



Research Note

Gamified inoculation reduces susceptibility to misinformation from political ingroups

Psychological inoculation interventions, which seek to pre-emptively build resistance against unwanted persuasion attempts, have shown promise in reducing susceptibility to misinformation. However, as many people receive news from popular, mainstream ingroup sources (e.g., a left-wing person consuming left-wing media) which may host misleading or false content, and as ingroup sources may be more persuasive, the impact of source effects on inoculation interventions demands attention. In this experiment, we find that although news consumers are more susceptible to (non-political) misinformation from political ingroup publishers pre-intervention, gamified inoculation successfully improves veracity discernment and reduces susceptibility to misinformation from both political ingroup and outgroup publishers.

Authors: Cecilie Steenbuch Traberg (1), Jon Roozenbeek (1,2), Sander van der Linden (1)

Affiliations: (1) Department of Psychology, University of Cambridge, UK, (2) Department of War Studies, King's College London, UK.

How to cite: Traberg, C. S., Roozenbeek, J., & van der Linden, S. (2024). Gamified inoculation reduces susceptibility to misinformation from political ingroups. *Harvard Kennedy School (HKS) Misinformation Review*, 5(2).

Received: January 3rd, 2024. Accepted: April 9th, 2024. Published: April 30th, 2024.

Research questions

- Are news consumers more susceptible to misinformation from political ingroup publishers?
- Can gamified inoculation reduce susceptibility to misinformation from real news publishers?
- Can gamified inoculation reduce susceptibility to misinformation from political ingroup publishers?
- Are news consumers more susceptible to misinformation from ingroup publishers post-intervention?
- Does gamified inoculation improve veracity discernment or the ability to tell apart misinformation and “real news”?

Research note summary

- News consumers are more susceptible to misinformation if it comes from political ingroup publishers (which align with their ideological preferences).
- Gamified inoculation reduces susceptibility to misinformation even if the source of the misinformation is a well-known (mainstream) news publisher.

¹ A publication of the Shorenstein Center on Media, Politics and Public Policy at Harvard University, John F. Kennedy School of Government.

- Gamified inoculation reduces susceptibility to misinformation from political ingroup publishers.
- News consumers are no longer more susceptible to ingroup (compared to outgroup) attributed misinformation after a gamified inoculation intervention.

Implications

Research has shown that it is possible to cognitively immunize individuals against misinformation (which we define as any information that is false or misleading), much like how people can be vaccinated against viral infections (Compton et al., 2021; McGuire, 1964; Traberg et al., 2022; van der Linden, 2022). Gamified inoculation interventions (Roozenbeek & van der Linden, 2019; Maertens et al., 2024)—which both forewarn participants about the threat posed by misinformation and cognitively prepare them to refute deceptive tactics and strategies—have shown promise in terms of both their effects and potential for scalability (Roozenbeek, van der Linden, et al., 2022). However, many people receive news from popular (mainstream) political ingroup news sources (Grieco, 2020; Matsa, 2021), which have been found to increase misinformation susceptibility (Traberg & van der Linden, 2022). This prompts the question of whether inoculation interventions can reduce susceptibility to misinformation if it is produced by popular and well-known news sources, particularly sources that share news consumers' political ideology.

This study addresses these questions by testing whether a popular gamified inoculation intervention—the *Bad News* game (Roozenbeek & van der Linden, 2019)—can reduce susceptibility to misinformation from (mainstream) political ingroup sources. Previous studies have shown that people who play *Bad News* are significantly better at correctly identifying misinformation as being unreliable either post-gameplay (Roozenbeek, Traberg, et al., 2022) or compared to a control condition (Basol et al., 2020).

However, to date, inoculation interventions have largely been tested using misinformation from bogus sources or without source attributions (Traberg et al., 2022). However, as news consumers often rely on and seek out partisan news sources (Grieco, 2020; Matsa, 2021), and popular (mainstream) news sources have also been found to publish misleading and sometimes false content (Motta et al., 2020), it is crucial to understand whether inoculation interventions also protect people even in the presence of potentially powerful source effects. Examining whether interventions can overcome source effects is relevant and important not just for inoculation interventions, but any cognitive or psychological intervention aimed at reducing misinformation susceptibility. This is in part because misinformation attributed to mainstream sources that news consumers may find particularly credible can be more persuasive than complete falsehoods from bogus sources (Traberg, 2022; Tsfati et al., 2020).

To answer this question, our study employed a mixed pre-post between-subjects design, in which participants were assigned to view a series of misinformation headlines either from left-leaning US news outlets (e.g., *The New York Times*, *The Washington Post*), right-leaning U.S. news outlets (e.g., *Fox News*, *The Wall Street Journal*), or a control condition (with no source information), both before and after playing *Bad News* (total $N = 657$). We recoded source conditions into political ingroups or outgroups based on participants' own indicated political ideology, to examine whether or not misinformation susceptibility would be differentially reduced based on whether the sources of misinformation were political in- or outgroups, and further assessing this difference relative to the control condition, in which participants saw no information about the news sources.

Promisingly, our results show that across conditions, playing *Bad News* successfully reduces misinformation susceptibility; that is, participants are better at correctly identifying misleading headlines as unreliable post-gameplay, with no major source effects. Our results further show that participants were indeed more likely to perceive misinformation attributed to political ingroup publishers as being more reliable pre-intervention compared to when it was attributed to an outgroup or with no source

information present. This was the case when groups were categorised based on the match or mismatch between a participant's political ideology and a previous crowdsourced content-based analysis of the news source slant (Budak et al., 2016). In addition, however, we found similar results when we categorised the data based on participants' own perceptions of the source's media slant. At the same time, participants were also significantly more likely to reject factual headlines when they were attributed to an outgroup source pre-inoculation.

