

NBER WORKING PAPER SERIES

THE ROLE OF ADVERTISERS AND PLATFORMS IN MONETIZING MISINFORMATION:
DESCRIPTIVE AND EXPERIMENTAL EVIDENCE

Wajeeha Ahmad
Ananya Sen
Chuck Eesley
Erik Brynjolfsson

Working Paper 32187
<http://www.nber.org/papers/w32187>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2024

We are grateful to our data partners (NewsGuard, the Global Disinformation Index, and Oracle) for sharing data with us. We thank the Executive Education Programs at the Stanford Graduate School of Business and Heinz College at Carnegie Mellon University for partnering with us and Tran Le for providing research assistance on part of this work. We gratefully acknowledge feedback and comments on earlier versions of this work from Ginger Jin, Joy Wu, Meng Liu, Manuella Collis, Rafael Jimenez Duran, Ryan Moore, Ruth Appel, Saharsh Agarwal, Sameer Borwankar, Wyatt Lee and from participants at the Queen's Workshop on the Economics of Media 2022, 2022 MIT Conference on Digital Experimentation, Workshop on Information Systems Economics 2022, Rotman School of Management PhD Strategy Seminar 2023, ISMS Marketing Science Conference 2023, The Sumantra Ghoshal Conference at London Business School 2023, Platform Strategy Research Symposium at Boston University 2023, NBER SI 2023 Digital Economics and Artificial Intelligence, Social Science Research Council Workshop on the Economics of Social Media 2023, West Coast Research Symposium 2023, Stanford University Trust and Safety Research Conference 2023, Conference on Information Systems and Technology 2023, HKU Business School Management and Strategy Seminar 2023, and Boston University Online Research Seminar on Digital Businesses 2024. This research was supported in part by the Stanford Digital Economy Lab, Stanford McCoy Family Center for Ethics in Society, Stanford Impact Labs, Stanford Technology Ventures Program, and the Economics of Digital Services initiative at the University of Pennsylvania. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w32187>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Wajeeha Ahmad, Ananya Sen, Chuck Eesley, and Erik Brynjolfsson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Role of Advertisers and Platforms in Monetizing Misinformation: Descriptive and Experimental Evidence

Wajeeha Ahmad, Ananya Sen, Chuck Eesley, and Erik Brynjolfsson

NBER Working Paper No. 32187

March 2024

JEL No. D8,M30,O3

ABSTRACT

The financial motivation to earn advertising revenue by spreading misinformation has been widely conjectured to be among the main reasons misinformation continues to be prevalent online. Research aimed at reducing the spread of misinformation has so far focused on user-level interventions with little emphasis on how the supply of misinformation can itself be countered. In this work, we show how online misinformation is largely financially sustained via advertising, examine how financing misinformation affects the advertisers and ad platforms involved and outline ways of reducing the financing of misinformation. First, we find that advertising on misinformation outlets is pervasive for companies across several industries and is amplified by digital ad platforms that automatically distribute companies' ads across the web. Using an information provision survey experiment, we show that people decrease their demand for a company's products or services upon learning about its role in monetizing misinformation via online ads. To shed light on why misinformation continues to be monetized despite the potential backlash for the advertisers involved, we survey decision-makers at companies. We find that most decision-makers are unaware of their companies' ads appearing on misinformation websites but have a strong preference to avoid appearing on such websites. Moreover, those uncertain and unaware about their role in financing misinformation increase their demand for a platform-based solution to reduce monetizing misinformation upon learning about how platforms amplify ad placement on misinformation websites. We identify low-cost, scalable information-based interventions that digital platforms could implement to reduce the financial incentive to misinform and counter the supply of misinformation online.

Wajeeha Ahmad
Department of Management Science
and Engineering,
Stanford University
Stanford, CA 94305
wajeehaa@stanford.edu

Ananya Sen
Heinz College of Information Systems
and Public Policy
Carnegie Mellon University
Pittsburgh, PA 15217
ananyase@andrew.cmu.edu

Chuck Eesley
Department of Management Science
and Engineering
Stanford University
cee@stanford.edu

Erik Brynjolfsson
Stanford Digital Economy Laboratory
353 Jane Stanford Way, Office 136
Stanford, CA 94305
and NBER
erik.brynjolfsson@gmail.com

1 Introduction

The prevalence of online misinformation can have significant social consequences, such as contributing to greater fatalities during the COVID-19 pandemic (Bursztyrn et al., 2020), exacerbating the climate crisis (Van der Linden et al., 2017), and sowing political discord (McCarthy, 2021). Yet, the supply of misinformation is often financially motivated. The economic incentive to produce misinformation has been widely conjectured by academics and practitioners alike as one of the main reasons misinformation websites, masquerading as legitimate news outlets, continue to be prevalent online (Blumberg, 2023; Guess et al., 2019; Lazer et al., 2018; Mosseri, 2017). During the 2016 U.S. Presidential election, one operator of a misinformation outlet openly stated: “For me, this is all about income” (Higgins et al., 2016).

Several media reports have anecdotally observed that companies and digital platforms contribute towards financially sustaining misinformation outlets via advertising (Hao, 2021; Giansiracusa, 2021). Advertising companies can directly place their ads on specific websites or use digital ad platforms to distribute their ads across the internet. The vast majority of online display advertising today is done via digital ad platforms that automatically distribute ads across millions of websites (Austin et al., 2019), which may include misinformation outlets. According to an industry estimate, for every \$2.16 in digital ad revenue sent to legitimate newspapers, U.S. advertisers send \$1 to misinformation sites, thereby financing such outlets (NewsGuard, 2021).

Existing work to counter the proliferation of misinformation online has primarily focused on empowering individual consumers of online news (Arechar et al., 2023; Lazer et al., 2018). In particular, previous studies have evaluated interventions such as fact-checking news articles (Aslett et al., 2022), providing crowd-sourced labels (Pennycook and Rand, 2019), and nudging users to share more accurate content (Pennycook and Rand, 2022). However, a vital question remains as to how the incentive to produce misinformation may be countered. Indeed, recently, academics have proposed ‘supply-side’ policies for steering platforms away from the revenue models that might contribute towards sustaining harmful content (Romer, 2019). Digital platforms have also attempted to decrease ad revenue going to some misinformation websites at times (Love and Cooke, 2016; Hiar, 2021). However, despite these attempts, ads from well-known companies and organizations continue to appear on misinformation websites (Hsu and Tracy, 2021; Grant and Myers, 2023). The financing of misinformation is likely to exacerbate as generative AI technologies lower the barrier to creating large volumes of content without meeting journalistic standards in order to maximize potential advertising revenue (Ryan-Mosley, 2023; Milmo and Hern, 2023).

In this paper, we attempt to provide a first step in understanding how to limit the financing of online misinformation by providing descriptive and experimental evidence about advertising on misinformation outlets. To tackle the problem of financing online misinformation, it is important first to understand the role of different entities within this ecosystem. In particular, we need to establish whether companies directly place ads on misinformation outlets or do so by automating ad placement through digital ad platforms. Although several mainstream digital platforms generate the vast majority of their revenue via advertising (Lazer et al., 2018; Casadesus-Masanell and Zhu, 2013), little is understood about the role of ad-driven platforms in financing misinformation. To evaluate the relative roles played by advertising companies and digital ad platforms in

monetizing misinformation, we construct a unique large-scale dataset by combining data on websites publishing misinformation with advertising activity per website over a period of three years.

Next, the extent to which companies can be dissuaded from advertising on misinformation websites depends on how their customers respond to information about the prevalence of companies' ads on such websites. As people find out about companies advertising on misinformation websites via news and social media reports (DeGeurin, 2023; Gomes Ribeiro et al., 2022; Dua, 2021; Crovitz, 2020), they may reduce their demand for such companies or voice concerns against such practices online (Hirschman, 1970; Gans et al., 2021). To measure these effects, we conduct a randomized survey experiment with a representative sample of the U.S. population. Since peoples' responses to companies advertising on misinformation websites could vary depending on the roles played by advertising companies and digital ad platforms, we randomly vary the pieces of factual information we provide to participants. By simultaneously measuring how people change their consumption of a company's products and the types of actors (i.e. advertisers or digital ad platforms) people voice concerns about, we capture how peoples' reactions change as the degree to which advertisers and ad platforms are held responsible varies. We also study how consumer responses may vary depending on the intensity of a company's advertising on misinformation websites by providing company rankings on this dimension.

Finally, whether decision-makers within companies are aware of their company's ads appearing on misinformation outlets and prefer to avoid doing so can play an important role in curbing the financing of misinformation. In recent years, advertisers have often participated in boycotts of ad-driven platforms such as YouTube, Facebook and Twitter for placing their ads next to problematic content (D'Onfro, 2019; Hsu and Lutz, 2020; Hsu, 2022). However, there is little systematic measurement of the knowledge and preferences of key decision-makers within companies regarding advertising on misinformation websites. To address this gap, we surveyed executives and managers by reaching out to the alumni of executive education programs. Moreover, we conduct an information provision experiment to examine whether decision-makers would increase their demand for a platform-based solution to avoid advertising on misinformation outlets when informed about the role played by digital ad platforms in monetizing misinformation.

We report three sets of findings from our descriptive and experimental analyses. First, our descriptive analysis suggests that misinformation websites are primarily monetized via advertising revenue with a substantial proportion of digital advertisers across several industries appearing on such websites. We also show that the use of digital ad platforms amplifies the financing of misinformation. Second, we find that advertising on misinformation websites can impose substantial costs on the companies and platforms involved once consumers find out about the roles they play in financing misinformation online. We find that consumers switch away from using companies whose ads appear on misinformation outlets. This switching effect persists even when consumers are informed about the role played by digital ad platforms in placing companies' ads on misinformation websites and the role played by other advertising companies in financing misinformation. These findings indicate that advertisers may have an incentive to avoid advertising on misinformation websites given the potential for substantial consumer backlash as people find out about their ads appearing on misinformation websites. Third, our survey of decision-makers suggests that most of them are ill-informed about the roles played by their own company and the digital ad platforms they use in financing misinformation out-

lets. However, decision-makers report a high demand for information on learning whether their ads appeared on misinformation outlets and solutions to avoid doing so. Those uncertain and unaware about where their ads appeared also increased their demand for a platform-based solution to reduce advertising on misinformation websites upon learning how platforms amplify ad placement on such websites. These results suggest that several advertising companies may be financing misinformation inadvertently. Upon access to relevant information, decision-makers within companies show interest in reducing the monetization of misinformation.

Altogether, our results indicate that there is room for decreasing the financing of misinformation using two low-cost, scalable interventions that provide greater transparency to consumers and advertisers about where companies' ads appear online. First, improving transparency for advertisers about where their ads appear could by itself reduce unintended advertising on misinformation websites, especially among companies who were unaware of their ads appearing on such outlets. Second, while it is currently possible for consumers to find about advertising companies financing misinformation through news and social media reports, platforms could make advertising on misinformation outlets more easily and continuously traceable to the advertising companies involved for consumers. Our results show that both information labels and company rankings could be used to reduce consumer demand away from companies advertising on misinformation websites. This could provide a stronger incentive for companies to steer their ads away from such outlets given the potential for consumer backlash resulting from increased visibility.

The rest of this paper proceeds as follows. Section 2 describes our contributions to the literature. We outline the empirical context, data and descriptive findings in Section 3. In Section 4, we describe the design and results of our consumer experiment. Section 5 presents the design and results of our decision-maker survey. Finally, Section 6 concludes with a discussion of our findings.

2 Related literature

Our paper contributes to several strands of academic literature. Prior research has alluded to the importance of the advertising business model in sustaining news outlets (Guess et al., 2019; Lazer et al., 2018; Casadesus-Masanell and Zhu, 2010, 2013). Researchers have examined the types of ads appearing on misinformation websites (Papadogiannakis et al., 2023; Kohno et al., 2020), the infrastructure supporting misinformation websites (Han et al., 2022), and the structure of the programmatic advertising ecosystem (Braun and Eklund, 2019). These papers provide an important first step but are often based on a small sample of websites or advertisers and are limited to a short time window of a few days. Relative to this, we provide the first large-scale evidence of the ecosystem sustaining online misinformation by constructing a unique dataset of thousands of news outlets and advertisers over a period of three years. Moreover, our data allows us to analyze the relative roles of advertising companies and digital ad platforms in placing ads on misinformation websites.

Previous work has examined the conditions under which people react against companies for failing to operate up to their expectations (Hirschman, 1970; Broccardo et al., 2022; Du et al., 2011). For example, studies have examined the consequences on consumer responses when service quality deteriorates (Gans et al., 2021) or when a company or its leadership takes a political stance (Liaukonytė et al., 2022; Chatterji and Toffel,

2019). In the online advertising context, prior research has examined the effects of advertising on sales and consumer behavior under various scenarios, such as companies reaching out to people during high-stress times, debunking false claims in ads, and posting sensational content in ads (Fong et al., 2022; Bellman et al., 2018; Lull and Bushman, 2015). Gomes Ribeiro et al. (2022) suggest that when an activist group targets companies for advertising on misinformation websites, tweets mentioning the company temporarily become more toxic and less positive. Similar to (Gomes Ribeiro et al., 2022), we analyze how consumers respond when they learn about companies advertising on misinformation websites, and do so in an incentive-compatible manner. Incentive compatibility is captured by measuring behavioral outcomes at the individual level to capture both “exit” and “voice”, the two types of potential consumer responses theorized in the literature (Hirschman, 1970). Additionally, we test how consumer reactions differ when informed about the different actors involved in financing misinformation, such as digital ad platforms and other advertising companies, which has direct implications for managers and policy makers.

We also contribute to a literature strand showing how digital platforms create externalities for different players within the platform ecosystem. Prior research shows that platform and advertiser incentives are not aligned regarding ad effectiveness (Johnson and Lewis, 2015; Frick et al., 2022; Agarwal and Mukhopadhyay, 2016). More broadly, digital platforms can create other negative indirect externalities in the advertising context, e.g., when more advertisers on a search engine platform decrease its value for searchers of independent advice (De Reuver et al., 2018). We extend this literature to show that ad platforms could create reputational externalities for advertisers since they are about ten times more likely to appear on misinformation when using digital ad platforms which, according to our experimental results, could alienate their consumers. Even when informed about the role played by digital ad platforms in placing companies’ ads on misinformation websites, people switch their consumption away from companies whose ads appear on such websites 2.5 times more than the control group.

