News from Generative Artificial Intelligence Is Believed Less

Chiara Longoni*
clongoni@bu.edu
Assistant Professor of Marketing
Questrom School of Business
Boston University
Boston, MA, USA

Andrey Fradkin
fradkin@bu.edu
Assistant Professor of Marketing
Questrom School of Business
Boston University
Boston, MA, USA

Luca Cian
cianl@darden.virginia.edu
Killgallon Ohio Art Associate Professor of Business Administration
Darden School of Business
University of Virginia
Charlottesville, VA, USA

Gordon Pennycook
grpennycook@gmail.com
Associate Professor
Hill/Levene Schools of Business
University of Regina
Regina, Saskatchewan, Canada

ABSTRACT

Artificial Intelligence (AI) can generate text virtually indistinguishable from text written by humans. A key question, then, is whether people believe news headlines generated by AI as much as news headlines generated by humans. AI is viewed as lacking human motives and emotions, suggesting that people might view news written by AI as more accurate. By contrast, two pre-registered experiments on representative U.S. samples (N = 4,034) showed that people rated news headlines written by AI as less accurate than those written by humans. People were more likely to incorrectly rate news headlines written by AI (vs. a human) as inaccurate when they were actually true, and more likely to correctly rate them as inaccurate when they were indeed false. Our findings are important given the increasing adoption of AI in news generation, and the associated ethical and governance pressures to disclose it use and address standards of transparency and accountability.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Computing methodologies** → Cognitive science; • **General and reference** → **Empirical studies**; *Experimentation*; • **Applied computing** → *Psychology*.

KEYWORDS

generative artificial intelligence, algorithmic transparency, fairness, news, news generation

*Corresponding author

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06…$15.00
https://doi.org/10.1145/3531146.3533077

ACM Reference Format:
Chiara Longoni, Andrey Fradkin, Luca Cian, and Gordon Pennycook. 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

1 INTRODUCTION

One of the applications of Artificial Intelligence (AI) that has shown the most promising advances in the last decade is that of generative AI: AI algorithms capable of producing textual, visual, and auditory content with little no human intervention. To illustrate, in 2020 the Generative Pre-trained Transformer 3 (GPT-3) has been lauded as the most advanced neural network capable of producing text—a fictional story, a poem, an answer to a math problem, or even a programming code—virtually indistinguishable from text written by a person [5, 26, 29]. Applications of generative AI are more pervasive than one might think. For instance, leading media companies such as the *Associated Press*, *Forbes*, the *New York Times*, the *Washington Post*, and *ProPublica*, use AI to generate entire articles from scratch and automatically report on crimes, financial markets, politics, sporting events and foreign affairs [14, 35]. Generative AI is also increasingly used as an input in the writing process across several domains—from user- and company-generated content, to institutional communications from public organizations and governments.

When AI is used to generate content, its role is typically not disclosed. In news articles, for example, the byline is rarely attributed to an AI algorithm even when an algorithm was used. Without this disclosure, readers cannot determine whether an AI was used from the text alone [5, 26, 29]. However, given the potential misuse or unintended consequences of this new technology [1, 28], ethicists and policymakers have argued that the use of AI should be disclosed [14, 37]. Indeed, it's possible that such disclosure will be mandated by law [48] as advocated, for instance, in the Algorithmic Justice and Online Platform Transparency Act of 2021.

How will people perceive news generated by AI once it's labeled as such? At the moment, we do not know the answer to this question. Existing research in the area of generative AI has either focused on the technical aspects of text generation (e.g., [5, 26, 29]) or on the risks and benefits of AI to publishers (e.g., [35]), thereby neglecting to consider how people will perceive news from AI [22]. News perceptions play a critical role for the civil society, as trustworthy news reporting can provide a check on misconduct and corruption [16] and influence important societal outcomes, from elections to finance, public policy, and political economy[20, 30]. It is therefore both timely and important to understand how people will perceive news generated by AI, which is the focus of the present research.

2 THEORETICAL DEVELOPMENT

We focus on a specific dimension of news perception: that of accuracy — judgments of the veracity of a news item. Accuracy is an important dimension of news perception because the extent to which a news item is initially accepted as true determines the extent to which that item is processed and later remembered [15, 25], and the extent to which it will influence subsequent judgments [27] even in the face of retractions and corrections [31]. As such, accuracy perceptions have received considerable attention in the literature [41].

There are two opposing theoretical perspectives related to the perceived accuracy of AI-generated news. The first predicts that news from AI would be perceived as more accurate than news from human reporters. AI is typically viewed as lacking human desires, motives, and emotions [18, 19]. These qualities might be incorporated into how people judge content generated by AI—as being relayed impartially and dispassionately or, in other words, as more truthful. Furthermore, people appreciate algorithms more than humans for tasks that are impersonal [3], require objectivity [8] or impartiality [24]. As people want journalism to be impartial and neutral [42], this *AI appreciation* account predicts that people would perceive news from AI as more accurate than news from human reporters, with higher trust ascribed to AI than human reporters.

The alternative account is grounded in people's resistance toward replacement of humans by automated systems [11, 21, 36]. This line of research has shown that people are averse to AI supplanting humans, on the grounds that AI is unable to adapt to mutable, unpredictable, or unique contexts [13, 32, 33] and that it lacks empathy and experiential abilities [8, 34]. Moreover, news production might be considered a morally laden task, and people view it inappropriate for AI to make moral decisions [4]. Thus, the competing *AI aversion* account predicts that people would perceive news from AI as less accurate than news from human reporters, with lower trust ascribed to AI than human reporters.

3 EXPERIMENTS 1 AND 2

3.1 Method

We tested these competing predictions in two large experiments on samples recruited on the platform Lucid, which uses quota sampling to provide respondents representative of the U.S. population on age, gender, ethnicity and geographic region. Experiment 1 employed a between-subject/separate evaluation paradigm: participants saw *either* news items tagged as written by an AI or by a human. Experiment 2 employed a within-subject/joint evaluation paradigm: participants saw *both* news items tagged as written by AI and by a human. We conducted the within-subject experiment to refine the measurement of accuracy, as the evaluability of the writer should be more salient in a joint paradigm [23]. Our main dependent variable was perception of news accuracy, and our secondary dependent variable was trust toward the reporter.

