

Using Large Language Models to ‘Lighten the Mood’: Satirically Reframing News Recommendations to Reduce News Avoidance

Tobias J. Wessel¹, Christoph Trattner¹ and Alain Starke¹

¹MediaFutures, University of Bergen, Lars Hilles Gate 30, Bergen, Vestland, 5008, Norway

Abstract

News avoidance is a growing issue that leads to less informed citizens and endangers democratic processes. This also poses problems in news recommender environments, as ‘unpleasant’ news content could be avoided through personalized algorithms. To ‘lighten the user’s mood’, this paper investigates whether satirical re-framing of news article summaries, generated by Large Language Models (LLMs), can mitigate news avoidance by making news content more engaging. Through two online experiments ($N = 89$; $N = 151$), we tested various prompting techniques, assessing the impact on user perception, humor, understanding, and news consumption choices. Results indicate that satirically re-framed summaries were perceived to be engaging and informative. Less frequent news consumers showed a stronger preference for satirical content, suggesting that satire could be a tool for reconnecting with disengaged audiences. These findings show the promise of AI-generated personalized satire as an innovative approach to reducing news avoidance.

Keywords

News Avoidance, Computational Humor, Large Language Models, News Recommender Systems, Personalization

1. Introduction

News avoidance has become an increasingly pressing issue in recent years. It can be defined as the action of intentional disconnection or rejection from news content [1, 2]. Global surveys document a sharp increase in the proportion of people who *often* or *sometimes* avoid the news. For example, 38% of respondents to the 2022 Reuters Report indicated to actively avoid the news, up from 29% in 2017 [3, 4]. In various countries, the proportion of news avoiders are now over 40% [3]. Younger audiences are particularly disengaged, with over a third of respondents under the age of 35 stating that news consumption negatively impacts their emotions [4].

This leads to the question: Why do people avoid news? Prior research suggests a combination of psychological and content-based drivers [5]. News avoidance may be caused by the perception that news is too negative, depressing, or emotionally heavy [2]. Skovsgaard and Andersen (2020) conceptualize this intentional news avoidance as: *users opt out of news to guard their mood*. For example, the tone of news articles, often focused around a societal issue, is a common reason consumers keep avoiding it [2]. In addition, people may avoid specific topics, because they are perceived as boring, complex or upsetting. Political news, for example, may be interpreted to be partisan or overly negative, while climate news may induce a sense of powerlessness [6, 7].

A news category that tends to be perceived as more positive is satire. In general, satire intends to entertain and critique, and these dimensions should interact [8]. News satire tends to incorporate more journalistic dimensions, taking current affairs as its topic, while also including elements that are clearly dissociated from the truth [9], which raises the entertainment value.

This paper examines to what extent satirical re-framing of news content can persuade readers to increase their news consumption. Satire has a long history in journalism and entertainment as a tool to actively engage audiences by reporting events in a lighter, more humorous way. Examples include *The*

Proceedings of the 13th International Workshop on News Recommendation and Analytics (INRA 2025), co-located with the 19th ACM Conference on Recommender Systems (RecSys 2025), September 22–26, 2025, Prague, Czech Republic.

✉ tobias.wessel@student.uib.no (T. J. Wessel); christoph.trattner@uib.no (C. Trattner); alain.starke@uib.no (A. Starke)

id 0009-0008-4310-9555 (T. J. Wessel); 0000-0002-1193-0508 (C. Trattner); 0000-0002-9873-8016 (A. Starke)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Daily Show and *Last Week Tonight*, that have demonstrated how blending serious events with humor can attract viewers [8]. Prior studies have found that political satire can increase the viewers interest and understanding of a topic, by mixing the negative aspect of news with humor [2]. However, it may also lead to more selective exposure for partisan political news [10].

We will examine the potential of using large language models for the re-framing task. We use such LLMs to summarize existing news articles in a more satirical tone, applying it to a news recommender system interface. The combination of satire and news recommender is an underexplored research area, which leads to the following research questions:

- **RQ1a:** How does consuming news in a satirical format affect the user’s sense of feeling informed compared to traditional news formats?
- **RQ1b:** What prompting style performs best on humor when re-framing news articles to be more satirical?
- **RQ2a:** To what extent can satirical news article summaries created by AI help reduce the occurrence of news avoidance?
- **RQ2b:** How does news consumption frequency influence the effectiveness of AI-generated satirical summaries in reducing the avoidance of news content?

Recent advances in AI have made our strategy of using satirical re-framing feasible at scale. Powerful large language models (LLMs) such as GPT-4 enable the automation of human-like text generation, as also shown by Jeng et al. [11]. However, employing AI for humor introduces specific challenges. Previous research has highlighted audience biases toward AI-generated content, indicating that users may approach such material with skepticism or low expectations, especially due to the subjective nature of humor [12]. However, when users are unaware of the source, AI-created jokes can be rated equally as human-created [12, 13]. To better understand this dynamic, we explore in detail how different prompt styles impact humor perception and user engagement (RQ1b). Specifically, we develop and compare three distinct strategies for creating satirical news summaries based on insights from previous studies.

2. Background

2.1. News Avoidance: Types and Drivers

News Avoidance is typically defined as the practice of limiting ones exposure to news, either in a more general manner or within certain topics [5]. A distinction is made between intentional and unintentional news avoidance [2], where intentional avoidance can be caused by the overwhelming volume of news content, leading to avoidance to escape stress [2]. Unintentional avoidance, on the other hand, occurs without an explicit decision to avoid the news, which can be caused by users being distracted, for example by entertainment on social media feeds [2]. Recommendation algorithms and user habits can eventually lead to the consumption of a “softened” media diet.

