

The Attention–Information Tradeoff

Marta Serra-Garcia*

This version: December 2025

Abstract

How does information transmission change when it requires attracting the attention of receivers? This paper combines an experiment that varies freelance professionals' incentives to attract attention about scientific findings, with several online experiments that exogenously expose receivers to the content created. Attention incentives lead to significantly less information being transmitted, but not more factually inaccurate content. These incentives increase information demand and the knowledge of interested receivers. However, among the majority of receivers who do not demand more information, attention incentives lower knowledge and increase biases in beliefs, revealing that *missing* information can be a channel through which misperceptions arise.

JEL Classification: D83, D91, C72, C91.

Keywords: Attention, incentives, information, experiment.

*Serra-Garcia: Rady School of Management, UC San Diego and CESifo (email: mserragarcia@ucsd.edu.). This research was conducted under IRBs #806238 and #161827, at UC San Diego. This paper has benefited from highly insightful comments from Jose Apesteguia, Russell Golman, Ian Krajbich, Sam Lindquist, George Loewenstein, Michel Marechal, Kirby Nielsen, John-Henry Pezzutto, Sally Sadoff, Joel Sobel, Hannah Trachtman, Isabel Trevino, Manu Vespa, as well as several seminar and conference participants. Wenxuan Cao, Haowei Li, Siqi Li, Wanting Zhou, and Jerry Zhu provided excellent research assistance.

“Knowledge isn’t free. You have to pay attention.”

Richard P. Feynman

1 Introduction

Attention is a scarce cognitive resource that is fundamental to decision making in many settings (e.g., Simon, 1971; Kahneman, 1973; Krajbich et al., 2010; Bordalo et al., 2022; Loewenstein and Wojtowicz, 2023). One such setting is the Internet, where an important goal of stakeholders is to attract individuals’ attention. The Internet has been portrayed as “a battleground for clicks and eyeballs” (Wu, 2017) among content producers, who seek clicks and the sharing of the information they post. The incentive to grab attention has raised concerns because of its potential impact on the accuracy of the information that citizens are exposed to and pay attention to online (Pariser, 2011). Given these concerns, it is critical to carefully examine the relationship between attention and information transmission.

Important existing work has studied what kind of content attracts individuals’ attention in online settings. It has shown that emotional and arousing content can receive more clicks and shares (e.g., Berger and Milkman, 2012; Serra-Garcia and Gneezy, 2021; Qiu and Golman, 2022; Robertson et al., 2023), but also documented that predicting attention is difficult (Bakshy et al., 2011; Banerjee and Urminsky, 2023).

This paper contributes novel evidence by taking a step back and asking: How does the incentive to grab attention affect the content that is generated and, subsequently, the information that is transmitted? In an experiment that exogenously varies the incentives of freelance professional writers, the paper first studies how incentives to grab attention affect the supply of content, studying the content’s accuracy, focus, and style. Then, leveraging online experiments that control individuals’ exposure to content, the paper examines how content generated with incentives to grab attention affects individuals’ knowledge, attention allocation and subsequent beliefs.

The paper studies information transmission about scientific findings, with a focus on four topics: cancer, sleep, vaccines, and climate. Information about these topics can influence important beliefs and behaviors, such as immunization choices or actions against climate change. But information about them can be false, incomplete, and misleading (e.g., Chou

et al., 2018; Allen et al., 2024), and spread virally online (e.g., Vosoughi et al., 2018).

How attention incentives affect content and receiver behavior is ex-ante unclear. First, these incentives could distort the information transmitted or reduce the amount of information transmitted. Distortion could occur if, to generate curiosity among readers, the content generated includes more false information about scientific findings (e.g., exaggeration). Alternatively, a reduction in the amount of information could arise if the content highlights what readers do not yet know and introduces information gaps (Loewenstein, 1994; Golman and Loewenstein, 2018). Information distortions or selective disclosures could in turn bias receivers’ beliefs regarding the implications of the findings. Second, a scarcely examined effect of the incentive to grab attention is to increase readers’ attention, which leads to the consumption of more information. This compensating effect of increased information demand is central to the attention-information tradeoff.

To study the relationship between attention and information, the experimental paradigms in this paper are designed to align with important features of online attention markets, while preserving experimental control over the incentives of content producers and over the exposure to content of those receiving information. In the pre-registered Sender experiment freelance professionals who specialize in writing social media copy or research summaries write a total of 595 summaries about recent research findings. These professionals receive written materials about four research articles and write summaries, consisting of a headline and a three-sentence description, with the incentive to inform receivers (Info-Incentives), or to attract the attention of receivers in one of two ways. The first is to have receivers demand more information, by clicking on a button that delivers detailed information about the research (Clickbait-Incentives). The second is to create a summary which receivers believe is likely to be shared virally on social media, generating the “eyeballs” or views of many others (Viral-Incentives). These two incentives focus on fundamentally different forms of attention, both of which are often discussed in the context of online communication.

To measure the impact of incentives on the content produced by senders, we adapt the method developed for political communication in Carlson (2018) to science communication. We provide senders with materials about several research articles. These articles have a consistent structure (introduction, methodology, findings, conclusions; see e.g., Schimel, 2012),

and the materials serve as a benchmark to establish accuracy in information transmission. Content is classified according to three dimensions: its accuracy, relative to the materials; its focus, whether it relates to the introduction, methodology, findings, or conclusion of the research; and its writing style, by quantifying changes in summary length, readability, sentiment and punctuation use.

The Sender experiment reveals three main effects of Clickbait- and Viral-Incentives relative to Info-Incentives. First, Clickbait- and Viral-Incentives significantly reduce the amount of accurate information in summaries. The amount of inaccurate information, that wrongly describes the research, does not change. Rather, the amount of new content added by professionals, often to motivate and make the research relatable to readers, increases. Second, summaries increase their focus on introducing the research and reduce the amount of information about the methodology used. Third, summaries become more readable, shorter, with more than double the amount of exclamation and question marks.

In exploratory analyses, we examine *how* content changes. Clickbait- and Viral-Incentives lead to content that is more likely to directly appeal to readers (using expressions such as “you will learn how to sleep better”). With Viral-Incentives, we also observe a marginally significant increase in incorrect implications being drawn in the summary, claims that are not mentioned or follow from the materials (e.g., the reader should change their diet, although the materials note the findings are correlational).

The findings in the Sender experiment suggest that attention incentives lead to a reduction in the amount of information about the research, and to the introduction of new content that makes the research relevant to readers and could lead readers to derive incorrect impressions about its implications for behavior. An important limitation is that naturally occurring communication about research in online settings could differ from that within our experiment, in which professionals are hired to promote scientific findings for a broad audience. Among others, professionals may not feel as free to introduce distortions as individuals on their personal social media accounts. To shed light on the validity of the context, we explore the features of headlines in online content about the research articles. While the incentives of news, blogs, and social media content producers are heterogeneous and difficult to identify, we study how they describe the research and compare the distributions of

features of headlines.

We find that the headlines written by professionals align in accuracy, focus, and writing style with those observed online. For example, the online headlines of news and blogs about the research articles are rarely classified as inaccurate relative to the research materials, but often as new content that relates to the topic more broadly, which is consistent with the findings for the headlines written by professionals.

We examine the impact of senders' content using pre-registered online Receiver experiments. We measure receiver knowledge, attention allocation, and beliefs. We exogenously assign receivers to summaries written under different incentives and assess knowledge by asking the same multiple-choice questions about the research articles. Receivers are either only exposed to summaries or can click on summaries to learn more about the research. When receivers click to learn more, they receive the research materials that senders had and answer knowledge questions of which they had no prior awareness. We also vary whether receivers are asked to recall the information they received or have it available when measuring their knowledge, to capture whether incentives to grab attention affect recall (e.g., Zimmermann, 2020; Graeber et al., 2024a).

We document three main results. First, when receivers are only exposed to summaries, those written with Clickbait- and Viral-Incentives lead to significantly lower receiver knowledge (6-7 percentage points, or 10-11%). Simultaneously, Clickbait- and Viral-Incentives increase click rates significantly, from 20% to 25% and 24%, respectively, which lead interested receivers to be exposed to more information about the research and thereby achieve significantly higher knowledge. This illustrates the presence of a trade-off: less information in the summaries but higher click rates. These compensatory effects result in small losses in overall knowledge in our experiments, but could result in overall gains or losses in knowledge depending on the context.

Second, we document that recall may be an additional channel through which research information is “lost,” which reduces the impact of the incentive to grab attention. When receivers must recall the information they received, they make significantly more mistakes, especially about the sample and methodology, across all treatments, and the effect of the incentive to grab attention becomes significantly smaller.

Third, we focus on the large majority of receivers who were only exposed to summaries about the research. An open question is whether summaries written with Clickbait- and Viral-Incentives are more likely to induce biased beliefs about the research and its implications. We address this question in two ways. First, we explore receivers' mistakes about the sample involved in the study: They are more likely to erroneously think that the sample is representative of the US population when summaries were written with Clickbait- and Viral-Incentives. Second, we conduct an experiment that measures receivers' beliefs about doctors' official recommendations and their intentions to suggest behavioral changes to others and apply them to themselves. The data show that receivers exposed to summaries written with Clickbait-Incentives are more likely to believe doctors' official recommendations align closely with the research findings, although in four out of five cases the findings were suggestive and not officially recommended. Receivers are also more likely to self-report intended behavioral changes based on the research summaries. The findings with Viral-Incentives are directionally similar, but not statistically significant.

Taken together, these results reveal that Clickbait- and Viral-Incentives reduce the amount of information about research findings. While they increase information demand, they can leave receivers who do not demand information with incorrect impressions, thinking that the research is more applicable to them than it actually is.

This paper contributes to the existing literature on information transmission (e.g., Crawford and Sobel, 1982) and a growing body of theoretical work studying limited attention (e.g., Bordalo et al., 2016; Galperti and Trevino, 2020; Chen and Suen, 2023), by complementing existing experimental work that studies how attention scarcity affects decision-making and tests models of rational inattention (e.g., Caplin et al., 2011; Dean and Neligh, 2023) and choice in complex environments (e.g., Kang et al., 2009; Smith and Krajbich, 2018; Abeler et al., 2023; Esponda et al., 2023; Guan et al., 2023). By focusing on the incentive to attract the attention of receivers, the paper complements prior research using sender-receiver games, in which receivers' attention is assumed and senders are incentivized to influence the beliefs and actions of receivers (e.g., Sanchez-Pages and Vorsatz, 2007; Abeler et al., 2019; Deversi et al., 2021; Jin et al., 2021; Thaler, 2024), to study incentives such as persuasion and reputation concerns.

Growing concerns regarding the spread of misinformation – intentionally and unintentionally false or misleading information–, make it important to better understand the fundamental drivers of online behavior. Related work has shown that individuals share false headlines, often inadvertently (e.g., Osmundsen et al., 2021; Pennycook and Rand, 2021; Roozenbeek et al., 2022; Serra-Garcia and Gneezy, 2021; Altay et al, 2022; Acemoglu et al., 2024), and that suppliers of content exhibit an intrinsic desire for clicks and attention (Srivanasan, 2023). The findings in this paper show that incentives to attract attention need not lead to more inaccurate content when communicating about research findings. But, they decrease the amount of detailed information about the research and bring in their broader relevance, which can leave receivers with inaccurate beliefs about its implications for their behavior. This *missing* information about the research could be harmful and may represent a new channel through which misperceptions arise.

2 Experimental Design

The experimental design consists of a Sender experiment and several Receiver experiments.

