Is Earth a Perfect Square?

Repetition Increases the Perceived Truth of Highly Implausible Statements

Doris Lacassagne, Jérémy Béné & Olivier Corneille

UCLouvain, Belgium

[Preprint]

Word count: 2999 (without the title page, acknowledgments, references, footnotes, and figure captions)

Running title: Truth-by-Repetition effect for highly implausible statements

Conflict of interest statement: The authors declare no conflict of interest.

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Correspondence concerning this article should be addressed to Jérémy Béné, UCLouvain, PSP IPSY, 10 Place du Cardinal Mercier, 1348, Louvain-la-Neuve, Belgium. Email:

jeremy.bena@uclouvain.be

CRediT authorship contribution statement

DL: Conceptualization, Formal analysis, Investigation, Methodology, Software, Writing – original draft. JB: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. OC: Conceptualization, Methodology, Supervision, Writing – review & editing.

Abstract

A single exposure to statements is typically enough to increase their perceived truth. This Truth-by-Repetition (TBR) effect has long been assumed to occur only with statements whose truth value is unknown to participants. Contrary to this hypothesis, recent research found a TBR effect with statements known to be false. Of note, a recent model even posits that repetition could increase the perceived truth of highly implausible statements. As for now, however, no empirical evidence has reported a TBR effect for highly implausible statements. Here, we reasoned that one may be found provided a sensitive truth measure is used and statements are repeated more than just once. In a preregistered experiment, participants judged the truth of highly implausible statements on a 100-point scale, and these statements were either new to them or had been presented five times before the judgment task. We observed a TBR effect: truth judgments were higher for repeated statements than for new ones - even if all statements were still judged as false. Exploratory analyses additionally suggest that all participants were not equally prone to this TBR effect: about half the participants showed no or even a reverse effect. Overall, the results provide direct empirical evidence to the claim that repetition can increase perceived truth even for highly implausible statements, although not equally so for all participants and not to the point of making the statements look true.

Keywords: truth effect; illusory truth; repetition; multiple exposures; low plausibility

Is Earth a Perfect Square?

Repetition Increases the Perceived Truth of Highly Implausible Statements

1. Introduction

People are exposed every day to a large amount of information, some of which is false (Marsh & Rajaram, 2019; Vosoughi et al., 2018), and sometimes seems highly implausible. In general, repeated exposure to a given piece of information increases its believability. This Truth-by-Repetition (TBR) effect is robust and has been demonstrated hundreds of times (for a meta-analysis, see Dechêne et al., 2010; for early demonstrations, see Hasher et al., 1977; Bacon, 1979; Schwartz, 1982; for a recent overview, see Unkelbach et al., 2019). A prominent theory is based on processing fluency (Reber & Schwarz, 1999; Unkelbach, 2007): repeated statements are easier to process than new ones, and this processing fluency is used as a cue for truth. Typically, the TBR effect is demonstrated using only one previous exposure, but additional exposures (e.g., 4, 8, 16) may lead to a larger TBR effect (Arkes et al., 1991; DiFonzo et al., 2016; Fazio et al., 2021; Hassan & Barber, 2021).

A critical question, from both a theoretical and practical standpoint, is whether the TBR effect is observed for statements known to be false. A widespread assumption is that a statement's truth status must be ambiguous for a TBR effect to be found. In their meta-analysis of the TBR effect, Dechêne et al. (2010, p.239) noted:

“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.”

This “truth ambiguity” assumption has been formally implemented in Multinomial Processing Tree (MPT) models. For instance, in Unkelbach and Stahl’s (2009) MPT model, a statement’s factual truth is known with a probability k (for knowledge), and it is only in the absence of this knowledge ($1-k$) that enhanced fluency should increase the probability f of judging this statement as true (see also Hilbig, 2012, for a related model of truth judgments).

