

Can Children Detect Fake News?

Shaio Xu (sxu@oxy.edu)

Department of Psychology, Occidental College

Andrew Shtulman (shtulman@oxy.edu)

Department of Psychology, Occidental College

Andrew G. Young (ayoung20@neiu.edu)

Department of Psychology, Northeastern Illinois University

Abstract

Fake news has permeated online media, presenting consumers with the challenge of detecting it. At what age are we capable of undertaking this challenge? And what factors predict success? We explored these questions with elementary-school-aged children ($n = 86$), who were asked to judge the veracity of ten news stories, five fake and five real. Children also completed a developmental version of the cognitive reflection test (CRT-D; Young & Shtulman, 2020a). As a group, children were at chance at differentiating fake news from real news, and their individual performance did not vary by age or cognitive reflection. Adults ($n = 271$) given the same materials succeeded at detecting fake news, especially those high in cognitive reflection. These results suggest that children lack the knowledge or skill needed to evaluate news credibility and that cognitive reflection predicts fake news detection only after we have attained some baseline level of information literacy.

Keywords: cognitive development, cognitive reflection, information literacy

Introduction

With the development of social media and the internet, people are increasingly exposed to fake news and conspiracy theories. According to Vosoughi et al. (2018), fake news reaches more people and spreads more quickly on the internet than real news. News consumers must actively discriminate fake news from real news, which raises several epistemological challenges: assessing the authenticity of the reporting in relation to prior factual knowledge, assessing the plausibility of the event in relation to prior conceptual knowledge, and assessing the credibility of the source in relation to prior social knowledge. Detecting fake news is a cognitive task that has both practical and theoretical significance.

A common explanation for the recent proliferation of fake news is partisanship and politically motivated reasoning (Kahan, 2017; Jardina & Traugott, 2019). Fake news is typically created by agents with a political agenda and spread by people who share that agenda. However, research on the cognitive underpinning of fake-news detection suggests that fake news may be appealing for different reasons, namely, a *lack* of reasoning and relevant knowledge (Pennycook & Rand, 2019; Pennycook & Rand, 2021). For example, Pennycook and Rand (2019) used the Cognitive Reflection

Test (CRT; Frederick, 2005)—a measure of individual differences in analytic thought—to investigate how people evaluate politically partisan news headlines, some real and some fake. They found that CRT performance was strongly associated with detecting fake news and the ability to discern fake news from real news. Indeed, CRT performance predicts accurate reasoning in general, including more accurate reasoning about science and increased skepticism towards religious, supernatural, and conspiratorial beliefs (Shtulman & McCallum, 2014; Stecula & Pickup, 2021; Pennycook et al., 2012).

Previous research has focused exclusively on adults' propensity to believe in fake news. Here, we explore children's perception of fake news and the factors that predict their ability to detect it. According to the 2019 American Community Survey, more than 95% of children (aged 3 to 18) have home access to the internet through computers or smartphones, yet children may be particularly vulnerable to misinformation on the internet given their limited understanding of source reliability (OfCom, 2019; Olafsson et al., 2014; Einav et al., 2020), limited cognitive abilities (UNICEF, 2021), and inadequate critical literacy skills (Flanagin & Metzger, 2008).

Several studies have explored children's perception of misinformation in terms of their ability to evaluate information sources (Flanagin & Metzger, 2008; Einav et al., 2020; Danovitch & Lane, 2021) or their ability to identify hoaxes (Loos et al., 2018; Dumitru, 2020). These studies have found that children do not generally use source information in a consistent way, nor are they generally able to discriminate hoaxes from real events. While these studies suggest that children would be poor at detecting fake news, they have not explored children's evaluation of fake news itself, as it might appear on social media, nor have they examined the relationship between fake-news detection and individual differences in analytic thinking, as has been done with adults.

We explored these issues utilizing a developmental version of the CRT—the CRT-D—developed by Young and Shtulman (2020a). This instrument predicts rational thinking and normative thinking dispositions in elementary-school-aged children, similar to how the CRT predicts the same constructs in adults (Gong et al., 2021). Performance on the

CRT-D also predicts children's understanding and learning of counterintuitive math and science concepts (Young & Shtulman, 2020b). Given this pattern of developmental continuity, we expected that children with higher cognitive reflection scores would be better at differentiating fake news from real news, as is true of adults (Pennycook & Rand, 2019).