2.1 Sender Experiment

2.1.1 Setting

The Sender experiment was conducted with freelance professionals, who participated in the role of senders and were asked to submit four summaries of a research article. The freelance professionals were recruited on Fiverr, a large online platform for freelance services with over 4 million buyers in 2022. We recruited 149 freelancers specialized in writing services, within “social media copy” and “research and summaries,” and collected a total of 595 summaries (Serra-Garcia, 2025). The recruitment was done in two waves (108 freelancers in January-February 2023 and 41 in February 2024, due to the recruitment process hitting limits to the availability of freelancers within budget during the first wave). Since there were no differential treatment effects by wave of data collection, the data for both waves is pooled

and an indicator for wave is included in the empirical analysis.¹

Emulating a typical writing task, freelancers were first contacted about their interest in writing about cancer, sleep, climate and vaccine research, and offered a payment based on their posted prices on the platform. Interested freelancers started the task and received detailed instructions, which depended on the treatment (described in Section 2.1.2). On average, they received \$19.1 ($SD = \4.5), with \$20 being the median and most frequent payment (56% of cases). In what follows, the freelancers are referred to as “professionals.”

Professionals’ public profiles (available for 138 of them) reveal that 90% had experience in writing, with 67% mentioning experience in work to attract attention, and 56% describe working in the digital domain.² These writers are mainly focused on writing anonymous online person-to-person or business-to-person content. They are not journalists and reputational incentives or outlet-focused incentives (e.g., Gentzkow and Shapiro, 2008) are not central to them. Their typical incentive is to write appealing and attention-grabbing content, which is the focus of this paper.

Since Large Language Models (LLMs) could produce content like professionals, in a complementary experiment we use an LLM (GPT-3.5) that acts as sender to generate 360 summaries with different objectives. The resulting content and its consequences on receiver beliefs and attention are qualitatively similar to those documented in the main experiment with professionals, as discussed in Section 3.3 and Online Appendix F.1.

2.1.2 Treatments

The experiment consisted of three between-subjects treatments. The wording of the instructions and incentives in each treatment were chosen to align closely to wording used on the platform (instructions are provided in Online Appendix Section A.2). In all treatments, professionals were told that the goal of the task was to inform the general public about

¹One sender wrongly submitted the same summary for two articles. Detailed information on the implementation, recruitment waves, instructions, and materials given is provided in Online Appendix A. Freelancers were restricted to be based in the US, Canada, and the UK.

²For example, one of their profiles stated “*I’ve written over 1000 articles in my time as a writing specialist, and I can help you find the right, creative words to keep your readers engaged.*” The content analysis was done by two independent coders. In addition to expertise categorization, they coded the reported experience in writing. 50% of profiles report experience and 35% report an exact number in years. The average experience reported is 7.5 years ($SD=4.1$).

research findings, which may be posted on social media. They also received explicit goals and financial incentives, as follows:

- In the *Info-Incentives* treatment, senders were asked to “*provide accurate information to potential readers*”. They received an additional \$5, described as “*I’ll ask someone else to read one of your summaries and if he/she answers several questions about the research correctly, as a pre-test, I’ll tip you.*”
- In the *Clickbait-Incentives* treatment, senders were asked to “*stimulate readers to read more about these research articles*”. They received an additional \$5, described as “*I’ll also ask someone else to read one of your summaries and several others. If he/she chooses to read more about the research you summarized, as a pre-test, I’ll tip you.*”
- In the *Viral-Incentives* treatment, senders were asked to “*maximize the number of times ... (their writing) is shared.*” They received an additional \$5, described as “*I’ll also ask someone else to read one of your summaries and others that I am collecting and if he/she believes yours would be one that is shared most often, as a pre-test, I’ll tip you.*” Beliefs rather than actual sharing were chosen to avoid the spread of misleading or incorrect information. Sharing beliefs are in line with online sharing as discussed in Online Appendix E.5 (see also, Mosleh et al., 2020).³

When initially contacted, senders were told that summaries would consist of approximately 70-100 words. Upon starting the task, they received 1-2 pages about each research article which included, at the bottom, two spaces labeled “Headline (with a maximum of 140 characters)” and “3-sentence summary.” These guidelines provided clear expectations but longer summaries were approved. Consistent with senders understanding this flexibility, the length of summaries varies significantly across treatments. Posts on social media platforms also vary significantly in length, depending on the platform. Summaries in this study are comparable in length (96 words, on average) to those in LinkedIn (50-100 words) and shorter than a typical Twitter thread (120-180 words; Pitzalis, 2022; LinkedIn, 2023).

³After the experiment, a subset of summaries was posted on Twitter, to act on what senders were told. The subset of summaries was selected carefully such that the summaries would not contain any incorrect information and reflected closely the materials provided to the senders. The posts were also accompanied by a link to the original research article, and the accounts informed users that the content posted was part of a research study.

2.1.3 Articles

The research articles were selected among articles featured in the New York Times in 2019, to feature topics, such as public health or climate, that could interest a broad audience. They were described to the senders via materials that had three elements: the abstract, containing specialist language; a press release, written for a broad audience to communicate with journalists and the public (e.g., National Academies of Sciences, Engineering, and Medicine, 2017; Peters, 2020);⁴ and a link to the entire paper. Figure 1 summarizes the eight articles included in the experiment.

A potential concern in asking senders to summarize research relevant to public health behavior is that they could incorrectly report about it, and this could lead to harmful behavior among readers. To minimize this risk, the articles focused on how a healthy diet is related to lower skin cancer risk (Kim et al., 2019), potential cancer treatments (Lee et al., 2019), and reduced insomnia (Gangwisch et al., 2020); the benefits and safety of the MMR vaccine (Mina et al., 2019 and Hviid et al., 2019); and the importance of reducing deforestation and of wildlife protection (Maxwell et al., 2019, and Berzaghi et al., 2019). After a job started, senders received the materials, including a brief consent form, which informed them that their jobs were part of a research study. If this information increased concern for accuracy, the measured impact of Clickbait- and Viral-Incentives could represent a lower bound in the types of inaccuracies that may be expressed in other settings, which motivates the comparison to naturally occurring online mentions of the articles (in Section 3.2). Further information about the research articles, an example of the sender’s materials about the research, and an example of professionals’ summaries are provided in Online Appendices A.4 and A.5.

2.1.4 Outcomes of Interest

We measure the impact of incentives on senders’ writing across three sets of outcomes – accuracy, focus, and style – using text-based analyses. These outcomes were pre-registered,

⁴Existing work has descriptively documented “spin” in such communication (e.g., Yavchitz et al., 2012; Sumner et al., 2016), although it is also associated with better knowledge, in the context of cancer (e.g. Stryker et al., 2008). In this paper we exogenously vary senders’ incentives to identify how attention incentives affect content, keeping the materials used by senders constant.

Figure 1: Articles

Topic	Reference	Main finding	Sample	Analysis	Conclusions	
1	Cancer	Kim et al. (2019)	Individuals with diets high in vitamin A have a 17% lower risk of squamous cell carcinoma (SCC), a form of skin cancer, than those with low vitamin A consumption.	Nurses' Health (N=121,700) & Health Professionals Follow-Up (N=51,529) Studies.	Cohort study, longitudinal	The study authors suggest that the next step would be to conduct a clinical trial to test whether higher intake of vitamin A reduces the risk of SCC.
2	Cancer	Lee et al. (2019)	A molecule (I3C) in cruciferous vegetables reduces tumor growth because it inactivates a gene found in many cancers, the WWP1 gene.	Cancer-prone mice and human cells	Laboratory study	Individuals should not change their diets yet, as too large amounts of cruciferous vegetables (over 6 lbs a day) would be needed to prevent cancer.
3	Sleep	Gangwisch et al. (2020)	A diet high in glycemic index is associated with higher odds of insomnia among postmenopausal women.	Women's Health Initiative Observational Study, baseline N=77,869 & follow-up N=53,069, on postmenopausal women	Cohort study, longitudinal	The authors suspect that reducing refined carbohydrates in one's diet may be a low-cost way of preventing insomnia, but RCTs are needed.
4	Sleep	Perrault et al. (2019)	Participants fell asleep faster and had better memory the next day when sleeping on a rocking bed compared to a stationary bed.	18 healthy young adults	Laboratory study, within-subjects	Rocking is important for sleep quality and memory consolidation.
5	Vaccine	Mina et al. (2019)	Measles leads unvaccinated children to lose up to 73% of their protective antibodies for other diseases.	77 unvaccinated children in a community in the Netherlands	Event study, using VirScan to test blood samples	Measles vaccination is important to protect children from measles and other diseases.
6	Vaccine	Hviid et al. (2019)	The MMR vaccine is not associated with autism.	Danish population registry, N=657,461 children born in 1999-2010	Cohort study, longitudinal	Showing no association between the MMR vaccine and autism is important due to increase in vaccine hesitancy and measles outbreaks in Europe.
7	Climate	Maxwell et al. (2019)	Deforestation of intact tropical forests has a 626% stronger impact on carbon emissions than previously estimated.	Intact forest loss from 2000 to 2013	Computational	Intact forest retention needs to receive more attention and more funding from initiative such as REDD+.
8	Climate	Berzaghi et al. (2019)	Forest elephants affect tree growth and aboveground biomass, and thereby affect the carbon emissions of forests.	Field measurements in the Congo Basin in forests affected by elephants and in forests in which they had been eradicated	Field measures and computer simulations	Forest elephants should receive more protection and expand to help fight climate change.

and we add an index to summarize each set (not pre-registered) following Anderson (2008). After examining the impacts on writing, we further examine the impact of incentives on receiver knowledge and behavior in the Receiver Experiments (Section 2.2).⁵ Throughout, to correct for multiple hypotheses testing, we use sharpened FDR-adjusted q -values, as in Benjamini et al. (2006), Anderson (2008), Braghieri (2024) and Graeber et al. (2024c),

⁵Pre-registered and additional analyses are described in Online Appendix C for the Sender Experiment and within Online Appendix E for the Receiver Experiments.

applied to each set of outcomes.⁶

To analyze accuracy and focus, we build on the approach by Carlson (2018), developed to study political communication. First, each summary is divided into units of information. A unit of information is a statement that conveys a single, identifiable piece of information. It is often a sentence or a clause within the summary. Then, the content of the unit is separately analyzed with regards to its accuracy and its focus.⁷

Accuracy. Establishing accuracy in naturally-occurring communication is challenging, as it is difficult to identify the exact ground truth behind every statement. We measure accuracy by comparing senders’ writing to the materials. Units are classified as:

1. Accurate-In text: information from the materials that is correctly described.
2. Inaccurate-In text: inaccurate information of two kinds: wrong, i.e., explicitly incorrectly reporting what is in the materials, such as stating causation when correlation is described, or incomplete, leaving out important information. We pool both types of inaccuracy into “Inaccurate-In text,” as they were observed in less than 10% of the cases.
3. Not in text: information that is not provided in the materials.⁸

We construct an index to represent summary accuracy (*Accuracy index*), in which the number of Accurate-In text units increase the index, while the variables capturing the number of Inaccurate-In text and Not in text units decrease the index value.

⁶Throughout, sharpened FDR q -values are computed using the two-stage BKY (2006) procedure (Anderson 2008). Each q -value is the smallest false-discovery-rate threshold q at which the coefficient is rejected. Unlike adjusted p -values, these q -values are FDR-rejection thresholds and may be smaller than these for p -values below 0.15 (Braghieri, 2024).

⁷The analyses were conducted by independent coders, who were blind to senders’ incentives. Each summary was analyzed by two coders who discussed and reached an agreement on their coding. Coders’ instructions are shown in Online Appendix G. To verify the robustness of the coding results, GPT was used for coding. It provided qualitatively the same results, shown in Online Appendix F.2.