“Demand-side” interventions to counter online misinformation studied in prior work have focused on reducing the consumption and spread of misinformation among consumers of news on online platforms. While interventions such as accuracy prompts and digital literacy tips can increase the quality of news that people share (Arechar et al., 2023), this line of work has found limited support for news credibility signals in increasing the demand for credible news (Chopra et al., 2022) or in reducing misperceptions among users Aslett et al. (2022). Moreover, such interventions are only effective for the small subset of users who are exposed to misinformation (Allen et al., 2020). In our ‘supply-side’ approach, we target entities and individuals who might not necessarily consume or spread misinformation themselves. Moreover, the consumer-focused intervention we identify targets a broader, potentially more persuadable set of people who use products from companies that finance misinformation, which is arguably much larger than the limited set of the general public who consume misinformation themselves.¹

Existing supply-side policy interventions proposed to curb misinformation include social media platforms banning the promotion of false news. For instance, Facebook’s ban on the advertising of fake news resulted in

¹While less than 30% of the participants in our consumer experiment had exposure to one or more misinformation websites based on self-reported data, over 95% had consumed one or more products from our six gift card companies during the past year.

a decline in the subsequent sharing of fake news on Facebook relative to Twitter (Chiou and Tucker, 2018). While such approaches can prevent misinformation outlets from using specific social media platforms to promote their content, they do not prevent misinformation sites from generating advertising revenue. Another proposed supply-side approach advocates for imposing taxes on ad revenue as a means of incentivizing digital platforms to shift their business model from advertising towards a “healthier, more traditional” subscription-based model (Romer, 2019). Relative to interventions that involve platforms banning specific news outlets or changing their ad-driven business model altogether, we identify a middle path that accounts for advertisers’ and consumers’ preferences to help counter the financing of online misinformation.

3 Descriptive evidence

3.1 Background on Digital Advertising

The predominant business model of several mainstream digital media platforms relies on monetizing attention via advertising (Lazer et al., 2018). While these platforms typically offer free content and services to individual consumers, they generate revenue by serving as an intermediary or ad exchange connecting advertisers with independent websites that want to host ads. To do so, platforms run online auctions to algorithmically distribute ads across websites, known as “programmatic advertising”. For example, Google distributes ads in this manner to over 2 million non-Google sites in what is known as the Google Display Network. In this way, the websites receive payment from advertisers for hosting ads, and they share a percentage of this payment with the platform. In the U.S., more than 80% of digital display ads are programmatic ads (Austin et al., 2019). We refer to these ad exchanges as digital ad platforms and use the term digital platforms to collectively refer to all the services offered by such media platforms.²

While in other forms of (offline) media, advertisers typically have significant control over where their ads appear, ad placement through digital ad platforms is mainly automated. Since most companies do not have the capacity to participate in high-frequency ad auctions that require them to place individual bids for each ad slot they are interested in, they typically outsource the bidding process to an ad platform (Frick et al., 2022). Such programmatic advertising gives companies relatively less control over where their ads end up online. However, advertising companies can take steps to reduce advertising on misinformation websites, such as by only being part of ad auctions for a select list of credible websites or blocking ads from appearing on specific lists of misinformation outlets.

3.2 Data

To categorize whether a website contains misinformation, we compiled a list of misinformation domains using three different sources. First, we use a dataset maintained by NewsGuard. This company rates all the news and information websites that account for 95% of online engagement in each of the five countries where it operates. Journalists and experienced editors manually generate these ratings by reviewing news and information

²Our empirical context is similar to other papers studying digital advertising such as Cowgill and Dorobantu (2018), Grewal et al. (2022), and Frick et al. (2022).

websites according to nine apolitical journalistic criteria.³ Recent research has used this dataset to identify misinformation websites (Edelson et al., 2021; Aslett et al., 2022; Bhadani et al., 2022; Moore et al., 2023). In this paper, we consider each website that NewsGuard rates as *repeatedly publishing false content* between 2019 and 2021 to be a misinformation website and all others to be non-misinformation websites, leading to a set of 1,546 misinformation websites and 6,499 non-misinformation websites.⁴ Table A1 summarizes the characteristics of this dataset. Our NewsGuard dataset contains websites across the political spectrum, including left-leaning websites (e.g., palmerreport.com, occupydemocrats.com), politically neutral websites (e.g., rt.com, nationalenquirer.com), and right-leaning websites (e.g., thegatewaypundit.com, theconservativetreehouse.com).

In addition to the NewsGuard dataset, we use a list of websites provided by the Global Disinformation Index (GDI). This non-profit organization identifies disinformation by analyzing both the content and context of a message, and how they are spread through networks and across platforms (Decker, 2019). In this way, GDI maintains a list of monthly-updated websites, which it also shares with interested ad tech platforms to help reduce advertising on misinformation websites. The GDI list allows us to identify 1,869 additional misinformation websites. Finally, we augment our list of misinformation websites with 396 additional ones used in prior work (Guess et al., 2020; Allcott et al., 2019). Altogether, our website dataset consists of 10,310 websites, including 3,811 misinformation and 6,499 non-misinformation websites.⁵

Similar to prior work (Moore et al., 2023; Aslett et al., 2022), our final measure of misinformation is at the level of the website or online news outlet. The different sources we use employ article-level information as well as website-level metadata to provide an aggregated metric at the website level. This is a meaningful approach since it reduces the potential for noise due to aggregation over a number of articles from each news outlet. Both NewsGuard and the GDI use a combination of automated and manual methods to source websites to evaluate, but each website in their data is rated manually by expert professionals who apply journalistic standards to evaluate online news outlets in a neutral, transparent and independent manner.

We use novel data on advertiser behavior from Oracle’s Moat Pro platform, which includes data collected by crawling approximately ten thousand websites daily to create a snapshot of the advertising landscape and the players in the space. Moat’s web crawlers mirror a normal user experience and attempt to visit a representative sample of pages for each website at least once a day. To the best of our knowledge, this data is the gold standard and is used by a large number of industry players. We use the Moat platform to collect data from 2019 to 2021. For all the websites in our sample that get non-zero traffic throughout this period

³These criteria include four metrics for assessing a site’s credibility (“Does not repeatedly publish false content”, “Gathers and presents information responsibly”, “Regularly corrects or clarifies errors”, “Handles the difference between news and opinion responsibly”, “Avoids deceptive headlines”) and three transparency-related metrics (“Website discloses ownership and financing”, “Clearly labels advertising”, “Reveals who’s in charge, including possible conflicts of interest”, “The site provides the names of content creators, along with either contact or biographical information”). More information can be found at the NewsGuard website: <https://www.newsguardtech.com/ratings/rating-process-criteria/>

⁴NewsGuard began collecting data in 2019. To get coverage throughout our study period, we sample websites provided by NewsGuard from the start, middle and end of each year from 2019 to 2021. Additionally, we also sample websites from January 2022 and June 2022 to account for websites that may have existed during our study period and discovered later.

⁵Among the websites that NewsGuard rated as non-misinformation (at any point in our sample), 310 websites were considered to be misinformation websites by our other sources or by NewsGuard itself (during a different period in our sample). We categorize these websites as misinformation websites given their risk of producing misinformation.

and have advertising data available,⁶ we collected monthly data on the advertising companies appearing on each website and digital ad platforms used by each website.

Our final dataset, which has data on advertising and misinformation, consists of 5,485 websites, of which 1,276 are misinformation websites and the remaining 4,209 are non-misinformation websites. Additionally, for the most active 100 advertisers each year as identified by Moat Pro, we collected weekly data on the websites they appeared on and the digital ad platforms they used.

3.3 Descriptive results

Most misinformation websites in our sample (74%) were supported by advertising revenue between 2019 and 2021.⁷ Moreover, among websites rated by NewsGuard, a much smaller percentage of misinformation websites had a paywall (2.7% in the U.S. and 3.2% globally) relative to non-misinformation websites (25.0% in the US and 24.0% globally).⁸ These findings suggest that relative to other websites, misinformation websites less often rely on subscription-based business models that require paywalls. Given that advertising appears to be the dominant business model sustaining misinformation outlets, it merits a closer look. Next, we examine the roles played by advertising companies and digital ad platforms in financing misinformation outlets.

3.3.1 The role of advertising companies

To examine the level of advertising on misinformation websites, we collect data on advertisers appearing on each of the 5,485 websites in our dataset. Of the 42,595 unique advertisers on these websites, about 44% appear on misinformation websites. Focusing on the one hundred most active advertisers each year, we find that 55% of these appear on misinformation websites weekly.

Figure 1 shows the proportion of companies by industry and the number of times they appear on the websites in our dataset between 2019 and 2021. As shown in Figure 1, advertising companies that appear on misinformation websites span a wide range of industries. These include several well-known brands among commonly used household products, technology products, and business services as well as finance, health, government, and educational institutions among other industries. A substantial proportion of companies in each industry appear on misinformation websites with the majority of companies advertising on misinformation websites in most industries ads (Figure 1a). Further, the intensity of advertising on misinformation sites is similar to that on non-misinformation sites for companies across several industries (Figure 1b).

3.3.2 The role of digital ad platforms

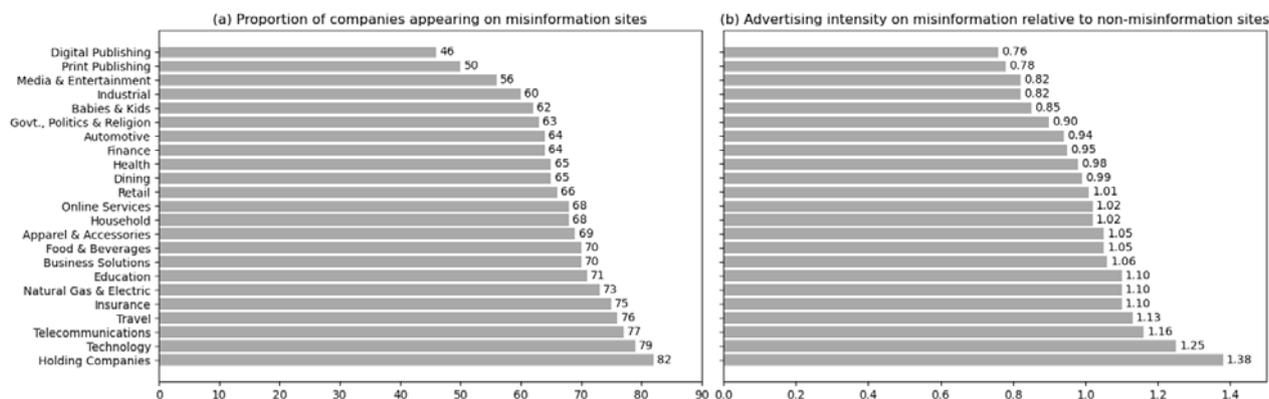
For the one hundred most active advertisers in each year, we collected weekly data on which websites their ads appeared on and their use of digital ad platforms. About 80% of advertisers that used digital ad platforms appeared on misinformation websites. In contrast to this figure, among companies that do not use digital

⁶We use data from SEMRush, a leading online analytics platform, to determine the level of monthly traffic received by each website from 2019 to 2021.

⁷Most non-misinformation websites in our sample (94%) also received advertising revenue during this period.

⁸As shown in Table A1, relative to non-misinformation websites, misinformation websites were also more likely to be operated by individuals as opposed to corporate or non-profit entities.

Figure 1: Advertising Companies Appearing On Misinformation Websites by Industry



Notes: From 2019 to 2021, we record the number of times companies in a given industry appeared on all 5485 websites in our sample per month. We remove industries where the number of ad appearances by all companies combined was below the fifth percentile of the total number of ad appearances. Figure 1(a) shows the proportion of companies in each industry that appear on misinformation websites at least once in our sample. In Figure 1(b), we calculate the relative advertising intensity on misinformation sites by dividing the proportion of ads in a given industry that appear on all misinformation sites with the proportion of ads appearing on all non-misinformation sites. Therefore, values lower than indicate less, values close to 1 represent similar and values higher than 1 represent greater advertising intensity on misinformation sites relative to non-misinformation websites.

ad platforms in a given week, only approximately 8% appear on misinformation websites. In other words, companies that used digital ad platforms were about 10 times more likely to appear on misinformation websites than companies that did not use digital ad platforms. Moreover, we account for industry and time trends to find that the use of digital ad platforms by companies substantially amplifies the likelihood of a company's ads appearing on misinformation websites (Table A3).

4 Consumer experiment

Our descriptive results show evidence of companies across several industries advertising on misinformation websites with digital ad platforms amplifying the financing of misinformation online. We now examine how advertising on misinformation websites affects the advertisers and platforms involved once consumers find out about their role in financing misinformation. People can find out about companies advertising on misinformation through news reports that specifically name which companies' ads appear on misinformation outlets (Crovitz, 2020; Dua, 2021). Additionally, several activist and non-profit organizations such as Sleeping Giants, CheckMyAds.org, the #StopHateforProfit campaign and the Global Disinformation Index periodically mention companies contributing towards financing misinformation through social media and published reports on their websites. Therefore, it is important to measure the preferences of the people who consume a company's products or services regardless of whether these consumers visit misinformation websites themselves.

4.1 Research design

Our survey experiment aims to determine potential changes in consumer behavior based on experimentally varied information about the roles of companies and platforms in financing misinformation via advertising.

Using the framework of Hirschman (1970), we measure how people 1) exit, i.e. decrease their consumption and 2) voice concerns about company or platform practices via online petitions in response to the information provided in an incentive-compatible manner.⁹

This study was reviewed by the Stanford University Institutional Review Board (Protocol No. IRB-63897) and the Carnegie Mellon University Institutional Review Board (Protocol No. IRB0000603). Our study was preregistered at the American Economic Association's Registry under AEARCTR-0009973.¹⁰ Informed consent was obtained from all participants at the beginning of the survey.

