-0.939	0.348	[-0.704, 0.247]
	High Outlier * 100% Cue (trial-level) * Detect-First	0.024	0.232	0.103	0.918	[-0.430, 0.478]
	R^2_{McF}	0.093				
Exp4	Low Outlier - Control	1.467	0.109	13.520	< 0.001	[1.255, 1.681]
	High Outlier - Control	1.165	0.108	10.830	< 0.001	[0.955, 1.376]
	Health Improvement - Side Effect	-0.368	0.089	-4.150	< 0.001	[-0.542, -0.194]
	Low Outlier * Health-Improvements	0.283	0.160	1.770	0.077	[-0.031, 0.597]
	High Outlier * Health-Improvements	0.549	0.162	3.400	< 0.001	[0.232, 0.866]
	R^2_{McF}	0.053				
Exp5	Low Outlier - Control	1.514	0.107	14.095	< 0.001	[1.304, 1.725]
	High Outlier - Control	1.249	0.105	11.893	< 0.001	[1.044, 1.456]
	100 % Cue (report-level)	-1.201	0.084	-14.242	< 0.001	[-1.366, -1.036]
	Health Improvement	-0.276	0.083	-3.334	< 0.001	[-0.438, -0.114]
	Low Outlier * 100% Cue (report-level)	1.984	0.160	12.401	< 0.001	[1.671, 2.298]
	High Outlier * 100% Cue (report-level)	2.415	0.159	15.162	< 0.001	[2.103, 2.728]
	Low Outlier * Health-Improvements	-0.128	0.150	-0.852	0.394	[-0.422, 0.166]
	High Outlier * Health-Improvements	0.049	0.149	0.328	0.743	[-0.242, 0.340]
	100 % Cue (report-level) * Health-Improvements	0.563	0.118	4.779	< 0.001	[0.332, 0.794]
	Low Outlier * 100% Cue (report-level) * Health-Improvements	0.074	0.223	0.333	0.739	[-0.363, 0.512]
	High Outlier * 100% Cue (report-level) * Health-Improvements	-0.290	0.222	-1.306	0.191	[-0.725, 0.145]
	R^2_{McF}	0.135				

Table 5

Estimated coefficients from the mixed effects linear regression models for error scores as a function of key manipulations in each experiment, after accounting for by-subject intercepts.

