Local News in the Crosshairs: Audience Perceptions of Algorithmic News Sites*

David A. Beavers† Kevin DeLuca‡

August 4, 2025

Word Count: 9,878

Abstract

Amid years of declines in American local news media, recent advances in artificial intelligence have brought about a host of algorithmic news sites purporting to cover local issues. Can audiences distinguish between these algorithmic sites and real local news outlets? And can a digital literacy intervention improve discernment? We report results from a series of survey experiments in which participants are shown live websites from matched pairs of journalistic and real-world algorithmic news outlets. Participants randomly assigned to receive a digital media literacy intervention paid more attention to credibility-related features, such as bylines and "About Us" pages. However, the intervention had little effect on site evaluations or preference, as participants continued to base their choices on heuristics like layout and perceived bias. As generative AI lowers the cost of producing increasingly sophisticated imitation news sites, these results raise urgent concerns about democratic accountability and journalism's place in the digital environment.

*The authors thank Erik Peterson, Christopher Wlezien, Johanna Burger, Richard Fletcher, Andrew Trexler, Maureen Craig, Jamie Druckman, and Jeremy Freese, and audience members at the 120th American Political Science Association annual meeting and the 82nd Midwest Political Science Association annual meeting for their helpful comments and feedback. We thank Nat Markey, Zoe Kava, and Patricia Hughes for their research assistance. We also thank Eric Gober, Nicole Korp, Amy Shin, Ashley Marsh and the team at NORC; Matthew Dardet, Ethan Miles, and the DLABSS team; and the team at CloudResearch, all for their help with data collection.
†Harvard University, Department of Government. dbeavers@g.harvard.edu
‡Yale University, Department of Political Science. kevin.deluca@yale.edu

Years of newspaper closures (Abernathy, 2018) and cuts to reporting staffs (Peterson, 2021) have left a void in the American local news media. Existing research has demonstrated the far-reaching effects of this decline. The loss of a local paper leads to decreased political engagement (Hayes and Lawless, 2018, 2021; Rubado and Jennings, 2020; Schulhofer-Wohl and Garrido, 2013; Shaker, 2014); increased polarization (Darr, Hitt, and Dunaway, 2018); diminished sense of local identity (Mathews, 2022); fewer prosecutions for political corruption (Usher and Kim-Leffingwell, 2023); and even higher municipal borrowing costs (Gao, Lee, and Murphy, 2020).

As traditional journalism has contracted, a new class of online outlets has emerged to fill the informational void: algorithmically generated, low-cost "pink slime" news sites that mimic the appearance of local journalism but lack its institutional grounding, editorial oversight, and commitment to factual reporting (Bengani, 2019). These sites are often politically motivated, centrally produced, and strategically designed to appear trustworthy by leveraging familiar local branding and presentation cues (Anderson-Davis, 2024). Advances in artificial intelligence (AI) provide tools to create networks of such algorithmic sites at scale. In fact, ratings company NewsGuard at the time of this writing has identified 1,254 AI-generated news sites operating with little or no human oversight (Sadeghi et al., 2025). Recent survey evidence finds that more than 4 in 5 Americans are concerned about AI being used to spread misinformation in U.S. elections (Yan et al., 2025), and rightly so: in the lead-up to the 2024 U.S. presidential election, Russia reportedly used AI tools to create fake local news sites with innocuous sounding names like the Chicago Chronicle and New York News Daily as part of its influence campaign (Myers, 2024).

The rapid transformation in the local news environment and AI technologies poses urgent questions about the public's ability to discern real news from imitation. Existing research shows that individuals rely on heuristics like familiarity, source cues, and ideological congruence to evaluate news sources (Flanagin and Metzger, 2007; Sundar, 2008). However, less is known about how these heuristics operate in the emerging context of algorithmic news sites and how those evaluations translate into behavioral preferences. In this paper, we ask two core questions. First, how do audiences' evaluations of the credibility of journalist-produced local news sites differ from those of algorithmically generated pink slime sites that mimic their appearance? Second, does a digital media literacy intervention shift those evaluations and influence audience preferences? These questions matter for elections and democratic politics more broadly because misleading local news sites can shape public opinion and influence vote choices if they are seen as credible and trustworthy. If voters are unable to distinguish real news from algorithmic imitation, these sites risk distorting the information environment on which political judgment depends, undermining informed voting and weakening democratic accountability.

To evaluate how audiences respond to algorithmic news sites, we focus on Metric Media, a network that operates over 1,100 purportedly local news outlets across the United States (Bengani, 2021). The vast majority of content produced on these sites is algorithmically generated and low-quality, drawing on publicly available press releases, social media posts, and localizable data (Royal and Napoli, 2022). While these sites may offer some informational value, they have also produced politically motivated and at times misleading content for a fee (Alba and Nicas, 2020; Quéré and Jakesch, 2022).

We examine audiences' perceptions of real and pink slime news sites via a series of survey experiments: two pilot studies on convenience samples (n = 1,096) and a pre-registered experiment on a nationally representative sample (n = 1,613). In our pilot studies, we investigated whether evaluations differed when directing participants to live website links as opposed to a homepage screenshot more similar to the type of static stimuli used in existing research (e.g. a mock social media feed or series of headlines). Having found the live links approach to be both technically feasible and to better simulate the experience of encountering unfamiliar news while browsing online, we proceeded to use this approach in our main study, a survey experiment on NORC's AmeriSpeak panel fielded through Time-sharing Experiments for the Social Sciences (TESS). Participants were randomly assigned to see either a brief digital media literacy tipsheet or no intervention. They were then shown live links to a matched pair of news sites from their state, one a reputable journalistic outlet and the other a Metric Media site. Participants evaluated each site, gave open-ended responses on what features of the site they used to make their evaluations, and indicated which they preferred as a news source. To complement the experiment, we simultaneously scraped the homepages of the journalistic and Metric Media outlets used in the survey. This research design allowed us to assess how the intervention shapes credibility perceptions in the presence of various site attributes, such as bylines, layout, advertising, and the specific topics covered.

We find that, on average, participants evaluate the journalistic sites as more trustworthy, accurate, and reliable than their algorithmic counterparts. These differences are modest but consistent: the average evaluation gap favors the journalistic sites, and over half of all participants indicated preference for the journalistic outlet. Still, the margin was narrow: just 54% to 46% overall. Participants exposed to the digital media literacy intervention were around 4 percentage points more likely to prefer the journalistic sites than those in the control group. However, these effects were small in magnitude and did not reach statistical significance in multivariate models. Instead, participants' preferences appeared to be driven more by superficial features: whether the site looked familiar, had fewer ads, or covered topics they found engaging or personally relevant.

Open-ended responses suggest that the intervention worked as intended in terms of cognitive engagement: it significantly increased the salience of important site features such as named bylines, informative "About Us" pages, author biographies, and ethics policies, while reducing reliance on superficial site features. Still, attention to these considerations were overwhelmed by references to the topics covered and perceived bias. These results point to a key challenge for combating misinformation in the digital news environment. Even when audiences recognize legitimate journalism and are prompted to attend to meaningful signals of credibility, their actual choices are guided by shallow or misleading cues. This suggests that efforts to improve news discernment through light-touch interventions may have limited effectiveness unless paired with broader platform, policy, or structural reforms. As generative AI makes it easier to produce content that looks like professional journalism, this problem is likely to intensify.

Overall, our results paint a rather pessimistic picture for local journalism, particularly in the face of algorithmic competitors. Without addressing the long-term decline in newsroom resources, many local outlets may struggle to maintain their perceived value among audiences. At the same time, our findings highlight ways journalists might more effectively engage skeptical audiences, not only by emphasizing accuracy, but by reconsidering how they communicate the distinct role and value of journalistic content.

The Case of Metric Media

Metric Media operates a network of more than 1,100 algorithmic local "news" sites – the largest such network, to our knowledge, in operation. According to reporting from the *New York Times* and the Tow Center for Digital Journalism, Metric Media is part of a sprawling network of companies largely overseen by Brian Timpone, a former television reporter and Republican staffer in the Illinois statehouse (Alba and Nicas, 2020; Bengani, 2021). Metric Media's websites are designed to mimic traditional local news outlets, but the bulk of the sites' content is generated with little or no human intervention, drawing on press releases, tweets from local politicians, and local public data on subjects like gas prices, professional licenses, and FEC donations (Royal and Napoli, 2022; Tow Center for Digital Journalism, 2024).

While these sites mostly publish algorithmically generated content, they have also earned lucrative contracts from conservative campaigns, right-leaning causes, and companies to produce politically motivated content, written by freelancers, in what the *New York Times* calls a "pay-for-play" scheme (Alba and Nicas, 2020; Rafsky, 2022). A content analysis of Metric Media sites by Royal and Napoli (2022) found that the vast majority (97%) of content was algorithmically generated, but websites in Georgia, Wisconsin, and Arizona ramped up human-generated content about President Donald Trump's election fraud claims in the weeks following the 2020 election. The same constellation of groups behind Metric Media has blanketed swing-state voters with real-looking physical newspapers with names such as the Grand Canyon Times, Tucson Standard, and Kansas Catholic Tribune, all postmarked from the same Chicago address and all pushing partisan messaging. For instance, the Kansas Catholic Tribune emphasized anti-abortion messaging ahead of a 2022 state referendum, while the Grand Canyon Times promoted Republican Arizona Senate nominee Blake Masters (Tow Center for Digital Journalism, 2024). Leading up to the 2024 presidential election, voters in Michigan, Wisconsin, and Nevada reported receiving mailers from a Metric Media-owned publication branded as The Catholic Tribune (Richards and O'Matz, 2024; Kremer, 2024). The Catholic Tribune featured sensationalized and partisan headlines such as "How Many 'Sex Change' Mutilation Surgeries Occurred on Wisconsin Kids?" and repeatedly emphasized Robert F. Kennedy Jr.'s Catholicism in coverage of his endorsement of Donald Trump (Richards and O'Matz, 2024).

To more concretely illustrate how Metric Media sites operate, Figure 1 presents the homepage of Keystone Today, a Metric Media site covering Pennsylvania (Panel A), alongside the homepage of the *Harrisburg Patriot-News* (PennLive), a reputable journalistic outlet covering the same state (Panel B). The lead article on Keystone Today uses the site name in place of a byline and closely mirrors a press release issued by the Pennsylvania Department of Transportation three days earlier (see Appendix A for a side-by-side comparison). The "Opinions" section on the right reflects the site's ideological orientation, featuring op-eds authored by a lawyer for the Republican National Lawyers Association, a Republican gubernatorial candidate, and Trump's current secretary of agriculture. Notably, these are the only homepage articles with named bylines and most are more than two years old. By comparison, the *Patriot-News* homepage contains a wider range of articles on crime, sports, and local happenings, and less emphasis on politics.

Keystone Today's repurposing of a press release from the Pennsylvania Department of

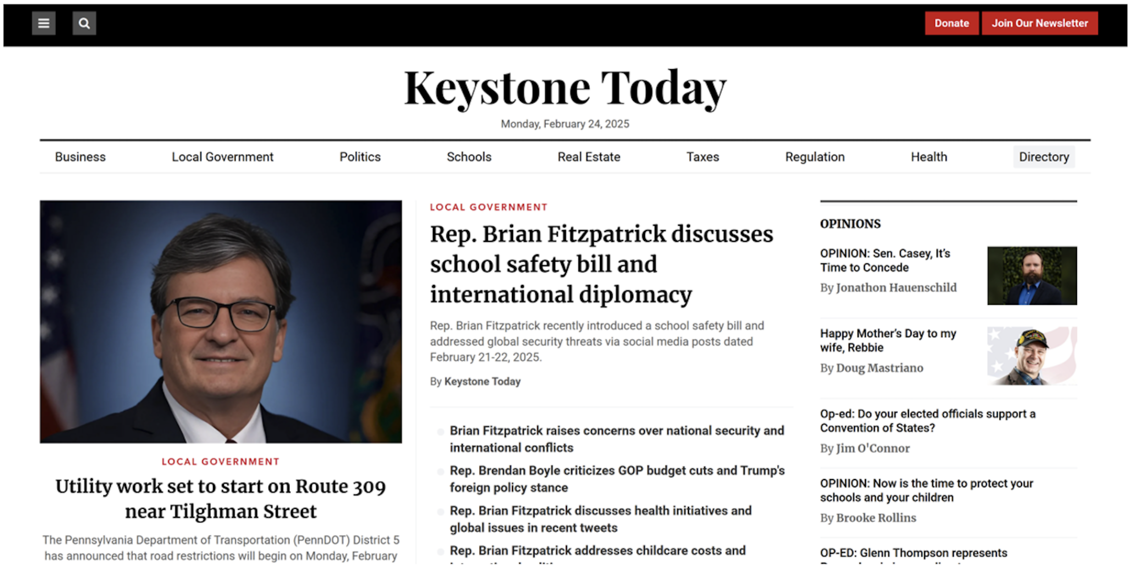

(a) Keystone Today (Metric Media)

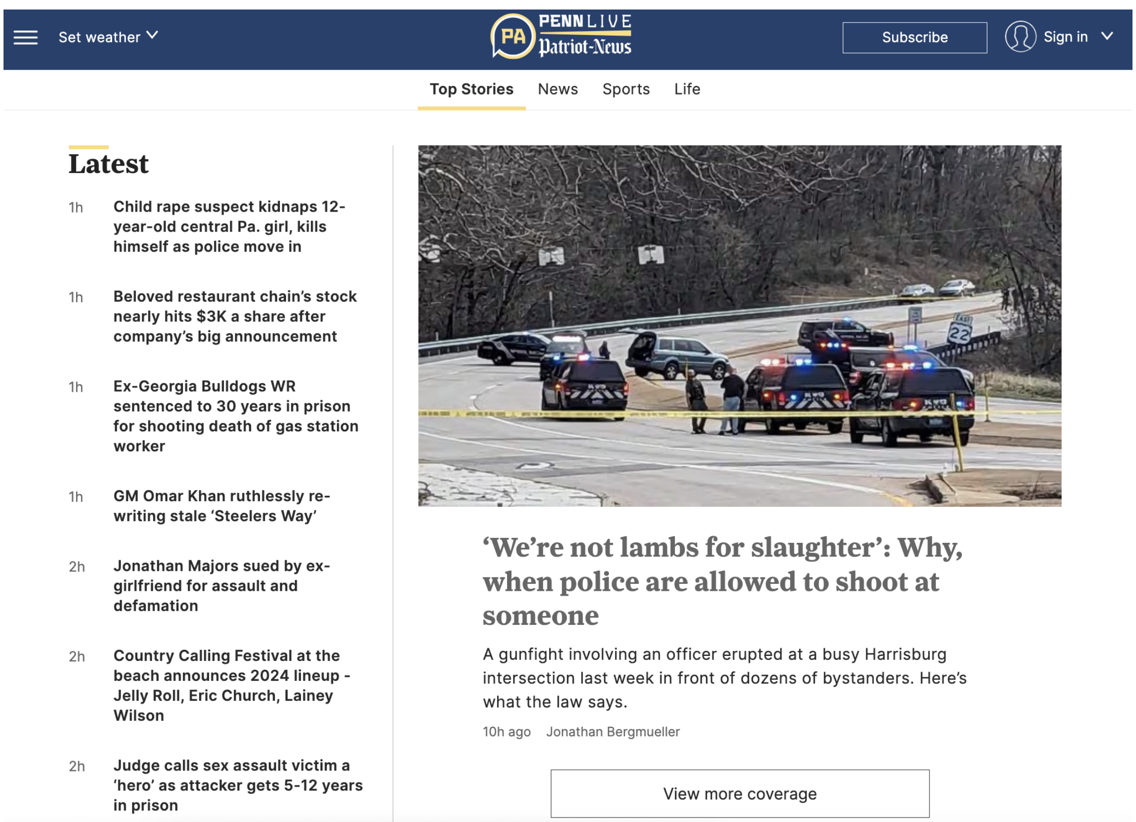

(b) Patriot-News (PennLive)

Figure 1: Metric Media and Journalistic Local News Website Appearances

Transportation as original reporting is typical of how Metric Media generates the bulk of its content. For example, on Old North News, a Metric Media site that covers North Carolina, the top stories on January 31, 2025 – the first day our survey was in the field – included a mix of government announcements and political statements, each lifted directly from press releases or recent tweets. Stories on Hurricane Helene relief aid and the reconvening of the General Assembly drew heavily on press releases from the National Federation of Independent Business, an association of small businesses that is generally aligned with the Republican Party; another quoted a trio of tweets from the Charlotte-Mecklenburg school district; and another drew on a press release from the state Department of Agriculture & Consumer Services about wildfire preparedness.