Encouragingly, however, across the ingroup, outgroup, and control conditions, reliability judgements of misinformation were lower post-intervention. In fact, this reduction was (descriptively, not significantly) largest for misinformation from ingroup sources. We also find that inoculation lowered the perceived reliability of factual headlines post-gameplay, but only at an effect size of $d = 0.11$, considered to be a negligible effect size (Sullivan & Feinn, 2012). Furthermore, veracity discernment, that is, participants' ability to discern between factual information and misinformation significantly improved post-gameplay.

This work builds on our understanding of the impact of source cues on misinformation interventions. Our results highlight that sources have a significant impact on misinformation susceptibility, but also show that gamified inoculation interventions can significantly improve misinformation susceptibility despite the presence of these effects. In fact, we also find that although group-based source effects are present in pre-intervention judgements of misinformation, these effects disappear post-intervention, suggesting that inoculation interventions can contribute to reducing source biases, despite the intervention not being explicitly designed to reduce them.

We further show that clear source effects appear pre-intervention based on participants' own perceived slant of the source. Previous research has to date only investigated the impact of researcher-categorised in- and outgroups, defining an "ingroup" based on whether participants' indicated political ideology matches a pre-defined categorisation of source slant—e.g., defining *Fox News* as an ingroup for a participant who identifies as conservative. We here show the same effects for sources that participants *themselves* believe to share their political ideology.

Our results also impact our understanding of how we should test and develop interventions tailored at reducing misinformation susceptibility. Although we find that the *Bad News* game reduces susceptibility to misinformation from political ingroup sources, we also find clear ingroup/outgroup source effects. This suggests that interventions might benefit from incorporating training in the impact of contextual cues that may mislead news consumers, in addition to content-based clues.

This study is one of the first to test the efficacy of a misinformation intervention with a variety of contextual cues, but there are several limitations. First, although we find promising results, the results are based on a limited selection of (U.S. mainstream) sources and test items and, as with most intervention research, are limited in the sense that online participants are self-selected and may be those who are already willing to learn about misinformation. Furthermore, there are far more contextual cues and biases in the news environment than sources alone and publishers are not the only sources that may play a role in the news consumption process. The "sharer" or "messenger" of headlines may play an equal if not larger role. Secondly, because a burgeoning literature is now pointing to a significant impact of sources on misinformation susceptibility, our work suggests that inoculation interventions may benefit from including training and refutational material that specifically seeks to reduce this bias to achieve effects more relevant to the real world.

Findings

Finding 1: News consumers are more susceptible to misinformation from political ingroup sources pre-intervention.

While we did not pre-register this analysis, our goal was to assess whether participants would be more susceptible to misinformation from political ingroups pre-inoculation. After all, for it to be interesting to examine whether inoculation interventions can specifically reduce misinformation susceptibility to misinformation from ingroup sources, this bias must be observed prior to administering inoculation interventions in the first place. While previous research has found this to be the case (Traberg & van der Linden, 2022), it was necessary to assess whether this was also the case in the current study.

We ran a generalized linear model (GLM) with *perceived reliability of misinformation pre-gameplay* as the dependent variable (DV) and *group identity* (ingroup, outgroup, control/no source) as the independent variable (IV). Results showed an overall model effect, $F(2,654) = 9.97, p < .001, \eta_p^2 = 0.03$: participants were significantly more likely to judge misinformation headlines as reliable when they were attributed to political ingroup sources ($M = 3.36, SE = 0.09$) compared to when attributed to both political outgroups ($M = 3.02, SE = 0.09, d = 0.29, p < .001$) and a control condition ($M = 2.91, SE = 0.09, d = 0.40, p < .001$). There was no significant difference between misinformation attributed to outgroups and the control condition ($p = .61, d = 0.10$), a finding which suggests it is *ingroups* that are powerful in making information more persuasive.

As a robustness check, we also ran the same analysis with source group level categorized based on the participants' own perception of the sources' slant (e.g., if the participant identified as right-wing and they judged the source to be a right-leaning source, this was categorized as a perceived ingroup). In this categorization, we categorized sources as "neither" if participants rated it a 4 (neither left-leaning nor right-leaning on the 7-point scale). A GLM with perceived reliability of misinformation pre-gameplay as DV and perceived group identity as IV was significant, $F(3, 644) = 7.15, p < .001, \eta_p^2 = 0.03$, and mirroring the results above, participants were significantly more likely to judge misinformation headlines as reliable when they were attributed to sources they *perceived* as ingroup sources ($M = 3.36, SE = 0.08$) compared to when attributed to both the perceived outgroup ($M = 3.02, SE = 0.09, d = 0.30, p < .05$) and the control ($M = 2.91, SE = 0.09, d = 0.40, p < .001$). There was also a significant difference between the control ($M = 2.91, SE = 0.09$) and the "neither" conditions ($M = 3.40, SE = 0.12, d = 0.43, p < .01$), and between the perceived outgroup ($M = 3.02, SE = 0.09$) and "neither" ($M = 3.40, SE = 0.12, d = 0.33, p < .05$) conditions. No other significant differences between conditions were observed. These findings suggest that sources are powerfully persuasive cues when it comes to misinformation and that news consumers are more likely to judge misinformation as reliable if it is attributed to a political ingroup news publisher. Findings are displayed in Figure 1.