In the current study, we asked two groups of participants—elementary-school-aged children and adults—to decide whether each of ten news stories was true or false. The stories were presented as headlines, with an accompanying image, one-sentence summary, and news source. Our primary questions were whether children could differentiate fake news stories from real ones and whether this ability was predicted by their age and by their performance on the CRT-D. Adults were included as a comparison group to verify that our fake news stories were, in fact, discriminable from our real news stories.

The adult comparison group also allowed us to assess the role of source information in discriminating fake news from real news, as half of our adult participants were not provided with this information. Source information has been shown to have mixed effects on adults' evaluation of news stories. Some studies find that source information improves adults' accuracy at identifying fake news (Nadarevic et al., 2020; Kim et al., 2019), while others find no influence (Austin & Dong, 1994; Shen et al., 2019; Pennycook & Rand, 2019; Dias et al., 2020). The value of source information is likely linked to the plausibility of the information itself, as sources should matter more when plausibility is harder to gauge. By manipulating the availability of source information, we sought to assess whether source familiarity might provide a distinct advantage at discriminating fake news from real news or whether the content of the news stories was sufficient on its own.

We expected that children would be worse than adults at discriminating fake news from real news but were uncertain whether this difference could be explained by source familiarity. That is, we were uncertain whether adults deprived of source information would perform comparably to children, comparably to adults given source information, or somewhere in between. We further expected that cognitive reflection scores would predict fake-news detection in both children and adults, though, for children, we were less certain about whether the relation between cognitive reflection and fake-news detection would be consistent across children of different ages (5 to 12) or between children and adults

Method

Participants

Eighty-six school-aged children (i.e., kindergarten to 6th grade; M age = 7.14 years, SD = 2.0; 56% female) were recruited from public playgrounds in Southern California. Child participants completed the tasks on-site with the consent of their guardians. Children were recruited over the course of two semesters, with the goal of recruiting 100

children. The resultant sample of 86 was sufficiently powered ($1-\beta > .85$) to detect a medium-sized difference between two dependent means ($d = .33$) and a medium-sized correlation ($r = .28$).

An additional 271 college undergraduate students participated (M age = 20 years; 72.6% female) for course credit or a \$7.50 Amazon gift card. Adults were recruited from introductory psychology courses over two semesters, with the goal of recruiting 100 per condition (source vs. no-source). The resultant samples ($n = 147$ and $n = 124$) were sufficiently powered ($1-\beta > .85$) to detect a medium-sized difference between two dependent means ($d = .23$ and $d = .28$) and a medium-sized correlation ($r = .22$ and $r = .24$).

Procedure and Materials

A battery of tests was administered to both children and adult participants in the order shown below. Children participants completed the study one-on-one with trained research assistants. The research assistants would read each question and record the child's verbal responses on a tablet. Adult participants completed the study online through Qualtrics.

Cognitive Reflection Test (CRT) Adult participants completed the original 3-item CRT (Fredrick, 2005). We used the number of correct responses as participants' scores, with higher scores indicating greater cognitive reflection.

Cognitive Reflection Test - Developmental (CRT-D) Adult and child participants answered 8 items from the children-friendly CRT-D (Young & Shtulman, 2020a). Example items are "I have 3 apples and you take away 2 from me, how many do you have?" and "If you're running a race and you pass the person in second place, what place are you in?". Prior research suggests the CRT-D functions well as a cognitive reflection test for adults (Gong et al., 2021). We used the number of correct responses as participants' scores, with higher scores indicating greater cognitive reflection.

Perceived Accuracy of News Headlines Participants judged 5 factually accurate headlines (real news) and 5 false headlines (fake news). The real news headlines were selected from mainstream news sources including FoxNews.com and NBCNews.com. The fake news headlines were selected from well-known fake news and satirical websites including PrettyCoolSite.com and WorldNewsDailyReport.com. All headlines featured child-friendly and politically neutral content (see Table 1). All were also culled from the internet and represent stories that news consumers could actually have encountered.