Datasets, preregistrations of sample sizes, dependent variables, primary analyses, and all experimental materials are available on the Open Science Framework at https://bit.ly/2YVyh8z. These experiments were conducted under the approval by the Boston University Institutional Review Board, Protocol No. 4408E.

3.2 Participants

The preregistered samples were 3,000 participants for Experiment 1 and 1,000 for Experiment 2, recruited on the platform Lucid. In total, we recruited 3,029 participants (1,469 female, 1,560 male) in Experiment 1, and 1,005 (490 female, 515 male) in Experiment 2.

3.3 Design

3.3.1 Experiment 1. In a 2-cell, between-subject design and across three experimental waves, participants were randomly assigned to a condition in which they saw news items tagged as written by an AI reporter or a condition in which they saw news items tagged as written by a human reporter.

3.3.2 Experiment 2. In a 2-cell, within-subject design, participants saw both news items tagged as written by an AI and by a human reporter.

3.4 Materials

All news items comprised a text headline and an accompanying photo. We used real news headlines and real photos that appeared in news outlets at the time of the experiment. We focused on headlines rather than full articles because news consumption largely occurs at the level of headlines, as is often the case on social media [17]. We omitted information about the outlet given research indicating that publisher information has no effect on accuracy perceptions [12]. To assess whether the effect of AI disclosure was moderated by the news' actual veracity (whether the news was objectively accurate or inaccurate), we predetermined whether each news item was true or false by relying on the fact-checking site Snopes.com. As it was important to compare subjective ratings of accuracy (a continuous measure of accuracy and our dependent variable) with an objective measure of accuracy (a binary measure of veracity), we did not consider news that Snopes.com had rated as "mostly false," "mostly true," "mixture," "unproven," "miscaptioned," or "misattributed." We report in the appendix in Table A.1 and Table A.2 the list of news headlines by experiment, experimental wave, and date of fact-checking.

In Experiment 1, participants saw a total of 30 news items in wave one (15 true news; 15 false news); a total of 36 news items in wave two (15 true news from wave one, plus 3 novel true news; 15 false news from wave one, plus 3 novel false news); and a total of 42 news items in wave three (15 true news from wave one, plus 3 true news from wave two, plus 3 novel true news; 15 false news from wave one, plus 3 false news from wave two, plus 3 novel false news). We leveraged the three waves to add novel news to account for the potential effect of the "age" of the news—the amount of time a certain news had been accessible to the public—as novelty affects reactivity to news [46]. In Experiment 2, participants saw a total of 20 news (10 true; 10 false). Stimuli are available at https://bit.ly/2YVyh8z.

3.5 Procedure

Participants read that they would be presented with news headlines that had appeared on various outlets and social media platforms, and that these headlines may had been written by human reporters or by artificial intelligence (AI) reporters. We described AI reporters as "algorithmic processes that convert data into narrative news texts with limited to no human intervention beyond the initial programming choices" [7]. To rule out the potential role of believing that AI reporters had differential access to sources or were differentially skilled at data mining, we clarified that "human or AI reporters wrote these news headlines based on information available to them and to the news outlets employing them at the time in which the news headline was written."

Participants were then randomly assigned to one of two conditions as per the respective experimental design (between-subject in Experiment 1, and within-subject in Experiment 2). Then, participants viewed the news items one at a time and in random order. Our pre-registered and primary dependent variable was news accuracy, measured after each news item as follows: "To the best of your knowledge, how accurate is this news headline?" 1 - *Not at all accurate*, 4 - *Very accurate* [38]. Our pre-registered and secondary dependent variable was trust toward the reporter, measured at the end of the survey with two items: "To what extent do you trust [a human reporter / an AI reporter] to write accurate news headlines?" (1 - *Do not trust at all*, 5 - *Trust completely*), "To what extent do you think [a human reporter / an AI reporter] is capable of writing accurate news headlines?" (1 - *Not at all*, 5 - *Very much*).

As pre-registered, we also collected the following measures as prior research indicated their relevance for news discernment: risk appraisals, information seeking behavior, political orientation, religiosity, and demographic variables. Although included in the pre-registration, we made no a-priori hypotheses with respect to the potential interactions of these measures with the independent variable, and therefore acknowledge the exploratory nature of the respective analyses.

Based on research on the relationship between accuracy judgments and threats [40], we measured risk appraisals by use of severity ratings ("In your opinion, how severe is Coronavirus (COVID-19)?" 1 - *Not at all severe*, 5 - *Extremely severe*), comparative likelihood estimates ("How likely do you think it is that you will get infected by the Coronavirus (COVID-19) in the next year?" and "How likely do you think it is that an average American person will get infected by the Coronavirus (COVID-19) in the next year?" 0% = *Impossible*, 100% = *Certain*), and negative affect ("How concerned are you about COVID-19?" 1 - *Not at all*, 5 - *Extremely*). Based on research on news accuracy and COVID-19 [26] we assessed information seeking behavior ("How often do you check the news about COVID-19?" 1 Never, 5 - Very often) and recency of acquisition of information related to COVID-19 ("Have you checked the news related to COVID-19 in the last five hours?" *Yes*, *No*). We also measured political orientation ("Which of these best describes your political view?" *Democrat*, *Republican*, *Independent*, *Other*) and religiosity ("Would you say you are:..." *A religious person*, *Not a religious person*, *An atheist*; "How strongly do you believe in the existence of a God or Gods?" 1 - *Very little*, 5 - *Very much*) because prior relevant work on has indicated their relevance for news discernment [2]. Finally, we collected demographic information (age, gender, ethnicity, geographic region, marital status, education, employment).

At the end of the survey, and in line with prior research [38], participants were asked (*i*) if they had searched for any of the headlines while responding to the survey, and (*ii*) if they had responded randomly at any point. A manipulation check in Experiment 1 assessed whether participants correctly recalled that the headlines they had viewed had been generated by a human or an AI reporter. Upon completion of the survey, participants could follow a link to know which of the news items were true and which were false.