Both factors play a role in the overall decline in news consumption. Researchers describe this as the “News Avoidance Paradox”: As access to news has never been easier, significant proportions of the population are actively disengaging from it [2, 14]. Furthermore, topical avoidance can also occur, as users tend to avoid certain topics they find unpleasant or uninteresting [2]. This might be critical in the context of news recommender systems, as they can assist in mitigating unintentional avoidance and topical avoidance.

2.2. Political Satire and News Engagement

Satire tends to be higher in entertainment [8]. Many studies in communication science and media studies have examined how satirical news and political comedy affect audience engagement, knowledge, learning, and attitudes. Satire can act as a gateway for ‘light’ news consumers to engage with societal issues and increase attention toward ‘hard news’ [15]. Humorous or more interest-oriented framing of serious topics has been found to draw the attention of light consumers and, in some cases, even lead

them to seek out more information from traditional outlets [15]. The idea is that the humor lowers the barriers and audiences that might have found the straight news to repetitive or boring can still be informed when consuming news in a more entertaining format. This showed how invoking humor and absurdity tends to generate a form of amusement that can counteract the otherwise negative tone of the news [16]. Secondly, satire is shown to support comprehension and learning. Becker and Bode [15] show how participants that watched satirical news learned just as much factual information as those who watched a traditional news report. In some cases, satire even increased the viewers' recall of details, possibly by representing the information in a way that was more memorable. However, results on learning are mixed, as Burgers and Brugman [16] note that while satire improved knowledge compared to no information, it performed similarly to straight news when it comes to factual learning. In summary, political and news satire has shown to be a valuable complement to traditional news by maintaining informational content while wrapping it in entertainment.

2.3. Personalization & News Recommenders

The role of satire in a news recommender system is very limited, as it is not mentioned in literature surveys and reviews [17, 18]. What stands out in recommender research is that by adapting the selection and ranking of news stories to each user's interest and behavior, this improves the user experience by increasing relevance [17, 19, 18]. Important challenges are news avoidance and selective exposure, for algorithms may simply reinforce user preferences without considering what is important to inform a user. The failure to consider normative aspects of news recommendations may lead to less informed citizens [20, 21]. For example, only considering personalization may lead to the omission of harder news categories, such as economic news or war coverage, and may marginalize certain voices in the public debate [22].

Our research takes a more novel angle on personalization: rather than recommending only content that users like, we recommended content we presume they tend to avoid, but in a format they might like. This approach still respects preferences at some level, as we use personalization to identify avoided topics and then tailor how we present these topics. Instead of completely disregarding a user preference, we re-frame the content fit more to the users liking. The idea builds on the intuition that news avoiders are not indifferent to important issues but rather turned away by how its presented. By re-presenting the same factual content in a more engaging frame, we seek to explore ways to reconnect with segments of the lost audiences. Consequently, we had the following behavioral hypothesis:

- **H1:** Users are more likely to choose satirical versions of news articles in non-preferred categories, compared to those in preferred categories.

2.4. Computational Humor

LLMs have been used to a limited extent to re-frame news articles in the context of a simple recommender system. Jeng et al. [11] show that using ChatGPT to write news articles in a different emotional tone can change user preferences. In contrast, the application of LLMs for generating news satire is new. Computational humor has historically posed significant challenges due to humor's dependency on creativity, context, and nuanced language [23]. With recent advances in natural language processing, LLMs such as GPT-4 now possess the ability to generate more sophisticated humor, closely approximating human-written jokes when prompted effectively [13, 12].

Previous research indicates that participants often struggle to distinguish between AI- and human-generated humor, indicating that advanced models can produce convincing comedic content [13]. However, transparency influences a user's perception, as they tend to judge AI-generated humorous content more critically. Building on these findings, our prompting strategies incorporate proven humor techniques from previous studies. One condition leverages lighter satire, known to maximize engagement while maintaining user trust [16]. Another employs more aggressive humor, aligning with previous observations that ChatGPT performs particularly well with bold comedic styles [13]. All prompting strategies were carefully structured, as structured prompts consistently yield the best

humor output from LLMs [13]. Given the role trust plays in user reactions to AI-generated humor, we formulate the following hypothesis:

- **H2:** A user’s general trust in generative AI positively affects the reported trust in satirically manipulated articles.

2.5. Contribution

We present two studies in which we examine the effects of satirically re-framing news articles and how they compare to their original counterparts. This is the first study to use this strategy to reduce news avoidance, as well as the first study to do so in relation to news preferences and recommender systems. It is also one of the first to apply LLMs to re-frame a news article to present an alternative, new version of an existing one. We show how satirical versions of news articles could have positive effects on people who are more likely to avoid news consumption. In the first study, we tested which type of prompting technique is effective in this context, while the second examined the user’s evaluation of both original and satire-based news articles.

3. Study 1

In order to address RQ1a and RQ1b, we performed an online experiment designed to evaluate how satirical re-framing of news summaries using Large Language Models (LLMs) influences reader perceptions, enjoyment, and comprehension. Participants were asked to evaluate nine different news article summaries, each manipulated to be more satirical. We also assessed which prompting techniques are most effective to generate satirical news articles.