⁸For example, we can consider the article by Hviid et al. (2019), which shows a lack of correlation between the MMR vaccine and autism using registry data in Denmark from children born between 1999-2010. A unit of information stated “this researcher diagnosed children with autism,” which is classified as wrong because researchers did not diagnose autism. Another unit stated “analyzed data from over 650,000 kids,” which is classified as incomplete, because it lacks the country and time range during which the data were collected. Finally, a unit stating “You must have heard of rumors about vaccines causing autism...” is Not in text, because the materials did not refer to any rumors.

Focus. Each unit is classified (Schimel, 2012; Munch, 2023) as: (1) Introduction, motivates the research and introduces the research question; (2) Method, describes the sample used and the analysis approach; (3) Finding, reports the findings; (4) Conclusion, describes the research implications and the authors’ conclusions; and (5) Call to Action (CTA), encourages the reader to click, read, or share. We summarize the focus of a summary using an index that combines the number of units in each category (*Focus index*). Because the index aims to capture the amount of detailed information about the research within the summary, the number of Method and Finding units increase its value. The number of Introduction, Conclusion and CTA units, which provide motivation and context, decrease its value.

Writing Style. The third main outcome is writing style, which measures *how* the summary is written. It consists of four elements of style: (1) length, measured by word count; (2) readability, based on Flesch’s Reading Ease Score (Flesch 1948), according to which high scores (90-100) indicate that the text is very easy to read; (3) sentiment, which captures the tone of text (positive, neutral, or negative), using sentimentr (Rinker, 2023); and (4) use of punctuation, which is an indicator variable if the summary uses exclamation marks, question marks, ellipsis, hashtags as well as icons (such as emojis).⁹

We construct an index that combines these style variables (*Style index*). A higher index value indicates a more attention grabbing style, with higher reading score and punctuation increasing the value of the index, and higher word count and more positive sentiment decreasing it (indicated as “x (-1)”).

We also use these measures to analyze naturally occurring online mentions of the research. For the Altmetric data, we use GPT to code accuracy and focus of headlines, after validating it with human coding (as described in Online Appendix D). We use the same measures of style as above.

In exploratory analyses, we examine four additional outcomes that help in interpreting the main effects. First, we explore whether the summary mentions the sample used in the research article and if so, how accurately it is reported. Then, we explore whether

⁹Additional analyses that measure the effects of incentives on the presence of each type of punctuation mark, and on the count of each one, are shown in Online Appendix C.2.

the summary directly appeals to readers, with words such as “you”, and whether it contains implications for receivers’ behavior and their accuracy. Online Appendix C.3 provides details.

2.1.5 Research Questions

The impacts of Clickbait- and Viral-Incentives on the outcomes of interest vary depending on what motivates clicks and sharing. Clicks are a form of information demand that stems from igniting an individual’s curiosity (Loewenstein, 1994), while sharing can be driven by a range of motives (see, e.g., Metzger et al., 2021; Osmundsen et al., 2021). For example, people may share altruistically to benefit others (e.g., Berger and Milkman, 2012), to signal their identity to a wider audience (Barasch and Berger, 2014), to build their own reputation and recognition online (e.g., Toubia and Stephen, 2013; Smirnova et al., 2022; Niitsuma et al., 2025), or to ignite others’ emotions (e.g., Brady et al., 2017).

To induce clicks and sharing, senders could write summaries that violate expectations, providing information that deviates from what the reader knows or expects (Hebb, 1949; Loewenstein, 1994). Then, both Clickbait- and Viral-Incentives could lead senders to introduce distortion, by writing summaries with more Inaccurate-In text units (e.g., exaggerations). Or, they could lead to selection in the information reported, by increasing the number of (surprising) Findings and Conclusion units, with more negative emotions to arouse readers (e.g., Berger and Milkman, 2012), but not more inaccuracy.

Alternatively, Clickbait-Incentives could have different consequences than Viral-Incentives. Senders with Clickbait-Incentives could choose to create information gaps that make salient what the reader does not know, as proposed in a model by Golman and Loewenstein (2018).¹⁰ Readers would then click to gain information that closes the information gap. In that case, Clickbait-Incentives could lead to summaries that contain fewer units with detailed information about the research, i.e., fewer Accurate-In text units, and fewer Findings and Conclusion units, but more Introduction units to pose questions but not provide answers.

In both cases, we would expect Clickbait- and Viral-Incentives to lead senders to focus on the broader relevance of the research, writing about their implications for readers, making them easy to understand (Kahneman, 1973; Clark, 2008), with higher readability scores and

¹⁰This model was tested experimentally using puzzles and facial recognition tasks in Golman et al. (2022).

fewer (detailed and technical) Methods units, relative to Info-Incentives. By focusing on broader topics and selecting among features of the research findings, Clickbait- and Viral-Incentives could induce biased beliefs in readers regarding how the research was conducted and how it is relevant to them.

The research questions are therefore: Do Clickbait-Incentives lead to “missing information” (information gaps) or “distorted information” (violating expectations)? Do Viral-Incentives lead to “distorted information” or rather selection in reporting, by emphasizing findings and broad implications, without inaccuracy?

2.2 Receivers: Experimental Design

The Receiver experiments measure the effects of senders’ incentives on receivers’ knowledge, attention allocation, and beliefs. Detailed instructions are in Online Appendix B.2.

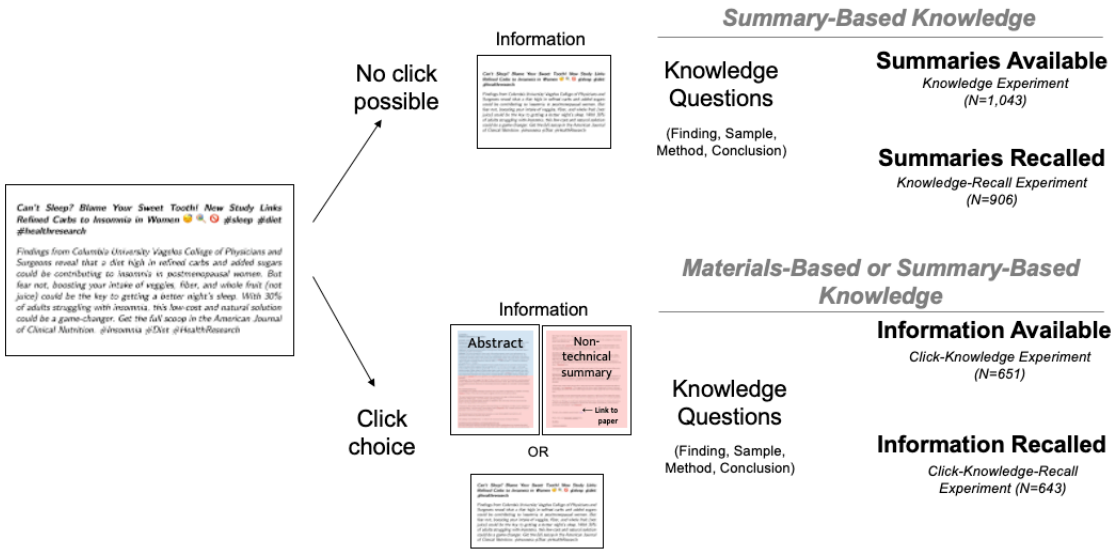
2.2.1 Knowledge and Attention Measurement

The goal of the Knowledge and Click experiments is to measure how the incentive to grab attention affects receiver knowledge and attention. To measure knowledge transmission separately from attention allocation, we implement experiments in which no clicks are possible (Knowledge Experiments) and receivers are only exposed to senders’ summaries, and experiments in which clicks are possible (Click Experiments) and receivers choose to learn more about the research, beyond senders’ summaries.

The incentive to grab attention could affect receiver knowledge over time in two ways. It could lead to content with differential amounts of information (e.g., less information with Clickbait- and Viral-Incentives) or to content that results in differential recall (e.g., higher recall of content with Clickbait- and Viral-Incentives). For that reason, we vary the need to recall the content across the Knowledge and Click experiments, as outlined in Figure 2 (a detailed design chart is shown in Online Appendix B.1).

Knowledge Experiments. In these experiments, receivers are assigned summaries written under one incentive treatment (Info-, Clickbait-, or Viral-Incentives) in the Sender experiment. In the *Knowledge* experiment, they answer knowledge questions having the summaries

Figure 2: Design of Receiver Experiments - Knowledge Measurement



available on the same page. In the *Knowledge-Recall* experiment, they are first exposed to the summaries and, to ensure they pay attention to what they read, they write one sentence about each summary.¹¹ Then, they complete a second part of the experiment, in which, unannounced, they are asked knowledge questions, without having the summaries available.

Click Experiments. In these experiments, receivers are presented with a list of eight summaries, one for each research article (as a list of posts on social media would appear). They can click on any number of articles, from zero to eight, to receive more information. There is no financial incentive to click and receivers are not told what will happen after they receive more information. In this click stage, receivers are either presented with four summaries written with Info-Incentives and four with Clickbait-Incentives, or four summaries written with Info-Incentives and four with Viral-Incentives. This design allows to measure whether summaries written with Clickbait- or Viral-Incentives are clicked on more often, relative to summaries written with Info-Incentives. And, to address selection concerns, the receiver's choice is implemented probabilistically, with a 75% chance. Receivers are aware

¹¹The survey was programmed such that the summaries would appear as images and could not be easily copy-pasted elsewhere, such as into an LLM. We check the content of these sentences and find that the average sentence length is 16.2 words. There is significant variation, with the 10th percentile having 8 words and the 90th percentile having 27 words. We then check for the presence of one keyword per article in each sentence. For example, the keyword is "cancer" if the research article is about cancer. Between 71% and 97% of sentences contain the keyword that closely relates to the research article.

that, with a 25% chance, the computer selects one article at random for the receiver to read about.

In the *Click-Knowledge* experiment, after receivers click, they are shown the materials about the research and asked to answer the knowledge questions, of which they had no prior notice. Both if the click choices are implemented or the computer implements one article, receivers only answer questions about articles they read more about. In the *Click-Knowledge-Recall* experiment, there are three steps after receivers make their click choices. First, depending on their choices and random assignment, receivers receive more information and write one sentence about the article. In the second step, receivers are exposed to the summaries they did not receive more information for, and write one sentence about them (as in the *Knowledge-Recall* experiment). In the third step, without previous announcement and without having any information available, receivers answer the knowledge questions about all articles. Therefore, in the *Click-Knowledge-Recall* experiment, knowledge is based on recall, and it is elicited for all articles (both if the receiver clicked and did not click to read more about them).

Outcomes of Interest. Receiver knowledge is measured as the number of correct answers to multiple-choice questions about four main components of each research article: its main empirical finding, the sample used, its analysis methodology and its conclusions. These components capture the main elements of scientific papers, as described in Munch (2013) and Schimel (2012), excluding the introduction, due to its overlap with conclusions and its lower specificity to the research article.

In all Receiver experiments, receivers were exposed to summaries, and potentially more information, about 8 research articles and incentivized to correctly answer the same 32 questions (receiving \$0.10 per correct answer). We measure the following outcomes:

- Summary-based knowledge, measured after receivers have read only a summary.
- Click rate, an indicator for clicks to learn more after being exposed to a summary.
- Materials-based knowledge, measured after receivers have read 1-2 pages about the research (the same as the senders' materials).

- Overall knowledge, the weighted average of summary-based knowledge, for those only exposed to a summary, and materials-based knowledge, for those who received more information.

These measures capture the potential, countervailing effects of Clickbait- and Viral-Incentives. Attention incentives could reduce summary-based knowledge, based on the effects discussed in Section 2.1.5.¹² Simultaneously, they could also increase click rates and the fraction of people with (higher) materials-based knowledge. Hence, Clickbait- and Viral-Incentives could *increase* or *decrease* overall knowledge, depending on the relative strength of these effects.

