



The global effectiveness of fact-checking: Evidence from simultaneous experiments in Argentina, Nigeria, South Africa, and the United Kingdom

Ethan Porter^{a,1,2} and Thomas J. Wood^{b,1}

^aSchool of Media and Public Affairs, The George Washington University, Washington, DC 20052; and ^bDepartment of Political Science, The Ohio State University, Columbus, OH 43210

Edited by Gordon Pennycook, University of Regina, Regina, SK, Canada, and accepted by Editorial Board Member Margaret Levi July 23, 2021 (received for review March 4, 2021)

The spread of misinformation is a global phenomenon, with implications for elections, state-sanctioned violence, and health outcomes. Yet, even though scholars have investigated the capacity of fact-checking to reduce belief in misinformation, little evidence exists on the global effectiveness of this approach. We describe fact-checking experiments conducted simultaneously in Argentina, Nigeria, South Africa, and the United Kingdom, in which we studied whether fact-checking can durably reduce belief in misinformation. In total, we evaluated 22 fact-checks, including two that were tested in all four countries. Fact-checking reduced belief in misinformation, with most effects still apparent more than 2 wk later. A meta-analytic procedure indicates that fact-checks reduced belief in misinformation by at least 0.59 points on a 5-point scale. Exposure to misinformation, however, only increased false beliefs by less than 0.07 points on the same scale. Across continents, fact-checks reduce belief in misinformation, often durably so.

misinformation | fact-checking | multicountry experiment

The spread of misinformation is a global phenomenon (1). Misinformation is said to have played a role in the Myanmar genocide (2), national elections (3), and the resurgence of measles (4). Scholars have investigated various means of reducing belief in misinformation, including, but not limited to, fact-checking (5–8). Yet, despite the global scope of the challenge, much of the available evidence about decreasing false beliefs comes from single-country samples gathered in North America, Europe, or Australia. The available evidence also pays scant attention to the durability of accuracy increases that fact-checking may generate. Prior research has shown that fact-checking can reduce false beliefs in single countries (9, 10). Yet, whether fact-checking can reduce belief in misinformation around the world and whether any such reductions endure are unknown.

We describe simultaneous experiments conducted in four countries that help resolve both questions. In partnership with fact-checking organizations, we administered experiments in September and October 2020 in Argentina, Nigeria, South Africa, and the United Kingdom. The four countries are diverse along racial, economic, and political lines, but are unified by the presence of fact-checking organizations that have signed on to the standards of the International Fact-Checking Network. The experiments evaluated the effects of fact-checks on beliefs about both country-specific and global misinformation.

In total, we conducted 28 experiments, evaluating 22 distinct fact-checks. To limit the extent to which differences in timing may have been responsible for differential effects, particularly on the global misinformation items, we fielded all experiments in each country at the same time. In each experiment, participants were randomly assigned to misinformation; misinformation followed by a fact-check; or control. All participants then immediately answered outcome questions about their belief in

the false claim advanced by the misinformation. Fact-checking stimuli consisted of fact-checks produced by fact-checking organizations in each country, while misinformation stimuli consisted of brief summaries of the false claims that led to the corresponding fact-checks. This allowed us to estimate misinformation effects (the effect of misinformation on belief accuracy compared to control) and correction effects (the effect of corrections on belief accuracy compared to misinformation).

The fact-checks targeted a broad swath of misinformation topics, including COVID-19, local politics, crime, and the economy. In Argentina, South Africa, and the United Kingdom, we were able to evaluate the durability of effects by recontacting subjects approximately 2 wk after the first survey. In the second wave, we asked subjects outcome questions once again, without reminding them of earlier stimuli or providing any signal about each claim's truthfulness.

The tested fact-checks caused significant gains in factual accuracy. A meta-analytic procedure indicated that, on average, fact-checks increased factual accuracy by 0.59 points on a 5-point scale. In comparison, the same procedure showed that misinformation decreased factual accuracy by less than 0.07 on the same scale and that this decrease was not significant. The observed accuracy increases attributable to fact-checks were durable, with most detectable more than 2 wk after initial exposure to the fact-check. Despite concerns that fact-checking can “backfire” and increase false beliefs (11), we were unable to identify any instances of such behavior. Instead, in all countries studied, fact-checks reduced belief in misinformation, often for a time beyond immediate exposure.

Significance

Little evidence exists on the global effectiveness, or lack thereof, of potential solutions to misinformation. We conducted simultaneous experiments in four countries to investigate the extent to which fact-checking can reduce false beliefs. Fact-checks reduced false beliefs in all countries, with most effects detectable more than 2 wk later and with surprisingly little variation by country. Our evidence underscores that fact-checking can serve as a pivotal tool in the fight against misinformation.