3.1. Methods

3.1.1. Dataset

We created a dataset consisting of 90 news article summaries, which were evenly distributed across nine different categories (10 articles per category). For representativeness for U.S. audiences, categories were based on the most frequently read categories in the Washington Post (2012-2018) and the most published article categories by Huffington post (2012-2022) [24, 25, 26, 27]. News article summaries were selected to represent the U.S. media landscape, based on their popularity and political alignment, using rankings provided by the news aggregator Ground News [28]. Summaries of news articles, along with essential metadata, were retrieved in real time using the NewsCatcher API [29], which has been previously used for research purposes. Article summaries were manipulated into a satirical format by OpenAI’s ChatGPT-4 model. This was an advanced large language model selected due to its accessibility, cost-effectiveness, and prior success in generating humorous content in related studies [13]. Each summary was approximately 80 words in length, truncated consistently across original and manipulated versions to maintain comparability.

Three distinct satirical re-framing techniques were tested across the dataset. Condition A (*gentle satire*) featured subtle humor and mild wordplay without significant exaggeration. Condition B (*bold satire*) employed exaggerated and playful humor inspired by traditional satirical outlets such as The Onion (cf. [8]), guided by humor-focused cues. Condition C (*free satire*) allowed greater creative freedom, instructing the LLM to adopt a persona of a satirical columnist, with fewer restrictions on style and delivery. All prompting conditions were explicitly told to “keep all factual content intact”. The full dataset used in Study 1 along with the exact prompts used to re-frame the content can be found in our repository¹.

¹<https://anonymous.4open.science/r/Study1-anon-01EE/README.md>.

3.1.2. Participants

A total of 103 participants were recruited via the crowdsourcing platform Prolific. They were residents of the United States, fluent in English, and maintained at least a 95% approval rating on the platform. The study took 12 minutes to complete, for 9 GBP per hour. We eventually retained $N = 89$ participants, removed those who did not pass the attention check or provided incomplete data.

3.1.3. Research Design

Participants were randomly assigned to different conditions. We used a mixed 2 (between) x 3 (within) design to examine the impact of satirical re-framing of news article summaries on the perceptions and comprehension of the participants. The between factor involved transparency of LLM use, where participants were told about the use of AI for summaries or not. The within factor concerned the prompting style, which varied in terms of the level and nature of humor employed. The temperature of the ChatGPT 4 model was set to a more balanced 0.6 to ensure engagement and creativity, but also in an attempt to give the different prompting techniques more control over the outputted content. Each participant evaluated a total of nine summaries, three from each prompting style. The presentation order of article categories and prompting styles was fully randomized for each participant to control for potential order effects.

3.1.4. Procedure

We first inquired on participants' demographics, news frequency use and possibly reasons for news avoidance. Then, participants in the transparent condition received a message that the summaries were manipulated to be more satirical by generative AI. All participants then evaluated nine news summaries, each corresponding to one distinct topical category. For each summary, participants responded to five questionnaire items designed to measure their perceptions and engagement with the content. Each of these items was measured using 7-point Likert scales. An attention check item appeared randomly during evaluations to confirm participants' attention. After evaluating all summaries, participants in the transparent condition completed an additional measure assessing their beliefs about the human- or AI-generated nature of the summaries.

3.1.5. Measures

Participants evaluated each news article summary using five distinct measures, each designed to capture a specific aspect of their experience with the content. All measures utilized 7-point Likert scales. Perceived Enjoyment was assessed with the statement "I found this article summary enjoyable," while Perceived Fun was measured by "I found this article summary funny." Understanding was evaluated by asking participants to rate their agreement with the statement "I clearly understood the key message of this article summary." The Intention to Share was measured using "I would be likely to share this article summary with a friend." Lastly, Perceived Summary Quality was captured through the statement "This article summary is well-written."

3.2. Results

We analyzed the users' experience with the AI-based re-framed news article summaries. In the full sample ($N = 792$ trials), participants' perceived fun rating ($M = 3.93$, $SD = 1.90$) did not differ from the neutral midpoint of 4, $t(791) = -1.05$, $p = .29$. In contrast, we found that perceived summary quality ratings were significantly above neutral ($M = 4.90$, $SD = 1.68$), $t(791) = 15.01$, $p < .001$, the transparency condition did not have an impact $t(758.1) = -0.56$, $p = .57$. Perceived enjoyment ratings also exceeded the midpoint ($M = 4.55$, $SD = 1.76$), $t(791) = 8.76$, $p < .001$, indicating that most participants found the satirical summaries engaging and entertaining despite not reporting very high humor scores. Lastly, intention to share, was significantly below the neutral midpoint of 4 ($M = 3.70$,

$t(791) = -3.98, p < .001, 95\%-CI [3.56, 3.85]$). Assessing these four dimensions could in other words not find that users found the content very humorous or the opposite, but did indicate that participants viewed the content to be engaging and of high quality.

3.2.1. RQ1b: Prompt-Evaluation

One-way ANOVAs revealed only non-significant differences across the three prompt styles. There were small differences for perceived fun: $F(2, 789) = 2.82, p = .060, \eta^2 \approx .007$. Perceived engagement showed a similar non-significant trend, $F(2, 789) = 2.53, p = .080$, while summary perceived quality neither varied across prompt types, $F(2, 789) = 2.27, p = .105$. However, when restricting our analysis to our “High” news-consumption cluster, consisting of a considerable portion of our data (81%), prompting condition C *Free satire* led to significantly more perceived fun than the conventional condition A, *gentle satire* (mean difference = 0.58, 95% CI [0.16, 0.99], $p = .003$, Tukey-adjusted). We observed no differences in relation to Condition B, *bold satire*. This subgroup finding provided the primary empirical basis for selecting *free satire* as our preferred LLM-prompt moving forward.