4.1.1 Setting and sample size

We recruit a representative sample of U.S. internet users via CloudResearch.¹¹ CloudResearch screened respondents for our study so that they are representative of the US internet population in terms of age, gender and race based on the US Census (2020). To ensure data quality, we include a screener in our survey to check whether participants pay attention to the information provided. Only participants who pass this screener can proceed with the survey. Our total sample includes approximately 4,000 participants, who are randomized into five groups with about 800 participants per group.

The flow of the survey study is shown in Figure A7. We begin by asking participants to report demographics such as age, gender and residence. From a list of trustworthy and misinformation outlets, we then ask participants questions about their behaviors in terms of the news outlets they have used in the past 12 months, their trust in the media (on a 5-point scale), the online services or platforms they have used and the number of petitions they have signed in the past 12 months.

4.1.2 Initial gift card preferences

We then inform participants that one in five (i.e. 20% of all respondents) who complete the survey will be offered a \$25 gift card from a company of their choice out of six company options. Respondents are asked to rank the six gift card companies on a scale from their first choice (most preferred) to their sixth choice (least preferred). These six companies belong to one of three categories: fast food, food delivery and ride-sharing. All six companies appeared on the misinformation websites in our sample during the past three years (2019-2021), offer items below \$25, and are commonly used throughout the U.S. The order in which the six companies are presented is randomized at the respondent level.¹² We then ask participants to confirm which gift card they would like to receive if they are selected to ensure they have consistent preferences regardless of how the question is asked.¹³

⁹Incentive compatibility is captured by measuring behavioral outcomes at the individual level.

¹⁰<https://www.socialscienceregistry.org/trials/9973>

¹¹CloudResearch is a data provider used in survey research that is more diverse and provide higher data quality than other providers such as MTurk (Chandler et al., 2019; Eyal et al., 2021).

¹²As a robustness check, we also ask respondents to assign weights to each of the six gift card options. This question gives respondents greater flexibility by allowing them to indicate the possibility of indifference (i.e., equal weights) between any set of options.

¹³At this initial elicitation stage, the respondents did not know that they will get another chance to revise their choice. Hence, these choices can be thought of as capturing revealed preference.

4.1.3 Information treatments

All participants in the experiment are given baseline information on misinformation and advertising as shown in Figure A10. This is meant to ensure that all participants in our experiment are made aware of how we define misinformation along with examples of a few misinformation websites (including right-wing, neutral and left-wing misinformation websites), how misinformation websites are identified, and how companies advertise on misinformation websites (via an illustrative example) and use digital platforms to automate placing ads.

Participants are then randomized into one control and four treatment groups, in which the information treatments are all based on factual information from our data and prior research. We use an active control design to isolate the effect of providing information relevant to the practice of specific companies on people's behavior (Haaland et al., 2023). Participants in the control group are given generic information based on prior research that is unrelated to advertising companies or platforms but relevant to topic of news and misinformation.

In our first “company only” treatment group (T1), participants are given factual information stating that ads from their top choice gift card company appeared on misinformation websites in the recent past. Based on their preferences, people may change their final gift card preference away from their initial top-ranked company after receiving this information. It is unclear, however, whether advertising on misinformation websites would cause a sufficient change in consumption patterns and which sets of participants may be more affected.

Our second “platform only” treatment group (T2) informs participants that companies using digital ad platforms were about 10 times more likely to appear on misinformation websites than companies that did not use such platforms in the recent past. This information treatment measures the effects of digital ad platforms in financing misinformation news outlets. Since it does not contain information about advertising companies, it practically serves as a second control group for our company-level outcome and aims to measure how people may respond to our platform-related outcome.

Because our descriptive data suggest that the use of digital ad platforms amplifies advertising revenue for misinformation outlets, we are interested in measuring how consumers respond to a specific advertising company appearing on misinformation websites when also informed of the potential role played by digital ad platforms in placing companies' ads on misinformation websites. It is unclear whether consumers will attribute more blame to companies or ad platforms for financing misinformation websites when informed about the role of the different players in this ecosystem. For this reason, our third “company and platform” treatment (T3) combines information from our first two treatments (T1 and T2). Similar to T1, participants are given factual information that ads from their top choice gift card company appeared on misinformation websites in the recent past. Additionally, we informed participants that their top choice company used digital ad platforms and companies that used such platforms were about ten times more likely to appear on misinformation websites than companies that did not use digital ad platforms, as mentioned in T2.

Finally, since several advertising companies appear on misinformation websites, we would like to determine whether informing consumers about other advertising companies also appearing on misinformation websites

changes their response towards their top choice company. In our fourth “company ranking” treatment (T4), participants are given factual information that ads from all six gift card companies appeared on misinformation websites in the recent past, along with a ranking based on the order of their intensity of advertising on misinformation websites. We personalize these rankings by providing truthful information based on data from different years in the recent past such that the respondents’ top gift card choice company does not appear last in the ranking (i.e., is not the company that advertises least on misinformation websites) and in most cases, advertises more intensely on misinformation websites than its potential substitute in the same company category (e.g., fast food, food delivery or ride-sharing).¹⁴ Such a treatment allows us to measure potential differences in the direction of consumers switching their gift card choices, such as switching towards companies that advertise more or less intensely on misinformation websites. It could also give consumers reasonable deniability such as “everyone advertises on misinformation websites” leading to ambiguous predictions about the exact impact of the treatment effect.

4.1.4 Outcome Measures

We measure two behavioral outcomes that collectively allow us to measure how people respond to our information treatments in terms of both voice and exit (Hirschman, 1970). After the information treatment, all participants are asked to make their final gift card choice from the same six options they were shown earlier. To ensure incentive compatibility, participants are (truthfully) told that those randomly selected to receive a gift card will be offered the gift card of their choice at the end of our study. As mentioned above, the probability of being randomly chosen to receive a gift card is 20%. We choose a high probability of receiving a gift card relative to other online experiments since prior work has shown that consumers process choice-relevant information more carefully as realization probability increases (Cao and Zhang, 2021). Our main outcome of interest is whether participants “exit” or switch their gift card preference, i.e., whether they select a different gift card after the information treatment than their top choice indicated before the information treatment.

Secondly, participants are given the option to sign one of several real online petitions that we made and hosted on Change.org. Participants can opt to sign a petition that advocates for either blocking or allowing advertising on misinformation or choose not to sign any petition. Further, participants could choose between two petitions for blocking ads on misinformation websites, suggesting that either 1) advertising companies or 2) digital ad platforms need to block ads from appearing on misinformation websites.¹⁵ To track the number of petition signatures across our randomized groups, we provide separate petition links to participants in each randomized group. We record several petition-related outcomes. First, we measure participants’ intention to sign a petition based on the option they select in this question. Participants who pass our attention check and opt to sign a petition are later provided with a link to their petition of choice. This allows tracking whether

¹⁴As depicted in Figure A15, respondents are told that “In the recent past, ads from all six companies below repeatedly appeared on misinformation websites in the following order of intensity” and provided with a ranking from one of three years in our study period, i.e., 2019, 2020, or 2021.

¹⁵Participants select among the following five choices: 1. “Companies like X need to block their ads from appearing on misinformation websites.”, where X is their top choice gift card company; 2. “Companies like X need to allow their ads to appear on misinformation websites.”, where X is their top choice gift card company; 3. “Digital ad platforms used by companies need to block ads from appearing on misinformation websites.”; 4. “Digital ad platforms used by companies need to allow ads to appear on misinformation websites.”; and 5. *I do not want to sign any petition.*

participants click on the petition link provided. Participants can also self-report whether they signed the petition. Finally, for each randomized group, we can track the total number of actual petition signatures.

Our petition outcomes serves two purposes. While our gift card outcome measures how people change their consumption behavior in response to the information provided, people may also respond to our information treatments in alternative ways, e.g. by voicing their concerns or supplying information to the parties involved (Hirschman, 1970; Gans et al., 2021; Lenox and Eesley, 2009). Given that the process of signing a petition is costly, participants' responses to this outcome would constitute a meaningful measure similar to petition measures used in prior experimental work (Grigorieff et al., 2020; Haaland and Roth, 2020). Second, since participants must choose between signing either company or platform petitions, this outcome allows us to measure whether or not, across our treatments, people hold advertising companies more responsible for financing misinformation than the digital ad platforms that automatically place ads for companies.

In addition to our behavioral outcomes, we also record participants' stated preferences. To do so, we ask participants about their degree of agreement with several statements about misinformation on a seven-point scale ranging from "strongly agree" to "strongly disagree". These include whether they think 1) companies have an important role in reducing the spread of misinformation through their advertising practices, and whether 2) digital platforms should give companies the option to avoid advertising on misinformation websites.

4.1.5 Dealing with Experimenter Demand Effects

In our incentivized, online setting where we measure behavioral outcomes, we expect experimenter demand effects to be minimal as has been evidenced in the experimental literature (De Quidt et al., 2018). We take several steps to mitigate potential experimenter demand effects, including incorporating several suggestions by Haaland et al. (2023).

First, our survey experiment has a neutral framing throughout the survey since the recruitment of participants. While recruiting participants, we invite them to "take a survey about the news, technology and businesses" without making any specific references to misinformation or its effects. While introducing misinformation websites and how they are identified by independent non-partisan organizations, we include examples of misinformation websites across the political spectrum (including both right-wing and left-wing sites) and provide an illustrative example of misinformation by foreign actors (Figure A10). In drafting the survey instruments, the phrasing of the questions and choices available were as neutral as possible.¹⁶ In presenting our information interventions and measuring our behavioral outcomes, we take special care to not highlight the names of the specific entities being randomized across groups to avoid emphasizing what is being measured (Appendix 6). We do, however, highlight our gift card incentives by putting the gift card information in bold text to ensure incentive compatibility since prior work has found that failing to make incentives conspicuous can vastly undermine their ability to shift behavior (John et al., 2022).

In our active control design, participants in all randomized groups are presented with the same baseline

¹⁶For example, while introducing our online petitions, we presented participants with the option to sign real petitions that suggest both blocking and allowing advertising on misinformation sites.

information about misinformation, given misinformation-related information in the information intervention and asked the same questions after the information intervention to emphasize the same topics and minimize potential differences in the understanding of the study across treatment groups.

To maximize privacy and increase truthful reporting (Ong and Weiss, 2000), respondents complete the surveys on their own devices without the physical presence of a researcher. We also do not collect respondents' names or contact details (with the exception of eliciting emails to provide gift cards to participants at the end of the study). Apart from making the above design choices to minimize experimenter demand effects, we measure their relevance using a survey question. Since demand effects are less likely a concern if participants cannot identify the intent of the study (Haaland et al., 2023), we ask participants an open-ended question, i.e., "What do you think is the purpose of our study?". Following Bursztyn et al. (2021) and Song (2022), we then analyze the responses to this question to examine whether they differ across treatment groups.

4.2 Results

4.2.1 Average Treatment Effects: Exit

Our primary outcome is whether respondents exit by switching their top gift card choice, which takes the value one for people who switch and the value zero for all other participants. To observe exit outcomes, we focus on company-related information treatments, i.e., all treatments (T1, T3 and T4) where respondents are informed that ads from their top choice gift card company recently appeared on misinformation websites.

Table 1 shows the regression results for our behavioral outcomes measured after participants receive the information treatment. Column 1 of Table 1 shows that respondents increasingly exit (i.e., increase switching away or decrease demand from) their first choice company by 13 percentage points ($p < 0.001$) relative to control in response to learning about their top choice gift card company's ads appearing on misinformation websites (T1). This effect persists when we control for participants' demographic and behavioral characteristics ($p < 0.001$, Column 2 of Table 1). We also use text analysis of the responses to a free-form question which helps identify the impact of the information intervention more directly. Respondents' text responses explaining their choice of gift card as shown in Figure 2 (a) reveal that misinformation concerns drive this switching behavior.¹⁷

Switching behavior also increases relative to the control group by ten percentage points ($p < 0.001$) when respondents are told about the substantial role played by digital ad platforms in placing companies' ads on misinformation websites (T3). This switching behavior persists even though respondents are more likely to state that digital ad platforms are responsible for placing companies' ads on misinformation websites by four percentage points ($p < 0.001$) relative to the control group (Figure 2b). This suggests that advertising companies can continue to experience a decline in demand for their products or services despite consumers knowing that digital ad platforms play a substantial role in placing companies' ads on misinformation websites.

When provided with a ranking of companies in order of their intensity of appearance on misinformation websites (T4), respondents switch away from opting for their top choice gift card company by seven percentage

¹⁷For more sample text responses and details about the text analysis methodology, see Table A6 in Appendix C.

Table 1: Average Treatment Effects on Exit

	Switch in preference		Switch to lower preference		Switch in category		Switch to lower misinformation	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Company (T1)	0.13*** (0.01)	0.13*** (0.01)	0.08*** (0.01)	0.08*** (0.01)	0.05*** (0.01)	0.05*** (0.01)	1.03** (0.48)	0.69* (0.38)
Platform (T2)	0.03*** (0.01)	0.03** (0.01)	0.01 (0.01)	0.01 (0.01)	0.02* (0.01)	0.01 (0.01)	0.52 (0.54)	0.23 (0.48)
Company and Platform (T3)	0.10*** (0.01)	0.10*** (0.01)	0.06*** (0.01)	0.06*** (0.01)	0.04*** (0.01)	0.04*** (0.01)	0.69 (0.49)	0.28 (0.38)
Company Ranking (T4)	0.08*** (0.01)	0.08*** (0.01)	0.06*** (0.01)	0.06*** (0.01)	0.03*** (0.01)	0.02** (0.01)	1.57*** (0.50)	0.95** (0.39)
Control group mean	0.04	0.04	0.02	0.02	0.03	0.03	0.61	0.61
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	4039	4039	4039	4039	4039	4039	430	430

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results for each of the four treatment groups (T1, T2, T3 and T4). In columns (1) and (2), the dependent variable is a binary variable that takes the value 1 when a participant switches their gift card choice from their top choice company after receiving the information treatment and is zero otherwise. In columns (3) and (4), the dependent variable is a binary variable that takes the value 1 when a participant switches their gift card choice from their top choice company to a company they prefer less (as measured by how participants assign weights to each of the six gift card choices that must all sum up to 100) and is zero otherwise. In columns (5) and (6), the dependent variable is a binary variable that takes the value 1 when a participant switches their gift card choice across product categories (e.g. from ride-sharing gift cards to a fast food gift card) and is zero otherwise. Columns (7) and (8) show regressions for the sub-sample of participants who switch their gift card choice; the dependent variable is the difference in the intensity of advertising misinformation between the participant's top choice gift card company and the company they finally choose after receiving the information treatment. In Columns (2), (4), (6) and (8), we include demographic controls (including the respondent's age, gender, region of residence within the US, race, education level, employment status, household income and whether the respondent voted for President Joseph Biden in the 2020 U.S. Presidential election) and behavioral controls (including the types of news sources consumed, whether the respondent had low trust in the news media, the number of online services used, whether the respondent had signed a petition in the past 12 months, whether the respondent reported using one or more misinformation news outlets from a list of 26 popular news outlets in the past 12 months, the respondent's top choice gift card and whether the respondent frequently uses their top choice gift card company). Robust standard errors in parentheses.

points ($p < 0.001$). This result shows that advertising companies can expect to face a decrease in consumption for financing misinformation despite other companies also advertising on misinformation outlets. Respondents are less likely to mention product features relevant to the companies they are interested in, e.g. healthy food, good prices, availability in local area, etc. by 7 percentage points ($p < 0.001$, Figure 2a). Examining the direction of consumer switching shows that among those who switch their gift card preference, those provided with company ranking information in T4 made the most switches towards companies that less frequently advertised on misinformation websites (Columns 7-8, Table 1). This result suggests that providing a ranking of advertising companies transparently could steer consumer demand towards companies that advertise less frequently on misinformation websites.