Recent studies finding a TBR effect with statements known to be false, however, question the “truth ambiguity” assumption (Brashier et al., 2020; Fazio et al., 2015; Fazio et al., 2019; Fazio, 2020). Fazio et al. (2015) compared the fit of two MPT models. One model was knowledge-conditional; in that model, fluency drives judgments with a probability f only in the absence of knowledge. The other model was fluency-conditional; in that model, fluency drives judgments with a probability f , and knowledge drives judgments only in the absence of fluency. In two experiments where some statements’ truth status was presumably known, these authors found that the fluency-conditional model fitted the data well, while this was not the case for the knowledge-conditional model. Further challenging the “truth ambiguity” assumption, Pennycook et al. (2018) recently found a TBR effect with *fake news*, which commonly lacks plausibility¹.

But how false and implausible can a statement be for resisting a TBR effect? Surprisingly enough, very little research to date has informed this question of high theoretical and practical relevance. Pennycook et al. (2018) did not find a TBR effect with extremely implausible statements (e.g., “The earth is a perfect square”), which suggests a plausibility constraint on the

¹Relatedly, TBR effects were found even when declarative information about statements’ truth status was available when performing truth judgments (Unkelbach & Greifeneder, 2018). The truth ambiguity assumption is also questioned by studies finding TBR effects on statements commonly known as true (Unkelbach & Speckmann, 2021).

TBR effect. However, Fazio et al. (2019) proposed a model where repetition-induced belief is constant across all levels of plausibility. In this model, repeated exposure to unambiguously false, unambiguously true, and uncertain statements should lead to the same increase in belief in the statements. Yet, because unambiguously false statements are unlikely to be judged as true regardless of their repetition, the TBR effect is less likely to be observed with such statements.

Not categorizing as true a highly implausible statement, however, does not mean that repetition did not influence its perceived truth value. It may just be the case that typical truth measures are insufficiently sensitive to a repetition effect. A TBR effect may be observed when a large gradient of falsehood is allowed. In the present study, we focused on highly implausible statements and tested whether we could find empirical evidence for an effect of repetition on their perceived truth. We choose to use only highly implausible statements because we believe finding a TBR effect with such statements is of greatest concern from an applied perspective, as misinformation and disinformation (e.g., fake news; hoaxes; conspiracy theories) can spread on a large scale (e.g., Del Vicario et al., 2016; Vosoughi et al., 2018). In addition, a TBR effect on highly plausible (but not highly implausible) statements could be constrained by ceiling effects.

Finding a TBR effect with highly implausible statements would suggest that even blatantly false and preposterous information may benefit from repeated exposure. Here, we departed from the typical TBR effect paradigm in two regards. First, repeated statements were displayed five times (not only one) in the exposure phase. Second, we used a 100-point scale in the truth judgment task (not a dichotomous rating task or a 7-point Likert scale). In doing so, we may be able to detect a TBR effect for highly implausible statements, thereby giving credence to Fazio et al.'s (2019) model and questioning the “truth ambiguity” assumption.

2. Method

We report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study. The preregistration, experiment program, data, and analyses are publicly available on <https://osf.io/pe4g9/>.

2.1 Participants and design

The design was a 2 (Repetition: New vs. Repeated five times) $\times$ 2 (Counterbalancing: Set A repeated and Set B new vs. Set B repeated and Set A new), with the first factor within participants.

We recruited 240 participants (our targeted sample size) on Prolific (51.25% female, one not reported; $M_{\text{age}} = 35.28$; $SD_{\text{age}} = 11.96$, one not reported). Participants (1) were English speakers, (2) declared to live in the United States, (3) had an approval rate of at least 95%, and (4) had already completed at least 100 previous submissions. Participants were paid US\$1.41 for completing the study. We excluded eight participants (seven participants did not end the study, resulting in no data; one participant declared that they did not take their response seriously). This resulted in a final sample size of 232 participants.