3.2.2. RQ1a: Sense of Feeling Informed

Participants’ overall understanding was shown to be well above the neutral benchmark of 4 through the understanding dimension ($M = 5.61, SD = 1.45$): $t(791) = 31.29, p < .001$. This indicated that participants still felt informed after receiving the content in a satirical format. Transparency was also found to affect understanding between the conditions without ($M = 5.49$) and with transparency ($M = 5.76$) disclosure: Welch’s $t(780.5) = -2.64, p = .0085$. This suggested that participants’ understanding was increased upon being notified about the use of AI for a news article summary. Interestingly, this indicates that transparency about the use of AI and the satirical nature of the article modestly boosts users’ sense of feeling informed. Participants who found the articles more funny also tended to report higher understanding ($r = .36, 95\% CI [.30, .42], p < .001$) as well as greater perceived summary quality ($r = .61, 95\% CI [.56, .65], p < .001$). Perceived enjoyment was also positively linked with both perceived summary quality ($r = .77, p < .001$) and understanding ($r = .60, p < .001$). These results suggested that satirical re-framing not only preserved understanding and learning as shown in earlier studies [2, 30, 15], but also suggested that it can help enhance those aspects for articles across different news categories.

3.3. Conclusion

Our findings from Study 1 indicate that while participants did not consistently rate satirical news summaries as highly humorous, they nonetheless found them engaging, of high quality, and effective at maintaining or even enhancing their sense of being informed. Importantly, transparency regarding the AI-generated nature of the summaries did not negatively impact user perceptions. Instead, it modestly improved comprehension. Given these promising results, particularly with the *free satire* prompting style, we proceeded to Study 2 to investigate whether this satirical approach could actively influence user choices and more directly mitigate news avoidance behaviors.

4. Study 2

To address RQ2a and RQ2b, Study 2 was designed as an online experiment to investigate how satirical summaries generated by Large Language Models (LLMs) could influence users’ selection of news content and potentially mitigate news avoidance behaviors. Participants were asked to actively choose between original and satirical versions of news article summaries based on their topical preferences.

4.1. Methods

4.1.1. Dataset

We used the same dataset as in Study 1 (cf. Section 3.1.1), consisting of 90 news article summaries consisting of 80 words. Each summary was re-framed into a satirical version using OpenAI's ChatGPT-4 model, leveraging the most effective prompting style identified in Study 1 (*free satire*), characterized by greater creative freedom and a columnist persona. The complete dataset, including all prompts and original summaries, can be accessed in our repository² for Study 2.

4.1.2. Participants

A total of 160 U.S. residents were recruited through the crowdsourcing platform Prolific. They were fluent in English, maintained at least a 95% approval rating on the platform, and were compensated at a rate of 9 GBP per hour. The study took approximately 15 minutes to complete. Following the application of exclusion criteria, such as failing attention checks or submitting incomplete data, the final sample for analysis included $N = 151$ participants.

4.1.3. Research Design

The design of Study 2 was subject to 2 (between) x 2 (within) mixed-factorial design. As in Study 1, the between factor involved Transparency, either disclosing the use of AI to participants or not. The within factor was News Preference, re-framing news articles from either preferred (top-2 selection) or non-preferred (3 least preferred) categories. Each participant evaluated five pairs of articles, which each consisted of one original and one satirical summary. The temperature for generating satirical summaries via ChatGPT-4 was consistently set at 0.6, promoting balance between creativity and coherence. The presentation order of article summaries and categories was randomized to counterbalance order effects.

4.1.4. Procedure

We first inquired on participants' demographic details and news consumption frequency. Afterwards, they rated their preferences across the nine topical news categories using 7-point Likert scales. Based on their responses, the system selected five articles: two from their most preferred categories and three from their least preferred categories. Each participant was then presented with pairs of summaries (original vs. satirical) and was asked to select their preferred version for each of the five articles. Participants were not initially informed about the satirical manipulation. Following their choices, participants in all conditions were informed of the manipulation in a subsequent screen and responded to additional questionnaire items designed to explore their reasons for selecting each summary and their attitudes toward the satirical summaries. Attention-check items were embedded to ensure participants' attentiveness and validity of responses.


4.1.5. Measures

Each participant's motivation was captured using several measures, all employing 7-point Likert scales. For each chosen summary, participants rated their perceptions across four dimensions: Perceived Veracity ("The article summary I chose seemed more factually accurate"), Perceived Humor ("I found the article summary I chose to be more humorous"), Perceived Entertainment ("I found the article summary I chose more entertaining"), and Perceived Understanding ("The article summary I chose was easier to understand"). Additionally, participants evaluated the satirical summaries specifically, using four measures that later formed a single factor labeled "Satire-based Engagement": Emotional Lightness ("The satirical article summary made this topic feel less heavy or overwhelming"), Future Interest ("I would like to read about this topic in a satirical format in the future"), Topical Engagement ("The


²<https://anonymous.4open.science/r/Study2-anon-9C8E/README.md>.

