



Repetition increases perceived truth equally for plausible and implausible statements

Lisa K. Fazio¹ · David G. Rand² · Gordon Pennycook³

Published online: 16 August 2019
© The Psychonomic Society, Inc. 2019

Abstract

Repetition increases the likelihood that a statement will be judged as true. This illusory truth effect is well established; however, it has been argued that repetition will not affect belief in unambiguous statements. When individuals are faced with obviously true or false statements, repetition should have no impact. We report a simulation study and a preregistered experiment that investigate this idea. Contrary to many intuitions, our results suggest that belief in all statements is increased by repetition. The observed illusory truth effect is largest for ambiguous items, but this can be explained by the psychometric properties of the task, rather than an underlying psychological mechanism that blocks the impact of repetition for implausible items. Our results indicate that the illusory truth effect is highly robust and occurs across all levels of plausibility. Therefore, even highly implausible statements will become more plausible with enough repetition.

Keywords Truth · Repetition · Illusory truth · Plausibility

How do people decide if a statement is true or false? Over three decades of research indicate that repeated statements are more likely to be judged true than novel statements (see Dechêne, Stahl, Hansen, & Wänke, 2010, for a review). Termed the illusory truth effect, these findings are particularly significant in the modern world where falsehoods are often repeated by politicians, advertisers, and public figures in the news and on social media. One prevalent explanation for the illusory truth effect is that repeated statements are more easily processed and understood, and this processing fluency is used as a signal for truth (Unkelbach, 2007). Thus, other manipulations that increase processing fluency (e.g., presenting statements in easy-to-read font colors) also increase truth ratings (Reber & Schwarz, 1999; Unkelbach, 2007).

An open question, however, is whether repetition provides a consistent boost to perceived truth for all statements, or if it is particularly powerful for plausible or implausible statements. Originally, it was believed that people only relied on repetition as a cue for truth if they did not have any other information available (such as prior knowledge). For example, the authors of an illusory truth meta-analysis wrote that “the only constraint seems to be that the statements have to be ambiguous, that is participants have to be uncertain about their truth status because otherwise the statements’ truthfulness will be judged on the basis of their knowledge and not on the basis of fluency” (Dechêne et al., 2010, p. 239).

However, recent studies have shown that participants give higher truth ratings to repeated statements, even when the statements contradict their prior knowledge (Fazio, 2019; Fazio, Brashier, Payne, & Marsh, 2015). For example, even among participants who subsequently provided evidence that they knew that a one-eyed giant is called a cyclops, participants who read “The Minotaur is the legendary one-eyed giant in Greek mythology” twice gave the statement higher truth ratings than participants who read it only once. This implies that both plausible and implausible statements may show an illusory truth effect.

Nonetheless, there is some evidence to suggest that the illusory truth effect does *not* occur for extremely implausible statements (Pennycook, Cannon, & Rand, 2018). When

✉ Lisa K. Fazio
lisa.fazio@vanderbilt.edu

¹ Department of Psychology and Human Development, Vanderbilt University, 230 Appleton Place #552, Jesup 105, Nashville, TN 37203, USA

² Sloan School and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

³ Hill/Levene Schools of Business, University of Regina, Regina, SK, Canada

shown statements such as “The earth is a perfect square” or “A single elephant weighs less than a single ant,” participants rated them as equally false whether they were novel or repeated. However, it is possible that the illusory truth effect occurs for these statements as well, but the increase is masked by the extreme disbelief. That is, the statements are so disbelieved initially that even with an increase in belief due to repetition they are still rated as definitely false.

We conducted two studies to systematically examine the relationship between statement plausibility and the size of the illusory truth effect. First, we did a simulation study to examine what the relationship would look like if repetition affects all statements similarly versus if it has a greater effect on plausible statements, a greater effect on implausible statements, or if the effect is greatest for items near the midpoint of the scale. Next, we conducted an empirical study examining how 500 participants rated novel and repeated statements from across the full range of plausibility (highly implausible to highly plausible) to determine which simulation most closely matched actual behavior.

