

ORIGINAL RESEARCH REPORT

Repetition Increases Perceived Truth Even for Known Falsehoods

Lisa K. Fazio

Repetition increases belief in false statements. This illusory truth effect occurs with many different types of statements (e.g., trivia facts, news headlines, advertisements), and even occurs when the false statement contradicts participants' prior knowledge. However, existing studies of the effect of prior knowledge on the illusory truth effect share a common flaw; they measure participants' knowledge after the experimental manipulation and thus conditionalize responses on posttreatment variables. In the current study, we measure prior knowledge prior to the experimental manipulation and thus provide a cleaner measurement of the causal effect of repetition on belief. We again find that prior knowledge does not protect against the illusory truth effect. Repeated false statements were given higher truth ratings than novel statements, even when they contradicted participants' prior knowledge.

Keywords: illusory truth; truth judgments; knowledge; truth; misinformation

Every day we are exposed to numerous claims about the world. Claims such as, "The Washington Nationals won the 2019 World Series", "The pygmy jerboa is the world's smallest rodent" and "NASA's Mars 2020 rover is named Curiosity". Some of these claims are true and others are false (the rover's actual name is Perseverance). Sometimes false statements are simple mistakes or misstatements (oops, I forgot that it is the current Mars rover who is named Curiosity), but other times the falsehoods are intended to mislead and misinform (e.g. political disinformation). Thus, people must determine whether what they hear is likely true or false.

Current evidence suggests that people use a number of different cues to judge truth (see Brashier & Marsh, 2020 for a review). One cue is prior knowledge. Participants are less likely to accept falsehoods that contradict well-known facts (e.g. "Elephants weigh less than ants") as compared to false statements that contradict more obscure knowledge (e.g. "Oslo is the capital of Finland") (Fazio, Brashier, Payne, & Marsh, 2015; Fazio, Rand, & Pennycook, 2019). Another cue is the source of the information. People are less likely to believe statements that come from unreliable sources (Begg, Anas, & Farinacci, 1992; Unkelbach & Greifeneder, 2018). However, people also rely on other more proximal cues for truth. One such cue is repetition. People are more likely to think that false statements are true when they are repeated (Hasher, Goldstein, & Toppino, 1977; Unkelbach, Koch, Silva, & Garcia-Marques, 2019).

In fact, people continue to rely on these proximal cues, such as repetition, even when they have access to more direct signals of truth such as prior knowledge and source reliability (Fazio et al., 2015; Fazio et al., 2019; Fazio & Sherry, in press; Unkelbach & Greifeneder, 2018). That is, while people pay attention to cues such as prior knowledge and source credibility when judging truth, they are also affected by proximal cues such as repetition. In one example, participants were given advice on which statements were true or false by an advisor (Unkelbach & Greifeneder, 2018). Respondents paid attention to the advisor's credibility and were more likely to follow the advice of an advisor that was described as being 100% accurate than one described as being 50% accurate. However, with both advisors, repetition also affected participants' truth judgements. Repeated statements were given higher truth ratings than novel statements (Unkelbach & Greifeneder, 2018).

For decades, there was an assumption in the field that people only used repetition or fluency as a cue for truth when they could not rely on more reliable signals (see Dechêne, Stahl, Hansen, & Wänke, 2010 for a review). However, we now know that this is false. As described above, people are affected by repetition even with information about source reliability. Similarly, they are affected by repetition even when they have relevant prior knowledge. Fazio and colleagues (2015) first demonstrated that repetition increases truth ratings for known falsehoods. In their study, participants first rated the truth of a series of new and repeated general knowledge facts (e.g., "A prune is a dried plum"). Then, they were asked a series of multiple-choice questions as a knowledge check (e.g., "What is the name of a dried plum? Prune, Date, I

don't know). The authors then sorted the statements into those that were known by the participants (they answered the question correctly on the knowledge check) and those that were unknown (they were answered incorrectly or with "I don't know"). Repetition increased participants' truth ratings not only for unknown falsehoods, but also for known falsehoods.