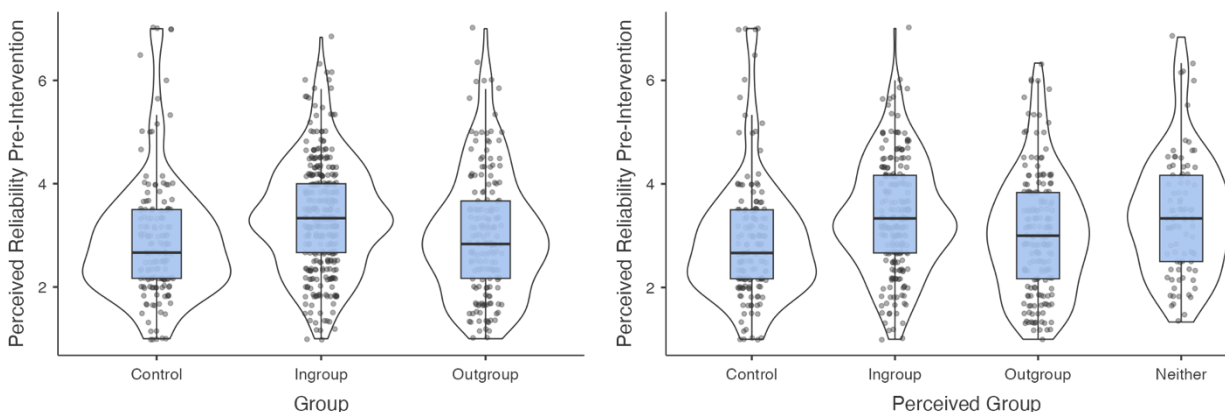


Figure 1. 1A (left): Mean perceived reliability of headlines pre-inoculation in the control condition, the ingroup source condition and the outgroup source condition; 1B (right): Mean perceived reliability of headlines pre inoculation in the control condition, the perceived ingroup condition, the perceived outgroup and the "neither" conditions. The figure shows boxplots (with the black horizontal lines indicating the median and the blue areas indicating the inter-quartile range) as well as violin plots (showing the data distribution) and data jitter (individual data points).

Finding 2: Gamified inoculation reduces susceptibility to misinformation from real news publishers.

As hypothesized in the pre-registration, we find a significant overall reduction in perceived reliability after playing *Bad News* ($M = 2.55, SD = 1.38$) compared to pre-gameplay ($M = 3.16, SD = 1.16, t(656) = 14.10, p < .001, d = 0.55$).

To further examine the impact of the source conditions compared to the control condition in which no sources were present, we ran a GLM with *perceived reliability of misinformation* as the DV and *inoculation* (pre-/post-intervention) and *source condition* (left, right, control/no source) as IVs. Based on previous research (Traberg & van der Linden, 2022), we hypothesized in our pre-registration that the inoculation effect (the pre-/post-difference in reliability) would be larger for misinformation from left-leaning sources (as our sample skewed liberal and as the left-leaning sources included have previously been judged to be more credible, on average; see Table S3 for the sample composition). The results showed a main effect of inoculation across conditions, $F(1, 1308) = 53.84, p < .001, \eta_p^2 = 0.04$, such that the perceived reliability of our misinformation items was significantly lower post-gameplay ($M = 2.43, SE = 0.05$) compared to pre-gameplay ($M = 2.99, SE = 0.05, p < .001, d = 0.45$). We also find a significant main effect of news source condition, $F(2, 1308) = 38.84, p < .001, \eta_p^2 = 0.06$, such that misinformation headlines from left-leaning sources were judged as significantly more reliable ($M = 3.11, SE = 0.05$) compared to misinformation headlines from right-leaning sources ($M = 2.43, SE = 0.08, p < .001, d = 0.55$) and without source attributions ($M = 2.57, SE = 0.07, p < .001, d = 0.43$). No significant differences were observed between reliability judgements of headlines from right-leaning sources and without source attributions ($p = .36, d = 0.12$). However, we find no significant interaction between the source conditions and the inoculation effect, $F(2, 1308) = 2.07, p = .13, \eta_p^2 < 0.01$, suggesting that there was no significant difference between the inoculation effect based on the source conditions. As such, our hypothesis was not confirmed. Findings are displayed in Figure 2 (see Table S2 and Figure S1 in the Appendix for a full overview).

This finding has important implications for our understanding of the efficacy of inoculation interventions in the modern news environment where even popular and mainstream media sources have been found to publish misinformation (not necessarily false but rather misleading; see Roozenbeek & van der Linden, 2024, chapter 1). False or misleading information coming from such publishers may have longer-lasting effects than if the source is unknown or generally known to be unreliable (Traberg, 2022; Tsfati et al., 2020) and the post hoc effects may be harder to eliminate, given that these mainstream sources benefit from significantly higher levels of perceived credibility.

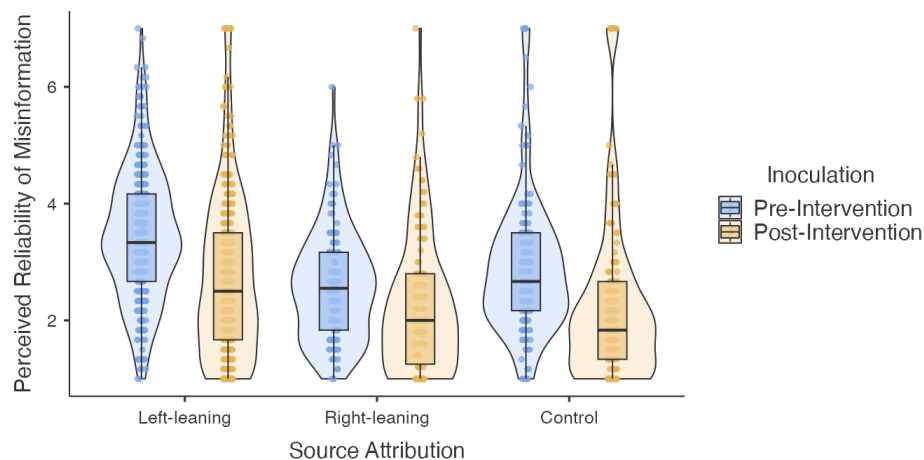


Figure 2. Mean perceived reliability of headlines pre-inoculation vs post-inoculation across source conditions. The figure shows boxplots (with the black horizontal lines indicating the median, the black dot indicating the mean, and the boxed areas indicating the inter-quartile range) as well as violin plots (showing the data distribution) and data jitter (individual data points).