These examples help to motivate the normative implications of algorithmic news in general and Metric Media specifically. As an easily navigable repository of government press releases, these site offer something of a public service and could even be used as a tool by journalists. But while some of Metric Media's material is informative, it lacks the accountability and investigative scrutiny central to local journalism. Furthermore, by taking their raw data from ideologically-motivated groups and at times publishing overtly partisan content, typically without disclosure, algorithmic sites like Metric Media can bake political motivations into their content-generating process and can be used to manipulate public opinion at scale. As these types of outlets expand and generative AI lowers the cost of producing content for these sites, the ability of audiences to distinguish between traditional local news outlets and pink slime becomes politically consequential.

Source Credibility, Audience Perceptions of News Sites, and AI-Generated Content

Pink slime sites like Metric Media blur the boundary between authentic journalism and political propaganda, raising concerns about the public's ability to distinguish between trustworthy and deceptive sources. Understanding how audiences perceive and evaluate such outlets is therefore critical, especially in a digital environment where generative AI is lowering the cost of imitation.

A long line of research in political communication has investigated how source characteristics influence information retention and persuasive effects (Hovland and Weiss, 1951; Von Hohenberg and Guess, 2023), focusing particularly on perceived expertise and trustworthiness. Researchers have operationalized trust or credibility using a wide range of items including concern for public interest, facticity, fairness, lack of bias, respect for privacy, selectivity of topics, accuracy of depictions, honesty, and timeliness (Gaziano and McGrath, 1986; Kohring and Matthes, 2007; Yale et al., 2015). Furthermore, existing work has shown that credibility perceptions of news sites are based on a combination of source cues, content, and presentation (Flanagin and Metzger, 2007; Sundar, 2008).

These traits serve as the foundation for many digital media literacy efforts (Guess and Munger, 2023). Digital media literacy is the ability "to reliably assess the credibility of information encountered online ... [and] verify claims and look up answers to questions using a variety of strategies" (Guess et al., 2020, p. 113). Digital media literacy is positively associated with the ability to detect deepfakes (Barari, Lucas, and Munger, 2024), and individuals low in digital media literacy engage in poor web search practices, increasing their risk of exposure to low-quality and misinforming sources online (Aslett et al., 2024). Experimental evidence generally finds that digital media literacy interventions improve discernment between real and fake news headlines, articles, and social media posts (Guess et al., 2020; Tully, Vraga, and Bode, 2020; Epstein et al., 2021; Hameleers, 2022; Moore and Hancock, 2022; Altay, Angelis, and Hoes, 2024). A related literature finds that pre-bunking, or psychological inoculation, against misinformation is generally effective at reducing perceived believability of subsequently encountered false information (Traberg, Roozenbeek, and van der Linden, 2022; Kuru, 2025).

Advances in generative AI further complicate audiences' task of evaluating the credibility of unfamiliar sources online. Though local newsrooms for years have used algorithmic tools to facilitate journalists' work – for instance, in automating publication of earthquake notifications or combing through vast troves of bureaucratic documents (Quéré and Jakesch, 2022) – generative AI tools have the potential to displace journalists. Indeed, research shows that people generally struggle to discern between human- and AI-generated texts (Brown et al., 2020; Clark et al., 2021; Spitale, Biller-Andorno, and Germani, 2023). Experimental research has independently manipulated the content source (human- or computer-generated) and the attribution of the source (human or computer) to assess their effects on source credibility. Wölker and Powell (2021) find that source credibility did not differ for human- or algorithmically generated sports and finance content. Graefe et al. (2018) find that audiences evaluated algorithmically generated articles as more credible, but less readable, regardless of attribution. Clerwall (2014), van der Kaa and Krahmer (2014), and Henestrosa, Greving, and Kimmerle (2023) find muted differences between content attributed to a human or an algorithm. This body of work underscores a growing concern: surface-level credibility cues are increasingly easy to fake, and audiences often rely on them when forming judgments about news content.

Recent research in political science has begun to shed light on how audiences assess the credibility of unfamiliar pink slime sites. Much of this work focuses on the initial stage of credibility judgments, such as how users might evaluate an unfamiliar site name they encounter on a social media feed. Peterson and Allamong (2022) exposed participants to an identical text attributed to a prominent state or national newspaper or to an unfamiliar source – RT (Russia Today) or a local site affiliated with a Metric Media subnetwork. They found that participants prefer familiar sources but, conditional on exposure to unfamiliar sources, they have an equivalent influence on public opinion. Darr (2023) and Peterson et al. (2025) conducted a series of survey experiments focused on trust in local news outlets in which they manipulated the name of fictional news sites and asked participants to evaluate their credibility. They found that fictional news sites which convey a local affiliation in their title are evaluated at least as favorably as national news sites (people rely on a "local trust heuristic"). Peterson et al. (2025) specifically found that Metric Media site names, which tend to include local cues, are evaluated as more trustworthy relative to national news sites.

Overall, current research provides valuable insights into the initial stage of credibility assessments, such as how individuals react to seeing a Metric Media link in a social media feed. Our study focuses on the next stage: what happens if they decide to click that link? Existing experimental research has tended to isolate and test the effects of individual surface-level features, eliciting credibility perceptions after exposure to only a site name or a decontextualized piece of content. Although this is advantageous in securing experimental control, it lacks mundane realism, since in everyday practice an internet user would not need to evaluate the credibility of an unfamiliar website without having the ability to examine other features of the site that would be indicative of its legitimacy. To more closely simulate real-world browsing, we employ a "live links" design. Participants are directed to functional homepages of news sites, where they can interact with actual site content. We believe this design choice maximizes mundane realism (Druckman, 2022, pp. 51-61) without sacrificing external or internal validity, given the standardized design of Metric Media sites. Additionally, we believe by giving participants a more cognitively engaging task (live websites vs. isolated content or site names), our research design is more likely to achieve high experimental realism.

Experimental Design and Data Collection

To assess whether people can distinguish between local journalism and algorithmic pink slime, we conducted a series of survey experiments. First, we conducted a pair of pilot studies to explore how participants respond when shown live, interactive website homepages rather than static screenshots, examining whether the format of exposure influences their perceptions and choices. Second, we conducted a pre-registered national representative survey experiment to test the efficacy of a digital media literacy intervention. Across all experiments, participants were asked to evaluate whether each website was trustworthy, reliable, accurate, and informative and indicate which site they would prefer to use to learn about public affairs in their state.

Pilot Studies: Testing the Live Links Approach

Though credibility perceptions of news sites are based on a combination of source cues, content, and presentation (Flanagin and Metzger, 2007; Sundar, 2008), most existing research directs participants to atomized pieces of content or site names. To better simulate the kind of decision-making that occurs when individuals encounter unfamiliar news sites online, we conducted a pair of pilot studies, one via CloudResearch ($n = 323$) on June 13, 2024, and one via the Harvard Digital Lab for the Social Sciences (DLABSS; $n = 773$) on July 5-16, 2024.[1,2] Participants were randomly assigned either to a live links condition or a condition in which they were shown screenshots of the same two sites' homepages taking shortly before fielding the experiment. Similar to the atomized stimuli used in existing research, static screenshots of a site homepage obscure important features that participants use to assess site credibility, such as variety and depth of content, clickable bylines, navigability, external links, ads, or the presence of and details on "About Us" pages. These elements are central to how users navigate and interpret novel digital information, yet it is hard to convey the extent of their presence in a screenshot and impossible to do so with only site names.

We measured three sets of outcome variables to assess how participants evaluated the

[1]CloudResearch has been found to have high response quality relative to other online convenience panels (Stagnaro et al., 2024). Harvard DLABSS has also been used in a variety of research studies (see https://dlabss.harvard.edu/results) as a source of high-quality survey participants (Strange et al., 2019).

[2]The pilot studies and survey experiment were conducted in compliance with relevant laws. The research designs were approved by both authors' institutions' IRBs, and informed consent was obtained from all participants included in the study.

two websites. First, participants rated each site's trustworthiness, reliability, accuracy, and informativeness on a five-point Likert scale ranging from "strongly disagree" to "strongly agree." We recoded these items to a 0-1 scale, averaged the values for each type of site (journalistic or algorithmic), and computed the difference ("evaluation index").[3] The resulting index ranges from -1 to +1, where positive values indicate greater favorability toward the journalistic site. Second, participants were asked which of the two sites they would prefer as a source of news "if you wanted to learn about public affairs in your state in the future." We coded as 1 if participants preferred the journalistic outlet and 0 otherwise ("site preference").[4] Third, immediately after the evaluation task, we asked participants an open-ended prompt: "Still thinking about the two sites you just browsed, what elements of the sites did you have in mind when evaluating them?" One of the authors inductively developed a coding scheme (e.g. topics covered, perceived political bias, quality of reporting, presence of advertising) and hand-coded all responses. We provide more details about the coding scheme in the section discussing the results from the nationally representative survey experiment.[5]

Because directing participants outside the survey platform sacrifices some experimental control, we measured compliance in the live links condition by embedding JavaScript into the survey instrument to record whether participants clicked one or both links. 77.1% in the pooled sample did so. To account for the possibility that participants could have viewed the stimulus sites by copying and pasting the URLs without actually clicking them, we added to our nationally representative survey a second JavaScript code – adapted from Graham (2023)

[3]These indices were highly reliable. In the nationally representative survey experiment, the Cronbach's α values were 0.910 for the journalistic sites and 0.927 for the algorithmic sites.)

[4]In the pilot studies, we included a "no preference" option to capture ambivalence. We dropped this option on the main study.

[5]For the full survey instrument, see Appendix B. For additional details on the pilot study results, see Appendix C.

– that recorded when participants entered a new browser tab. Furthermore, we did not see a significant increase in survey drop-off in the live links condition, which could happen if participants clicked the links and never returned to the survey.[6]

Pilot Results

We found sizable and statistically significant differences in how participants evaluated the same news outlets depending on whether they viewed a static image or visited the live site. These results reinforced our decision to prioritize ecological validity over full experimental control, suggesting that the live links design better captures how individuals assess unfamiliar news sources in real-world settings.

Table C1 presents a summary of the pilot study results. Participants in the live links condition expressed more favorable evaluations of the journalistic sites (0.591 vs. 0.538) and slightly less favorable evaluations of the Metric Media sites (0.491 vs. 0.508) compared to those in the image-only condition, resulting in a substantially larger evaluation index difference (+0.100 vs. +0.030). Similarly, site preferences were more polarized in the live links condition: 42.8% of participants preferred the journalistic site, compared to 29.4% in the image condition. Participants in the image condition were more likely to prefer the Metric Media site (36.3% vs. 31.3%), illustrating how screenshots obscure important cues that help users detect low-credibility sources. Multivariate OLS regression results in Table C1 confirm that these differences are statistically significant controlling for individual-level demographics.

The open-ended responses also reveal meaningful differences in what participants attended to when evaluating news sites across conditions. Participants in the live links condition were more likely to mention indicators like bias (38.7% vs. 22.0%), topic coverage (39.4% vs. 33.0%), and quality (17.7% vs. 10.8%), whereas those in the images condition were somewhat more likely to mention site layout (17.4% vs. 14.9%). Most significantly, attention

[6]Overall, 87.4% of participants completed the survey in the live links condition, compared to 89.0% in the screenshots condition.

Table 1: Summary of Pilot Study Results

	Images	*Live Links*
Evaluation Index		
Journalistic site	0.538	0.591
Metric Media site	0.508	0.491
Difference (Journo – MM)	0.030	0.100
Site Preference		
Prefer Journalistic	29.4%	42.8%
Prefer Metric Media	36.3%	31.3%
No Preference	33.9%	25.4%
Net Difference (Journo – MM)	-6.8 pp	+11.4 pp
OLS Estimate of Live Links Effect (vs. Images)		
Evaluation Index Difference	0.054**	
	(0.017)	
Journalistic Site Preference	0.100**	
	(0.034)	
Factors Mentioned in Open-Ended Responses		
Topics	33.0%	39.4%
Bias	22.0%	38.7%
Layout	17.4%	14.9%
Familiarity	15.8%	8.4%
Quality	10.8%	17.7%
Metadata	5.9%	5.6%
Advertising	1.4%	6.5%

Notes: The reported means are pooled (weighted) across CloudResearch and DLABSS pilot samples. Evaluation Index is the mean of four items (trustworthy, reliable, accurate, informative), rescaled from 0–1. Site Preference indicates the percentage of participants who preferred each type of website (or had no preference between them). Site Preference is the proportion of participants selecting each option. OLS estimates control for demographics; standard errors in parentheses. *$= p < .05$, **$= p < .01$. The bottom section of the table shows the proportion of respondents whose open-ended comments mentioned in each selected category, based on manual hand-coding. Total participants $= 1,096$ (CloudResearch $n = 323$; DLABSS $n = 773$; see Appendix C for results broken down by sample).

to metadata – the label we give to mentions of factors such as external reviews, "About Us" pages, and informative bylines – were low in both conditions (5.9% vs. 5.6%). In the main experiment, we test a digital literacy tipsheet designed to increase attention to these cues.

Together, the pilot results reinforce the value of the live links design in approximating real-world browsing behavior and support the potential of targeted interventions to redirect user attention toward features like site metadata that could improve discernment and alter preferences.

Nationally Representative Experiment: Assessing the Digital Media Literacy Intervention

We fielded a pre-registered survey experiment on NORC's AmeriSpeak panel from January 31 to February 16, 2025.[7] Our sample is broadly representative of the adult American population.[8] Participants were instructed to follow two links – one to a Metric Media site and one to a traditional journalistic news site based in their state – one at a time, and in random order. As in the pilot studies, we measured site evaluations (trustworthy, reliable, accurate, and informative), a forced-choice site preference item, and an open-ended prompt asking what site elements participants considered when forming their evaluations. We also passively measured compliance (link-clicking and entering a new browser tab) and asked participants whether they were familiar with each site.

Before visiting the websites, participants were randomly assigned to see a digital media literacy tipsheet, as in Figure 2, or to a control group which saw only the first two introductory sentences of the prompt, and none of the tips. The intervention was lightly adapted from, and attributed to, the American Library Association (ALA).[9] This tipsheet has not been used in

[7]Our pre-analysis plan was submitted to OSF on February 2, 2025, prior to receipt of any data (https://osf.io/drpuv/).

[8]See Appendix D.I for additional sample characteristics and survey details.

[9]Our only changes were to expand on items 5 and 8 in the original tipsheet, as we felt these were lacking in specificity in the original. See 'summary of tips' at right: https://libguides.ala.org/InformationEvaluation

existing research, to our knowledge. Because the digital media literacy interventions tested in previous studies were aimed at improving discernment of true from false information, that approach is less well-suited to the present context, since Metric Media sites contain mostly real (if low-quality and at times partisan) information packaged as local news. Nonetheless, the success of digital media literacy tipsheets in adjacent domains leads us to hypothesize that it will improve discernment between journalistic and algorithmic local news sites, as well. We expect the intervention to be effective by shifting attention from superficial site features (such as layout) toward cues (such as site ownership, existence of bylines, and sources cited in articles) that more reliably differentiate authentic journalistic outlets from algorithmically generated ones – but which were mentioned by very few participants in our pilot studies.

Next, we are going to ask you to browse two news sites. One or both sites may be unfamiliar to you. The American Library Association offers the following tips for evaluating unfamiliar news sites:

1. **Consider the source.** Click away from the story to investigate the site, its mission and its contact info.
2. **Read past the headline.** Headlines can be outrageous in effort to get clicks. Go beyond headlines.
3. **Assess the credibility of the author.** Do a quick Google search on the author. What is their expertise? What organization do they represent?
4. **Look at the links and sources supporting the article.** Click those links. Determine if the subsequent information supports the story. Consider the reliability of the sources.
5. **Check the date.** Is the story out of date or is it current?
6. **Consider that the item might be satire.** If it seems too outlandish, it might be satire. Do some quick research on the site and author to find out.
7. **Consider that it might be promotional.** Is the purpose of the site to sell a product?
8. **Check your biases.** Try to evaluate the story on its merits, rather than whether you want it to be true or not.
9. **Search other news outlets** to see if the news is widely reported.

