



HAL
open science

**I know It's false, but I keep thinking as if it were true:
A replication study of Johnson and Seifert's (1994)
continued influence effect**

Victor Laurent, Thierry Kosinski, Stephane Rusinek

► **To cite this version:**

Victor Laurent, Thierry Kosinski, Stephane Rusinek. I know It's false, but I keep thinking as if it were true: A replication study of Johnson and Seifert's (1994) continued influence effect. *Acta Psychologica*, 2023, *Acta Psychologica*, 241, pp.104094. 10.1016/j.actpsy.2023.104094 . hal-04312002

HAL Id: hal-04312002

<https://hal.univ-lille.fr/hal-04312002>


Submitted on 28 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0
International License



I know It's false, but I keep thinking as if it were true: A replication study of Johnson and Seifert's (1994) continued influence effect

Victor Laurent^{*}, Thierry Kosinski, Stéphane Rusinek

Univ. Lille, ULR 4072 – PSITEC – Psychologie: Interactions, Temps, Emotions, Cognition, F-59000 Lille, France

ARTICLE INFO

Keywords:

Continued Influence Effect
False information
Replication

ABSTRACT

False Information (FI) is a critical societal issue, made even more pressing by our inability to mitigate its influence through correction. Researchers Johnson and Seifert (1994, Experiment 1A) penned a seminal paper on this “Continued Influence Effect” (CIE), which they observed in English-speaking participants. In their experiments, one group read a text containing FI that was later retracted (FI group), while another read the same text without FI (control group). Interestingly, even after receiving corrections, participants who read the FI were more likely than their peers to form FI-related inferences about the text. To the best of our knowledge, this finding has never been successfully directly replicated. Given the current replicability crisis plaguing the human sciences, the influence of culture on CIE and the importance of Johnson and Seifert's paradigm in this literature, the reassessment of their findings within a non-English-speaking population appears crucial. The present research investigated the direct replicability of their study with a French-speaking sample, comparing the inferences drawn by an FI group ($n = 21$) to those made by a control group ($n = 23$). The results confirm those of the original study, supporting the validity of Johnson and Seifert's paradigm (1994) and extending its applicability to a French-speaking population.

1. Introduction

False Information (FI) – defined here as inaccurate information presented as true – is a pervasive and vividly topical societal issue. Given the widespread use of the internet (Bronner, 2011; Lazer et al., 2018), FI has come to influence numerous aspects of public life including health (Sénéchal et al., 2004), the environment (Drummond et al., 2020), justice (Black, 2018), politics (Vosoughi et al., 2018), military affairs (Lewandowsky et al., 2009) and the economy (Cavazos, 2019). For instance, exposure to misleading climate change statistics can deter belief in the existence of climate change (Ranney & Clark, 2016), subsequently diminishing pro-environmental behavior (van der Linden, 2015). Similarly, exposure to FI discrediting a politician can increase negative attitudes towards both the politician and their party (Dobber et al., 2021).

One common strategy to counteract the influence of FI involves its correction, a topic which has garnered recent academic interest (Courchesne et al., 2021). From 2014 to 2018, the number of fact-checking organizations reportedly tripled worldwide. Institutions such

as the UN (<https://shareverified.com>), NATO (NATO, 2020, 2022) and the European Union (<https://euvsdisinfo.eu/>) have also recently developed initiatives to halt the spread of FI. Similarly, social networks such as Facebook have recently implemented systems to notify users when they are exposed to misleading information (Luengo & García-Marín, 2020).

However, ample research asserts that the correction of FI fails to completely negate its influence (Lewandowsky et al., 2012; Walter & Murphy, 2018; Walter & Tukachinsky, 2020). The term “Continued Influence Effect” (CIE) (Johnson & Seifert, 1994) describes situations in which misinformation persists even after its correction. The CIE might, in part, be explained by a cognitive competition between the FI and the correction (Ecker et al., 2022; Ecker et al., 2011). Even after retraction, the FI may persist in memory, remaining automatically activable when triggered by certain environmental cues. Therefore, individuals must proactively retrieve the correction from memory to counteract its influence. The CIE may manifest when the correction is not retrieved or when its activation is insufficient to suppress that of the FI (Kendeou et al., 2014).

^{*} Corresponding author at: Université de Lille, Domaine Universitaire du Pont-de-Bois, laboratoire PSITEC (ULR 4072), Rue du Barreau, BP 60149, 59653 Villeneuve d'Ascq, France.