An important additional outcome to consider is what are the beliefs of the individuals who do not click and are only exposed to summaries. Summaries written with Clickbait- and Viral-Incentives could not only contain less information, but the information provided could potentially mislead readers, rather than lead to random errors. For that reason, we first consider what *kinds* of mistakes receivers make when asked about the research sample, after being exposed to a summary. We code incorrect answer options that indicate a larger sample than the study actually used for all research articles. For articles on health (cancer, sleep and vaccines), we also code incorrect answer options that indicate the study used a more representative sample, defined as being closer to a sample of adults in the United States (US), than it actually did. The questions asked, all answer options, and how they were coded (if applicable) are shown in Online Appendix B.3.

2.2.2 Beliefs Measurement

We next examine what kinds of impressions summaries written with Clickbait- and Viral-Incentives, relative to Info-Incentives, leave on readers in the *Beliefs & Intentions* experiment. We measure receivers’ beliefs about the implications of the research for what doctors recom-

¹²Additionally, we coded whether the summary clearly identified the correct answer option in the knowledge questions of the 595 summaries, using a human research assistant and GPT (version 3.5). In Online Appendix C.4 we describe this coding and provide descriptive statistics for how frequently summaries contained information that clearly identified the correct answer option in the knowledge questions. We observe similar effects of both Clickbait- and Viral-Incentives on the likelihood that the correct answer is available in the summary, according to each source of coding (human and GPT) separately and combining the coding sources, including an analysis that only considers the cases with coding agreement (in 71% of the cases).

mend, what they would recommend to their family and friends and their own behavior. The research articles studied five behaviors: using rocking beds to improve sleep, consuming a low GI diet to reduce insomnia, eating cruciferous vegetables to reduce the risk of cancer, taking vitamin A to reduce the risk of skin cancer, and vaccinating children with the MMR vaccine. The first four behaviors are based on research that provides first evidence of a relationship or impact, and are *not* official doctor recommendations, while the last one (vaccination) is.

The experiment consisted of two main parts. In the first part, receivers were exposed to eight summaries and wrote a one-sentence summary about them. In the second part, receivers answered two sets of five questions, presented in random order. For each of the five behaviors, they answered yes or no to “is it an official doctor recommendation?” and “would you encourage your family and friends to engage in the behavior?”. They received \$1 if their answer to one randomly selected question regarding official doctor recommendations was correct. Receivers also answered the open-ended question “*How does the information you read about these scientific studies concern you or affect you and your future behavior?*” Two independent research assistants coded the answers as “no effect,” if the receiver stated it would not affect them, “increase interest,” if the receiver indicated they would like to learn more, and “behavior change,” if the receiver indicated they would change their diet or other behaviors. We build an *Intentions index* (not pre-registered), which considers these three measures (the average response to questions about doctor recommendations and encouragement, and an indicator for whether receivers indicated an intended behavior change themselves), using the approach in Anderson (2008).¹³

Receivers were exposed to a subset of 120 sender messages. The same number of summaries was drawn from each treatment in the Sender experiment. We used the knowledge of receivers, obtained in the *Knowledge* experiment, to build two categories based on whether receivers scored above-median and below-median in the knowledge questions. These categories were an additional criterion used for the draw of the 120 summaries. Since we do not

¹³At the end of the experiment, to fully inform receivers about the research, they were presented with the full materials about the research and they could click to access the research article. Receivers were also asked about their trust in scientists and media (“How much do you trust scientists?” and “How much do you trust science that you read about in news articles and on social media?”, with a scale 1 to 4 (from “Do not trust at all” to “Trust completely”) either before or after reading the materials. Trust was not different depending on the order (p -value > 0.05 for trust in media and scientists, see Online Appendix Table E.18).

observe differences in receivers’ beliefs depending on the knowledge categories, we pool the data to focus on differences in beliefs based on senders’ incentives.¹⁴

2.2.3 Additional experiments

We conduct two additional experiments, measuring sharing beliefs and donation decisions, reported in detail in Online Appendices E.5 and E.6.

Sharing is a form of online attention that increases the number of people who see a summary. We elicit receivers’ sharing beliefs, by asking them which summary would be shared most frequently on social media, with and without incentives for accuracy. Receivers read summaries written with Clickbait-Incentives and Info-Incentives, or Viral-Incentives and Info-Incentives, as in the Click experiments. The results show that summaries written with Viral-Incentives are selected significantly more often as the ones that would be most shared on social media.

We elicited receivers’ support for vaccination and climate causes in an additional experiment in which receivers made two donation decisions towards vaccine and climate conservation nonprofit organizations. The results indicate that summaries written with Clickbait-Incentives and Viral-Incentives directional, but not statistically significant, increase donation rates.¹⁵

2.3 Experimental Procedures

All Receiver experiments were conducted on Prolific Academic. They were pre-registered on aspredicted.org (details in Online Appendix B.1). A total of 3,743 participants were exposed to summaries by professionals (1,043 in the Knowledge, 905 in the Knowledge-Recall, 651 in the Click-Knowledge, 643 in the Click-Knowledge-Recall, and 501 in the Beliefs & Intentions

¹⁴Among summaries written under Clickbait- and Viral-Incentives, half resulted in below-median knowledge and half resulted in above-median knowledge. They all received above-median clicks (for Clickbait-Incentives) and above-median sharing belief (for Viral-Incentives). Among summaries written with Info-Incentives, all received above-median knowledge. Half of the summaries received below-median click rate and half received below-median sharing belief. We pool these groups in the main analyses, for comparability with other results and because differences between the groups are not observed (p -value > 0.05 in all comparisons), and provide detailed results in Online Appendix E.3.

¹⁵With our sample of 299 participants, the increase in donations is 0.04-0.07 percentage points (s.e. 0.07) with a rate of giving between 26 and 33 percent for summaries written with Info-Incentives.

Experiment).¹⁶ Receivers earned a payment for participation equivalent to \$12 per hour, based on the expected length of each experiment (5-25 minutes). The approval rate required was 95-100% in previous studies, and receivers were required to be located in the United States (checked via their IP address).

3 Results of the Sender Experiment


This section shows the impact of senders’ incentives on summary accuracy, focus and writing style. It explores how summaries differ in two dimensions of content, sample information and implications, and compares senders’ writing to online writing about the research. It concludes by examining the summary writing by an LLM, rather than human senders. While the effects of Clickbait- and Viral-Incentives do not differ significantly in the Sender Experiment among professionals, we report them separately—deviating from the pre-registration—to be consistent in the presentation of results across all Sender and Receiver Experiments.

A. Accuracy. Summaries written with Info-Incentives have a total of 5.5 units, on average. Within each summary, 4.3 units are Accurate-In text, 0.35 units are Inaccurate-In text, and 0.84 units are Not in text. Figure 3 and Table 1 show that the accuracy index drops by 0.31 SD with Clickbait-Incentives and 0.38 with Viral-Incentives (q -value ≤ 0.01).

The number of Accurate-In text units decreases by 0.61 units (or, equivalently, 11% of total units, and 0.35 standard deviations (SD)) under Clickbait-Incentives, and 0.82 units (14% or 0.47 SD) under Viral-Incentives, as detailed in Online Appendix Table C.1. This drop is matched with an increase in the number of Not in text units (0.62 units for Clickbait-Incentives, and 0.87 units for Viral-Incentives), leaving the total number of units unchanged. At the same time, Clickbait and Viral-Incentives do not increase the number of Inaccurate-In text units, which are rare (between 0.35 and 0.39 units).

¹⁶Recruitment occurred in two waves for the Knowledge and Click-Knowledge experiments, and controls for wave of data collection are included throughout in the analyses. Additionally, 659 participants were exposed to summaries by GPT (360 in the Knowledge experiment, 299 in the Click experiment). There are 878 participants in the incentivized Sharing experiments (611 exposed to professionals’ summaries and 267 to GPT’s summaries), 404 in the unincentivized Sharing experiment (198 exposed to professionals’ and 206 exposed to GPT’s summaries), and 505 in the Donation experiment, reported in the Online Appendices E.5 and E.6.

Figure 3: Average Effect of Sender Incentives on Message Production



Notes: This figure shows the average treatment effect of Clickbait-Incentives and Viral-Incentives, relative to Info Incentives, on the three main outcomes, capturing senders' communication. The Accuracy, Focus and Style Indices show the effects of the treatments on each index, constructed following Anderson (2008). The separate effects on each variable are shown separately (light-shaded icons). The accuracy variables (Accurate-In text, Inaccurate-In text, Not in text) and the focus variables (Introduction, Method, Finding, Conclusion and Call to Action) are the number of units in each category in a message. All variables are standardized. Error bars correspond to 95% confidence intervals. Detailed results in Online Appendix Tables C.1, C.2 and C.3.

B. Focus. Summaries written with Info-Incentives feature 1.5 units classified as Introduction, 2.2 units as Findings, and 1.4 units as Conclusion. On average, only 0.4 units are classified as Methods, indicating that methodology receives little coverage in research summaries. Clickbait- and Viral-Incentives decrease the focus on research details (Methods and Findings), increasing the focus on motivating the research and drawing implications, as detailed in Online Appendix Table C.2. The Focus index drops by 0.31 and 0.45 SD, respectively (Table 1, q -value < 0.01).

Table 1: Incentive Effects on Sender Summaries

	(1)	(2)	(3)
	Accuracy Index	Focus Index	Style Index
Clickbait-Incentives	-0.308** (0.124) [0.010]	-0.305*** (0.112) [0.008]	0.402*** (0.121) [0.004]
Viral-Incentives	-0.375*** (0.133) [0.002]	-0.454*** (0.114) [0.001]	0.676*** (0.127) [0.001]
Constant	-0.378* (0.193)	0.584*** (0.149)	-0.669*** (0.173)
Observations	595	595	595

Notes: This table displays the estimated average marginal effects from linear regressions on the Accuracy, Focus and Style indices, constructed as in Anderson (2008). All regressions include fixed effects for each article and sender characteristics. Standard errors clustered at the sender level are reported in parentheses, with * $p < .10$; ** $p < .05$; *** $p < .01$, based on unadjusted p -values. Sharpened FDR q -values are computed using the Anderson (2008) two-stage BKY procedure, over all outcomes in the Table for a treatment, and are shown in square brackets.

C. Style. Clickbait- and Viral-Incentives lead to shorter and more readable summaries, that have more exclamation and question marks. While the average summary is 102 words long under Info-Incentives, it is 7.5 words shorter with Clickbait-Incentives and 13.3 words shorter with Viral-Incentives. The Flesch Reading Ease score under Info-Incentives is 25.55, equivalent to a *college graduate level* – meaning it is very difficult to understand and suitable for university graduates. The score goes up by 6.7 and 9.8 points, respectively, the equivalent of going down one grade level (10 points), decreasing the reading difficulty to *college level*. The frequency of exclamation marks more than quadruples, increasing from 5% by 23pp, and that of question marks more than doubles, from 7% by 14-16pp with Clickbait- and Viral-Incentives (see Online Appendix C.2 for more detail). Sentiment decreases with Viral-Incentives too, by 0.12 SD – details on each outcome are in Online Appendix Table C.3 and Section C.2. Reflecting these changes, the Style index increases by 0.40 SD and 0.68 SD with Clickbait- and Viral-Incentives, respectively (Table 1, q -value < 0.01).

The effects are summarized in Result 1.

Result 1. Clickbait- and Viral-Incentives lead to summaries that: (1) contain a lower amount of Accurate-In text information, and more Not in text content, (2) focus more on

the motivation and less on the methodology of the research, and (3) are shorter, easier to read and with punctuation that exhorts attention, such as exclamation marks.