Author contributions: E.P. and T.J.W. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. G.P. is a guest editor invited by the Editorial Board.


Published under the PNAS license.

¹E.P. and T.J.W. contributed equally to this work.

²To whom correspondence may be addressed. Email: evporter@gwu.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2104235118/-/DCSupplemental>.

Published September 10, 2021.



Scholars are perennially concerned that their conclusions about human behavior are overly reliant on samples of Western, educated, industrialized, rich, and democratic, or “WEIRD,” populations (12, 13). WEIRD populations may be distinct from other populations, minimizing the external validity of psychological findings (14). Our evidence suggests that, when it comes to the effects of fact-checking on belief in misinformation, this is not the case. Although the countries in our study differ starkly along educational, economic, and racial lines, the effects of fact-checking were remarkably similar in all of them.

Misinformation and Fact-Checking in Global Context

Exposure to misinformation is widespread (3). On social media, misinformation appears to be more appealing to users than factually accurate information (15). However, research has identified various ways of rebutting the false beliefs that misin-

formation generates. Relying on crowd-sourcing (8), delivering news-literacy interventions (7), and providing fact-checks (9, 16) have all been shown to have sharp, positive effects on factual accuracy. Our experiments in the present study evaluated fact-checking efforts; for this reason, we hypothesized that exposure to factual corrections would increase subjects’ factual accuracy (H1). (We preregistered our hypotheses, research questions, and research design with the Open Science Framework [OSF]. The preregistration is included in *SI Appendix*.)

Little prior work of which we are aware has examined whether national setting affects the size and direction of correction effects. Critical for our purposes, a previous meta-analysis of the effects of attempts to correct misinformation (9) includes only WEIRD samples, none of which attempted to compare the effects of corrections across countries, let alone non-WEIRD countries. The populations of the four countries