E-mail addresses: victor.laurent2.etu@univ-lille.fr (V. Laurent), thierry.kosinski@univ-lille.fr (T. Kosinski), stephane.rusinek@univ-lille.fr (S. Rusinek).

The CIE has been observed across various disciplines and contexts. For example, Johnson and Seifert (1994) observed CIE in five experiments involving American psychology students. The sessions were run in groups, with participants receiving course credits for their participation. Provided with a booklet containing instructions and several page-length “messages”, participants were instructed to read at their own pace, as if reading a newspaper, and informed that they would be asked to recall elements from the text. In experiment 1A (Johnson & Seifert, 1994, Exp. 1A), the narrative message, featuring content derived from Wilkes and Leatherbarrow’s (1988) study, recounted the investigation of a warehouse fire. Participants were informed that the fire may have resulted from a short circuit in a closet’s wiring. Some subjects (FI group) were shown an additional message suggesting that volatile flammable materials had been carelessly stored in the closet. This FI was subsequently rectified by yet another message clarifying that the closet had been empty before the fire broke out. Participants in the control group received the true information (i.e., the closet was empty) straight away without exposure to the FI. Participants were then administered open-ended questions about their recall of the text and the inferences they made about the event. Despite comparable levels of recall between the groups, participants in the FI group tended to reference the FI more directly in their text explanations than their peers in the control group (Johnson & Seifert, 1994, Exp. 1A). The effect size was large ($d = 1.521$) (Lee, 2016; Maher et al., 2013).

A sizeable body of CIE literature has built upon the foundational work of Johnson and Seifert (1994). As far as we are aware, their research was the first to offer a detailed description of this effect. Their research is widely acknowledged and highly cited, with many studies adopting a similar experimental paradigm, albeit with some modifications (Brydges et al., 2018; Ecker et al., 2014, 2017; Kan et al., 2021; O’Rear & Radvansky, 2020; Sanderson et al., 2022; Xu et al., 2020). For instance, Xu et al. (2020), in comparison to Johnson and Seifert (1994), employed different texts and recruited participants from different cultures. Despite these deviations from the original protocol, the CIE persisted, although with a less pronounced effect size. This substantial body of literature (for a meta-analysis, see Walter & Tukachinsky, 2020) supports the conceptual replicability of this effect and findings of the Johnson and Seifert (1994).

A conceptual replication is the reevaluation of a hypothesis or a result from previous research using different methods (Romero, 2019). Conceptual replications are essential for assessing the validity of an effect since they examine its applicability in various situations while providing multiple evidences for its occurrence. They also play a vital role in identifying the variables that contribute to effect variations. For instance, a stronger manifestation of this effect has been observed in a Chinese population as compared to an Australian one (Xu et al., 2020), suggesting that the CIE may not manifest uniformly across cultures. However, as noted by Pashler and Harris (2012), one limitation of conceptual replication is its inability to effectively challenge the validity of an effect in the event of replication failure. Indeed, due to changes in the methodology, any differences in the results can be attributed to the methodological adjustments rather than the effect itself.

An additional method for evaluating the validity of an effect is to conduct direct (or close) replications. The fundamental aim of a direct replication is to recreate an experimental procedure to determine whether the initial findings can be reproduced. This goal is put into practice by striving to reproduce the initial protocol as closely as possible (Brandt et al., 2014), while acknowledging that it is practically impossible for an entire procedure to be strictly replicated (i.e., the location, the timing, and the participants are inevitably different) (Romero, 2019; Simons, 2014). Thus, depending on extent and number of modified elements, a replication may fall at varying points along a continuum ranging from “close” to “conceptual” (Brandt et al., 2014).

Close replications offer a significant advantage as they inherently involve only minor methodological modifications compared to the original protocol. Consequently, in contrast to conceptual replications, a

failed close replication has the potential to effectively challenge the validity of an effect (Pashler & Harris, 2012), whereas a successful one may provide supporting evidence (Simons, 2014). Recent extensive literature has delved into the close replicability of various effects but has encountered numerous failures in replication (Pashler & Wagenmakers, 2012). For instance, an investigation into the replicability of nearly a hundred psychology studies (Open Science Collaboration, 2012, 2015) found that only 35 of 97 replication attempts successfully reproduced the original findings. The current prevalence of false positives emphasizes the significance of studying replication.