Figure 2: American Library Association Tipsheet Intervention

Where available, we directed participants to their state's Metric Media "hub" site, the set of sites that serve as centralized sources of content for other Metric Media publications

(Royal and Napoli, 2022). For the two states that lack such a site (Michigan and North Dakota), we chose a site purporting to cover the state's biggest city. As a matched site for each state, we chose the largest newspaper in the state capital, except where that paper is the largest newspaper in the state (e.g. the *Boston Globe*). In that case, we picked the second-largest newspaper in the state capital, where one exists, or another medium-circulation newspaper that covers state politics. The intent was to avoid picking newspapers that nearly all participants would be familiar with (e.g. the *New York Times*) as a comparable set of matched state sites. We believe this is especially important in our experimental setting, as the goal is to evaluate how participants assess local news sources that are unfamiliar to them. Using well-known outlets would not only create an unbalanced comparison with the relatively obscure Metric Media sites, but would also undermine our ability to observe how people rely on site features (rather than familiarity) when making credibility judgments. Appendix B.II lists all stimulus sites used in the experiment. Appendix B.IV lists ownership data for the non-Metric Media sites.

Content Analysis

While our survey experiment was in the field, we manually scraped content from the 102 stimulus websites every other day to collect a simulated week of data (e.g. one Monday, one Tuesday, one Wednesday, etc.). We collected the five most prominently displayed articles each day, recording the headline, full text, byline, date of publication, and URL, for a total of 3,570 articles.

To classify the topic of each article, we use the Political DEBATE model, a natural language inference (NLI) model designed for text classification (Burnham et al., 2024).[10] We use the model to generate a binary label indicating whether each "premise" (headline)

[10]Specifically, we use the untrained "zero-shot" large (v1.0) Political DEBATE model. Prior research demonstrates that even untrained NLI models perform well on basic topic classification tasks (Burnham et al., 2024).

supports a set of hypotheses (e.g., "This text is about immigration"). Headlines may be assigned multiple topics. The topics we evaluate are: business, climate/environment, crime, the economy, education, foreign policy, gender/gender issues, healthcare, immigration, local government, race/racial issues, reproductive rights, and sports. The three most common predicted classes for the journalistic sites were sports (21.6%), local government (21.5%), and education (16.1%), while the three most common classes for Metric Media sites were the economy (23.6%), education (20.6%), and local government (20.1%). We next use these topic labels to estimate the proportion of soft news on each site, following Baum (2002, p. 92), who operationalizes soft news as "a set of story characteristics, including the absence of a public policy component, sensationalized presentation, human-interest themes, and emphasis on dramatic subject matter, such as crime and disaster." In accordance with this definition, we categorize content labeled as being about crime or sports as soft news.

Table 2 provides descriptive results from our content analysis. A typical prominently displayed article on the Metric Media sites was 3 days old, 260 words long, and had no named byline. By comparison, a typical article on the journalistic sites was less than 24 hours old, 645 words long, and had a named byline. The journalistic sites had an average of 4.86 unique articles per day out of a possible 5 that we collected, implying that the five most prominent articles turned over within no more than 48 hours on these sites. On the Metric Media sites, more than half of the five most prominently placed articles had been in the same placement two days before. On the other hand, Metric Media sites contained a much lower proportion of news labeled as soft (20.8% vs. 35.4%).

The topic distribution also reveals notable differences between journalistic and Metric Media sites. Journalistic outlets devote more coverage to sports and crime, while Metric Media sites feature more articles on the economy, business, and race. Both types of outlets cover local government and education at similar rates, but the overall topic mix suggests that Metric Media sites emphasize political and economic content over community-oriented or lifestyle reporting.

Table 2: Content Analysis of Journalistic and Algorithmic News Sites

	Journalistic	**Algorithmic**
Median article age (days)	0 days	3 days
	(1.34)	(10.40)
Median word count	645	260
	(578)	(227)
Unique articles per day	4.86	2.29
	(0.60)	(1.08)
Proportion with named byline	0.902	0.004
	(0.297)	(0.063)
Proportion of soft news	0.354	0.208
	(0.478)	(0.406)
Article Topics		
Sports	21.6%	13.6%
Local Government	21.5%	20.1%
Education	16.1%	20.6%
Crime	15.2%	7.2%
Race	11.3%	17.2%
Business	8.2%	15.5%
Healthcare	6.0%	6.8%
Economy	5.5%	23.6%
Gender	5.2%	9.9%
Immigration	4.7%	2.6%
Environment	3.8%	2.5%
Foreign Policy	2.9%	3.3%
Reproductive Rights	0.5%	0.4%

Notes: Measures are based on content collected during the experimental survey window. Values in parentheses are standard deviations. Topic proportions are the percentage of all coded articles from each site type that mention each topic, and articles can be assigned multiple topics. "Soft news" includes articles classified as sports or crime.

Results

We begin by assessing whether the digital media literacy intervention succeeded in its intended goal: directing participants' attention toward features that distinguish authentic journalistic outlets from algorithmic pink slime sites. To do so, one of the authors and an undergraduate research assistant hand-coded all open-ended evaluations using the inductive coding scheme, first developed in the pilot studies. This included 12 codes: (1) Advertising (the amount of pop-up ads, paywalls, prompts to subscribe); (2) Bias (whether the content was politically biased, accurate, separated fact from opinion, objective); (3) Blanket distrust (statements expressing broad cynicism about the media rather than specific evaluations of the two stimulus sites); (4) Clickbait (gimmicky or sensationalized coverage appealing to emotions); (5) Familiarity (prior knowledge of the sites, reputation or popularity, comments on the sites' names); (6) Layout (ease of navigation, aesthetically pleasing); (7) Metadata (references to bylines, about-us pages, dates of the articles, ethics policies, author biographies, and mentions of looking elsewhere on the web for information about the site's credibility, including reviews and social media mentions); (8) Multimedia (references to images or video content); (9) Norms (adherence to journalistic practices such as proper sourcing, hyperlinking to documents or original reporting, fact-checking); (10) Quality of writing and reporting (lack of spelling or grammatical mistakes, thorough reporting); (11) Topics (references to the specific content covered, choice and diversity of topics); and (12) Usefulness (whether the content was interesting, informative, or personally relevant). Intercoder reliability was high; the average Krippendorff's α across all codes was 0.826 (Lacy et al., 2015). Intercoder reliability was over 0.7 for all codes except for Norms, which was dropped from the analysis. We created a composite measure of each code that takes a value of 1 if both coders assigned that code, 0.5 if only one coder assigned that code, and 0 if neither coder assigned that code. Appendix B.III lists the full set of coding instructions and reports intercoder reliability for each code.

Table 3 summarizes respondents' open-ended responses about what site elements influenced their evaluations. The "Treatment" and "Control" columns in Table 3 display the percentage of respondents in the treatment and control groups, respectively, coded as mentioning each element, weighted by composite coder agreement. Consistent with adjacent research exploring the considerations people have in mind when answering survey items about trust in the media (Ladd, 2012; Newman and Fletcher, 2017), perceived site bias was top of mind for many respondents. References to the topics covered, the quality of writing and reporting, the presence of advertising, and site layout were also frequently mentioned.

Table 3: Impact of Treatment on Open-Ended Response Evaluation Criteria

Open-Ended Response	Treatment (%)	Control (%)	Difference	p-value
Topics	25.5	29.5	**-4.0**	0.054^{+}
Bias	22.5	21.8	0.7	0.741
Quality	8.2	7.3	0.9	0.472
Advertising	8.0	6.1	1.9	0.130
Layout	7.7	12.1	**-4.4**	0.002^{**}
Metadata	7.2	2.5	**4.7**	$<0.001^{**}$
Usefulness	6.4	5.1	1.3	0.235
Familiarity	6.2	3.8	**2.4**	0.018^{*}
Clickbait	4.2	4.2	0.0	0.964
Multimedia	3.4	6.8	**-3.4**	0.002^{**}
General Distrust	2.0	2.7	-0.7	0.270

Notes: Entries report the percentage of participants in each group whose open-ended responses mentioned the listed theme. p-values are from two-sample tests of proportions. $^{+}$ $p < .10$, * $p < .05$, ** $p < .01$.

Compared to the control group, treated participants were significantly more likely to mention metadata cues – such as bylines, author information, or article dates – which may reflect editorial standards that can indicate source credibility and legitimacy. Mentions of metadata rose from 2.5% in the control group to 7.2% among treated respondents ($p <$ .001). Likewise, references to site familiarity increased significantly following exposure to the intervention. Notably, numerous participants in the treated group mentioned performing external searches to learn more about the unfamiliar sites or to verify the information they reported, comments that were largely absent from control-group participants. One treated respondent wrote: "I did a Google search for the one I was less familiar with, Old North

News, and let the reviews and low number of social media interactions influence my opinion of it." Another wrote: "In looking up the stories, the Colorado Sun's articles are verifiable across other outlets and Google, whereas the Centennial State News articles have no listed dates of authorship nor verifiable accounting of the stories from other outlets nor Google." Conversely, participants in the control condition were more likely to comment on superficial elements like layout or multimedia features, and also placed relatively greater emphasis on the specific topics covered by each site.

That nearly three times as many participants in the treatment group mentioned metadata as those in the control group provides compelling evidence that the tipsheet intervention successfully shifted attention toward source credibility cues like named bylines and "About Us" pages, as well as independent verification. However, even among treated participants, references to these cues remained relatively infrequent. Mentions of metadata rose from the least common code in the control group to only the sixth most common in the treatment group, still behind more surface-level site features such as layout, advertising, and perceived bias. This highlights both the promise and the limits of light-touch interventions such as the one used here. While the tipsheet moved participants' attention in the intended direction, people continue to rely heavily on features that are easier to notice but less diagnostic of source legitimacy. Given the ubiquity of elite criticism about the media (Ladd, 2012; Archer, 2020) and the well-established tendency for partisans to find bias where they look for it (Vallone, Ross, and Lepper, 1985), reducing perceptions of site bias is easier said than done. Still, these results confirm that the digital media literacy tipsheet functioned as intended: it successfully shifted attention away from surface-level design elements and toward cues more indicative of source credibility.

The open-ended response data provide additional insight into which site characteristics most strongly predict participant preferences. As shown in Table 4, participants who referenced certain features – metadata, multimedia, and bias – were significantly more likely to prefer the journalistic site over the algorithmic alternative. Specifically, participants who mentioned metadata were 11.2 percentage points more likely to prefer the human-produced site than those who did not, a statistically significant difference ($p = 0.047$). This is further evidence of the efficacy of the digital media literacy intervention: the tipsheet increased the frequency with which participants cited metadata, and these participants were also more likely to prefer the journalistic outlet. Statistically significant differences also emerged for mentions of bias and multimedia features. The bias effect may have been amplified by charged content circulating on Metric Media sites during the survey window, including a widely featured story about transgender high school athletes, which several respondents flagged as partisan. Multimedia mentions, meanwhile, often referenced the repeated or low-quality images on Metric Media homepages, which reduced the perceived professionalism of those outlets.

Table 4: Difference in Site Preference by Evaluation Criteria

Positive Effects (Journalistic Site Preferred)	Δ Preference	*p*-value
Multimedia	**0.141**	0.008**
Metadata	**0.112**	0.047*
Bias	**0.102**	<0.001**
Norms	0.108	0.258
Familiarity	0.091	0.093
Distrust	0.082	0.303
Quality	0.071	0.111
Layout	0.023	0.563
Usefulness	0.028	0.584
Topics	0.017	0.536
Negative Effects (Metric Media Site Preferred)		
Advertising	**-0.197**	<0.001**
Clickbait	-0.085	0.161

Notes: Δ Preference is the difference in the percentage of participants who preferred the human-produced site between those who mentioned each code and those who did not. Positive values indicate greater preference for the journalistic site; negative values indicate greater preference for the algorithmic site. * $p < .05$, ** $p < .01$.

At the other end of the spectrum, the largest negative preference effect was for mentions of advertising. Participants who flagged ads, paywalls, or prompts to subscribe were 19.7 percentage points less likely to prefer the journalistic site, a striking and highly significant difference. In open-ended comments, many participants expressed frustration with the journalistic sites' advertisements. As one respondent summed it up: "The first [site] has some stuff which requires subscription. I don't like it. Why should I pay [for] some news while there are many free resources on [the] internet[?]" These types of barriers are absent on the pink slime sites, which feature minimal advertising and no paywalls. Though the Metric Media sites have a "Donate" button at the top right of their home pages (see Figure 1, Panel A), this is unobtrusive compared to the paywalls, banner ads, pop-ups, and prompts to subscribe that many of the human-produced sites use.

This highlights a crucial disadvantage for traditional journalistic operations. Algorithmic sites have low operational costs because they don't rely on humans to produce content. It is not feasible for most traditional news outlets to drop advertising from their sites, but these features may drive away audiences who are not accustomed to paying for information online. Journalists can, however, do a better job of conveying their business model to potential audiences, particularly to digital natives with limited exposure to the norms and costs associated with professional reporting.

Tipsheet Effects on Evaluations and Preferences

Though the digital media literacy intervention was effective in steering users toward cues associated with source credibility, it had only modest effects on site evaluations and preferences. The increased attention to metadata features – author information, article dates,

[11] Values throughout are for the full sample, rather than the subset of "compliers." Per our pre-analysis plan, we label as compliers those who either clicked both links or entered a new browser tab on both pages where we instructed participants to evaluate a news site. 56.6% of our sample pass this compliance check – a good bit lower than in our pilot studies. Though this is on its face somewhat concerning, the results are consistent whether analyzing the full sample or the subset of compliers. We provide analyses on the subset of compliers in Appendix D.II.

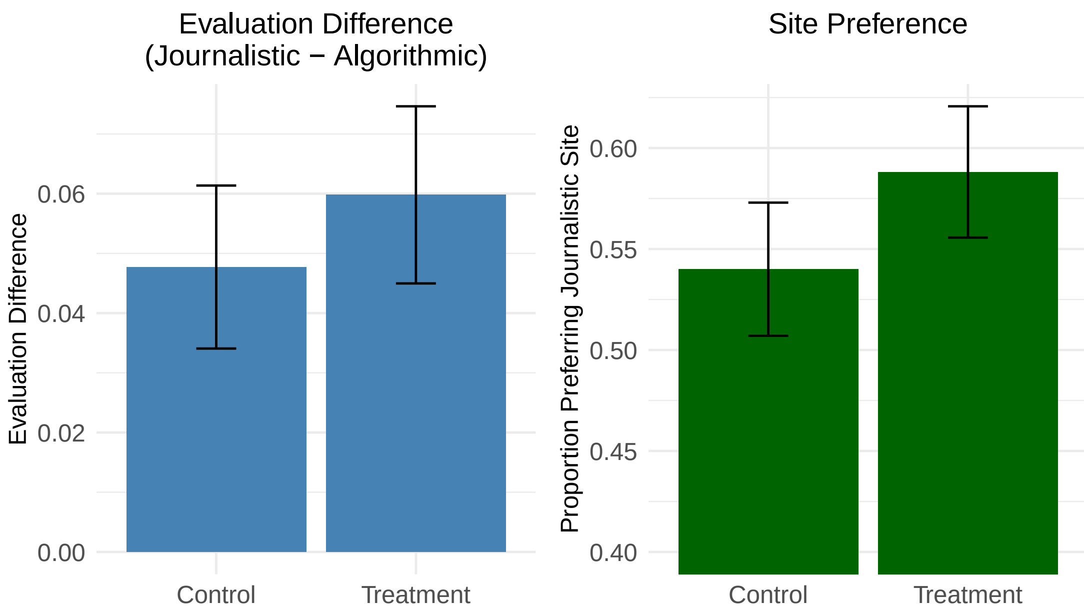

Figure 3: Evaluation Differences and Site Preferences by Treatment Group

Notes: Left panel bars represent the mean difference in evaluations of the journalistic and algorithmic sites by treatment condition. Right panel bars represent the proportion of respondents who preferred Journalistic sites over Metric Media sites in the forced-choice question. Error bars are 95% confidence intervals.

etc. – did not translate into statistically significant shifts in either the evaluation index or site preference measures. As shown in Figure 3, in the full sample, the average evaluation difference between the journalistic and algorithmic sites was slightly larger in the treatment group than in the control group, and treated participants were more likely to indicate preference for the journalistic site.[11] But neither difference reaches conventional thresholds for statistical significance.