Experiment	Term	β	SE	t	df	p-value	95% CI
Exp 1	Intercept: Control, Estimate-Only	0.055	0.043	1.284	89	0.202	[-0.029, 0.139]
	Low Outlier	-0.229	0.050	-4.578	2483	< 0.001	[-0.328, -0.131]
	High Outlier	0.371	0.050	7.421	2483	< 0.001	[0.273, 0.468]
	Estimate-Detect	-0.233	0.086	-2.717	89	0.008	[-0.401, -0.065]
	Low Outlier * Estimate-Detect	-0.117	0.100	-1.171	2483	0.242	[-0.314, 0.079]
	High Outlier * Estimate-Detect	-0.002	0.100	-0.019	2483	0.985	[-0.198, 0.194]
	Conditional R^2	0.142					
Exp 2	Intercept: Control, 70% Cue (trial-level)	0.025	0.044	0.562	84	0.575	[-0.062, 0.112]
	Low Outlier	-0.190	0.047	-4.057	2494	< 0.001	[-0.281, -0.098]
	High Outlier	0.370	0.047	7.898	2494	< 0.001	[0.278, 0.461]
	100 % Cue (trial-level)	0.108	0.089	1.217	84	0.227	[-0.066, 0.282]
	Low Outlier * 100 % Cue (trial-level)	-0.037	0.094	-0.393	2494	0.694	[-0.220, 0.146]
	High Outlier * 100 % Cue (trial-level)	0.053	0.094	0.568	2494	0.570	[-0.130, 0.237]
	Conditional R^2	0.157					
Exp 3	Intercept: Control, No Cue, Estimate-First	0.080	0.026	3.137	323	0.002	[0.030, 0.130]
	Low Outlier	-0.085	0.027	-3.128	9223	0.002	[-0.138, -0.032]
	High Outlier	0.437	0.027	16.125	9216	< 0.001	[0.383, 0.490]
	70 % Cue (trial-level)	0.067	0.062	1.086	324	0.278	[-0.054, 0.189]
	100 % Cue (trial-level)	0.089	0.063	1.418	323	0.157	[-0.034, 0.212]
	Detect-First	-0.038	0.051	-0.750	323	0.454	[-0.138, 0.062]
	Low Outlier * 70% Cue (trial-level)	-0.046	0.066	-0.698	9227	0.485	[-0.176, 0.083]
	High Outlier * 70% Cue (trial-level)	-0.002	0.066	-0.024	9216	0.981	[-0.131, 0.128]
	Low Outlier * 100% Cue (trial-level)	0.014	0.067	0.216	9226	0.829	[-0.116, 0.145]
	High Outlier * 100% Cue (trial-level)	-0.010	0.067	-0.147	9219	0.883	[-0.140, 0.121]
	Low Outlier * Detect-First	-0.044	0.054	-0.812	9223	0.417	[-0.150, 0.062]
	High Outlier * Detect-First	0.052	0.054	0.957	9216	0.339	[-0.054, 0.158]
	70 % Cue (trial-level) * Detect-First	0.027	0.124	0.219	324	0.827	[-0.216, 0.271]
	100 % Cue (trial-level) * Detect-First	0.254	0.126	2.025	323	0.044	[0.008, 0.500]
	Low Outlier * 70% Cue (trial-level) * Detect-First	-0.012	0.132	-0.092	9227	0.927	[-0.271, 0.247]
	High Outlier * 70% Cue (trial-level) * Detect-First	-0.012	0.132	-0.088	9216	0.930	[-0.271, 0.247]
	Low Outlier * 100% Cue (trial-level) * Detect-First	-0.104	0.133	-0.781	9226	0.435	[-0.366, 0.157]
	High Outlier * 100% Cue (trial-level) * Detect-First	0.216	0.133	1.620	9219	0.105	[-0.045, 0.477]
	Conditional R^2	0.155					
Exp 4	Intercept: Control, Side-Effects	-0.126	0.044	-2.898	104	0.005	[-0.211, -0.041]
	Low Outlier	-0.320	0.051	-6.305	3047	< 0.001	[-0.419, -0.220]
	High Outlier	0.418	0.051	8.197	3050	< 0.001	[0.318, 0.517]
	Health Improvement	-0.064	0.087	-0.736	104	0.463	[-0.234, 0.106]
	Low Outlier * Health-Improvements	-0.123	0.101	-1.215	3047	0.225	[-0.322, 0.076]
	High Outlier * Health-Improvements	-0.161	0.102	-1.580	3050	0.114	[-0.361, 0.039]
	Conditional R^2	0.138					
Exp 5	Intercept: Control, No-Cue, Side-Effects	-0.253	0.061	-4.144	230	< 0.001	[-0.373, -0.134]
	Low Outlier	-0.255	0.035	-7.358	6801	< 0.001	[-0.322, -0.187]
	High Outlier	0.161	0.035	4.678	6801	< 0.001	[0.094, 0.229]
	100 % Cue (report-level)	-0.107	0.122	-0.876	230	0.382	[-0.347, 0.133]
	Health Improvement	-0.264	0.122	-2.160	230	0.032	[-0.504, -0.024]
	Low Outlier * 100% Cue (report-level)	0.141	0.069	2.043	6801	0.041	[0.006, 0.277]
	High Outlier * 100% Cue (report-level)	-0.156	0.069	-2.263	6801	0.024	[-0.291, -0.021]
	Low Outlier * Health-Improvements	0.003	0.069	0.036	6801	0.971	[-0.133, 0.138]
	High Outlier * Health-Improvements	-0.046	0.069	-0.662	6801	0.508	[-0.181, 0.089]
	100% Cue (report-level) * Health-Improvements	-0.071	0.245	-0.288	230	0.773	[-0.550, 0.409]
	Low Outlier * 100% Cue (report-level) * Health-Improvements	0.057	0.138	0.408	6801	0.683	[-0.215, 0.328]
	High Outlier * 100% Cue (report-level) * Health-Improvements	-0.021	0.138	-0.152	6801	0.879	[-0.291, 0.249]
	Conditional R^2	0.384

Figure 5

Mean errors for Experiments 4 and 5. Panel A: Estimates in Experiment 4 were biased towards outliers, indicating participants were not able to fully ignore them. This was true regardless of whether participants were estimating the underlying rates of health improvements or negative side effects. Panel B: Mean errors for Experiment 5 for both cue groups combined. Panel C: Results for the No-Cue group in Experiment 5. Panel D: Results for the Cue-Present group in Experiment 5. Note that the presence of warning cues in Experiment 5 was at the report level. Error bars represent $\pm$ standard error of the means.