Given its significance in the CIE literature, Johnson and Seifert’s (1994) study, including the previously described Experiment 1A, serves as a suitable work for assessing the close replicability of this effect. To the best of our knowledge, two experiments offer insights into this question. Although their research didn’t aim to assess CIE replicability, Ecker et al. (2011, Exp. 1) exposed participants to an adaptation of the warehouse fire scenario and queried them with open-ended inferential questions to evaluate whether repeating a correction could reduce the CIE. In one experimental condition, FI was initially provided and then retracted, while a control group did not receive any FI. A comparison of these groups revealed evidence of a CIE. This result supports the close replicability of Johnson and Seifert’s (Exp 1A, 1994) findings, even though the adaptations to the scenario and measures have not been reported. A second paper providing information about the close replicability of the Johnson and Seifert’s (1994) findings is that of Connor Desai and Reimers (2019). In this study, the authors conducted a computerized version of their paradigm in four experiments (Exp. 1A, 1B, 2A, 2B). Specifically, in experiments 2A and 2B, they replicated the scenario, measures’ content and experimental design from Johnson and Seifert’s original protocol (1994). One group received FI that was later corrected, while a control group received the correction without the initial FI. However, instead of simply stating that the closet was empty, the correction for the control group explicitly mentioned that “*a closet reportedly containing cans of oil paint and gas cylinders had actually been empty before the fire.*” The results showed that participants in control and FI conditions referred to the FI in the same proportions. Despite this, the authors interpreted these findings as indicative of the CIE, suggesting that the presence of FI-related inferences demonstrated the latter’s influence. Their argument was that participants could have produced FI-related inferences due to exposure to its content – oil paint and gas cylinders. Furthermore, subjects in the FI group may have mirrored the control group in referencing the FI due to a failure to retrieve the correct information from memory (Connor Desai & Reimers, 2019).

In spite of the significance of the presented studies, they do not enable the close replicability of Johnson and Seifert’s (1994) findings to be fully determined. Indeed, there is room to question the claim that Connor Desai and Reimers’ (Exp. 2A, 2B, 2019) direct replication was successful. The mere presence of FI-related inferences does not conclusively prove FI’s influence; participants may spontaneously refer to it even without prior exposure. Additionally, the lack of differences between the two groups could justify arguments that the correction was successful in eliminating the influence of FI, thereby casting doubt on the replication attempt. Furthermore, although Ecker et al. (2011) demonstrated a CIE within their study, the limited information about scenario and measure modifications hinders a precise assessment of how closely their replication aligns with the original paradigm.

The state of the literature regarding the close replicability of the CIE, the current trend of false positives and the cornerstone role of Johnson and Seifert’s (1994) paradigm in the CIE literature together underlie the necessity of studying the replicability of their results. In line with this aim, the present research sought to examine the replicability of their findings. We exposed two participant groups to the original text used by the aforementioned authors. One group engaged with a version of the scenario containing the FI, which was later corrected (FI group), while the other group read a version devoid of FI (control group). Participants’ recall of the text and the inferences they generated about the event were

then measured. We opted for an online, computerized and individual paradigm, as opposed to a paper-based group format. Furthermore, we opted to use closed-ended questions rather than open-ended questions in this study to reduce participant dropout (Connor Desai & Reimers, 2019). Finally, the current study recruited French-speaking participants to examine the validity of this effect within a distinct cultural context. While this modification may introduce an aspect of ‘conceptual replication’ into our close replication attempt, as suggested by previous studies that indicate culture’s potential influence on the CIE (Walter & Murphy, 2018; Xu et al., 2020), we deem it of particular relevance. This relevance stems from the fact that research on the fight against misinformation has predominantly focused on American populations (Courchesne et al., 2021) and because, to the best of our knowledge, the CIE has never been explored with a French-speaking population.

2. Method

The study carried out adheres to the American Psychological Association’s Ethical Principles in the Conduct of Research with Human Participants and the Helsinki declaration. It also follows the recommendations for best practices in conducting a replication provided by Brandt et al. (2014). Also, as suggested by Simmons et al. (2012), “we have reported how we determined our sample size, all data exclusions (if any), all manipulations, and all measures in the study.”

2.1. Materials

All materials are provided in the [Supplementary material](#).

2.1.1. Text

The text employed a French translation of the text originally used by Johnson and Seifert (1994). The translation was conducted by two researchers, one of whom was bilingual. The text consisted of 15 messages written in black against a white background. Each message was displayed individually and sequentially on the screen. The text could be either an FI or a control version. In the FI version, the sixth message contained a piece of information that was subsequently corrected in the thirteenth message (i.e., message 6: “[...] they have reports that cans of oil paint and pressurized gas cylinders had been present in the closet before the fire.”; message 13: “[...] A second message received from Police Investigator Lucas [...] stated that the closet reportedly containing cans of oil paint and gas cylinders had actually been empty before the fire.”). Thus, the correction was provided by a credible source and was unambiguous. On the other hand, the control version replaced the FI with the correct information and the correction with neutral information (i.e., message 6: “[...] the closet was empty [...]”; message 13: “[...] the two firefighters taken to the hospital had been released.”). The remaining thirteen messages were the same in both versions.