Finding 3: Gamified inoculation reduces susceptibility to misinformation from political ingroup sources.

Playing the *Bad News* game successfully reduced misinformation susceptibility across all conditions. We hypothesized that there would be no interaction between political source congruence and the inoculation effect (pre-/post-difference). Testing null hypotheses can be complicated, as the absence of a significant effect (i.e., $p < .05$ under a frequentist framework) does not always mean a *true* absence of effects, as there may be small yet meaningful effects that were not found due to, for instance, a lack of sufficient sample size (Lakens et al., 2020). We, therefore, tested this hypothesis under a Bayesian framework, which allows for more intuitive null hypothesis testing. Specifically, we conducted a Bayesian ANOVA with the perceived reliability of misinformation as the dependent variable, group (ingroup/outgroup/control) and time (pre-/post-intervention) as fixed factors, and participant ID as a random effect.²

Doing so shows that the interaction group * time (i.e., the interaction between experimental condition and time, i.e., pre vs post) does not add meaningful predictive power under a Bayesian framework ($BF_{\text{Inclusion}} = .153$; $P(\text{Model} \mid \text{Data}) = .037$).³ In other words, the posterior values of the perceived reliability of misinformation are not meaningfully different if the model controls for this interaction. This means that although both experimental condition (group) and time (pre-post) separately strongly predict the perceived reliability of misinformation ($BF_{10} = 2.67 \cdot 10^{38}$ for the group + time model), we found support for the null hypothesis that source congruence does not influence the inoculation effect, i.e., the pre-post difference in reliability ratings of misinformation (see supplementary analyses and Tables S1A and S1B in the Appendix for a full overview of this analysis).

In fact, examining the difference in reductions across conditions, we find a greater reduction in perceived reliability of misinformation in the ingroup-source misinformation condition ($M_{\text{diff}} = 0.64$, $SE = 0.10$) compared to the outgroup-source misinformation condition ($M_{\text{diff}} = 0.52$, $SE = 0.13$).

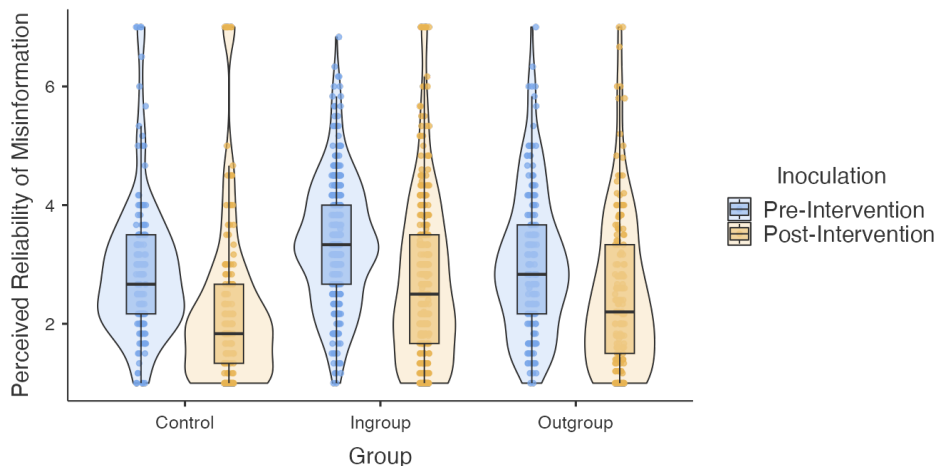


Figure 3. Mean perceived reliability of headlines pre-inoculation and post-inoculation across the three source conditions. The figure shows boxplots (with the black horizontal lines indicating the median and the boxed areas indicating the inter-quartile range) as well as violin plots (showing the data distribution) and data jitter (individual data points).

² To do so, we used Jamovi (<https://www.jamovi.org/>), specifically the “JSQ” library which allows for Bayesian analyses; see the analysis script on our Harvard Dataverse page.

³ A Bayes Factor of $< .200$ is considered sufficient support for the null hypothesis H_0 , the equivalent of a Bayes Factor of > 5 being considered sufficient support for the alternative hypothesis H_1 (Lakens et al., 2020; van Doorn et al., 2021)

Finding 4: News consumers are no longer more susceptible to ingroup (compared to outgroup) attributed misinformation post-gameplay.

Given that we find statistically significant differences on the perceived reliability of misinformation pre-intervention, we ran an exploratory analysis to examine whether these differences were also present post-intervention. An ANOVA with group identity (ingroup, outgroup, control) and perceived reliability of misinformation post-intervention shows a significant impact of group, $F(2,654) = 6.55, p < .01$. However, post hoc comparisons show that the only significant difference is between the control and ingroup ($M_{diff} = -0.48, SE = 0.13, p < .01$), with no significant differences between the ingroup and outgroup ($M_{diff} = 0.22, SE = 0.13, p = .20$) or between the outgroup and control ($M_{diff} = -0.26, SE = 0.15, p = .19$). This suggests that, although there was a significant group-level bias on misinformation judgements prior to inoculation, this bias was no longer present post-gameplay. It is, therefore, possible that inoculating individuals against misinformation also helps reduce ingroup-sourced misinformation bias. However, further research is necessary to confirm this hypothesis. Furthermore, source biases may be more pervasive in the real world and when it comes to political misinformation. As such, it may still be beneficial for future inoculation interventions to also aim to inoculate specifically against source effects on top of inoculating against content effects.