Since then, numerous studies have replicated that repetition increases perceived truth, even for statements that contradict people's prior knowledge. Repetition increases truth ratings for implausible false news headlines (De keersmaecker et al., 2019; Pennycook, Cannon, & Rand, 2018), for facts that contradict common knowledge (Brashier, Eliseev, & Marsh, 2020; Brashier, Umanath, Cabeza, & Marsh, 2017; Fazio & Sherry, in press), and there is some evidence that it increases truth ratings equally for all statements regardless of their plausibility (Fazio et al., 2019).

However, very few studies have explicitly measured individuals' prior knowledge, rather than relying on norming studies or general implausibility, and they all share the same methodological flaw. In past research, participants first completed the illusory truth study and then completed the knowledge check (Brashier et al., 2017; Fazio et al., 2015). This ensures that the knowledge check does not interfere with measurement of the illusory truth effect. However, conditioning on posttreatment variables can also bias estimates of causal effects (see Montgomery, Nyhan, & Torres, 2018 for further explanation).

The main issue is that the experimental manipulation may affect participants' later responses. For example, imagine that researchers are interested in the effects of study skills training on participants' study habits. In a randomized trial, half of the participants receive the intervention (Group A) and half do a control activity (Group B). On a post-test, the researchers measure their study habits and also ask the participants to rate how important they think it is to learn new study skills. The researchers may want to examine if the intervention is still effective among students who do not think that learning new habits is important. However, by selecting on a post-treatment variable, the researchers are now comparing two disparate groups – participants who think learning new study skills is unimportant *despite* just learning about new study skills (Group A) and participants who think it was unimportant without the intervention (Group B). Those two groups of participants may not be the same, thus nullifying the benefits of random assignment.

Within the current paradigm, since the knowledge check occurred after the experimental randomization of statements, it is possible that what people were able to remember during the knowledge check was affected by whether the statement was new or repeated in the experiment. That is, they may have been more or less likely to answer correctly on the knowledge check when a statement was repeated as compared to when it was novel. Thus, in the current study we had participants first complete the knowledge check and then, at least 2 weeks later, complete the typical illusory truth experiment. Replicating the effect when knowledge is measured prior to the experimental session is essential in order to verify

this important finding. We predict that repetition will again increase truth ratings for both known and unknown falsehoods.

Method

As detailed below, the experiment consisted of two sessions. During session 1, participants completed the knowledge check to measure their individual prior knowledge. Then, at least two weeks later, participants completed session 2. This session consisted of an exposure phase where participants rated their interest in half of the statements followed by a truth-rating phase where participants rated the accuracy of all of the statements (both new statements and ones they had previously seen in the exposure phase).

Participants

Participants were recruited in two phases using Amazon's Mechanical Turk and both surveys were completed online. Our goal was to recruit 150 participants for the second session. As a conservative prediction, we thought that half of the participants who completed the knowledge check would also complete the illusory truth study. Thus, we recruited 300 participants for the knowledge check. Then, two weeks later, a new study was made available only to those who had completed the knowledge check. Vastly exceeding our expectations, 258 participants completed both studies. Using TurkPrime (Litman, Robinson, & Abberbock, 2017), we restricted the sample to participants in the United States, and blocked duplicate IP addresses.

Two catch trials were included (one during the exposure phase and one during the truth-rating phase) that asked participants to choose a specific answer. All participants answered at least one catch trial correctly and thus were included in the final sample. However, 2 participants listed ages that differed by more than one year across the two surveys and they were removed. This left 256 participants ($M_{\text{age}} = 37.16$, $SD = 10.48$) in the analyses below.

Design

The experiment featured a 2 (repetition: repeated, new) \times 2 (knowledge: known, unknown) within-subjects design. Across two separate analyses, the statements were sorted into those that were "known" and "unknown" based on either estimated knowledge from preexisting population norms or by participants' demonstrated knowledge on the knowledge check. For the demonstrated knowledge analyses, questions that were previously answered correctly were included as known statements and questions that were answered incorrectly or with "don't know" were included as unknown statements.