Still, the direction of these effects is consistent with the theory behind the intervention. Participants overall preferred the human-produced news sites: 56.4% chose the journalistic outlet, compared to 43.6% who preferred the algorithmic one. This gap widened in the treatment group (58.8% vs. 41.2%) relative to the control group (54.0% vs. 46.0%). Similarly, the average evaluation difference (how much higher the journalistic site was rated relative to the Metric Media site) between the two site types was 0.062 in the treatment group versus 0.048 in the control. Treated participants on average rated *both* sites more positively: evaluation index values for treated participants were 0.603 vs. 0.542 for journalistic and algorithmic sites, compared to 0.580 vs. 0.533 in the control group.[12] These are small but meaningful directional effects. The intervention moved participants' judgments in the predicted direction, but by too modest a magnitude to conclude that the tipsheet significantly improved audience discernment.

To further assess the treatment effects, we performed weighted OLS regressions with standard errors clustered at the state level and state fixed effects, controlling for individual-level characteristics (sex, age, education, race, household income, ideology, party ID, main source of news, media trust, and home internet access) as well as self-reported familiarity with each site, ownership type of the journalistic site, an indicator for the presence of a paywall on the journalistic site, an indicator variable for stimulus order (which site was displayed first in the survey), and the proportion of soft news on the journalistic and Metric Media sites. Appendix D.II has more details on the primary regression specification and additional results with alternative estimation strategies and controls.

The point estimates in Table 5 indicate that the tipsheet treatment increased preference for the journalistic sites by 3.8 percentage points (column 1) and increased the difference in site evaluations for journalistic over Metric Media sites by 0.007 points (column 2), though neither estimate is statistically significant. Importantly, site evaluations (and the differences between them) are highly predictive of site choice, as shown in column 3 of Table 5. Given that the standard deviation for the difference in site evaluations is equal to about 0.22, a one standard deviation increase in the evaluation index (the difference in the evaluation index of journalistic sites and the evaluation index of Metric Media sites) is associated with an approximately 21 percentage point increase in the proportion of respondents who prefer the journalistic sites. This underscores a key implication of our findings: the intervention

[12] It is noteworthy that we did not find any negative spillover effect on reputable sources from the digital literacy intervention, as some previous research has found (Guess et al., 2020; Kuru, 2025).

plausibly could have shifted site preferences, had it more substantially altered participants' evaluations of the sites' trustworthiness, reliability, accuracy, and informativeness. In that sense, the problem may not lie with the logic of the intervention, but rather with its modest influence on site evaluations. A stronger or more effective treatment might yet hold promise for improving audience discernment between real and algorithmic news sources.

Table 5: Main Treatment Effects on Evaluation and Choice

	(1) Site Choice	**(2) Evaluation Difference**	**(3) Site Choice**
Treatment (Tipsheet)	0.038	0.007	0.025
	(0.040)	(0.016)	(0.035)
Evaluation Difference	—	—	0.988**
			(0.049)
Observations	1,428	1,493	1,419
Mean of Outcome	0.564	0.053	0.564
Adjusted R^2	0.098	0.114	0.261
Adjusted-Within R^2	0.058	0.051	0.230
State Fixed Effects	✓	✓	✓
Demographic Controls	✓	✓	✓
Clustered SE (State)	✓	✓	✓

Notes: Demographic controls include: gender, education level, 5-point ideology, 3-point income, and race. Additional controls include: main source of news, self-reported familiarity with each website, general media trust, difference in the proportion of hard news stories between sites, and site ownership type. Evaluation difference is the difference in a standardized index of four credibility items between journalistic sites and Metric Media sites. Standard errors clustered at the state level. * $p < .05$, ** $p < .01$.

Table 6 provides insight into other predictors of participants' evaluations of journalistic and algorithmic news sites. The outcome variables are participants' evaluations of the journalistic site (column 1), evaluations of the Metric Media site (column 2), and the difference between the two (column 3). Participants who reported recognizing a site, whether journalistic or algorithmic, consistently rated that site more favorably. Although it is likely that, due to flawed recall or social desirability bias, some participants mistakenly report familiarity with the Metric Media site, this result is nonetheless consistent with prior research showing that familiarity boosts perceived credibility and trustworthiness (Metzger, Flanagin, and Medders, 2010).[13] The powerful role of familiarity underscores a key challenge in combating misinformation: audiences may form positive evaluations of low-quality or deceptive outlets simply because they appear recognizable or resemble traditional media brands.[14]

Another important factor shown in Table 6 is general media trust. Higher levels of media trust are associated with more favorable evaluations of *both* journalistic and algorithmic news sites. That is, people who are dispositionally more trusting of the media appear to extend that trust indiscriminately across outlets, regardless of their production quality or ownership structure. This finding aligns with work suggesting that media trust acts more as a generalized orientation than as a discriminating tool for evaluating specific sources (Prochazka and Schweiger, 2019), and it may help explain why even low-quality outlets like Metric Media can earn reasonably high marks from the public.[15] While media literacy interventions may

[13]Participants on average were three times as likely to report being familiar with the human-produced news site than the Metric Media site (36.0% vs. 11.8%). Existing research using three months of digital trace data found that only 3.7% of respondents had visited a set of domains including Metric Media sites (Moore et al., 2023). The higher rate of familiarity in our sample could be caused by social desirability bias, by the local geographic cues in the Metric Media website names (Peterson et al., 2025), or due to flawed recall. It is also possible that Metric Media sites have become more popular since Moore et al. (2023) collected their data in fall 2020.

[14]We also include a binary indicator for "Reads Local Newspaper," which is 1 if a respondent indicates local newspapers as their preferred source of news. This is not statistically significant, likely due to such a small proportion (2.9%) of respondents and also due to the fact that the journalistic website they were presented with was often not their particular "local" newspaper, since papers were generally chosen from the state capital.

[15]See also Peterson and Allamong (2022), which finds that people are often less negative toward

help redirect attention toward credibility-relevant cues, these findings suggest that baseline levels of trust and familiarity remain key influences on how people evaluate news content.

Table 6: Predictors of Site Evaluations and Evaluation Differences

	(1) Journalistic Evaluation	**(2) Algorithmic Evaluation**	**(3) Evaluation Difference**
Treatment (Tipsheet)	0.019	0.012	0.006
	(0.014)	(0.011)	(0.015)
Familiar w/ Algorithmic Site	0.012	0.047^{**}	-0.033^{+}
	(0.013)	(0.016)	(0.018)
Familiar w/ Journalistic Site	0.069^{**}	0.015	0.055^{**}
	(0.011)	(0.015)	(0.015)
Reads Local Newspaper	0.023	0.033	-0.009
	(0.026)	(0.023)	(0.039)
Media Trust	0.106^{**}	0.052^{**}	0.053^{**}
	(0.012)	(0.013)	(0.014)
Observations	1,500	1,495	1,493
Adjusted R^2	0.174	0.128	0.112
Adjusted-Within R^2	0.167	0.061	0.049
State Fixed Effects	✓	✓	✓
Demographic Controls	✓	✓	✓
Clustered SE (State)	✓	✓	✓

Notes: Dependent variables are evaluation scores for journalistic and algorithmic sites (cols. 1–2), and their difference (col. 3). Demographic controls include: gender, education level, 5-point ideology, 3-point income, and race. Additional controls include: main source of news, self-reported familiarity with each website, general media trust, difference in the proportion of hard news stories between sites, and site ownership type. Evaluation difference is the difference in a standardized index of four credibility items between Journalistic sites and Metric Media sites. Standard errors clustered at the state level. $^{+}$ $p < .10$, * $p < .05$, ** $p < .01$.

unfamiliar news sources than expected, a dynamic that may be mediated by overall trust in media, and which can also help to explain the relatively high ratings and selection of Metric Media sites.

Additional Analyses: Effect Heterogeneity and Topic Prevalence

To assess whether small overall treatment effects mask large effects among certain subsets of the sample, we repeat regressions using models stratified by ideology, party identification, education, and media trust. Treatment effects remain small and mostly non-significant across all evaluated subgroups and outcome variables, as shown in Figure 4. Treatment effects on site preference are largest for moderate respondents ($p < 0.01$) and independents ($p < 0.10$), though the statistically significant effect on moderates disappears in the compliers-only subsample.[16] Given independents' typically lower levels of media consumption (Toff, Palmer, and Nielsen, 2024) – indeed, independents in our sample were 6.7 percentage points less likely than partisans to indicate familiarity with the journalistic news sites and were more likely to indicate their preferred source of news was social media or "other" – it makes sense that a light-touch intervention would have a greater effect on their site preferences.

Finally, because we collected content from the sites concurrently to our survey being fielded, we can examine how the topics covered by the journalistic and Metric Media sites affected participants' site preferences. Appendix D.II has the details. Overall, the topics covered explain about 6.5% of the variation in site preference, with the largest effects on Democrats. Consistent with participant self-reports that the choice of topics impacted their site evaluations, we find some evidence that the content of the sites they viewed impacted their site preferences. In particular, we find that Democrats' relative preference for the journalistic sites increased when the proportion of Metric Media coverage of education and immigration increased relative to the journalistic sites. This is perhaps unsurprising, as these issues have been highly salient and highly controversial during the early months of the second Trump administration. Furthermore, the previously mentioned stories about transgender high school athletes were all labeled as education.

[16] See Appendix D.III.

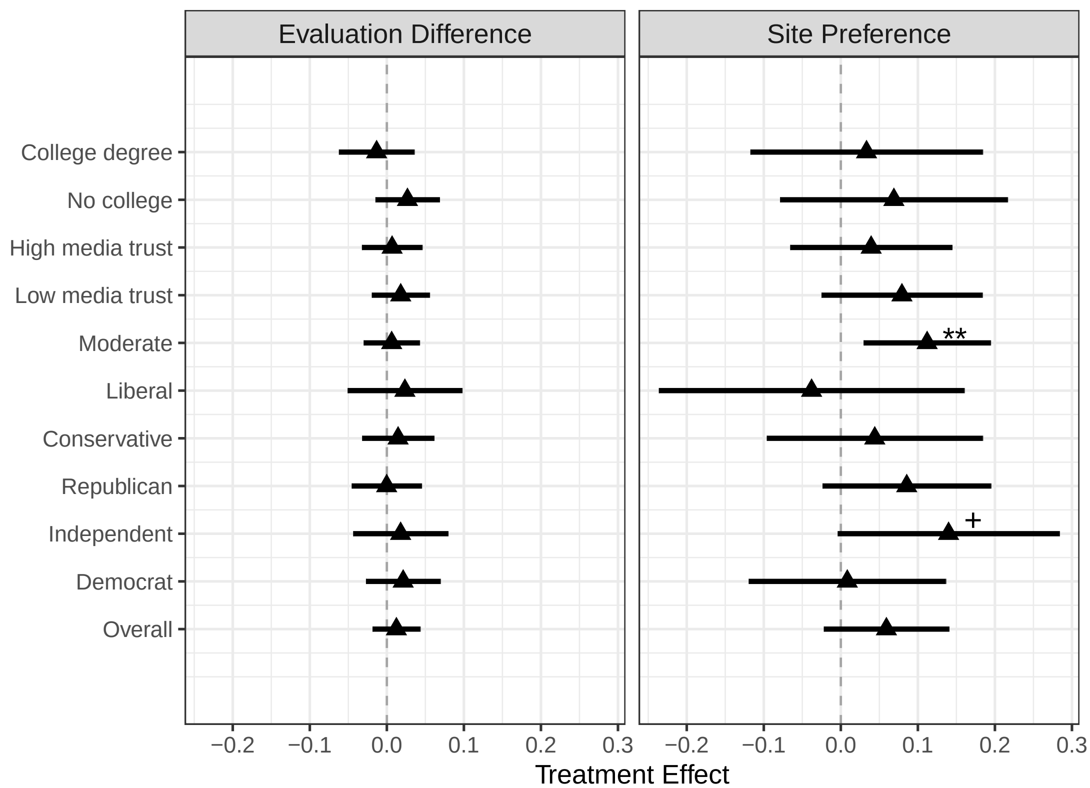

Figure 4: Estimated Treatment Effects by Subgroup

Notes: Heterogeneous treatment effects. Estimated treatment effects on site preference and evaluation difference by subgroup, with 95% confidence intervals. $^{+}$ $p < .10$ * $p < .05$, ** $p < .01$.

Conclusion

As the American news media grapples with dwindling trust and advances in artificial intelligence threaten to further disrupt the relationships between news organizations and their audiences, it is critical that we advance our understanding of news audiences' considerations in choosing content to consume online, their susceptibility to AI-augmented, biased stories or misinformation, and the efficacy of media literacy interventions designed to improve audiences' ability to protect themselves against low-quality digital content.

This paper contributes to that end. We tested whether a novel digital media literacy intervention improves audience discernment between real local news sites and algorithmically generated pink slime sites designed to imitate them. Using a nationally representative survey experiment with live website stimuli, we find that the intervention worked as intended in one sense: it successfully redirected participant attention toward more diagnostic cues of source credibility such as informative "About Us" pages, named bylines, and detailed ownership information. However, these changes in attention did not translate into meaningful differences in evaluations of the sites or in participants' choices between them.

Our analysis of open-ended responses suggests why. While participants in the treatment group mentioned credibility-relevant cues more often, those considerations were still overshadowed by other features, such as the topics covered, perceived ideological slant, site layout, and the presence of advertisements. These features may not be indicative of journalistic quality (layout, advertisements) or may be more a function of consumers' own interests and partisan preferences (topics, bias), but they heavily shaped audience preference nonetheless. As a result, an algorithmic site that aligns with a user's interests or political views may be preferred over a high-quality, investigative journalistic site that does not. This dynamic presents a fundamental challenge for efforts to improve news discernment: even when people are shown tips for detecting reliable sources, they may still prioritize alignment and relevance over journalistic integrity.

Methodologically, our primary innovation was to demonstrate the feasibility and value of using live websites as experimental stimuli. Our pilot studies demonstrated that individuals systematically evaluate sites differently when exposed to the full set of cues about source, content, and presentation than when they are provided with a low-information proxy as is common in existing work. We hope future research will build on this approach to better capture the complexity of real-world online news evaluation.

Of course, our study is not without limitations. Our sample size, while nationally representative, was likely underpowered to detect small treatment effects. We also focused only on a single, right-leaning network of algorithmic news sites. While liberal activists have created some similar sites of their own (Merrill and Kozlowska, 2020), these appear to be isolated sites and not a nationwide network, limiting our ability to make generalizable inferences. Future work should test whether similar dynamics hold for other types of partisan or ideologically ambiguous algorithmic sites. In addition, we did not disclose to participants (prior to debriefing) that one of the sites they visited was algorithmically generated or forewarn them about the existence of algorithmic sites. It is possible that priming participants to the existence of algorithmic news sites might have affected their evaluations of both types of sites. It is also possible that some individuals may prefer an algorithmic site over a journalistic one if they perceive algorithms to be an impartial alternative. Future work should explore these questions.

Finally, the pace of advancement in AI technology means that this space requires continuous scholarly attention. The algorithmic sites we study are relatively rudimentary, and yet they were preferred by over 40% of participants. These findings raise concerns about how well journalistic institutions are communicating their distinct value to audiences, particularly in an environment where much online content is free and professionally produced reporting often sits behind paywalls. While structural challenges such as declining newsroom resources limit the capacity for outreach, our results highlight the potential importance of increasing public awareness around what distinguishes legitimate journalism from automated alternatives.

References

Abernathy, Penelope Muse. 2018. “The Expanding News Desert.” *Center for Innovation and Sustainability in Local Media, University of North Carolina at Chapel Hill.*

Alba, Davey, and Jack Nicas. 2020. “As Local News Dies, a Pay-for-Play Network Rises in Its Place.” *New York Times* (October).

Altay, Sacha, Andrea De Angelis, and Emma Hoes. 2024. “Media literacy tips promoting reliable news improve discernment and enhance trust in traditional media.” *Communications Psychology* 2 (74).