Figure 1 provides examples of the general presentation format, which included an image, the headline, and a brief one-sentence summary of the news story. This format was adapted from Pennycook and Rand (2019) and includes the information news consumers would most likely see in a social media post. Adult participants were randomized to source ($n = 147$) and no-source ($n = 124$) conditions. For children and adults in the source condition, headlines additionally

Table 1: Mean (SD) truth ratings (1 = “not at all true” to 4 = “very true”) for real (#1-5) and fake (#6-10) news headlines.

Headline	Children	Adults No Source	Adults Source
1. Germs In Your Gut are Talking to Your Brain. Scientists Want to Know What They are Saying	2.36 (1.21)	3.23 (.76)	3.33 (.75)
2. Is Milk Healthy? Canada’s New Food Guide Says Not Necessarily	2.47 (1.17)	2.82 (.83)	3.12 (.72)
3. Hotel Dishwasher Awarded \$21 Million After Boss Made Her Work on Sundays	2.40 (1.13)	2.73 (.90)	2.90 (.92)
4. United Flight Attendant Walks onto Biplane Wing Mid-Air in Birthday Stunt	2.29 (1.14)	2.44 (.99)	2.54 (.98)
5. Gnarly! 6-Story Wave Is Biggest Ever Recorded	2.55 (1.13)	2.49 (.83)	2.45 (1.01)
6. Man Kicked Out of All-You-Can-Eat Buffet After Eating More Than 50 LBS of Food, Sues for \$2-Million	2.12 (1.09)	2.42 (.85)	2.36 (.85)
7. Man Hospitalized After His AirPods Exploded in His Ears	2.14 (1.17)	2.61 (.88)	2.28 (.89)
8. Yellowstone Evacuated: Experts Claim ‘Super Volcano’ Could Erupt Within Weeks	2.87 (1.16)	2.13 (.91)	2.16 (.96)
9. Winner of ‘The Voice; Excited to Use \$50 Chili’s Gift Card	2.69 (1.07)	2.02 (.94)	1.96 (.91)
10. California Newborn Becomes First Baby to Be Named an Emoji: Her Name Is 🍌🍌🍌🍌	1.90 (1.12)	1.46 (.68)	1.41 (.80)

displayed news source information (i.e., a website), whereas adults in the no-source condition were not presented source information (see Figure 1).

Headlines were presented in random order. For each headline, participants answered “how true do you think this story is?” on a 4-point scale (1 = “not at all true”, 2 = “not very true”, 3 = “kind of true”, 4 = “very true”). Table 1 presents the mean truth ratings (on the 4-point scale) for each headline across condition.

To better achieve measurement equivalence between adult and child responses, we coded participants’ truth ratings into true judgments (“kind of true” or “very true”) and not true judgments (“not at all true” or “not very true”). However, we find the same overall pattern of results when data are analyzed on the 4-point scale.

Additional Tasks To facilitate efficient data collection, participants completed several additional tasks related to other research questions during the study session (e.g., a covariation reasoning task and a verbal fluency task). We do not consider these unrelated measures in the following analyses.

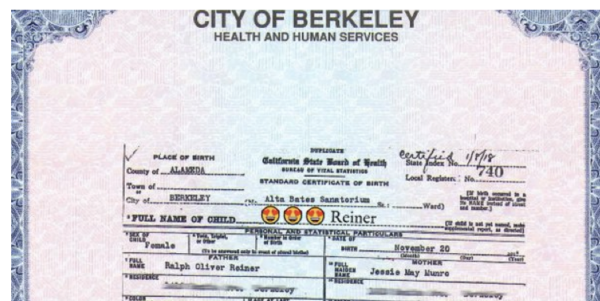
Results

We analyzed participants’ evaluations of news stories using binomial generalized linear mixed models (GLMM) on truth judgments (1 = true, 0 = not true). Following procedures recommended by Matuscheck et al. (2017), each model includes a parsimonious by-participant and by-item random effects structure generated from removing random effects from a maximal model that were not supported by the data. Inference for fixed effects was carried out via Type 2 likelihood ratio test (LRT) model comparison.