Study 1

Model

In our simulations, N simulated subjects judge the accuracy of M simulated statements under either novel or repeated conditions. Each statement i has plausibility level p_i such that when that statement is novel, subject j 's perceived accuracy of that statement $a_{i,j}$ is drawn from a normal distribution centered at p_i with standard deviation σ . We then model the effect of repetition by increasing the mean of the sampling distribution: When statement i is repeated rather than novel, $a_{i,j}$ is drawn from a normal distribution centered at $p_i + f_i$ with standard deviation μ . Thus, f_i describes the increase in average perceived accuracy due to repetition for statement i . Finally, we convert subject j 's perceived accuracy of statement i $a_{i,j}$ (which is a continuous variable ranging from negative infinity to positive infinity) into a binary true/false judgment based on whether $a_{i,j}$ is greater or less than 0.5. We then simulate data for $N = 100,000$ subjects evaluating statements with plausibility p_i ranging from -1 (highly implausible) to 2 (highly plausible) in increments of 0.01 , using $\mu = 0.5$ (the specific range of p_i values and value of σ do not qualitatively affect the results).

This modeling framework allows us to investigate what would be expected when the repetition effect is independent of baseline plausibility, when repetition has a greater or lesser effect for more plausible statements and when the repetition effect is largest for statements near the midpoint. Specifically, if the repetition effect is independent of plausibility, then f_i would not vary with p_i (e.g., $f_i = m$ for all statements i). Alternatively, f_i could vary with p_i . We consider two simple

cases in which that variation is linear, such that $f_i = mp_i$ when repetition increases with plausibility or $f_i = m(1 - p_i)$ when repetition decreases with plausibility (both bounded such that $f_i > 0$). Or, f_i might peak at the midpoint of plausibility (0.5), such that $f_i = 2m(\frac{1}{2} - |\frac{1}{2} - p_i|)$.

Results

To examine the link between plausibility and the magnitude of the illusory truth effect, we begin with the case in which the repetition effect is independent of plausibility and simulate the results for $m = 0.1$, $m = 0.2$, and $m = 0.4$. Following item response theory (Chapman & Chapman, 1988), we then plot the size of the illusory truth effect (proportion rated true when repeated minus proportion rated true when novel) for each statement against the average proportion rated true. As shown in Fig. 1a, despite the equivalent repetition effect across levels of plausibility, we observe a symmetric inverted U-shaped curve centered at 0.5 (the plausibility midpoint). Repetition has less of an effect the closer the average perceived truth is to either extreme. Conceptually, this is because when an item's plausibility value is very low, then even when it is increased by f , the perceived accuracy is still very likely to be less than 0.5 and the item is still judged to be false. Conversely, when an item's plausibility is very high, there is a ceiling effect that prevents repetition from increasing truth ratings.

We next examine the predictions when the repetition effect increases linearly with p_i (see Fig. 1b), decreases linearly with p_i (Fig. 1c), or peaks at $p_i = 0.5$ (Fig. 1d). In each case, we simulate the results for $m = 0.1$, $m = 0.2$, and $m = 0.4$. In all cases, we observe that the symmetry seen in Fig. 1a under the assumption of a constant repetition effect is broken: When the repetition effect decreases linearly with plausibility, the inverted U is centered (i.e., reaches the maximal effect size) below 0.5, and when the repetition effect increases linearly with plausibility or peaks at the plausibility midpoint, the inverted U is centered above 0.5.

Having established that (a) an inverted U-shaped distribution is expected even when repetition increases belief equally for all statements, and (b) the center of that inverted U-shaped distribution indicates the presence and direction of the relationship between repetition and plausibility, Study 2 assesses this relationship empirically.

Study 2

Method

Participants Five hundred and three participants completed the full study online via Amazon's Mechanical Turk. An additional 43 participants started, but did not finish, the study.