Anderson-Davis, Stuart. 2024. ““Pink Slime Journalism” and a history of media manipulation in America.” *Columbia Journalism Review* (February).

Archer, Allison M. N. 2020. *Attacking the fourth estate: the nature and effects of political leaders’ war with the press.* Edward Elgar Publishing pp. 129–147.

Aslett, Kevin, Zeve Sanderson, William Godel, Nathaniel Persily, Jonathan Nagler, and Joshua A. Tucker. 2024. “Online searches to evaluate misinformation can increase its perceived veracity.” *Nature* 625: 548-556.

Barari, Soubhik, Christopher Lucas, and Kevin Munger. 2024. “Political Deepfakes are as Credible as Other Fake Media and (Sometimes) Real Media.” *Journal of Politics.*

Baum, Matthew A. 2002. “Sex, Lies, and War: How Soft News Brings Foreign Policy to the Inattentive Public.” *American Political Science Review* 96: 91-109.

Bengani, Priyanjana. 2019. “Hundreds of ‘pink slime’ local news outlets are distributing algorithmic stories and conservative talking points.” *Columbia Journalism Review* (December).

Bengani, Priyanjana. 2021. “The Metric Media network runs more than 1,200 local news sites. Here are some of the non-profits funding them.” *Columbia Journalism Review* (October).

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla

Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alect Radford, Ilya Sutskever, and Dario Amodei. 2020. "Language Models are Few-Shot Learners." In *34th Conference on Neural Information Processing Systems*.

Burnham, Michael, Kayla Kahn, Ryan Yank Wang, and Rachel X. Peng. 2024. "Political DEBATE: Efficient Zero-shot and Few-shot Classifiers for Political Text.".

Clark, Elizabeth, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A. Smith. 2021. "All That's 'Human' Is Not Gold: Evaluating Human Evaluation of Generated Text.".

Clerwall, Christer. 2014. "Enter the Robot Journalist: Users' perceptions of automated content." *Journalism Practice* 8: 519-531.

Darr, Joshua P. 2023. "How Sticky Is Pink Slime? Assessing the Credibility of Deceptive Local Media." *Annals of the American Academy of Political and Social Science* 707: 109-124.

Darr, Joshua P., Matthew P. Hitt, and Johanna L. Dunaway. 2018. "Newspaper Closures Polarize Voting Behavior." *Journal of Communication* 68: 1007-1028.

Druckman, James N. 2022. *Experimental Thinking: A Primer on Social Science Experiments*. Cambridge University Press.

Epstein, Ziv, Adam J. Berinsky, Rocky Cole, Andrew Gully, Gordon Pennycook, and David G. Rand. 2021. "Developing an accuracy-prompt toolkit to reduce COVID-19 misinformation online." *Harvard Kennedy School Misinformation Review* 2.

Flanagin, Andrew J., and Miriam J. Metzger. 2007. "The role of site features, user attributes, and information verification behaviors on the perceived credibility of web-based information." *New Media and Society* 9: 319-342.

Gao, Pengjie, Chang Lee, and Dermot Murphy. 2020. "Financing Dies in Darkness? The Impact of Newspaper Closures on Public Finance." *Journal of Financial Economics* 135: 445-467.

Gaziano, Cecilie, and Kristin McGrath. 1986. "Measuring the Concept of Credibility." *Journalism and Mass Communication Quarterly* 63 (3): 451–462.

Graefe, Andreas, Mario Haim, Bastian Haarmann, and Hans-Bernd Brosius. 2018. "Readers' perception of computer-generated news: Credibility, expertise, and readability." *Journalism* 19: 595-610.

Graham, Matthew H. 2023. "Detecting and Deterring Information Search in Online Surveys." *American Journal of Political Science*: 1–20.

Guess, Andrew M., and Kevin Munger. 2023. "Digital literacy and online political behavior." *Political Science Research and Methods* 11: 110-128.

Guess, Andrew M., Michael Lerner, Benjamin Lyons, Jacob M. Montgomery, Brendan Nyhan, Jason Reifler, and Neelanjan Sircar. 2020. "A digital media literacy intervention increases discernment between mainstream and false news in the United States and India." *Proceedings of the National Academy of Sciences of the United States of America* 117: 15536-15545.

Hameleers, Michael. 2022. "Separating truth from lies: comparing the effects of news media literacy interventions and fact-checkers in response to political misinformation in the US and Netherlands." *Information, Communication Society* 25: 110-126.

Hayes, Danny, and Jennifer L. Lawless. 2018. "The Decline of Local News and Its Effects: New Evidence from Longitudinal Data." *Journal of Politics* 80: 332-336.

Hayes, Danny, and Jennifer L. Lawless. 2021. *News Hole: The Demise of Local Journalism* *and Political Engagement*. Cambridge University Press.

Henestrosa, Angelica Lermann, Hannah Greving, and Joachim Kimmerle. 2023. "Automated journalism: The effects of AI authorship and evaluative information on the perception of a science journalism article." *Computers in Human Behavior* 138: 107445.

Hovland, Carl I., and Walter Weiss. 1951. "The Influence of Source Credibility on Communication Effectiveness." *Public Opinion Quarterly* 15 (4): 635–650.

Kohring, Matthias, and Jörg Matthes. 2007. "Trust in News Media: Development and Validation of a Multidimensional Scale." *Communication Research* 34 (2): 231–252.

Kremer, Rich. 2024. "Catholic leaders say Wisconsin Catholic Tribune is political mail not authorized by the church." *Wisconsin Public Radio* (October).

Kuru, Ozan. 2025. "Literacy training vs. psychological inoculation? Explicating and comparing the effects of predominantly informational and predominantly motivational interventions on the processing of health statistics." *Journal of Communication* 75: 64-78.

Lacy, Stephen, Brendan R. Watson, Daniel Riffe, and Jennette Lovejoy. 2015. "Issues and best practices in content analysis." *Journalism and Mass Communication Quarterly* 92 (4): 791–811.

Ladd, Jonathan M. 2012. *Why Americans Hate the Media and How It Matters.* Princeton University Press.

Mathews, Nick. 2022. "Life in a news desert: The perceived impact of a newspaper closure on community members." *Journalism* 23: 1250-1265.

Merrill, Jeremy B., and Hanna Kozlowska. 2020. "A well-heeled liberal group is filling the local news void with Trump-bashing coronavirus ads." *Quartz* (April).

Metzger, Miriam J., Andrew J. Flanagin, and Ryan B. Medders. 2010. "Social and Heuristic Approaches to Credibility Evaluation Online." *Journal of Communication* 60.

Moore, Ryan C., and Jeffrey T. Hancock. 2022. "A digital media literacy intervention for older adults improves resilience to fake news." *Scientific Reports* 12: 6008.

Moore, Ryan C., Ross Dahlke, Priyanjana Bengani, and Jeffrey T. Hancock. 2023. "The Consumption of Pink Slime Journalism: Who, What, When, Where, and Why?".

Myers, Steven Lee. 2024. "Spate of Mock News Sites With Russian Ties Pop Up in U.S." *New York Times* (March).

Newman, Nic, and Richard Fletcher. 2017. "Bias, Bullshit and Lies: Audience Perspectives on Low Trust in the Media."

Peterson, Erik. 2021. "Paper Cuts: How Reporting Resources Affect Political News Coverage." *American Journal of Political Science* 65: 443-459.

Peterson, Erik, Joshua P. Darr, Maxwell B. Allamong, and Michael Henderson. 2025. "Can Americans' trust in local news be trusted? The emergence, sources, and implications of the local news trust advantage." *American Journal of Political Science* n/a (n/a): 1-16.

Peterson, Erik, and Maxwell B. Allamong. 2022. "The Influence of Unknown Media on Public Opinion: Evidence from Local and Foreign News Sources." *American Political Science Review* 116: 719-733.

Prochazka, Fabian, and Wolfgang Schweiger. 2019. "How to Measure Generalized Trust in News Media? An Adaptation and Test of Scales." *Communication Methods and Measures* 13 (1): 26–42.

Quéré, Marianne Aubin Le, and Maurice Jakesch. 2022. "Trust in AI in Under-resourced Environments: Lessons from Local Journalism." In *CHI Workshop on Trust and Reliance in AI-Human Teams*.

Rafsky, Sara. 2022. "Tow Center audience study: Reader perspectives on partisan local news sites." (September).

Richards, Jennifer Smith, and Megan O'Matz. 2024. "Who's Mailing the Catholic Tribune?

It's Not the Church, It's Partisan Media." *ProPublica* (October).

Royal, Asa, and Philip M. Napoli. 2022. "Local Journalism without Journalists? Metric Media and the Future of Local news." *Journal of Creative Industries and Cultural Studies* 8: 119-147.

Rubado, Meghan E., and Jay T. Jennings. 2020. "Political Consequences of the Endangered Local Watchdog: Newspaper Decline and Mayoral Elections in the United States." *Urban Affairs Review* 56: 1327-1356.

Sadeghi, McKenzie, Dimitris Dimitriadis, Lorenzo Arvanitis, Virginia Padovese, Giulia Pozzi, Sara Badilini, Chiara Vercellone, Macrina Wang, Natalie Huet, Zack Fishman, Leonie Pfaller, Natalie Adams, and Miranda Wollen. 2025. "Tracking AI-enabled Misinformation: 1,254 'Unreliable AI-Generated News' Websites (and Counting), Plus the Top False Narratives Generated by Artificial Intelligence Tools." *NewsGuard* (February).

Schulhofer-Wohl, Sam, and Miguel Garrido. 2013. "Do Newspapers Matter? Short-Run and Long-Run Evidence From the Closure of The Cincinnati Post." *Journal of Media Economics* 26: 60-81.

Shaker, Lee. 2014. "Dead Newspapers and Citizens' Civic Engagement." *Political Communication* 31: 131-148.

Spitale, Giovanni, Nikola Biller-Andorno, and Federico Germani. 2023. "AI model GPT-3 (dis)informs us better than humans." *Science Advances* 9: 1-9.

Stagnaro, Michael Nicholas, James N. Druckman, Adam J. Berinsky, Antonio A. Arechar, Robb Willer, and David G. Rand. 2024. "Representativeness versus Response Quality: Assessing Nine Opt-In Online Survey Samples.".

Strange, Austin M., Ryan D. Enos, Mark Hill, and Amy Lakeman. 2019. "Online volunteer laboratories for human subjects research." *PLoS One* 14 (8): e0221676.

Sundar, S. Shyam. 2008. *The MAIN Model: A Heuristic Approach to Understanding Tech-*

nology Effects on Credibility. The MIT Press pp. 73–100.

Toff, Benjamin, Ruth Palmer, and Rasmus Kleis Nielsen. 2024. *Avoiding the News: Reluctant Audiences for Journalism*. New York: Columbia University Press.

Tow Center for Digital Journalism. 2024. "“Pink Slime”: Partisan journalism and the future of local news.".

Traberg, Cecilie S., Jon Roozenbeek, and Sander van der Linden. 2022. "Psychological Inoculation against Misinformation: Current Evidence and Future Directions." *Annals of the American Academy of Political and Social Science* 700: 136-151.

Tully, Melissa, Emily K. Vraga, and Leticia Bode. 2020. "Designing and Testing News Literacy Messages for Social Media." *Mass Communication and Society* 23: 22-46.

Usher, Nikk, and Sanghoon Kim-Leffingwell. 2023. "How Loud Does the Watchdog Bark? A Reconsideration of Losing Local Journalism, News Nonprofits, and Political Corruption." *International Journal of Press/Politics*: 1-23.

Vallone, Robert P., Lee Ross, and Mark R. Lepper. 1985. "The Hostile Media Phenomenon: Biased Perception and Perceptions of Media Bias in Coverage of the Beirut Massacre." *Journal of Personality and Social Psychology* 49: 577-585.

van der Kaa, Hille, and Emiel Krahmer. 2014. "Journalist versus news consumer: The perceived credibility of machine written news." In *Proceedings of the Computation+Journalism conference*. https://research.tilburguniversity.edu/en/publications/journalist-versus-news-consumer-the-perceived-credibility-of-mach.

Von Hohenberg, Bernhard Clemm, and Andrew M. Guess. 2023. "When Do Sources Persuade? The Effect of Source Credibility on Opinion Change." *Journal of Experimental Political Science* 10 (3): 328–342.

Wölker, Anja, and Thomas E. Powell. 2021. "Algorithms in the newsroom? News readers' perceived credibility and selection of automated journalism." *Journalism* 22: 86-103.

Yale, Robert N., Jakob D. Jensen, Nick Carcioppolo, Ye Sun, and Miao Liu. 2015. "Examining First- and Second-Order Factor Structures for News Credibility." *Communication Methods and Measures* 3 (9): 152–169.

Yan, Harry Yaojun, Garrett Morrow, Kai-Cheng Yang, and John Wihbey. 2025. "The origin of public concerns over AI supercharging misinformation in the 2024 U.S. presidential election." *Harvard Kennedy School Misinformation Review* 1: 1-13.

Online Appendix: Supporting Information for

Local News in the Crosshairs: Audience Perceptions of Algorithmic News Sites

David A. Beavers and Kevin DeLuca

Contents

A Additional Examples of Metric Media Content

This appendix provides additional detail on the operations and content strategies of Metric Media, the pink slime network at the center of our study. It includes documented examples of overtly partisan political campaigning, corporate advocacy disguised as local reporting, and direct content replication from government press releases. Together, these cases illustrate the strategic use of algorithmic local news sites to shape public opinion, influence local policy debates, and potentially affect electoral outcomes.

Reporting by *The New York Times* and *Columbia Journalism Review* reveals that Metric Media outlets have boosted Republican gubernatorial candidates in Illinois and Ohio, publishing hundreds of articles in the lead-up to elections, often containing candidates' own quotes in headlines (Alba and Nicas, 2020; Tow Center for Digital Journalism, 2024). Some campaigns have reportedly used Metric Media's online portals to pitch stories directly (Stanley-Becker and Dawsey, 2023).

Beyond partisan politics, Metric Media outlets have also been used to advance corporate agendas. In North Dakota, for example, Central ND News distributed mailers that resurfaced anti-pipeline protest narratives from 2017, just months before a major jury trial involving the Dakota Access Pipeline (Steurer, 2025). Greenpeace, an activist group the backed the protesters, alleged the coverage was intended to sway potential jurors in favor of Energy Transfer, the Texas-based pipeline company behind the project. In Ohio, Metric Media's *Mount Vernon News* has regularly published content favoring natural gas development and opposing solar energy, often reprinting messaging from industry-aligned groups like The Empowerment Alliance (Green et al., 2024). These campaigns reflect a broader strategy of influencing public opinion and local government decision-making under the guise of community-based reporting.

These patterns point to a strategic use of pink slime media to shape voter attitudes and potentially influence elections and public policy, particularly in environments where voters only source of information on the topics may be from local news outlets. By publishing manipulative or misleading content disguised as local journalism, pink slime networks exploit the very cues that typically signal trustworthiness. This underscores the urgency of understanding how well voters can actually discern between algorithmic pink slime outlets and legitimate local journalism, and whether interventions can meaningfully alter their perceptions or choices.

The Pennsylvania Department of Transportation (PennDOT) District 5 has announced that road restrictions will begin on Monday, February 24, for a project aimed at improving the interchange of Route 309 and Tilghman Street in South Whitehall Township, Lehigh County.

Starting Monday, lane restrictions will be in place on Tilghman Street between 40th Street and just west of Hausman Road. These restrictions are expected to occur between 9 a.m. and 3 p.m., and from 6 p.m. to 6 a.m. on weekdays, with possible changes anytime during weekends. Motorists should anticipate traffic pattern changes and allow extra travel time.

Additionally, Broadway will be closed between Parkway and Hausman roads for approximately four weeks starting March 3 due to work around the Route 309 bridge over Broadway. During this period, traffic will be detoured via Parkway Road, Tilghman Street, and Hausman Road.

The project aims to reconfigure the Route 309 at Tilghman Street interchange to enhance safety and mobility. It includes replacing two bridges—one at Route 309 over Tilghman Street and another at Route 309 over Broadway Street—rehabilitating the Route 309 culvert over Little Cedar Creek, improving roadway drainage, repairing road bases, paving, and adding sidewalks on Tilghman Street. Two new traffic signals will also be installed at the end of the reconfigured Route 309 ramps to Tilghman Street.