3.6 Results

As pre-registered, we performed the main analyses at the level of individual accuracy perceptions ratings (i.e., one data point per news item per participant) using a linear regression with robust standard errors clustered on participant.

3.6.1 Experiment 1. Perceived Accuracy. A linear regression with robust standard errors clustered on participant on the individual accuracy perceptions ratings (i.e., one data point per news item per participant) revealed that tagging news items as written by an AI was associated with an average reduction in perceived accuracy of 7.6pp (*SE* 1.5pp, $p < .001$) compared to the control condition, in which news items were tagged as written by a human (Column 1; Table 1). Regression specifications with news item fixed effects, which allowed for baseline accuracy rating to vary with the news (Column 2), and separately for true (Column 3) and false news (Column 4), yielded the same conclusions. The effect was directionally larger for false news than true news, but the two effects were not statistically different ($p = .359$). The negative effect of AI disclosure emerged for 41 out of the 42 news items tested and ranged from -20pp to 3pp across news. These results are based on the entire dataset: we did not remove responses by those who (i) reported searching on Google (15% of the sample), (ii) reported responding randomly 22% of the sample), or (iii) failed the manipulation check (i.e., if they incorrectly recalled whether the reporter was AI or human; 18% of the sample). Statistical conclusions do not change if we restrict analysis to those who did not search on Google, did not respond randomly, or passed the manipulation check. Thus, these results, shown in Table 1 and Figure 1, strongly supported aversion toward algorithmic reporting.

Trust. A t-test on the average of the two items measuring trust ($r = .635$, $p < .001$) revealed that participants trusted AI reporters less than human reporters ($M_{AI} = 2.57$, $SD = 0.90$; $M_{Human} = 3.30$, $SD = 0.97$; $t(3016.7) = 21.51$, $p < .001$), again supporting aversion toward algorithmic reporting.

Mediation. To further probe aversion toward algorithmic reporting, we explored whether trust mediated the effect of reporter on perceived accuracy. To do so, we averaged all the responses for each participant, and then conducted a mediation using linear models for the mediation and outcome equations (1,000 simulations;

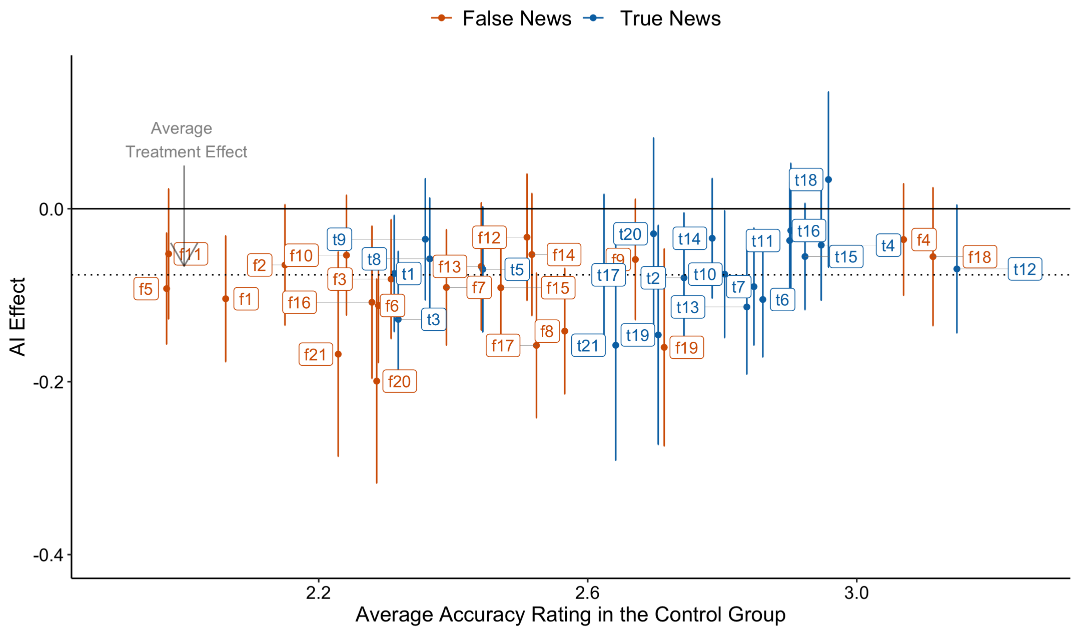

Figure 1: Results of Experiment 1: Negative effect of AI disclosure on perceptions of news accuracy by news item

Table 1: Results of Experiment 1: Negative effect of AI disclosure on perceptions of news accuracy by regression specification

	Perceptions of News Accuracy			
	(1)	(2)	(3)	(4)
AI reporter condition	-0.076*** (0.015)	-0.076*** (0.015)	-0.068*** (0.018)	-0.085*** (0.017)
M	2.56	2.56	2.72	2.41
SD	1.04	1.04	1.03	1.02
Sample	All	All	True News	False News
Item FE	No	Yes	Yes	Yes
Observations	109,068	109,068	54,534	54,534
Adjusted R^2	0.001	0.093	0.059	0.085

Mediation Package for R; [45]). Note that this analysis requires the assumption of sequential ignorability, which is quite stringent and implies that there are no unmodeled variables that affect both the trust and accuracy. The effect of reporter (human vs. AI) on trust was significant (β = -.73, p < .001, CI: [-.80, -0.67]), and trust (specifically, α*reporter + β*trust + ϵ) had a significant effect on perceived accuracy (β = .169, p = < .001, CI: [.155, .183]). The indirect effect of reporter on perceived accuracy through trust was significant (β = -.12, p = < .001, CI: [-.14, -.11]). These results further corroborated aversion toward AI reporting.