A - Select preference


Select Your Preferences
Rate each category based on your interest.
1 = No interest | 4 = Neutral | 7 = High Interest



Sports
Interest Level:
1 2 3 4 5 6 7
Please select an option




Politics
Interest Level:
1 2 3 4 5 6 7
Please select an option



Climate
Interest Level:
1 2 3 4 5 6 7
Please select an option

B - Make choice



Version 1

President Donald Trump plans to create commission on college athletics: Sources

Former Alabama head coach Nick Saban and Texas Tech board of regents chairman Cody Campbell would co-chair the commission President Donald Trump is interested in forming to examine the long list of issues facing college sports, a source briefed on the plans told The Athletic. The source said Trump is expected to be 'very engaged' with the commission because he sees the current state of college sports as an issue of national importance.

Version 2

Trump's Next Great Adventure: College Sports Commission, Starring Nick Saban and Cody Campbell

Former Alabama coach Nick Saban and Texas Tech's regent, Cody Campbell, are set to co-chair Trump's latest hobby: a commission on college sports. Our source, clearly thrilled, says Trump is ready to dive into this 'national issue' with the same passion he reserves for Twitter feuds. Because who needs healthcare reform when you've got college sports to worry about?

C - Evaluate choice & satirical version

President Donald Trump plans to create commission on college athletics: Sources

Former Alabama head coach Nick Saban and Texas Tech board of regents chairman Cody Campbell would co-chair the commission President Donald Trump is interested in forming to examine the long list of issues facing college sports, a source briefed on the plans told The Athletic. The source said Trump is expected to be 'very engaged' with the commission because he sees the current state of college sports as an issue of national importance.

Original version

Trump's Next Great Adventure: College Sports Commission, Starring Nick Saban and Cody Campbell

Former Alabama coach Nick Saban and Texas Tech's regent, Cody Campbell, are set to co-chair Trump's latest hobby: a commission on college sports. Our source, clearly thrilled, says Trump is ready to dive into this 'national issue' with the same passion he reserves for Twitter feuds. Because who needs healthcare reform when you've got college sports to worry about?

Satirical version

You chose the satirical version of this article.

Figure 1: (A) Shows the preference selection page of the study. (B) Shows the choosing experiment between the re-framed and original summary, while (C) shows the choice evaluation page.

Table 1

Results of the factor analysis with loadings for the factor of Satire-based Engagement

Aspect	Item Measure	Item	Loading
Satire-based Engagement	Emotional Lightness	The satirical article summary made this topic feel less heavy or overwhelming.	0.691
	Future Interest	I would like to read about this topic in a satirical format in the future.	0.937
	Topical Engagement	The satirical version of this article summary increased my interest in the topic.	0.932
	Mood Impact	The satirical version of this article summary had a positive impact on my mood.	0.924

satirical version of this article summary increased my interest in the topic”), and Mood Impact (“The satirical version of this article summary had a positive impact on my mood”). Satire-based engagement was formed using a principal component factor analysis; see Table 1: All items loaded strongly on that single factor ($.69 \leq \lambda \leq .94$) and the factor accounted for 77% of the common variance ($SS = 3.08$). Sampling adequacy was excellent ($KMO = .83$) and internal consistency was high (Cronbach’s $\alpha = .93$). Eventually, we formed a factor that consisted of the average across all of the items.

To study news avoidance, we divided participants into two groups based on their news consumption habits: High-frequency and low-frequency news consumers. This was based on participants’ self-reported news consumption frequency, with those reporting consuming news at least once a day categorized as high-frequency consumers, and those consuming news less often categorized as low-frequency consumers (i.e., multiple times a week or less). We deliberately set the cutoff to once daily consumption, rather than using a more strict median split, as categorizing daily news consumers into the low-frequency group would have inaccurately represented their consumption behavior. In our sample, 64 participants (42%) consumed news *several times a day*, another 44 (29%) consumed news *once*

a day, while only 41 participants (27%) reported consuming news less than daily. The daily threshold resulted in an unequal split: *high* consumers $n = 108$ vs. *low* consumers $n = 43$. We have accounted for the unequal group sizes with the models used in subsequent analyses.

4.2. Results

In study 2, we examined to what extent AI-generated satirical summaries could influence the selection of news content of users and potentially mitigate news avoidance. Across the entire data set (755 trials), participants chose the satirical summary 257 times, 34% of all choices ($M = 0.34$, $SD \approx 0.47$), indicating that roughly one in three summaries selected overall was the AI-generated satirical version.

4.2.1. Manipulation Check

Before performing the main part of the analysis, we verified that the two types of summary versions were perceived as intended. Across all 755 trials, participants perceived the *original* summary ($M_{\text{orig}} = 5.88$, $SE = 0.06$) to be higher in veracity than its satirical counterpart $M_{\text{sat}} = 4.36$, $SE = 0.10$: Wilcoxon $W = 99,174$, $p < .001$). The pattern was reversed for humor, where the *satirical* version ($M_{\text{sat}} = 5.25$, $SE = 0.10$) was considered funnier than the original version ($M_{\text{orig}} = 2.48$, $SE = 0.08$): $W = 19,697$, $p < .001$).

4.2.2. RQ2a: How can Satire Created by AI help mitigate News Avoidance

We compared the two groups of news users ($n = 43$ low-frequency vs $n = 108$ high-frequency users). Overall, low-frequency readers selected the satirical version in 43% of their 215 total trials, compared to 31% of the 540 trials among daily news consumers, which was significantly higher: $\chi^2(1) = 9.71$, $p = .002$. This is also depicted in Figure 3.