Taken together, the effects of Clickbait-Incentives are consistent with senders creating information gaps, by increasing their focus on introducing the research and providing less information about the research. The effects of Viral-Incentives are qualitatively similar to those of Clickbait-Incentives. Contrary to sharing motives based on violating expectations, there is no more inaccuracy or an increase in Conclusion units, rather more new content (Not in text units). An open question is whether the new content could mislead receivers through *how* it describes the context and the broader importance of the research. We therefore conduct exploratory analyses that examine how the research is described.

3.1 Effects of Incentives on *How* the Research is Described

First, we examine sample information provided in summaries, which is quantifiable. Consistent with the drop in units describing the research methods, senders are less likely to report sample information with Clickbait- and Viral-Incentives (39% and 34%, respectively), compared to Info-Incentives (48%). If the sample is reported, under Viral-Incentives, but not Clickbait-Incentives, senders used sample descriptions that were not precise, without reporting the exact group of individuals or specimens involved (Online Appendix Table C.6, $q\text{-val} = 0.03$).¹⁷ These patterns suggest that, if information is provided, more vagueness is present under Viral-Incentives.

Second, we explore whether senders use language that increases the relevance of the research to readers. We consider two measures: whether the research directly appeals to the reader (e.g., “you will learn how to sleep better”), and whether the summary draws incorrect implications (e.g., the reader should change their diet to reduce insomnia, although the research caveats that the findings are thus far correlational). Both Clickbait- and Viral-Incentives increase direct appeals to readers in summaries (by 17 percentage points (pp) and 26 pp, respectively, relative to a 0.15 mean under Info-Incentives, $q\text{-val} < 0.01$, Online

¹⁷It is rare for the sample size or exact origin (e.g., group or specimens involved) to be reported incorrectly (1% to 6% of the cases). It is more frequent for the sample size and origin to be described in vague terms (8% to 38%).

Appendix Table C.6). Simultaneously, with Viral-Incentives there is a marginally significant increase in the likelihood of incorrect implications (by 12 pp, relative to a 0.29 mean under Info-Incentives, q -val= 0.054, Online Appendix Table C.6), which could mislead readers regarding the relevance of the research for them.

These findings suggest that, in addition to the reduction in the amount of detailed information about the research with Clickbait- and Viral-Incentives, the likelihood that summaries contain content that could mislead readers is higher.

3.2 Naturally-Occurring Attention to the Research Articles

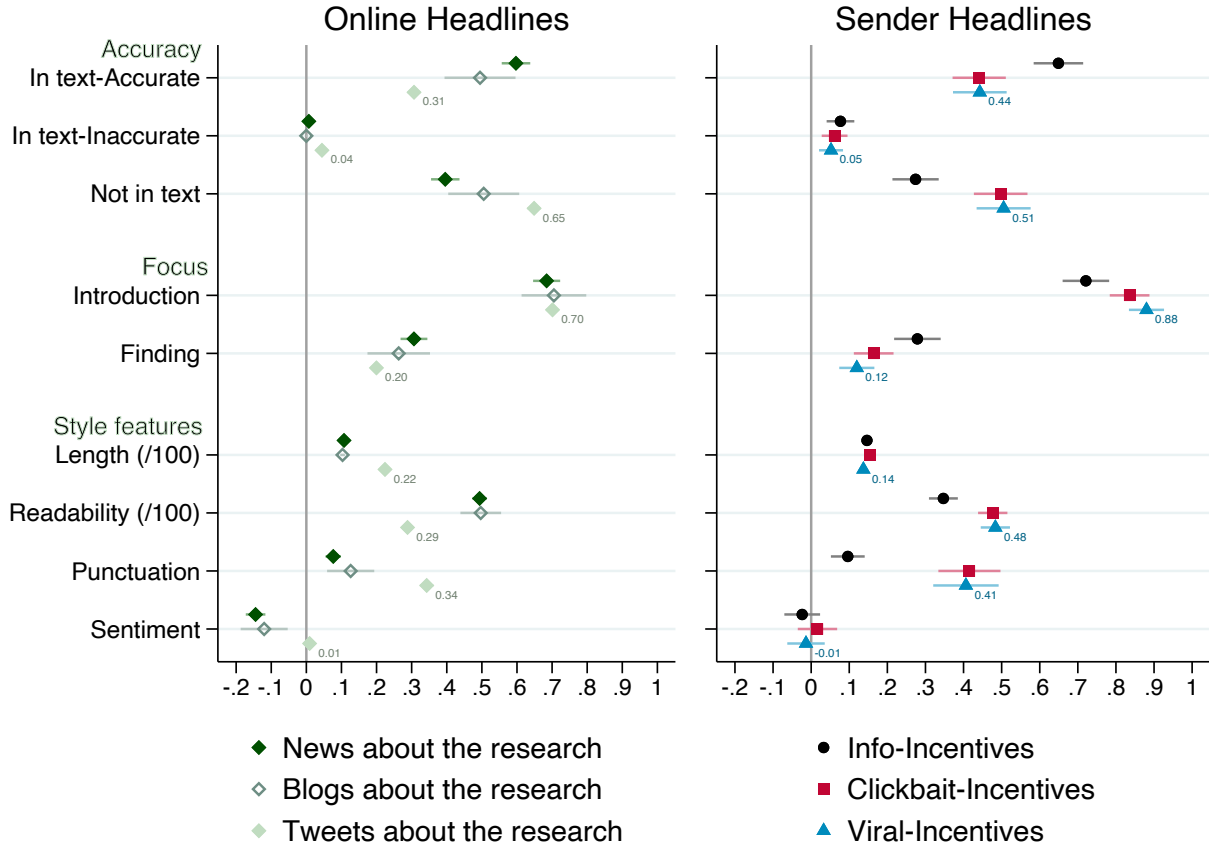
Next, we examine whether the content written by senders in the Sender experiment is similar to that appearing online, in news, blogs and Twitter, collected by Altmetric (details in Online Appendix D). Figure 4 reveals strong similarities in accuracy and focus between headlines written in the Sender experiment and those appearing online. Headlines written with Clickbait and Viral-Incentives feature Accurate-In text content 44% of the time, which is between the percentage found in tweets (31%) and online blogs (50%) and less than that found in news (60%). In both settings, the remainder of content is mainly classified as Not in text. Headlines rarely feature Inaccurate-In text content (at most 4% of the time for tweets).

In terms of focus, a majority of headlines focus on introducing the research, over 70% of the time. The rate of such units is between 84% and 88% with Clickbait- and Viral-Incentives, compared to 68% for news, 70% for tweets, and 71% for blogs.

In terms of writing style, the average length of headlines written by senders is in between the length of headlines in news and blogs (11 words) and that of tweets (22 words). The readability of headlines in the Clickbait- and Viral-Incentives (a score of 48) is similar to that in news (49) and blogs (50). The sentiment of senders' headlines (between -0.02 and 0.02) is similar to that of tweets, which is neutral (0 sentiment score) and somewhat more positive than the sentiment of news and blogs headlines (-0.15 and -0.12 sentiment score). Further, the use of punctuation under Clickbait- and Viral-Incentives is similar to that observed in tweets.

Taken together, although the incentives of online content producers are heterogeneous

Figure 4: Headline Features of Sender Summaries and Naturally-Occurring Mentions



Notes: This figure compares the headlines written in news, blogs and tweets about the research (left panel) with the headlines written by senders (right panel). The figure shows the fraction of headlines classified by accuracy (Accurate-In text, Inaccurate-In text, Not in text) and focus (Introduction and Finding). The figure then shows the average values for features of writing style in headlines (or tweets). Words and readability scores are divided by 100 for comparability purposes. Error bars correspond to 95% confidence intervals. Details shown in Online Appendix Tables D.1 and D.2.

and unknown, the findings indicate that the distribution of features of senders' summaries in the Sender experiment is consistent with those found online.

3.3 Large Language Model as Sender


The LLM reveals qualitatively similar results to our experiment with professionals. The effects of incentives (instructions) on content creation by the LLM are summarized in Figure 5. Additional regression analyses are shown in Online Appendix F.1.

A main difference between professionals and the LLM is that the LLM is significantly

more likely to only use the materials it was provided. This increases the accuracy of summaries created by the LLM in all treatments. Under Info-Incentives, summaries contain almost only Accurate-In text information: 5.2 out of 5.4 units are Accurate-In text. On average, only 0.15 units are Inaccurate-In text and 0.1 units are Not in text (compared to 0.4 and 0.8 with professionals).

For the LLM, Viral-Incentives lead to the strongest changes in accuracy, focus, and style. With regards to accuracy, Clickbait- and Viral-Incentives reduce the number of Accurate-In texts units. However, the LLM does not substitute them with more Not in text units, like professionals did. Clickbait- and Viral-Incentives increase the number of Introduction units, and decrease Methods and Findings units, as they did for professionals.¹⁸ Likewise, summaries become shorter and more readable, and with Viral-Incentives, with more punctuation that aims to attract attention.

Figure 5: Average Effects on Message Production: LLM as Sender



Notes: This figure shows the average treatment effect of Clickbait-Incentives and Viral-Incentives, relative to Info Incentives, on message production by the LLM (GPT). Error bars correspond to 95% confidence intervals. $N = 360$ summaries.

¹⁸Since the LLM does not include any units that are calls to action, there are no effects for them.

Overall, the effects of instructing the LLM to create content that grabs attention are qualitatively similar to those observed for professionals. However, they are smaller in magnitude, which can be explained by the fact that the LLM produces summaries that more closely follow the materials than professionals.

4 Results of the Receiver Experiments

We next examine the impact of senders’ incentives in the Sender experiment on the knowledge of receivers. We consider the four main outcome variables described in Section 2.2: the knowledge of receivers when exposed to the summaries only, their click rates, knowledge when exposed to detailed research information, and the resulting overall knowledge.

A. Summary-Based Knowledge: Summaries created with Clickbait- and Viral-Incentives result in lower knowledge among receivers. Table 2 column (1) shows that, when the summaries are available, the decrease in knowledge is of 5.5 to 6.8 pp, respectively. And, Panel A of Figure 6 shows that knowledge with Info-Incentives stochastically dominates that with Clickbait- and Viral-Incentives.

Panel B of Figure 6 shows the largest drop is in knowledge about the study sample. Receivers correctly answer questions about the study sample in 38% of the cases when exposed to summaries written with Clickbait and Viral-Incentives, compared to 51% with Info-Incentives. With Clickbait-Incentives, knowledge about the findings, the methods and the conclusion of the research decreases to a lesser extent, 3-4 pp. With Viral-Incentives, knowledge decreases more for questions about the findings and methods, by 6 and 7pp, respectively, and shows a smaller drop for questions about the conclusion (2 pp, see Online Appendix Table E.8 for details).

When summaries must be recalled, knowledge drops across all conditions and the impact of senders’ incentives is reduced, as shown in column (2) of Table 2.¹⁹ Panel B of Figure 6 shows that the drop in knowledge about the sample and method becomes smaller when

¹⁹To analyze summary-based knowledge with recall, we pool receivers’ knowledge in the Knowledge-Recall Experiment and in the Click-Knowledge-Recall Experiment, when answering questions about articles they did not click to read more about (four subjects are dropped because they click to read all articles), because there is no significant difference in accuracy, as shown in Online Appendix Table E.7.