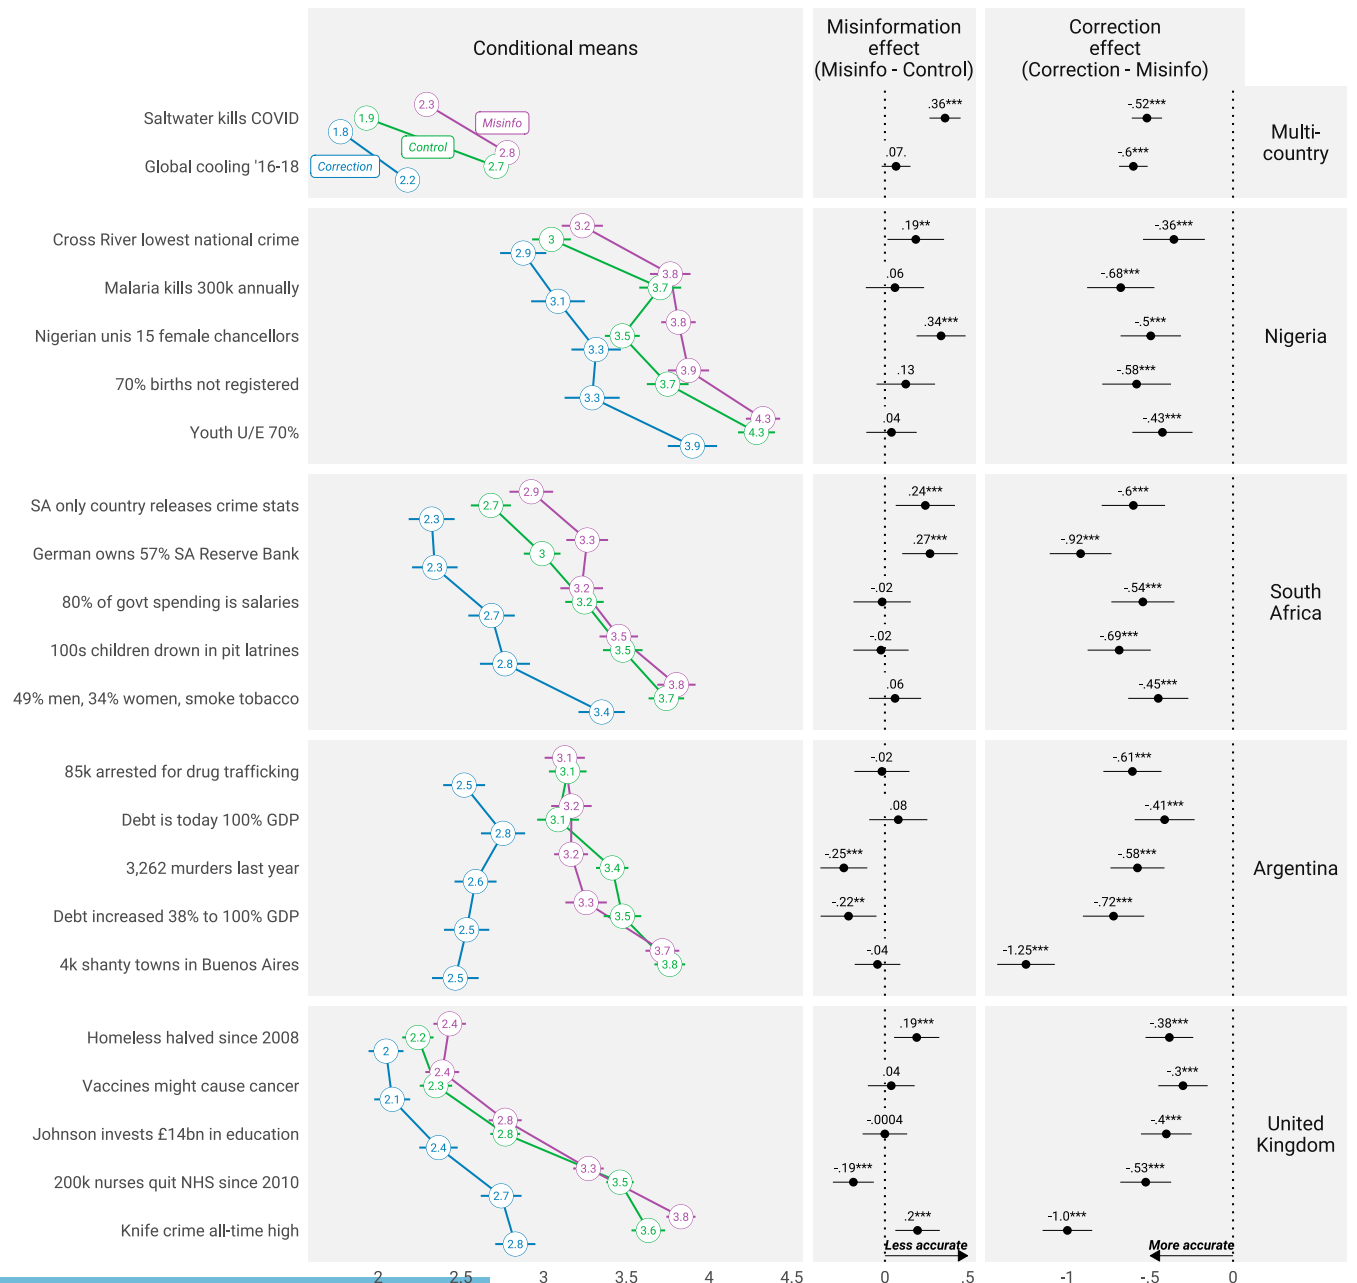


Fig. 1. Conditional means, correction, and misinformation effects. Horizontal lines report 95% CIs. ** $P < 0.01$; *** $P < 0.005$ (two-sided).

studied here are distinct along numerous lines, including aggregate ideological orientation and traditional demographics. It may be the case that the size of accuracy increases generated by fact-checks are different in different national settings. The size of any increases may also vary with different demographics. For these reasons, we studied research questions concerning the relationship between national setting and correction effects; the relationship between participants' ideology and correction effects; and the relationship between other demographics and correction effects.

The existing literature is also unclear on the duration of accuracy increases that may follow factual corrections. Even when fact-checks bring about greater accuracy, the initial misinformation can continue to affect reasoning over time (17). Differences in the duration of effects may be attributable to differences among the topics of misinformation and fact-checks. If the accuracy increases that follow fact-checks are only temporary, this suggests that the increases do not represent meaningful gains in accurate knowledge (18). While the effects of factual information in general can endure (19, 20), the durability of accuracy increases prompted by fact-checking in particular is not known. Given the uncertainty of existing findings, we investigate a research question pertaining to the duration of accuracy increases.

Finally, the existing literature does not systematically investigate whether different topics of political misinformation are more (or less) susceptible to factual correction. Scholars have studied a wide range of misinformation topics, including healthcare (21), climate change (22), and political candidates (23). Some issues may be "easy" to correct; others might prove more difficult (24, 25). So far as we are aware, there is no comprehensive evidence concerning how response to fact-checks differs between issues. The large number of fact-checks investigated here, spanning a broad array of issues, led us to examine a research question concerning any differences in accuracy increases across different topics.