Beyond behavior, low-frequency consumers consistently rated satire more positively. Satire was judged to make the topic feel less heavy (4.74, $SD = 1.78$ vs 4.27, $SD = 1.85$, Wilcoxon $p = .005$), increased interest in the topics overall (4.36, $SD = 1.94$ vs 3.65, $SD = 2.10$, $p < .001$), positively affected a participant's mood (4.41, $SD = 1.94$ vs 3.77, $SD = 2.02$, $p < .001$), and sparked a desire to read more satirical news summaries on the given topic (median 5 vs 4, $p \approx 4 \times 10^{-6}$) among low-frequency consumers.

Ratings for perceived understanding were found to be similar across both groups of frequency of news use. However, both means were higher than the neutral benchmark (4) (5.50, $SD = 1.20$ vs 5.41, $SD = 1.47$, test for difference: $p = .88$). Collapsing the four items into a reliable 'satire-based engagement' scale ($\alpha = .93$) led to similar results: Low-frequency participants averaged 4.51, $SD = 1.66$ versus 3.85, $SD = 1.86$ for daily readers (Wilcoxon $p < .001$). This can also be seen in Figure 2. In addition, Figure 4 shows how this trend continues across all categories in the study, with low frequency readers reporting larger positive effects of satire across the news spectrum. This pattern reinforces our findings even further.

Taken together, the behavior and attitudinal data conveyed a clear answer to RQ2a. AI-generated satirical summaries can measurably reduce news avoidance by making more disengaged readers more inclined to read news articles.

4.2.3. H1 - Preferred topics vs. avoided topics

We compared 'preferred' and 'avoided' categories across participant responses. A single-parameter logistic model (prefScore = +1 for preferred, -1 for avoided) showed a reliable effect on the odds of choosing the satirical version ($\beta = -0.16$, $SE = 0.08$, $z = -2.00$, $p = .045$; OR ≈ 0.85). Put plainly, participants were 7 percentage points less likely to pick satire for topics they already liked (30%) than for topics they normally avoid (37%). A confirmatory one-way ANOVA on the two groups repeated the finding ($F = 4.04$, $p = .045$), and a mixed-effect model with random intercepts for participants and




Figure 2: Mean factor scores (1 = low impact ... 7 = high impact) for participants who read news less than daily (Low, $n = 215$) and those who read at least daily (High, $n = 540$). Error bars show $\pm 1.96 \times SE \approx 95\%$ confidence interval.




Figure 3: Percentage of trials where satirical summaries were preferred, by news-consumption group. Error bars show $\pm 1.96 \times SE \approx 95\%$ confidence interval.

categories again found the avoided > preferred contrast ($\beta = -0.21$, $SE = 0.09$, $z = -2.35$, $p = .019$). Although, an exploratory interaction revealed that the effect is carried almost in its entirety by the high frequency news group ($\beta = -0.52$, $SE = 0.20$, $p = .008$), whereas low-frequency readers showed no reliable difference ($\beta \approx 0$, $p > .70$). In other words, the data supports [H1], as satire was found to not be particularly attractive when a user preferred a topic, instead it was more persuasive for less favorable categories.

4.2.4. H2 - AI Trust Correlation with Choice

We tested whether trust in generative AI differed across transparency conditions. It was found to be consistent across the non-transparent ($M = 4.84$, $SD = 1.56$; $n = 80$) and transparent conditions ($M = 4.72$, $SD = 1.58$; $n = 71$). Hence, further testing considered both conditions to be a single group when testing [H2], examining whether trust in AI correlated with choice. The tests revealed that higher levels of trust in AI among participants correlated with a higher probability to select the satirical version of an article: A logistic model with z-scored AI-trust showed that every 1 SD increase in trust raised the odds of choosing the satirical summary by about 4% ($\beta = 0.21$, $SE = 0.08$, $z = 2.65$, $p = .008$, $OR \approx 1.24$). Performing a median split showed the same pattern: high-trust participants selected satire in 40% of their 275 trials, compared to the 30% of 480 trials among low-trust readers ($\chi^2(1) = 7.27$, $p = .007$). A full three-way model confirmed that neither the transparency, AI-trust nor news-consumption frequency effected the relationship (all interaction p -values $> .14$). These findings suggested that greater confidence in AI reliably predicts a higher chance of embracing satirical articles.




Figure 4: Means \pm 95% CI for low- and high-frequency news consumers, shown separately for low-frequency (red) and high-frequency (blue) readers.

5. Discussion

We have investigated to what extent satirical summaries of news articles, generated by LLMs, could help news avoiders to read more news. Our findings suggest that LLM-generated news satire can help re-engage individuals that might otherwise avoid news, or some news topics. Participants in our studies generally reported the re-framed summaries to be of high quality and engaging, an effect we saw strengthened withing the users who consumed news less frequently in the second study. Notably, the appeal seemed to be strongest for articles the participants reported to dislike, and we saw positive effects when recommending disliked articles in a more satirical format. This implies how humor can

in fact lower some of the psychological barriers that have been shown to lead to intentional news avoidance [14]. By presenting these disliked topics in a more lighthearted way, the satirical framing seemed to draw readers into content they might normally skip. With the common reason for avoidance being due to feeling overwhelmed or disinterested, or results indicate that more positive humorous framing might counteract those feelings.

Our outcomes resonate with previous media research on the power of humor. Prior studies had shown that adding humor to news can boost audience engagement and in some studies was shown to even improve the retention of information [31]. In our online experiments, participants reported improved understanding when reading the satirical summaries, suggesting that the use of satire improved how informed participants felt without sacrificing the informational content of the article summary. This also suggests that information and entertainment can indeed co-exist, which adheres to some definitions of satire [8].