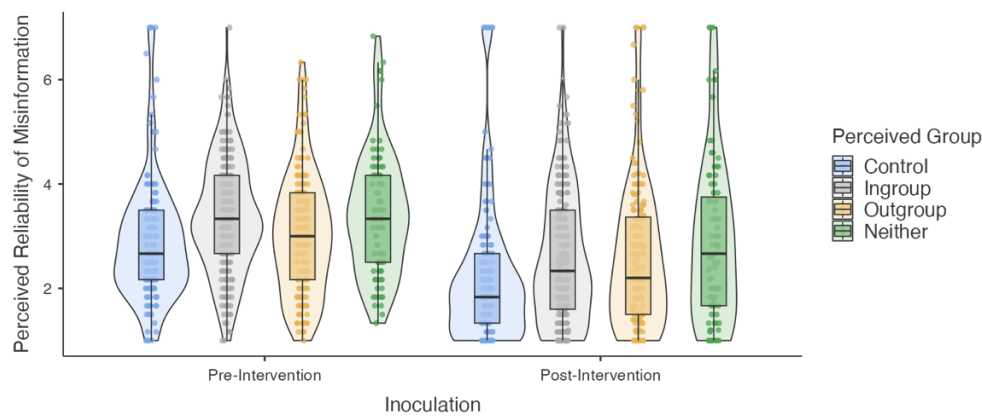


Figure 4. Mean perceived reliability of headlines pre-inoculation in the perceived group conditions. The figure shows boxplots (with the black horizontal lines indicating the median and the boxed areas indicating the inter-quartile range) as well as violin plots (showing the data distribution) and data jitter (individual data points).

Finding 5: Gamified inoculation significantly improves veracity discernment.

Given the open question regarding the efficacy of gamified inoculation interventions in improving veracity discernment (van der Linden, 2024), we analyzed whether participants' veracity discernment was significantly increased post-gameplay. Veracity discernment was calculated as the average perceived reliability of factual news minus the average perceived reliability of misinformation. We conducted a paired-samples t -test comparing veracity discernment pre- and post-inoculation intervention; this showed a significant improvement, $t(651) = 8.48, p < .001, M_{Diff} = 0.48, SE = 0.06, d = 0.33$. That is, participants were significantly better at discerning between the reliability of factual news and misinformation post-gameplay compared to before. Figure 5 illustrates this result. In line with our pre-registration, we also conducted a Bayesian paired-samples t -test on the pre- and post-inoculation reliability judgments of true news headlines, showing a small reduction in the perceived reliability of true news ($d = 0.12$) (see supplementary analyses in the Appendix).

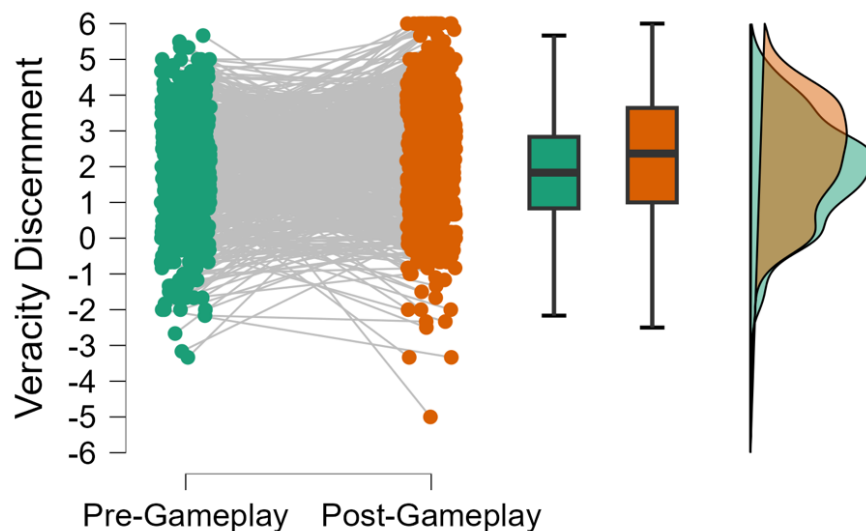


Figure 5. Mean veracity discernment (perceived reliability of factual news minus perceived reliability of misinformation) pre- and post-inoculation. The figure shows the data jitter (individual data points), boxplots (with the black horizontal lines indicating the median and the green/red areas indicating the inter-quartile range) as well as violin plots (showing the data distribution).

Methods

The purpose of this study was to investigate whether source cues impact the efficacy of the *Bad News* game. The study was approved by the University of Cambridge Psychology Research Ethics Committee (PRE.2019.104). All statistical analyses were performed using Jamovi (www.jamovi.org) or R. Visualisations were created using Jamovi and JASP (www.jasp-stats.org). Please see our Harvard Dataverse page for the datasets and analysis scripts.

Participants

Data was collected on the *Bad News* platform (see more information under “inoculation intervention”). Data collection was set to run for as long as the game platform was available to host the study—between November 1, 2020, and January 15, 2021. We pre-registered the study and main analyses on the Open Science Framework (see https://osf.io/uqwd2?view_only=f1030b5dcefe4278b13d092cb0d88676).

This study employed a between-subjects design and manipulated source slant across three conditions (left-wing source false/right-wing source facts, $n = 374$; right-wing source/left-wing facts, $n = 126$; and control, $n = 157$). In line with our pre-registration, we recoded conditions such that we had three group identity categories (ingroup, $n = 321$; outgroup, $n = 179$; control, $n = 157$) based on whether the source slant matched participants’ own identified political ideology.

The final sample was thus $N = 657$ participants (43% female, 84% with at least some college education, 71% between 18–29; 21% conservatives). As such the sample was unbalanced on these demographic variables (see Table S3 in the Appendix for more details). This skew is in line with previous studies that relied on data collected from the *Bad News* game’s in-game survey tool (e.g., Roozenbeek & van der Linden, 2019; Roozenbeek, Traberg, et al., 2022). As we relied on voluntary survey participation by people who navigated to the *Bad News* website, we were unable to reduce sample skew for this study. However, we note that our sample nonetheless contains $n = 137$ conservatives, which is substantial and comparable to previously published studies asking similar research questions around political ideology and misinformation susceptibility (e.g., Roozenbeek & van der Linden, 2020; Traberg & van der Linden, 2022).