Treatment Effect Heterogeneity. To analyze treatment effect heterogeneity based on the additional variables we collected, we used the causal forest approach of Wager and Athey [47], which relies on random forests to estimate the conditional treatment effect (CATE; i.e., the treatment effect conditional on observed covariates). We employed this machine learning method because we did not have strong a-priori hypotheses on which variable would predict treatment effects and in which direction, and this method does not overfit the data while allowing to include many covariates in the estimation of the treatment effect [47]. We used this method as follows. Our outcome metric was the average accuracy rating across all news items per participant. We used 50% of our sample to train the forest and the other 50% to estimate the heterogeneous treatment effects. We included as covariates the following variables: risk appraisals (risk for self, risk for others, severity, concern), information seeking behavior, recency of acquisition of information related to COVID-19, political orientation, religiosity, demographics (age, gender, ethnicity, geographic region, marital status, education, employment), and experimental wave. We treated categorical variables as dummy variables (one-hot encoding). We added additional columns to denote missing variable status for the cases when risk appraisals were not observed; the corresponding missing column was filled in as a 0. For the estimation, we used 10,000 trees and a minimum node size of 5. Although there was variation in these CATEs, the confidence intervals were wide enough to include the average treatment effect for all observations, which means that we were not able to precisely estimate heterogeneous treatment effects based on the variables examined.

We also conducted the test of Chernozhukov et al. [9] to see whether the estimate of the CATE was predictive of the true treatment effects. The coefficient on 'mean calibration' was close to 1 and precisely estimated, meaning that the mean prediction of the causal forest was correct. However, the coefficient on Heterogeneous Treatment Effect was negative and not precisely estimated. For the CATE estimates to pick up meaningful heterogeneity, the coefficient should instead have been positive and statistically different from 0. This provides further evidence that the variables examined did not precisely predict heterogeneity in treatment effects.

Table 2: Results of Experiment 2: Negative effect of AI disclosure on perceptions of news accuracy by regression specification

	Perceptions of News Accuracy			
	(1)	(2)	(3)	(4)
AI reporter condition	-0.145*** (0.015)	-0.142*** (0.015)	-0.140*** (0.020)	-0.143*** (0.019)
M	2.62	2.62	2.71	2.52
SD	1.01	1.01	1.01	1.00
Sample	All	All	True News	False News
Item FE	No	Yes	Yes	Yes
Observations	20,120	20,120	10,060	10,060
Adjusted R^2	0.005	0.093	0.059	0.085

3.6.2 Experiment 2. Perceived Accuracy. A linear regression with robust standard errors clustered on participant on the individual accuracy perceptions ratings (i.e., one data point per news item per participant) revealed that tagging news items as written by AI was associated with an average reduction in perceived accuracy of 14.5pp (*SE* 1.5pp, $p < .001$) compared to the control (Column 1, Table 2). A regression specification with news item fixed effects, which allowed for the baseline response to vary with the news item (Column 2), and separately for true (Column 3) and false news (Column 4) yielded the same results. The negative effect of AI disclosure emerged for all news items, ranging from -24 to -7. The results of the analysis are based on the entire dataset: we did not remove responses by those who (i) searched on Google (17% of sample) or (ii) responded randomly (18% of sample). Statistical conclusions do not change if we restrict the analyses to those who did not search on Google or responded randomly. Overall, Experiment 2 replicated Experiment 1 and pointed to an even larger effect, again supporting aversion towards algorithmic reporting. The results are shown in Table 2 and Figure 2.

Trust. A t-test on the average of the two items measuring trust ($r = .607$, $p < .001$) revealed that participants trusted AI reporters less than human reporters ($M_{AI} = 2.75$, $SD = .89$; $M_{Human} = 3.39$, $SD = 1.00$; $t(1975.6) = 15.19$, $p < .001$), again supporting aversion toward algorithmic reporting.

Mediation. To account for the fact that participants were treated for some news items and not others (i.e., viewed some news items tagged as written by an AI reporter but not others), we used a multi-level mediation model to analyze Experiment 2. Specifically, we modeled both the mediation and outcome equations as mixed models with participant random effects (1,000 simulations; Mediation Package for R; [45]). The effect of reporter (human vs. AI) on trust was significant ($\beta = -.64$, $p < .001$, CI: [-0.66, -0.63]), and trust had a significant effect on perceived accuracy ($\beta = .14$, $p < .001$, CI: [.12, .16]). The indirect effect of the reporter on perceived accuracy through trust was significant ($\beta = -.09$, $p < .001$, CI: [-.10, -.08]). These results replicated those of Experiment 1.

Treatment Effect Heterogeneity. As Experiment 2 used a joint evaluation paradigm and had fewer observations then Experiment 1, we conducted a different test of heterogeneity than the one we used in Experiment 1. Specifically, we computed an individual treatment effect by computing the difference in average responses between treated items (AI) and control items (human) for each participant. We then regressed the individual treatment effect estimate on the same covariates as we did in Experiment 1: risk appraisals (risk for self, risk for others, severity, concern), information seeking behavior, recency of acquisition of information related to Covid-19, political orientation, religiosity, and demographic variables (gender, age, income, religion, region, ethnicity, marital status, education). This regression had an F-statistic of 1.1, indicating that we could not reject the null hypothesis that all coefficients in the regression are 0. We also performed a 10-folds cross-validated penalized generalized linear model (cv.glmnet in R) procedure using the individual treatment effect estimate as the outcome and the same covariates. The selected glmnet model placed zero weight on all the explanatory variables. Together, these results indicate that none of the covariates had a strong enough correlation with the treatment effect for us to be able to precisely estimate conditional average treatment effects given our sample size.

4 GENERAL DISCUSSION

Across two large pre-registered experiments on representative U.S. samples, we examined how disclosing use of AI in news generation affected news accuracy perceptions. The results strongly corroborated the AI aversion account: disclosing the use of AI led people to believe news items substantially less, a negative effect explained by lower trust toward AI reporters. The effect was robust to experimental paradigm (separate and joint evaluations), respondents' characteristics (risk appraisals, information seeking behavior, political orientation, religiosity, and demographics), and emerged for every single news item that we included (with one exception in Experiment 1), indicating robustness across actual news veracity and age/novelty.