2.1.2. Measures

The design for the questions, the corresponding answers and the scoring methodology were based on the work of Connor Desai and Reimers (2019).

2.1.2.1. Recall questions. Participants were presented with eight recall questions. For each question, four answer options were provided, and participants had to select the one they believed to be correct. Six of the questions asked about general facts mentioned in the text (e.g., “What business was the firm in?”), while two questions specifically focused on the information that differed between the groups (e.g., “What was the point of the second message of the police?”). The correct answer to the latter type of questions depended on which version of the text the participant read (e.g., FI group: “The closet [...] was empty.”; control group: “The two firefighters taken to the hospital had been released.”). Each correct response was assigned one point. Therefore, the maximum score

for the recall questions was eight.

2.1.2.2. Inference questions. Six inference questions probed the participants’ interpretations of the text. For each question, (e.g., “Where was the probable location of the explosions?”), four inference responses were provided. Among the four answers, one option corresponded to an FI-related inference (e.g., “The storage closet.”). Participants were instructed to distribute a total of ten points across the four answer options based on the plausibility of each. Therefore, the more points a participant assigned to the FI-related response, the more they considered the FI to be a possible explanation for the text.

2.2. Procedure

The procedure was computerized and supported by the web version of the Inquisit software. Participants were recruited through social networks or posters and accessed the study via an internet link. The task began with an informational letter outlining the study’s focus on text-based reasoning. Participants were told that they would read a text about a fire and then answer questions related to it. If they agreed to proceed, participants were requested to disclose their gender, age and level of education. Following this, a page of instructions prompted them to read a text the way they would a newspaper article and answer questions about it, emphasizing that they were not allowed to backtrack in their reading. The Inquisit software randomly assigned the text version to be read. The control group read by the control version, and the FI group the FI version. Participants read the texts at their own pace, moving to the next text segment by pressing the spacebar. Once they finished reading, participants answered inference questions followed by recall questions. A debriefing text revealed the experiment’s true purpose, after which participants were asked for consent to include their data in the study. Participants who declined were excluded from the data analysis.

2.3. Participants

The sample size was determined a priori by averaging the number of participants recruited in Johnson and Seifert’s (1994) five experiments ($n = 49$). A power analysis was conducted using G*Power software (Faul et al., 2009), with $\alpha = 0.05$ and $1 - \beta = 0.80$. The analysis suggested that detecting a difference between two groups (two-sided Student *t*-test) with a large effect size ($d = 1.521$) (Johnson & Seifert, 1994, study 1A) would require a sample size of eight participants per condition.

A total of 52 participants were recruited and consented to have their data collected. The only inclusion criterion was that participants had to be over 18 years old. Eight participants were excluded from the data analysis due to a failure to follow the instructions (i.e., the number of points assigned to the inference questions did not consistently add up to ten).

The final sample size was 44 participants (30 female, 1 other), randomly assigned to either the FI ($n = 21$) or the control group ($n = 23$). The mean age was 23.63 years old ($SD = 9.49$). A sensitivity power analysis using G*Power software (Faul et al., 2009) reported that this sample size provided 80 % power to detect an effect size of $d = 0.397$, at a 5 % false-positive rate, in an independent-samples two-tailed *t*-test.

3. Results

The complete dataset is available in the [Supplementary material](#). For data analysis, parametric tests were used when their conditions were satisfied; otherwise, non-parametric tests were applied.

3.1. Recall questions

Participants were awarded one point for selecting the correct answer

and zero points for incorrect responses. The total points for recall questions were compared between the two groups. Due to the non-normal distribution of scores in the control group ($W = 0.919$, $p = 0.062$), a Mann-Whitney test was employed. No significant differences were noted in the recall question scores between the FI group ($M = 5.333$; $SD = 1.528$) and the control group ($M = 5.217$; $SD = 1.783$) ($U = 237$, $p = 0.924$).