To determine sample size, we set α to .05, and we aimed for a statistical power of 80% to detect an effect as small as Cohen's $d = 0.2$ in a paired samples t -test. An analysis with G*Power (Version 3.1.9.6, Faul et al., 2007) found that we would need 199 participants. To avoid a final sample smaller than the targeted sample size, we have increased this estimate by 20%, resulting in a targeted sample size of 240 participants.

2.2 Materials and procedure

We programmed the experiment with lab.js (Henninger et al., 2019), and we used JATOS (Lange, Kühn, & Filevich, 2015) to run the study online on Prolific. We used the 16 most

implausible statements used by Fazio et al. (2019) (only 0% to 20% of the participants rated them as true).

Upon clicking on the study link, participants were told that we were interested in the role of various variable in the evaluation of statements. After providing their informed consent, participants entered the exposure phase. The instructions were as follow: “In this first part, you will judge how interesting some statements are to you on a scale from 1 (very uninteresting) to 6 (very interesting). For each statement, please select the value that most closely represents what you think. Some of these statements will appear several times. When you are ready, please start by pressing the spacebar.” Participants rated their perceived interest on eight statements (either from Set A or from Set B). The eight statements were presented five times in a random order, resulting in 40 trials. To rate their interest, participants used a 6-point Likert scale (1: Very uninteresting; 6: Very interesting) with no time limit. Upon clicking on their response, the next statement was displayed with an intertrial blank screen of 500 milliseconds.

Right after the exposure phase, participants began the truth judgment task. The instructions were as follow: “In the following part, we are interested in how true you think some statements are. Some statements will be repeated from the previous task. To evaluate the truth of the statements, you will use a scale from -50 (Definitely false) to 50 (Definitely true). For each statement, please select the value that most closely represents what you think. When you are ready, please start by pressing the spacebar.” Participants rated their perceived truth of 16 statements (half seen in the exposure phase) displayed in a random order on a scale from -50 (Definitely false) to +50 (Definitely true – we recoded these values between 0 and 100), without time limit.

After the truth judgment task, we administered a seriousness check (based on Aust et al., 2013) to exclude participants that did not take their responses seriously (participants were informed that their response to this question would not affect their payment). We then thanked and debriefed participants.

3. Results

3.1 Preregistered analyses

First, we conducted a Shapiro-Wilk test on differences between mean truth judgments on repeated vs. new statements to determine whether we should use parametric or non-parametric tests. Because data distribution was not normal, we used a Wilcoxon signed-rank test. However, as the distribution was close to normal (see Figure 1B), we also report the frequentist and default Bayesian preregistered paired samples *t*-tests.

Indicative of a TBR effect, participants judged repeated statements as more true (i.e., less false) than new statements (see Figure 1; see Figure 2 for results by statements). This result is apparent in both a Wilcoxon signed-rank test, $Z = 4.65$, $p = 3.344\text{e-}06$, $r_{\text{rb}} = .216$, and a paired samples *t*-test, $t(231) = -4.49$, $p = 1.14\text{e-}5$, $d = 0.278$, $95\% \text{CI}_d = [0.154; 0.402]$. The default Bayesian *t*-test yielded substantial evidence for the TBR effect, $\text{BF}_{10} = 1023.89$. Repeated ($M = 8.79$; $SD = 9.22$; Median = 6.25) and new ($M = 6.43$; $SD = 7.48$; Median = 3.31) statements were all judged as false (as evidenced by ratings below 50, the mid-point of the scale) – so repetition

did not per se increase “true” judgments (as no statements were rated as true), but made statements judged as slightly less false than new ones².