Table 6

*Estimated coefficients from the mixed effects linear regression model for error scores combined across experiments: Error Scores = (1|participants) + Trial Type + Warning Cue + Task Scenario + Trial Type*Warning Cue + Trial Type * Task Scenario.*

Term	β	SE	t	df	p-value	95% CI
Intercept: No Cue, Control, Side Effects	-0.128	0.031	-4.107	778	< 0.001	[-0.189, -0.067]
Low Outlier	-0.199	0.024	-8.212	22564	< 0.001	[-0.246, -0.151]
High Outlier	0.291	0.024	12.023	22568	< 0.001	[0.244, 0.338]
70 % Cue (trial-level)	0.093	0.064	1.440	782	0.15	[-0.033, 0.219]
100 % Cue (trial-level)	0.136	0.064	2.125	781	0.034	[0.011, 0.261]
100 % Cue (report-level)	-0.164	0.068	-2.429	782	0.015	[-0.296, -0.032]
Health Improvements	-0.234	0.063	-3.744	780	< 0.001	[-0.356, -0.111]
Low Outlier * 70% Cue (trial-level)	0.081	0.050	1.614	22584	0.107	[-0.017, 0.180]
High Outlier * 70% Cue (trial-level)	0.017	0.050	0.332	22575	0.74	[-0.082, 0.115]
Low Outlier * 100% Cue (trial-level)	0.103	0.050	2.058	22580	0.04	[0.005, 0.201]
High Outlier * 100% Cue (trial-level)	0.030	0.050	0.589	22576	0.556	[-0.069, 0.127]
Low Outlier * 100% Cue (report-level)	0.083	0.053	1.562	22580	0.118	[-0.021, 0.186]
High Outlier * 100% Cue (report-level)	-0.263	0.053	-4.968	22589	< 0.001	[-0.366, -0.159]
Low Outlier * Health Improvements	-0.101	0.049	-2.075	22569	0.038	[-0.197, -0.006]
High Outlier * Health Improvements	-0.103	0.049	-2.116	22577	0.034	[-0.199, -0.008]
Conditional R^2	0.254					

Table 7

Mean and standard deviation of error scores estimated from the subset of trials where participants were certain about encountering outliers in each experiment.

Experiment	Low-Outlier Trials		High-Outlier Trials	
	M	SD	M	SD
Exp1	-0.164	1.180	0.335	1.560
Exp2	-0.178	0.971	0.282	1.110
Exp3	-0.092	1.130	0.357	1.260
Exp4	-0.449	1.250	0.149	1.530
Exp5	-0.429	1.460	-0.027	1.740

Table 8

Estimated coefficients from the mixed linear models for error scores estimated from the outlier trials where participants were maximally certain of their outlier detection. For each model, we also reported the number of trials and the conditional R^2 .