3.2. Inference questions

The total points assigned to FI-related answers were compared between the two groups. Data are shown in Fig. 1. The FI group recorded significantly higher scores than the control group ($t(42) = 2.28$, $p = 0.028$), implying that participants exposed to false information assigned more points to FI-related inferences ($M = 21.667$; $SD = 9.525$) compared to those who were not ($M = 15.783$; $SD = 7.556$). The effect size ($d = 0.688$; 95 % CI = [0.075, 1.294]) was moderate (Lee, 2016) and notably lower than that observed in Johnson and Seifert's study (1994, study 1A, $d = 1.521$).

4. Discussion

The primary objective of the present study was to examine the replicability of Johnson and Seifert's (1994) findings. Two groups of French-speaking participants were exposed to a computerized version of their paradigm. Despite similar recall between the groups and the reception of a credible correction, the participants who received the FI version were more inclined to leverage the false information to interpret the text, compared to those exposed to the control version. This observation indicates the presence of a CIE, i.e., persistent belief in FI despite correction. Such evidence for the replication of Johnson and Seifert's (1994) results in a non-American population holds significance for several reasons.

Firstly, considering the central role of the Johnson and Seifert's (1994) paradigm in the CIE literature, the successful replication of their findings bolsters support for the validity of this literature and the paradigm that undergirds it. Secondly, our research suggests the existence of a CIE among French-speaking individuals who encounter FI on the internet. This is significant given previous research highlighting

cultural factors as influencers of this effect (Walter & Murphy, 2018; Xu et al., 2020). Finally, the successful replication of Johnson and Seifert's (1994) findings lends credence to arguments that their results are not mere false positives. This successful replication is especially noteworthy in an era in which researchers grapple with issues of replication (Pashler & Wagenmakers, 2012), and given that previous literature presented mixed evidence on this topic (Connor Desai & Reimers, 2019). Some minor modifications in the methodology employed by Connor Desai and Reimers (2019) may explain these divergent findings. Specifically, in the present research and the original study (Johnson & Seifert, 1994), participants read each message at their own pace. In contrast, Connor Desai and Reimers (2019) set a ten-second minimum reading time, which could have prompted participants to contemplate the information, including the correction, with greater and more sustained attention. This change may have facilitated a more qualitative processing of the correction, thereby reducing the influence of the FI. Another potential explanation might relate to the handling of responses from participants who did not follow the instructions. In the current experiment, such data were excluded, whereas the authors of the previous study did not articulate their approach to this issue. The potential inclusion of such data could have influenced the results.

In practical terms, this successful replication of Johnson and Seifert's (1994) serves as a reminder that simply correcting FI may not fully mitigate its influence. Notably, the present study suggests that efforts to combat FI in French-speaking countries should take into account that correction alone might not be sufficient to eliminate their influence. For practitioners seeking to combat FI, two strategies are worth exploring. The first involves improving the quality of corrections in order to reduce the CIE. This approach draws upon a substantial body of literature that has explored several efficient strategies, such as repeating the correction (Ecker et al., 2011) or providing an alternative explanation to the facts that initially conveyed the FI (Johnson & Seifert, 1994). For example, in the present experiment, the correction (i.e., "[...] *the closet reportedly containing cans of oil paint and gas cylinders had actually been empty before the fire*") was a mere retraction provided only once. The potential for reducing the CIE could have been enhanced by repeating the correction and offering an alternative explanation to the described facts (i.e., "we know the fire was not due to negligence but was the result of a criminal act"). A second approach to addressing the societally relevant issue of FI




Fig. 1. Total points assigned to FI-related answers across conditions. Error bars show standard errors.

involves exploring alternative strategies other than correction. For instance, future research could focus on investigating novel technological solutions (e.g., Molina et al., 2021) or further developing studies on critical thinking (e.g., Roozenbeek et al., 2020) and inoculation theory (e.g., Roozenbeek et al., 2020).

One limitation of this experiment could be attributed to the relatively modest sample size. A larger sample size might have strengthened the data's validity. However, we consider this limitation acceptable since the sample size was determined through a power analysis, mirroring the sample used in the initial study, and yielding an effect size that approximates what is commonly observed in the CIE literature (Walter & Tukachinsky, 2020). These arguments support the validity of the results.