Allentown, PA – The Pennsylvania Department of Transportation (PennDOT) District 5 announced road restrictions are anticipated to commence Monday (February 24) for a project to improve the interchange of Route 309 and Tilghman Street in South Whitehall Township, Lehigh County.

Starting Monday, motorists can expect lane restrictions on Tilghman Street between 40th Street and just west of Hausman Road. Restrictions are anticipated between 9 a.m. and 3 p.m. and 6 p.m. and 6 a.m. weekdays and anytime on weekends. Motorists are advised to anticipate changing traffic patterns and allow for extra time traveling through the area.

Also, Broadway will be closed between Parkway and Hausman roads for approximately four weeks starting March 3 for work around the Route 309 bridge over Broadway. During this closure Broadway traffic will be detoured on Parkway Road, Tilghman Street and Hausman Road.

The project includes the reconfiguration of the Route 309 at Tilghman Street interchange to improve safety and mobility. It will include two bridge replacements - one at Route 309 over Tilghman Street and another at Route 309 over Broadway Street, the rehabilitation of the Route 309 culvert over Little Cedar Creek, roadway drainage improvements, road base repairs, paving, and sidewalks on Tilghman Street. Two new traffic signals will also be installed at the end of the reconfigured Route 309 ramps to Tilghman Street that will be coordinated with existing signals at the intersections of Hausman Road at Cetronia Road, Tilghman Street at Parkway Road, and Tilghman Street at 40th Street.

Figure A1: Example of direct content replication: Keystone Today article (Feb. 24, 2025) at left shows strong textual similarity to the PennDOT press release (Feb. 21, 2025) at right. The highlighted passages illustrate that the Metric Media article is nearly a verbatim copy of the press release, with only minimal editing and no named byline—demonstrating the network's reliance on unattributed official material rather than original reporting.

B Survey Instrument Details

B.I Survey Instrument

Informed Consent

What is the purpose of this research?
In this survey, we are interested in understanding audience attitudes about news websites. You will be asked to view a pair of news sites and then provide your thoughts on the sites.

What can I expect if I take part in this research?
In this survey, we will first ask you a few questions about yourself and what type of news you read or watch. Then, we will ask you to take a few minutes to look at a pair of news sites. Finally, we will ask you a few questions about your impressions of the sites. The entire survey should take between 5 and 10 minutes to complete.

What should I know about a research study?

- You must be 18 years of age or older.
- Whether or not you take part is up to you.
- Your participation is completely voluntary.
- You can choose not to take part.
- You can agree to take part and later change your mind.
- Your decision will not be held against you.
- Your refusal to participate will not result in any consequences or any loss of benefits that you are otherwise entitled to receive.

You may not be told everything or may be misled
As part of this research design, you may not be told or may be misled about the purpose or procedures of this research. However, the purpose or procedures of the research will be disclosed to you following your participation.

Anticipated benefits and risks
There are no direct benefits to participants for taking part in this study. However, the research may contribute to a broader understanding of audience attitudes toward news websites. Participants will be asked to provide their opinions about news websites and answer general demographic and media consumption questions. No sensitive or personal information will be collected beyond what is necessary for the research, and all responses will remain confidential. You may experience slight discomfort from answering survey questions.

How will data be published?
This study is being conducted through Time-sharing Experiments for the Social Sciences (TESS). Per TESS policies, one year after data have been collected, the data will become available to others via TESS' partnership with the Open Science Foundation. We will not collect or publish any directly identifiable information like your contact information, name, address, or birth date.

Are there any costs to participation? Will I be paid for participation?
You will not have to pay to take part in this study. The only costs may include the time it takes to complete the survey. You will be compensated in AmeriPoints for taking part in this study.

How will you keep my data safe and private?
All of your responses will be anonymous. Only the researchers involved in this study and those responsible for research oversight (such as representatives of the Harvard and Yale University Institutional Review Boards, and others) will have access to any information that could identify you that you provide. When we publish the results of the research or talk about it in conferences, we will not use your name. We will also share information about you with other researchers for future studies, but we will not use your name or any other identifiers. We will not ask you for any additional permissions.

Who can I talk to?
If you have questions, concerns, or complaints, or think the research has hurt you, talk to the survey team at support@AmeriSpeak.org. If you have questions about your rights as a research participant, or you would like to speak with someone other than the study team to discuss problems, concerns, or questions, or to obtain information or offer suggestions, you can call the Yale Institutional Review Boards at (203) 785-4688 or email hrpp@yale.edu.

By selecting "I agree", you consent to participate in this study and to provide the most honest answers you can.

Pre-Treatment Questions (Randomize Order)

Q1. What is your main source of news?

- National broadcast network news (ABC, CBS, NBC, PBS)
- Cable news (CNN, Fox News, MSNBC)
- Local television news
- National newspapers (New York Times, Washington Post, Wall Street Journal, etc.)
- Local newspapers
- Radio
- Social media
- Other

Q2. In general, how much trust and confidence do you have in the mass media – such as newspapers, TV, and radio – when it comes to reporting the news fully, accurately and fairly? (Randomize response order)

- A great deal
- A fair amount
- Not very much
- None at all

Q3. What state do you live in? (Dropdown menu including None/Other)

Intervention (Random Assignment)

Q4 (Treatment): Next, we are going to ask you to browse two news sites. One or both sites may be unfamiliar to you. The American Library Association offers the following tips for evaluating unfamiliar news sites:

1. **Consider the source.** Click away from the story to investigate the site, its mission and its contact info.
2. **Read past the headline.** Headlines can be outrageous in effort to get clicks. Go beyond headlines.
3. **Assess the credibility of the author.** Do a quick Google search on the author. What is their expertise? What organization do they represent?
4. **Look at the links and sources supporting the article.** Click those links. Determine if the subsequent information supports the story. Consider the reliability of the sources.
5. **Check the date.** Is the story out of date or is it current?
6. **Consider that the item might be satire.** If it seems too outlandish, it might be satire. Do some quick research on the site and author to find out.
7. **Consider that it might be promotional.** Is the purpose of the site to sell a product?
8. **Check your biases.** Try to evaluate the story on its merits, rather than whether you want it to be true or not.
9. **Search other news outlets** to see if the news is widely reported.

Q4 (Control): Next, we are going to ask you to browse two news sites. One or both sites may be unfamiliar to you.

Stimulus (Random Order)

Q5. Please take some time to browse the site linked below. When you click the link, it will open the site in a new internet browser tab. Once you have browsed the site, please return to this browser tab to complete the survey.

[Metric Media site hyperlink]

Q6-9. Please indicate how well you believe the following words describe the site you just browsed: [grid layout]

- Trustworthy

- Reliable
- Accurate
- Informative

Responses:

- Strongly agree
- Agree
- Neither agree nor disagree
- Disagree
- Strongly disagree

Q10. Please take some time to browse the site linked below. When you click the link, it will open the site in a new internet browser tab. Once you have browsed the site, please return to this browser tab to complete the survey.

[Legitimate site hyperlink]

Q11-14. Please indicate how well you believe the following words describe the site you just browsed: [grid layout]

- Trustworthy
- Reliable
- Accurate
- Informative

Responses:

- Strongly agree
- Agree
- Neither agree nor disagree
- Disagree
- Strongly disagree

Outcome Measures

Q15. Still thinking about the two sites you just browsed, what elements of the sites did you have in mind when evaluating them? (Open-ended)

Q16. If you wanted to learn about public affairs in [state] in the future, which website would you prefer as a source of news? (Randomize response order)

- Strongly prefer [Metric Media site name]
- Somewhat prefer [Metric Media site name]
- Somewhat prefer [journalistic site name]
- Strongly prefer [journalistic site name]

Q17. Had you previously heard of either or both of the sites you just browsed?

- Yes, both
- Yes, one
- No
- Not sure

Q18. (If "Yes, one") Which of the two sites were you previously familiar with? (Randomized order)

- [Metric Media site name]
- [journalistic site name]

Debrief Text

Thank you for taking part in this research study. There was some information about the study that we did not share with you at the beginning of your participation. We would now like to fully inform you of the nature of this research and provide some additional details about the study.

What was this study about? In this project, we are interested in understanding what considerations people have in mind when evaluating news websites. We are also interested in seeing if and how peoples' perceptions of different types of news websites vary. Specifically, we want to compare the evaluations participants give of typical news websites with evaluations they give of algorithmically-produced news sites that involve minimal human intervention to generate content.

What information was withheld from you? We withheld two pieces of information from you at the beginning of the study. First, we did not tell you that one of the sites we asked you to visit is produced by Metric Media, which operates a network of algorithmically-generated news sites run with little or no direct involvement of a local journalist. For more information about Metric Media, you may consider visiting their website (https://www.metricmedia.org/) or reading about them in the Columbia Journalism Review (https://www.cjr.org/tow_center_reports/metric-media-lobbyists-funding.php). Second, we did not tell you that we would randomly assign you to view a set of tips that we lightly adapted from the American Library Association (https://libguides.ala.org/InformationEvaluation) or not. Those tips are:

1. **Consider the source.** Click away from the story to investigate the site, its mission and its contact info.
2. **Read past the headline.** Headlines can be outrageous in effort to get clicks. Go beyond headlines.
3. **Assess the credibility of the author.** Do a quick Google search on the author. What is their expertise? What organization do they represent?
4. **Look at the links and sources supporting the article.** Click those links. Determine if the subsequent information supports the story. Consider the reliability of the sources.
5. **Check the date.** Is the story out of date or is it current?
6. **Consider that the item might be satire.** If it seems too outlandish, it might be satire. Do some quick research on the site and author to find out.
7. **Consider that it might be promotional.** Is the purpose of the site to sell a product?
8. **Check your biases.** Try to evaluate the story on its merits, rather than whether you want it to be true or not.
9. **Search other news outlets** to see if the news is widely reported.

Why didn't we share this information with you before? In order to get your authentic evaluation of the site, as if you had just come across them while surfing the web, we didn't tell you in advance that one of the sites was an algorithmically-produced news website.

Why weren't you shown the American Library Association tips until now? We want to be able to compare responses from individuals who were shown the American Library Association tips to responses from individuals who didn't see them in order to estimate the effectiveness of these types of tipsheets as an intervention for improving discernment of high- from low-quality information online. (for control group only)

How was the study conducted? Based on your state of residence, we assigned you to view two websites: a legitimate state news site and a state-based website from Metric Media. We then randomly assigned you to view the American Library Association tips or not. All participants were then asked the same set of questions to evaluate the sites.

Did you see an algorithmically-generated news site? Yes, you were assigned to see an algorithmically-generated news site operated by Metric Media.

Why is this study important? With recent advances in artificial intelligence and continued economic difficulties facing local newspapers, we believe this study will help us understand how news audiences come to trust, or not trust, information they find online.

How to contact the survey team: If you have questions or concerns about your participation, or want to request a summary of research findings, please contact the survey team at support@AmeriSpeak.org.

Whom to contact about your rights as a participant in this research: For questions, concerns, suggestions, or complaints that have not been or cannot be addressed by the researcher, or to report research-related harm, please contact the Committee on the Use of Human Subjects at Harvard University, 1350 Massachusetts Avenue - Smith Campus Center, Suite 645, Cambridge, MA 02138. Email: cuhs@harvard.edu; Phone: (617) 496-2847; or the Yale Institutional Review Boards. Phone: (203) 785-4688; Email: hrpp@yale.edu.

—End of Survey Instrument—

B.II Website Links

Table B1 below lists the exact website pairs participants were shown based on their state of residence. For each state, respondents were randomly assigned to browse both a Metric Media (algorithmic) site and a legitimate local journalistic news site. This pairing was designed to reflect realistic local news options and to test respondents' ability to evaluate and compare sites that cover the same geographic area. These specific site links were also used to construct our site-level control variables (e.g., ownership type, paywall presence, soft news share – see Table B3) in the regression analyses.

Table B1: Stimulus Website Links

State	Metric Media site	Journalistic news site
Alabama	Yellowhammer Times	Montgomery Advertiser
Alaska	Last Frontier News	Juneau Empire
Arizona	PHX Reporter	Arizona Capitol Times
Arkansas	Natural State News	Arkansas Times
California	Golden State Today	The Sacramento Bee
Colorado	Centennial State News	The Colorado Sun
Connecticut	Constitution State News	New Haven Register
Delaware	First State Times	Delaware State News
D.C.	Washington D.C. Business Daily	Washington City Paper
Florida	Sunshine Sentinel	Tallahassee Democrat
Georgia	Peach Tree Times	Gwinnet Daily Post
Hawaii	Aloha State News	The Maui News
Idaho	Gem State Wire	Idaho Press

State	Metric Media site	Journalistic news site
Illinois	Prairie State Wire	The State Journal-Register
Indiana	Hoosier State Today	Indiana Capital Chronicle
Iowa	Hawkeye Reporter	Quad-City Times
Kansas	Sunflower State News	The Topeka Capital-Journal
Kentucky	Bluegrass Times	The State-Journal
Louisiana	Pelican State News	Shreveport Times
Maine	Pine State News	Central Maine
Maryland	Maryland State Wire	Capital Gazette
Massachusetts	Bay State News	Boston Herald
Michigan	Detroit City Wire	Lansing State Journal
Minnesota	Minnesota State Wire	Pioneer Press
Mississippi	Magnolia State News	The Meridian Star
Missouri	Show-Me State Times	News Tribune
Montana	Big Sky Times	Independent Record
Nebraska	Cornhusker State News	Lincoln Journal Star
Nevada	Silver State Times	Nevada Appeal
New Hampshire	Granite State Times	Concord Monitor
New Jersey	Garden State Times	The Trentonian
New Mexico	Enchantment State News	Santa Fe New Mexican
New York	Empire State Today	Albany Times-Union
North Carolina	Old North News	The News & Observer
North Dakota	Fargo Standard	Minot Daily Times
Ohio	Buckeye Reporter	The Columbus Dispatch
Oklahoma	Sooner State News	The Oklahoman
Oregon	Beaver State News	Statesman Journal
Pennsylvania	Keystone Today	Patriot-News
Rhode Island	Ocean State Today	Southern Rhode Island Newspapers
South Carolina	Palmetto State News	The State
South Dakota	Rushmore State News	Capital Journal
Tennessee	Volunteer State News	The Tennessean
Texas	Austin News	Austin American-Statesman
Utah	Beehive State News	Deseret News
Vermont	Green Mountain Times	The Barre-Montpelier Times Argus
Virginia	Old Dominion News	Daily Press
Washington	Evergreen Reporter	The Olympian
West Virginia	Mountain State Times	The Register-Herald
Wisconsin	The Sconi	Wisconsin State Journal
Wyoming	Equality State News	Wyoming Tribune Eagle

B.III Coding Instructions for Open-Ended Responses

Below are the coding instructions used to annotate the open-ended responses. One of the authors and an undergraduate research assistant independently coded all open-ended responses. Because simple agreement is high due to the sparsity of many of these codes, we report Krippendorff's α as our metric for intercoder reliability (Lacy et al., 2015). Reliability was over 0.7 for all codes except for Norms, which we drop from the analysis. We created a composite measure of each code that takes a value of 1 if both coders assigned that code, 0.5 if only one coder assigned that code, and 0 if neither coder assigned that code. The intercoder reliability (Krippendorff's α) for each site element is displayed below in Table B2.

1. **Layout** — e.g. ease of navigating site, attractive colors and fonts, eye-catching, comments on the amount of white space, references to "presentation"

2. **Topics** — e.g. the specific topics covered, diversity of topics. In addition, this code should include any comments on having read the actual content on the sites that don't specifically evaluate the content along another dimension in this coding scheme – e.g. "I read the articles and thought they were biased" would be coded as bias, but "I read the articles" would be coded as topics. Also apply this code if they say they evaluate sites based on whether they contain specific types of content, like local news or sports. Comments that refer to the sites as interesting or relevant should be coded as both Topics and Usefulness, since these references are inherently about the topics/content.

3. **Bias** — e.g. reliability of facts, truthfulness, accuracy, separation of facts and opinions, objective reporting/objectivity, accusations that the site is slanted or has a liberal or conservative bias, references that the sites seem to have an "agenda", etc. Note that you should apply this code regardless of the accusation is that the site is or isn't biased. In other words, you should apply this code whenever bias comes up as a consideration, whether or not the speaker alleges bias is present.