While our primary exit outcome is the switch in gift card choice, our results are robust to alternative measures as shown in Table 1. These exit outcomes, which include whether participants switch to a product they prefer less than their top choice one (Columns 3-4, Table 1) and whether they switch their choice across product categories (Columns 5-6, Table 1), further indicate that our measures of exit are incentive-compatible since participants incur a real cost of switching to a company that is not equivalent to their top-ranked one.

4.2.2 Average Treatment Effects: Voice

Next, we examine how participants respond to our information treatments by signing an online petition to voice their concerns about advertising on misinformation websites. While we observe actual petition signatures at the group level, we use clicks on petition links as our primary voice outcome, since this information is available at the individual level. Our results are robust to using alternative petition outcomes, such as intention to sign a petition, self-reported petition signatures, and actual signatures, as shown in Table A5 in Appendix C.

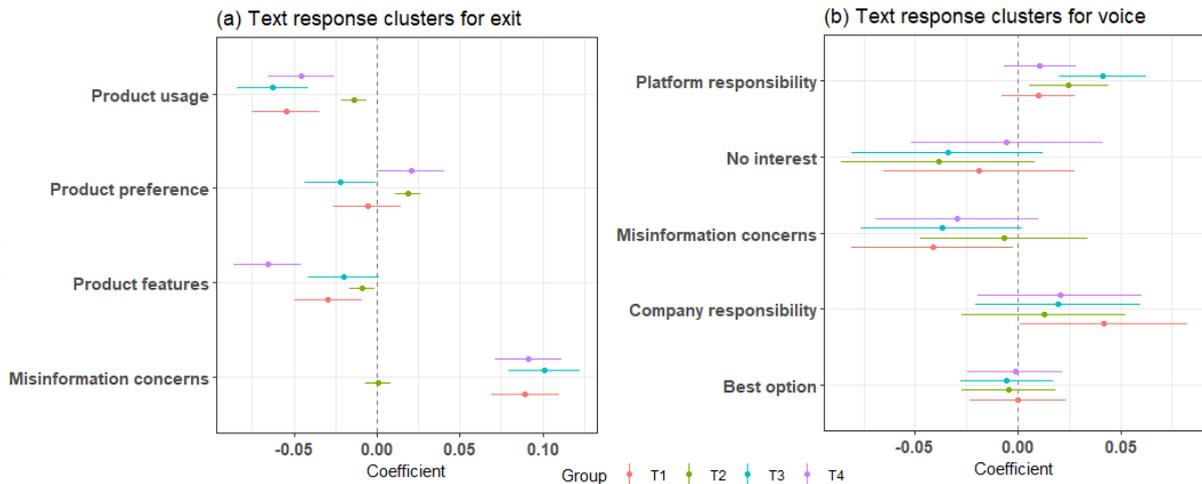
Table 2: Average Treatment Effects on Voice

	Company		Platform	
	(1)	(2)	(3)	(4)
Company (T1)	0.02 (0.02)	0.02 (0.02)	-0.02 (0.02)	-0.02 (0.02)
Platform (T2)	-0.01 (0.02)	-0.01 (0.02)	0.05*** (0.02)	0.05*** (0.02)
Company and Platform (T3)	-0.00 (0.02)	-0.00 (0.02)	-0.01 (0.02)	-0.01 (0.02)
Company Ranking (T4)	0.04** (0.02)	0.04** (0.02)	-0.03* (0.02)	-0.03 (0.02)
Control group mean	0.15	0.15	0.14	0.14
Controls	No	Yes	No	Yes
Observations	4039	4039	4039	4039

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results for each of the four treatment groups (T1, T2, T3 and T4). In columns (1) and (2), the dependent variable is clicking on a link to sign a petition that suggests that companies like the respondent's top choice gift card company need to block their ads from appearing on misinformation websites. In columns (3) and (4), the dependent variable is clicking on a link to sign a petition that suggests that digital ad platforms used by companies need to block ads from appearing on misinformation websites. We include the same baseline demographic and behavioral controls as those detailed in Table 1. Robust standard errors in parentheses.

Figure 2: Text Explanation Clustering by Randomized Treatment Group



Notes: This figure plots regression coefficients from OLS regressions of an indicator for cluster membership on each randomized group. The horizontal bars represent 95% confidence intervals. The topics along the y-axes are binary variables that take value 1 if a participant's response is classified into the given topic and zero otherwise. Details about the text analyses are mentioned in Appendix C and sample text responses are shown in Tables A6 and A7. Figure (a) shows OLS regression results for text analysis on the open-ended reasons participants mentioned while explaining their choice of gift card. Figure (b) shows OLS regression results for text analysis on the open-ended reasons participants mentioned while explaining their choice of online petition to sign. In all specifications above, we control for the same baseline demographic characteristics and behavioral characteristics as in Table 1. Robust standard errors in parentheses.

Relative to the control group, participants were 36% significantly more likely to click on the platform

petition link when given information about the role of digital ad platforms in automatically placing ads on misinformation websites in the Platform (T2) treatment group (Columns 3-4, Table 2). Text analysis from respondents' explanation of their petition choice confirms that respondents hold digital ad platforms more responsible for financing misinformation in T2 relative to the control group (Figure 2b). For example, one respondent stated, "Door Dash is not the only ad being put on misinformation sites. It is a larger issue that has to do with the platforms used to place ads."

Additionally, upon receiving information about all six gift card companies' ads appearing on misinformation websites (T4), participants are significantly more likely to click on petition links suggesting that advertising companies need to block their ads from appearing on misinformation websites (Columns 1-2, Table 2). Based on their open-ended text responses, respondents increasingly highlight misinformation-related concerns and place less emphasis on product usage and product features (Figure 2a).

4.2.3 Heterogeneous Treatment Effects

Next, we explore heterogeneity in treatment effects along four pre-registered dimensions: gender, political orientation, frequency of use of the company's products or services, and consumption of misinformation.

Prior research recognizes differences in the salience of prosocial motivations across gender (Croson and Gneezy, 2009; Falk et al., 2018) with women being more affected by social-impact messages than men (Guzman et al., 2020) and more critical consumers of new media content (Xiao et al., 2021). Given these findings, we could expect female participants to be more strongly affected by our information treatments. Indeed, while we observe positive treatment effects for both male and female participants, female participants exhibit greater switching or exit behavior by 5 percentage points ($p = 0.01$) in response to information about advertising on misinformation websites (Table 3, Column 1).

Responses to our information treatments may also differ by respondents' political orientation. According to prior research, conservatives are especially likely to associate the mainstream media with the term "fake news". These perceptions are generally linked to lower trust in media, voting for Trump, and higher belief in conspiracy theories (Van der Linden et al., 2020). Moreover, conservatives are more likely to consume misinformation (Guess et al., 2019) and the supply of misinformation has been found to be higher on the ideological right than on the left (Garrett and Bond, 2021; Benkler et al., 2018).¹⁸ Consequently, we might expect stronger treatment effects for left-wing respondents. Our results show that respondents who voted for both candidates (Joseph Biden and Donald Trump) reduced demand for their top choice in response to our information treatments. However, consistent with our prediction, respondents who voted for President Biden in the 2020 US Presidential election are 3 percentage points more likely to exit ($p = 0.06$) and 5 percentage points more likely to voice concerns against company practices ($p = 0.04$) as shown in Columns 2 and 6 of Table 3, respectively.

Consumers who more frequently use a company's products or services could be presumed to be more loyal towards the company or derive greater utility from its use, which could limit changes in their behavior

¹⁸Our data also finds that the proportion of right-wing outlets is higher among misinformation outlets identified by third-party journalists relative to left-wing outlets. See Appendix Table A1.

Table 3: Heterogeneous Treatment Effects for Exit and Voice

	Switch in gift card from top choice company (“exit”)				Petition clicks on company petition (“voice”)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.07*** (0.01)	0.07*** (0.01)	0.12*** (0.01)	0.10*** (0.01)	0.02 (0.02)	0.00 (0.02)	0.03** (0.02)	0.03** (0.01)
Treatment × Female	0.05** (0.02)				0.00 (0.02)			
Treatment × Biden voter		0.03* (0.02)				0.05** (0.02)		
Treatment × Frequent user			-0.05*** (0.02)				-0.02 (0.02)	
Treatment × Consumes misinformation				-0.04* (0.02)				-0.03 (0.03)
Female	0.00 (0.01)	0.03*** (0.01)	0.03*** (0.01)	0.03*** (0.01)	0.00 (0.02)	0.01 (0.01)	0.01 (0.01)	0.00 (0.01)
Biden voter	0.01 (0.01)	-0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02* (0.01)	-0.01 (0.02)	0.02* (0.01)	0.02* (0.01)
Frequent user	-0.04*** (0.01)	-0.04*** (0.01)	-0.01 (0.01)	-0.04*** (0.01)	0.03** (0.01)	0.03** (0.01)	0.04** (0.02)	0.03** (0.01)
Consumes misinformation	0.03** (0.01)	0.03** (0.01)	0.03** (0.01)	0.05*** (0.02)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.02 (0.02)
Observations	4039	4039	4039	4039	4039	4039	4039	4039
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results where *Treatment* is a binary variable that takes a value of 1 if a respondent is randomized into any of the company-specific treatment groups (T1, T3 or T4). In columns 1 to 4, the dependent variable is switch in gift card choice from the respondent’s top choice company (i.e. “exit”). In columns 5 to 8, the dependent variable is clicking on a link to sign a petition that suggests that companies like the respondent’s top choice gift card company need to block their ads from appearing on misinformation websites. *Female* is a binary variable that takes a value of 1 if a respondent reports being female and zero otherwise. *Biden voter* is a binary variable that takes a value of 1 if a respondent reported voting for President Biden in the 2020 US Presidential election and zero otherwise. *Frequent user* is a binary variable that takes a value of 1 if a respondent reported using their top choice gift card at least once a month. *Consumes misinformation* is a binary variable that takes a value of 1 if a respondent reported using one or more misinformation news outlets (out of a list of 26 popular news outlets) in the past 12 months and zero otherwise. In all specifications above, we control for baseline demographic characteristics and behavioral characteristics as in Table 1. Robust standard errors in parentheses.

(Liaukonytė et al., 2022). Alternatively, more frequent consumers may be more strongly affected by our information treatments as they may perceive their usage as supporting such company practices to a greater extent than less frequent consumers. In our results, both frequent and infrequent users of a company’s products or services exit in response to our information treatments. Still, we observe a negative and statistically significant interaction term for frequent users, revealing that frequent users were about 5 percentage points less likely to exit ($p = 0.01$) as shown in Column 3 of Table 3.

Finally, we measure whether people’s responses differ by whether they consume misinformation themselves based on whether they reported using misinformation outlets in the initial question asking them to select which news outlets they used in the past 12 months. We find that both types of participants (those who report using the misinformation outlets we identify and those who do not) exit in response to our information treatments, but participants who consume misinformation are 4 percentage points less likely to exit, a decrease in demand that is significant at the 10% level ($p = 0.10$) as shown in Column 4 of Table 3. Respondents who do not consume misinformation also have a positive treatment effect of 3 percentage points higher than the control ($p = 0.01$) in terms of clicking on petitions suggesting that companies block their ads from appearing on misinformation websites. In contrast, respondents who do not report using misinformation outlets have a

positive but statistically insignificant increase in clicks on the same petition links.

Overall, we believe these heterogeneity results bolster the external validity of our experimental estimates. In particular, we highlight that product-specific factors such as frequency of use can play an important role in the decision to switch or not apart from ideological reasons such as political leaning.

4.2.4 Comparing Stated and Revealed Preferences

We find stark differences between consumers' stated preferences as measured by their degree of agreement with specific statements and revealed preferences as measured by their behaviors. While 11% of our participants exit, a much larger percentage (68%) agree that companies have an important role in reducing the spread of misinformation through their advertising practices (Figure A1-a). Similarly, while 23% of our participants sign petitions suggesting changes in company or platform practices, 76% agree that digital ad platforms should allow companies to avoid advertising on misinformation websites (Figure A1-b). Our stated preferences are comparable to recent industry reports, which found that nearly two-thirds of consumers state that they would stop using a brand if its ad appeared next to fake or offensive content (DoubleVerify, 2019), and 62-70% of consumers want companies to take a stand on social, cultural, environmental and political issues (SproutSocial, 2018). However, consistent with prior research documenting hypothetical bias in the measurement of stated preferences (Athey et al., 2017; List et al., 2001; Cummings et al., 1995), the contrast between our stated and revealed preference measures underscores the importance of eliciting revealed preferences.