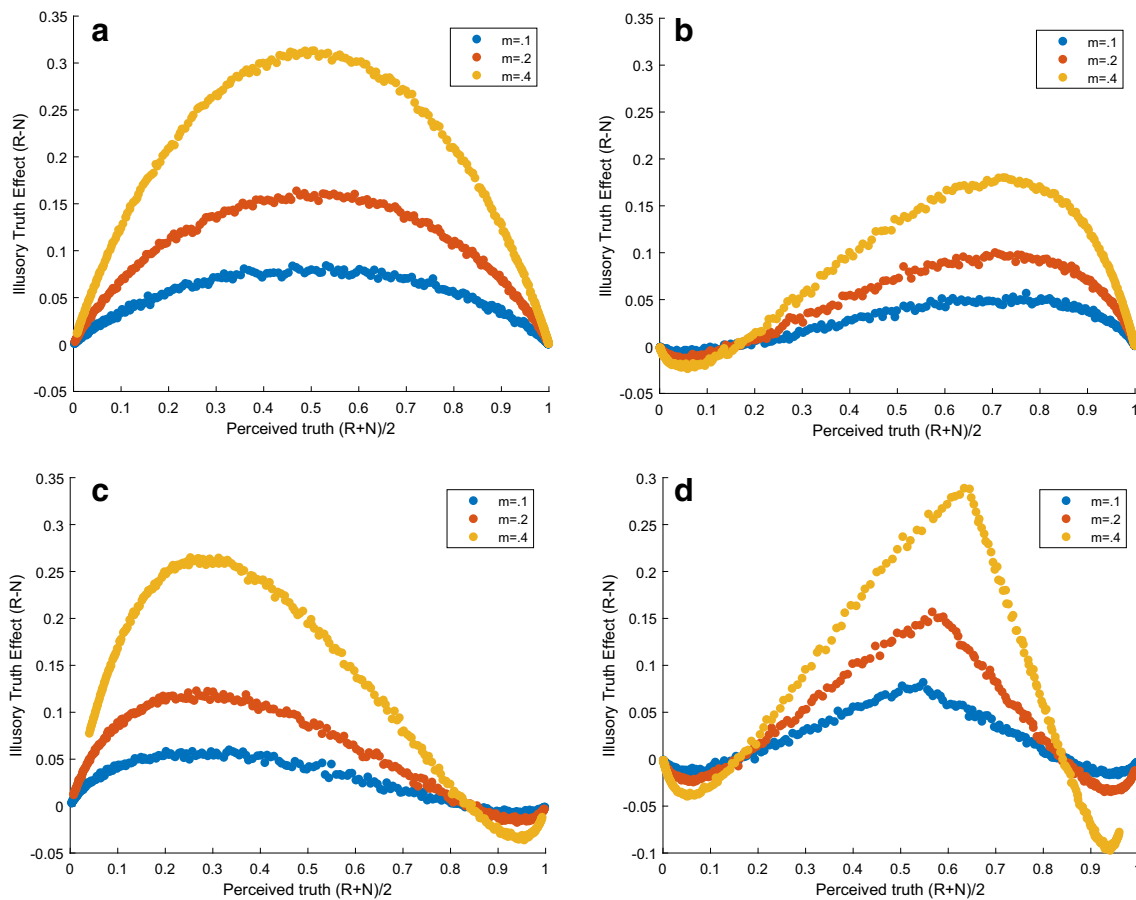


Fig. 1 Simulation results for $N = 100,000$ subjects evaluating statements with plausibility p_i ranging from -1 (highly implausible) to 2 (highly plausible) in increments of 0.01 , sampled from normal distributions with $\mu = 0.5$. **a** Repetition effect is a constant factor m regardless of statement plausibility, yielding a symmetric curve. **b** Repetition effect increases with

plausibility, yielding a right-skewed curve. **c** Repetition effect decreases with plausibility, yielding a left-skewed curve. **d** Repetition effect is maximized at plausibility of 0.5 , yielding a right-skewed curve. (Color figure online)

Materials We created a set of 80 true-and-false statements designed to cover the full range of plausibility (definitely false to definitely true). Forty statements were true and were rated as true by 50%–100% of participants in a previous study, the other 40 statements were false and were rated as true by 0%–49% of participants. In addition, all of the statements were unique such that participants never saw both a correct and incorrect version of the same statement (the full set of stimuli are available at <https://osf.io/w4k2c/>).

The majority of the statements were previously used in Experiment 3 of Fazio (2019). In that study, 102 control participants were asked to decide if each statement was “true” or “not true.” As in a typical illusory truth study, half of the statements were repeated from an earlier phase of the experiment and half were new. We used the proportion of participants who rated the statement as true (averaged across new and repeated statements) in order to select our stimuli. Within each decile of belief, we selected eight statements. That is, for the 61%–70% bin we selected eight statements that 61%–70% of the participants in the prior study had rated as true (e.g., “Napoleon was born on the island of Corsica.”).