Table 2: Clicks and Knowledge - Summary-level Results


	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Summary-Based Knowledge		Click Rate	Materials-Based Knowledge		Overall Knowledge	
	Avail.	Recall		Avail.	Recall	Avail.	Recall.
Clickbait-Incentives	-0.055*** (0.013) [0.001]	-0.028* (0.015) [0.078]	0.050*** (0.012) [0.001]	0.013 (0.023) [0.197]	-0.030 (0.025) [0.112]	-0.028** (0.013) [0.063]	-0.024* (0.013) [0.078]
Viral-Incentives	-0.068*** (0.013) [0.001]	-0.021 (0.015) [0.075]	0.041*** (0.012) [0.002]	-0.043 (0.026) [0.071]	-0.042* (0.025) [0.071]	-0.062*** (0.013) [0.001]	-0.025** (0.012) [0.040]
Constant	0.374*** (0.018)	0.330*** (0.022)	0.350*** (0.020)	0.785*** (0.029)	0.565*** (0.034)	0.507*** (0.019)	0.417*** (0.020)
Observations	595	589	595	395	451	595	595
Info-Incentives: Mean	0.57	0.49	0.20	0.82	0.64	0.62	0.53
Info-Incentives: SD	0.18	0.17	0.12	0.19	0.21	0.15	0.14

Notes: This table displays the estimated coefficients from linear regressions on: (1) Summary-based knowledge when the summary is available (Knowledge experiment) and when it must be recalled and the preferred click choices of the receiver are implemented (Knowledge-Recall and Click-Knowledge-Recall experiment, for summaries not clicked on, which due to random implementation results in N=589 being observed by receivers in this group) (columns (1) and (2)); (2) click rate (column (3)); (3) Materials-based knowledge (for those whose preferred click choices were implemented), when the materials were seen by the receiver (Click-Knowledge Experiment) and when they had to be recalled (Click-Knowledge-Recall Experiment) (columns (4) and (5)) – note that, since ca. 6% of summaries were never clicked on and ca. 25% of click choices were not implemented, by design, the number of observations in columns (4) and (5) is lower; (4) Overall knowledge (columns (6) and (7)). Regressions include article fixed effects and sender characteristics. Robust (HC3) standard errors shown in parentheses, where * $p < .10$; ** $p < .05$; *** $p < .01$ are based on unadjusted p-values. Sharpened FDR q -values are computed using the Anderson (2008) two-stage BKY procedure, over all outcomes in the Table for a treatment, and shown in square brackets.

summaries must be recalled. These reduced drops suggest that receivers pay less attention to these dimensions of the research and that senders with Clickbait- and Viral-Incentives focus their summaries on dimensions that receivers are most likely to recall.

The differences in receiver knowledge closely correspond to the effects of senders’ incentives on the accuracy and focus of summaries documented in Section 3. An important part of this correspondence stems from the summary information: If the summary contains a clear answer to a knowledge question, receivers’ likelihood of answering it correctly strongly increases. At the same time, exploratory analyses show that information in the summaries is not the only feature that matters. If summaries are available, a higher style index lowers the likelihood that a receiver answers the question correctly, as we report in Online Appendix Table E.5 column (4).

Figure 6: Receiver Knowledge When Exposed to Summaries



Notes: Panel A (left) shows the distribution of summary-based knowledge (fraction of correct answers after reading summaries) when receivers had the summaries available (Knowledge Experiment). Panel B (right) shows summary-based knowledge by question type. The filled markers indicate “Summary Available”, for receivers who saw the summary while answering the questions. The empty markers indicate “Summary Recalled”, for receivers who had to answer questions based on recall. Within Panel B, there are 4 subpanels: “All” panel (left-most), for the average over all questions, followed by the average knowledge for Finding, Sample, Methods and Conclusion questions separately. Estimates from responses at the *question* level based on $N = 2,588$ receivers. Regression estimates shown in Online Appendix Tables E.5 and E.6, column (1), for all questions, and Tables E.8 and E.9, column (1), by question type. Error bars correspond to 95% confidence intervals.

B. Click rate: Summaries written with Info-Incentives are clicked on 20% of the time. Table 2 column (3) shows that click rates increase by 5 pp and 4 pp (q -value < 0.01 in both cases) for summaries written with Clickbait- and Viral-Incentives, respectively. These results are obtained when comparing click rates at the message level, and also when examining receivers’ individual-level choices. The increased click rate toward summaries written with Clickbait-Incentives is observed in both the Click-Knowledge and Click-Knowledge-Recall experiment (5pp in both cases). As shown in Online Appendix Table E.10, the increased click rate on

summaries written with Viral-Incentives is not observed in the Click-Knowledge experiment, but it is observed in the Click-Knowledge-Recall experiment, suggesting that the differences in click rates between Viral- and Info-Incentives may be less robust.

C. Materials-based Knowledge: When receivers click to learn more about the research, they are all exposed to the same (detailed) information that senders read. In all treatments, their knowledge increases. The percentage of correct responses is 79%-85% when the materials are available, and 57%-60% when the materials must be recalled (for detail, see Online Appendix Tables E.12 and E.13, and Online Appendix Figure E.1).

Columns (4) and (5) of Table 2 show that sender incentives do not affect materials-based knowledge, when summaries are written with Clickbait-Incentives. By contrast, there is a 4pp marginally significant decrease when they are written with Viral-Incentives, suggesting that receivers' materials-based knowledge was still weakly affected by what they read in the summaries.

D. Overall Knowledge: When receivers answer knowledge questions having (summaries or materials) information available, overall knowledge exhibits a marginally significant drop with Clickbait-Incentives (q -value= 0.063) and a significant drop with Viral-Incentives, by 6pp (q -value< 0.01). When receivers must recall what they learned about the research, the knowledge losses are smaller. Overall knowledge exhibits a marginally significant drop of 2.4 pp (q -value= 0.078) with Clickbait-Incentives, and a significant drop of 2.5pp (q -value= 0.040), with Viral-Incentives.

These findings imply that the loss in knowledge when receivers are exposed to summaries written with Clickbait- and Viral-Incentives can be compensated by an increase in clicks, resulting in small changes in overall knowledge – an illustration of the attention-information tradeoff.

Result 2. Receiver knowledge decreases with Clickbait- and Viral-Incentives, when they are asked about the research and have the summaries available. If summaries must be recalled, the effect of senders' incentives is smaller and knowledge does not drop significantly. Simultaneously, summaries written with Clickbait- and Viral-Incentives attract significantly

more attention. Hence, accounting for clicks, overall knowledge losses with Clickbait- and Viral-Incentives are small in magnitude.

4.1 Impressions: Receivers' Errors, Beliefs and Intentions

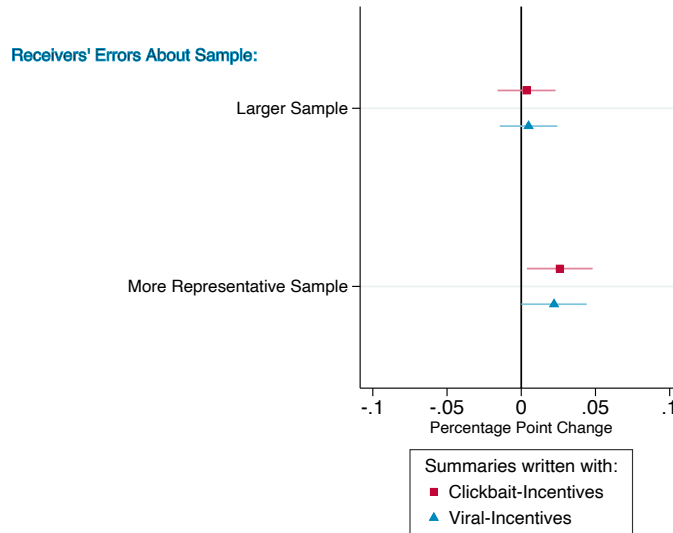
While Clickbait- and Viral-Incentives do not lead to large losses in knowledge on average, a large majority of receivers (75-80%) do *not* click and are only exposed to summaries. The evidence in Section 3 shows that these summaries contain significantly less information about the research. An open question is whether the summaries induce receivers to hold beliefs about the research that are systematically incorrect. In this section, we examine two types of beliefs about the research: beliefs about the sample involved in the research, measured by the size and representativeness of the sample, and beliefs about the relevance of the research to receivers' behavior outside the experiment.

4.1.1 Receiver Errors About the Sample

Since senders provided less information about the sample, receivers could form noisy beliefs about it or, alternatively, because of how summaries are written and their new (Not in text) content, they could form biased beliefs about it. We examine whether receivers were more likely to believe that the sample was larger or more representative of a US citizen, than the one used in the study. Figure 7 shows that receivers are not more likely to believe that the sample was larger. But, when exposed to summaries written with Clickbait- and Viral-Incentives, they are more likely to believe that it is representative of a US citizen than it actually is. The effects are of limited magnitude, 3pp, though, suggesting potential, small biases in receivers' impressions about the similarity of the research sample to the reader.²⁰

²⁰Detailed regression results are provided in Online Appendix E. A related question is whether mistaken beliefs are more likely to arise when no information is presented at all, or rather when some sample information is included, but it is vague. As shown in Online Appendix Table E.4 when receivers are exposed to imprecise sample information in summaries written with Clickbait- and Viral-Incentives, they are significantly more likely to hold biased beliefs about the representativeness of the sample. This analysis suggests that a complete lack of information may lead to random errors. But, when some information is provided under incentives to grab attention, this partial information may lead receivers to believe a more representative sample participated in the research than it actually did.

Figure 7: Receiver Errors



Notes: This figure displays the estimated coefficients from linear regressions on the likelihood that receivers incorrectly answer the question about the sample of a research article, selecting one that states a larger sample than the one used (“Larger Sample”) or one that states a group being studied that is closer to a more representative group, defined as closer to a US citizen (“More representative sample”). The regression includes receivers from the Knowledge and Knowledge-Recall experiments who incorrectly answered the questions about the study sample. Robust standard errors clustered at the receiver level are used. A detailed description for how each answer option is coded is provided in Online Appendix B.2. Error bars correspond to 95% confidence intervals. Details shown in Online Appendix Tables E.2 and E.3.

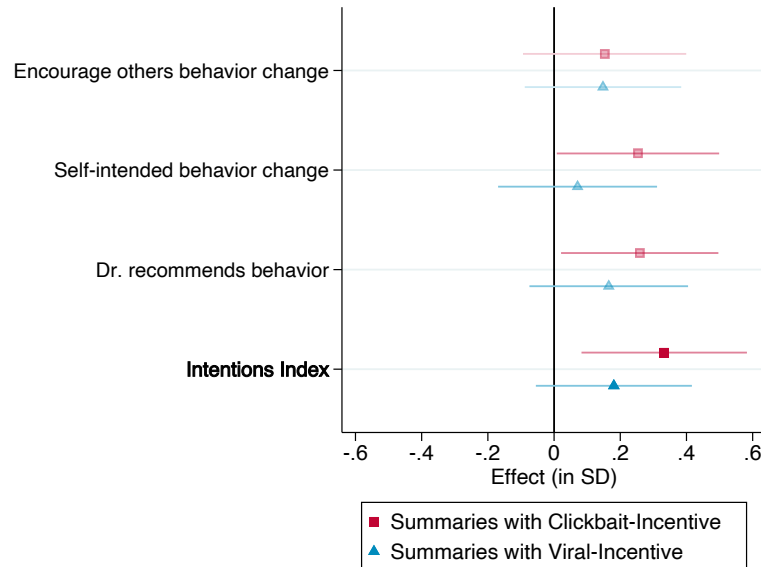
4.1.2 Beliefs and Intentions

We examine whether receivers formed different beliefs about the implications of the research depending on senders’ incentives when writing the summaries they were exposed to in the Beliefs & Intentions experiment. While the studies on sleep and cancer are based on first evidence that showed changes that *might* help, the studies on MMR vaccines further confirm the already known protective properties of the vaccine and the lack of evidence of harm (specifically, autism risk).