Results

When compared to misinformation, every fact-check produced more accurate beliefs in the first wave. Misinformation, on the other hand, did not always lead to less accurate beliefs when compared to control in this wave. Results from the first wave for all items appear in Fig. 1. Effects are displayed on the mean outcome scale, with larger numbers corresponding with greater belief in factually inaccurate information. The first column displays conditional means. The next column displays misinformation effects, or the contrasts between exposure to control and

exposure to misinformation only. The third column displays correction effects, or the contrasts between exposure to factual corrections and exposure to misinformation. In the top row, we display effects for the two global items, pertaining to COVID-19 and climate change. Although corrections to both items led to greater accuracy, the misinformation effect for COVID-19 was the largest of all misinformation effects. The largest correction effect concerned the number of shanty towns in Buenos Aires, Argentina.

In the three countries for which we collected second-wave data, the correction effects of most country-specific misinformation items were significant in the second wave, as were the correction effects on both global items. Specifically, 9 of 15 country-specific correction effects, and both cross-country items, all in the direction of greater accuracy, were still significant. In Fig. 2, we depict the duration of the country-specific and global correction effects. (To address concerns about attrition, here, we present effects only for subjects who completed both waves.)

Much of the concern about negative responses to fact-checking has focused on the possibility that individuals' political views might lead them to reject fact-checking that conflicts with those views. To investigate this possibility, in all four countries, we gathered subjects' responses to the World Values Survey 10-point question regarding ideology prior to treatment (the full text is in *SI Appendix*).

While corrections did not yield identical effects across the ideological spectrum, they also did not provoke any instances of backfire. Fig. 3 presents results by ideological affiliation. Although some corrections failed to improve accuracy for some ideological groups, adherents of the left, middle, and center alike were made more accurate by fact-checks. This was the case for most of the country-specific items and for both global items. As we show in *SI Appendix*, meta-analyses of effects by subjects' ideology indicate that, globally, misinformation sans correction has a smaller effect on those who report being on the left than those on the right and center. Fig. 3 demonstrates that our large collection of misinformation topics featured policy areas of importance to subjects on the ideological left and right (although none specifically invoked ideological terms). That corrections worked even when the topic of misinformation was politically charged constitutes powerful evidence for fact-checks' efficacy, across countries and across ideologies.

Accuracy increases generated by the common global items were similar across countries and items. There was, however, a discrepancy between misinformation effects and correction effects, as made apparent by Fig. 4. While exposure to

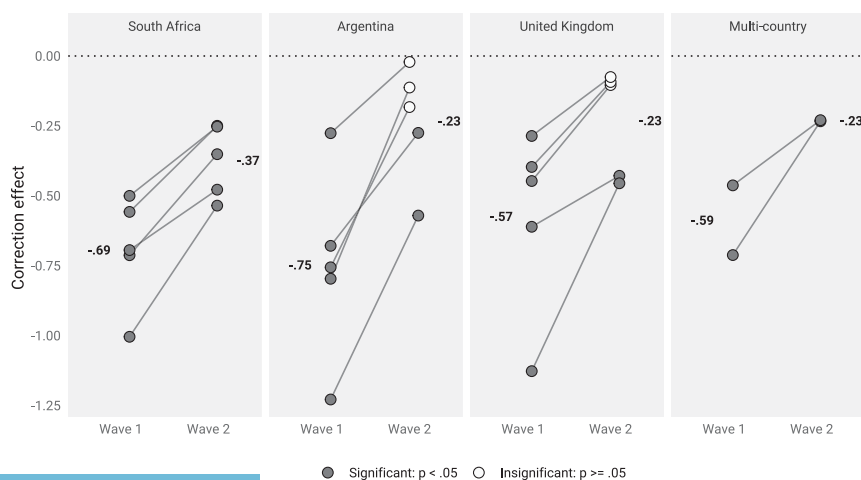


Fig. 2. Over-time effects.




Fig. 3. Correction effects and ideology (wave 1). Mod., moderate; S., South.

climate-change misinformation did not uniformly lead to accuracy degradation, exposure to a climate-change-related correction uniformly improved accuracy. In contrast, and of special relevance at the present moment, misinformation regarding COVID-19 degraded accuracy about COVID in three of the four countries. At the same time, fact-checks increased accurate beliefs about COVID-19 in all countries. The discrepancy between misinformation and correction effects may be attributable to features of our stimuli, as we discuss below.

