“Or they could just not use it?”: The Paradox of AI Disclosure for Audience Trust in News

Pre-Print – Not Peer-Reviewed

Benjamin Toff[1], Felix M. Simon[2]

[1]Hubbard School of Journalism & Mass Communication, University of Minnesota, USA, bjtoff@umn.edu
[2]Oxford Internet Institute, University of Oxford, Oxford, United Kingdom, felix.simon@oii.ox.ac.uk

December 1st, 2023

Abstract

The adoption of artificial intelligence (AI) technologies in the production and distribution of news has generated theoretical, normative, and practical concerns around the erosion of journalistic authority and autonomy and the spread of misinformation. With trust in news already low in many places worldwide, both scholars and practitioners are wary of how the public will respond to news generated through automated methods, prompting calls for labeling of AI-generated content. In this study, we present results from a novel survey-experiment conducted using actual AI-generated journalistic content. We test whether audiences in the US, where trust is particularly polarized along partisan lines, perceive news labeled as AI-generated as more or less trustworthy. We find on average that audiences perceive news labeled as AI-generated as less trustworthy, not more, even when articles themselves are not evaluated as any less accurate or unfair. Furthermore, we find that these effects are largely concentrated among those whose pre-existing levels of trust in news are higher to begin with and among those who exhibit higher levels of knowledge about journalism. We also find that negative effects associated with perceived trustworthiness are largely counteracted when articles disclose the list of sources used to generate the content. As news organizations increasingly look toward adopting AI technologies in their newsrooms, our results hold implications for how disclosure about these techniques may contribute to or further undermine audience confidence in the institution of journalism at a time in which its standing with the public is especially tenuous.

Keywords: Artificial Intelligence, AI, Generative AI, Journalism, LLM, News, Trust

Introduction

Use of Generative AI in News and Journalism

The use of generative AI within the news industry has grown significantly in recent years. Since the public launch of ChatGPT, a Large Language Model (LLM), in November 2022 by the US-based start-up OpenAI, and similar LLMs such as Google's "Bard" and Meta's LLaMA, many news organizations have shown strong interest in this sub-branch of artificial intelligence. LLMs equipped with the function to generate authentic, multi-modal content spanning textual and visual elements are perceived as having the potential to profoundly reshape the dynamics of news production, distribution, and consumption as well as the information ecosystem and publisher's business models (Simon & Isaza-Ibarra, 2023). Consequently, growing numbers of publishers—many of whom had already adopted AI in the past—have been experimenting with LLMs for a variety of tasks, such as generating summaries, creating illustrations, re-formatting content, performing copy-editing, increasing workflow efficiency, SEO-related tasks, and a range of other tasks in editorial, product, and distribution (Beckett, 2023). Many of these applications are consistent with prior uses of AI (mostly forms and applications of machine learning)[1] within news organizations, which continue apace.

While it is too early to say if some of the loftier promises and expectations around this set of technologies will be borne out by reality, a growing number of questions have been raised around what effects they may have on the public, particularly with respect to safety and accuracy, bias, or privacy. One overarching concern involves effects of AI use in news audience trust. With trust in the news media already low and declining in many places worldwide (Hanitzsch, Van Dalen, & Steindl, 2018; Kalogeropoulos et al., 2019), scholars and practitioners are wary of how the public will respond to the proliferation of news generated through automated methods, and LLMs in particular. Many fear that the use of AI in news production could further damage trust, with related knock-on effects on publishers' credibility with the audiences they seek to serve. While a growing number of publishers have begun responding to these concerns by adding labels to AI-generated content, there is no shared consensus about what such disclosure should look like nor agreement over what level of AI-involvement should trigger labeling (Becker, Simon, Crum, 2023).

At the same time, there is also reason to suspect that some audiences may greet AI-generated news *more* positively precisely because of the low esteem that many in the public already hold professional journalism. Previous surveys from the Reuters Institute for the Study of Journalism, for example, have shown that many say they prefer having

[1] We define AI following Mitchell (2019) as the computational simulation of human capabilities in tightly defined areas, most commonly through the application of machine learning approaches, a subset of AI in which machines learn, e.g. from data or their own performance.

news selected and curated for them algorithmically over relying on journalists' decisions about newsworthiness (Fletcher, 2023; Thurman et al., 2019). Likewise, a recent experimental study (Cloudy et al., 2023) examining attitudes toward journalists on social media when presented as "NewsBots" versus humans found that the former was associated with reductions in perceptions of bias. This notion has also been raised in industry circles where the potential for these technologies' and their perceived "mechanical accuracy" to actually engender rather than undermine trust has been seen as a potentially promising application of these technologies, at least when it comes to serving segments of the public who are distrusting about the editorial judgments of professional journalists.

In this study, we present results from a novel experiment conducted using actual AI-generated journalistic content created by a California-based technology start-up. We test whether audiences in the US, where trust in news is particularly polarized along partisan lines, perceive news content labeled as having been generated through automated processes as more or less trustworthy, varying the degree to which the story's topic is political in nature. We find on average that audiences perceive news labeled as having been generated with the help of AI as less trustworthy, not more, even where AI-generated news articles themselves are not evaluated as any less accurate or unfair. Furthermore, we find that these effects are largely concentrated among those whose pre-existing levels of trust in news are higher to begin with and among those who exhibit higher levels of knowledge about journalism. We also find that negative effects associated with perceived trustworthiness are largely counteracted when articles disclose the list of sources used to generate the content. As news organizations increasingly look toward adopting AI technologies in their newsrooms, our results hold implications for how disclosure about these techniques may contribute to or further undermine audience confidence in the institution of journalism at a time in which its standing with the public is especially tenuous.

Literature Review

Audience Attitudes about Algorithmic Technologies and AI

Despite concerns regarding audience attitudes towards AI in news production, there remains a lack of concrete evidence on the subject. *User attitudes towards algorithmic technologies have so far predominantly been investigated within the domain of news recommendations* (Mitova et al., 2023). As Mitova et al. aruge, the theoretical foundation for much of this work is the concept of the "machine heuristic" (Sundar & Kim, 2019) which posits that algorithms are perceived positively by many due to the belief that they exhibit greater neutrality and fairness compared to humans (Mitova et al., 2023, p. 92) due to their lack of human motives and emotions.

While certain studies indicate that news recommenders are viewed as fair and useful as human editors (Araujo et al., 2020), the broader perspective remains ambiguous. Fletcher's work (2023) suggests a general skepticism towards all forms of news selection, whether executed by humans or algorithms, referred to as "generalized skepticism" (Fletcher and Nielsen, 2018). Notably, higher interest in news correlates with increased approval for various forms of news selection, including algorithmic methods. Fletcher further contends that "approval for both algorithmic news selection and editorial news selection is significantly higher for those with higher levels of trust." In a survey across 10 European countries, Araujo et al. (2023) also explore audience attitudes towards algorithmic technologies, particularly in the media sector. The study reveals that trust and political attitudes influence perceptions of AI, with higher institutional trust correlating positively with attitudes towards AI in news recommendations and moderation, although they found that higher media trust tends to yield more negative views on AI in news-related tasks. Political orientation also plays a role, as right-wing individuals express more positivity towards AI in news. Notably, higher income and privacy orientation were identified as additional factors influencing optimism or pessimism about AI adoption at the societal level.

While some scholarship has specifically examined journalists' perceptions of automatically generated (textual) news, much of this work has focused on how well these tools are able to approximate the work of trained journalists. Milosavijevic and Vobic (2019), for example, find that journalists are often critical of AI-generated content when it comes to narrative and editorial quality, thus pointing to the importance of human oversight (Diakopoulos, 2020; Thurman et al. 2017). Other work has highlighted the need for editorial interventions around ensuring coverage includes meaningful connections with e.g. local readership (Thäsler-Kordonouri & Barling, 2023, p. 16). Yet few studies to date have considered *audience perceptions of automated, AI-informed production of news,* and the body of existing literature on these topics is scant and inconsistent (Tandoc et al., 2020) or observational in nature with inevitable questions around temporal validity (Munger, 2023) given such a fast-changing phenomenon. A recent representative study of the Swiss population (Vogler et al., 2023) showed that acceptance of AI-generated articles in journalism was low, with only 29.1% willing to read news entirely generated by AI, while the figure for non-AI generated news was 84.3%. Acceptance, however, varied by topic; it was highest for routine reporting such as weather or stock market developments and lowest for hard-news areas like culture, science, and politics. The majority of respondents believed that AI use in journalism should be declared and transparent. They also viewed its impact on the quality of journalistic content somewhat negatively, with 61.3% agreeing that reporting quality would deteriorate due to AI use. A representative survey of US-adults conducted by Monmouth University showed similar results, with 78%

of respondents expressing a negative view about the prospect of news articles being written by AI, deeming it a "bad thing."[2]

This evidence aside, other studies have tended to highlight the potential for news generated with the help of AI to elicit more favorable attitudes from the public. Wölker and Powell (2018) found little difference in user perceptions around message and source credibility for various degrees of automated journalism (AJ). Thurman et al. (2023) likewise studied audience perceptions of short-form news videos produced with varying degrees of AI involvement. Findings indicate that automated news videos require human supervision to attain comparable likability to human-made videos. On average, consumers expressed similar levels of liking for automated news videos and manually produced ones, but only when the automation process involved human post-editing. Longoni et al. (2022) looked at whether individuals believe news headlines generated by generative AI as much as those crafted by humans using survey experiments with representative US samples. They found that users rated news headlines written by AI as less accurate than those composed by humans. Notably, individuals were more prone to inaccurately rate news headlines generated by AI (compared to those written by humans) as incorrect when they were, in fact, true. Conversely, they were more likely to accurately rate AI-generated headlines as incorrect when they were, indeed, false. Finally, Jang et al (2022), explored the impact of knowledge about automated journalism on people's evaluations of algorithmically generated news. In two experiments, they found that individuals with little AJ knowledge prefered human-authored news over algorithmically generated news, while those with high AJ knowledge had an equal or stronger preference for algorithmically generated news. They further found age-specific effects at lower levels of AJ knowledge, where machine-like characterizations enhanced evaluations for younger users, and human-like traits enhanced evaluations for older users. One editor from a German digital news site that has experimented with publishing articles generated through automated processes said in an interview, citing an internal survey of its audiences, "It seems that some readers trust the mechanical accuracy of technology more in certain topics than the error-prone or ideologically shaped person" (Franz, 2023).