No statements in the prior study were rated as true by less than 14% of participants. We therefore completed the full set by taking items from an unpublished follow-up to Pennycook et al.’s (2018) Experiment 1. The study featured a set of 13 highly implausible statements (e.g., “The Earth is a perfect square”) and a set of 11 highly plausible statements (e.g., “Most Americans have ridden in a vehicle of some sort”). Four hundred and ninety-two participants from Mechanical Turk rated the truth of the statements on the following scale: 1 (*not at all accurate*), 2 (*not very accurate*), 3 (*somewhat accurate*), 4 (*very accurate*). The statements were presented as in Experiment 1 of Pennycook et al. (2018). Half of the items were presented in a familiarization stage where participants were asked if they had seen or heard the claim before. Subsequently, participants were presented with the full set of items and asked to assess their accuracy. We selected four statements with an average rating from 1.13–1.44 to fill out the 11%–20% bin and eight statements rated from 1.02–1.11 for the 0%–10% bin.

For counterbalancing purposes, we divided the statements within each bin into two sets. Each participant saw one of the

two sets during the exposure phase (40 statements) and both sets during the truth phase (80 statements).

Procedure Participants began with the exposure phase. Forty statements were presented individually and participants were asked to rate how interesting each statement was on a scale from 1 (*very interesting*) to 6 (*very uninteresting*). Participants were correctly informed that some of the statements were true and others were not true.

Immediately after the exposure phase, participants began the truth rating phase. They saw a series of 80 statements and were asked to judge if each statement was true or not true. Participants were told that some of the statements were true while others were false and that some of them would be repeated from the previous task.

Results

All data are available online, along with our preregistration of the primary analyses and sample size (<https://osf.io/w4k2c/>).

Overall effect We first checked for a typical illusory truth effect. As predicted, repeated statements ($M = .52$) were more likely to be rated as true than were novel statements ($M = .48$), $t(502) = 9.19$, $p < .001$, $d = 0.41$. In addition, we were successful in sampling across the full range of belief. As shown in Table 1, the proportion of participants rating the statements as true increased across the bins.

Effect by perceived accuracy As described above, given the basic psychometric properties of the task, we expect an inverted U-shaped relationship to exist between the size of the illusory truth effect, accuracy rating for repeated minus new, and perceived truth, accuracy rating averaged over repeated and new, (e.g., Chapman & Chapman, 1988). That is, low variability in responding at the extreme ends of the spectrum of plausibility should restrict the size of the illusory truth effect in the same way that occurs for item difficulty in other tasks (such as intelligence testing; Gulliksen, 1945). Moreover, the simulations in Study 1 suggest that the location of the peak of the inverted U shape is diagnostic of the relationship between plausibility and the illusory truth effect. If the repetition effect is equivalent across all levels of plausibility the curve should peak at 0.5 (the plausibility midpoint), if the effect decreases with plausibility the curve will peak below 0.5 and if the effect increases with plausibility or is largest in the middle of the scale then the curve will peak above 0.5. Here we examine the relation between plausibility and the illusory truth effect in the empirical data.

Following our preregistration, we operationally defined each statement's perceived truth as the proportion of "true" responses averaged across new and repeated items. The size of the illusory truth effect was computed by subtracting the

Table 1 Proportion of statements rated "true" across the different bins of plausibility

Bin	New	Repeated
1%–10%	.13 (.28)	.13 (.27)
11%–20%	.20 (.29)	.23 (.29)
21%–30%	.27 (.31)	.29 (.32)
31%–40%	.36 (.31)	.41 (.32)
41%–50%	.44 (.31)	.50 (.32)
51%–60%	.53 (.32)	.59 (.32)
61%–70%	.58 (.30)	.66 (.29)
71%–80%	.68 (.27)	.72 (.27)
81%–90%	.77 (.26)	.81 (.25)
91%–100%	.88 (.22)	.89 (.19)
<i>M</i>	.48 (.17)	.52 (.16)

Standard deviations are in parenthesis

proportion of "true" responses when the statement was new from the proportion of "true" responses when the statement was repeated.¹

As shown in Fig. 2, we do observe the predicted inverted U-shaped relationship between perceived truth and the illusory truth effect. A regression predicting the size of illusory truth effect shows a significant positive linear effect of perceived truth, $\beta = 1.90$, $t(77) = 3.98$, $p < .001$, and a significant negative quadratic effect of perceived truth, $\beta = -1.84$, $t(77) = -3.85$, $p < .001$. Overall, perceived truth predicted 17% of the variance in the size of the illusory truth effect, $F(2, 77) = 7.94$, $p = .001$. Adding a cubed component did not increase the variance explained by the model, $\Delta R^2 = 0$, $F = 0.01$, $p = .935$.