Another limitation in this replication's success could be its lack of perfect fidelity to the original study. For instance, the experiment was conducted individually over the internet. Considering the internet's influential role in the proliferation of FI (Bronner, 2011; Lazer et al., 2018), conducting the study online seemed preferable for enhancing ecological validity. However, although previous studies have indicated that these elements did not seem to hinder the CIE and reported similar effect size (e.g., Ecker et al., 2020; Susmann & Wegener, 2021; Xu et al., 2020), we cannot definitively conclude that they had no impact on this effect. In a similar vein, we used closed-ended questions, relying on the work of Connor Desai and Reimers (2019), in contrast to the initial experiment which employed open-ended questions. Several studies have demonstrated a CIE with closed-ended inferential questions, suggesting that this format does not impede this effect (e.g., McIlhiney et al., 2022; Sanderson et al., 2022). However, despite the questions we used being an adaptation of their original open-ended version, we cannot precisely gauge the extent to which this modification influenced the CIE. Other minor differences from the original study are that not all participants were psychology students, nor did they receive course credits for their participation. Despite these changes, most important aspects of the studies, such as the text, the procedure, the content of the measures and the instructions, remained similar. Given the centrality of these elements to the paradigm, it seems reasonable to state that the present work provides a direct or close replication (Brandt et al., 2014; Open Science Collaboration, 2015; Romero, 2019) of the original study.

Finally, it's worth noting that the moderate effect size observed in our study was smaller than the large effect size reported in the initial study by Johnson and Seifert (Exp 1A, 1994). This suggests that the CIE may have a lesser impact than indicated by their findings. This aligns with the broader CIE literature, which generally indicates a small effect size that can vary depending on multiple variables (Walter & Tukachinsky, 2020). In our case, factors such as using closed-ended inferential questions instead of open-ended ones, employing computerized administration rather than a paper-based method, and the inclusion of French-speaking participants may have contributed to the reduction in the effect size.

Thus, this replication provides support for the validity of the CIE and its implications in a French-speaking population. In conclusion, it is crucial to emphasize that the persistence of FI after correction does not imply that the provision of corrections is an entirely futile act. A recent meta-analysis examining the effect of corrections found a moderate and significant decrease in belief in corrected information (see Walter & Murphy, 2018). Additionally, as mentioned before, several strategies which increase the quality of the correction have already been identified (Ecker et al., 2011; Johnson & Seifert, 1994). Therefore, while additional approaches are necessary to further mitigate the influence of FI in society (see Bronner et al., 2022; Treen et al., 2020), scholars must continue to research and identify the conditions in which corrections are most effective.

Open practices

Data and materials are provided in the [Supplementary Material](#).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data are provided in the Supplementary material.

Acknowledgements

We thank Claire Hofer and Brian Stacy for English corrections, Kristin Laurin for her comments on the first version of this manuscript, Tessa Loonis, Vlad-Alin Rusu, Younes Yagoubi, Ilya Veshchunov, and Fanny Waigriez for help with data collection, and the people who volunteered to participate in this study.

Funding

This article has been supported by the European center for humanities and social sciences (MESHS-Lille, France).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.actpsy.2023.104094>.