Disclosure of AI-generated Content

Despite uncertainty surrounding whether audiences are likely to be favorably or unfavorably predisposed toward journalism generated with the help of AI, even less clear is how news organizations ought to handle the question of disclosing how and when they are using these technologies in their own workflows. A growing number of news organizations are considering or mandating such disclosure when they use AI tools to assist in the creation of content (Becker, Simon, Crum, 2023). Many do so because of

[2] https://www.monmouth.edu/polling-institute/reports/monmouthpoll_US_021523/

longstanding professional norms around transparency (Deuze, 2005; Karlsson, 2010) and a desire to make journalistic routines discernible to the public (Tuchman, 1972). Likewise, outside of journalism, there are growing calls for disclosure around the use of AI generally due to the opaque nature of AI systems and the ways by which these systems arrive at results (see, e.g. Grant et al., 2023). However, at present, the available literature around modes of disclosure of AI-generated content is limited. Even fewer studies examine the effects of disclosure.

Some research has grappled with how media organizations perceive and handle the authorship of algorithm-generated stories, along with the needs, expectations, and rights of their audiences. Early work by Montal and Reich (2017), for example, found that many organizations lacked a solid policy for bylining and disclosuring automated news stories, although respondents emphasized the importance of informing readers about the automated nature of content for credibility. More recent work by Becker et al. (2023) found that out of 52 international news organizations who had published AI guidelines, 90% of organizations mandated that the use of AI has to be disclosed, although 82% did not explicitly specify how this transparency should be communicated, with further difference for the situations in which disclosure would be required (p. 17). In terms of empirical studies looking at the effects of disclosure and labeling of AI content, to our knowledge only Epstein et al. (2023) investigate the same to date, exploring the challenges of effectively labeling AI-generated content in the context of images and videos, considering the ambiguity surrounding AI terminology. They found varying interpretations of AI-related terms proposed as labels. Terms explicitly mentioning AI, such as ‘AI Generated,’ ‘Generated with an AI Tool,’ and ‘AI Manipulated,’ were consistently associated with content generated using AI, while ‘Deepfake’ and ‘Manipulated’ were associated with misleading content. The study underscores the importance of aligning labeling objectives with chosen terms and understanding user perceptions before implementing effective content disclosure policies.

Hypotheses and Research Questions

Given the lack of previous research on the effects of disclosure around the use of AI in generating journalistic content, we designed the present study to test whether transparency and disclosure about the use of such technologies was associated with increases or decreases in trust. On average, we predict that news stories labeled as having been generated with the help of AI will be viewed as (H1a) less trustworthy and (H1b) less accurate given the descriptive findings from prior studies outlined above. We also ask (RQ1) whether news stories labeled as having been generated with the help of AI will be viewed as more or less fair or unbiased. We did not formulate a specific hypothesis on this question because prior research has been too limited to generate specific expectations pertaining to these evaluative dimensions.

Despite negative effects associated with disclosure about the use of AI, we predict that attitudes may be conditional on two factors. First, given that skepticism about the inherent subjectivity of journalists is likely highest among those who are distrusting of news in general (Mont'Alverne et al., 2023; Ojala, 2021), we theorize that prior levels of trust in news will also shape the way audiences perceive disclosures about the use of AI in non-linear ways. Specifically those who are least trusting may be most favorably predisposed to see advantages from "mechanical accuracy" as an alternative to the judgments of professional journalists. Therefore, we predict heterogeneous effects of disclosure: those who are less [more] trusting toward news will view news labeled as having been generated with the help of AI as (H2a) more [less] trustworthy, (H2b) more [less] accurate, and (H2c) less [more] biased or unfair. At the same time, we also theorize that individuals who are more literate about the processes involved in newsgathering and vetting information will also react differently to disclosure than those who are less knowledgeable about journalistic practices—what scholars have called "procedural news knowledge" (PNK) or familiarity with "what legitimate news production and reporting entails" (Amazeen & Bucy, 2019: 419; see also Maksl et al., 2015; Schulz, Fletcher, & Nielsen, 2022). In addition to trust in news, we therefore also predict heterogeneous effects of disclosure related to levels of PNK. Specifically, those who exhibit lower [higher] levels of PNK will view news labeled as having been generated with the help of AI as (H3a) more [less] trustworthy, (H3b) more [less] accurate, and (H3c) less [more] biased/unfair.

Finally, prior research leads us to formulate two additional expectations. First, we predict that disclosure about the list of sources used to generate news stories may be an important factor in shaping the way audiences perceive disclosures about the use of AI. That is, when the underlying list of sources is provided to readers, we predict that respondents will perceive such content as (H4a) more trustworthy, accurate, and fair, while at the same time (H4b) attenuating effects associated with stories labeled as AI-generated. We also hypothesize that treatment effects associated with AI labels will vary in magnitude depending on the topic of the news story. We predict (H5) that heterogeneous treatment effects will be larger [smaller] for more [less] politically contentious news stories since perceptions of bias in news tend to be highest when stories intersect with political partisanship (Tully, Vraga, & Smithson, 2020).

Research Question/Hypothesis	Description
H1a	News stories labeled as AI-generated will be viewed as less trustworthy.
H1b	News stories labeled as AI-generated will be viewed as less accurate.
RQ1	Will news stories labeled as AI-generated be perceived as more or less fair or unbiased?

H2a	Individuals less [more] trusting of news will view AI-generated news as more [less] trustworthy.
H2b	Individuals less [more] trusting of news will view AI-generated news as more [less] accurate.
H2c	Individuals less [more] trusting of news will view AI-generated news as less [more] biased or unfair.
H3a	Individuals with lower [higher] procedural news knowledge (PNK) will view AI-generated news as more [less] trustworthy.
H3b	Individuals with lower [higher] PNK will view AI-generated news as more [less] accurate.
H3c	Individuals with lower [higher] PNK will view AI-generated news as less [more] biased or unfair.
H4a	Disclosure of the list of sources will make AI-generated news more trustworthy, accurate, and fair.
H4b	Disclosure of the list of sources will attenuate effects associated with AI-generated labels.
H5	Heterogeneous treatment effects for AI labels will be larger for more politically contentious news stories.

Methods and Data

To test our hypotheses and investigate our research questions, we conducted a pre-registered, between-subject experiment using actual AI-generated journalistic content as stimulus material provided by a California-based start-up called HeyWire AI, which bills itself as "the industry's first self-prompting, and fully autonomous AI news and content generation engine."[3] Respondents who consented to enroll in the study completed a 5-minute self-administered survey in which they provided basic demographic information and pre-treatment attitudes and were then asked to read one of three HeyWire-created news stories on topics that varied in political contentiousness. One focused on the release of the Hollywood film "Barbie," another on the international BRICS summit of world leaders from the Global South, and a third on the criminal investigation into wrongdoing by the US president's son Hunter Biden. The full text of each of these articles is provided in the supplementary appendices. Respondents were randomly assigned to see versions of these stories either with or without labels disclosing the use of AI to generate this content. In this section we provide more information about the sample, procedures, and relevant measures used in this study.

[3] A pre-registered analysis plan is available at: REDACTED.

Sample

A sample of English-speaking, U.S.-based participants (N = 1,483) was recruited using the Prolific platform, whose panel of respondents has been validated in previous studies (Douglas, Ewell, & Brauer, 2023; Peer et al., 2022). Each was compensated with a small fee for their time. Respondents were somewhat less racially diverse, more educated, and politically engaged than the public at large but balanced across a range of characteristics summarized in Table 1. As in previous research on opt-in online panels (Guess & Munger, 2023), respondents in our study were somewhat more knowledgeable about and familiar with digital media technologies, including the use of AI. In a post-treatment question, when we asked respondents how much they had heard or read about the use of AI to "to write articles that report on news events and information," more than a quarter of the sample said they had heard "a lot" (27.6%) compared to 63% who had heard "a little" and just 9.4% who said they had heard "nothing at all."[4]

While the proliferation of bots on similar platforms has been a source of some concern (Chmielewski & Kucker, 2020), we are generally confident in the quality of our data based on optionally provided open-ended responses. Most respondents (59.8%) left a comment in the open-ended box when asked to "share any other additional thoughts or opinions you may have on the example news article you just read." The median response was 14 words although some participants wrote lengthier comments (the maximum was 163 words). About a third of those who left a comment and saw a label describing the article as generated with the help of AI (33.6%) also specifically referred to "AI" or "Artificial Intelligence" in their response, which suggests that respondents were both actual human respondents and likely paying attention to the treatment stimulus.[5]

[4] Percentages did not differ when we separately analyzed respondents in the control versus treatment conditions.

[5] A much smaller but not insignificant percentage (18.6%) of those who left a comment and were not assigned to see a label also mentioned "AI" or "Artificial Intelligence," which suggests that the subject itself may be particularly salient for many Prolific users.