Materials

As in Fazio and colleagues (2015), the materials were derived from a normed set of trivia questions. We selected 80 questions from Tauber, Dunlosky, Rawson, Rhodes, & Sitzman's (2013) general knowledge norms. Half of the questions were likely to be known by the participants (answered correctly by 60% of the norming participants, range 42%–80% correct) and half were likely to be unknown (answered correctly by 3% of the norming participants,

range 1%–11% correct). These general knowledge norms are likely an underestimation of the participants' actual knowledge since the norming test required participants to respond to an open-ended question. For each question, we created a true (e.g., "Insomnia is an inability to sleep") and false (e.g., "Narcolepsy is an inability to sleep") statement. The falsehoods used plausible, but incorrect alternatives to the correct answer. Sample statements are shown in **Table 1**. The full set of stimuli are available at <https://osf.io/kt2p8>.

Both the known and unknown statements were divided into four groups of ten statements. Each group had the same average accuracy in the norming data as the stimuli set as a whole. These eight sets of items were then used for counterbalancing purposes. For each participant, half of the statements they saw were truths, half were falsehoods, half were known, half were unknown, and half appeared in both the exposure and truth rating phases (repeated statements), while half were only presented during the truth rating phase (new statements). As in Fazio et al. (2015), we limited our analyses to the falsehoods and treated the correct statements as fillers.

In addition, each statement was turned into a multiple-choice question for the knowledge check. The alternatives for the multiple-choice question included the correct answer, the incorrect answer used in the false statement and a "don't know" option. For example, "What is the name for a cyclone that occurs over land? Hurricane, Tornado, I Don't Know". For half of the questions the correct answer was presented first and for the other half the incorrect answer was presented first. The last option was always "I don't know."

Procedure

During the knowledge check, participants were asked to answer the 80 multiple-choice questions. They were instructed to choose "I don't know" if they did not know the answer to a question and were explicitly instructed not to look up the answers to any of the questions.

Two weeks later, participants were invited to complete the main experiment. There was no explicit connection made between the two studies, other than the fact that the requests came from the same research group. Participants began the study with the exposure phase and were asked to rate how interesting (1 = very interesting, 2 = interesting, 3 = slightly interesting, 4 = slightly

uninteresting, 5 = uninteresting, 6 = very uninteresting) they found a series of 40 facts. The participants were told that the researchers were developing stimuli for a new experiment and were interested in their opinions about the statements. In addition, they were told that some of the statements were true and others were false.

Immediately following the exposure phase, participants began the truth rating phase. Eighty statements were presented one-by-one and participants were asked to rate how true or false each statement appeared to them (1 = definitely false, 2 = probably false, 3 = possibly false, 4 = possible true, 5 = probably true, 6 = definitely true). They were again warned that some of the statements were true, while others were false, and that some of the statements would be repeated from the previous task.

Results

Alpha was set at .05 for all statistical tests. As mentioned above, we focused our analysis on the falsehoods.

Knowledge check

On average, participants correctly answered 61% (*SD* = 14) of the questions on the knowledge check. They choose the incorrect answer on 12% (*SD* = 10) of the questions and responded "I don't know" on the remaining 27% (*SD* = 16). Thus, 61% of the items were known and 39% were unknown. As would be expected, participants correctly answered more of the questions estimated to be known by the norms (*M* = 86%, *SD* = 14), than those estimated to be unknown (*M* = 36%, *SD* = 21), *t*(255) = 37.00, *p* < .001, *d* = 2.31.

Truth ratings as a function of demonstrated knowledge

Analyses focused on participants' ratings for the false items. Known statements referred to questions that the participant answered correctly on the knowledge check and unknown statements were ones that were answered incorrectly or with "I don't know". The number of known and unknown items differed for each participant depending on their responses during knowledge check. Thirty participants were excluded because they had fewer than three responses in a given cell (e.g., of the 20 repeated falsehoods only two referred to questions that were answered correctly on the knowledge check). This exclusion criterion was established before analyzing

Table 1: Sample known and unknown statements.