4.2.5 Measuring the Experimenter Demand Effect

To minimize concerns about experimenter demand effects, we take several steps during our experimental design (Section 4.1), including using a neutral framing throughout our survey. We find that the vast majority of participants believe that the information provided in the survey was unbiased as shown in Figure A2.¹⁹ Participants' text responses also indicate that they believed their choices to be consequential (Tables A6 and A7). We now consider the extent to which experimenter demand effects may be relevant in driving the results.

To measure potential differences in the respondents' perceptions of the study, we examine their open-ended text responses about the purpose of the study using a Support Vector Machine classifier.²⁰ We predict treatment status using the classifier, keeping 75% of the sample for the training set and the remaining 25% as the test set. We find that the classifier predicts treatment status similar to chance for our main treatment groups relative to the control group, as shown in Table A8. These results, which are similar in magnitude to those of Bursztyn et al. (2021) and Song (2022), suggest that our treatments do not substantially affect participants' perceptions about the purpose of the study. Overall, this analysis gives us confidence that our main experimental findings are unlikely to be driven by experimenter demand effects.

¹⁹About 80% of survey participants chose "unbiased" when asked to rate the political bias of the survey information provided from a seven-point scale ranging from "very right-wing biased" to "very left-wing biased".

²⁰This classifier incorporates several features in text analysis, including word, character and sentence counts, sentiments, topics (using Gensim) and word embeddings.

5 Decision-maker study

Given that advertising on misinformation websites is both pervasive for companies across several industries and provokes consumer backlash in terms of exit and voice, what explains the prevalence of this phenomenon among companies that advertise online? To examine the beliefs and preferences of key decision-makers within companies relevant to advertising on misinformation websites, we survey executives and managers at companies. Throughout this study, we use the preferences of senior decision-makers (e.g., CEOs) as a proxy for company-level preferences since people in such roles shape the outcomes of their companies through their strategic decisions (Bertrand and Schoar, 2003; Porter, 1980; Drucker, 1967). This study followed the same IRB review, pre-registration and consent procedures as those used for our consumer study.

5.1 Research design

This study addresses two research questions. First, we aim to measure the existing beliefs and preferences decision-makers have about advertising on misinformation websites. This will help inform whether companies may be inadvertently or willingly sustaining online misinformation. Secondly, we ask: how do decision-makers update their beliefs and demand for a platform-based solution to avoid advertising on misinformation websites in response to information about the role of platforms in amplifying the financing of misinformation? This will suggest whether companies may be more interested in adopting ad platforms that reduce the financing of misinformation. To this end, we conduct an information provision experiment (Haaland et al., 2023). While past work has examined how firm behavior regarding market decisions changes in response to new information (Kim, 2021; Hanna et al., 2014; Bloom et al., 2013), it is unclear how information on the role of digital ad platforms in amplifying advertising on misinformation would affect decision-makers' non-market strategies (Lenox and Eesley, 2009; Eesley and Lenox, 2006).

5.1.1 Setting and sample size

We conduct an online survey experiment targeting key decision-makers such as managers and executives within who play a key role in strategic decision-making within their organizations. Our sample of respondents mainly comes from the executive education alumni of the Graduate School of Business at Stanford University with a smaller sample of executive alumni from Heinz College at Carnegie Mellon University.

Figure A18 shows the design of the survey study. We first elicit participants' current employment status. All those working in some capacity are allowed to continue the survey, whereas the rest of the participants are screened out. After asking for their main occupation, all participants in the experiment are provided with baseline information on misinformation and advertising similar to that provided in the consumer experiment.

5.1.2 Eliciting baseline beliefs

We measure participants' baseline beliefs. Specifically, participants are asked to estimate the number of companies among the most active 100 advertisers whose ads appeared on misinformation websites during the past three years (2019-2021). Additionally, we ask participants to report whether they think their company

or organization had its ads appear on misinformation websites in the past three years. Finally, we measure participants' beliefs about the role played by digital ad platforms in placing ads on misinformation websites. To do so, we first inform participants that during the past three years (2019-2021), out of every 100 companies that did not use digital ad platforms, eight companies appeared on misinformation websites on average. We then asked participants to provide their best estimate for the number of companies whose ads appeared on misinformation websites out of every 100 companies that did use digital ad platforms.

5.1.3 Measuring preferences

In addition to recording participants' stated preferences using self-reported survey measures, we measure participants' revealed preferences. To ensure incentive compatibility, participants are asked three questions in a randomized order: 1) Information demand about consumer responses, i.e. whether they would like to learn how consumers respond to companies whose ads appear on misinformation websites (based on our consumer survey experiment), 2) Ad check, i.e. whether they would like to know about their own company's ads appearing on misinformation websites in the recent past, and 3) Demand for a solution, i.e. whether they would like to sign up for a 15-minute information session on how companies can manage where their ads appear online. Participants are told they can receive information about consumer responses at the end of the study if they opt to receive it whereas the ad check and solution information are provided as a follow-up after the survey.²¹ Since all three types of information offered are novel and otherwise costly to obtain, we expect respondents' demand for such information to capture their revealed preferences.

5.1.4 Information intervention

Participants are then randomized into a treatment group, which receives information about the role of digital ad platforms in placing ads on misinformation websites, and a control group, which does not receive this information. Based on the dataset we assembled, participants are given factual information that companies that used digital ad platforms were about ten times more likely to appear on misinformation websites than companies that did not use such platforms in the recent past. This information is identical to the information provided to participants in the T2 (i.e. platform only) group in the consumer experiment.

5.1.5 Outcomes

Following the information intervention, we first measure participants' posterior beliefs about the role played by digital ad platforms in placing ads on misinformation websites. Participants are told about the average number of advertising companies whose ads appear per month on misinformation websites that are not monetized by digital ad platforms. They are then asked to estimate the average number of advertising companies whose ads appear monthly on misinformation websites that use digital ad platforms. This question measures whether participants believe that the use of digital ad platforms amplifies advertising on misinformation websites.

²¹Participants are required to provide their emails and company name for the ad check. To sign up for an information session from our industry partner on a potential solution to avoid advertising on misinformation websites, participants sign up on a separate form by providing their emails.

We record two behavioral outcomes. Our main outcome of interest is the respondents' demand for a platform-based solution to avoid advertising on misinformation websites. Participants can opt to learn more about two different types of information, i.e. 1) which platforms least frequently place companies' ads on misinformation websites and 2) which types of analytics technologies are used to improve ad performance, or opt not to receive any information (Figure A27). Since participants can only opt to receive one of the two types of information, this question is meant to capture the trade-off between respondents' concern for avoiding misinformation outlets and their desire to improve ad performance, respectively. Participants are told that they will be provided with the information they choose at the end of this study. Following the literature in measuring information acquisition (Capozza et al., 2021), we measure respondents' demand for solution information, which serves as a revealed-preference proxy for their interest in implementing a solution for their organization (Hjort et al., 2021).

Additionally, to measure whether the information treatment increases concern for financing misinformation in general, we record a second behavioral measure. Participants are told that the research team will donate \$100 to one of two organizations after randomly selecting one of the first hundred responses: 1) The Global Disinformation Index (GDI), and 2) DataKind, which helps mission-driven organizations increase their impact by unlocking their data science potential ethically and responsibly.

5.2 Results

We received 567 total complete responses, of which 90% are from currently employed respondents. To ensure data quality, we drop an additional 13% of responses which suggested that the participants did not read through and carefully answer the survey, resulting in a total sample of 442 responses.²² About 49% of the participants in our study were those currently serving in a top executive role (e.g., chief executives, general and operations managers of multiple departments or locations, etc.). Table A9 summarizes the descriptive characteristics, beliefs, and preferences of our study participants.

5.2.1 Baseline beliefs and characteristics

The vast majority of decision-makers in our sample believe it is important to control the spread of misinformation in society and that digital platforms should give companies a way to avoid advertising on misinformation websites (Table A9). There is a wide dispersion in decision-makers' beliefs about the role of companies and platforms in financing misinformation as shown in Figures A5 and A6. Decision-makers largely overestimate the overall proportion of companies advertising on misinformation websites and underestimate the role of digital ad platforms in placing companies' ads on misinformation websites.²³ Only 41% of decision-makers believe that consumers react against companies whose ads appear on misinformation websites. These results suggest that decision-makers believe that advertising on misinformation websites is likely commonplace but has little to do with the use of digital ad platforms and little consequences for the companies involved.

²²We dropped responses where participants provided an answer greater than 100 when asked to estimate a number out of 100.

²³On average, respondents estimated that about 64% of companies had their ads appear on misinformation websites (Table A9). However, our data shows that 55% of the 100 most active advertisers appeared on misinformation websites (see Section 3). Regarding the role of digital ad platforms, respondents estimated that about 44.5% of companies using digital ad platforms appear on misinformation

Table 4: Decision-makers' Beliefs and Characteristics about Advertising on Misinformation Websites

	All (1)	Executives (2)	Managers (3)	Other (4)
Belief about Advertising on Misinformation	0.20	0.18	0.24	0.17
Certainty of Belief about Advertising on Misinformation	0.79	0.81	0.75	0.83
Advertised on Misinformation*	0.81	0.72	0.86	1.00
% of Correct Beliefs about Own Company *	0.36	0.41	0.33	0.25
Observations	442	248	147	47

Notes: This table shows respondents' beliefs and characteristics about their own company advertising on misinformation outlets. Column 1 shows results for the full sample, Column 2 for the sub-sample of executives, Column 3 for the sub-sample of managers, and Column 4 for the remaining individuals. The proportions in rows marked with an asterisk (*) are calculated based on the subsample of participants who requested an ad check and whose companies appeared in our advertising data (N = 106).

However, as shown in Table 4, respondents substantially underestimate their own company's likelihood of appearing on misinformation websites with only 20% of respondents believing that their own company's ads appeared on misinformation websites in the three years prior to the study. Among the subsample of participants who requested an ad check (by providing their company name and contact details) and whose companies appeared in our advertising data, approximately 81% of companies appeared on misinformation websites. These figures illustrate that decision-makers are largely uninformed about the high likelihood of their company's ads appearing on misinformation websites. We further segment our results by type of role within the company. While our sub-samples are small, the results for beliefs are largely similar across the full sample, executives and marketers. Given that key decision-makers within companies ranging from marketers to executives are largely unaware of their companies' ads appearing on misinformation websites, our results suggest that companies may be financing misinformation inadvertently.

5.2.2 Preferences

The vast majority of participants requested an ad check by providing their company name and email address (74%). The demand for an ad check was high regardless of respondents' beliefs, suggesting a substantial interest in learning about whether their company's ads appeared on misinformation websites. Despite only 41% of respondents agreeing that consumers react against companies whose ads appear on misinformation websites, most participants opted to receive information on how consumers respond to companies whose ads appear on misinformation websites (73%). This suggests that while decision-makers may be unaware of how advertising on misinformation websites can provoke consumer backlash, most decision-makers are interested in learning about the degree of potential backlash.²⁴

Finally, for our most costly revealed preference measure, i.e. signing up to attend a 15-minute expert-led information session on how companies can avoid advertising on misinformation websites, 18% of decision-makers clicked to sign-up, an arguably high rate given the value of decision-makers' time and the opportunity cost of attending the session.²⁵ The difference in demand for our lower-cost information (73-74%) and higher-

websites (Table A9) as opposed to the 79.8% of companies among the 100 most active advertisers that do so

²⁴Most participants inquired about "exit" (58%), with only 15% inquiring about "voice".

²⁵For receiving an ad check, i.e., finding out whether their own company's ads appeared on misinformation websites in the recent past, we observed a close match between respondents' stated and revealed preferences with 76% of respondents stating that they would like to find out whether their company's ads appeared on misinformation websites and 74% of respondents providing their company's name and contact details to request an ad check. For solution information demand, however, there was a substantial gap between decision-makers'

cost information (18%) suggests that providing lower-cost interventions such as allowing advertisers to easily steer their ads across different types of news outlets could be more fruitful in aligning advertiser preferences with their algorithmically-driven ad placements.

5.2.3 Information intervention results

We report the results of our information treatment in Table 5. For the full sample of participants, we estimate positive and statistically significant effects on participants’ posterior beliefs about the role of ad platforms in placing ads on misinformation websites (Column 1), which is mainly driven by respondents who believe their company’s ads did not appear on misinformation websites in the recent past (Column 3).

Table 5: Average Treatment Effects of Information Intervention

	Posterior platform belief			Platform solution demand		
	All (1)	Yes (2)	No (3)	All (4)	Yes (5)	No (6)
Treatment	48.99*** (15.79)	6.11 (43.43)	54.15*** (17.77)	-0.03 (0.05)	-0.10 (0.13)	-0.03 (0.05)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	96.10	124.47	88.65	0.37	0.34	0.38
Observations	442	88	354	442	88	354

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results where the dependent variables are posterior beliefs (columns 1 to 3), demand for platform solution (columns 4 to 6). We winsorize the posterior beliefs to remove outliers. Our platform solution demand outcome variable is a binary variable that takes a value of one when participants choose to receive information on which platforms least frequently place companies’ ads on misinformation websites and zero otherwise. Columns 1 and 4 show results for the full sample of participants. Columns 2 and 5 show results for the sub-sample of participants who reported “yes” to the question “Do you think your company or organization had its ads appear on misinformation websites during the past three years (2019-2021)?”. Columns 3 and 6 show results for the sub-sample who reported “No” in response to the same question. We control for decision-makers’ characteristics and prior beliefs. Robust standard errors in parentheses.

Table 6: Treatments Effects Based On Prior Beliefs

	Posterior platform belief		Platform solution demand	
	Certain (1)	Uncertain (2)	Certain (3)	Uncertain (4)
Treatment	39.98** (20.23)	144.25** (60.23)	-0.07 (0.06)	0.36*** (0.13)
Controls	Yes	Yes	Yes	Yes
Control group mean	90.23	80.80	0.37	0.43
Observations	286	68	286	68

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results for the sub-sample of participants who reported “No” to the question “Do you think your company or organization had its ads appear on misinformation websites during the past three years (2019-2021)?”. The dependent variables are posterior beliefs (Columns 1-2) and demand for platform solution (Columns 3-4) from Table 5. Columns 1 and 3 show results for participants who report being certain about their response to the aforementioned question (choosing “Somewhat sure”, “Sure” or “Very sure”). Column 4 shows results for participants who report being uncertain about their response to the aforementioned question (choosing “Unsure” or “Very unsure”). Robust standard errors in parentheses.