4. **Quality** -– e.g. of writing and reporting; well-written articles (or not), thoroughly reported articles (or not), presence of absence of grammatical and typographical errors, etc. This can include references to the articles as clear, concise, etc. This includes relatively minimal references to "wording" or "language" on the sites, even if not accompanied by an evaluative statement (e.g. "the wording was concise").

5. **Blanket distrust** — e.g. this applies to any statement that doesn't specifically evaluate the two sites used in the experiment, but rather refers to evaluating these sites based on a blanket distrust of the media as a whole. You can also apply this code if a participant says they're completely disinterested in politics or refers to the media as propaganda or fake news.

6. **Clickbait** -– e.g. are the headlines and topics inflammatory and sensationalized or not, is the writing professional or gimmicky and tabloid-style, does the content appeal to emotions rather than reason. You can also apply this code if a respondent talks about the financial motivation behind the sites, e.g. "this site just seems to want to drive clicks to get more ad revenue rather than really report the news"

7. **Familiarity** — e.g. has or hasn't heard of the site, comments on the site's reputation or popularity; also include comments on the site's name, e.g. the name made them think the site is biased.

8. **Advertising** — e.g. comments on the amount of advertising on the site, the presence of pop-up ads, prompts to subscribe, or being unable to read articles without first having to subscribe. Do not include under this references to their malware software blocking access to the site or the site failing to load properly.

9. **Multimedia** — e.g. references to photos, videos

10. **Usefulness** -– e.g. comments on whether the site's coverage is relevant, important, useful, or informative, whether to themselves personally or to people in their community. Comments that refer to the sites as interesting or relevant should be coded as both Topics and Usefulness, since these references are inherently about the topics/content.

11. **Metadata** -– e.g. comments on how dated the articles are (i.e. are the most prominently placed articles from days or weeks ago?), whether there are named bylines on the articles, mentions of the site's about-us pages or author bios, whether the site has an ethics policy they could identify, etc. Also include under this code any mentions of searching the web for information about the site or for others' reviews about the site. And include references for searching the web to verify the information or claims made on the site, e.g. "I Googled to see if the information being reported on the site was appearing on other websites"

12. **Norms** — e.g. references to the professionalism of the journalists (purportedly) behind the sites. This includes references to practices that journalists are expected to do, e.g. fact-checking, citing sources, and proper hyperlinking to source documents and original reporting. Do not apply this code to vague mentions of "sources" unless the context is citing sources as a practice.

B.IV Site Characteristics and Ownership

Table B3 lists all of the journalistic (human-produced) news sites included in our experiment, along with their owners and owner type categories. This information is important for two reasons. First, it highlights the substantial heterogeneity in local news site ownership across states, from large publicly traded chains to small independent or nonprofit outlets, as well as hedge-fund and private-equity-backed groups known for aggressive cost-cutting. Second, these ownership types were used as control variables in our regression analyses to account for potential variation in quality, resources, or audience perceptions that might systematically differ by ownership structure.

Table B3: Stimulus Site Ownership Information

State	**Site**	**Owner**	**Owner Type**
Alabama	Montgomery Advertiser	Gannett/New Media	private equity
Alaska	Juneau Empire	Carpenter Media	privately owned
Arizona	Arizona Capitol Times	State Affairs	information services

Table B3 continued

State	Site	Owner	Owner Type
Arkansas	Arkansas Times	independent	independent
California	The Sacramento Bee	McClatchy/Chatham	hedge fund
Colorado	The Colorado Sun	independent	independent
Connecticut	New Haven Register	Hearst Communications	privately owned
Delaware	Delaware State News	Independent Newsmedia	privately owned
D.C.	Washington City Paper	independent	independent
Florida	Tallahassee Democrat	Gannett/New Media	private equity
Georgia	Gwinnett Daily Post	Times-Journal Inc.	privately owned
Hawaii	The Maui News	Ogden Newspapers	privately owned
Idaho	Idaho Press	Adams Publishing	privately owned
Illinois	The State Journal-Register	Gannett/New Media	private equity
Indiana	Indiana Capital Chronicle	independent	independent
Iowa	Quad-City Times	Lee Enterprises	publicly traded
Kansas	The Topeka Capital-Journal	Gannett/New Media	private equity
Kentucky	The State-Journal	Carpenter Media	privately owned
Louisiana	Shreveport Times	Gannett/New Media	private equity
Maine	Central Maine	Maine Trust for Local News/National Trust for Local News	nonprofit
Maryland	Capital Gazette	independent	independent
Massachusetts	Boston Herald	MediaNews/Digital First Media/Alden Global Capital	hedge fund
Michigan	Lansing State Journal	Gannett/New Media	private equity
Minnesota	Pioneer Press	MediaNews/Digital First Media/Alden Global Capital	hedge fund
Mississippi	The Meridian Star	Carpenter Media	privately owned
Missouri	News Tribune	WEHCO Media	privately owned
Montana	Independent Record	Lee Enterprises	publicly traded
Nebraska	Lincoln Journal Star	Lee Enterprises	publicly traded
Nevada	Nevada Appeal	Eagle Valley Publishing	privately owned
New Hampshire	Concord Monitor	Newspapers of New England	privately owned
New Jersey	The Trentonian	MediaNews/Digital First Media/Alden Global Capital	hedge fund
New Mexico	Santa Fe New Mexican	independent	independent
New York	Albany Times-Union	Hearst Communications	privately owned
North Carolina	The News & Observer	McClatchy/Chatham	hedge fund
North Dakota	Minot Daily Times	Ogden Newspapers	privately owned
Ohio	The Columbus Dispatch	Gannett/New Media	private equity
Oklahoma	The Oklahoman	Gannett/New Media	private equity
Oregon	Statesman Journal	Gannett/New Media	private equity
Pennsylvania	Patriot-News	Advance Publications	privately owned

Table B3 continued

State	Site	Owner	Owner Type
Rhode Island	Southern Rhode Island Newspapers	Rhode Island Suburban Newspapers	privately owned
South Carolina	The State	McClatchy/Chatham	hedge fund
South Dakota	Capital Journal	Wick Communications	privately owned
Tennessee	The Tennessean	Gannett/New Media	private equity
Texas	Austin American-Statesman	Gannett/New Media	private equity
Utah	Deseret News	independent	independent
Vermont	The Barre-Montpelier Times Argus	Sample News/Brunswick	privately owned
Virginia	Daily Press	Tribune Publishing/Alden Global Capital	hedge fund
Washington	The Olympian	McClatchy/Chatham	hedge fund
West Virginia	The Register-Herald	Community Newspaper Holdings	privately owned
Wisconsin	Wisconsin State Journal	Lee Enterprises	publicly traded
Wyoming	Wyoming Tribune Eagle	Adams Publishing	privately owned

C Pilot Study Results

To test the viability of our design and assess whether participants could meaningfully distinguish between local journalism and algorithmic pink slime content, we conducted two pilot studies. These pilots allowed us to evaluate whether linking participants to live, functioning homepages (rather than showing static screenshots) affected credibility perceptions and site preferences. We also used the pilots to refine question wording and identify potential sources of confusion in the experimental protocol.

C.I Pilot Study Samples

We recruited a total of 1,096 adult participants across two samples. The first sample (n = 323) was drawn from CloudResearch's Connect platform, a high-quality online panel. Cloud Connect respondents were paid $2 for completing a 5-to-10-minute survey. The second sample (n = 773) was recruited from DLABSS (the Digital Lab for the Social Sciences), who are volunteers. All participants were U.S. adults aged 18 or older. Full sample characteristics are shown in Table C1.

The pooled pilot sample skews older, more educated, and more white than the general U.S. population. As shown in the table, there are substantial differences across the two sources, with CloudResearch participants tending to be younger, more racially diverse, and more Democratic-leaning than those from DLABSS.

Table B2: Krippendorff's α for Open-Ended Response Codes

Code	**Krippendorff's α**
Advertising	0.971
Bias	0.811
Distrust	0.704
Clickbait	0.847
Familiarity	0.829
Layout	0.855
Metadata	0.844
Multimedia	0.949
Norms	0.548
Quality	0.737
Topics	0.758
Usefulness	0.788
Average	0.826

Notes: Krippendorff's α assesses intercoder reliability for each hand-coded theme in open-ended responses. "Average" row excludes "Norms" α value since we drop it from the analysis.

C.II Results

To verify the feasibility of our experimental design, we ran two identical pilot studies, via CloudResearch (n = 323) on June 13, 2024 and via the Harvard Digital Lab for the Social Sciences (DLABSS; n = 773) on July 5-16, 2024. Participants were asked to evaluate two state-specific news sites. Participants were randomly assigned with equal probability to either view two screenshots of those sites' homepages (low-information "images" condition) or to follow two live links to evaluate the sites in new browser tabs (full-information "links" condition).

We measured compliance in the links condition in two ways. First, we embedded JavaScript into the survey instrument to record whether participants clicked one or both links. Second, we examine time spent on the survey overall and on the particular page in which we ask participants to click the links and return to the survey once they're content with their investigation of the two sites. By the first metric, 83.6% of participants assigned to the links condition complied in the DLABSS sample and 62.7% complied in the CloudResearch sample (77.1% in the pooled sample). To account for the possibility that participants could have viewed the stimulus sites by copying and pasting the links without actually clicking them, we added a second JavaScript code from Graham (2023) to our NORC survey experiment.

A second concern with using live links is that it could impact the rate of survey drop-off. If drop-off is high – or if drop-off is different across treatment conditions – it would threaten internal validity. Overall, 88.2% of pilot study participants who started the survey completed the survey. Reassuringly, the rate of survey completion was similar in the links (87.4%) and images (89.0%) conditions.

Table 1 in the main body text presents the results, by treatment condition, for the eval-

Table C1: Sample Characteristics by Pilot Study

	CloudResearch	**DLABSS**	**Pooled**
Sex			
Female	40.6%	28.6%	32.1%
Education			
Some college or less	37.0%	29.7%	31.9%
Bachelor's degree	46.0%	32.7%	36.8%
Post-graduate degree	17.1%	37.6%	31.3%
Race			
White	64.1%	86.9%	79.8%
Black	16.4%	1.1%	5.9%
Hispanic	6.5%	2.8%	3.9%
Other	13.0%	9.2%	10.4%
Age			
18–44	75.2%	16.6%	34.6%
45–64	23.9%	32.0%	29.5%
65+	0.9%	51.4%	35.8%
Household Income			
Less than $50,000	29.3%	32.5%	31.4%
$50,000–$99,999	44.2%	33.4%	37.0%
$100,000 or more	26.5%	34.1%	31.6%
Partisan Identification			
Democrat	58.1%	41.0%	46.6%
Independent	10.2%	13.5%	12.4%
Republican	31.7%	45.5%	41.0%

Notes: Total sample size = 1,096; CloudResearch sample size = 323; DLABSS sample size = 773. Percentages may not sum to 100 due to rounding or missing responses.

uations index. Table C2 reports forced-choice preferences across treatment conditions and platforms. Participants in the low-information *images* condition were more likely to select Metric Media sites: by 21 percentage points on CloudResearch and 1 point on DLABSS. In contrast, participants in the *links* condition preferred the journalistic news sites by 16.3 percentage points (CloudResearch) and 9.4 points (DLABSS). Notably, those in the images condition were more equivocal overall: a significantly higher share selected the "no preference" option. This is further proof of the importance of using live links; the extra affordances of the live websites, compared to the limited information that can be gleaned from screenshots, allow participants to more confidently evaluate their relative preference between the sites.

Table C2: Forced-Choice Preferences by Platform and Treatment Condition

Platform	**Condition**	**Journalistic**	**Metric Media**	**No Preference**
CloudResearch	Images	24.2%	45.2%	30.6%
	Links	50.0%	33.7%	16.3%
DLABSS	Images	31.7%	32.7%	35.4%
	Links	39.8%	30.4%	29.3%

Table C3 shows that these patterns persist when adjusting for respondent characteristics. Across both outcome measures, being in the links condition significantly increased the likelihood of favoring the journalistic site. This effect holds in both the full sample and the restricted set of "compliers." Additional covariates, such as partisanship and media trust, show expected relationships.

Table C4 summarizes results from hand-coding participants' open-ended responses, based on the coding scheme described in Appendix section B.III. Mentions of topic content and bias were most common, followed by references to quality and layout. Importantly, participants in the links condition were significantly more likely to cite indicators of quality and bias, features that are more apparent when browsing an interactive site. The differences by treatment condition revealed in Table C3 are meaningful. It is easier to assess whether a site is biased, the quality of its writing and reporting, and whether the existence of pop-up ads or prompts to subscribe interrupt the browsing process when viewing a live site. Without these additional cues, participants in the images condition fell back on considerations more easily ascertained from the screenshots – e.g. site layout, familiarity with the name, and even blanket distrust of media – all elements that scholars would say have less clear normative value in informing news consumers.

Many responses point to structural challenges for traditional journalism. Participants frequently cited intrusive advertisements or paywall prompts as off-putting. As one respondent put it: "I just dislike [the Raleigh *News & Observer*] because they do the ads and the bait-and-switch of requiring registration and payment. Give me information, don't extort me for it." Another noted: "I found that Site 2, The *Wisconsin State Journal*, was absolutely LOADED with irritating advertisements. Because of this, I would immediately exit a site such as this." Some also critiqued content bundling, where soft news (like sports or local features) is prioritized on the homepage. Traditional newspapers attempt to attract a diverse readership by bundling together different types of news – an economic incentive

Table C3: Treatment Effects from Multivariate Regression Models

	Evaluation Index		Site Preference	
	(1) Full	(2) Compliers	(3) Full	(4) Compliers
Links treatment	0.058**	0.054**	0.112**	0.100**
	(0.015)	(0.017)	(0.030)	(0.034)
Independent ID	−0.061*	−0.058	−0.205**	−0.197**
	(0.027)	(0.031)	(0.052)	(0.060)
Republican ID	−0.101**	−0.104**	−0.173**	−0.146**
	(0.019)	(0.021)	(0.037)	(0.042)
Low media trust	−0.071**	−0.071**	−0.109**	−0.122**
	(0.018)	(0.020)	(0.035)	(0.040)
Prefer local papers	0.031	0.024	0.084	0.090
	(0.047)	(0.057)	(0.092)	(0.112)
Prefer national papers	0.062**	0.059*	0.141**	0.143**
	(0.021)	(0.023)	(0.041)	(0.046)
Medium political interest	−0.038*	−0.029	−0.076*	−0.074
	(0.019)	(0.021)	(0.037)	(0.042)
Low political interest	−0.058*	−0.062	−0.072	−0.077
	(0.029)	(0.034)	(0.057)	(0.067)
CloudResearch source	−0.035	−0.043	−0.003	−0.054
	(0.023)	(0.025)	(0.044)	(0.050)
Observations	952	812	953	813
"Compliers" only		✓		✓
Controls included	✓	✓	✓	✓

Notes: Demographic controls include: gender, education level, 5-point ideology, 3-point income, and race. Additional controls include: main source of news, self-reported familiarity with each website, general media trust, difference in the proportion of hard news stories between sites, and site ownership type. Evaluation difference is the difference in a standardized index of four credibility items between journalistic sites and Metric Media sites. Standard errors clustered at the state level. * $p < .05$, ** $p < .01$.

Table C4: Open-Ended Response Coding by Treatment Condition

Theme	**Links Condition**	**Images Condition**
Topics covered	39.4%	33.0%
Partisan bias	38.7%	22.0%
Reporting quality	17.7%	10.8%
Site layout	14.9%	17.4%
Name familiarity	8.4%	15.8%
Clickbait	7.6%	6.3%
Usefulness	7.2%	6.8%
Advertising	6.5%	1.4%
Metadata (e.g., bylines, ownership)	5.6%	5.9%
Multimedia content	5.2%	6.1%
Norms of journalism	4.6%	4.8%
Blanket distrust	2.2%	2.9%

that algorithmic sites don't face. Several participants rated the journalistic news sites poorly when they led with sports content or other soft news content that one participant described as "local flower show content." One participant dismissed such content as "local flower show content."