Source

We manipulated the publisher logo on news headlines across three conditions: a traditionally left-wing publisher, a right-wing publisher, and a control condition where the source was blurred out. News sources were selected using the same method as Traberg & van der Linden (2022) relying on a previous crowdsourced content analysis of news media (Budak et al., 2016). The left-wing publishers were *The New York Times*, *The Washington Post*, and *CNN*; the right-wing publishers were *Fox News*, *The Wall Street Journal*, and *Breitbart*.



Figure 6. Examples of the source manipulation across conditions. (A) Left-wing source misinformation, (B) Right-wing source misinformation, (C) Control condition misinformation.

Headlines

To select unreliable news headlines, we used Hoaxy, a platform developed by researchers at Indiana University that visualizes the spread of claims and fact checking. We carried out a search to identify misleading headlines based on their use of one of three misleading strategies used by misinformation producers identified by previous research (Roozenbeek & van der Linden, 2020). These three misinformation tactics included: 1) using exaggeratedly emotional language to distort the news story to generate strong emotional responses, 2) creating or inspiring conspiratorial thinking to rationalise current events, and 3) discrediting otherwise reputable individuals, institutions, or facts to instil doubt in audiences. The non-misleading headlines did not make use of any manipulation techniques and were based on factual news.

Measures

Perceived reliability of misinformation

To measure misinformation susceptibility, participants were asked to rate each item's reliability on a standard 7-point Likert scale: "How reliable is the above news headline?" (1 = very unreliable, 7 = very reliable).

Source similarity

Source similarity was coded according to whether or not the participants' reported political ideology matched the political slant of the source. As political affiliation was originally collected on a 7-point scale,

this was converted to a binary (liberal vs conservative) measure and for the analysis which focused specifically on source similarity, participants who had reported a “moderate” political stance (that is, a 4 on the 1–7 scale) were excluded from analyses.

In a secondary analysis and as a manipulation check, we also analysed participants’ own perceptions of the source slant and categorised sources based on this. That is, if a participant identified as being “left-wing” and rated *CNN* as “left-wing,” this was categorised as an ingroup source.

Demographics

Age was measured categorically (18–29, 30–49, and over 50). Highest level of completed education was categorized as “high school or less,” “college or university,” and “higher degree;” gender was grouped as “female,” “male,” and “other.” Finally, political affiliation was measured on a 7-point scale where 1 is very left-wing and 7 is very right-wing (see Table S3 in the Appendix for a sample overview).

Inoculation intervention

The inoculation intervention used was the *Bad News* game (<https://www.getbadnews.com/>; Basol et al., 2020; Roozenbeek & van der Linden, 2019), an online “serious” game in which players learn about six manipulation techniques commonly used in misinformation, such as impersonation, conspiratorial reasoning, or ad hominem attacks (see Roozenbeek & van der Linden, 2024, chapter 8). Individuals roleplay as a “fake news” producer in a simulated social media environment, where they are exposed to weakened doses of misinformation. Within this environment, they are actively prompted to generate their own content and gain a following by spreading misinformation. The game is free and was created in collaboration with the Dutch media platforms DROG (*DROG*, 2023) and TILT (<https://www.tilt.co/>).

Procedure

Participants who visited the *Bad News* game website and played the game between November 1, 2020, and January 15, 2021, were automatically asked if they would like to take part in a scientific study prior to gameplay. If participants agreed to take part in this study, they were asked to provide informed consent. Following this, they answered the socio-demographic information. Participants were assigned to a condition by selecting a number from one to three, which then redirected them to one of the three source conditions. Once they were assigned a condition, they were exposed to new headlines, told that these were screenshots of real headlines, and asked to report on the above measures (perceived reliability). All participants subsequently played the *Bad News* game, after which they were asked to judge the reliability of the same headlines post-gameplay. Finally, they answered questions about the various source slants and credibility. After completing the study, participants were debriefed regarding the true purpose of the study and informed that the headlines were fictitious.

Bibliography

- Basol, M., Roozenbeek, J., & van der Linden, S. (2020). Good news about Bad News: Gamified inoculation boosts confidence and cognitive immunity against fake news. *Journal of Cognition*, 3(1), 2. <https://doi.org/10.5334/joc.91>
- Budak, C., Goel, S., & Rao, J. M. (2016). Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly*, 80(S1), 250–271. <https://doi.org/10.1093/poq/nfw007>