To study whether the size of effects differed by topic of the misinformation, we grouped misinformation items into the following topics: government spending, health, crime, and economic data. (The specific items used in each group can be found in the preanalysis plan.) Pooling the correction and misinformation effects by these topics does not suggest that effects differ by topic. Correction and misinformation effects are indistinguishable from one another when items are grouped this way.

Meta-Analysis. To better understand our effects in aggregate, we performed meta-analysis with random effects of the 28 exper-

iments included in this study. Fig. 5 displays results of this exercise. Using this approach, we find that corrections reduced belief in falsehoods by 0.59 point on our 5-point scale ($P < 0.01$). On the same scale, misinformation only increased belief in falsehoods by 0.07 ($P > 0.05$).^{*} Fact-checks thus increase factual accuracy by more than eight times the amount that misinformation degrades factual accuracy.

Discussion and Conclusion

Scholars, governments, and civil society have investigated a variety of potential tools to combat the crisis of misinformation. Prior work has indicated that fact-checking is one effective such tool, capable of reducing false beliefs (9). Yet, much of the evidence gathered previously has focused on samples from a small handful of countries, limiting the generalizability of subsequent conclusions. The existing evidence also pays little

^{*}When we rely on standardized outcome variables in meta-analysis, we observe in Wave 1 that corrections decreased false beliefs by 0.45 SDs, while misinformation increased false beliefs by only 0.05 SDs.




Fig. 4. Global experimental effects (wave 1). ** $P < 0.01$; *** $P < 0.005$ (two-sided).

attention to how long the accuracy increases that fact-checks generate endure. This calls into question whether fact-checks are meaningfully improving accurate knowledge, as understood by prior scholars (18), or changing survey responses only ephemeral. Across a wide variety of national contexts, are fact-checks effective at reducing false beliefs? And are any of the effects detectable after immediate exposure?

Our evidence answers both questions in the affirmative. Experiments conducted simultaneously in Argentina, Nigeria, South Africa, and the United Kingdom reveal that fact-checks increase factual accuracy, decreasing belief in misinformation. This was the case across a broad array of country-specific items, as well as two items investigated in all countries. Meta-analysis demonstrates that fact-checks reduced belief in misinformation by 0.59 point on a 5-point scale, while exposure to misinformation without a fact-check increased belief in that misinformation by less than 0.07. The factual accuracy increases generated by fact-checks proved robust to the passage of time, with most still evident approximately 2 wk later. Although responses to the fact-checks differed along ideological lines, as prior literature would anticipate (26, 27), in no case did an ideological group become more inaccurate because they were exposed to a correction.

Our study makes clear that, in four diverse countries, fact-checking can help mitigate the threat that misinformation poses to factual accuracy. While fact-checks improved factual accuracy more than misinformation degraded it, our results may approximate the lower boundary of misinformation's effects. The misinformation stimuli did not include source cues or provide other signals that might have heightened its impact on participants. The fact-checks were also lengthier than the misinformation. In addition, our estimate of misinformation's effects may approximate the lower boundary because of a fundamental difference between our study and the world outside the laboratory. While our participants were compelled to see either fact-checks,

misinformation, or neither, observational research shows that misinformation is more intrinsically appealing on social media than accurate information (15). On their own accord, social media users could choose to expose themselves repeatedly to misinformation while avoiding fact-checks, inflating the effects of the former beyond what we find here. All that having been said, our evidence shows that, at least when it comes to fact-checking, responses are generally similar across diverse samples.

What, specifically, makes the effects of fact-checks durable is a topic deserving of further research. Fact-checks may be cognitively demanding and invite active processing, similar to other interventions that have generated longer-term effects (28). Prior work has shown that cognitive style is related to susceptibility to misinformation (29, 30). The inverse might be true as well, with subjects who perform comparatively well on Cognitive Reflection Tests more responsive to factual corrections. Future research should vary the extent to which fact-checks are cognitively demanding, while paying careful attention to respondents' cognitive styles.

Although the present study expands the geographic scope of research into misinformation and fact-checking, it hardly exhausts possible avenues of inquiry. Our experiments were all administered in countries that currently have reputable fact-checking organizations. Had we fielded our experiments in different countries, we may have observed different results, particularly if those countries had levels of political polarization or trust different from the countries tested here; if they had different political institutions; or if they had no existing fact-checking organizations within them. In addition, while we strove to present fact-checks in as realistic of a form as possible, we nonetheless relied on online panels that may limit the external validity of these findings. Future research into these topics should not only encompass more diverse participants, but, to the greatest extent possible, partner with fact-checking organizations to deliver interventions in close proximity to sources of misinformation. Doing so may account for subjects more willing to believe misinformation than subjects on online panels. To test this possibility, researchers should administer studies modeled on patient-preference trials, wherein participants' preferences for false claims and fact-checks can be accounted for (31).