3.2 Not preregistered analyses looking at subsets of participants.

We computed a TBR effect score for each participant (the difference between truth judgment scores on repeated and new statements). In doing so, we were able to identify which participants descriptively showed larger truth judgment scores on repeated vs. new statements (positive scores), larger truth judgment scores on new vs. repeated statements (negative scores), and no difference between repetition conditions (null scores). The results computed on the total sample are driven by a subset of participants as only 52.58% ($n = 122$) had a positive TBR effect score ($M = 7.6$, $SD = 6.36$), 19.4% ($n = 45$) had a null truth effect score, and interestingly 28.02% ($n = 65$) had a *negative* TBR effect score ($M = -5.82$; $SD = 5.67$), the latter of which is significantly different from 0, $t(64) = -8.28$, $p = 1.05\text{e-}11$, $d = 1.027$. Consistent with recent findings (see below), these complementary analyses suggest that all participants are not equally prone to the TBR effect.

² In a non-preregistered 2 (Repetition: New; Repeated five times – within-participants) $\times$ 2 (Counterbalancing: Set A repeated and Set B new vs. Set B repeated and Set A new – between participants) mixed ANOVA, we did not find a significant main ($F(1, 230) = 0.01$, $p = .927$, $\eta^2_G < .001$) or interactive ($F(1, 230) = 0.64$, $p = .425$, $\eta^2_G < .001$) effect of the Counterbalancing condition.

Figure 1. Truth judgments as a function of Repetition (dashed horizontal line: no bias toward the "False" or the "True" side) (A) and TBR effect scores (dashed horizontal line: no difference between truth judgments scores on repeated and new statements; positive scores: higher truth judgments for repeated than new statements; negative scores: higher truth judgments for new than repeated statements) (B). The dots are the participants' scores (jittered). The lower and upper limits of the boxplots are the 95% confidence intervals, with the mean in between. The distributions represent the kernel probability density of the data.

Figure 2. Mean truth judgments as a function of Statements and Repetition. The error bars are the 95% confidence intervals. The statements were abbreviated for the sake of readability; the complete statements are available at <https://osf.io/pe4g9/>.

4. Discussion

Using five repetitions and a 100-point judgment scale allowing for finer-grained gradients in perceived falsehood, we found a TBR effect on highly implausible statements: participants started giving credence to statements as highly implausible as “Earth is a perfect square” or “Benjamin Franklin lived 150 years” after repeating them for just five times. Beyond its practical significance, this finding informs current models of the TBR effect. Contrary to the “truth ambiguity” assumption endorsed in knowledge-conditional multinomial processing tree models.

Fazio et al.’s (2019) model posits that repetition increases the perceived truth of statements at all gradients of plausibility. Although Fazio et al. found results consistent with this model, their empirical evidence was only indirect. These authors reported an inverted U-shaped relation between the size of the truth effect and the statements’ level of plausibility. However, they did not find a TBR effect at the extremes of the distribution (i.e., with either highly implausible or highly plausible statements; see Table 1 in Fazio et al., 2019).

Because we used five repetitions and a 100-point scale, our study simultaneously deviates in two ways from usual TBR procedures where one repetition is implemented, and dichotomous ratings/Likert scales are typically used. Therefore, we cannot say which feature (or a combination thereof) is responsible for the current findings. There are reasons to think that increasing the number of repetitions from one to five increased the likelihood to find a TBR effect, as suggested by research on ambiguous statements (e.g., Fazio et al., 2021; Hassan & Barber, 2021). Likewise, using a sensitive scale increased the probability of detecting small differences of perceived truth between repeated and new statements that may be difficult to find with less sensitive dichotomous ratings or Likert scales. Future research could orthogonally vary the number of

repetitions and the truth judgment scale to quantify whether and how much each feature contributes to the TBR effect with implausible statements.

Future research may also clarify the reason for interindividual variations in the TBR effect. Almost half of the participants had a null TBR score or even showed *lower* perceived truth on repeated than new statements. The variability between participants in the TBR effect was greater here than the one estimated by Schnuerch et al. (2020) using Bayesian analyses. In their study, about 20% of the participants did not demonstrate the typical TBR effect (see also Henderson et al., 2021, who found that 14.8% of their sample did not show a TBR effect). We can speculate on two non-exclusive explanations: the link between fluency and truth judgments and demand characteristics.