	Truth	Falsehood
Known	A tornado is a cyclone that occurs over land.	A hurricane is a cyclone that occurs over land.
	Tennis is the sport associated with Wimbledon.	Rugby is the sport associated with Wimbledon.
	An ostrich is a bird that cannot fly and is the largest bird on Earth.	An emu is a bird that cannot fly and is the largest bird on Earth.
Unknown	Mozart is the composer who wrote the opera "Don Giovanni."	Bach is the composer who wrote the opera "Don Giovanni."
	Napoleon was born on the island of Corsica.	Napoleon was born on the island of Sardinia.
	Angel Falls is located in Venezuela.	Angel Falls is located in Brazil.

the data. This left 226 participants in the analyses below.¹ Within each cell, the median participant had 12 statements that were known and 8 that were unknown. Full cell counts are available in the online supplement.

As shown in **Figure 1**, the results replicated those of Fazio et al. (2015). Repetition increased perceived truth for both known and unknown falsehoods. To examine these patterns statistically, we conducted a 2 (repetition: repeated, new) \times 2 (demonstrated knowledge: known, unknown) ANOVA on participants' truth ratings for the falsehoods. As would be expected, participants gave lower truth ratings for known falsehoods ($M = 2.80$) than unknown falsehoods ($M = 4.01$), $F(1, 225) = 547.48$, $p < .001$, $\eta_p^2 = .71$. In addition, there was a typical illusory truth effect, $F(1, 225) = 57.18$, $p < .001$, $\eta_p^2 = .20$. Repeated falsehoods ($M = 3.55$) were given higher truth ratings than novel falsehoods ($M = 3.27$). Replicating Fazio and colleagues (2015), there was no interaction between repetition and prior knowledge, $F(1, 225) = 0.67$, $p = .415$, $\eta_p^2 = .003$. Repetition increased truth ratings for both unknown falsehoods (New $M = 3.89$, Repeated $M = 4.14$, $t(225) = 6.49$, $p < .001$, $d = 0.43$) and for false statements that contradicted participants' prior knowledge (New $M = 2.65$, Repeated $M = 2.95$, $t(225) = 5.45$, $p < .001$, $d = 0.37$).

Truth ratings as a function of knowledge estimated by norms

As shown in **Figure 2**, the pattern was very similar when the statements were sorted into known and unknown categories by the norming data. We conducted a 2 (repetition: repeated,

new) \times 2 (estimated knowledge: known, unknown) ANOVA on participants' truth ratings (with the full sample of 256 participants). There was again a main effect of knowledge, $F(1, 255) = 422.43$, $p < .001$, $\eta_p^2 = .62$. Participants rated unknown falsehoods ($M = 3.65$) as more accurate than known falsehoods ($M = 2.82$). In addition, repetition again affected truth judgements, $F(1, 255) = 46.52$, $p < .001$, $\eta_p^2 = .15$. Repeated falsehoods ($M = 3.36$) were rated as more accurate than novel statements ($M = 3.11$). As with split by demonstrated knowledge, there was again no interaction between repetition and prior knowledge, $F(1, 255) = 1.06$, $p = .305$, $\eta_p^2 = .004$. Repetition increased truth ratings for falsehoods that would be unknown to most participants (New $M = 3.54$, Repeated $M = 3.76$, $t(255) = 4.79$, $p < .001$, $d = 0.30$) and for false statements that likely contradicted participants' prior knowledge (New $M = 2.67$, Repeated $M = 2.96$, $t(255) = 5.05$, $p < .001$, $d = 0.32$).

Discussion

As in previous research (Brashier et al., 2017; Fazio et al., 2015), prior knowledge did not protect participants from the illusory truth effect. Repeated falsehoods were rated as being more true than novel falsehoods, even when they both contradicted participants' prior knowledge. By measuring prior knowledge before the experimental session, this study avoids conditioning on posttreatment variables and provides cleaner evidence for the effect (Montgomery et al., 2018). Whether prior knowledge is measured before or after the manipulation, it is clear that repetition increases belief in falsehoods that contradict existing knowledge.