We have also observed that humor can make the news more memorable. Another notable insight we gained, is the role of user attitudes towards AI in mediating these effects. We observed how individuals with higher trust in AI were significantly more likely to prefer the re-framed satirical versions. In direct contrast, those more skeptical towards AI tended to be less receptive to the manipulated content. Our interpretation for this would be that those who report more trust in AI are more open to its more creative and non-traditional outputs, while the skeptics might question the quality and the trustworthiness. This does highlight that beyond the content itself, audience bias towards AI will have an effect on the overall success of similar applications of LLMs on online news content. These insights show us how applications of similar services will need to take into account the individual differences of users, while some might embrace this untraditional way of receiving news content, others might not.

Another interesting finding was how our transparency condition about the use of AI, was in most cases found to not have an effect. In both studies, clearly disclosing that the summaries was re-framed by AI to be more satirical had minimal effect of any of our main measurement dimensions. These results to some extent counters earlier research, which found that when the readers know AI is involved in producing news, their trust in the content often drops [32]. One study reported that simply suspecting that the content was generated by AI lowered readers' credibility judgments [32]. However, in our context, informing participants that the content was "re-framed to be more satirical by a large language model such as ChatGPT" has not significantly affected their engagement or trust. In a similar vein, we have also found that the transparency condition has had no effect on the users' reported humor, also contradicting earlier research that found that users reported jokes and humorous content to be less funny when they were told that it was created by AI [12].

One possible explanation for this could be the nature of our sample and content. Participants sourced online could perhaps be more accustomed to AI in their daily lives, as it is in the time the studies were conducted highly relevant and spoken about. This could lead to the AI-disclosure tag not being as alarming as originally anticipated. The satirical tone of the articles might also have signaled some form of departure from traditional journalism, effectively "resetting" expectations to some degree. Another potential impact factor could be the fact that the LLMs role was only to amplify the entertainment value, not to deliver hard facts on its own, maybe making users less critical of the use of AI. Lastly, it should be noted that perceptions are prone to evolving over time. The use of AI and LLMs has become much more common for the average user in recent years. As the field evolves at a staggering speed, it is natural to expect corresponding increase in users trust and acceptance, a trend our results indicate may already have begun within certain groups.

5.1. Limitations

We would like to discuss a few limitations to our study. First and foremost, our participant sample mostly consisted of active news consumers, while very persistent news avoiders were underrepresented. We encouraged participation by news avoiders in the second study by explicitly mentioning this in the crowdsourcing invitation, but this has likely only made a small difference. However, the fact that low-frequency news readers in our sample found the satire to be especially engaging illuminates the

positive effects of satire, but it could be that this strategy is less effective among even less frequent news consumers (i.e., less than once a week).

Another important issue is the factual accuracy of the manipulations. While we explicitly instructed the LLM to maintain “every factual detail”, we lacked a formal method to verify that every important detail was in fact preserved and no misinformation slipped in. Satire, by its nature, exaggerates and twists aspects of reality to raise the entertainment value [8]. Although we have not found any substantial errors in the manual review of the content, a more systematic audit would have been useful to further validate the veracity of the re-framed content.

5.2. Future Work

To validate our results, we would propose a study where users can actively choose to consume ‘regular’ or satirical news. We envision a news recommender website where users can choose between reading an original news article or a satirical version, and even switch back and forth. Multiple news recommendations like this could be combined into a single interface, presenting a more realistic, out-of-lab scenario. Ultimately, determining whether news avoiders would like to remain on such a news website is a long-term objective.

Acknowledgments

This research was supported by the Research Council of Norway with funding to MediaFutures: Research Centre for Responsible Media Technology and Innovation, through the Centre for Research-based Innovation scheme, project number 309339.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication’s content.

References

- [1] A. Blekesaune, E. Elvestad, T. Aalberg, Tuning out the world of news and current affairs—an empirical study of europe’s disconnected citizens, *European Sociological Review* 28 (2012) 110–126.
- [2] M. Skovsgaard, K. C. Andersen, Conceptualizing news avoidance: Towards a shared understanding of different causes and potential solutions, *Journalism Studies* 21 (2020) 459–476. doi:10.1080/1461670X.2020.1812002.
- [3] A. Kalogeropoulos, News avoidance, in: Reuters Institute Digital News Report 2017, Reuters Institute for the Study of Journalism, University of Oxford, 2017.
- [4] H. Coster, More people are avoiding the news, and trusting it less, report says, *Reuters News* (2022). Published June 14, 2022.
- [5] M. Villi, T. Aharoni, K. Tenenboim-Weinblatt, P. J. Boczkowski, K. Hayashi, E. Mitchelstein, A. Tanaka, N. Kligler-Vilenchik, Taking a break from news: A five-nation study of news avoidance in the digital era, *Digital Journalism* 10 (2022) 148–164.
- [6] F. Comba, A. Starke, Climate change or catastrophe? examining the use of fear appeals by climate scientists, *Tijdschrift voor Communicatiewetenschap* 52 (2024) 408–433.
- [7] L. Feldman, P. S. Hart, Is there any hope? how climate change news imagery and text influence audience emotions and support for climate mitigation policies, *Risk Analysis* 38 (2018) 585–602.
- [8] D. Declercq, A definition of satire (and why a definition matters), *The Journal of Aesthetics and Art Criticism* 76 (2018) 319–330.