We find an overall null effect of our information treatment on participants’ demand for a platform-based solution, as measured by their information demand for which platforms least frequently place companies’ ads

stated and revealed preferences with 71% of respondents stating that they would recommend that their company adopt a product to avoid advertising on misinformation websites, but only 18% of respondents clicking on the form to sign up for a 15-minute information session to learn how to do so.

on misinformation websites (Columns 4-6 in Table 5). However, this result masks substantial heterogeneity based on participants' prior beliefs. Since our information treatment changes beliefs for the subset of participants who believe their company's ads did not recently appear on misinformation websites (Column 3, Table 5), we further investigate and report results based on participant's prior beliefs for this sub-sample in Table 6. We find that only participants who were uncertain about their own company's ads appearing on misinformation websites responded positively and significantly to our information treatment by increasing their demand for a platform-based solution by 40 percentage points ($p = 0.003$) as shown in Table 6 (Column 4).

Our results imply that the way participants respond to information about the role played by digital ad platforms in financing misinformation is highly dependent on their prior beliefs about their own company. For those uncertain about their beliefs, providing such information can increase their demand for a platform-based solution to reduce advertising on misinformation outlets. Such information could make companies switch ad platforms or pressure the platforms they currently use to allow them to easily steer their ads away from misinformation outlets.

We did not find meaningful treatment effects for our donation preference outcome for the full sample or any subsamples based on participants' self-reported beliefs (Table A10). As previously mentioned, this outcome measures the proportion of respondents who prefer that we donate to the Global Disinformation Index (GDI) instead of DataKind. Since both organizations have similar goals of advancing technology's ethical and responsible use, respondents may have considered their missions interchangeable. Moreover, unlike our first behavioral outcome, respondents could have considered donating to the GDI less relevant to their own organizations' needs and more a matter of personal preference.

6 Discussion

In this paper, we show how advertising financially sustains online misinformation. We extend prior work on the infrastructure sustaining misinformation websites by combining data on misinformation websites identified from three different sources with novel digital advertising data over a period of three years. We find that a substantial proportion of companies across several industries advertise on misinformation websites and that digital ad platforms amplify advertising on misinformation websites. Our experimental results suggest that both advertising companies and digital ad platforms can face consumer backlash for monetizing misinformation outlets. Previous work has examined the conditions under which people react against companies for failing to operate up to their expectations (Hirschman, 1970; Broccardo et al., 2022). Our research design contributes to this literature in two key ways by (1) measuring both types of potential consumer responses, i.e., "exit" and "voice", theorized in the literature, and (2) doing so using incentive-compatible behavioral outcomes at the individual level. Finally, we show that most key decision-makers within companies are unaware of their company's ads appearing on misinformation websites, but have a high demand to find out about such information. Our findings complement prior work that shows wide dispersion in decision-makers' beliefs about key economic conditions (Coibion et al., 2018; Link et al., 2023). Some decision-makers also increased their demand for a platform-based solution to reduce advertising on misinformation websites once informed

about the role of platforms in financing misinformation, which is in line with a lack of attention describing decision-makers' behaviors across various settings (Kim, 2021; Hanna et al., 2014; Ocasio, 1997).

Given the potential for a substantial decline in demand advertising companies may wish to account for consumer preferences in placing their ads across various online outlets and exercise caution while incorporating automation in their business processes via digital ad platforms. For instance, given that consumers switched to other products upon learning about a company's ads appearing on misinformation websites, companies could use lists of misinformation outlets provided by independent third-party organizations such as NewsGuard and the Global Disinformation Index to limit ad dollars going to misinformation outlets through digital ad platforms. Moreover, since consumer backlash was particularly strong for women and politically left-leaning consumers, companies targeting such audiences may wish to exercise greater caution.

Based on our results, we identify two interventions that could reduce the financing of online misinformation. First, digital ad platforms that run automated auctions could enable advertisers to more easily access data on whether their ads appear on misinformation outlets. This would enable advertisers to make ad placement decisions consistent with their preferences rather than inadvertently financing misinformation (Ada et al., 2021). Second, while it is currently possible for consumers to find out about companies financing misinformation through media reports, digital platforms could improve transparency for consumers about which companies advertise on misinformation outlets. Platforms could provide such information to consumers when they are viewing with an ad using simple information labels (as in our "Company only" information treatment) similar to the "Sponsored by" and "Paid for by X" labels that are presently common on various digital media platforms. Similarly, rank-based information provided in our "Company ranking" information treatment (T4) could be provided as a ranking of companies in order of intensity of appearing on misinformation websites where customers are selecting products from a menu of choices while shopping. Platforms have provided similar contextual information about companies in other settings, e.g., Google Flights displays carbon emissions data alongside flight prices when people select a flight to purchase among several options (Holden, 2021). Enabling consumers to view such information at the point of purchase could provide a stronger incentive for companies to steer their ads away from such outlets, especially since the impact of negative information can persist for several months (Babar et al., 2023; Spampatti et al., 2022). Overall, these interventions could decrease the inadvertent ad revenue going towards misinformation outlets, which could eventually lead to such sites ceasing to operate, as observed anecdotally in prior work (Han et al., 2022).

These interventions can serve to ensure that both consumers and advertisers are provided with information about the consequences of their respective purchasing and ad placement decisions so that they can make choices consistent with their preferences. Having access to such information is necessary for an efficiently functioning economic system per the first fundamental theorem of welfare economics. However, while digital platforms are uniquely well-positioned in the ecosystem of consumers, advertisers, and publishers to implement information interventions in the form of disclosures and rankings (Boudreau and Hagiu, 2009; Rietveld et al., 2021), they may not have incentives to implement such interventions. In the backdrop of mounting pressure from advertisers (Hsu, 2022; Hsu and Lutz, 2020; D'Onfro, 2019) and calls for transparency in the programmatic ad business (DOJ, 2023; Horwitz and Hagey, 2021), such interventions could be incorporated

into existing legislation to improve transparency. Importantly, in recent years, policy proposals aiming to reduce the prevalence of misinformation such as Australia’s Combating Misinformation and Disinformation bill and Germany’s bill against fake news have faced backlash over posing risks to free speech.(Remeikis, 2023; Faiola and Kirchner, 2017) While such policy proposals all face the challenge of striking the right balance between combating misinformation and protecting freedom of expression, the information interventions we identify could help counter the financial incentive to produce misinformation in the first place by reducing the unintended ad revenue going towards misinformation outlets. There are many parallels for regulation by information provision to address externalities in other industries, including chemicals (toxic release inventory reporting requirements), automobiles (miles per gallon information), food (nutrition and content labels), airlines (greenhouse gas emissions), of which several have been demonstrated to be effective in prior work (Doshi et al., 2013; Johnson, 2020; Jin and Leslie, 2009).

“Demand-side” interventions to counter online misinformation studied in prior work have focused on reducing the consumption and spread of misinformation among news consumers on online platforms. While interventions such as accuracy prompts and digital literacy tips can increase the quality of news that people share (Arechar et al., 2023), this line of work has found limited support for news credibility signals in increasing the demand for credible news (Chopra et al., 2022) or in reducing misperceptions among users (Aslett et al., 2022). Such constraints in changing user behavior may also apply to credibility signals like watermarks for detecting AI misinformation. Moreover, while such interventions are only effective for the small subset of users who are exposed to misinformation (Allen et al., 2020), the complementary ‘supply-side’ approach identified in this paper targets entities and individuals who might not necessarily consume or spread misinformation themselves.

Relative to existing proposals of supply-side interventions to curb the production of misinformation, which involve social media platforms banning the advertising of false news (Chiou and Tucker, 2018) or changing their ad-driven business model altogether (Romer, 2019), the information interventions we identify take a middle path by accounting for advertisers’ and consumers’ preferences to help counter the financing of online misinformation. Moreover, the information interventions we identify are also an improvement on the status quo, whereby advertisers and consumers can only implement their preferences by participating in boycotts of digital platforms over their inability to contain misinformation. Allowing advertisers to more easily observe and control whether their ads appear on misinformation websites could also limit backlash by enabling advertisers to better implement their preferences rather than participating in one-off short-term ad boycotts (Hsu, 2022; Hsu and Lutz, 2020; D’Onfro, 2019). Additionally, since consistently providing negative information can create lasting associations for consumers (Lougee and Wallace, 2008), providing information disclosures on every ad for whether the advertiser places ads on misinformation websites could have a substantial effect on consumer demand over time, further incentivizing advertisers to remove their ads from misinformation websites.

Given our findings, we suggest three promising avenues for future research. Future work could evaluate the effectiveness of our information interventions in the field over a longer time period to quantify the decline in revenue generated by misinformation outlets resulting from increasing transparency for consumers or advertisers. Relatedly, future work could also target a wider set of advertisers to validate the robustness of our

interventions which would allow for broader generalizability. Secondly, our results on whether companies are willing to adopt solutions to avoid monetizing misinformation are based on their existing (incorrect) beliefs about the prevalence of advertising on misinformation websites in general and for their own company. More research is needed to understand how advertising companies would respond in the context of correct beliefs. Finally, while our research identifies potential interventions that digital platforms can adopt to curb the monetization of online misinformation, it is unclear whether it is in the interest of digital ad platforms to do so. Moreover, whether the potential monetary and societal benefits of the information interventions we identify outweigh the revenue platforms generate by serving ads on misinformation websites remains to be studied. Overall, the effectiveness of platforms in mitigating misinformation will depend on a multi-pronged approach. Given that misinformation is largely financially motivated and that financially sustaining online misinformation can be substantially harmful for the advertising companies involved, simple low-cost informational interventions such as the ones we identify could go a long way in curbing the supply of online misinformation.

References

- (2023). Justice Department Sues Google for Monopolizing Digital Advertising Technologies. *Office of Public Affairs*. Retrieved from: <https://www.justice.gov/opa/pr/justice-department-sues-google-monopolizing-digital-advertising-technologies>.
- Ada, S., Nabout, N. A., and Feit, E. M. (2021). Context Information Can Increase Revenue in Online Display Advertising Auctions: Evidence from a Policy Change. *Journal of Marketing Research*, 59(5):1040–1058.
- Agarwal, A. and Mukhopadhyay, T. (2016). The Impact of Competing Ads on Click Performance in Sponsored Search. *Information Systems Research*, 27(3):538–557.
- Allcott, H., Gentzkow, M., and Yu, C. (2019). Trends in the diffusion of misinformation on social media. *Research and Politics*, 6(2).
- Allen, J., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14).
- Angelov, D. (2020). Top2Vec: Distributed Representations of Topics. Retrieved from: <https://arxiv.org/abs/2008.09470v1>.
- Arechar, A. A., Allen, J., Berinsky, A. J., Cole, R., Epstein, Z., Garimella, K., Gully, A., Lu, J. G., Ross, R. M., Stagnaro, M. N., Zhang, Y., Pennycook, G., and Rand, D. G. (2023). Understanding and combatting misinformation across 16 countries on six continents. *Nature Human Behaviour*.
- Aslett, K., Guess, A. M., Bonneau, R., Nagler, J., and Tucker, J. A. (2022). News credibility labels have limited average effects on news diet quality and fail to reduce misperceptions. *Science Advances*, 8(18):3844.
- Athey, S., Catalini, C., and Tucker, C. E. (2017). The Digital Privacy Paradox: Small Money, Small Costs, Small Talk. *National Bureau of Economic Research Working Paper 23488*. Retrieved from: <http://www.nber.org/papers/w23488.pdf>.
- Austin, A., Barnard, J., and Hutcheon, N. (2019). Programmatic Marketing Forecasts. Technical report, Zenith. Retrieved from: <https://s3.amazonaws.com/media.mediapost.com/uploads/ProgrammaticMarketingForecasts2019.pdf>.
- Babar, Y., Mahdavi Adeli, A., and Burtch, G. (2023). The effects of online social identity signals on retailer demand. *Management Science*.
- Bellman, S., Abdelmoety, Z. H., Murphy, J., Arismendez, S., and Varan, D. (2018). Brand safety: the effects of controversial video content on pre-roll advertising. *Heliyon*, 4(12):e01041.
- Benkler, Y., Faris, R., and Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford University Press.
- Bertrand, M. and Schoar, A. (2003). Managing with Style: The Effect of Managers on Firm Policies. *The Quarterly Journal of Economics*, 118(4):1169–1208.
- Bhadani, S., Yamaya, S., Flammini, A., Menczer, F., Ciampaglia, G. L., and Nyhan, B. (2022). Political audience diversity and news reliability in algorithmic ranking. *Nature Human Behaviour* 2022 6:4, 6(4):495–505.
- Blei, D. M., Ng, A. Y., and Edu, J. B. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3:993–1022.
- Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? Evidence from india. *Quarterly Journal of Economics*, 128(1):1–51.
- Blumberg, D. L. (2023). 3 ways the ‘splinternet’ is damaging society. *Science Advances*.
- Boudreau, K. and Hagiu, A. (2009). *Platforms, Markets, and Innovation*. Edward Elgar Publishing, Inc., Northampton, MA.
- Braun, J. A. and Eklund, J. L. (2019). Fake News, Real Money: Ad Tech Platforms, Profit-Driven Hoaxes, and the Business of Journalism. *Digital Journalism*, 7(1):1–21.
- Broccardo, E., Hart, O., and Zingales, L. (2022). Exit versus Voice. *Journal of Political Economy*, 130(12).
- Bursztyn, L., Haaland, I. K., Rao, A., and Roth, C. P. (2021). Disguising Prejudice: Popular Rationales as Excuses for Intolerant Expression. Retrieved from: <https://ideas.repec.org/p/wrk/warwec/1340.html>.
- Bursztyn, L., Rao, A., Roth, C. P., and Yanagizawa Drott, D. H. (2020). Misinformation During a Pandemic. *National Bureau of Economic Research Working Paper 27417*. Retrieved from: <https://www.nber.org/papers/w27417>.
- Cao, X. and Zhang, J. (2021). Preference learning and demand forecast. *Marketing Science*, 40(1):62–79.