Finally, the paucity of prior research in these countries limited our ability to target groups that may be especially susceptible to misinformation. In the United States, for example, research has identified demographic groups (32) and cognitive styles (30) that are associated with susceptibility to misinformation. This is also critical for the study of the duration of correction effects; research into the "continued influence effect" has found that some demographic groups are more likely to be influenced by false beliefs over time, even following an effective correction (33). Despite the absence of pronounced heterogeneous effects in the present study, researchers should conduct more cross-national studies of factual corrections and misinformation to determine whether the patterns we observe here are common.

While these findings illustrate the potential of fact-checking to rebut false beliefs around the world, fact-checking alone is likely insufficient to address the scope of the misinformation




Fig. 5. Meta-analysis of corrections and misinformation effects.

problem. Fact-checks may undo the effects of misinformation on factual beliefs, but whether they can also affect related political attitudes is unclear. As others have shown, even when corrections succeed in reducing false beliefs, they may, nonetheless, benefit politicians who disseminate them, with worrisome consequences for the incentives that politicians face (23). Our study does not address this question. For now, the evidence shows that, around the world, fact-checking causes durable reductions in false beliefs, mitigating one of the central harms of misinformation.

Materials and Methods

The selection of misinformation items and corresponding fact-checks was made in consultation with the fact-checking organizations on the ground. For the country-specific items, all factual corrections tested were genuine corrections previously used by the fact-checking organization in response to misinformation. To maximize external validity, participants exposed to a country-specific fact-check were shown the fact-check as it appeared on the fact-checking organization's website, with accompanying text and graphics as they originally appeared. The two cross-national items, pertaining to global warming and COVID-19, referred to genuine misinformation, but, to maintain consistency across countries, necessitated the generation of two novel fact-checks. Translation was conducted in partnership with the fact-checking organizations and survey vendors. The complete text of all misinformation items and factual corrections can be found in *SI Appendix*.

In each of the four countries, participants were randomly exposed to between zero and seven misinformation items and between zero and seven fact-checks. Randomization occurred at the item level. Of the seven misinformation items, two were common across all four states, while the remaining five were country-specific. This resulted in tests of 22 distinct fact-checks, evaluated with 28 experiments. For each item, participants were either exposed only to the misinformation item; the false item followed by a fact-check; or only answered outcome questions. The order of misinformation items was also randomized. A graphical depiction of the full factorial design can be found in the *SI Appendix*.

To measure outcomes, we asked respondents two questions about each misinformation item, both of which prompted the respondent to assess the veracity of the misinformation. We relied on two questions in order to minimize measurement error that might result from relying on only one question, a concern that prior work in this area has raised (34). First, respondents were asked to agree with a statement that summarized the false claim. Agreement was measured on a 1-to-5 agree-disagree scale, with larger numbers corresponding to greater accuracy. Then, participants were asked whether they regarded the statement as true or false, with responses measured on a 1-to-5 true-false scale, again with larger numbers corresponding with greater accuracy. The statement appeared twice, so that subjects could read it before responding to each question. To evaluate the effects of fact-checking and misinformation on beliefs, we modeled ordinary least-squares regressions of the following type:

$$\text{factual accuracy} = b_0 + b_1(\text{experimental condition}),$$

where experimental condition is a three-value factorial variable, with possible values for misinformation and outcome items; misinformation, fact-check, and outcome items; and outcome items only. Outcomes consist of the average response of the two items described above. To measure over-time effects, as depicted in Fig. 2, we estimated the same model on the second wave, holding constant those who completed both waves. In this second wave, participants were asked outcome questions only, receiving no reminders of their earlier treatment. While we cannot rule out the possibility that subjects who completed the second wave were distinct from those who only completed the first, our concerns are mitigated by the similarity in first-wave effects between those who completed both waves and those

who did not. Indeed, as we show in *SI Appendix*, effects in the first wave among those who completed both waves are indistinguishable from first-wave effects of those who completed both waves. To evaluate ideology, we estimated a linear ideological term, with outcomes consisting of the 10-point World Values Survey ideology scale. Specifically, this question asks subjects: "In political matters, people talk of 'the left' and 'the right.' How would you place your views on this scale, generally speaking?" A 1-to-10 scale appears below, with "left" above 1 and "right" above 10. The results shown in Fig. 3 are the regression contrasts, comparing misinformation and correction conditions, when we vary ideology from 1 to 10, in 0.01 increments. While at first blush, the randomized provision of misinformation may raise ethical concerns, it is important to note that factual corrections on their own almost always reiterate the misinformation being corrected. Experimental tests of factual corrections that aim to achieve a modicum of realism thus also effectively randomize misinformation, as we do here. With our approach, we separate out what other research sometimes collapses.