Table 1. Descriptive information about the Prolific sample.

Age	
18 to 24	8.0%
25 to 34	27.8%
35 to 44	24.0%
45 to 54	17.7%
55 to 64	13.5%
65 years old or older	9.0%
Gender	
Identifying as male	49.5%
Race and Ethnicity	
Identifying as Black	11.4%
Identifying as Hispanic	9.1%
Education	
College degree	52.4%
Graduate degree	12.1%
Politics	
Not very/at all interested	19.3%
Somewhat interested	35.7%
Very/extremely interested	45.0%
Identifying as Democrat	48.7%
Identifying as Republican	21.1%
News Use (Sometimes/Often)	
Online	49.1%
Social Media	70.2%
TV	54.5%
Cable TV	41.1%
Print	76.9%
Audio	24.3%
Trust in News	
Trust completely	1.8%
Trust somewhat	39.0%
Neither trust nor do not trust	17.9%
Do not trust very much	27.9%
Do not trust at all	13.4%
Procedural News Knowledge (PNK)	
Answered 4 questions correctly	46.5%
Answered 3 questions correctly	28.2%
Answered 2 questions correctly	15.4%
Answered 1 question correctly	8.0%
Answered 0 questions correctly	1.8%
N (respondents)	1483

Experimental Design

Respondents were randomly assigned to two treatment conditions associated with their exposure to the stimulus articles. A control group saw the story alone by itself, attributed under the byline "Intelligent Press," a made-up news organization named in the byline. A random subset also saw a label preceding the stories disclosing that the news organization had used AI to generate the article. The text of this label (see Figure 1) was adapted using language from the HeyWire AI website about its "content engine" and mirrors language used by various news organizations to disclose the use of AI in content production.

> *Intelligent Press relies on a fully autonomous Artificial Intelligence (AI) content engine to identify, curate, and produce newsworthy stories, creating content without any human prompting required.*
>
> **Barbie Takes Hollywood: How Margot Robbie Reinvented the Iconic Toy**
>
> BY *INTELLIGENT PRESS*
> *Updated 4:42 PM CDT, September 5, 2023*
>
> Once upon a time in Hollywood, Margot Robbie was faced with the daunting task of taking on the lead role of Barbie for the highly anticipated summer film. As she delved into the intricacies of the iconic toy, questions arose regarding beauty and sexiness, which she had to ponder on quite a bit.

Figure 1. Label preceding an AI-generated news story used in the study. Full text of the stimuli articles is provided in the supplementary appendices.

In addition to labels disclosing the use of AI to generate news stories, a randomized subset of respondents also saw a list of sources at the end of each story linking to the underlying news reports that the story was based on. An example of this source list is provided in Figure 2. These lists of sources differed for each story as they are specific to the actual news articles HeyWire AI used to generate stimulus materials, but otherwise the composition of these lists of sources were held constant across experimental conditions.


Sources:

- The Daily Mail (UK)
- Variety
- People
- CNN
- The Independent (UK)

Figure 2. Example list of sources accompanying the Barbie article.

Dependent and Independent Variables

Three sets of questions were asked following exposure to the news article stimuli. These items were designed to capture attitudes about the news organization that had produced the content. These include (1) a question specifically about trust: "How trustworthy would you say the news organization is that published this article?" Respondents were asked to place themselves on a scale ranging from 0 ("Not at all trustworthy") to 10 ("Completely trustworthy"). Respondents were also asked (2), "To the best of your knowledge, how accurate is the article you just read?" with responses provided on a 4-point Likert scale ranging from "Not at all accurate" to "Very accurate." Lastly, respondents were asked about (3) how fair or biased they perceived the story to be. This underlying construct was measured by asking respondents for their level of agreement or disagreement with four separate statements: "The article is fair", "The article is unbiased", "The article tells the whole story", "The article separates facts from opinions." These items, captured on a 5-point scale from "Strongly agree" to "Strongly disagree," were adapted from Strömbäck et al. (2020). As responses to these items were reliably consistent with one another (Cronbach's alpha = 0.9), responses were averaged together in a composite index.

For our two key independent variables—prior levels of trust in news and PNK—we used pre-treatment measures drawn from other studies. For trust in news, we asked a question again adapted from Strömbäck et al. (2020) meant to capture generalized views about the trustworthiness of information in the news: "Generally speaking, to what extent do you trust or not trust information from the news media?" with responses coded on a 5-point scale ranging from "trust completely" to "do not trust at all." To measure PNK, we used a 4-item battery of questions designed to measure literacy about the editorial procedures used to generate content and gather news in the U.S. media system. These items were adapted from a previous study by Amazeen & Bucy (2019) (see the supplementary appendices for the items). Those who answered all four items correctly were coded as a 1 on the PNK scale with others assigned fractional scores accordingly. In addition we asked a series of pre-treatment questions used as control variables about age, gender, race and ethnicity, political interest and partisanship, and frequency of different types of news media use.

Analytic approach

To evaluate our hypotheses and research questions, we calculated both average treatment effects, comparing mean responses to each of these dependent variables between treatment and control groups, as well as testing for heterogenous effects by estimating linear models interacting indicators corresponding to treatment conditions with trust in news and/or PNK where relevant. For most analyses, we pool across the three different story topics, with the exception of our test of H5, where we include an additional interaction effect to test for differences in treatment effects associated with the political contentiousness of each story.

Findings

In this section, we begin by presenting average treatments effects across the full sample. Next, we consider heterogeneous effects for prior levels of trust in news in general and PNK. In our last set of analyses, we consider the ways in which these results interact with the inclusion of a list of original sources and/or the degree to which results vary depending on the topic of the underlying story.

Overall, *we find evidence consistent with the theory that audiences do generally perceive news labeled as having been generated with the help of AI as less trustworthy*, even though we do not find that they evaluate the content of these articles as less accurate or more biased. W*e also find evidence that where sources are provided alongside the text, labels disclosing the use of AI do not reduce trust in news*. We also provide exploratory findings about audience attitudes on the use of labels, which aid interpretation of these results.

News labeled as generated with AI are perceived as less trustworthy

Our first hypothesis (H1a) predicted that on average, respondents would evaluate the news organization as less trustworthy when exposed to a label disclosing the use of AI to generate the content. We find evidence consistent with this expectation. Respondents in the control group with no labels evaluated the news organization just above the midpoint on an 11-point scale (mean = 5.9, sd = 2.3), whereas respondents randomly assigned to see a label disclosing the use of generative AI evaluated the organization as less trustworthy (mean = 5.5, sd = 2.4), a statistically significant difference ($t = -3.7$, $p < 0.001$). These results are summarized in Figure 3 below.

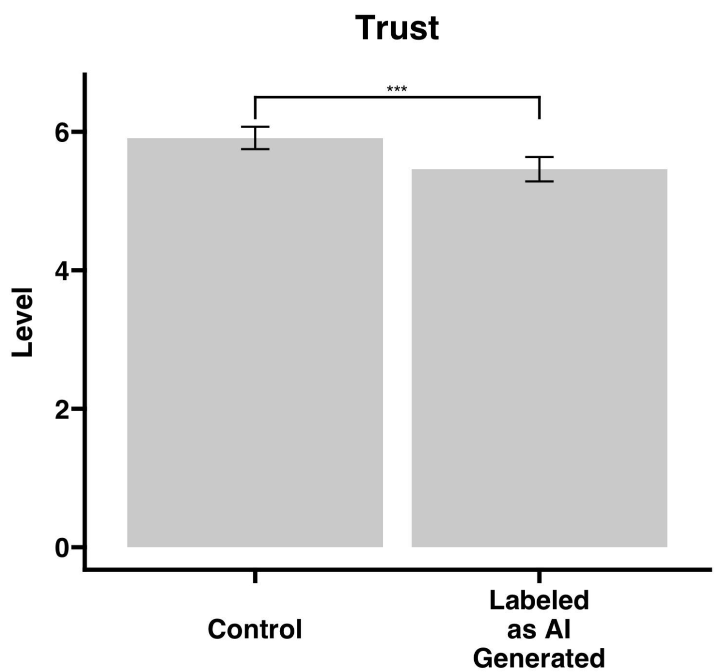

Figure 3. Differences in mean levels of perceived trustworthiness when comparing treatment and control groups.

These significant differences in perceptions about trustworthiness, however, did not extend to evaluations of the accuracy (H1b) or fairness (RQ1) of the coverage itself. Respondents in the control group perceived the story approximately as accurate (mean = 2.06, sd = 0.68) as those in the treatment group (mean = 2.03, sd = 0.62). Likewise, differences in levels of perceived fairness between the control group (mean = 0.37, sd = 0.88) and treatment groups (mean = 0.34 , sd = 0.86) were not distinguishable from one another.

Differences were conditional on prior levels of trust in news and procedural news knowledge

Our second set of hypotheses examined heterogeneity in how respondents reacted to labels disclosing the use of generative AI. We find evidence that prior levels of trust in news in general is associated with differences in these responses. In Figure 4, we summarize results of our analysis testing for non-linear treatment effects related to trust in news (H2a). We find the largest gaps in perceived trustworthiness among those at the highest end of the scale in terms of prior levels of trust in news but no differences among those who otherwise do not trust news at all—findings that are consistent in models both with and without control variables.[6] As we found with regards to average treatment effects, we did not find

[6] See Appendix Table C-1 for the full regression output for both sets of results.

any significant differences in treatment effects for perceived accuracy (H2b) or fairness (H2c).