Some respondents cited what we label metadata – aspects of the site such as ownership, duration of existence, and list of reporting staff. Relatively few responses mentioned the fact that most Metric Media articles do not list an author in the byline and instead list the site's name as the byline. We had *a priori* expected this to be a commonly sited red flag. The fact that it was little mentioned (though one participant did ask "Do the news articles have real people authoring them?") led us to believe that the American Library Association intervention would be effective at improving discernment between human- and algorithmically generated websites. One of the recommendations on that tipsheet is to "Assess the credibility of the author. Do a quick Google search on the author. What is their expertise? What organization do they represent?"

D TESS Survey Details

This section provides additional information about the nationally representative survey fielded through Time-sharing Experiments for the Social Sciences (TESS). The experiment was conducted by NORC at the University of Chicago using their AmeriSpeak Panel, a probability-based panel designed to be representative of the U.S. adult population. Participants were randomly assigned to treatment and control conditions. Table D1 displays demographic characteristics of the sample participants. The survey was administered online from January 31 to February 16, 2025.

All experimental procedures, including randomization and question wording, were pre-registered prior to fielding. Our pre-analysis plan was submitted to OSF on February 2, 2025, prior to receipt of any data (https://osf.io/drpuv/). The survey instrument included measures of participants' perceptions of credibility, informativeness, and trustworthiness for each news source shown, as well as their preferred source for local political information. The structure and wording of these questions mirrored those used in our pilot studies, allowing for comparability across samples.

Table D1: Unweighted Sample Characteristics (NORC)

Characteristic	**Unweighted Proportion**
Gender	
Female	51.5%
Male	48.5%
Education	
Bachelor's degree or higher	34.0%
High school or less	25.9%
Some college/Associate's degree	40.1%
Race/Ethnicity	
White	61.7%
Black	12.2%
Hispanic	18.0%
Other	8.1%
Age	
18–29	19.5%
30–45	30.1%
46–62	25.8%
63+	24.6%
Household Income	
$100k or more	28.5%
$50k–$99,999	35.3%
Less than $50k	36.2%
Region	
Midwest	26.4%
Northeast	13.8%
South	34.7%
West	25.2%
Partisan Identification	
Democrat	43.6%
Independent/Other/Unknown	20.2%
Republican	36.1%

D.I Empirical Specifications

To test for treatment effects, we run regressions of the following form:

$$Y_{i,s} = \alpha + \beta T_i + \gamma_s + \delta \mathbf{X}_i + \epsilon_{i,s}$$

where $Y_{i,s}$ is the outcome variable (forced choice or difference in evaluations index) for individual i in state s; T_i is an indicator variable equal to 1 if the participant is in the treatment condition; γ_s are state fixed effects; $\mathbf{X}_i$ are individual-level demographic controls (sex, age, education, race, household income, ideology, party ID, main source of news, media trust, and home internet access); and $\epsilon_{i,s}$ is the error term. In addition to the individual-level demographic characteristics, we also control for several additional variables related to the experimental design: familiarity with the Metric Media site and the journalistic site (binary variables equal to 1 if familiar, 0 otherwise); a binary indicator for stimulus order (1 if the journalistic site is presented first), and three variables related to the stimulus sites: owner type collapsed into four categories (financial services, privately owned newspaper chain, publicly traded newspaper chain, and independent/other); a binary variable equal to 1 if a human-produced site had articles labeled subscriber-only, pop-up prompts to subscribe, a metered paywall, or pop-ups appearing indicating the number of free articles remaining after several minutes of scrolling and 0 otherwise; and two variables indicating the proportion of soft news on each state's human-produced and algorithmic site over the course of data collection. We run all analyses using weighted OLS regressions with standard errors clustered by state.

Figure 4 in the main text breaks down the estimated treatment effects. Figure D1 shows the raw evaluation index scores for both Journalistic and Metric Media sites, across treatment and control groups, along with 95% confidence intervals. Table D2 shows various regression specifications estimating the effects of the literary tipsheet intervention, with column 5 being the preferred specification with all relevant control variables. While point estimates are generally positive and in the expected direction across in all specifications, they are small and imprecise, with wide confidence intervals.

Table D3 reports how topic-level differences in content between paired sites affect participants' site preferences, disaggregated by party identification. Overall, approximately 6.5% of the variation in journalistic site preference is attributable to content differences, with the largest explanatory power observed among Democrats (within $R^2 = 0.088$) and the smallest among Independents (within $R^2 = 0.026$). Democrats were especially responsive to differences in coverage of education, immigration, and race—favoring journalistic sites when those topics were covered more extensively. Republicans showed smaller and less consistent effects, while Independents exhibited weak associations across all topics.

To better understand which aspects of credibility matter most for audience choices, we disaggregate the overall Evaluation Index and examine how specific evaluation components predict site preference. As shown in Table D4, perceived informativeness and accuracy emerge as the strongest predictors of preferring the Journalistic site. The coefficient on informativeness was particularly large and highly statistically significant ($p < 0.001$), while accuracy also reached conventional significance levels ($p < 0.05$). In contrast, the estimated effects

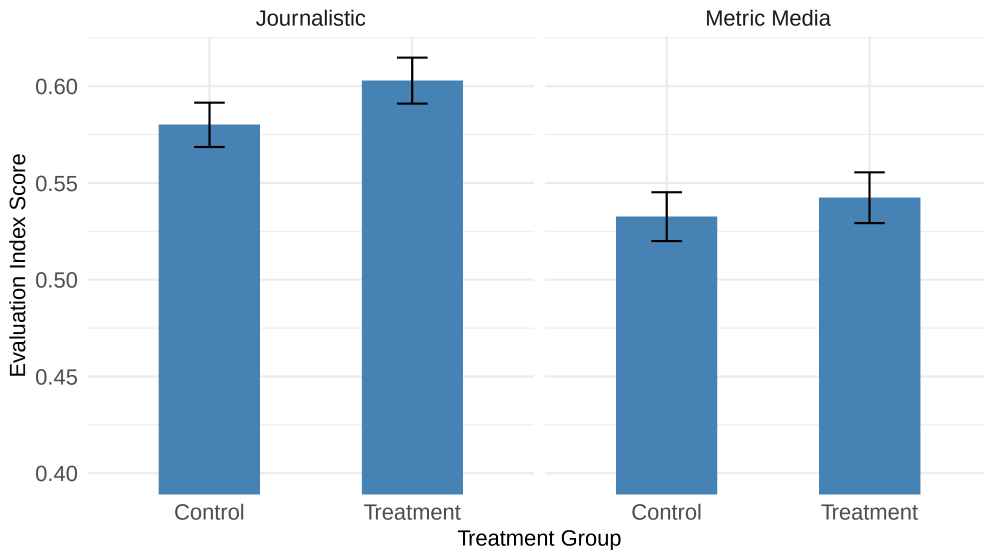

Figure D1: Evaluation of Journalistic and Algorithmic Sites by Treatment Group

Notes: Bars represent mean evaluations of the journalistic and algorithmic sites by treatment condition, with 95% confidence intervals. Evaluation scores range from 0 to 1, where higher values indicate greater perceived credibility.

for perceived reliability and trustworthiness were smaller and less precise, with p-values of 0.067 and 0.149, respectively. These results suggest that interventions aiming to boost news discernment may be most effective if they increase attention to the informational content and accuracy of the source – features that appear to drive audience preferences the most – rather than to the trustworthiness of a site, which is more related to the idea of source credibility.

D.II Heterogeneous Treatment Effects

To explore treatment effects among theoretically relevant subgroups, we preregistered a series of heterogeneous treatment effects analyses using regressions of the following form:

$$Y_{i,s} = \alpha + \beta_1 T_i + \beta_2 C_i + \beta_3 (T_i \times C_i) + \gamma_s + \delta \mathbf{X}_i + \epsilon_{i,s}$$

where C_i is the relevant covariate and $(T_i \times C_i)$ is the interaction of being in the treated group with the relevant covariate. Note that C_i terms are removed from the $\mathbf{X}_i$ term for these regressions.

Figure 4 in the main text displays the results of the heterogeneous treatment effects analyses for the full sample, and Figure D2 shows the results for compliers only. Broadly, we do not find strong evidence of heterogeneous effects across subgroups. We do see statistically significant effects for "moderates" in the full sample, and marginally significant ($p < .10$) effects for independents in both the full sample and compliers sample, though given the multiple hypotheses being tested, this is far from conclusive..

Table D2: Regression Models Predicting Preference for Journalistic Site

	Dependent Variable: Prefer Journalistic Site				
	(1)	(2)	(3)	(4)	(5)
Treatment	0.048	0.043	0.040	0.040	0.038
	(0.035)	(0.045)	(0.042)	(0.042)	(0.040)
Familiar w/ Metric Site			-0.091^{*}	-0.091^{*}	-0.096^{*}
			(0.053)	(0.053)	(0.052)
Familiar w/ Journalistic Site			0.220^{**}	0.220^{**}	0.214^{**}
			(0.038)	(0.038)	(0.038)
Reads Local Newspaper				0.181^{**}	0.178^{**}
				(0.083)	(0.085)
Media Trust					0.047
					(0.037)
Observations	1,432	1,431	1,431	1,431	1,428
Adjusted R^2	0.002	0.066	0.098	0.098	0.098
Within Adjusted R^2	–	0.025	0.058	0.058	0.058
State FEs	No	✓	✓	✓	✓
Demographic Controls	No	✓	✓	✓	✓
Clustered SE	No	✓	✓	✓	✓

Notes: Demographic controls include: gender, education level, 5-point ideology, 3-point income, and race. Additional controls include: main source of news, self-reported familiarity with each website, general media trust, difference in the proportion of hard news stories between sites, and site ownership type. Evaluation difference is the difference in a standardized index of four credibility items between journalistic sites and Metric Media sites. Standard errors clustered at the state level. * $p < .05$, ** $p < .01$.

Table D3: Content Effects on Preference for Journalistic Site, by Party ID

	All	**Democrats**	**Independents**	**Republicans**
Δ Business	0.038	0.127	−0.179	0.028
	(0.088)	(0.158)	(0.277)	(0.110)
Δ Crime	0.031	0.042	0.110	0.129
	(0.078)	(0.134)	(0.163)	(0.127)
Δ Education	−0.142**	−0.220**	0.004	−0.167
	(0.059)	(0.080)	(0.156)	(0.102)
Δ Gender	0.109	0.049	0.008	0.175
	(0.094)	(0.169)	(0.286)	(0.138)
Δ Healthcare	0.093	0.122	−0.212	0.148
	(0.086)	(0.111)	(0.204)	(0.183)
Δ Immigration	−0.339**	−0.601**	−0.359	−0.067
	(0.128)	(0.171)	(0.364)	(0.218)
Δ Local Government	0.032	−0.160**	0.062	0.172
	(0.063)	(0.076)	(0.141)	(0.106)
Δ Race	0.044	0.211**	−0.219	−0.027
	(0.057)	(0.093)	(0.129)	(0.086)
Δ Sports	0.027	−0.017	−0.215	−0.031
	(0.055)	(0.103)	(0.174)	(0.108)
Δ Economy	−0.100	−0.049	−0.019	−0.180
	(0.079)	(0.118)	(0.190)	(0.121)
Observations	1,203	531	215	434
Adjusted R^2	0.087	0.128	0.013	0.115
Within Adjusted R^2	0.065	0.088	0.026	0.047
State FEs	✓	✓	✓	✓
Demographic Controls	✓	✓	✓	✓
Clustered SE	✓	✓	✓	✓

Notes: Demographic controls include: gender, education level, 5-point ideology, 3-point income, and race. Additional controls include: main source of news, self-reported familiarity with each website, general media trust, difference in the proportion of hard news stories between sites, and site ownership type. Evaluation difference is the difference in a standardized index of four credibility items between journalistic sites and Metric Media sites. Standard errors clustered at the state level. * $p < .05$, ** $p < .01$.

Table D4: Predictors of Site Choice: Evaluation Components

Variable	**Coefficient**	**Standard Error**	***p*-value**
Treatment (Tipsheet)	0.037	0.039	0.356
Trustworthiness Difference	0.175	0.119	0.149
Reliability Difference	0.229	0.122	0.067
Accuracy Difference	0.205	0.085	0.019^{*}
Informativeness Difference	0.421	0.080	$< 0.001^{**}$

Notes: Outcome is a binary indicator for preferring the journalistic site. Differences are calculated as the evaluation of the journalistic site minus the algorithmic site for each item. Model includes state fixed effects and robust standard errors clustered by state. $N = 1{,}403$. * $p < .05$, ** $p < .01$.

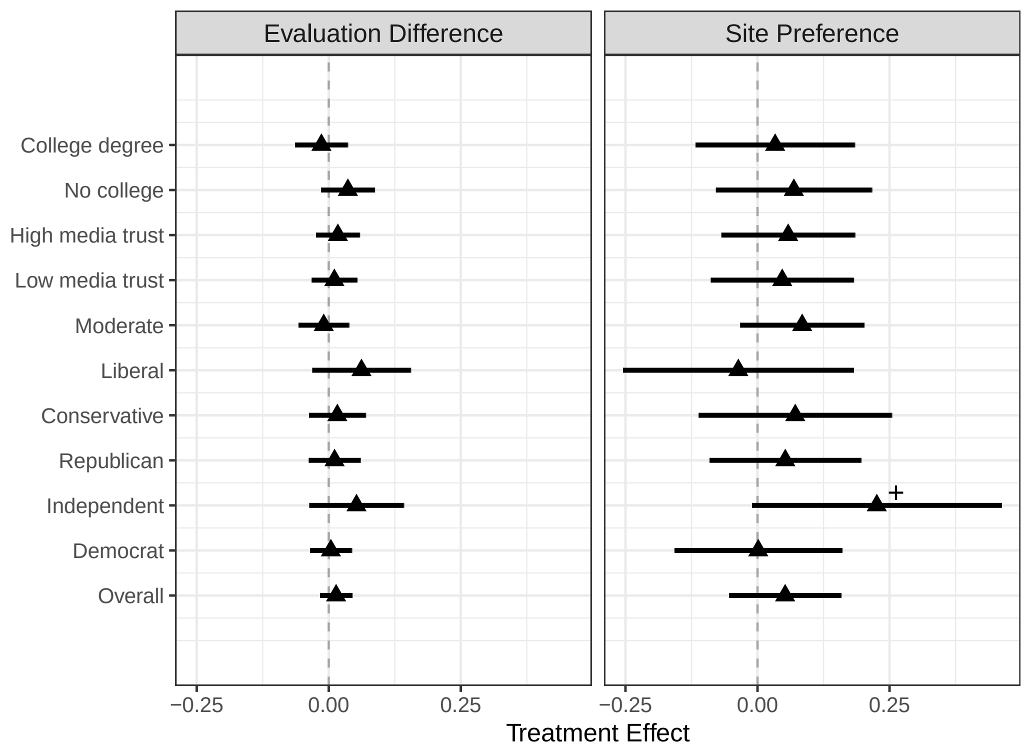

Figure D2: Estimated Treatment Effects by Subgroup - Compliers Only

Notes: Estimated treatment effects on site preference and evaluation difference by participant subgroup, with 95% confidence intervals. Results shown for the "compliers" sample only. $^{+}$ $p < .10$ * $p < .05$, ** $p < .01$.

Appendix References

Alba, Davey, and Jack Nicas. 2020. "As Local News Dies, a Pay-for-Play Network Rises in Its Place." *New York Times* (October).

Green, Miranda, Jennifer Smith Richards, Priyanjana Bengani, and Sarahbeth Maney. 2024. "Fossil Fuel Interests Are Working to Kill Solar in One Ohio County. The Hometown Newspaper Is Helping." *ProPublica* (October).

Lacy, Stephen, Brendan R. Watson, Daniel Riffe, and Jennette Lovejoy. 2015. "Issues and best practices in content analysis." *Journalism and Mass Communication Quarterly* 92 (4): 791–811.

Stanley-Becker, Isaac, and Josh Dawsey. 2023. "Inside a private portal from GOP campaigns to local news sites." *Washington Post* (April).

Steurer, Mary. 2025. "Judge denies Greenpeace request to investigate mailer critical of DAPL protests." *North Dakota Monitor* (January).

Tow Center for Digital Journalism. 2024. ""Pink Slime": Partisan journalism and the future of local news.".