This research makes two theoretical contributions. First, it builds on and extends the literature on psychological responses to automated systems and on human-machine interaction [11, 21, 36], indicating new – and perhaps surprising - depths to this AI aversion given that past scholarship has shown cases of both AI appreciation [3, 8, 34] and aversion [6, 13, 32–34]. Though this literature spans over decades, it has mostly focused on outcome variables such as stated or revealed preference. Our research is the first to empirically examine the effect of disclosing use of generative AI on accuracy perceptions in a context of everyday importance. Even though the effect we document falls in the domain of news accuracy, the robustness of the negative effect of AI implies that it is likely to generalize to the many other domains in which AI systems are increasingly used to generate text—from social media posts, podcast notes, to institutional communications from corporations and governments.

We make a second theoretical and applied contribution to research on perceptions of news accuracy [15, 25] by documenting a negative effect of AI disclosure on news perceptions. This effect is novel, as research on generative AI has largely focused on technical improvements and people's ability to discriminate between human and AI-generated text [5, 26, 29]. Thus, people's perceptions of the output of generative AI have largely been neglected. Shifting the perspective from the technical aspects of text production to the public is paramount. Adoption of generative AI is likely to further accelerate and become an industry norm in the coming years due to its many advantages—efficiencies in coverage, speed,

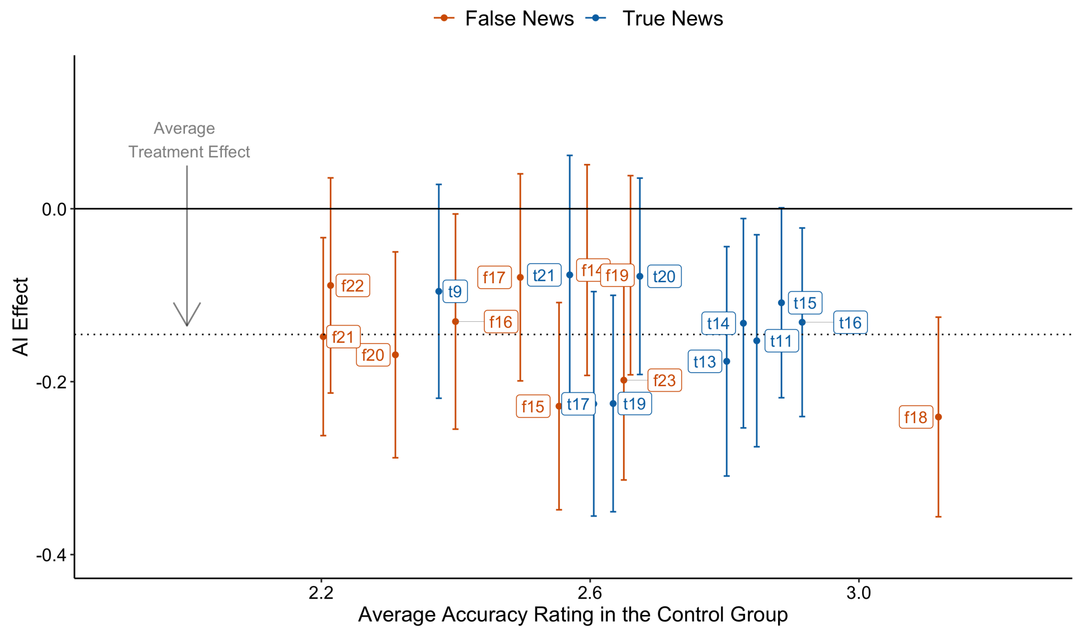

Figure 2: Results of Experiment 2: Negative effect of AI disclosure on perceptions of news accuracy by news item

and costs, and tremendous leaps in the capability to generate text indistinguishable from text generated by humans. As adoption of AI widens, however, so will the pressure from governance bodies to disclose its use and address standards of transparency and accountability [14, 37], and to remedy the potential for bias and misuse [1, 28]. Examining how the public will respond to disclosure of use of generative AI is thus a somewhat neglected priority [22].

An important implication of our experiments is that calls for transparency in the use of AI may backfire. Our results point to a potentially detrimental consequence of disclosing use of generative AI, which may further exacerbate the already declining public trust in news outlets [10, 43]. However, we wish to note that this implication may hinge on the assumption that the public views a human reporter as the default. If this assumption were to shift, and the public started viewing AI as the default reporter, disclosing use of AI may have beneficial effects on trust. Future research could build on our findings and test whether shifts in priors represent boundary conditions of where our results apply.

Our research has limitations that offer several opportunities for future research. First, future research could map out the theoretical boundaries of aversion to algorithmic reporting. Even though our results point to trust as mediating variable, we acknowledge the limitations of this approach and refrain from making strong causal claims. Future research could build on our findings both by testing alternative psychological processes leading to a negative AI effect, and by exploring whether there are circumstances where the effect reverses, with people appreciating (rather than derogating) AI reporters. For instance, recent research has shown that threat of inequality in medical outcomes or hiring decisions increases preference for algorithms [4, 24].

A second set of limitations pertains to the stimuli we employed in our experiments. We focused on news related to the COVID-19 pandemic. Future research could systematically examine whether relying on AI to generate text may be associated with higher (rather than lower) accuracy perceptions and overall evaluation as a function of the type of domain. For instance, it is possible that generative AI could have a positive effect in domains characterized by high consequentiality for the self (e.g., frivolous versus consequential) and base-rate uncertainty (e.g., well-known versus unknown). Similarly, although political ideology is correlated with perceptions of information related to COVID-19 [39], there is nonetheless imperfect overlap, and individuals—who may be polarized ideologically or engage in motivated reasoning on other issues—may not respond homogeneously to issues related to COVID. Future research could explore perceptions of generative AI for news unrelated to COVID, and test whether our results vary depending on political orientation and other demographic variables.

Another limitation is that we tested news headlines rather than whole articles. We made this choice based on considerations pertaining to both external and construct validity. In terms of external validity, we used headlines because the public largely consumes news at the level of story headline [44]. In terms of construct validity, it is unclear how people would assess the accuracy of an article comprising several statements, each statement potentially varying in perceived accuracy (e.g., it is unclear whether people would average the perceived accuracy of each statement or weigh the perceived accuracy of the first statement more, etc.). As our key variable was accuracy perceptions, we decided to assess perceptions of accuracy of one statement (i.e., a headline), which we could predetermine to be objectively entirely true or entirely false. Even though this approach has been adopted by other research on accuracy judgments [41], future research could investigate the effect of AI disclosure on perceptions of accuracy of full articles or text. For instance, it is possible that AI's perceived objectivity might matter more if participants were to read full articles rather than headlines. Future research could test whether our results replicate or reverse in the context of full news articles produced by AI.