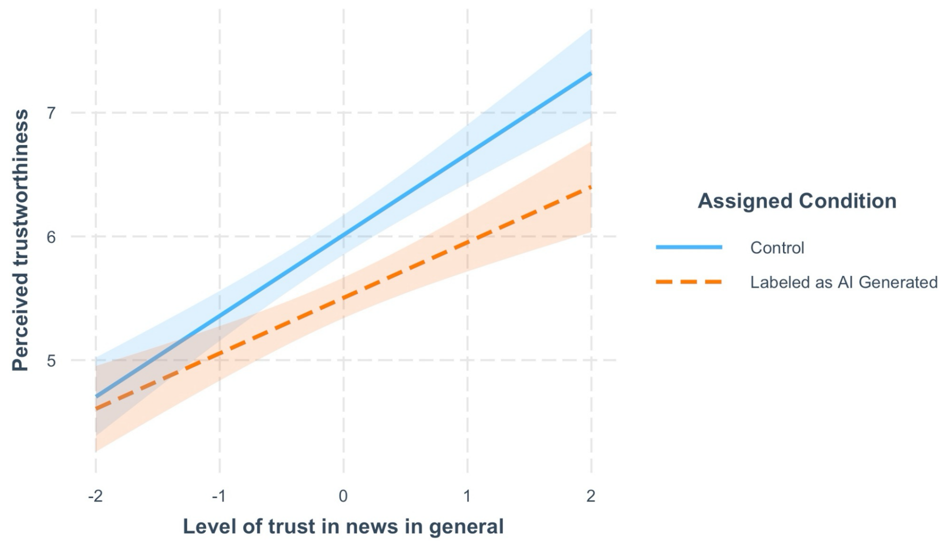

Figure 4. Gaps in perceived trustworthiness associated with disclosure about the use of generative AI varied as a function of prior levels of trust in news, holding all other variables at their mean values.

Likewise, we also found evidence of heterogeneity with respect to respondents' levels of procedural news knowledge (PNK). When we interacted PNK with exposure to the treatment labels, we found that the additional information about stories being generated with the help of AI was associated with lower levels of trustworthiness only for those exhibiting higher levels of PNK. There were no differences in perceived trustworthiness for those with lower levels of PNK, confirming H3a. No non-linear differences in treatment effects were found for perceived accuracy (H3b) or fairness (H3c).[7] In Figure 5 we plot differences in trustworthiness perceptions as a function of exposure to treatment and levels of PNK.

[7] Full regression model output is provided in Appendix Table C-2.

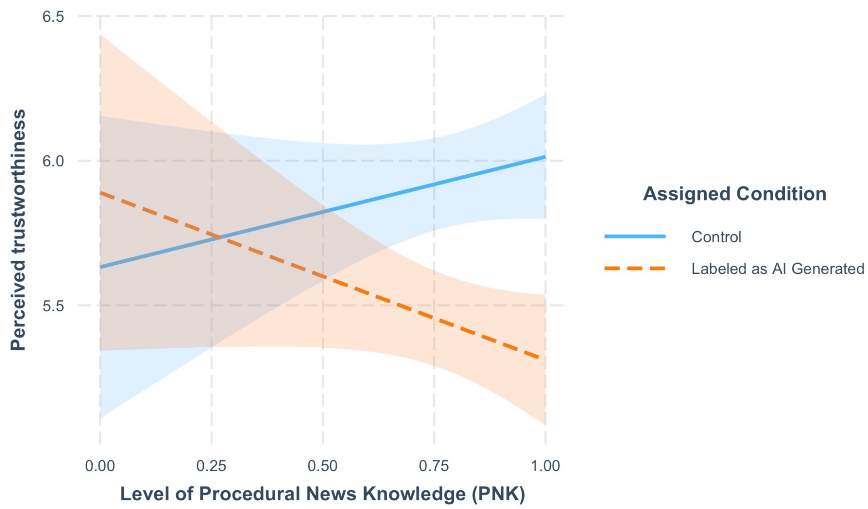

Figure 5. Gaps in perceived trustworthiness associated with disclosure about the use of generative AI varied as a function of procedural news knowledge (PNK), holding all other variables at their mean values.

Inclusion of underlying sources reduced effects on trustworthiness with minimal variation found by topic

Our final set of expectations concerned the inclusion of a list of sources used to generate articles and the degree to which results varied by topic. We first tested for average treatment effects associated with the inclusion of the source list (H4a) and found mixed evidence that this information affected attitudes about the content respondents were shown. We found no statistically significant differences in respondent perceptions of the trustworthiness of the source (mean = 5.8, sd = 2.3) compared to the control condition without the source list (mean = 5.6, sd = 2.4). Likewise, we found no differences in how fair respondents evaluated the content when comparing stories with the source list provided (mean = 0.38, sd = 0.87) versus the control condition (mean = 0.33, sd = 0.87). However, we did find slight differences in evaluations of the accuracy of stories when the source list was provided (mean = 2.07, sd = 0.65) compared to the control condition (mean = 2.01, sd = 0.65), just reaching conventional thresholds of statistical significance ($t = 1.81$, $p < 0.1$). Effects associated with the inclusion of these source lists, however, were largely conditional on whether respondents were also provided with labels about the use of generative AI. In other words, when we tested for an interaction between the label disclosure and the source list disclosure conditions (H4b), we found treatment effects from labels were found mainly only when respondents were *not* provided with a list of sources (see Figure 6 and Table C-3 in the appendix).

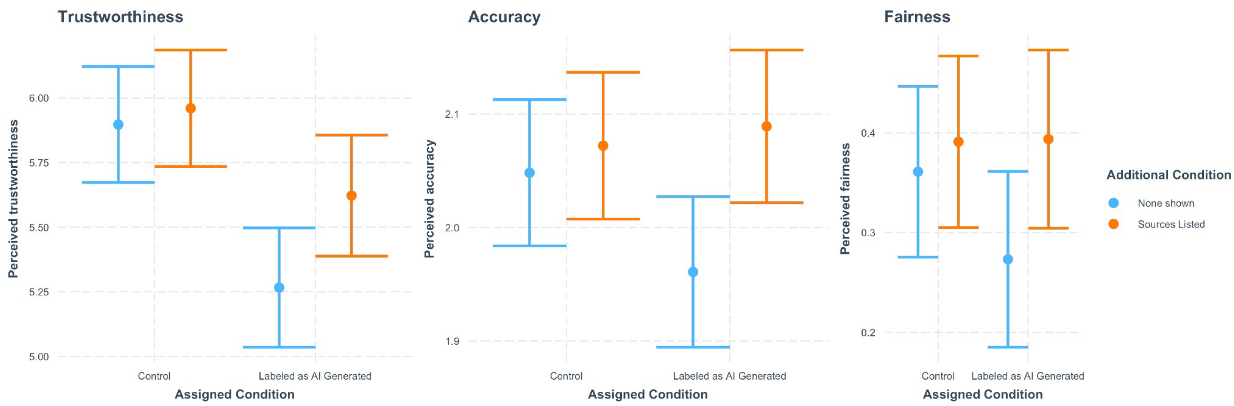

Figure 6. Providing a list of sources at the end of the article moderated treatment effects associated with stories labeled as having been generated with AI.

Finally, to test for differences in story topic (H5), we estimated additional models examining interactions between exposure to labels and story topic. We found minimal differences associated with topic, although admittedly statistical power substantially limits our ability to assess these differences (see Appendix D in the the supplementary appendices).

Exploratory Findings: Attitudes about Labels on AI-Generated Content

To aid our interpretation of the experimental findings in this study, we also included a series of questions at the end of the survey intended to directly capture respondent attitudes about disclosure about the use of generative AI in journalistic content. We specifically included these questions post-treatment after we collected our dependent measures in order to avoid priming respondents to think about the subject of AI.[8] Given emergent interest in this subject, we thought it important to report these results even though they were not strictly speaking part of the experimental design. In this subsection, we present these exploratory descriptive findings before we conclude the article with a discussion of our overall results.

First, not surprisingly given the experimental results, we find that participants in our study tended to perceive generative AI more negatively compared to professional journalists when it comes to writing news articles. A plurality (39.6%) said they believed AI technologies did a "worse job than humans" compared to a third (33.3%) who said they did "about the same job." Just one-in-ten thought they did "a better job than humans" (11.0%) with the remainder saying "don't know" (16.0%). There were no statistical differences on this question between those in different treatment conditions in the main experiment. As these technologies improve, and audiences become more familiar with these tools, it is possible that these perceptions may change, which could also affect

[8] Respondents were not able to revise their responses to the treatment once they reached this section of the survey, ruling out that our post-treatment questions affected them.

perceptions about the trustworthiness of organizations that use these tools to help write news articles. In fact, respondents in our sample who said they had heard or read "a lot" about news organizations using generative AI were more likely to say they thought AI did a better job than humans in writing news articles (16.4% versus 8.9%), a statistically significant difference ($t = 2.54$, $p < 0.05$).

Attitudes about disclosure and labeling are more complex to evaluate as here we did see differences in these attitudes depending on whether respondents were assigned to treatment conditions in which they saw labels disclosing the use of AI. In general, however, overwhelming majorities of respondents in both treatment and control groups said they thought news organizations should "alert readers or viewers that AI was used" (81.3% of those in the control group and 84% of those who saw such a label). The remainder (15.0% of the control group and 11.3% of the treatment group) responded that the statement "I don't need to know how they used AI as long as they stand by their reporting" came closest to their view.