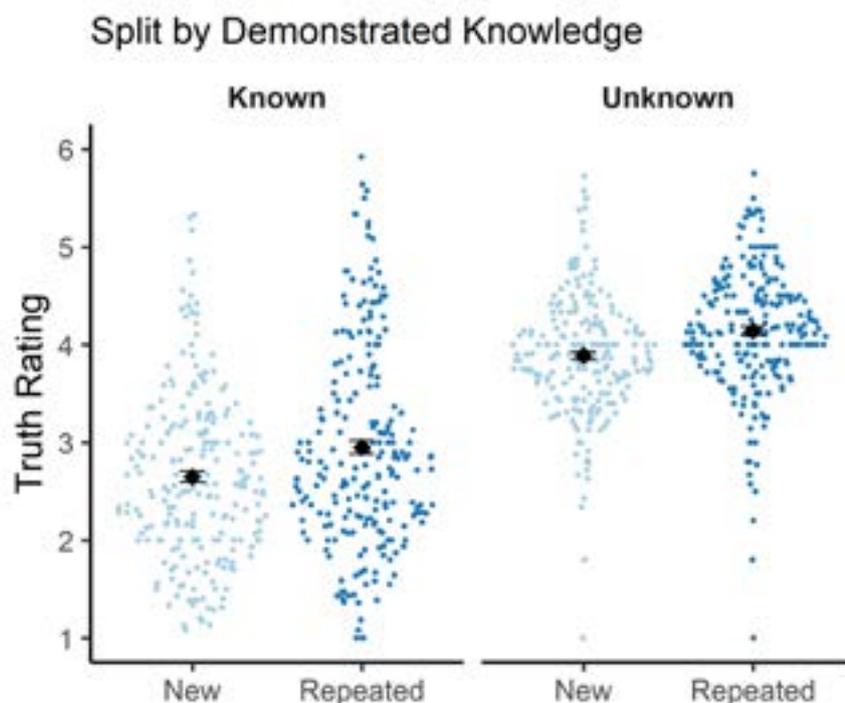


Figure 1: Mean truth ratings for falsehoods as a function of repetition and demonstrated knowledge on the knowledge check. The scale ranged from 1 = Definitely False to 6 = Definitely True. Each dot represents one participant ($N = 226$). In order to show the distribution of the data, points are shifted horizontally to represent the density distribution. The black diamond represents the mean value and the error bars reflect standard error of the mean.