References

- Black, S. (2018). The continued influence of retracted witness testimony at jury trials. *Asia Pacific Journal of Advanced Business and Social Studies*, 4(1), 283–290. <https://doi.org/10.25275/apjabssv4i1ss3>
- Brandt, M. J., IJzerman, H., Dijksterhuis, A., Farach, F. J., Geller, J., Giner-Sorolla, R., ... van 't Veer, A. (2014). The Replication Recipe: What makes for a convincing replication? *Journal of Experimental Social Psychology*, 50, 217–224. <https://doi.org/10.1016/j.jesp.2013.10.005>
- Bronner, G. (2011). Ce qu'Internet fait à la diffusion des croyances. *Revue européenne des sciences sociales. European Journal of Social Sciences*, 49-1, Article, Article 49-1. <https://doi.org/10.4000/ress.805>
- Bronner, G., Cayrol, R., Cordonier, L., Douzet, F., Farinella, R.-M., Favre, A., Garrigues, J., Harfoush, R., Khan, R., Muxel, A., Reichstadt, R., Roder, I., Warusfel, B., & Wiewiorka, A. (2022). Les Lumières à l'ère numérique. <https://www.elysee.fr/admin/upload/default/0001/12/0f50f46f0941569e780ffc456e2f6aa5c59a9e3b7.pdf>.
- Brydges, C. R., Gignac, G. E., & Ecker, U. K. H. (2018). Working memory capacity, short-term memory capacity, and the continued influence effect: A latent-variable analysis. *Intelligence*, 69, 117–122. <https://doi.org/10.1016/j.intell.2018.03.009>
- Cavazos, R. (2019). The economic cost of bad actors on the Internet. CHEQ. https://info.cheq.ai/hubfs/Research/the_economic_cost_fake_news_final.pdf.
- Connor Desai, S., & Reimers, S. (2019). Comparing the use of open and closed questions for web-based measures of the continued-influence effect. *Behavior Research Methods*, 51(3), 1426–1440. <https://doi.org/10.3758/s13428-018-1066-z>
- Courchesne, L., Ilhardt, J., & Shapiro, J. N. (2021). Review of social science research on the impact of countermeasures against influence operations. *Harvard Kennedy School Misinformation Review*. <https://doi.org/10.37016/mr-2020-79>
- Dobber, T., Metoui, N., Trilling, D., Helberger, N., & de Vreese, C. (2021). Do (microtargeted) deepfakes have real effects on political attitudes? *The International Journal of Press/Politics*, 26(1), 69–91. <https://doi.org/10.1177/1940161220944364>
- Drummond, C., Siegrist, M., & Árvai, J. (2020). Limited effects of exposure to fake news about climate change. *Environmental Research Communications*, 2(8), Article 081003. <https://doi.org/10.1088/2515-7620/abae77>
- Ecker, U. K. H., Hogan, J. L., & Lewandowsky, S. (2017). Reminders and repetition of misinformation: Helping or hindering its retraction? *Journal of Applied Research in Memory and Cognition*, 6(2), 185–192. <https://doi.org/10.1016/j.jarmac.2017.01.014>
- Ecker, U. K. H., Lewandowsky, S., Cook, J., Schmid, P., Fazio, L. K., Brashier, N., ... Amazeen, M. A. (2022). The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1(1), 13–29. <https://doi.org/10.1038/s44159-021-00006-y>
- Ecker, U. K. H., Lewandowsky, S., Fenton, O., & Martin, K. (2014). Do people keep believing because they want to? Preexisting attitudes and the continued influence of misinformation. *Memory & Cognition*, 42(2), 292–304. <https://doi.org/10.3758/s13421-013-0358-x>
- Ecker, U. K. H., 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. <https://doi.org/10.3758/s13423-011-0065-1>
- Ecker, U. K. H., O'Reilly, Z., Reid, J. S., & Chang, E. P. (2020). The effectiveness of short-format refutational fact-checks. *British Journal of Psychology*, 111(1), 36–54. <https://doi.org/10.1111/bjop.12383>
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>
- 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), 1420–1436. <https://doi.org/10.1037/0278-7393.20.6.1420>
- Kan, I. P., Pizzonia, K. L., Drumme, A. B., & Mikkelsen, E. J. V. (2021). Exploring factors that mitigate the continued influence of misinformation. *Cognitive Research: Principles and Implications*, 6(1), 76. <https://doi.org/10.1186/s41235-021-00335-9>
- Kendeou, P., Walsh, E. K., Smith, E. R., & O'Brien, E. J. (2014). Knowledge revision processes in refutation texts. *Discourse Processes*, 51(5–6), 374–397. <https://doi.org/10.1080/0163853X.2014.913961>
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., ... Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Lee, D. K. (2016). Alternatives to P value: Confidence interval and effect size. *Korean Journal of Anesthesiology*, 69(6), 555–562. <https://doi.org/10.4097/kjae.2016.69.6.555>
- Lewandowsky, S., Ecker, U. K. H., 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. <https://doi.org/10.1177/1529100612451018>
- Lewandowsky, S., Stritzke, W. G. K., Oberauer, K., & Morales, M. (2009). Misinformation and the “war on terror”: When memory turns fiction into fact. In W. G. K. In, S. Stritzke, D. Lewandowsky, J. C. Denemark, & F. Morgan (Eds.), *Terrorism and torture* (pp. 179–203). Cambridge University Press. <https://doi.org/10.1017/CBO9780511581199.010>
- Luengo, M., & García-Marín, D. (2020). The performance of truth: Politicians, fact-checking journalism, and the struggle to tackle COVID-19 misinformation. *American Journal of Cultural Sociology*, 8(3), 405–427. <https://doi.org/10.1057/s41290-020-00115-w>