Hotel Dishwasher Awarded \$21 Million After Boss Made Her Work On Sundays

A jury has awarded a Miami hotel dishwasher \$21.5 million, finding that her employer violated her religious rights by repeatedly scheduling her for work on Sundays and ultimately firing her. NBCNEWS.COM



California Newborn Becomes First Baby To Be Named An Emoji: Her Name Is 🍌🍌🍌🍌

The United States is among the lightest of the naming laws worldwide governing given names and this freedom has given rise to a wide variety of names and naming trends.

Figure 1: Example headline items with source information (Top) and no source information (Bottom).

Discriminating Real from Fake News

To compare children and adults, we estimated a binomial GLMM on truth judgments with news type (real or fake), condition (child, adult no-source, or adult source), and their interactions as fixed effects. As seen in Figure 2, there was an interaction between news type and condition, $LRT \chi^2(2) = 8.83, p = .012$. Adults judged real news to be more true than fake news in the source condition, $OR = 6.61, 95\% CI [2.30, 18.97]$, and no-source condition, $OR = 4.06, 95\% CI [1.44, 11.39]$. Further, adults in the source condition were better at discriminating news types than the adults in the no source condition, $OR = 1.63, 95\% CI [1.01, 2.62]$. However, children did not judge real news to be more true than fake news, $OR = 1.08, 95\% CI [.58, 1.99]$. Indeed, they were no better than chance. Thus, children were unable to discriminate real news from fake news.

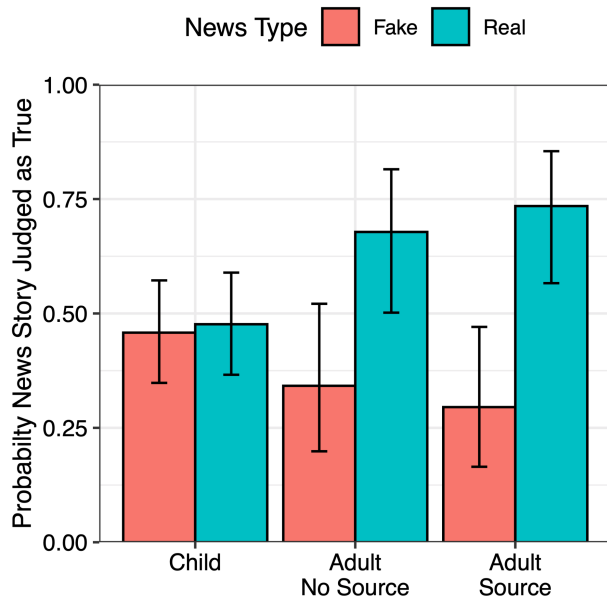


Figure 2: Estimated probabilities of true judgments by news type and condition. Error bars represent 95% CIs.

Cognitive Reflection, Age, & News Discrimination

To examine whether more reflective and/or older children were better able to discriminate real news from fake news, we estimated a binomial GLMM on truth judgments with news type (real or fake), age, CRT-D score, a news type by age interaction, and a news type by CRT-D score interaction as fixed effects. As can be seen in Figures 3A & 3B, age, CRT-D score, and their interactions with news type were unrelated to children's judgments. Thus, even the oldest and most reflective children in our study could not distinguish real news from fake news.

To confirm adult's cognitive reflection was related to news discrimination, we estimated a binomial GLMM on truth judgments with news type (real or fake), condition (no-source or source), CRT_{Total} (i.e., number correct across the original 3 item CRT and 8 item CRT-D), and their interactions as fixed