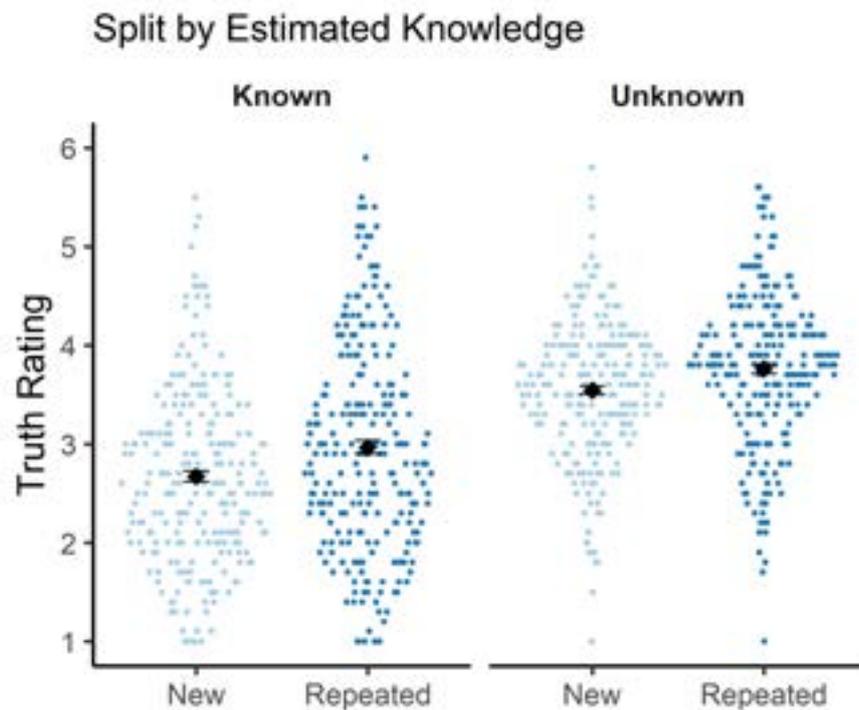


Figure 2: Mean truth ratings for falsehoods as a function of repetition and knowledge estimated from the norming data. The scale ranged from 1 = Definitely False to 6 = Definitely True. Each dot represents one participant ($N = 256$). In order to show the distribution of the data, points are shifted horizontally to represent the density distribution. The black diamond represents the mean value and the error bars reflect standard error of the mean.

This finding fits within a larger literature on “knowledge neglect.” That is, people often neglect knowledge that they possess and fail to notice when information contradicts their prior knowledge. For example, when reading a fictional story which mentions that “the Atlantic is the largest ocean on Earth”, many readers do not notice the error in that sentence and will later answer the question “What’s the largest ocean on Earth?” with “Atlantic” (e.g., Marsh, Meade, & Roediger, 2003). Importantly, some readers will still respond with “Atlantic”, even when they correctly answered the question with “Pacific” two weeks earlier (Fazio, Barber, Rajaram, Ornstein, & Marsh, 2013). This negative consequence from reading false information even exists among people who previously answered correctly with very high confidence (Fazio et al., 2013). The parallels to the current research are clear. In the current study, participants failed to rely on their prior knowledge when judging truth and instead relied (at least in part) on repetition as a signal of truth. Future studies should examine not only participants’ answers on the knowledge check but also their confidence in those answers. It is still possible that participants will not rely on repetition as a cue when they are judging statements that contradict prior high-confidence correct responses.

While it is possible that very well-known information may be immune to the effects of repetition, the current results are consistent with the idea that repetition increases belief equally for all statements, regardless of source credibility or prior knowledge (Fazio et al., 2019; Fazio & Sherry, in press; Unkelbach & Greifeneder, 2018). However, such a boost may not be inevitable. Recent

research suggests that having participants focus on the accuracy of the statement during initial exposure can reduce the effects of repetition on belief (Brashier et al., 2020). Similarly, thinking about the accuracy of a headline (through a prompt to “explain how you know the headline is true or false”) reduces participants’ intentions to share false news headlines, but the reduction is much larger if the accuracy prompt occurs during the first viewing of a headline rather than when it is repeated (Fazio, 2020). Careful processing during initial exposure can help people recognize the accuracy of what they are reading. In this current world, where false information is abundant and often repeated, people should actively consider the accuracy of what they read, rather than relying on their gut-feelings of truth.

Data Accessibility Statement

The stimuli and data are available at <https://osf.io/kt2p8>.

Note

¹ The pattern of results is identical for the sample that provided at least one response of each type ($N = 251$). This analysis is presented at <https://osf.io/kt2p8>.

Acknowledgements

We thank Deep Patel for his help with data collection and Carrie Sherry for her copyediting skills.

Funding Information

This research was partially funded by a gift from Facebook Research.

Competing Interests

The author has no competing interests to declare.

Author Contributions

LKF is the sole author of the paper.