-
- Capozza, F., Haaland, I., Roth, C., and Wohlfart, J. (2021). *Studying information acquisition in the field : a practical guide and review*. CEBI working paper series : working paper. Copenhagen : CEBI, Department of Economics, University of Copenhagen.
- Casadesus-Masanell, R. and Zhu, F. (2010). Strategies to fight ad-sponsored rivals. *Management Science*, 56(9):1484–1499.
- Casadesus-Masanell, R. and Zhu, F. (2013). Business model innovation and competitive imitation: The case of sponsor-based business models. *Strategic Management Journal*, 34(4):464–482.
- Chandler, J., Rosenzweig, C., Moss, A. J., Robinson, J., and Litman, L. (2019). Online panels in social science research: Expanding sampling methods beyond Mechanical Turk. *Behavior research methods*, 51(5):2022–2038.
- Chatterji, A. K. and Toffel, M. W. (2019). Assessing the Impact of CEO Activism:. *Organization & Environment*, 32(2):159–185.
- Chiou, L. and Tucker, C. (2018). Fake News and Advertising on Social Media: A Study of the Anti-Vaccination Movement. *National Bureau of Economic Research Working Paper 25223*. Retrieved from: <http://www.nber.org/papers/w25223.pdf>.
- Chopra, F., Haaland, I., and Roth, C. (2022). Do people demand fact-checked news? Evidence from U.S. Democrats. *Journal of Public Economics*, 205:104549.
- Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). How Do Firms Form Their Expectations? New Survey Evidence. *American Economic Review*, 108(9):2671–2713.
- Cowgill, B. and Dorobantu, C. (2018). Competition and Specificity in Market Design: Evidence from Geotargeted Advertising. *NET Institute Working Paper No. 18-09*. Retrieved from: <https://ideas.repec.org/p/net/wpaper/1809.html>.
- Crosan, R. and Gneezy, U. (2009). Gender Differences in Preferences. *Journal of Economic Literature*, 47(2):448–74.
- Crovitz, G. (2020). How Amazon, Geico and Walmart Fund Propaganda. *The New York Times*. Retrieved from: <https://www.nytimes.com/2020/01/21/opinion/fake-news-russia-ads.html>.
- Cummings, R. G., Harrison, G., Rutstrom, E., Cummings, R. G., Harrison, G., and Rutstrom, E. (1995). Homegrown Values and Hypothetical Surveys: Is the Dichotomous Choice Approach Incentive-Compatible? *American Economic Review*, 85(1):260–66.
- De Quidt, J., Haushofer, J., and Roth, C. (2018). Measuring and Bounding Experiment Demand. *American Economic Review*, 108(11):3266–3302.
- De Reuver, M., Sørensen, C., and Basole, R. C. (2018). The Digital Platform: A Research Agenda. *Journal of Information Technology*, 33(2):124–135.
- Decker, B. (2019). Adversarial narratives: A new model for disinformation, global disinformation index. *Global Disinformation Index*. Retrieved from: <https://www.disinformationindex.org/research/2019-4-1-adversarial-narratives-a-new-model-for-disinformation/>.
- DeGeurin, M. (2023). Twitter runs ads for disney, microsoft, and the nba next to neo-nazi propaganda. *Gizmodo*. Retrieved from: <https://gizmodo.com/twitter-ads-disney-microsoft-nba-neo-nazi-videos-elon-1850549017>.
- D’Onfro, J. (2019). Advertisers Boycott YouTube — And If History Is A Guide, They’ll Be Back Soon. *Forbes*. Retrieved from: <https://www.forbes.com/sites/jilliandonfro/2019/02/21/advertisers-boycott-youtube--and-if-history-is-a-guide-theyll-be-back-soon/?sh=5588eb367a96>.
- Doshi, A. R., Dowell, G. W., and Toffel, M. W. (2013). How firms respond to mandatory information disclosure. *Strategic Management Journal*, 34(10):1209–1231.
- DoubleVerify (2019). Study: Consumers Reject Brands That Advertise on ‘Fake News’ and Objectionable Content Online. Retrieved from: <https://doubleverify.com/newsroom/study-consumers-reject-brands-that-advertise-on-fake-news-and-objectionable-content-online/>.
- Drucker, P. (1967). *The effective executive*. Harper & Row, New York.
- Du, S., Bhattacharya, C. B., and Sen, S. (2011). Corporate Social Responsibility and Competitive Advantage: Overcoming the Trust Barrier. *Management Science*, 57(9):1528–1545.
- Dua, T. (2021). Pfizer, Walmart, CDC Ran Ads on Websites Peddling Vaccine Misinformation. Retrieved from: <https://www.businessinsider.com/pfizer-walmart-cdc-ran-ads-on-websites-peddling-vaccine-misinformation-2021-2>.
- Edelson, L., Nguyen, M. K., Goldstein, I., Goga, O., McCoy, D., and Lauinger, T. (2021). Understanding engagement with U.S. (mis)information news sources on Facebook. *Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC*, pages 444–463.

-
- Easley, C. and Lenox, M. J. (2006). Firm responses to secondary stakeholder action. *Strategic Management Journal*, 27(8):765–781.
- Eyal, P., David, R., Andrew, G., Zak, E., and Ekaterina, D. (2021). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, pages 1–20.
- Faiola, A. and Kirchner, S. (2017). How do you stop fake news? in germany, with a law. Retrieved from: https://www.washingtonpost.com/world/europe/how-do-you-stop-fake-news-in-germany-with-a-law/2017/04/05/e6834ad6-1a08-11e7-bcc2-7d1a0973e7b2_story.html.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics*, 133(4):1645–1692.
- Fong, J., Guo, T., and Rao, A. (2022). Debunking Misinformation about Consumer Products: Effects on Beliefs and Purchase Behavior. *Journal of Marketing Research*.
- Frick, T. W., Belo, R., and Telang, R. (2022). Incentive Misalignments in Programmatic Advertising: Evidence from a Randomized Field Experiment. *Management Science*, 69(3).
- Gans, J. S., Goldfarb, A., and Lederman, M. (2021). Exit, Tweets, and Loyalty. *American Economic Journal: Microeconomics*, 13(2):68–112.
- Garrett, R. K. and Bond, R. M. (2021). Conservatives’ susceptibility to political misperceptions. *Science Advances*, 23(7).
- Giansiracusa, N. (2021). Google needs to defund misinformation. *Slate*. Retrieved from: <https://slate.com/technology/2021/11/google-ads-misinformation-defunding-artificial-intelligence.html>.
- Gomes Ribeiro, B., Horta Ribeiro, M., Almeida, V., and Meira, W. (2022). Analyzing the "Sleeping Giants" Activism Model in Brazil. *ACM International Conference Proceeding Series*, pages 87–97.
- Grant, N. and Myers, S. L. (2023). Google Promised to Defund Climate Lies, but the Ads Keep Coming. *The New York Times*. Retrieved from: <https://www.nytimes.com/2023/05/02/technology/google-youtube-disinformation-climate-change.html>.
- Grewal, L., Stephens, A., and Vana, P. (2022). Brands In Unsafe Places: Effects of Brand Safety Incidents on Consumers’ Brand Attitudes. *Working Paper*.
- Grigorieff, A., Roth, C., and Ubfal, D. (2020). Does Information Change Attitudes Toward Immigrants? *Demography*, 57(3):1117–1143.
- Guess, A., Nagler, J., and Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, 5(1):1494–1504.
- Guess, A. M., Nyhan, B., and Reifler, J. (2020). Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour*, 4(5):472–480.
- Guzman, J., Oh, J. J., and Sen, A. (2020). What motivates innovative entrepreneurs? evidence from a global field experiment. *Management Science*, 66(10):4808–4819.
- Haaland, I. and Roth, C. (2020). Labor market concerns and support for immigration. *Journal of Public Economics*, 191:104256.
- Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing Information Provision Experiments. *Journal of Economic Literature*, 61(1):3–40.
- Han, C., Kumar, D., and Durumeric, Z. (2022). On the Infrastructure Providers That Support Misinformation Websites. In *Proceedings of the Sixteenth International AAAI Conference on Web and Social Media (ICWSM2022)*.
- Hanna, R., Mullainathan, S., and Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *Quarterly Journal of Economics*, 129(3):1311–1353.
- Hao, K. (2021). How Facebook and Google fund global misinformation. *MIT Technology Review*. Retrieved from: <https://www.technologyreview.com/2021/11/20/1039076/facebook-google-disinformation-clickbait/>.
- Hiar, C. (2021). Google Bans Ads That Spread Climate Misinformation. *E & E News*. Retrieved from: <https://www.scientificamerican.com/article/google-bans-ads-that-spread-climate-misinformation/>.
- Higgins, A., McIntire, M., and Dance, J. G. (2016). Inside a Fake News Sausage Factory: ‘This Is All About Income’. *The New York Times*. Retrieved from: <https://www.nytimes.com/2016/11/25/world/europe/fake-news-donald-trump-hillary-clinton-georgia.html>.
- Hirschman, A. O. (1970). *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states*. Harvard University Press.

-
- Hjort, J., Moreira, D., Rao, G., and Santini, J. F. (2021). How research affects policy: Experimental evidence from 2,150 Brazilian municipalities. *American Economic Review*, 111(5):1442–1480.
- Holden, R. (2021). Google Travel: Find flights with lower carbon emissions. *Google Blog*. Retrieved from: <https://blog.google/products/travel/find-flights-with-lower-carbon-emissions/>.
- Horwitz, J. and Hagey, K. (2021). Google’s Secret ‘Project Bernanke’ Revealed in Texas Antitrust Case. *Wall Street Journal*. Retrieved from: <https://www.wsj.com/articles/googles-secret-project-bernanke-revealed-in-texas-antitrust-case-11618097760>.
- Hsu, T. (2022). Twitter’s Advertisers Pull Back as Layoffs Sweep Through Company. *The New York Times*. Retrieved from: <https://www.nytimes.com/2022/11/04/technology/twitter-advertisers.html>.
- Hsu, T. and Lutz, E. (2020). More Than 1,000 Companies Boycotted Facebook. Did It Work? - The New York Times. *The New York Times*. Retrieved from: <https://www.nytimes.com/2020/08/01/business/media/facebook-boycott.html>.
- Hsu, T. and Tracy, M. (2021). Investors Push Home Depot and Omnicom to Steer Ads From Misinformation. *The New York Times*. Retrieved from: <https://www.nytimes.com/2021/01/18/business/media/investors-push-home-depot-and-omnicom-to-steer-ads-from-misinformation.html>.
- Jin, G. Z. and Leslie, P. (2009). Reputational Incentives for Restaurant Hygiene. *American Economic Journal: Microeconomics*, 1(1):237–67.
- John, L. K., Blunden, H., Milkman, K. L., Foschini, L., and Tuckfield, B. (2022). The limits of inconspicuous incentives. *Organizational Behavior and Human Decision Processes*, 172.
- Johnson, G. A. and Lewis, R. A. (2015). Cost Per Incremental Action: Efficient Pricing of Advertising.
- Johnson, M. (2020). Regulation by Shaming: Deterrence Effects of Publicizing Violations of Workplace Safety and Health Laws. *American Economic Review*, 110(6):1866–1904.
- Kim, H. (2021). The Value of Competitor Information: Evidence from a Field Experiment. In *Academy of Management Proceedings*, volume 2021, page 12714. Academy of Management Briarcliff Manor, NY 10510.
- Kohno, T., Zeng, E., and Roesner, F. (2020). Bad News, Clickbait and Deceptive Ads on News and Misinformation Websites. In *Workshop on Technology and Consumer Protection*.
- Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., and Zittrain, J. L. (2018). The science of fake news: Addressing fake news requires a multidisciplinary effort. *Science*, 359(6380):1094–1096.
- Lenox, M. J. and Eesley, C. E. (2009). Private environmental activism and the selection and response of firm targets. *Journal of Economics and Management Strategy*, 18(1):45–73.
- Liaukonytė, J., Tuchman, A., and Zhu, X. (2022). Frontiers: Spilling the Beans on Political Consumerism: Do Social Media Boycotts and Buycotts Translate to Real Sales Impact? *Marketing Science*.
- Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information Frictions among Firms and Households. *Journal of Monetary Economics*, 135:99–115.
- List, J., Gallet, C., List, J., and Gallet, C. (2001). What Experimental Protocol Influence Disparities Between Actual and Hypothetical Stated Values? *Environmental & Resource Economics*, 20(3):241–254.
- Lougee, B. and Wallace, J. (2008). The corporate social responsibility (csr) trend. *Journal of Applied Corporate Finance*, 20(1):96–108.
- Love, J. and Cooke, K. (2016). Google, Facebook move to restrict ads on fake news sites. *Reuters*. Retrieved from: <https://www.reuters.com/article/us-alphabet-advertising/google-facebook-move-to-restrict-ads-on-fake-news-sites-idUSKBN1392MM>.
- Lull, R. B. and Bushman, B. J. (2015). Do sex and violence sell? A meta-analytic review of the effects of sexual and violent media and ad content on memory, attitudes, and buying intentions. *Psychological Bulletin*, 141(5):1022–1048.
- McCarthy, B. (2021). Misinformation and the Jan. 6 insurrection: When ‘patriot warriors’ were fed lies. *Politifact*. <https://www.politifact.com/article/2021/jun/30/misinformation-and-jan-6-insurrection-when-patriot/>.
- Milmo, D. and Hern, A. (2023). Elections in UK and US at risk from AI-driven disinformation, say experts. *The Guardian*. Retrieved from: <https://www.theguardian.com/technology/2023/may/20/elections-in-uk-and-us-at-risk-from-ai-driven-disinformation-say-experts>.
- Moore, R. C., Dahlke, R., and Hancock, J. T. (2023). Exposure to untrustworthy websites in the 2020 us election. *Nature Human Behaviour*, pages 1–10.