Figure 8 shows the average impact of being exposed to summaries written with Clickbait- and Viral-Incentives, compared to Info-Incentives, on encouragement of others, intended own behaviors and beliefs about doctor recommendations, and the intentions index. Summaries written with Clickbait-Incentives lead to a significant increase in the index (q -value < 0.01). They also lead to a marginally significant increase in the individual outcomes (q -value = 0.07 for beliefs about doctor recommendations and own behaviors, and q -value = 0.08 for en-

couragement of others). Summaries written with Viral-Incentives do not have a statistically significant impact on these outcomes, but are directionally similar (see Online Appendix Table E.14).

Figure 8: Receiver Beliefs and Intentions



Notes: This figure displays the estimated coefficients from linear regressions on the effect of summaries written with Clickbait- and Viral-Incentives on questions about encouraging others, their own behavior, and doctors’ official recommendations. All outcomes are standardized. Encouraging others and doctor recommendations involved five behaviors: rocking to improve sleep, consume a low GI diet to reduce insomnia, to eat vegetables to reduce the risk of cancer, to take Vitamin A to reduce the risk of skin cancer, and to give their children the MMR vaccine. Self-intended change is an indicator variable that takes value one if receivers indicated the summaries increased their likelihood to change behavior in an open-text response to a question asking about their effects of the summaries. The intentions index aggregates follows Anderson (2008), computed on the sum over the sum of behaviors receivers would encourage others to take and believe are doctor recommendations. Detailed regression results are shown in Online Appendix Table E.14.

The data show that receivers are generally aware that the MMR vaccine is officially recommended and correctly indicate so in 96% of the cases. However, they hold incorrect beliefs about doctor’s official recommendations to improve sleep or reduce cancer risk. Under Info-Incentives, between 50% and 72% of receivers indicate changes in diets and rocking as officially recommended. These fractions increase directionally when summaries are written with Clickbait- and Viral-Incentives, as detailed in Online Appendix Table E.16.²¹

²¹To better understand what drives encouragement of others, we explore which summary features predict encouragement. Summaries that contain incorrect implications, as described in Section 3.1, increase the likelihood that the receiver intends to encourage behavior changes in others, providing suggestive evidence that implications drawn in summaries could affect behavior (see Online Appendix Table E.17). We also

Result 3. Exposure to summaries created under Clickbait- and Viral-Incentives increases the fraction of receivers who believe the sample used in the research is representative of the US population and, for Clickbait-Incentives, this exposure increases the likelihood that intended behaviors and beliefs change. Specifically, behaviors suggested as beneficial in the summaries are more likely to be perceived as official doctor recommendations and affect behavioral intentions, at a marginally significant level.

While receivers may infer the wrong implications from the initial evidence they see for cancer and sleep, it is important to highlight that these behaviors are health-improving behaviors. For ethical reasons, we did not expose receivers to research that could reduce willingness to vaccinate or increase the intake of harmful substances. However, recent evidence by Allen et al. (2024) shows that exposure to content on Facebook that discourages vaccination, but is not flagged as false, has significant negative impacts on vaccination intentions, consistent with the effects we document in the Beliefs & Intentions experiment.

4.2 Heterogeneity in Click Rates and Knowledge

The analysis thus far has focused on the average effect of summaries written with Clickbait- and Viral-Incentives on receiver knowledge and attention. The findings suggest that, while they contain less information, summaries also attract more attention when they are written with the incentive to generate clicks and shares. In this section, we examine the heterogeneity in the relationship between click rates and knowledge. We focus on heterogeneity at the summary level, rather than the sender level, because individual-level sender effects explain limited variation (2-15%) in knowledge, clicks, and the combination of both.²²

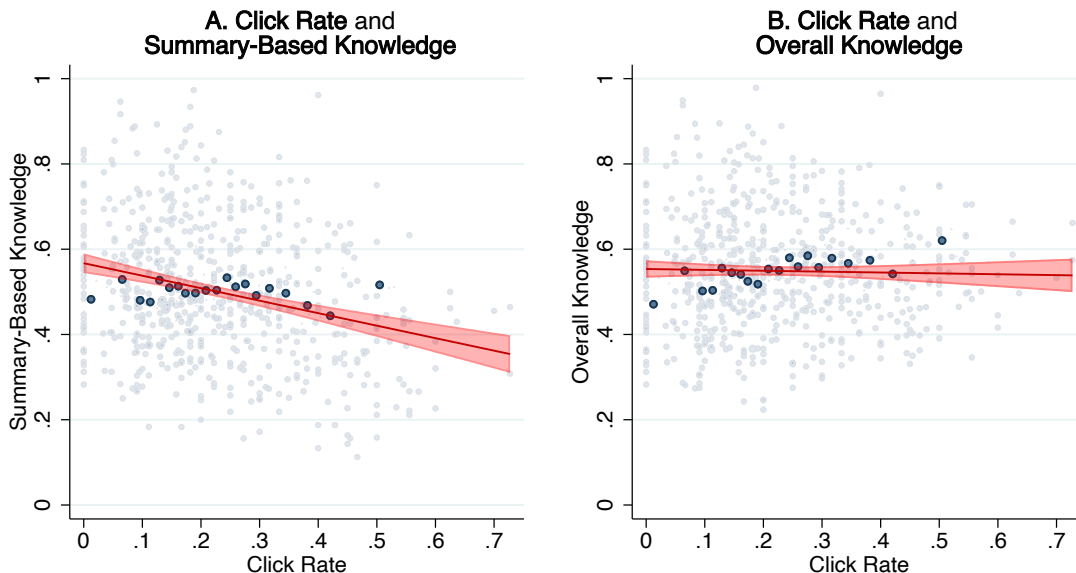
Figure 9 shows the relationship between click rates and summary-based knowledge (Panel A) or overall knowledge (Panel B). Consistent with summaries creating informa-

examine whether receivers’ trust in media and scientists changes. When exposed to summaries written with Viral-Incentives, relative to Info-Incentives summaries, their trust does not change. There is a marginally significant increase in trust in the media when receivers are exposed to summaries written with Clickbait-Incentives, as shown in Online Appendix Table E.18.

²²The fraction of explained variance is calculated with the intraclass correlation coefficient, based on a random effects panel regression, where senders are the individuals. A small fraction of the variance in clicks is explained by sender differences (2%). A higher fraction of the variance in knowledge and the likelihood that both knowledge and clicks are above median is explained by sender differences (5-15%).

tion gaps, as in Golman and Loewenstein (2018), Panel A shows that a one percentage-point increase in click rates is associated with an average 0.3 percentage-point decrease in summary-based knowledge. Panel B shows that, once we account for the knowledge that receivers gain after clicking, the negative relationship between click rates and knowledge disappears. The gain in knowledge among those who click offsets the knowledge losses among those who only read the summary.²³

Figure 9: Heterogeneity in Receivers’ Attention and Knowledge



Notes: Panel A shows the average click rate on each message (x-axis) and summary-based knowledge (y-axis) pooling all Sender Incentive treatments. Panel B shows the average click rate (x-axis) and the overall receiver knowledge, (y-axis). In the background, each light-colored dot is weighted by the number of observations (summaries) with the same values. Each dark-colored dot represents a demi-decile of the distribution. A fitted linear regression between the two variables is shown in each panel, with 95% confidence intervals. Detailed regression analyses and separate figures are provided in Online Appendix Table E.19 and Figure E.2.

Figure 9 reveals significant heterogeneity in click rates and knowledge. For example, the top 25th percentile of summaries achieve a click rate of more than 31%, while 6% are never clicked on. In Online Appendix E.4 (Table E.20) we examine which, if any, features of summaries are predictive of click rates and knowledge. Summaries with a higher accuracy index are associated with more knowledge among receivers, but fewer clicks. Considering clicks and knowledge together, we find that it is difficult to predict above-median knowledge

²³Online Appendix Table E.19 shows detailed regressions with further controls to examine the relationship between clicks and summary-based knowledge. This relationship does not change significantly depending on whether receiver must recall the summary.

and clicks (regressions achieve an R^2 of 0.09), a common finding in related work (e.g., Bakhsy et al., 2011).

5 Conclusion

Incentives to attract online attention, in the form of clicks and sharing, can have important impacts on information transmission. This paper combines several experiments to provide evidence on the impact of incentives to attract attention on the content that is generated, and the subsequent knowledge and attention allocation of readers in the context of science communication. The findings reveal that incentives to attract attention generate a tradeoff: the content contains less information but it is more attention-grabbing. Since these incentives generate information demand, making interested readers better informed, we find small impacts of attention incentives on receivers' knowledge about the research, on average.

However, there is a large share of receivers who do not demand more information and are left only with the information that is provided in short summaries. We find that these summaries can lead to misperceptions. While attention incentives do not increase the amount of content that is factually incorrect, they lead senders to introduce new content that attempts to increase the perceived relevance of the findings for the reader. Readers are more likely to mistakenly believe the study sample is representative of the US population and exhibit bias towards believing that the behaviors studied by the research are prescriptive and motivating to change behavior, although the research findings are suggestive and too preliminary to generate clear behavioral implications.

Taken together, these findings suggest that, when communicating information about science, a new channel can lead misperceptions to arise: incorrect beliefs may be generated through vague information, that lacks detail, combined with content focused on the relevance of the topic, rather than through explicitly inaccurate information about what scientists did and found.

It is important to highlight that this study sheds light on the impact of incentives to grab attention when professionals are hired to write about scientific articles, without reputational concerns. More research is needed to investigate whether adding reputational

incentives could result in different impacts, not only on how research is described, but also which research topics are chosen to be broadcasted. Further, the current study focuses on the first stage of communication, from a highly informed sender to a receiver, who could potentially share some of the content with others, creating further distortions in information transmission (e.g., Carlson, 2018; Graeber et al., 2024b). Documenting the impact of different incentives and studying the intensive and extensive margins of communication will be important to provide further insights into how attention and information transmission interact.

This paper focuses on science communication, but political communication is another setting of relevance. Some content producers (news outlets, politicians, and others) could have strong incentives to maximize their clicks and shares when they face intense competition (e.g., Djourelouva et al., 2024) and it is an open question whether incentives to grab attention can play a role in the polarization of content (e.g. Levy, 2021).

References

1. Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4), 1115-1153.
2. Abeler, J., Huffmann, D. B., and Raymond, C. (2023). Incentive Complexity, Bounded Rationality and Effort Provision. IZA Discussion Paper 16284.
3. Acemoglu, D., Ozdaglar, A., & Siderius, J. (2024). A model of online misinformation. *Review of Economic Studies*, 91(6), 3117-3150.
4. Allen, J., Watts, D. J., and Rand, D. G. (2024). Quantifying the impact of misinformation and vaccine-skeptical content on Facebook. *Science*, 384(6699), eadk3451.
5. Allcott, H., Gentzkow, M. , and Song, L. (2022). Digital addiction. *American Economic Review*, 112 (7): 2424–63.
6. Altay, S., De Araujo, E., and Mercier, H. (2022). “If this account is true, it is most enormously wonderful”: Interestingness-if-true and the sharing of true and false news. *Digital Journalism*, 10(3), 373-394.
7. Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American statistical Association*, 103(484), 1481-1495.
8. Bakshy, E., Hofman, J. M., Mason, W. A., and Watts, D. J. (2011). Everyone’s an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 65-74.