- Compton, J., van der Linden, S., Cook, J., & Basol, M. (2021). Inoculation theory in the post-truth era: Extant findings and new frontiers for contested science, misinformation, and conspiracy theories. *Social and Personality Psychology Compass*, 15(6). <https://doi.org/10.1111/spc3.12602>
- DROG. (2023). <http://drog.group>
- Grieco, E. (2020). *Americans' main sources for political news vary by party and age*. Pew Research Center. <https://www.pewresearch.org/short-reads/2020/04/01/americans-main-sources-for-political-news-vary-by-party-and-age/>
- Lakens, D., McLatchie, N., Isager, P., Scheel, A., & Dienes, Z. (2020). Improving inferences about null effects with Bayes factors and equivalence tests. *The Journals of Gerontology. Series B, Psychological Sciences and Social Sciences*, 75(1). <https://doi.org/10.1093/geronb/gby065>
- Maertens, R., Roozenbeek, J., Simons, J., Lewandowsky, S., Maturo, V., Goldberg, B., Xu, R., & van der Linden, S. (2024). Psychological booster shots targeting memory increase long-term resistance against misinformation. *Nature Communications*. <https://doi.org/10.31234/osf.io/6r9as>
- Matsa, M. W. and K. E. (2021, September 20). *News consumption across social media in 2021*. Pew Research Center. <https://www.pewresearch.org/journalism/2021/09/20/news-consumption-across-social-media-in-2021/>
- McGuire, W. J. (1964). Some contemporary approaches. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 1, pp. 191–229). Elsevier. [https://doi.org/10.1016/S0065-2601\(08\)60052-0](https://doi.org/10.1016/S0065-2601(08)60052-0)
- Motta, M., Stecula, D., & Farhart, C. (2020). How right-leaning media coverage of COVID-19 facilitated the spread of misinformation in the early stages of the pandemic in the U.S. *Canadian Journal of Political Science*, 53(2), 335–342. <https://doi.org/10.1017/S0008423920000396>
- Roozenbeek, J., & van der Linden, S. (2024). *The psychology of misinformation*. Cambridge University Press.
- Roozenbeek, J., & van der Linden, S. (2020). Breaking Harmony Square: A game that “inoculates” against political misinformation. *Harvard Kennedy School (HKS) Misinformation Review*, 1(8). <https://doi.org/10.37016/mr-2020-47>
- Roozenbeek, J., van der Linden, S., & Nygren, T. (2020). Prebunking interventions based on “inoculation” theory can reduce susceptibility to misinformation across cultures. *Harvard Kennedy School (HKS) Misinformation Review*, 1(2). <https://doi.org/10.37016/mr-2020-008>
- Roozenbeek, J., Traberg, C. S., & van der Linden, S. (2022). Technique-based inoculation against real-world misinformation. *Royal Society Open Science*, 9(5), 211719. <https://doi.org/10.1098/rsos.211719>
- Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, 5(1), <https://doi.org/10.1057/s41599-019-0279-9>
- Roozenbeek, J., van der Linden, S., Goldberg, B., Rathje, S., & Lewandowsky, S. (2022). Psychological inoculation improves resilience against misinformation on social media. *Science Advances*, 8(34), eabo6254. <https://doi.org/10.1126/sciadv.abo6254>
- Sullivan, G. M., & Feinn, R. (2012). Using effect size—Or why the p value is not enough. *Journal of Graduate Medical Education*, 4(3), 279–282. <https://doi.org/10.4300/JGME-D-12-00156.1>
- Traberg, C. S. (2022). Misinformation: Broaden definition to curb its societal influence. *Nature*, 606(7915), 653. <https://doi.org/10.1038/d41586-022-01700-4>
- Traberg, C. S., Roozenbeek, J., & van der Linden, S. (2022). Psychological inoculation against misinformation: Current evidence and future directions. *The Annals of the American Academy of Political and Social Science*, 700(1). <https://doi.org/10.1177/00027162221087936>

- Traberg, C. S., & van der Linden, S. (2022). Birds of a feather are persuaded together: Perceived source credibility mediates the effect of political bias on misinformation susceptibility. *Personality and Individual Differences, 185*(111269). <https://doi.org/10.1016/j.paid.2021.111269>
- Tsfati, Y., Boomgaarden, H. G., Strömbäck, J., Vliegenthart, R., Damstra, A., & Lindgren, E. (2020). Causes and consequences of mainstream media dissemination of fake news: Literature review and synthesis. *Annals of the International Communication Association, 44*(2), 157–173. <https://doi.org/10.1080/23808985.2020.1759443>
- van der Linden, S. (2024). Countering misinformation through psychological inoculation. In B. Garwonski (Ed.), *Advances in experimental social psychology* (Vol. 69, pp. 1–58). Elsevier. <https://doi.org/10.1016/bs.aesp.2023.11.001>
- van Doorn, J., van den Bergh, D., Böhm, U., Dablander, F., Derks, K., Draws, T., Etz, A., Evans, N. J., Gronau, Q. F., Haaf, J. M., Hinne, M., Kucharský, Š., Ly, A., Marsman, M., Matzke, D., Gupta, A. R. K. N., Sarafoglou, A., Stefan, A., Voelkel, J. G., & Wagenmakers, E.-J. (2021). The JASP guidelines for conducting and reporting a Bayesian analysis. *Psychonomic Bulletin & Review, 28*(3), 813–826. <https://doi.org/10.3758/s13423-020-01798-5>

Funding

This research was supported by the Harding Distinguished Postgraduate Scholarship and the Economic and Social Research Council Doctoral Training Programme.

Competing interests

The authors declare no competing interests. However, two authors (J. R. and S. v. d. L.) were involved in the development of the *Bad News* game. They have no financial stake in the game, DROG group, or TILT.

Ethics

The research protocol employed was approved by the University of Cambridge Psychology Research Ethics Committee (PRE.2019.104).

Copyright

This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided that the original author and source are properly credited.

Data availability

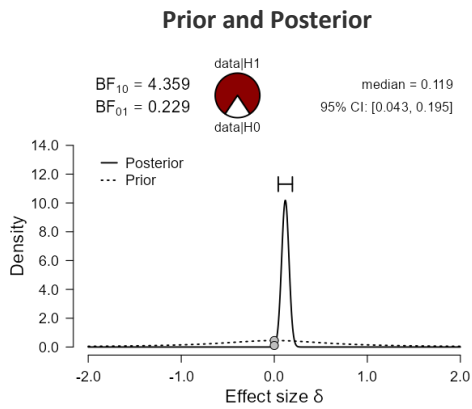
All materials needed to replicate this study are available via the Harvard Dataverse: <https://doi.org/10.7910/DVN/BOJP4D>

Appendix: Supplementary analyses


To test whether playing the *Bad News* game reduces the perceived reliability of reliable/true information (“real news”), we conducted a Bayesian paired-samples *t*-test on the pre- and post-inoculation reliability judgments of true news headlines. Doing so shows that the perceived reliability of true news is reduced post-inoculation compared to pre-inoculation ($BF_{10} = 4.359$, indicating moderate support for the alternative hypothesis). However, we note that this effect is very small ($\delta = .119$, which represents the population-level version of Cohen’s *d*). Furthermore, an exploratory Bayesian paired-samples *t*-test on the pre- and post-scores for the difference between average reliability ratings for misinformation and true news (i.e., veracity discernment) yields strong support for the alternative hypothesis ($BF_{10} = 2.49 \times 10^{13}$, $\delta = .332$), meaning that inoculation strongly reduces the perceived reliability of misinformation relative to true news. See Table S2 and Figure S1 for a full overview.