First, processing fluency may not always be interpreted as a cue for truth. The way in which fluency is interpreted can be reversed (Unkelbach, 2007; see also Silva et al., 2016). Particularly relevant here is a recent study by Corneille et al. (2020) showing that in judgment ecologies where misinformation is widespread (i.e., social media), a fakeness-by-repetition effect is found, whereby perception of fakeness increases after a statement repetition. Because only highly implausible statements were used in the present experiment, it may be that some participants relied on fluency as a proxy for falsehood.

Second, participants can act in ways that confirm or refute an experimenter's hypothesis (Orne, 1962; for an integrative model, see Corneille & Lush, 2021): if participants identified that we expected repeated statements to be judged as more true than new ones, they were able to respond in line or against the hypothesis, accounting for individual differences in the observed TBR effect. This could have been facilitated in our procedure, as recognition memory was likely to be high (five repetitions: truth judgment task right after the exposure phase). This issue is

alleviated, however, by the fact that we asked participants to judge their interest in statements in the exposure phase: participants may have mistakenly believed that we were interested in the link between perceived interest and perceived truth. Future research could assess whether the TBR effect with implausible statements depends on hypothesis awareness, which could account for individual differences in the demonstrated effect.

5. Conclusion

We found that even a very limited number of repetitions can increase the perceived truth of highly implausible statements. Complementary analyses reveal that this is not equally true for everyone, though: some individuals may be left unaffected by repetition and some others may even show decreased perceived truth following repetition. Better understanding when, whom, and why one may develop a sense of truth for implausible statements following their repetition is of both high theoretical and practical interest.

Acknowledgments

We thank Douglas Bancu, Olivier Desmedt, Antoine Mauclet, and Prof. Christian Unkelbach for their valuable comments during the preparation of the manuscript.

References

- Arkes, H. R., Boehm, L. E., & Xu, G. (1991). Determinants of judged validity. *Journal of Experimental Social Psychology*, 27(6), 576605. [https://doi.org/10.1016/0022-1031\(91\)90026-3](https://doi.org/10.1016/0022-1031(91)90026-3)
- Aust, F., Diedenhofen, B., Ullrich, S., & Musch, J. (2013). Seriousness checks are useful to improve data validity in online research. *Behavior Research Methods*, 45(2), 527535. <https://doi.org/10.3758/s13428-012-0265-2>
- Bacon, F. T. (1979). Credibility of repeated statements: Memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory*, 5(3), 241252. <https://doi.org/10.1037/0278-7393.5.3.241>
- Brashier, N. M., Eliseev, E. D., & Marsh, E. J. (2020). An initial accuracy focus prevents illusory truth. *Cognition*, 194, 104054. <https://10.1016/j.cognition.2019.104054>
- Corneille, O., & Lush, P. (2021). Sixty years after Orne's American Psychologist article: A conceptual analysis of "Demand Characteristics". <https://doi.org/10.31234/osf.io/jqyvz>
- Corneille, O., Mierop, A., & Unkelbach, C. (2020). Repetition increases both the perceived truth and fakeness of information: An ecological account. *Cognition*, 205, 104470. <https://doi.org/10.1016/j.cognition.2020.104470>

- 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), 238257. <https://doi.org/10.1177/1088868309352251>
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., ... Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3), 554–559. <https://10.1073/pnas.1517441113>
- Difonzo, N., Beckstead, J., Stupak, N., & Walders, K. (2016). Validity judgments of rumors heard multiple times: The shape of the truth effect. *Social Influence*, 11, 118. <https://doi.org/10.1080/15534510.2015.1137224>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175191. <https://doi.org/10.3758/BF03193146>
- Fazio, L. K. (2020). Repetition Increases Perceived Truth Even for Known Falsehoods. *Collabra: Psychology*, 6(38). <https://doi.org/10.1525/collabra.347>
- Fazio, L., Pillai, R., & Patel, D. (2021). The effects of repetition on belief in naturalistic settings. <https://doi.org/10.31234/osf.io/r85mw>
- 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), 9931002. <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), 17051710. <https://doi.org/10.3758/s13423-019-01651-4>

Fazio, L., Pillai, R., & Patel, D. (2021). *The effects of repetition on belief in naturalistic settings*.