Finally, we asked a follow-up question of those who said they wished to see a label detailing the use of AI about what type of disclosure they would like to see. Among those who said they wanted to see some disclosure, 78% said news organizations "should provide an explanatory note describing how AI was used." Likewise, half (50.0%) said they were in favor of including bylines on stories "attributing the work to AI." Others provided additional suggestions in an open-ended text box. These suggestions varied from the substantive—for example, "a universally accepted symbol" or "industry-wide labels... similar to how you have buttons in the car industry with pictograms" or the "standard way nutrition information is displayed on food products"—to the more general statement of disapproval, such as "or they could just not do this" or they should "pay actual human beings for the work."

Discussion

In this study, using a pre-registered experiment in the U.S. testing effects of disclosure about the use of AI in journalism, we find evidence consistent with our hypothesis that audiences perceive news as less trustworthy when articles are labeled as AI-generated. This is true even though respondents largely did not differ in their evaluations of the accuracy or fairness of the actual AI-generated news articles used as stimuli in the experiment. Furthermore, consistent with our hypotheses, we find evidence of non-linear differences in the way the public perceives such disclosures. While those who are more trusting toward news and more knowledge about journalistic practices (PNK) exhibit the largest treatment effects from exposure to labels about the use of AI, perceptions of trustworthiness largely do not appear to change among the least trusting or least knowledgeable segments of the public. In other words, while disclosure about the use of AI in generating news content does appear to be associated with decreases in trust for some audiences, it also did not appear to increase trust either where it was most lacking.

These results largely held regardless of the political content of the story but they were also largely counteracted where AI-generated news articles included the list of sources used to generate the content, which suggests that reservations about the use of AI may be alleviated where audiences are provided with links to the underlying news articles these tools are drawing from.

To be clear, our study contains several limitations, which we hope will point toward further avenues for studying these issues. A first limitation is our focus solely on the US context; future studies ought to consider audience attitudes across a range of media systems where attitudes about journalism and technology may differ. A second limitation concerns the nature of our Prolific sample, which is more educated, more liberal, and more digitally savvy than the public at large (Guess & Munger, 2023). Given the patterns we observed around heterogeneous effects associated with labeling news as having been generated with the help of AI, it is certainly plausible that the least trusting toward news and the least knowledge about journalistic procedures may well react more positively to such disclosures than we found in our study. Although we did find significant heterogeneous effects, our study may lack the statistical power to satisfactorily characterize subgroup differences for these segments of the public underrepresented in our sample. Likewise, while our study offers a level of realism by employing actual AI-generated news content as stimulus material, it is also predicated on an abstraction: We use a mock news organization rather than an actual brand name in article bylines. Perceptions about a hypothetical news outlet disclosing its use of AI may or may not be generalizable to the more specific case of a known news outlet disclosing its own use of these technologies. Presumably prior attitudes about the organization in question is likely to shape the way individuals think about the implementation of these methods.

Additionally, we see opportunities for further research that delves more deeply into *how* news organizations disclose their use of AI, not only effects of disclosure in general. Audiences may vary in how they respond to labeling depending on the way these technologies are described. While the present study tests a single label adapted from language used by the company whose articles we use in the experiment, future studies ought to consider varying these labels to assess under what conditions audiences may look more or less favorably upon the use of these technologies by news organizations. Similarly, given our findings that the inclusion of a list of original sources alongside labels helps to mitigate reductions in perceived trustworthiness, we hope our study fosters further research that considers how the particular sources listed might also alter these perceptions. For example, when these lists of sources include a mix of partisan brands, which tend to be highly polarizing (Ladd, 2012), it is possible that some audiences may perceive value from the use of AI in journalism and its presumed "mechanical accuracy" as many say they prefer to seek out a range of sources to guard against presumed bias (Nelson & Lewis, 2023).

Our findings hold implications for news organizations and media regulators who are grappling with how to draw lines around the use of these technologies in the news and under what circumstances and in what manner to disclose how such tools are being used (Becker, Simon, Crum, 2023). We find evidence here of what we call the "paradox of AI disclosure," On the one hand, we find an overwhelming consensus among most respondents that they would like to know in what way news outlets are using AI to formulate the content they are seeing, with some even calling for a common industry standard. On the other hand, we find potential reputational costs incurred to news organizations that do disclose their use of these technologies—even as audience evaluations of the accuracy or fairness of such content is unchanged. Rather than being rewarded for transparency, news organizations that disclose their use of these tools are perceived as less trustworthy and may therefore have fewer incentives to be so forthcoming.

These findings matter in the context of calls for journalism to be transparent. Historically, a commitment to (some) transparency over "how the news gets made" belongs to journalisms' authoritative rituals—the routines that allow journalism to distinguish itself from other forms of media work, stress its commitment to the truth, win and retain legitimacy and trust (Karlsson, 2010, p. 536; Tuchman, 1972). As publishing moved to digital formats and content became more ephemeral, calls for—and commitment to—transparency increased in what some have called the normative turn towards transparency (Diakopoulos & Koliska, 2017). With algorithmic systems and as of late AI systems becoming more entrenched in news work, the focus on transparency around their use has taken center stage in scholarly and industry discussions. As these systems can shape the work of journalists in unforeseen ways, can be inscrutable in their automated decision making, or introduce outside logics into news work (Simon, 2022), the worry is that these systems will ultimately limit the autonomy of news organizations and journalists, thus compromising their ability to inform publics in an unbiased manner and further eroding audiencetrust in the news.

From a normative point of view, the commitment to laying bare how the news gets made and which factors—including AI—shape its production is clear. But as our findings demonstrate, this normative commitment could sit uneasily with the reality that disclosing the use of AI systems could undermine what such measures are supposed to strengthen: the publics' trust in the institution that produces the news. For many news organizations (and perhaps especially commercial ones, where gaining and retaining subscribers is an added complication) this will of course raise the question: whither disclosure (and when)? For now, it is not clear how the "paradox of AI disclosure" can be easily resolved. While at this moment, a growing consensus has emerged among many news organizations toward disclosing the use of AI for fear that doing nothing will be worse. If, how, and in which guise these efforts will continue as AI comes to bear more strongly on the news, as well how different audiences may learn and adapt to changing norms around these technologies, remains to be seen.

Acknowledgments

The authors would like to thank the participants of the 2023 International Journal of Press and Politics conference in Edinburgh for their valuable feedback. Additionally, we are grateful to Von Raees and the HeyWire AI team for prompting us to study these questions and supporting the research by providing the stimuli texts, as well as Michelle Disser for general feedback.

Disclosure Statement

The authors have no conflicts of interest to disclose.

Funding

Benjamin Toff would like to thank the Hubbard School of Journalism and Mass Communication at the University of Minnesota for providing the funding to conduct this study. Felix M. Simon would like to thank the OII-Dieter Schwarz Scholarship for supporting his doctoral studies.

References

Amazeen, M. A., & Bucy, E. P. (2019). Conferring resistance to digital disinformation: The inoculating influence of procedural news knowledge. *Journal of Broadcasting & Electronic Media, 63*(3), 415-432. https://doi.org/10.1080/08838151.2019.1653101.

Araujo, T., Helberger, N., Kruikemeier, S., & Vreese, C. H. (2020). In AI we trust? Perceptions about automated decisionmaking by artificial intelligence. *AI & Society*, *35*(6), 611–623. https://doi.org/10.1007/s00146-019-00931-w.

Araujo, T., Brosius, A., Goldberg, A. C., Möller, J., & Vreese, C. de. (2023). Humans vs. AI: The Role of Trust, Political Attitudes, and Individual Characteristics on Perceptions About Automated Decision Making Across Europe. *International Journal of Communication*, *17*(0), 28. https://ijoc.org/index.php/ijoc/article/view/20612

Beckett, C., & Yaseen, M. (2023). Generating Change: A global survey of what news organizations are doing with AI. *JournalismAI, Polis, Department of Media and Communications, The London School of Economics and Political Science.* https://www.journalismai.info/research/2023-generating-change.

Chmielewski, M., & Kucker, S. C. (2020). An MTurk crisis? Shifts in data quality and the impact on study results. *Social Psychological and Personality Science, 11*(4), 464-473. https://doi.org/10.1177/1948550619875149.

Cloudy, J., Banks, J., & Bowman, N. D. (2023). The str(AI)ght scoop: Artificial intelligence cues reduce perceptions of hostile media bias. *Digital Journalism, 11*(9), 1577-1596. https://doi.org/10.1080/21670811.2021.1969974.

Deuze, M. (2005). What is journalism?: Professional identity and ideology of journalists reconsidered. *Journalism*, *6*(4), 442–464. https://doi.org/10.1177/1464884905056815.

Douglas, B. D., Ewell, P. J., & Brauer, M. (2023). Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *Plos one, 18*(3), e0279720.

Epstein, Z., Fang, M. C., Arechar, A. A., & Rand, D. G. (2023, July 28). What label should be applied to content produced by generative AI?. https://doi.org/10.31234/osf.io/v4mfz.

Fletcher, R. (2023, June 14). Attitudes towards algorithms and their impact on news. Oxford: Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2023/attitudes-towards-algorithms-impact-news.

Fletcher, R., & Nielsen, R. K. (2018). Are people incidentally exposed to news on social media? A comparative analysis. *New Media & Society*, *20*(7), 2450–2468. https://doi.org/10.1177/1461444817724170.

Franz, M. (2023, June 2). "Wir setzen ein, was auf dem Markt ist." Ippen Digital. Journalist.de. Retrieved from https://www.journalist.de/startseite/detail/article/wir-setzen-ein-was-auf-dem-markt-is.