9. Banerjee, A., and Urminsky, O. (2023). The Language That Drives Engagement: A Systematic Large-scale Analysis of Headline Experiments. Available at SSRN 3770366.
10. Barasch, A., and Berger, J. (2014). Broadcasting and narrowcasting: How audience size affects what people share. *Journal of Marketing Research*, 51(3), 286–299.
11. Berger, J., and K. Milkman (2012), What Online Content Viral? *Journal of Marketing Research* 49 (2), 192-205.
12. Berzaghi, F., Longo, M., Ciaia, P., Blake, S., Bretagnolle, F., Vieira, S., Scaranello, M., Scarasia-Mugnozza, G., and Doughty, C. E. (2019). Carbon stocks in central African forests enhanced by elephant disturbance. *Nature Geoscience*, 12(9), 725-729.
13. Bordalo, P., Gennaioli, N., and Shleifer, A. (2016). Competition for attention. *The Review of Economic Studies*, 83(2), 481-513.
14. Bordalo, P., Gennaioli, N. and Shleifer, A. (2022). Salience. *Annual Review of Economics*, 14, 521-544.
15. Brady, W. J., Wills, J. Jost, J. T., Tucker, J. A. and Van Bavel, J.A (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences* 114(28):7313–7318.
16. Carlson, T. (2018). Modeling Political Information Transmission as a Game of Telephone. *Journal of Politics* 80 (1), 348-352.
17. Chen, H., and Suen, W. (2023). Competition for attention and news quality. *American Economic Journal: Microeconomics*, 15(3), 1-32.
18. Chou, W.S., Oh, A., Klein, W.M.P. (2018). Addressing Health-Related Misinformation on Social Media. *JAMA*, 320 (23), 2417-2418.
19. Clark, R. P. (2008). *Writing Tools: 55 Essential Strategies for Every Writer* (10th Anniversary Edition). Little, Brown and Company. New York, USA.
20. Crawford, V.P., and Sobel, J. (1982). Strategic Information Transmission. *Econometrica* 50 (6), 1431-1449.
21. Caplin, A., Dean, M., and Martin, D. (2011). Search and satisficing. *American Economic Review*, 101(7), 2899-2922.
22. Dean, M., and Neligh, N. N. (2023). Experimental tests of rational inattention. *Journal of Political Economy*, forthcoming.
23. Deversi, M., Ispano, A., and Schwardmann, P. (2021). Spin doctors: An experiment on vague disclosure. *European Economic Review*, 139, 103872.
24. Djourelouva, M., Durante, R., and Martin, G. J. (2024). The impact of online competition on local newspapers: Evidence from the introduction of Craigslist. *Review of Economic Studies*, rdae049.
25. Esponda, I., Oprea, R., and Yuksel, S. (2023). Seeing What is Representative. *Quarterly Journal of Economics* 138 (4), 2607–2657.

26. Flesch, R. (1948). A New Readability Yardstick. *Journal of Applied Psychology*, 32(3), 221.
27. Galperti, S., and Trevino, I. (2020). Coordination motives and competition for attention in information markets. *Journal of Economic Theory*, 188, 105039.
28. Gangwisch, J. E., Hale, L., St-Onge, M. P., Choi, L., LeBlanc, E. S., Malaspina, D., Opler, M.G., Shadyab, A.H., Shikany, J.M., Snetselaar, L., Zaslavsky, O., and Lane, D. (2020). High glycemic index and glycemic load diets as risk factors for insomnia: analyses from the Women’s Health Initiative. *The American journal of clinical nutrition*, 111(2), 429-439.
29. Gentzkow, M., and J. Shapiro (2008). Competition and Truth in the Market for News. *Journal of Economic Perspectives* 22 (2), 133-154.
30. Golman, R., and Loewenstein, G. (2018). Information gaps: A theory of preferences regarding the presence and absence of information. *Decision*, 5(3), 143.
31. Golman, R., Loewenstein, G., Molnar, A., and Saccardo, S. (2022). The demand for, and avoidance of, information. *Management Science*, 68(9), 6454-6476.
32. Graeber, T., Zimmermann, F., and Roth, C. (2024). Stories, statistics, and memory. *Quarterly Journal of Economics*, forthcoming.
33. Graeber, T., Noy, S., and Roth, C. (2024). Lost in Transmission. Manuscript.
34. Graeber, T., Hussam, R., Pandey, K., Shonchoy, A., Yamauchi, C. (2024). Translating Information Into Action. Manuscript.
35. Guan, M, Oprea, R., and Yuksel, S. (2023). Too Much Information. Working Paper.
36. Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. Psychology press.
37. Hviid, A., Hansen, J. V., Frisch, M., and Melbye, M. (2019). Measles, mumps, rubella vaccination and autism: a nationwide cohort study. *Annals of internal medicine*, 170(8), 513-520.
38. Jin, G. Z., Luca, M., and Martin, D. (2021). Is no news (perceived as) bad news? An experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, 13(2), 141-173.
39. Kahneman, D. (1973). *Attention and effort*. Vol. 1063. Englewood Cliffs, NJ: Prentice-Hall.
40. Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T. Y., and Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20(8), 963-973.
41. Kim, J., Park, M. K., Li, W. Q., Qureshi, A. A., and Cho, E. (2019). Association of vitamin A intake with cutaneous squamous cell carcinoma risk in the United States. *JAMA dermatology*, 155(11), 1260-1268.

42. Krajbich, I., Armel, C. and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience* 13.10, 1292-1298.
43. Lee, Y. R., Chen, M., Lee, J. D., Zhang, J., Lin, S. Y., Fu, T. M., Hao, C., Ishikawa, S.C., Katon, J., Zhang, Y., Shulga, Y.V., Bester, A.C., Fung, J., Monteleone, E., Wan, L., Shen, C., Hsu, C., Papa, A., Clohessy, J.G., Teruya-Feldstein, J., Jain, S., Wu, H., Matesic, L., Chen, R., Wei, W., and Pandolfi, P. P. (2019). Reactivation of PTEN tumor suppressor for cancer treatment through inhibition of a MYC-WWP1 inhibitory pathway. *Science*, 364(6441), eaau0159.
44. Levy, R. E. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3), 831-870.
45. LinkedIn (2023). What is the optimal length of a social media post?
<https://www.linkedin.com/advice/1/what-optimal-length-social-media-post-gohwf>
46. Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, 116(1), 75.
47. Loewenstein, G., and Wojtowicz, Z. (2023). The Economics of Attention. Available at SSRN 4368304.
48. Maxwell, S. L., Evans, T., Watson, J. E., Morel, A., Grantham, H., Duncan, A., Harris, N., Potapov, P., Runting, R.K., Venter, O., Wang, S., and Malhi, Y. (2019). Degradation and forgone removals increase the carbon impact of intact forest loss by 626%. *Science Advances*, 5(10), eaax2546.
49. Metzger, M. J., Flanagin, A. J., Mena, P., Jiang, S., & Wilson, C. (2021). From dark to light: The many shades of sharing misinformation online. *Media and Communication*, 9(1), 134-143.
50. Mina, M. J., Kula, T., Leng, Y., Li, M., De Vries, R. D., Knip, M., Siljander, H., Rewers, M., Choy, D.F., Wilson, M.S., Larman, H. B., Nelson, A.N., Griffin, D.E., De Swart, R.L., and Elledge, S. J. (2019). Measles virus infection diminishes preexisting antibodies that offer protection from other pathogens. *Science*, 366(6465), 599-606.
51. Mosleh, M., Pennycook, G., & Rand, D. G. (2020). Self-reported willingness to share political news articles in online surveys correlates with actual sharing on Twitter. *Plos one*, 15(2), e0228882.
52. Munch, Joe (2023), Ten Types of Sentences Every Scientific Paper Should Have.
<https://www3.mdanderson.org/library/education/pdf/10-types-of-sentences.pdf>
53. National Academies of Sciences, Engineering, and Medicine (2017). *Communicating Science Effectively: A Research Agenda*. Washington, DC: The National Academies Press. doi: 10.17226/23674.
54. Niitsuma, K., Kinoshita, T., and Takahashi, H. (2025). Prestige bias drives the viral spread of content reposted by influencers. *Scientific Reports*, 15, 98955.
55. Osmundsen, M., Bor, A., Bjerregaard Vahlstrup, P., Bechmann, A., and M.B. Petersen (2021). Partisan polarization is the primary psychological motivation behind political fake news sharing on Twitter. *American Political Science Review* 115 (3), 999-1015.

56. Pariser, E. (2011). *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.
57. Pennycook, G., and Rand, D.G. (2021). The psychology of fake news. *Trends in cognitive sciences* 25 (5), 388-402.
58. Perrault, A. A., Khani, A., Quairiaux, C., Kompotis, K., Franken, P., Muhlethaler, M., Schwartz, S., and Bayer, L. (2019). Whole-night continuous rocking entrains spontaneous neural oscillations with benefits for sleep and memory. *Current Biology*, 29(3), 402-411.
59. Peters, E (2020). *Innumeracy in the wild: misunderstanding and misusing numbers*. New York: Oxford University Press.
60. Pitzalis, Josh (2022). "How to make a twitter thread (with examples that won't make you cringe." <https://getchirrapp.com/how-to/twitter-threads>
61. Qiu, J., and Golman, R. (2022). Headline Curiosity. Available at SSRN. <https://ssrn.com/abstract=4221372>
62. Rinker, T. (2023). Sentimentr. <https://cran.r-project.org/web/packages/sentimentr>
63. Robertson, C. E., Pröllochs, N., Schwarzenegger, K., Pärnamets, P., Van Bavel, J. J., and Feuerriegel, S. (2023). Negativity drives online news consumption. *Nature Human Behaviour*, 7(5), 812-822.
64. Roozenbeek, J., Suiter, J., and Culloty, E. (2022). Countering misinformation: Evidence, knowledge gaps, and implications of current interventions. <https://psyarxiv.com/b52um/>
65. Sanchez-Pages, S. and M. Vorsatz (2007). An experimental study of truth-telling in sender-receiver games'. *Games and Economic Behavior* 61: 86–112.
66. Schimel, Joshua (2012). *Writing Science: How to write papers that get cited and proposals that get funded*. Oxford University Press, New York, USA.
67. Serra-Garcia, M. (2025). Data and Code for: The Attention-Information Tradeoff. American Economic Association, Inter-university Consortium for Political and Social Research. <https://doi.org/10.3886/E241261V1>
68. Serra-Garcia, M. and U. Gneezy (2021). Mistakes, Overconfidence and the Effect of Sharing in Detecting Lies. *American Economic Review* 111(10), 3160-3183.
69. Simon, H. A. (1969). Designing organizations for an information-rich world. Brookings Institute Lecture.
70. Smith, S. M., and Krajbich, I. (2018). Attention and choice across domains. *Journal of Experimental Psychology: General*, 147(12), 1810.
71. Smirnova, I., Reitzig, M., and Sorenson, O. (2022). Building status in an online community. *Organization Science*, 33(6), 2519-2540.
72. Srivanasan, K. (2023). Paying Attention. Working paper.

73. Stryker, J. E., Moriarty, C. M., and Jensen, J. D. (2008). Effects of newspaper coverage on public knowledge about modifiable cancer risks. *Health communication*, 23(4), 380-390.
74. Sumner, P., Vivian-Griffiths, S., Boivin, J., Williams, A., Bott, L., Adams, R., ... , and Chambers, C. D. (2016). Exaggerations and caveats in press releases and health-related science news. *PloS one*, 11(12), e0168217.
75. Toubia, O., and Stephen, A. T. (2013). Intrinsic vs. image-related utility in social media: Why do people contribute content to Twitter? *Marketing Science*, 32(3), 368–392.
76. Thaler, M. (2024) The Supply of Motivated Beliefs. Manuscript.
77. Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359 (6380), 1146-1151.
78. Wu, T. (2017). *The Attention Merchants: The Epic Scramble to Get Inside our Heads*. Knopf, USA.
79. Yavchitz, A., Boutron, I., Bafeta, A., Marroun, I., Charles, P., Mantz, J., and Ravaud, P. (2012). Misrepresentation of randomized controlled trials in press releases and news coverage: a cohort study. *PLoS medicine*, 9 (9), e1001308.
80. Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review* 110 (2): 337–363.