Sample Composition. Demographic data on all waves in all countries are located in *SI Appendix*. As we show, across demographic lines measured by the World Values Survey (including age, employment status, gender, and ideology), the sample composition of the first wave is broadly similar to national data in each country. The first wave of experiments began simultaneously on the week of September 24, 2020, and concluded shortly thereafter. In South Africa, the United Kingdom, and Argentina, we were able to conduct a second wave of the study, for which data collection began on October 16. In this second wave, subjects were only asked to provide answers to the outcome questions; no additional treatments were administered.

In the United Kingdom, South Africa, and Argentina, Ipsos MORI recruited subjects and collected data. Recruitment efforts relied on targets matched to official statistics on age, gender, region, and working status for each country. In wave 1, 2,000 adults in each country were surveyed, with resulting data weighted by age, gender, region, and working status to match the profile of the adult population for the following age group in each market: 18 to 75 in the United Kingdom, 18 to 50 in South Africa, and 18 to 55 in Argentina. Wave 2 ($n = 1,000$ in each country) was conducted by recontacting respondents who completed wave 1.

In Nigeria, YouGov was responsible for recruitment and data collection. Subjects were recruited by using banners on websites, emails to a permission-based database, and loyalty websites. The sampling frame was based on the 2017 Afrobarometer's estimation of the internet population in Nigeria, with sampling based on age, gender, education, and the combination of age and gender. This data were matched to a sampling frame, with matched cases weighted by propensity score. The matched cases and sampling frame were combined, and logistic regression determined inclusion in the frame.

Additional information. Data files and scripts necessary to replicate the results in this article are available in the Dataverse repository. This study was deemed exempt by the George Washington University Institutional Review Board.

Data Availability. Data and replication code (35) have been deposited in the Harvard Dataverse, <https://doi.org/10.7910/DVN/Y8WPFR>.

ACKNOWLEDGMENTS. We are grateful to Full Fact, Africa Check, Chequeado, and the many fact-checking professionals who helped bring this project to fruition, including, but not limited to, Mevan Babakar, Amy Sippitt, Peter Cunliffe-Jones, Ariel Riera, Will Moy, Olivia Vicol, and Nicola Theunissen. We are also grateful to seminar audiences at George Washington University and to Brendan Nyhan for feedback. Ipsos MORI and YouGov were responsible for the fieldwork and data collection only. This research is supported by a grant from Luminata, and by the John S. and James L. Knight Foundation through a grant to the Institute for Data, Democracy & Politics at The George Washington University. All mistakes are our own.

1. D. M. J. Lazer *et al.*, The science of fake news. *Science* **359**, 1094–1096 (2018).
2. P. Mozur, A genocide incited on Facebook, with posts from Myanmar's military. *The New York Times*, 15 October 2018. <https://www.nytimes.com/2018/10/15/technology/myanmar-facebook-genocide.html>. Accessed 23 August 2021.
3. H. Allcott, M. Gentzkow, Social media and fake news in the 2016 election. *J. Econ. Perspect.* **31**, 211–236 (2017).
4. T. Burki, Vaccine misinformation and social media. *Lancet Digit. Health* **1**, e258–e259 (2019).

5. G. Pennycook, A. Bear, E. T. Collins, D. G. Rand, The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Manage. Sci.* **66**, 4944–4957 (2020).
6. G. Pennycook, D. G. Rand, Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 2521–2526 (2019).
7. A. M. Guess *et al.*, A digital media literacy intervention increases discernment between mainstream and false news in the United States and India. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 15536–15545 (2020).