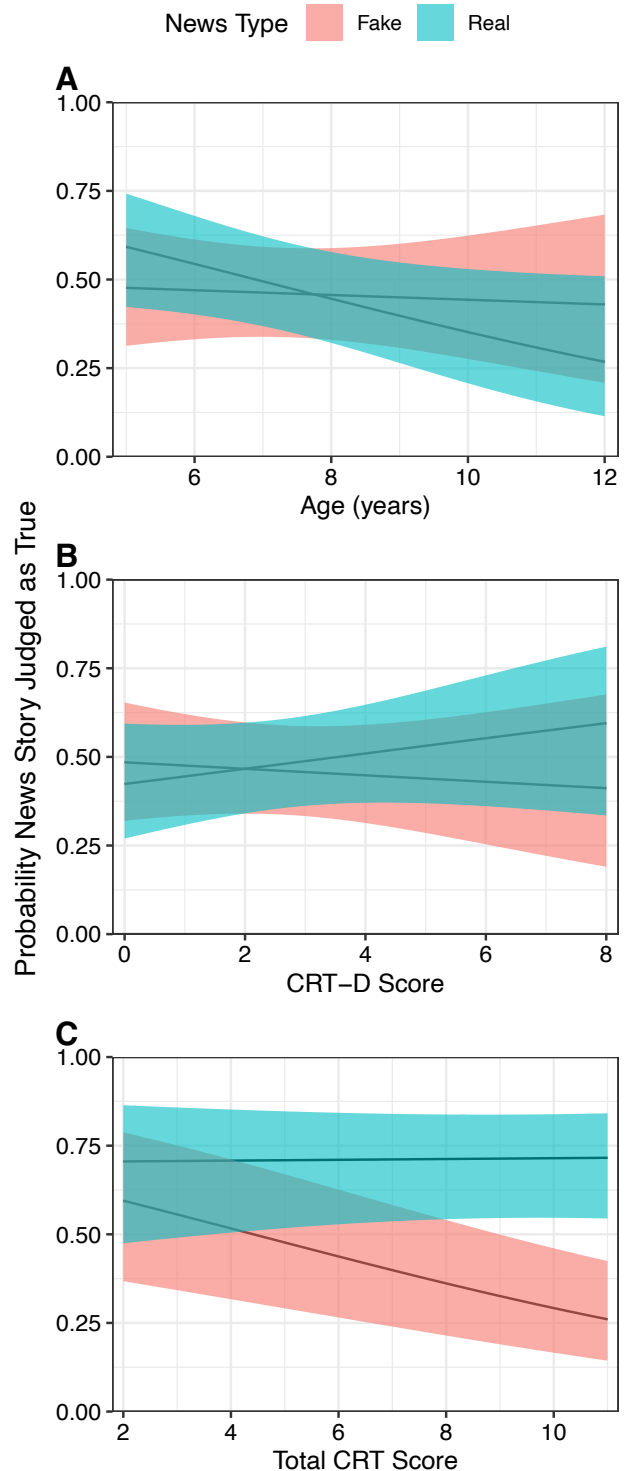


Figure 3: Estimated probabilities of true judgments by: A). children's age and news type, B). children's CRT-D score and news type, and C). adults' total CRT score by news type. Error bands represent 95% CIs.

effects. As seen in Figure 3C, there was a news type by CRT_{Total} interaction, $LRT \chi^2(1) = 7.86, p = .005$. CRT_{Total} did not predict adults' judgments of real news as true, $\text{logit } \beta = .01, 95\% \text{ CI } [-.08, .09]$. However, adults with greater CRT_{Total} were less likely to judge fake news as true, $\text{logit } \beta = -.16, 95\% \text{ CI } [-.08, -.24]$. When analyzed separately, similar interactions were observed for the original CRT, $LRT \chi^2(1) = 4.66, p = .031$, and CRT-D, $LRT \chi^2(1) = 6.58, p = .010$.

General Discussion

Are children able to detect fake news? Our results suggest that they are not. Children's judgments of whether a news story was true or false were at chance. Their judgments hovered around chance for different types of news (fake vs. real) and for most of the individual stories. Older children were no better at evaluating news stories than younger children, and children with high cognitive reflection scores were no better than those with low scores. The ability to decide whether a news story is true or false appears to develop sometime after elementary school, perhaps when children learn how to find and evaluate sources on their own, in the context of research projects.