Experiment	Term	β	SE	t	df	p-value	95% CI
Exp 1	Intercept: Low Outlier	0.043	0.136	0.316	33	0.754	[-0.224, 0.310]
	High Outlier	0.459	0.174	2.637	208	0.009	[0.118, 0.801]
	Number of Trials	218					
	Conditional R^2	0.183					
Exp 2	Intercept: 70% Cue, Low outlier	0.057	0.069	0.820	73	0.415	[-0.079, 0.193]
	High Outlier	0.445	0.081	5.463	564	< 0.001	[0.285, 0.605]
	100 % Cue	0.059	0.139	0.422	73	0.674	[-0.214, 0.331]
	High Outlier 100 % Cue	0.250	0.163	1.532	564	0.126	[-0.070, 0.569]
	Number of Trials	597					
Exp 3	Intercept: No Cue, Low Outlier, Estimate-First	0.114	0.043	2.680	208	0.008	[0.031, 0.197]
	High Outlier	0.456	0.052	8.816	1827	< 0.001	[0.354, 0.557]
	70 % Cue	0.039	0.108	0.358	224	0.721	[-0.173, 0.250]
	100 % Cue	-0.036	0.104	-0.343	207	0.732	[-0.240, 0.169]
	Detect-First	-0.029	0.085	-0.344	208	0.731	[-0.196, 0.137]
	High Outlier * 70% Cue	0.073	0.135	0.540	1833	0.589	[-0.191, 0.336]
	High Outlier * 100% Cue	0.035	0.126	0.274	1820	0.784	[-0.212, 0.281]
	High Outlier * Detect-First	-0.042	0.103	-0.407	1827	0.684	[-0.245, 0.161]
	70 % Cue * Detect-First	0.089	0.216	0.410	224	0.683	[-0.335, 0.512]
	100 % Cue * Detect-First	0.274	0.209	1.313	207	0.191	[-0.135, 0.683]
	High Outlier * 70% Cue * Detect-First	0.052	0.269	0.194	1833	0.846	[-0.475, 0.580]
	High Outlier * 100% Cue * Detect-First	0.622	0.251	2.474	1820	0.013	[0.129, 1.114]
	Number of Trials	1990					
	Conditional R^2	0.223					
Exp 4	Intercept: Low Outlier, Side-Effects	-0.166	0.081	-2.039	46	0.047	[-0.325, -0.006]
	High Outlier	0.628	0.145	4.321	355	< 0.001	[0.343, 0.913]
	Health Improvement	-0.327	0.162	-2.015	46	0.050	[-0.646, -0.009]
	High Outlier * Health-Improvements	0.206	0.291	0.708	355	0.479	[-0.364, 0.776]
	Number of Trials	366					
Exp 5	Intercept: No-Cue, Low Outlier, Side-Effects	-0.287	0.098	-2.915	178	0.004	[-0.479, -0.094]
	High Outlier	0.502	0.071	7.128	1486	< 0.001	[0.364, 0.640]
	100 % Cue	-0.081	0.197	-0.412	178	0.681	[-0.466, 0.304]
	High Outlier * 100% Cue	-0.444	0.141	-3.154	1486	0.002	[-0.721, -0.168]
	Health Improvement	-0.392	0.197	-1.993	178	0.048	[-0.777, -0.006]
	High Outlier * Health-Improvements	-0.024	0.141	-0.170	1486	0.865	[-0.300, 0.252]
	100 % Cue * Health-Improvements	0.129	0.393	0.328	178	0.743	[-0.642, 0.900]
	High Outlier * 100% Cue * Health-Improvements	-0.404	0.282	-1.434	1486	0.152	[-0.957, 0.148]
	Number of Trials	1616					
	Conditional R^2	0.533					

Control trials were excluded from this analysis.

References

- Allyn, B. (2020, March 29). *Fauci estimates that 100,000 to 200,000 americans could die from the coronavirus.*
<https://www.npr.org/sections/coronavirus-live-updates/2020/03/29/823517467/fauci-estimates-that-100-000-to-200-000-americans-could-die-from-the-coronavirus>
- André, Q., Reinholtz, N., & De Langhe, B. (2021). Can consumers learn price dispersion? evidence for dispersion spillover across categories. *Journal of Consumer Research*.
- Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121(4), 446.
- Brashier, N. M., Eliseev, E. D., & Marsh, E. J. (2020). An initial accuracy focus prevents illusory truth. *Cognition*, 194, 104054.
- Brashier, N. M., Pennycook, G., Berinsky, A. J., & Rand, D. G. (2021). Timing matters when correcting fake news. *Proceedings of the National Academy of Sciences*, 118(5).
- Chan, M.-p. S., Jones, C. R., Hall Jamieson, K., & Albarracín, D. (2017). Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological science*, 28(11), 1531–1546.
- Cook, J., & Lewandowsky, S. (2016). Rational irrationality: Modeling climate change belief polarization using bayesian networks. *Topics in cognitive science*, 8(1), 160–179.
- Ecker, U. K., Lewandowsky, S., Swire, B., & Chang, D. (2011). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic bulletin & review*, 18(3), 570–578.
- Ecker, U. K., Lewandowsky, S., & Tang, D. T. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & cognition*, 38(8), 1087–1100.