Overall, this research is only the first step toward understanding the complex phenomenon of how people perceive generative AI. Given the speed with which AI systems are being developed and adopted, we hope this research will spur further investigation of this important topic.

ACKNOWLEDGMENTS

We gratefully acknowledge the invaluable help of Asif Mehedi and the Batten Institute at the Darden School of Business, without whom this paper would not have been possible.

REFERENCES

[1] Abubakar Abid, Maheen Farooqi, and James Zou. [n.d.]. Large Language Models Associate Muslims with Violence. 3, 6 ([n. d.]), 461–463. https://doi.org/10.1038/s42256-021-00359-2
[2] Hunt Allcott and Matthew Gentzkow. [n.d.]. Social Media and Fake News in the 2016 Election. 31, 2 ([n. d.]), 211–36. https://doi.org/10.1257/jep.31.2.211
[3] Benedikt Berger, Martin Adam, Alexander Rühr, and Alexander Benlian. [n.d.]. Watch Me Improve—Algorithm Aversion and Demonstrating the Ability to Learn. 63, 1 ([n. d.]), 55–68. https://doi.org/10.1007/s12599-020-00678-5
[4] Yochanan E Bigman and Kurt Gray. [n.d.]. People Are Averse to Machines Making Moral Decisions. 181 ([n. d.]), 21–34. https://doi.org/10.1016/j.cognition.2018.08.003
[5] Tom B Brown, 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, Alec Radford, Ilya Sutskever, and Dario Amodei. [n.d.]. Language Models Are Few-Shot Learners. http://arxiv.org/abs/2005.14165
[6] Romain Cadario, Chiara Longoni, and Carey K Morewedge. [n.d.]. Understanding, Explaining, and Utilizing Medical Artificial Intelligence. 5, 12 ([n. d.]), 1636–1642. https://doi.org/10.1038/s41562-021-01146-0
[7] Matt Carlson. [n.d.]. The Robotic Reporter: Automated Journalism and the Redefinition of Labor, Compositional Forms, and Journalistic Authority. 3, 3 ([n. d.]), 416–431. https://doi.org/10.1080/21670811.2014.976412
[8] Noah Castelo, Maarten W Bos, and Donald R Lehmann. [n.d.]. Task-Dependent Algorithm Aversion. 56, 5 ([n. d.]), 809–825. https://doi.org/10.1177/0022243719851788
[9] Victor Chernozhukov, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. [n.d.]. Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India.
[10] Luca Cian, Chiara Longoni, and Aradhna Krishna. [n.d.]. Advertising a Desired Change: When Process Simulation Fosters (vs. Hinders) Credibility and Persuasion. 57, 3 ([n. d.]), 489–508. https://doi.org/10.1177/0022243720904758
[11] Robyn M Dawes. [n.d.]. The Robust Beauty of Improper Linear Models in Decision Making. 34, 7 ([n. d.]), 571. https://doi.org/10.1037/0003-066x.34.7.571
[12] Nicholas Dias, Gordon Pennycook, and David G Rand. [n.d.]. Emphasizing Publishers Does Not Effectively Reduce Susceptibility to Misinformation on Social Media. 1, 1 ([n. d.]). https://doi.org/10.37016/mr-2020-001
[13] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. [n.d.]. Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err. 144, 1 ([n. d.]), 114. https://doi.org/10.1037/xge0000033
[14] Konstantin Nicholas Dörr. [n.d.]. Mapping the Field of Algorithmic Journalism. ([n. d.]). https://doi.org/10.1080/21670811.2015.1096748
[15] Ullrich KH Ecker, Stephan Lewandowsky, and David TW Tang. [n.d.]. Explicit Warnings Reduce but Do Not Eliminate the Continued Influence of Misinformation. 38, 8 ([n. d.]), 1087–1100. https://doi.org/10.3758/mc.38.8.1087
[16] OECD (Organisation for Economic Cooperation and Development). [n.d.]. The Role of Media and Investigative Journalism in Combating Corruption. https://www.oecd.org/corruption/The-role-of-media-and-investigative-journalism-in-combating-corruption.htm
[17] Jeffrey Gottfried and Elisa Shearer. [n.d.]. News Use across Social Media Platforms 2016. https://www.pewresearch.org/journalism/2016/05/26/news-use-across-social-media-platforms-2016/
[18] Heather M Gray, Kurt Gray, and Daniel M Wegner. [n.d.]. Dimensions of Mind Perception. 315, 5812 ([n. d.]), 619–619. https://doi.org/10.1126/science.1134475
[19] Kurt Gray and Daniel M Wegner. [n.d.]. Feeling Robots and Human Zombies: Mind Perception and the Uncanny Valley. 125, 1 ([n. d.]), 125–130. https://doi.org/10.1016/j.cognition.2012.06.007
[20] Rainer Greifeneder, Mariela Elena Jaffé, Eryn Newman, and Norbert Schwarz (Eds.). [n.d.]. *The Psychology of Fake News: Accepting, Sharing, and Correcting Misinformation.* Routledge, Taylor & Francis Group.
[21] William M Grove and Paul E Meehl. [n.d.]. Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical–Statistical Controversy. 2, 2 ([n. d.]), 293. https://doi.org/10.1037/1076-8971.2.2.293
[22] Andrea L Guzman. [n.d.]. Prioritizing the Audience's View of Automation in Journalism. 7, 8 ([n. d.]), 1185–1190. https://doi.org/10.1080/21670811.2019.1681902
[23] Chrisopher K Hsee, George F Loewenstein, Sally Blount, and Max H Bazerman. [n.d.]. Preference Reversals between Joint and Separate Evaluations of Options: A Review and Theoretical Analysis. 125, 5 ([n. d.]), 576. https://doi.org/10.1037/0033-2909.125.5.576
[24] Arthur S Jago and Kristin Laurin. [n.d.]. Assumptions about Algorithms' Capacity for Discrimination. ([n. d.]), 01461672211016187. https://doi.org/10.1177/01461672211016187
[25] Asher Koriat, Morris Goldsmith, and Ainat Pansky. [n.d.]. Toward a Psychology of Memory Accuracy. 51, 1 ([n. d.]), 481–537. https://doi.org/10.1146/annurev.psych.51.1.481
[26] Sarah Kreps, R Miles McCain, and Miles Brundage. [n.d.]. All the News That's Fit to Fabricate: AI-generated Text as a Tool of Media Misinformation. 9, 1 ([n. d.]), 104–117. https://doi.org/10.1017/xps.2020.37
[27] Arie W Kruglanski. [n.d.]. The Psychology of Being" Right": The Problem of Accuracy in Social Perception and Cognition. 106, 3 ([n. d.]), 395. https://doi.org/10.1037/0033-2909.106.3.395
[28] Nils Köbis, Jean-François Bonnefon, and Iyad Rahwan. [n.d.]. Bad Machines Corrupt Good Morals. 5, 6 ([n. d.]), 679–685. https://doi.org/10.1038/s41562-021-01128-2
[29] Nils Köbis and Luca D. Mossink. [n.d.]. Artificial Intelligence versus Maya Angelou: Experimental Evidence That People Cannot Differentiate AI-generated from Human-Written Poetry. 114 ([n. d.]), 106553. https://doi.org/10.1016/j.chb.2020.106553
[30] David MJ Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, and David Rothschild. [n.d.]. The Science of Fake News. 359, 6380 ([n. d.]), 1094–1096. https://doi.org/10.1126/science.aao2998
[31] Stephan Lewandowsky, Ullrich KH Ecker, Colleen M Seifert, Norbert Schwarz, and John Cook. [n.d.]. Misinformation and Its Correction: Continued Influence and Successful Debiasing. 13, 3 ([n. d.]), 106–131. https://doi.org/10.1177/1529100612451018
[32] Chiara Longoni, Andrea Bonezzi, and Carey K Morewedge. [n.d.]. Resistance to Medical Artificial Intelligence. 46, 4 ([n. d.]), 629–650. https://doi.org/10.1093/jcr/ucz013
[33] Chiara Longoni, Andrea Bonezzi, and Carey K Morewedge. [n.d.]. Resistance to Medical Artificial Intelligence Is an Attribute in a Compensatory Decision Process: Response to Pezzo and Beckstead (2020). 15, 3 ([n. d.]), 446–448.
[34] Chiara Longoni and Luca Cian. [n.d.]. Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The "Word-of-Machine" Effect. ([n. d.]). https://doi.org/10.1177/0022242920957347
[35] Francesco Marconi. [n.d.]. Newsmakers. In *Newsmakers.* Columbia University Press.
[36] Paul E. Meehl. [n.d.]. *Clinical versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence.* University of Minnesota Press. x, 149 pages. https://doi.org/10.1037/11281-000
[37] Tal Montal and Zvi Reich. [n.d.]. I, Robot. You, Journalist. Who Is the Author? Authorship, Bylines and Full Disclosure in Automated Journalism. 5, 7 ([n. d.]), 829–849. https://doi.org/10.1080/21670811.2016.1209083
[38] Gordon Pennycook, Tyrone D Cannon, and David G Rand. [n.d.]. Prior Exposure Increases Perceived Accuracy of Fake News. 147, 12 ([n. d.]), 1865. https://doi.org/10.1037/xge0000465
[39] Gordon Pennycook, Jonathon McPhetres, Bence Bago, and David G Rand. [n.d.]. Beliefs about COVID-19 in Canada, the United Kingdom, and the United States: A Novel Test of Political Polarization and Motivated Reasoning. ([n. d.]), 01461672211023652. https://doi.org/10.1177/01461672211023652
[40] Gordon Pennycook, Jonathon McPhetres, Yunhao Zhang, Jackson G Lu, and David G Rand. [n.d.]. Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention. 31, 7 ([n. d.]), 770–780. https://doi.org/10.1177/0956797620939054
[41] Gordon Pennycook and David G Rand. [n.d.]. The Psychology of Fake News. 25, 5 ([n. d.]), 388–402. https://doi.org/10.1016/j.tics.2021.02.007