Perhaps the most surprising aspect of children's failure to differentiate fake news from real news was that cognitively reflective children also failed to do so. Cognitively reflective adults, on the other hand, were particularly successful at the task, replicating prior findings that cognitive reflection facilitates fake-news detection (Pennycook & Rand, 2019). Cognitive reflection was a significant predictor of fake-news detection in adults regardless of whether we measured cognitive reflection with the standard CRT or with the child-friendly CRT-D, which speaks to the validity of the latter measure. Even though CRT-D scores did not predict fake-news detection in children, the CRT-D has proven to be a reliable measure of early cognitive reflection in the context of rational thought and normative thinking dispositions (Young & Shtulman, 2020a; Gong et al., 2021), both for children and adults. Here, we confirm that finding in light of its predictive power among adults, tracking the original CRT in its ability to predict fake-news detection.

While adults successfully differentiated fake news from real news, their ability to do so was improved by source information. Sources allowed adults to detect fake news above and beyond the implausibility of its content. The added value of source information helps explain what children might be lacking in their evaluation of news sources. Children likely viewed the fake news stories as implausible, but implausibility alone was not a reliable guide to whether the story was made-up. A story about germs in the gut talking to the brain might seem implausible, even to an adult, but the story was published in *The New York Times*, and anyone familiar with this source, and its record of credibility, would have reason to override their incredulity and accept the story as true.

Conversely, a story about a super volcano threatening to destroy Yellowstone National Park may sound credible on its surface, but its publication source—globalnetwork.info—

could help a reader recognize that this story, if true, would have appeared in more mainstream venues as well. In short, source information may help readers validate, or invalidate, their intuitions about plausibility, and children lack this crucial cue to veracity (e.g., Tong et al., 2022).

That said, children's inability to detect fake news may reflect different intuitions about plausibility as well. Future research should explore the underpinnings of children's failure by using a wider variety of news stories and comparing children's relative success to additional measures of content knowledge and conceptual understanding. The more children know about a topic, the better they might be at differentiating fake news from real news, regardless of their familiarity with the publication venue.

Still, there are several reasons to take children's failure in the current task as informative and not a mere artifact of the stories we used. First, these stories were culled from the internet and provided an authentic snapshot of the content that news consumers encounter, beyond the politically partisan content used in previous studies. Second, the adults in our study were able to discriminate the fake news from the real news, indicating that these stories contained sufficient cues for accuracy (or lack thereof), even without source information. Finally, children's failure to detect fake news was surprisingly profound. Children as old as twelve were at chance in deciding whether the stories were true or false, despite several years of instruction in social studies and language arts. Children's poor performance across ages and cognitive abilities underscores the difficulty of fake-news detection. Adults may succeed at this task, but the skills that allow them to do so should not be taken for granted.

In conclusion, children in elementary school appear ill-equipped to discriminate fake news from real news, even when they exhibit high levels of cognitive reflection. While cognitive reflection may facilitate adults' ability to recognize stories that could not plausibly map onto reality, cognitive reflection does not yield the same benefits for children, who may need additional content knowledge or source knowledge to succeed at this epistemologically challenging task. Future research on the developmental prerequisites of fake-news detection promises to shed light on the interaction between cognitive reflection and evidential reasoning. It also promises to advance our understanding of how to improve information literacy in an age of questionable information and information sources.

Acknowledgements

We would like to thank the James S. McDonnell Foundation for supporting this research with an Understanding Human Cognition Scholar Award to Andrew Shtulman. We would also like to thank Kidspace Children's Museum, Julieta Alas, Connor Allison, Taleen Berberian, Yukimi Hiroshima, Ellen McDermott, Phoebe Patterson, Khetsi Pratt, and Emma Ragan for their assistance with data collection.