- Effron, D. A., & Raj, M. (2020). Misinformation and morality: Encountering fake-news headlines makes them seem less unethical to publish and share. *Psychological Science*, 31(1), 75–87.
- Eysenbach, G. (2002). Infodemiology: The epidemiology of (mis) information. *The American Journal of Medicine*, 113(9), 763–765.
- Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General*, 144(5), 993.
- Fazio, L. K., Rand, D. G., & Pennycook, G. (2019). Repetition increases perceived truth equally for plausible and implausible statements. *Psychonomic bulletin & review*, 26(5), 1705–1710.
- Filipowicz, A., Valadao, D., Anderson, B., & Danckert, J. (2018). Rejecting outliers: Surprising changes do not always improve belief updating. *Decision*, 5(3), 165.
- Gallucci, M. (2019). *Gamlj: General analyses for linear models* [jamovi module].
<https://gamlj.github.io/>
- Greenspan, R. L., & Loftus, E. F. (2021). Pandemics and infodemics: Research on the effects of misinformation on memory. *Human Behavior and Emerging Technologies*, 3(1), 8–12.
- Hara, K., Adams, A., Milland, K., Savage, S., Callison-Burch, C., & Bigham, J. P. (2018). A data-driven analysis of workers' earnings on amazon mechanical turk. *Proceedings of the 2018 CHI conference on human factors in computing systems*, 1–14.
- 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.
- Hassan, A., & Barber, S. J. (2021). The effects of repetition frequency on the illusory truth effect. *Cognitive Research: Principles and Implications*, 6(1), 1–12.
- Jamovi project. (2021). *Jamovi* (Version 1.6) [Computer Software]. <https://www.jamovi.org>
- Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of experimental psychology: Learning, memory, and cognition*, 20(6), 1420.

- Johnson, P. C. (2014). Extension of nakagawa & schielzeth's r2glmm to random slopes models. *Methods in Ecology and Evolution*, 5(9), 944–946.
- Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological science in the public interest*, 13(3), 106–131.
- London, K., & Nunez, N. (2000). The effect of jury deliberations on jurors' propensity to disregard inadmissible evidence. *Journal of Applied Psychology*, 85(6), 932.
- McFadden, D., et al. (1973). Conditional logit analysis of qualitative choice behavior.
- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*, 592(7855), 590–595.
- Prévost, J.-G., & Beaud, J.-P. (2015). *Statistics, public debate and the state, 1800–1945: A social, political and intellectual history of numbers*. Routledge.
- R Core Team. (2020). *R: A language and environment for statistical computing* (Version 4.0) [Computer Software (R packages retrieved from MRAN snapshot 2020-08-24)]. <https://www.R-project.org>
- Ripley, B., Venables, W., Bates, D. M., Hornik, K., Gebhardt, A., & Firth, D. (2018). *Mass: Support functions and datasets for venables and ripley's mass*. [R package]. <https://cran.r-project.org/package=MASS>
- Schaaf, J. M., Bederian-Gardner, D., & Goodman, G. S. (2015). Gating out misinformation: Can young children follow instructions to ignore false information? *Behavioral sciences & the law*, 33(4), 390–406.
- Stubenvoll, M., & Matthes, J. (2021). Why retractions of numerical misinformation fail: The anchoring effect of inaccurate numbers in the news. *Journalism & Mass Communication Quarterly*, 10776990211021800.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.

- Westlund, O., & Hermida, A. (2021). Data journalism and misinformation. In *The routledge companion to media disinformation and populism* (pp. 142–150). Routledge.
- Wilkes, A., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *The Quarterly Journal of Experimental Psychology*, 40(2), 361–387.
- Wistrich, A. J., Guthrie, C., & Rachlinski, J. J. (2004). Can judges ignore inadmissible information-the difficulty of deliberately disregarding. *U. Pa. L. Rev.*, 153, 1251.