Using a frequentist approach, a paired-samples *t*-test showed that the difference between perceived reliability of factual headlines pre- and post-intervention, $t(1303) = 4.31$, $M_{diff} = 0.13$, $SE = 0.03$, $p < .001$, but only at a small effect size of $d = 0.12$.


True News



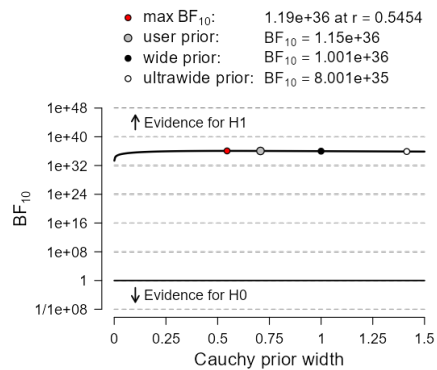
Bayes Factor Robustness Check



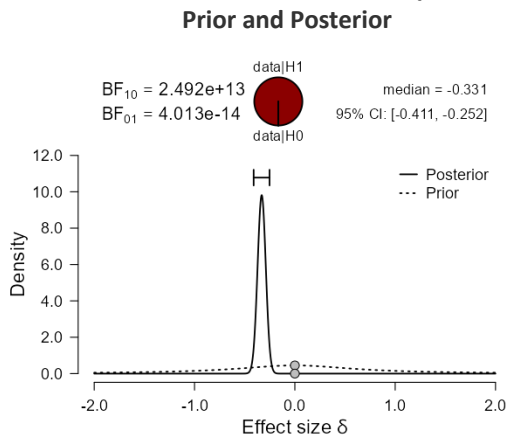
Misinformation



Bayes Factor Robustness Check



Discernment (difference: true news – misinformation)



Bayes Factor Robustness Check




Figure S1. Bayes factors and effect sizes (left panels) and robustness checks with various priors and levels of support for the alternative hypothesis H1 that there is a difference in perceived reliability when comparing pre- and post-gameplay scores (right panels) for true news, misinformation, and veracity discernment.

Table S1A. Bayesian ANOVA results for the interaction between experimental condition (group) and time (pre-post intervention): Model comparison – PerceivedReliability (perceived reliability of misinformation).

Models	P(M)	P(M data)	BF _M	BF ₁₀	Error %
Null model (incl. Session.ID)	0.2	3.59E-39	1.44E-38	1	
Group	0.2	8.21E-37	3.28E-36	228.86	1
Time	0.2	0.00362	0.0145	1.01E+36	1.16
Group + time	0.2	0.95949	94.7437	2.67E+38	2.85
Group + time + Group * time	0.2	0.03689	0.1532	1.03E+37	1.57

Note: All models include Session.ID. $P(M)$ represents the prior probabilities of each model; since we have five models, this is $1/5 = 0.2$ for each model. $P(M|data)$ signifies the posterior probability of each model after seeing the data. BF_m is a Bayes Factor that compares each model to the $P(M|data)$ of the other models. BF_{10} gives the probability of the data if H_1 is true, $P(D|H_1)$. The inclusion Bayes Factor ($BF_{Inclusion}$) represents the evidence in the data for including a predictor (the interaction group * time).

Table S1B. Bayesian ANOVA results for the interaction between experimental condition (group) and time (pre-post intervention): Analysis of effects – PerceivedReliability.

Analysis of Effects – PerceivedReliability		
P(incl)	P(incl data)	BF _{Inclusion}
0.6	0.9964	183.583
0.6	1	6.00E+15
0.2	0.0369	0.153

Note: Analysis was conducted using the “JSQ” package in Jamovi (<https://www.jamovi.org>). Specifically, we ran a Bayesian ANOVA with the perceived reliability of misinformation (PerceivedReliability) as the dependent variable, experimental condition (group) and time (pre-post) as fixed factors, and session.ID (i.e., participant ID) as a random effect. We used the long-format dataset available on the Harvard Dataverse.

Table S2. Paired-samples Bayesian t-tests for the pre-post difference in the perceived reliability of true news (FactsPre/FactsPost), misinformation (FakePre/FakePost), and the difference between ratings of true news and misinformation, i.e., “veracity discernment” (DiffPre/DiffPost). Bottom table shows descriptive statistics.

			BF ₁₀	Error %
FactsPre	-	FactsPost	4.36	0.00547
FakePre	-	FakePost	1.15E+36	4.71E-43
DiffPre	-	DiffPost	2.49E+13	7.82E-20

Descriptives

	N	Mean	SD	SE
FactsPre	657	4.98	1.14	0.0446
FactsPost	652	4.84	1.37	0.0538
FakePre	657	3.16	1.16	0.0451
FakePost	657	2.55	1.38	0.054
DiffPre	657	1.82	1.54	0.0601
DiffPost	652	2.29	1.76	0.0689

Table S3. *Sample composition in N (count) and percentages of total sample.*

Category	Groups
Education	High school or less (107; 16.3%)
	Higher degree (288; 43.8%)
	Some college (262; 39.9%)
Gender	Male (337; 51.3%)
	Female (284; 43.2%)
	Other (36; 5.5%)
Age	18-29 (464; 70.6%)
	30-49 (156; 23.7%)
	Over 50 (37; 5.6%)
Political ideology	1 (67; 10.2%)
	2 (211; 32.1%)
	3 (242; 36.8%)
	5 (75; 11.4%)
	6 (25; 3.8%)
	7 (37; 5.6%)

Note: Political ideology is coded from very liberal (1) to very conservative (7).