Grant, D. G., Behrends, J., & Basl, J. (2023). What we owe to decision-subjects: Beyond transparency and explanation in automated decision-making. *Philosophical Studies*. Advance online publication. https://doi.org/10.1007/s11098-023-02013-6.

Guess, A. M., & Munger, K. (2023). Digital literacy and online political behavior. *Political Science Research and Methods, 11*(1), 110-128. https://doi.org/10.1017/psrm.2022.17.

Hanitzsch, T., Van Dalen, A., & Steindl, N. (2018). Caught in the Nexus: A Comparative and Longitudinal Analysis of Public Trust in the Press. *The International Journal of Press/Politics*, *23*(1), 3-23. https://doi.org/10.1177/1940161217740695.

Jang, W. (Eric), Kwak, D. H., & Bucy, E. (2022). Knowledge of automated journalism moderates evaluations of algorithmically generated news. *New Media & Society*, *0*(0). https://doi.org/10.1177/14614448221142534.

Kalogeropoulos, A., Suiter, J., Udris, L., & Eisenegger, M. (2019). News media trust and news consumption: Factors related to trust in news in 35 countries. *International Journal of Communication, 13*, 3672–3693. https://ijoc.org/index.php/ijoc/article/viewFile/10141/2745.

Ladd, J. M. (2012). *Why Americans hate the media and how it matters.* Princeton University Press.

Longoni, C., Fradkin, A., Cian, L., & Pennycook, G. (2022). News from Generative Artificial Intelligence Is Believed Less. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3531146.3533077.

Maksl, A., Ashley, S., & Craft, S. (2015). Measuring news media literacy. *Journal of Media Literacy Education, 6*(3), 29-45. https://doi.org/10.23860/jmle-6-3-3.

Mitova, E., Blassnig, S., Strikovic, E., Urman, A., Hannak, A., de Vreese, C. H., & Esser, F. (2023). News recommender systems: A programmatic research review. *Annals of the International Communication Association*, *47*(1), 84-113. https://doi.org/10.1080/23808985.2022.2142149.

Mont'Alverne, C., Badrinathan, S., Ross Arguedas, A., Toff, B., Fletcher, R., & Nielsen, R. (2023). "Fair and Balanced": What News Audiences in Four Countries Mean When They Say They Prefer Impartial News. *Journalism Studies*, *24*(9), 1131–1148. https://doi.org/10.1080/1461670X.2023.2201864.

Montal, T., & Reich, Z. (2017). I, Robot. You, Journalist. Who is the Author? *Digital Journalism*, *5*(7), 829-849. https://doi.org/10.1080/21670811.2016.1209083.

Munger, K. (2023). Temporal validity as meta-science. *Research & Politics*, *10*(3). https://doi.org/10.1177/20531680231187271.

Nelson, J. L., & Lewis, S. C. (2023). Only "sheep" trust journalists? How citizens' self-perceptions shape their approach to news. *New Media & Society, 25*(7), 1522-1541. https://doi.org/10.1177/14614448211018160.

Ojala, M. (2021). Is the age of impartial journalism over? The neutrality principle and audience (dis) trust in mainstream news. *Journalism Studies, 22*(15), 2042-2060. https://doi.org/10.1080/1461670X.2021.1942150.

Peer, E., Rothschild, D., Gordon, A., Evernden, Z., & Damer, E. (2022). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, *54*(4), 1643–1662. https://doi.org/10.3758/s13428-021-01694-3.

Simon, F. M. (2022). Uneasy Bedfellows: AI in the News, Platform Companies and the Issue of Journalistic Autonomy. *Digital Journalism*, *10*(10), 1823–1854. https://doi.org/10.1080/21670811.2022.2063150.

Simon, F. M., & Isaza-Ibarra, L. F. (2023). *AI in the news: Reshaping the information ecosystem?* (p. 24). Oxford Internet Institute, University of Oxford. http://dx.doi.org/10.5287/ora-dx865edma.

Schulz, A., Fletcher, R., & Nielsen, R. K. (2022). The role of news media knowledge for how people use social media for news in five countries. *New Media & Society.* https://doi.org/10.1177/14614448221108957.

Strömbäck, J., Tsfati, Y., Boomgaarden, H., Damstra, A., Lindgren, E., Vliegenthart, R., & Lindholm, T. (2020). News media trust and its impact on media use: Toward a framework for future research. *Annals of the International Communication Association, 44*(2), 139-156. https://doi.org/10.1080/23808985.2020.1755338.

Sundar, S. S., & Kim, J. (2019). Machine heuristic: When we trust computers more than humans with our personal information. Proceedings of the 2019 CHI conference on human factors in computing systems, 1-9. https://dl.acm.org/doi/fullHtml/10.1145/3290605.3300768.

Thurman, N. J., Stares, S., & Koliska, M. (2023). Audience Evaluations of News Videos Made with Various Levels of Automation: A Population-Based Survey Experiment. Available at SSRN: https://ssrn.com/abstract=4304961 or http://dx.doi.org/10.2139/ssrn.4304961.

Thurman, N., Moeller, J., Helberger, N., & Trilling, D. (2019). My friends, editors, algorithms, and I: Examining audience attitudes to news selection. *Digital Journalism, 7*(4), 447-469. https://doi.org/10.1080/21670811.2018.1493936.

Tuchman, G. (1972) 'Objectivity as Strategic Ritual: an examination of newsmen's notion of objectivity'. *The American Journal of Sociology, 77*(4), pp. 660-79. https://www.jstor.org/stable/2776752.

Tully, M., Vraga, E. K., & Smithson, A. B. (2020). News media literacy, perceptions of bias, and interpretation of news. *Journalism, 21*(2), 209-226. https://doi.org/10.1177/1464884918805262.

Wölker, A., & Powell, T. E. (2021). Algorithms in the newsroom? News readers' perceived credibility and selection of automated journalism. *Journalism*, *22*(1), 86-103. https://doi.org/10.1177/1464884918757072.

Supplementary Appendices

Appendix A: Stimulus news articles provided by HeyWire AI and used in the experiment.

Barbie Takes Hollywood: How Margot Robbie Reinvented the Iconic Toy

BY INTELLIGENT PRESS
Updated 4:42 PM CDT, September 5, 2023

Once upon a time in Hollywood, Margot Robbie was faced with the daunting task of taking on the lead role of Barbie for the highly anticipated summer film. As she delved into the intricacies of the iconic toy, questions arose regarding beauty and sexiness, which she had to ponder on quite a bit.

"If she doesn't have reproductive organs, would she even feel sexual desire?" Robbie questioned in an interview with Vogue. "No, I don't think she could."

Robbie, who is both the lead actress and co-producer of the upcoming film, had quite a task on her hand. She had to honor the 60-year legacy of the brand while acknowledging that there were many people who weren't fans of Barbie. She kept this in mind as she sought actors who had "Barbie energy" to play different iterations of the character.

The movie boasts of an impressive star-studded cast, including names like Issa Rae, Hari Nef, Emma Mackey, Dua Lipa, Alexandra Shipp, Kate McKinnon, Rhea Perlman, and many more. Ryan Gosling, who plays Barbie's boyfriend Ken, also puts in an appearance, revealing that Robbie helped him find his inner Ken doll.

As the movie approached its release date, the buzz surrounding it continued to grow. One of the earliest choices for the role of the iconic doll, Gal Gadot, was unavailable, but she ended up serving as a template for further casting decisions.

While some questioned the decision to make a movie about Barbie, the Mattel brand had generated over $1 billion in gross sales the previous year, showing its enduring appeal. The CEO of Mattell, Ynon Kreiz, expressed his excitement about Barbie's long-term prospects, acknowledging that the brand had grown and evolved over time.

As the movie debuted, audiences were treated to a comical scene between Robbie and Gosling discussing spending the night together. The teaser also revealed a hyper-real Barbie-inspired world called "Barbieland," where there were many iterations of the iconic dolls in existence. However, their plastic-fantastic lives were turned upside down when they traveled to the real world in a car that spun up into the air. The fish-out-of-water comedy had fun with the quirky premise of what the Mattel dolls would get up to if left to their own devices.

Despite speculation that the movie would fail to capture the spirit of Barbie, the trailers went viral, with fans of the iconic toy eagerly anticipating the movie's release. Dua Lipa also played one of the Barbies, releasing a hit written specifically for the movie.
In the end, Margot Robbie and her production company, LuckyChap, managed to transform the Barbie brand into a movie that showed the iconic toy in a new and refreshing light. And so, in Hollywood, another classic brand had been given a new lease on life.

Sources:

- The Daily Mail (UK)
- Variety
- People
- CNN
- The Independent (UK)

The 15th Annual BRICS Summit: Shaping the Geopolitical Landscape

BY INTELLIGENT PRESS
Updated 4:42 PM CDT, September 5, 2023

The 15th annual BRICS summit, opening on August 22nd in Johannesburg, South Africa, underscored significant geopolitical shifts and interplay among Brazil, Russia, India, China and South Africa. The five nations have come together to collectively reshape the existing global architecture based on the shared values and interests of the global South. They aim to transform it into a more inclusive and equitable multipolar world, notwithstanding the complex events of the past year.

In the run-up to this significant gathering, President Ramaphosa of South Africa has been lauded for his leadership. Through a consultative process with all BRICS partners, he has managed to find a solution that does not diminish the gravity of the summit amidst the global turbulence. He has been instrumental in bringing about a renewed focus on BRICS, particularly regarding the prospective criteria for the expansion of this impactful consortium.

Interestingly, given the increased number of states seeking to be part of the BRICS group, it is now seen as responsible for creating a logical procedure to manage this potential extension. A ground-level approach has been adopted, targeting deliverables that directly impact their citizens' quality of life. This approach could be a game-changer for the complexities associated with global expansion.