We now turn to our key question of interest. To determine whether the curve is shifted in one direction or the other, we ask whether the perceived truth value corresponding to the peak illusory truth effect size is significantly different from 0.5 (scale midpoint). In the quadratic model presented in Fig. 2, the largest illusory truth effect occurs when perceived truth = 0.53. To determine whether this value is significantly different from 0.5, we use bootstrapping. Specifically, we construct 5,000 bootstrap samples by sampling our 80 items with replacement, fit the quadratic model to each sample, and calculate the plausibility value at which each sample's model reaches maximum illusory truth effect size. We then determine a 95% confidence interval on the perceived truth value yielding the maximum illusory truth effect size by sorting those 5,000 values from smallest to largest, and examining the 125th (2.5th centile) and 4,875th (97.5th centile) entries. Doing so yields a 95% confidence interval of [0.489, 0.593], which includes 0.5. Thus, our data do not suggest a significant

¹ At the item level, these differences are between subject since each item was either *only* new or *only* repeated for an individual participant.

asymmetry in the relationship between plausibility and the magnitude of the illusory truth effect, and therefore are consistent with a constant effect of fluency across varying levels of plausibility.

Discussion

This work demonstrates two important features of the illusory truth effect. First, the simulations in Study 1 demonstrate that even when internal belief in all statements is increased equally with repetition, the observed illusory truth effect will differ across different levels of statement plausibility. The basic psychometric properties of the task mean that one will observe an inverted U-shaped function with the largest repetition effect for statements near the midpoint of the scale. Thus, previous conclusions that the illusory truth effect does not occur for extremely implausible statements (e.g., Pennycook et al., 2018) are likely incorrect. It is true that there is no observable effect of repetition for extremely implausible statements, but participants' internal belief in the truth of a statement may still increase with repetition.

Second, the pattern of results in Study 2 is consistent with repetition providing a consistent boost to belief across all levels of plausibility. The simulations within Study 1 demonstrated that the midpoint of the inverted U function was diagnostic of the relation between plausibility and the illusory truth effect. When the repetition effect was larger for implausible statements, the midpoint was below 0.5, and when the repetition effect increased with plausibility or was largest for items in the middle of the plausibility scale, the midpoint was above 0.5. In contrast, we found no asymmetry in the results for

Study 2. The observed midpoint did not differ from 0.5, consistent with the repetition effect being equivalent for all statements.

These results fit with previous findings suggesting that fluency affects truth judgments independent of prior knowledge and other factors (Fazio et al., 2015; Unkelbach & Greifeneder, 2018). While participants can, and often do, judge the truth of a statement based on their prior knowledge or source credulity (Begg, Anas, & Farinacci, 1992; Unkelbach & Greifeneder, 2018), they are also influenced by low-level perceptual cues that impact fluency, such as font color and repetition (Reber & Schwarz, 1999; Unkelbach, 2007). Thus, even when participants are given advice on which statements are true or false from an advisor who is described as being 100% accurate, their truth judgements are still affected by repetition (Unkelbach & Greifeneder, 2018). In fact, in the same study, there was no evidence that the size of the illusory truth effect was affected by the reliability of the advisor. The increase in perceived truth with repetition was equivalent regardless of whether the advisor was described as being 50%, 60%, 70%, 80%, 90%, or 100% accurate (Unkelbach & Greifeneder, 2018). Similarly, we found that our results were best explained by a model where all statements show an identical increase as a function of repetition, regardless of plausibility.