- [9] V. L. Rubin, N. Conroy, Y. Chen, S. Cornwell, Fake news or truth? using satirical cues to detect potentially misleading news, in: *Proceedings of the Second Workshop on Computational Approaches to Deception Detection*, 2016, pp. 7–17.
- [10] N. J. Stroud, A. Muddiman, Selective exposure, tolerance, and satirical news, *International Journal of Public Opinion Research* 25 (2013) 271–290.
- [11] J. H. Jeng, G. Kasangu, A. Starke, K. M. A. Seddik, C. Trattner, The role of gpt as an adaptive technology in climate change journalism, in: *Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, 2025, pp. 214–223.
- [12] N. N. Joshi, Evaluating human perception and bias in ai-generated humor, in: *Proceedings of the 1st Workshop on Computational Humor (CHum)*, Association for Computational Linguistics, 2025, pp. 79–87.
- [13] D. Gorenz, N. Schwarz, How funny is chatgpt? a comparison of human- and a.i.-produced jokes, *PLOS ONE* 19 (2024) e0305364. doi:10.1371/journal.pone.0305364.
- [14] R. Vliegthart, K. de Bruin, Y. de Haan, S. Kruikemeier, The news avoidance paradox? exploring the relationship between news repertoires and intentional news avoidance, *Communications* (2025). doi:10.1515/commun-2024-0062, published online March 15, 2025.
- [15] A. B. Becker, L. Bode, Satire as a source for learning? the differential impact of news versus satire exposure on net neutrality knowledge gain, *Information, Communication & Society* 21 (2018) 612–625.
- [16] C. Burgers, B. Brugman, How satirical news impacts affective responses, learning, and persuasion: A three-level random-effects meta-analysis, *Communication Research* 49 (2022) 966–993.
- [17] M. Karimi, D. Jannach, M. Jugovac, News recommender systems—survey and roads ahead, *Information Processing & Management* 54 (2018) 1203–1227.
- [18] S. Raza, C. Ding, News recommender system: a review of recent progress, challenges, and opportunities, *Artificial Intelligence Review* (2022) 1–52.
- [19] X. Meng, H. Huo, X. Zhang, W. Wang, J. Zhu, A survey of personalized news recommendation, *Data Science and Engineering* 8 (2023) 396–416.
- [20] N. Helberger, On the democratic role of news recommenders, in: *Algorithms, Automation, and News*, Routledge, 2021, pp. 14–33.
- [21] S. Vrijenhoek, L. Michiels, J. Kruse, A. Starke, N. Tintarev, J. Viader Guerrero, Normalize: The first workshop on normative design and evaluation of recommender systems, in: *Proceedings of the 17th ACM Conference on Recommender Systems*, 2023, pp. 1252–1254.
- [22] S. Vrijenhoek, G. Bénédict, M. Gutierrez Granada, D. Odijk, Radio*—an introduction to measuring normative diversity in news recommendations, *ACM Transactions on Recommender Systems* 3 (2024) 1–29.
- [23] S. Kim, L. B. Chilton, Ai humor generation: Cognitive, social and creative skills for effective humor, *arXiv preprint arXiv:2502.07981* (2025). URL: <https://arxiv.org/abs/2502.07981>.
- [24] S. O. Larsen, Developing and Comparing Similarity Functions for the News Recommender Domain Using Human Judgments, Master’s thesis, University of Bergen, 2024. URL: <https://bora.uib.no/bora-xmlui/bitstream/handle/1956/23059/Master-thesis-final-sebastian-overhaug-larsen.pdf?sequence=1&isAllowed=y>, available via BORA (University of Bergen Open Research Archive).
- [25] R. Misra, News category dataset, *arXiv preprint arXiv:2209.11429v3*, 2022. doi:10.48550/arXiv.2209.11429, submitted September 23, 2022; revised October 6, 2022.
- [26] N. Newman, R. Fletcher, A. Schulz, S. Andi, C. T. Robertson, R. K. Nielsen, News readers really want to read: Relevance works—news audiences in an era of choice, <https://www.digitalnewsreport.org/publications/2019/news-readers-really-want-read-relevance-works-news-audiences/>, 2019.
- [27] A. D. Starke, V. R. Solberg, S. Øverhaug, C. Trattner, Examining the merits of feature-specific similarity functions in the news domain using human judgments, *User Modeling and User-Adapted Interaction* 34 (2024) 995–1042.
- [28] Ground News, Methodology: Bias ratings & factuality scores, <https://ground.news/rating-system#>, 2025. Retrieved June 25, 2025.
- [29] NewsCatcher API, Product overview, <https://www.newscatcherapi.com/product-overview>, 2025.

Web page. Retrieved June 25, 2025.

- [30] M. A. Xenos, A. B. Becker, Moments of zen: Effects of the daily show on information seeking and political learning, *Political Communication* 26 (2009) 317–332.
- [31] J. C. Coronel, M. B. O'Donnell, P. Pandey, M. X. Delli Carpini, E. B. Falk, Political humor, sharing, and remembering: Insights from neuroimaging, *Journal of Communication* 71 (2021) 129–161. doi:10.1093/joc/jqaa041.
- [32] H. Jia, A. Appelman, M. Wu, S. Bien-Aimé, News bylines and perceived ai authorship: Effects on source and message credibility, *Computers in Human Behavior: Artificial Humans* 2 (2024) 100093. doi:10.1016/j.chbah.2024.100093.