A potential query stemming from these developments is the Western world's potential negative reactions, particularly, in view of Britain's exit from the European Union and its stance towards multiple countries. However, BRICS nations remain steadfast, maintaining a single standard in conducting their business and foreign relations with the world.

Shaping the agenda for the summit earlier in June, the foreign ministers of BRICS countries met in Cape Town. Here, dignitaries critiqued institutions like the World Bank and Security Council for sidelining developing countries, termed 'global south.' India's foreign minister, Subrahmanyam Jaishankar, referred to the 'concentration' of global economic power that leaves many nations susceptible to a select few.

In a similar vein, Jair Bolsonaro, Brazil's former president, felt an increased need for BRICS after his ally, Donald Trump, left the White House. In an attempt to understand the bloc's relevance to Russia, an attempted interaction with the country's ambassador to South Africa was made at the meeting, further emphasizing the need for BRICS expansion.

With continued shifts in global relations, BRICS provides an opportunity for countries like Saudi Arabia and United Arab Emirates to reconsider their longstanding alliances. These nations view joining BRICS as a chance to realign their relationships with America, as well as establish closer ties with China.

Despite the complexities, BRICS has managed to set annual summits not just for political bigwigs but also for a vast network of academics, firms, ministers, ruling parties, and think-tanks. These meetings have been successful in establishing unconventional connections with officials, thereby fostering diplomatic dialogue.

As BRICS nations gather, anticipation builds regarding the impact this summit could have on the current global order, the expansion of the bloc, and the future dynamics of international relations. Time will tell whether these geopolitical shifts resurrect the bloc's original intention; to challenge the West's dominance in world affairs.

Sources:

- Reuters
- The Economist
- Bloomberg

Hunter Biden's Lead Defense Attorney Seeks Withdrawal from Controversial Case

BY INTELLIGENT PRESS
Updated 4:42 PM CDT, Sept 5, 2023

The ongoing drama surrounding Hunter Biden, son of U.S. President Joe Biden, saw a major breakthrough on Tuesday when his lead defense attorney, Christopher Clark, pleaded for his withdrawal from the case. Clark sought permission from a federal district court judge in Delaware, and if accepted, this move could herald a new twist in the controversial tax and gun offense case against Hunter Biden.

In his request, Clark specified that due to his close involvement with the now-failed plea agreement and several discussions with prosecutors, he could serve as a witness for numerous contested legal actions. The plea agreement and other discussions are at the heart of the new legal dispute that the Biden defense team plans to zealously challenge.

"There are no 'take backs' once the government signs it and delivers it to court," an anonymous member of Biden's team informed CBS News, highlighting their intent to keep intact the broad immunity provisions contained in a diversion agreement signed earlier. This agreement, they argue, needs to be upheld despite the breakdown of the former plea deal.

According to federal prosecutors in Delaware and Biden's defense team, a preliminary agreement on Hunter Biden's plea and diversion had been reached in July, addressing two misdemeanor tax charges and a felony gun charge. The original deals, however, began to crumble in the final moments due to concerns raised by the judge overseeing the case. The ensuing legal spectacle in an open court last month served to thrust the matter into the spotlight.

In a series of newer court filings, the prosecutors revealed that an impasse had been reached between the two parties. With the Biden defense team alleging that the government had backed down from earlier agreements, the likelihood of a trial has increased. The bone of contention lies in the signed and binding diversion agreement on the gun charge, which the government rejected on Tuesday.

Adding intrigue to this legal tussle, the relationship between Hunter Biden and his attorneys saw a reshuffling on Monday. Abbe Lowell, a seasoned Washington D.C. lawyer, formally entered the case's fray. He is known for his representation of top-flight clients, including Ivanka Trump and Jared Kushner, and has been guiding Hunter Biden through the ongoing congressional probe into his business dealings.

Additionally, a dramatic turn of events could see the case —originally involving tax crimes and a potential diversion agreement regarding a gun charge— go to trial. Despite potential wide-reaching consequences, Hunter Biden's lawyers remain steadfast in their claim that the prior diversion agreement —which they and prosecutor Weiss signed off on— has legal compulsion. Biden's attorneys recently argued that the pretrial diversion agreement for the gun charge, which they and Weiss both signed off on, is "valid and binding." Prosecutors are expected to respond to this argument soon.

As the tense legal negotiations between Hunter Biden and the Justice Department escalate into an open fight, the result could see Christopher J. Clark transform from defendant to firsthand witness, further complicating this long-standing legal controversy. The potential departure of Clark, coupled with the entry of Lowell into the defense team, suggests that this high-profile case is far from over.

Sources:

- Washington Examiner
- CNBC
- CNN
- CBS News
- New York Times

Appendix B: Items Used in Measuring Procedural News Knowledge (PNK)

PNK_Preface We would like to know more about how familiar you are with aspects of journalism. Please answer the following questions to the best of your ability.

PNK1 In what section does a newspaper's editorial staff endorse candidates and express their opinions about current issues?
- On the front page
- On the editorial page [CORRECT]
- In the business section
- In a special weekly advertising section
- Mainstream newspapers don't endorse candidates or take issue stands

PNK2 Which of the following best describes a press release?
- A short news piece written or produced by a reporter
- A written statement or short video about a newsworthy event given out to reporters by an official or public relations specialist [CORRECT]
- An opinion piece written by a syndicated columnist
- A paid advertisement that appears in newspapers and on news websites with the label "paid advertisement"

PNK3 Which of the following U.S. news outlets does NOT depend primarily on advertising for financial support?
- FOX News
- PBS [CORRECT]
- The New York Times
- Time Magazine
- Don't know

PNK4 When it comes to reporting the news, the main difference between websites like Google News and a website like CNN.com is that:
- Google doesn't have reporters who gather information, while CNN does. [CORRECT]
- Google focuses on national news, while CNN focuses on local news.
- Google has more editors than CNN does.
- Google charges money for the news, while CNN does not.
- Don't know

Appendix C: Output from Linear Regression Models

Table C-1. Linear regression models predicting attitudes about stimulus articles as a function of trust in news and exposure to AI labels.

	Perceived Trustworthiness (H2a)		Article Accuracy (H2b)		Article Fairness (H2c)	
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	6.03***	5.87***	2.07***	1.98***	0.39***	0.45***
	(0.08)	(0.29)	(0.02)	(0.08)	(0.03)	(0.11)
Treatment (AI label)	-0.54***	-0.51***	-0.04	-0.04	-0.04	-0.04
	(0.12)	(0.12)	(0.03)	(0.03)	(0.04)	(0.04)
Trust in news	0.68***	0.65***	0.10***	0.09***	0.16***	0.16***
	(0.07)	(0.08)	(0.02)	(0.02)	(0.02)	(0.03)
Trust X Treatment	-0.18$^{\#}$	-0.21*	-0.03	-0.04	-0.01	-0.01
	(0.10)	(0.10)	(0.03)	(0.03)	(0.04)	(0.04)
Demographic Controls						
Age	—	0.06	—	-0.01	—	-0.04*
		(0.05)		(0.01)		(0.01)
Gender (Male)	—	0.10	—	0.07$^{\#}$	—	0.01
		(0.12)		(0.04)		(0.05)
Race (Black)	—	-0.28	—	-0.11*	—	-0.07
		(0.12)		(0.06)		(0.07)
Ethnicity (Hispanic)	—	-0.27	—	-0.20***	—	-0.05
		(0.20)		(0.06)		(0.08)
Education	—	-0.05	—	0.01	—	-0.01
		(0.05)		(0.01)		(0.02)
Political Controls						
Political interest	—	-0.08	—	-0.02	—	0.01
		(0.06)		(0.02)		(0.02)
Party ID	—	-0.07		-0.05*		0.03
		(0.08)		(0.02)		(0.03)
News Consumption						
News use (Online)	—	-0.03	—	0.00	—	-0.03
		(0.07)		(0.01)		(0.02)
News use (Soc Med)	—	0.05		-0.01		0.03
		(0.06)		(0.02)		(0.02)
News use (Print)	—	-0.03	—	-0.00	—	-0.02

		(0.08)		(0.02)		(0.03)
News use (Audio)	—	-0.01	—	-0.00	—	0.00
		(0.08)		(0.02)		(0.03)
News use (TV)	—	0.03	—	0.03	—	0.06*
		(0.08)		(0.02)		(0.03)
News use (Cable TV)	—	0.17*	—	-0.02	—	-0.03
		(0.08)		(0.02)		(0.03)
News use (Fox News)	—	0.18*	—	0.02	—	0.03
		(0.08)		(0.02)		(0.03)
N (respondents)	1483	1483	1483	1483	1483	1483
R^2	0.21	0.11	0.02	0.21	0.04	0.05
Adj. R^2	0.21	0.10	0.02	0.20	0.04	0.04

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; $^{\#}$ $p < 0.10$.

Table C-2. Linear regression models predicting attitudes about stimulus articles as a function of PNK and exposure to AI labels.