PsyArXiv. <https://doi.org/10.31234/osf.io/r85mw>

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. <https://doi.org/10.1186/s41235-021-00301-5>

Henderson, E. L., Simons, D. J., & Barr, D. J. (2021). The trajectory of truth: A longitudinal study of the illusory truth effect. *Journal of cognition*, 4(1), 1-23.

<https://doi.org/10.5334/joc.161>

Henninger, F., Shevchenko, Y., Mertens, U., Kieslich, P. J., & Hilbig, B. E. (2019). *lab.js: A free, open, online study builder*. PsyArXiv. <https://doi.org/10.31234/osf.io/fqr49>

Hilbig, B. E. (2012). How framing statistical statements affects subjective veracity: Validation and application of a multinomial model for judgments of truth. *Cognition*, 125, 37-48.

<https://doi.org/10.1016/j.cognition.2012.06.009>

Lange, K., Kühn, S., & Filevich, E. (2015). Correction: “Just Another Tool for Online Studies” (JATOS): An Easy Solution for Setup and Management of Web Servers Supporting Online Studies. *PLOS ONE*, 10(7), e0134073.

<https://doi.org/10.1371/journal.pone.0134073>

Marsh, E. J., & Rajaram, S. (2019). The Digital Expansion of the Mind: Implications of Internet Usage for Memory and Cognition. *Journal of Applied Research in Memory and*

Cognition, 8(1), 114. <https://doi.org/10.1016/j.jarmac.2018.11.001>

Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, 17(11), 776–783. <https://10.1037/h0043424>

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. <https://10.1037/xge0000465>

Reber, R., & Schwarz, N. (1999). Effects of Perceptual Fluency on Judgments of Truth. *Consciousness and Cognition*, 8(3), 338342. <https://doi.org/10.1006/ccog.1999.0386>

Schnuerch, M., Nadarevic, L., & Rouder, J. N. (2020). The truth revisited: Bayesian analysis of individual differences in the truth effect. *Psychonomic Bulletin & Review*. Advance online publication. <https://doi.org/10.3758/s13423-020-01814-8>

Schwartz, M. (1982). Repetition and rated truth value of statements. *American Journal of Psychology*, 95(3), 393-407. <https://doi.org/10.2307/1422132>

Silva, R. R., Garcia-marques, T., & Mello, J. (2016). The differential effects of fluency due to repetition and fluency due to color contrast on judgments of truth. *Psychological Research*, 80(5), 821837. <https://doi.org/10.1007/s00426-015-0692-7>

Unkelbach, C. (2007). Reversing the Truth Effect: Learning the Interpretation of Processing Fluency in Judgments of Truth. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 33(1), 219230. <https://doi.org/10.1037/0278-7393.33.1.219>

Unkelbach, C., & Greifeneder, R. (2018). Experiential fluency and declarative advice jointly inform judgments of truth. *Journal of Experimental Social Psychology*, 79, 78–86. <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. <https://doi.org/10.1177/0963721419827854>
- Unkelbach, C., & Speckmann, F. (2021). Mere repetition increases belief in factually true COVID-19 related information. *Journal of Applied Research in Memory and Cognition*, Advance online publication. <https://doi.org/10.1016/j.jarmac.2021.02.001>
- Unkelbach, C., & Stahl, C. (2009). A multinomial modeling approach to dissociate different components of the truth effect. *Consciousness and Cognition*, 18(1), 223-238. <https://doi.org/10.1016/j.concog.2008.09.006>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146-1151. <https://doi.org/10.1126/science.aap9559>