[42] The Media Insight Project. [n.d.]. What the Public Expects from the Press (and What Journalists Think). https://www.americanpressinstitute.org/publications/reports/survey-research/public-expects-from-press/
[43] Michael Schudson. [n.d.]. The Fall, Rise, and Fall of Media Trust. 58, 1 ([n. d.]), 26.
[44] E Shearer and KE Matsa. [n.d.]. News Use across Social Media Platforms 2018. https://www.pewresearch.org/journalism/2018/09/10/news-use-across-social-media-platforms-2018/
[45] Dustin Tingley, Teppei Yamamoto, Kentaro Hirose, Luke Keele, and Kosuke Imai. [n.d.]. Mediation: R Package for Causal Mediation Analysis. ([n. d.]). https://doi.org/10.18637/jss.v059.i05
[46] Soroush Vosoughi, Deb Roy, and Sinan Aral. [n.d.]. The Spread of True and False News Online. 359, 6380 ([n. d.]), 1146–1151. https://doi.org/10.1126/science.aap9559
[47] Stefan Wager and Susan Athey. [n.d.]. Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests. 113, 523 ([n. d.]), 1228–1242. https://doi.org/10.1080/01621459.2017.1319839
[48] Kathleen Walch. [n.d.]. *AI Laws Are Coming*. Forbes. https://www.forbes.com/sites/cognitiveworld/2020/02/20/ai-laws-are-coming/