	Perceived Trustworthiness (H3a)		Article Accuracy (H3b)		Article Fairness (H3c)	
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	5.64***	5.87***	1.83***	1.85***	0.36***	0.40***
	(0.26)	(0.29)	(0.07)	(0.10)	(0.10)	(0.13)
Treatment (AI label)	0.22	0.26	0.03	0.02	-0.01	-0.00
	(0.38)	(0.37)	(0.11)	(0.11)	(0.14)	(0.14)
PNK	0.35	0.38	0.29**	0.26**	0.01	0.11
	(0.32)	(0.32)	(0.09)	(0.09)	(0.12)	(0.13)
PNK X Treatment	-0.87#	-0.95*	-0.07	-0.08	-0.02	-0.06
	(0.47)	(0.44)	(0.13)	(0.13)	(0.17)	(0.17)
Demographic Controls						
Age	—	0.06	—	-0.01	—	-0.04*
		(0.05)		(0.01)		(0.02)
Gender (Male)	—	0.10	—	0.06	—	0.00
		(0.12)		(0.05)		(0.05)
Race (Black)	—	-0.27	—	-0.09*	—	-0.06
		(0.19)		(0.05)		(0.07)
Ethnicity (Hispanic)	—	-0.29	—	-0.19**	—	-0.05
		(0.20)		(0.06)		(0.08)
Education	—	-0.05	—	0.00	—	-0.01
		(0.05)		(0.01)		(0.02)
Political Controls						
Political interest	—	-0.06	—	0.01	—	0.01
		(0.06)		(0.02)		(0.02)
Party ID	—	-0.06	—	-0.04#	—	0.03
		(0.09)		(0.03)		(0.03)
News Consumption						
News use (Online)	—	-0.03	—	0.00	—	-0.03
		(0.06)		(0.02)		(0.02)
News use (Soc Med)	—	0.05	—	-0.01	—	0.03
		(0.06)		(0.02)		(0.02)
News use (Print)	—	-0.06	—	-0.02	—	-0.02
		(0.07)		(0.02)		(0.03)
News use (Audio)	—	-0.01	—	-0.00	—	0.00

		(0.08)		(0.02)		(0.03)
News use (TV)	—	0.02	—	0.03	—	0.06^{*}
		(0.08)		(0.02)		(0.03)
News use (Cable TV)	—	0.17^{*}	—	-0.03	—	-0.03
		(0.08)		(0.02)		(0.03)
News use (Fox News)	—	0.17^{*}	—	0.03	—	0.03
		(0.08)		(0.02)		(0.03)
Trust in news	—	0.56^{***}	—	0.07^{***}	—	0.16^{***}
		(0.06)		(0.02)		(0.02)
N (respondents)	1483	1483	1483	1483	1483	1483
R^2	0.01	0.11	0.01	0.05	0.00	0.05
Adj. R^2	0.01	0.10	0.01	0.03	0.00	0.04

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; $^{\#}$ $p < 0.10$.

Table C-3. Linear regression models predicting attitudes about stimulus articles as a function of both treatment conditions interacted with each other (H4b).

	Perceived Trustworthiness		Article Accuracy		Article Fairness	
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	5.90***	5.86***	2.05***	1.87***	0.36***	0.40***
	(0.12)	(0.32)	(0.03)	(0.09)	(0.10)	(0.12)
Treatment (AI label)	-0.60***	-0.63***	-0.09#	-0.09#	-0.06	-0.09
	(0.17)	(0.16)	(0.05)	(0.05)	(0.06)	(0.06)
Treatment (Source list)	0.02	0.06	0.00	0.02	0.01	0.03
	(0.17)	(0.16)	(0.05)	(0.05)	(0.06)	(0.06)
Labels X Source	0.31	0.29	0.12#	0.10	0.09	0.09
	(0.24)	(0.23)	(0.07)	(0.07)	(0.09)	(0.09)
Demographic Controls						
Age	—	0.07	—	-0.01	—	-0.04*
		(0.05)		(0.01)		(0.02)
Gender (Male)	—	0.12	—	0.06	—	0.01
		(0.12)		(0.04)		(0.05)
Race (Black)	—	-0.29	—	-0.09*	—	-0.07
		(0.19)		(0.06)		(0.07)
Ethnicity (Hispanic)	—	-0.24	—	-0.18***	—	-0.04
		(0.21)		(0.06)		(0.08)
Education	—	-0.05	—	0.00	—	-0.01
		(0.05)		(0.01)		(0.02)
Political Controls						
Political interest	—	-0.07	—	0.01	—	0.01
		(0.06)		(0.02)		(0.02)
Party ID	—	-0.07	—	-0.05#	—	0.03
		(0.09)		(0.03)		(0.03)
News Consumption						
News use (Online)	—	-0.03	—	0.00	—	-0.03
		(0.07)		(0.02)		(0.02)
News use (Soc Med)	—	0.05	—	-0.00	—	0.03
		(0.06)		(0.02)		(0.02)
News use (Print)	—	-0.05	—	-0.02	—	-0.02
		(0.07)		(0.02)		(0.03)
News use (Audio)	—	-0.02	—	-0.00	—	0.00

		(0.08)		(0.02)		(0.03)
News use (TV)	—	0.03	—	0.03	—	0.06*
		(0.08)		(0.02)		(0.03)
News use (Cable TV)	—	0.15$^{\#}$	—	-0.03	—	-0.04
		(0.08)		(0.02)		(0.03)
News use (Fox News)	—	0.18*	—	0.03	—	0.03
		(0.08)		(0.02)		(0.03)
Trust in news	—	0.56***	—	0.08***	—	0.16***
		(0.06)		(0.02)		(0.02)
PNK	—	-0.06	—	0.23***	—	0.08
		(0.25)		(0.07)		(0.10)
N (respondents)	1483	1483	1483	1483	1483	1483
R^2	0.01	0.11	0.00	0.05	0.00	0.05
Adj. R^2	0.01	0.10	0.00	0.03	0.00	0.04

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; $^{\#}$ $p < 0.10$.

Appendix D: Tests for heterogeneous treatment effects by story topic

We estimated additional models (output available on request) testing for three-way-interactions when including story topic as an additional variable in the models. Below we visualize the results of these analyses, holding all other variables at their mean values.

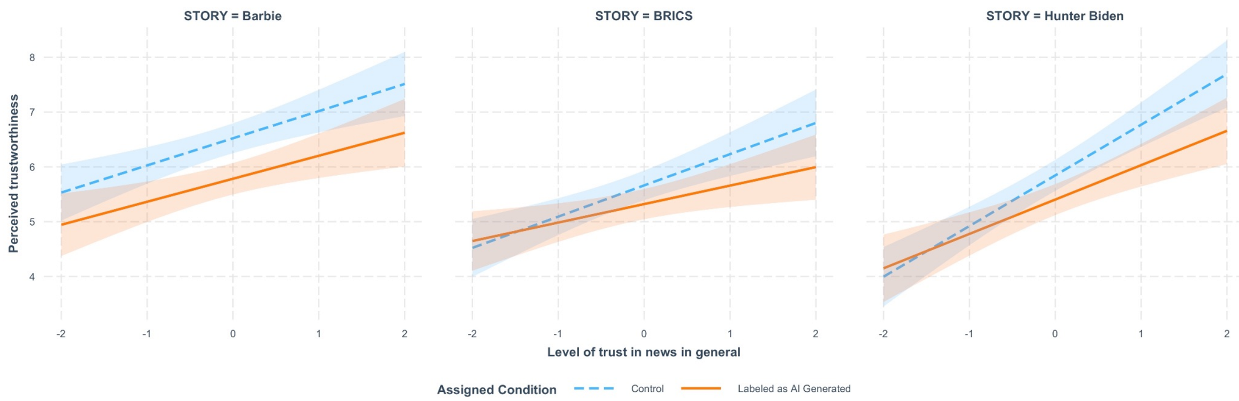

Figure D-1. Differences in heterogeneous treatment effects by prior levels of trust in news are plotted for each of the three story topics. The largest differences gaps in perceived trustworthiness are apparent for the most politically contentious of the three stories (Hunter Biden) although the largest treatment effects in general are apparent for the least politically contentious story (Barbie). Our sample size makes assessing the significance of these differences limited.

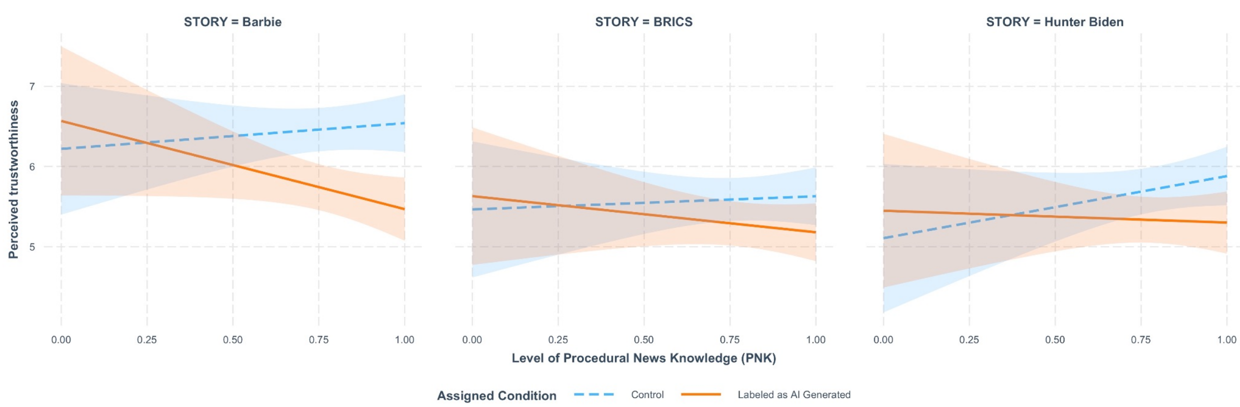

Figure D-2. Differences in heterogeneous treatment effects by PNK are plotted for each of the three story topics. Here we see the largest treatment effects associated with the least politically contentious of the three stories (Barbie). Our sample size makes assessing the significance of these differences limited.