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. DOI: <https://doi.org/10.1037/0096-3445.121.4.446>
- Brashier, N. M., Eliseev, E. D., & Marsh, E. J.** (2020). An initial accuracy focus prevents illusory truth. *Cognition*, *194*, 104054. DOI: <https://doi.org/10.1016/j.cognition.2019.104054>
- Brashier, N. M., & Marsh, E. J.** (2020). Judging truth. *Annual Review of Psychology*, *71*, 499–515. DOI: <https://doi.org/10.1146/annurev-psych-010419-050807>
- Brashier, N. M., Umanath, S., Cabeza, R., & Marsh, E. J.** (2017). Competing cues: Older adults rely on knowledge in the face of fluency. *Psychology and Aging*, *32*(4), 331–337. DOI: <https://doi.org/10.1037/pag0000156>
- De keersmaecker, J., Dunning, D., Pennycook, G., Rand, D. G., Sanchez, C., Unkelbach, C., & Roets, A.** (2019). Investigating the Robustness of the Illusory Truth Effect Across Individual Differences in Cognitive Ability, Need for Cognitive Closure, and Cognitive Style. *Personality and Social Psychology Bulletin*, *46*(2), 204–215. DOI: <https://doi.org/10.1177/0146167219853844>
- 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.** (2020). Pausing to consider why a headline is true or false can help reduce the sharing of false news. *The Harvard Kennedy School (HKS) Misinformation Review*. DOI: <https://doi.org/10.37016/mr-2020-009>
- Fazio, L. K., Barber, S. J., Rajaram, S., Ornstein, P. A., & Marsh, E. J.** (2013). Creating illusions of knowledge: Learning errors that contradict prior knowledge. *Journal of Experimental Psychology: General*, *142*(1), 1–5. DOI: <https://doi.org/10.1037/a0028649>
- 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. DOI: <https://doi.org/10.1037/xge0000098>
- 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. DOI: <https://doi.org/10.3758/s13423-019-01651-4>
- Fazio, L. K., & Sherry, C.** (in press). The Effect of Repetition on Truth Judgments across Development. *Psychological Science*.
- Hasher, L., Goldstein, D., & Toppino, T.** (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning & Verbal Behavior*, *16*, 107–112. DOI: [https://doi.org/10.1016/S0022-5371\(77\)80012-1](https://doi.org/10.1016/S0022-5371(77)80012-1)
- Litman, L., Robinson, J., & Abberbock, T.** (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*, 433–442. DOI: <https://doi.org/10.3758/s13428-016-0727-z>
- Marsh, E. J., Meade, M. L., & Roediger, H. L.** (2003). Learning facts from fiction. *Journal of Memory & Language*, *49*, 519–536. DOI: [https://doi.org/10.1016/S0749-596X\(03\)00092-5](https://doi.org/10.1016/S0749-596X(03)00092-5)
- Montgomery, J. M., Nyhan, B., & Torres, M.** (2018). How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science*, *62*(3), 760–775. DOI: <https://doi.org/10.1111/ajps.12357>
- 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>
- Tauber, S. K., Dunlosky, J., Rawson, K. A., Rhodes, M. G., & Sitzman, D. M.** (2013). General knowledge norms: Updated and expanded from the Nelson and Narens (1980) norms. *Behavior Research Methods*, *45*(4), 1115–1143. DOI: <https://doi.org/10.3758/s13428-012-0307-9>
- Unkelbach, C., & Greifeneder, R.** (2018). Experiential fluency and declarative advice jointly inform judgments of truth. *Journal of Experimental Social Psychology*, *79*, 78–86. DOI: <https://doi.org/10.1016/j.jesp.2018.06.010>
- Unkelbach, C., Koch, A., Silva, R. R., & Garcia-Marques, T.** (2019). Truth by Repetition: Explanations and Implications. *Current Directions in Psychological Science*, *28*(3), 247–253. DOI: <https://doi.org/10.1177/0963721419827854>

Peer review comments

The author(s) of this paper chose the Open Review option, and the peer review comments can be downloaded at:
<http://doi.org/10.1525/collabra.347.pr>

How to cite this article: Fazio, L. K. (2020). Repetition Increases Perceived Truth Even for Known Falsehoods. *Collabra: Psychology*, 6(1): 38. DOI: <https://doi.org/10.1525/collabra.347>

Senior Editor: Amy Perfors

Editor: Ullrich Ecker

Submitted: 23 March 2020 **Accepted:** 01 July 2020 **Published:** 28 July 2020

Copyright: © 2020 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.