It is not the case that plausibility and advisor reliability do not affect participants' truth ratings. They both play a large role. In our study, plausible statements were more likely to be judged "true" than implausible statements and in Unkelbach and Greifender (2018) participants were more likely to follow an advisor's advice when the advisor was more reliable. However, the increase in perceived truth due to repetition

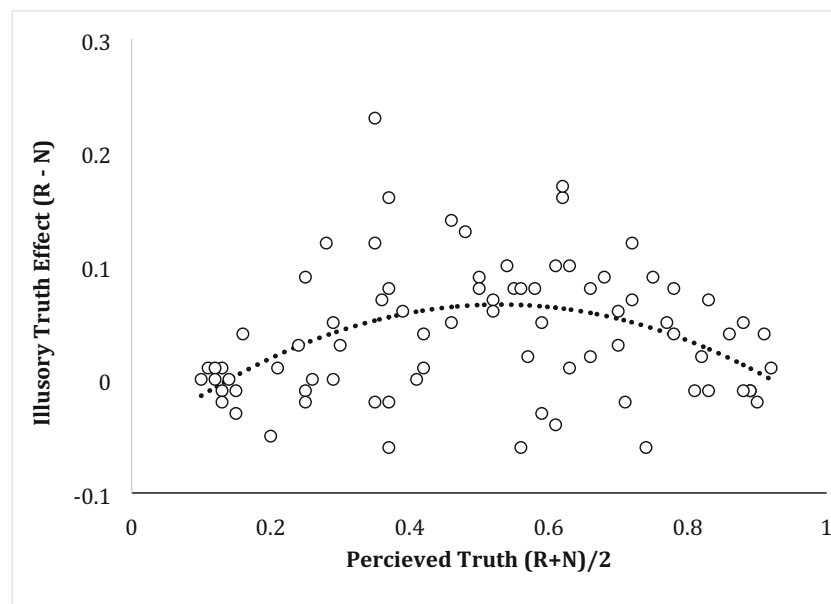


Fig. 2 Relation between perceived truth (average proportion “true”) and the illusory truth effect (proportion “true” when repeated minus proportion “true” when new) for each statement

was equivalent across all levels of reliability in Unkelbach's study and our results are consistent with repetition being equivalent across all levels of plausibility.

One important caveat for our results is that the analyses in Study 2 are conducted at the statement level by averaging across participants. While most people tended to agree that the statements in the 20%–30% bin were less likely to be true than the statements in the 50%–60% bin, there is individual variation and some participants rated the very implausible statements as true and/or the very plausible statements as false. Thus, the pattern may alter when examining statements that are very implausible to very plausible to a given participant rather than statements that vary in the aggregate. Future studies should measure participants' individual beliefs in each statement at baseline to ensure that the results hold within a single participant.

In conclusion, our findings are consistent with the hypothesis that repetition increases belief in all statements equally, regardless of their plausibility. However, there is an important difference between this internal mechanism (equal increase across plausibility) and the observable effect. The observable effect of repetition on truth ratings is greatest for items near the midpoint of perceived truth, and small or nonexistent for items at the extremes. While repetition effects are difficult to observe for very high and very low levels of perceived truth, our results suggest that repetition increases participants' internal representation of truth equally for all statements. These findings have large implications for daily life where people are often repeatedly exposed to both plausible and implausible falsehoods. Even implausible falsehoods may slowly become more plausible with repetition.

Author note We gratefully acknowledge funding from the Ethics and Governance of Artificial Intelligence Initiative of the Miami Foundation and the Social Sciences and Humanities Research Council of Canada.

Open practices statement The data, materials, and preregistration for Study 2 are available at <https://osf.io/w4k2c/>.

References

- Begg, I., 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*, 446–458.
- Chapman, L. J., & Chapman, J. P. (1988). Artifactual and genuine relationships of lateral difference scores to overall accuracy in studies of laterality. *Psychological Bulletin*, *104*(1), 127.
- Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review*, *14*(2), 238–257. doi:<https://doi.org/10.1177/1088868309352251>
- Fazio, L. K. (2019). *Preventing the illusory truth effect: When repetition does not increase perceived truth*. Manuscript under review.
- 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), 993–1002.
- Gulliksen, H. (1945). The relation of item difficulty and inter-item correlation to test variance and reliability. *Psychometrika*, *10*(2), 79–91.
- Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, *147*(12), 1865–1880. doi:<https://doi.org/10.1037/xge0000465>
- Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness & Cognition*, *8*, 338–342.
- Unkelbach, C. (2007). Reversing the truth effect: Learning the interpretation of processing fluency in judgments of truth. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *33*(1), 219–230.
- Unkelbach, C., & Greifeneder, R. (2018). Experiential fluency and declarative advice jointly inform judgments of truth. *Journal of Experimental Social Psychology*, *79*, 78–86.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.