A LIST OF NEWS HEADLINES BY EXPERIMENT, EXPERIMENTAL WAVE, AND DATE OF FACT-CHECKING

Table A.1: List of true news headlines

Code Name	Headline	Date it appeared on Snopes.com	Experiment
T1	Ivanka Trump Holds Variety of Trademarks in China, Including One For Coffins	14-Apr-20	Exp. 1 (wave 1))
T2	Obama Urged US Pandemic Preparedness in 2014	13-Apr-20	Exp. 1 (wave 1)
T3	Trump Praises China for Its 'Transparency' on COVID-19	16-Apr-20	Exp. 1 (wave 1)
T4	Experts and Officials Warned in 2018 US Couldn't Respond Effectively to a Pandemic	1-Apr-20	Exp. 1 (wave 1)
5T	Trump Administration Sends 18 Tons of Personal Protective Equipment to China in Early 2020	31-Mar-20	Exp. 1 (wave 1)
6T	Las Vegas Homeless Sleep in 'Social Distanced' Parking Lot	31-Mar-20	Exp. 1 (wave 1)
7T	Empire State Building Displays 'Siren' Lights During COVID-19 Pandemic	31-Mar-20	Exp. 1 (wave 1)
8T	Amazon Solicits Donations to Help Pay Worker Sick Leave	25-Mar-20	Exp. 1 (wave 1)
9T	World Wrestling Entertainment CEO Vince McMahon Advises Trump on Reopening the U.S. economy	16-Apr-20	Exp. 1 (wave 1) Exp. 2
10T	Trump Golfs And Holds Rallies After Learning About COVID-19 Threat	1-Apr-20	Exp. 1 (wave 1) Exp. 2
11T	Cities Closed Schools and Businesses During the 1918 Pandemic	31-Mar-20	Exp. 1 (wave 1) Exp. 2
12T	Trump's Name To Appear on COVID-19 Stimulus Checks	15-Apr-20	Exp. 1 (wave 1) Exp. 2
13T	Mass Graves Dug in New York's Hart Island For COVID-19 Deaths	10-Apr-20	Exp. 1 (wave 1) Exp. 2
14T	CBS News Use Footage from Italy for New York COVID-19 Report	9-Apr-20	Exp. 1 (wave 1) Exp. 2
15T	Time Magazine Warned About Global Warming and Pandemic Years Ago	9-Apr-20	Exp. 1 (wave 1) Exp. 2
16T	Ohio Man Who Called COVID-19 a 'Political Ploy' Dies from the Disease	22-Apr-20	Exp. 1 (wave 1) Exp. 2
17T	Photo Shows Flyer for an 'End the Lockdown' Rally	22-Apr-20	Exp. 1 (wave 1) Exp. 2
18T	Trump Suggests Injecting Disinfectants as COVID-19 Treatment	24-Apr-20	Exp. 1 (wave 2)
19T	President Trump Tweets That Reporters Should Return 'Noble' Prizes	27-Apr-20	Exp. 1 (wave 2) Exp. 2
20T	2008 Nobel Prize Winner Luc Montagnier Said That COVID-19 Coronavirus Disease Was Artificially Created in a Lab	29-Apr-20	Exp. 1 (wave 3) Exp. 2
21T	Trump Blames Obama for 'Bad' COVID-19 Tests	1-May-20	Exp. 1 (wave 3) Exp. 2

Table A.2: List of false news headlines

Code Name	Headline	Date it appeared on Snopes.com	Experiment
1F	Michigan's Governor Bans Sale of American Flags, Plants, Seeds	13-Apr-20	Exp. 1 (wave 1)
2F	Experts Recommend to Sanitize Fabric Masks in a Microwave	10-Apr-20	Exp. 1 (wave 1)
3F	Seven Children Die in Senegal from COVID-19 Vaccine	10-Apr-20	Exp. 1 (wave 1)
4F	Himalayas Visible from Northern India for First Time in 30 Years	9-Apr-20	Exp. 1 (wave 1)
5F	Bill Gates Sued by India Over Vaccination Deaths	10-Apr-20	Exp. 1 (wave 1)
6F	Kenyan Government Has Maasai Tribe Whip People To Enforce Curfew	7-Apr-20	Exp. 1 (wave 1))
7F	Plagues Repeated Exactly Every 100 Years	7-Apr-20	Exp. 1 (wave 1)
8F	Uninsured Teen Die of COVID-19 After Being Denied Treatment	3-Apr-20	Exp. 1 (wave 1)
9F	Ventilators Found 'Stashed' in a Warehouse in New York	31-Mar-20	Exp. 1 (wave 1)
10F	Financier George Soros Owns Lab in China Where COVID-19 Was "Developed."	2-Apr-20	Exp. 1 (wave 1)
11F	Your Coronavirus Stimulus Check Counts Against Your 2020 Tax Refund	7-Apr-20	Exp. 1 (wave 1))
12F	Trump Tweets in 2009 That He Would 'Never Let Thousands of Americans Die From a Pandemic'	17-Apr-20	Exp. 1 (wave 1)
13F	News Media Fake Photos of Jacksonville Beaches	20-Apr-20	Exp. 1 (wave 1)
14F	CDC Guidelines for Reporting COVID-19 Deaths Artificially Inflate Numbers	20-Apr-20	Exp. 1 (wave 1) Exp. 2
15F	The House Gives Itself a $25M Raise in Coronavirus Aid Bill	31-Mar-20	Exp. 1 (wave 1) Exp. 2
16F	Bill Gates and the ID2020 Coalition Are Using COVID-19 To Build Global Surveillance State	22-Apr-20	Exp. 1 (wave 1) Exp. 2
17F	Nancy Pelosi Visited Wuhan in November 2019	26-Apr-20	Exp. 1 (wave 2) Exp. 2
18F	Photo Shows Signs Carried by COVID-19 Anti-Lockdown Protesters	23-Apr-20	Exp. 1 (wave 2) Exp. 2
19F	Nobel Laureate Tasuku Honjo Says COVID-19 Was 'Man-Made'	27-Apr-20	Exp. 1 (wave 3) Exp. 2
20F	Elisa Granato, One of the UK's First Covid-19 Vaccine Trial Participants, Has Died	27-Apr-20	Exp. 1 (wave 3) Exp. 2
21F	1866 Court Case Bar States from Enforcing Social-Distancing Regulations	27-Apr-20	Exp. 1 (wave 3) Exp. 2
22F	CDC Readjusted the COVID-19 Death Toll From 60,000 Down to 37,000	27-Apr-20	Exp. 2
23F	Churches in Kansas City Required to Record List of Attendees	27-Apr-20	Exp. 2