- Maher, J. M., Markey, J. C., & Ebert-May, D. (2013). The other half of the story: Effect size analysis in quantitative research. *CBE—Life Sciences Education*, 12(3), 345–351. <https://doi.org/10.1187/cbe.13-04-0082>
- McIlhiney, P., Gignac, G. E., Ecker, U., Kennedy, B., & Weinborn, M. (2022). Executive function and the continued influence of misinformation: A latent variable analysis. *PsyArXiv*. <https://doi.org/10.31234/osf.io/pgsyd>
- Molina, M. D., Sundar, S. S., Le, T., & Lee, D. (2021). “Fake news” is not simply false information: A concept explication and taxonomy of online content. *American Behavioral Scientist*, 65(2), 180–212. <https://doi.org/10.1177/0002764219878224>
- NATO. (2020). NATO's approach to countering disinformation: A focus on COVID-19. NATO. <https://www.nato.int/cps/en/natohq/177273.htm>
- NATO. (2022). NATO-Russia relations: The facts. NATO. https://www.nato.int/cps/fr/natohq/topics_111767.htm?selectedLocale=en
- Open Science Collaboration. (2012). An open, large-scale, collaborative effort to estimate the reproducibility of psychological science. *Perspectives on Psychological Science*, 7(6), 657–660. <https://doi.org/10.1177/1745691612462588>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), Article aac4716. <https://doi.org/10.1126/science.aac4716>
- O'Rear, A. E., & Radvansky, G. A. (2020). Failure to accept retractions: A contribution to the continued influence effect. *Memory & Cognition*, 48(1), 127–144. <https://doi.org/10.3758/s13421-019-00967-9>
- Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7(6), 531–536. <https://doi.org/10.1177/1745691612463401>
- Pashler, H., & Wagenmakers, E. (2012). Editors' introduction to the special section on replicability in psychological science: A crisis of confidence? *Perspectives on Psychological Science*, 7(6), 528–530. <https://doi.org/10.1177/1745691612465253>
- Ranney, M. A., & Clark, D. (2016). Climate change conceptual change: Scientific information can transform attitudes. *Topics in Cognitive Science*, 8(1), 49–75. <https://doi.org/10.1111/tops.12187>
- Romero, F. (2019). Philosophy of science and the replicability crisis. *Philosophy Compass*, 14(11), Article e12633. <https://doi.org/10.1111/phc3.12633>
- Rozenbeek, J., Linden, S. v. d., & Nygren, T. (2020). Prebunking interventions based on “inoculation” theory can reduce susceptibility to misinformation across cultures. *Harvard Kennedy School Misinformation Review*, 1(2). <https://doi.org/10.37016/mr-2020-008>
- Sanderson, J. A., Bowden, V., Swire-Thompson, B., Lewandowsky, S., & Ecker, U. K. H. (2022). Listening to misinformation while driving: Cognitive load and the effectiveness of (repeated) corrections. *Journal of Applied Research in Memory and Cognition*. <https://doi.org/10.1037/mac0000057>
- Sénéchal, C., Larivée, S., Richard, E., & Robert, Y. (2004). Controverse Vaccin RRO et autisme: La désinformation à l'oeuvre. *Revue de psychoéducation*, 33(1), 205–227.
- Simmons, J., Nelson, L., & Simonsohn, U. (2012). A 21 word solution. *The Official Newsletter of the Society for Personality and Social Psychology*, 26(2), 4–7.
- Simons, D. J. (2014). The value of direct replication. *Perspectives on Psychological Science*, 9(1), 76–80. <https://doi.org/10.1177/1745691613514755>
- Susmann, M. W., & Wegener, D. T. (2021). The role of discomfort in the continued influence effect of misinformation. *Memory & Cognition*. <https://doi.org/10.3758/s13421-021-01232-8>
- Treen, K. M. d. L., Williams, H. T. P., & O'Neill, S. J. (2020). Online misinformation about climate change. *WIREs Climate Change*, 11(5), Article e665. <https://doi.org/10.1002/wcc.665>
- van der Linden, S. (2015). The conspiracy-effect: Exposure to conspiracy theories (about global warming) decreases pro-social behavior and science acceptance. *Personality and Individual Differences*, 87, 171–173. <https://doi.org/10.1016/j.paid.2015.07.045>
- 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>
- Walter, N., & Murphy, S. T. (2018). How to unring the bell: A meta-analytic approach to correction of misinformation. *Communication Monographs*, 85(3), 423–441. <https://doi.org/10.1080/03637751.2018.1467564>
- Walter, N., & Tukachinsky, R. (2020). A meta-analytic examination of the continued influence of misinformation in the face of correction: How powerful is it, why does it happen, and how to stop it? *Communication Research*, 47(2), 155–177. <https://doi.org/10.1177/0093650219854600>
- Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *The Quarterly Journal of Experimental Psychology Section A*, 40(2), 361–387. <https://doi.org/10.1080/02724988843000168>
- Xu, Y., Wong, R., He, S., Veldre, A., & Andrews, S. (2020). Is it smart to read on your phone? The impact of reading format and culture on the continued influence of misinformation. *Memory & Cognition*, 48(7), 1112–1127. <https://doi.org/10.3758/s13421-020-01046-0>