8. M. Tully, E. K. Vraga, L. Bode, Designing and testing news literacy messages for social media. *Mass Commun. Soc.* **23**, 22–46 (2020).
9. M. S. Chan, C. R. Jones, K. Hall Jamieson, D. Albarracín, Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychol. Sci.* **28**, 1531–1546 (2017).
10. E. Porter, T. J. Wood, *False Alarm: The Truth About Political Mistruths in the Trump Era* (Cambridge University Press, Cambridge, UK, 2019).
11. D. P. Christenson, S. E. Kreps, D. L. Kriner, Contemporary presidency: Going public in an era of social media: Tweets, corrections, and public opinion. *Presidential Studies Quarterly*, **51**, 151–165 (2020).
12. J. Henrich, S. J. Heine, A. Norenzayan, The weirdest people in the world? *Behav. Brain Sci.* **33**, 61–83, discussion 61–83 (2010).
13. M. S. Rad, A. J. Martingano, J. Ginges, Toward a psychology of *Homo sapiens*: Making psychological science more representative of the human population. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 11401–11405 (2018).
14. J. Friedman, WEIRD samples and external validity. *World Bank Blogs* (2012). <https://blogs.worldbank.org/impacetevaluations/weird-samples-and-external-validity>. Accessed 23 August 2021.
15. S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online. *Science* **359**, 1146–1151 (2018).
16. T. J. Wood, E. Porter, The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Polit. Behav.* **41**, 135–163 (2018).
17. S. Lewandowsky, U. K. H. Ecker, C. M. Seifert, N. Schwarz, J. Cook, Misinformation and its correction: Continued influence and successful debiasing. *Psychol. Sci. Public Interest* **13**, 106–131 (2012).
18. M. X. Delli Carpini, S. Keeter, *What Americans Know About Politics and Why It Matters* (Yale University Press, New Haven, CT, 1996).
19. P. Kowalski, A. K. Taylor, Reducing students' misconceptions with refutational teaching: For long-term retention, comprehension matters. *Scholarsh. Teach. Learn. Psychol.* **3**, 90–100 (2017).
20. P. R. Rich, M. H. Van Loon, J. Dunlosky, M. S. Zaragoza, Belief in corrective feedback for common misconceptions: Implications for knowledge revision. *J. Exp. Psychol. Learn. Mem. Cogn.* **43**, 492–501 (2017).
21. J. M. Carey, V. Chi, D. J. Flynn, B. Nyhan, T. Zeitzoff, The effects of corrective information about disease epidemics and outbreaks: Evidence from Zika and yellow fever in Brazil. *Sci. Adv.* **6**, eaaw7449 (2020).
22. J. Cook, S. Lewandowsky, U. K. H. Ecker, Neutralizing misinformation through inoculation: Exposing misleading argumentation techniques reduces their influence. *PLoS One* **12**, e0175799 (2017).
23. E. Thorson, Belief echoes: The persistent effects of corrected misinformation. *Polit. Commun.* **33**, 460–480 (2016).
24. R. M. Alvarez, J. Brehm, *Hard Choices, Easy Answers: Values, Information, and American Public Opinion* (Princeton University Press, Princeton, NJ, 2002).
25. E. G. Carmines, J. A. Stimson, The two faces of issue voting. *Am. Polit. Sci. Rev.* **74**, 78–91 (1980).
26. E. K. Vraga, L. Bode, Correction as a solution for health misinformation on social media. *Am. J. Public Health* **110**, S278–S280 (2020).
27. T. Bolsen, J. N. Druckman, F. L. Cook, The influence of partisan motivated reasoning on public opinion. *Polit. Behav.* **36**, 235–262 (2014).
28. D. Broockman, J. Kalla, Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science* **352**, 220–224 (2016).
29. M. Mosleh, G. Pennycook, A. A. Arechar, D. G. Rand, Cognitive reflection correlates with behavior on Twitter. *Nat. Commun.* **12**, 921 (2021).
30. G. Pennycook, D. G. Rand, Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition* **188**, 39–50 (2019).
31. D. Knox, T. Yamamoto, M. A. Baum, A. J. Berinsky, Design, identification, and sensitivity analysis for patient preference trials. *J. Am. Stat. Assoc.* **114**, 1532–1546 (2019).
32. A. Guess, J. Nagler, J. Tucker, Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Sci. Adv.* **5**, eaau4586 (2019).
33. B. Swire, U. K. H. Ecker, S. Lewandowsky, The role of familiarity in correcting inaccurate information. *J. Exp. Psychol. Learn. Mem. Cogn.* **43**, 1948–1961 (2017).
34. B. Swire, J. DeGutis, D. Lazer, Searching for the backfire effect: Measurement and design considerations. *PsyArXiv* [Preprint] (2020). <https://doi.org/10.31234/osf.io/ba2kc> (Accessed 23 August 2021).
35. E. Porter, T. J. Wood, Replication data for: The global effectiveness of fact-checking: Evidence from simultaneous experiments in Argentina, Nigeria, South Africa, and the U.K. Harvard Dataverse. <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DV/NY8WPFR>. Deposited 20 August 2021.