References

- Austin, E. W., & Dong, Q. (1994). Source v. content effects on judgments of news believability. *Journalism Quarterly*, 71, 973-983.
- Chan, C. C., & Tardif, T. (2013). Knowing better: the role of prior knowledge and culture in trust in testimony. *Developmental Psychology*, 49, 591-601.
- Danovitch, J. H., & Lane, J. D. (2020). Children's belief in purported events: When claims reference hearsay, books, or the internet. *Journal of Experimental Child Psychology*, 193, 104-808.
- Dias, N., Pennycook, G., & Rand, D. G. (2020). Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media. *Harvard Kennedy School Misinformation Review*, 1, 1-12.
- Dumitru, E. A. (2020). Testing children and adolescents' ability to identify fake news: a combined design of quasi-experiment and group discussions. *Societies*, 10, 71.
- Einav, S., Levey, A., Patel, P., & Westwood, A. (2020). Epistemic vigilance online: Textual inaccuracy and children's selective trust in webpages. *British Journal of Developmental Psychology*, 38, 566-579.
- Flanagin, A. J., & Metzger, M. J. (2008). *Digital media and youth: Unparalleled opportunity and unprecedented responsibility* (pp. 5-27). MacArthur Foundation Digital Media and Learning Initiative.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 25-42.
- Gong, T., Young, A. G., & Shtulman, A. (2021). The development of cognitive reflection in China. *Cognitive Science*, 45, 12966.
- Jardina, A. & Traugott, M. (2019). The genesis of the birther rumor: partisanship, racial attitudes, and political knowledge. *Journal of Race, Ethnicity and Politics*, 4, 60-80.
- Kahan, D. M. (2017). Misconceptions, misinformation, and the logic of identity-protective cognition. *SSRN Electronic Journal*.
- Kim, A., Moravec, P. L., & Dennis, A. R. (2019). Combating fake news on social media with source ratings: The effects of user and expert reputation ratings. *Journal of Management Information Systems*, 36(3), 931-968.
- Loos, E., Ivan, L., & Leu, D. (2018). "Save The Pacific Northwest Tree Octopus": a hoax revisited. *Information and Learning Science*, 119, 514-528.
- Metzger, M. J., Flanagin, A. J., Markov, A., Grossman, R., & Bulger, M. (2015). Believing the unbelievable: understanding young people's information literacy beliefs and practices in the United States. *Journal of Children and Media*, 9, 325-348.
- Nadarevic, L., Reber, R., Helmecke, A. J., & Köse, D. (2020). Perceived truth of statements and simulated social media postings: an experimental investigation of source credibility, repeated exposure, and presentation format. *Cognitive Research: Principles and Implications*, 5, 1-16.
- OfCom, U. K. (2019). Children and parents: Media use and attitudes report. *London: Office of Communications London*. <https://www.ofcom.org.uk/research-and-data/media-literacy-research/childrens/children-and-parents-media-use-and-attitudes-report-2018>
- Ólafsson, K., Livingstone, S., & Haddon, L. (2013). *Children's use of online technologies in Europe: A review of the European evidence base*. London, England: EU Kids Online.
- Pennycook, G., Cheyne, J. A., Seli, P., Koehler, D. J. & Fugelsang, J. A. (2012). Analytic cognitive style predicts religious and paranormal belief. *Cognition*, 123, 335-346.
- Pennycook, G. & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188, 39-50.
- Pennycook, G. & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 25, 388-402.
- Stecula, D. A., & Pickup, M. (2021). Social media, cognitive reflection, and conspiracy beliefs. *Frontiers in Political Science*, 3, 62.
- Shen, C., Kasra, M., Pan, W., Bassett, G. A., Malloch, Y., & O'Brien, J. F. (2019). Fake images: The effects of source, intermediary, and digital media literacy on contextual assessment of image credibility online. *New Media & Society*, 21, 438-463.
- Shtulman, A., & McCallum, K. (2014). Cognitive reflection predicts science understanding. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 36, 2937-2942
- Tong, Y., Wang, F., Danovitch, J., & Wang, W. (2022). When the internet is wrong: Children's trust in an inaccurate internet or human source. *British Journal of Developmental Psychology*.
- UNICEF (2021, August). *Digital misinformation / disinformation and children*. UNICEF. <https://www.unicef.org/globalinsight/media/2096/file/UNICEF-Global-Insight-Digital-Mis-Disinformation-and-Children-2021.pdf>
- Vosough, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359, 1146-1151.
- Young, A. & Shtulman, A. (2020a). Children's cognitive reflection predicts conceptual understanding in science and mathematics. *Psychological Science*, 31, 1396-1408.
- Young, A. & Shtulman, A. (2020b). How children's cognitive reflection shapes their science understanding. *Frontiers in Psychology*, 11, 1247.