-
- Mosseri, A. (2017). Working to Stop Misinformation and False News. *Meta Newsroom*. Retrieved from: <https://about.fb.com/news/2017/04/working-to-stop-misinformation-and-false-news/>.
- NewsGuard (2021). Special Report: Top brands are sending \$2.6 billion to misinformation websites each year. Technical report, NewsGuard. Retrieved from: <https://www.newsguardtech.com/special-reports/brands-send-billions-to-misinformation-websites-newsguard-comscore-report/>.
- Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(SPEC. ISS.):187–206.
- Ong, A. D. and Weiss, D. J. (2000). The Impact of Anonymity on Responses to Sensitive Questions. *Journal of Applied Social Psychology*, 30(8):1691–1708.
- Papadogiannakis, E., Papadopoulou, P., Markatos, E. P., and Kourtellis, N. (2023). Who Funds Misinformation? A Systematic Analysis of the Ad-related Profit Routines of Fake News sites. *The Web Conference 2023*. Retrieved from: <https://arxiv.org/abs/2202.05079>.
- Pennycook, G. and Rand, D. G. (2019). Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences of the United States of America*, 116(7):2521–2526.
- Pennycook, G. and Rand, D. G. (2022). Accuracy prompts are a replicable and generalizable approach for reducing the spread of misinformation. *Nature Communications 2022 13:1*, 13(1):1–12.
- Porter, M. E. (1980). Competitive Strategy Techniques for Analyzing Industries and Competitors. *The Free Press*, pages 1–422.
- Remeikis, A. (2023). Why is labor’s bill on combatting disinformation so controversial? Retrieved from: <https://www.theguardian.com/australia-news/2023/oct/01/why-is-labors-bill-on-combatting-disinformation-so-controversial>.
- Rietveld, J., Seamans, R., and Meggiorin, K. (2021). Market orchestrators: The effects of certification on platforms and their complementors. *Strategy Science*, 6(3):244–264.
- Romer, P. (2019). A Tax That Could Fix Big Tech. *The New York Times*. Retrieved from: <https://www.nytimes.com/2019/05/06/opinion/tax-facebook-google.html>.
- Ryan-Mosley, T. (2023). Junk websites filled with AI-generated text are pulling in money from programmatic ads. *MIT Technology Review*. Retrieved from: <https://www.technologyreview.com/2023/06/26/1075504/junk-websites-filled-with-ai-generated-text-are-pulling-in-money-from-programmatic-ads/>.
- Song, L. (2022). The Heterogeneous Effects of Social Media Content on Racial Attitudes. Retrieved from: <https://www.dropbox.com/s/f48vgfadd23226r/TwitterDiversity.pdf?dl=0>.
- Spampatti, T., Hahnel, U. J., Trutnevyte, E., and Brosch, T. (2022). Short and long-term dominance of negative information in shaping public energy perceptions: The case of shallow geothermal systems. *Energy Policy*, 167:113070.
- SproutSocial (2018). Creating Connections: What Consumers Want From Brands in an Increasingly Divided Society. Retrieved from: <https://sproutsocial.com/insights/data/social-media-connection/>.
- Van der Linden, S., Leiserowitz, A., Rosenthal, S., and Maibach, E. (2017). Inoculating the Public against Misinformation about Climate Change. *Global Challenges*, 1(2):756–784.
- Van der Linden, S., Panagopoulos, C., and Roozenbeek, J. (2020). You are fake news: political bias in perceptions of fake news. *Media, Culture & Society*, 42(3):460–470.
- Xiao, X., Su, Y., and Lee, D. K. L. (2021). Who Consumes New Media Content More Wisely? Examining Personality Factors, SNS Use, and New Media Literacy in the Era of Misinformation. *Social Media + Society*, 7(1).

Appendix

Appendix A: Additional results and descriptives

Descriptive statistics

Table A1: Summary statistics for NewsGuard data

	US		Global	
	Misinformation	Non-misinformation	Misinformation	Non-misinformation
Average score	17.3	73.7	17.9	76.6
% of trustworthy websites	6.1	70.1	6.0	76.4
% of websites with paywall	2.7	25.0	3.2	24.0
% of owned by individuals	25.3	4.0	27.1	3.8
% of owned by governments	1.1	0.4	2.1	1.1
% of owned by private companies	19.1	60.0	21.6	60.1
% of owned by public companies	1.4	24.7	1.3	25.2
% of owned by non-profits	4.8	7.2	6.3	6.4
% of neutral websites	19.5	68.8	23.3	73.3
% of right-wing websites	76.9	24.5	72.7	19.6
% of left-wing websites	3.5	6.8	4.1	7.1
Observations	1449	4838	1745	6499

Notes: These summary statistics are based on data provided by NewsGuard. NewsGuard assigns an aggregated score from 0 to 100 to each website based on a weighted average of how well it performs on its nine journalistic criteria, and considers websites that receive a rating below 60 to be untrustworthy websites. *Average score* shows the mean NewsGuard score and *% of untrustworthy websites* refers to NewsGuards' classification.

Table A2: Number and examples of companies whose ads appear on misinformation websites

Industry	N
Holding Companies	6767
Online Services	5347
Media	4749
Technology	4157
Govt., or Religion	3851
Business Solutions	3848
Household	3644
Travel	3484
Apparel	3373
Retail	3368
Insurance	3307
Telecommunications	3189
Digital Publishing	3111
Print Publishing	3103
Finance	3018
Health	2980
Babies & Kids	2344
Automotive	1766
Food or Beverages	1688
Industrial	1180
Education	1032
Dining	1028
Gas & Electric	457
Cosmetics	340
Arms	28

Notes: This table shows the number of unique companies whose ads appear on misinformation websites between 2019 and 2021 for each of the 25 industries in the Moat Pro dataset.

Table A3: Top 100 advertisers' activity on misinformation websites

	Appears on misinformation websites			Number of misinformation websites		
	(1)	(2)	(3)	(4)	(5)	(6)
Digital ad platform usage	3.36*** (0.06)	3.07*** (0.06)	3.41*** (0.08)	3.14*** (0.07)	2.82*** (0.07)	2.83*** (0.08)
Odds Ratio / Incidence Rate Ratio	28.71	21.63	30.35	23.19	16.76	17.03
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	Yes	Yes	No	Yes	Yes
Month FE × Industry FE	No	No	Yes	No	No	Yes
Number of Observations	11,969	11,969	11,969	11,969	11,969	11,969

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, $^{\circ}p < 0.1$

This table shows regression results for the top 100 most active advertisers in a given year based on weekly data collected for the years 2019 to 2021. In Columns (1)-(3), we show results from logit regressions where the dependent variable is *appears on misinformation websites*, a binary variable that takes value 1 when an advertiser appears on a misinformation website in a given week and zero otherwise. In Columns (4)-(6), we show results from poisson regressions where the dependent variable is *number of misinformation websites*, a continuous variable for the number of misinformation websites an advertiser appears on in a given week. *Digital platform usage* is a binary variable that takes value 1 when a company is using digital ad platforms in a given week and zero otherwise. Robust standard errors in parentheses.

Consumer study results

Table A4: Summary statistics and balance across treatment arms for the consumer survey.

	All	Information treatments					
	(1)	(2) Control	(3) T1	(4) T2	(5) T3	(6) T4	(7) p-value
Duration (in seconds)	1185	1005	1095	1032	1669	1122	0.14
Gender (Female)	0.52	0.53	0.50	0.53	0.55	0.49	0.17
Gender (Male)	0.47	0.46	0.49	0.46	0.45	0.51	0.13
Race (White)	0.78	0.82	0.78	0.77	0.77	0.77	0.07
Age (Below 45)	0.45	0.44	0.46	0.44	0.46	0.45	0.90
Residence (North East)	0.18	0.18	0.19	0.20	0.18	0.17	0.52
Residence (Midwest)	0.21	0.20	0.21	0.20	0.21	0.22	0.85
Residence (South)	0.40	0.40	0.41	0.39	0.39	0.40	0.90
Residence (West)	0.21	0.22	0.19	0.20	0.22	0.21	0.58
Household income (< 50K)	0.46	0.48	0.48	0.46	0.44	0.46	0.49
Education (No degree)	0.47	0.48	0.48	0.47	0.45	0.46	0.85
Education (At least college)	0.41	0.40	0.40	0.41	0.44	0.40	0.43
Employment (Working)	0.52	0.50	0.51	0.52	0.52	0.53	0.69
Employment (Not working)	0.47	0.49	0.48	0.46	0.47	0.46	0.76
Partisanship (Democrat)	0.44	0.42	0.43	0.47	0.47	0.42	0.12
Partisanship (Republican)	0.32	0.33	0.33	0.31	0.30	0.33	0.51
Vote (Trump)	0.32	0.33	0.33	0.31	0.30	0.35	0.14
Vote (Biden)	0.47	0.47	0.47	0.49	0.48	0.42	0.10
Vote (Other)	0.03	0.02	0.04	0.03	0.03	0.05	0.13
Vote (None)	0.18	0.17	0.17	0.17	0.19	0.18	0.75
Frequent user	0.57	0.53	0.57	0.58	0.56	0.61	0.03
Infrequent user	0.18	0.20	0.19	0.17	0.18	0.18	0.53
Prior petitions signed	0.35	0.35	0.33	0.37	0.35	0.36	0.64
Consumes misinformation	0.30	0.28	0.30	0.29	0.29	0.32	0.48
Media trust (Low)	0.34	0.36	0.33	0.33	0.33	0.34	0.78
Media trust (High)	0.25	0.25	0.27	0.25	0.24	0.24	0.50
Observations	4039	806	808	802	809	814	

Table A5: Comparison of responses across all petition outcomes.

	Company				Platform			
	Intention (1)	Clicks (2)	Reported (3)	Signed (4)	Intention (5)	Clicks (6)	Reported (7)	Signed (8)
Company (T1)	0.03* (0.02)	0.02 (0.02)	0.03 (0.02)	0.02 (0.02)	-0.01 (0.02)	-0.02 (0.02)	-0.02 (0.02)	-0.02 (0.01)
Platform (T2)	0.00 (0.02)	-0.01 (0.02)	0.00 (0.02)	-0.01 (0.02)	0.05** (0.02)	0.05*** (0.02)	0.04** (0.02)	0.04** (0.02)
Company and Platform (T3)	0.01 (0.02)	-0.00 (0.02)	0.01 (0.02)	-0.01 (0.02)	0.03 (0.02)	-0.01 (0.02)	0.03 (0.02)	0.00 (0.02)
Company Ranking (T4)	0.04* (0.02)	0.04** (0.02)	0.03 (0.02)		-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	
Controls	Yes	Yes	Yes	No	Yes	Yes	Yes	No
Control mean	0.22	0.15	0.21	0.12	0.21	0.14	0.20	0.11
Observations	4039	4039	4039	3225	4039	4039	4039	3225

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results for each of the four treatment groups (T1, T2, T3 and T4) across all of our petition outcomes. Columns (1) to (4) refer to company-specific petitions suggesting that companies like the respondent's top choice gift card company need to block their ads from appearing on misinformation websites. Columns (5) to (8) refer to platform-specific petitions suggesting that digital ad platforms used by companies need to block ads from appearing on misinformation websites. In columns (1) and (5), the dependent variable is the intention to sign a petition, a binary variable that takes the value 1 when a participant indicates wanting to sign a given petition and zero otherwise. In columns (2) and (6), the dependent variable is a click on the petition link that takes the user to the Change.org platform to sign a petition, a binary variable that takes the value 1 when a participant clicks on the link and zero otherwise. In columns (3) and (7), the dependent variable is the self-reported petition signature, a binary variable that takes the value 1 when a participant reports having signed a given petition and zero otherwise. We record actual petition signatures in columns (4) and (8). We omit signatures for the T4 group since these petitions were accidentally deleted by Change.org. Since we only observe actual signatures on the treatment group level, we cannot include controls and run regressions for these outcomes. To do testing, we calculate standard errors using the standard formula for (two-sided) proportion tests. For the remaining columns, we apply robust standard errors in parentheses.

Analysis of text responses

In our consumer survey, we ask our survey participants to briefly state the reason behind their choice of gift card and choice petition using an open-ended text field. We analyzed participants' text responses in order to understand their responses to each of these two behavioral outcomes. To do so, we first removed the names of companies from the text responses and then used the top2vec algorithm (Angelov, 2020), which automatically outputs the number of clusters and assigns. Top2vec uses word embeddings that account for the context of a word in a document, which is an advantage this method has over bag-of-word approaches like Latent Dirichlet Allocation (Blei et al., 2003).

For our exit outcome, we observe six topics emerge from the algorithm. We manually inspect the responses and find responses in these clusters correspond to responses that mainly mention misinformation-related concerns, how much they like a given company's products, how much they love products from a given company, how much they use a given company, how frequently they use a given company and specific features of a company's products. We further cluster together responses that mention how much they like a company and how much they love a given company's products together into a single "product preference" cluster. Similarly, we merge together responses in the use and frequency of use clusters into a single "product usage" cluster. We end up with four main clusters as shown in Figure 2. Table A6 shows sample text responses belonging to each cluster. These clusters are as follows:

1. Misinformation: a binary variable that takes a value of 1 if a participant indicated companies ads appearing on misinformation websites as being a factor contributing to their final gift card choice and zero otherwise.
2. Product usage: a binary variable that takes value 1 when a participant mentions their use or frequency of use of the product and zero otherwise.
3. Product preference: a binary variable that takes value 1 if a respondent mentions how they or their family like, enjoy or love the product they chose and zero otherwise.
4. Product features: a binary variable that takes a value of 1 if a participants refers to specific features such as how convenient, healthy or close to home the product they chose is and zero otherwise.

For our voice outcome, we take the same approach as above. This process results in five key clusters, which are shown in Figure 2 with select sample text responses in Table A7. These clusters are as follows:

1. Company responsibility: a binary variable that takes a value of 1 if a participant's response indicated that companies are responsible for their ads appearing on misinformation websites and zero otherwise.
2. Platform responsibility: a binary variable that takes value of 1 if a participant's response indicated that digital ad platforms are responsible for companies' ads appearing on misinformation websites and zero otherwise.
3. Misinformation concerns: a binary variable that takes value 1 if a participant's response mentions being concerned about misinformation and zero otherwise.
4. Best option: a binary variable that takes a value of 1 if a participant's explanation for their choice mentions the option they chose as being the best available option in their opinion and zero otherwise.
5. No interest: a binary variable that takes a value of 1 if a participant's response indicates that they would not like to sign an online petition and zero otherwise.

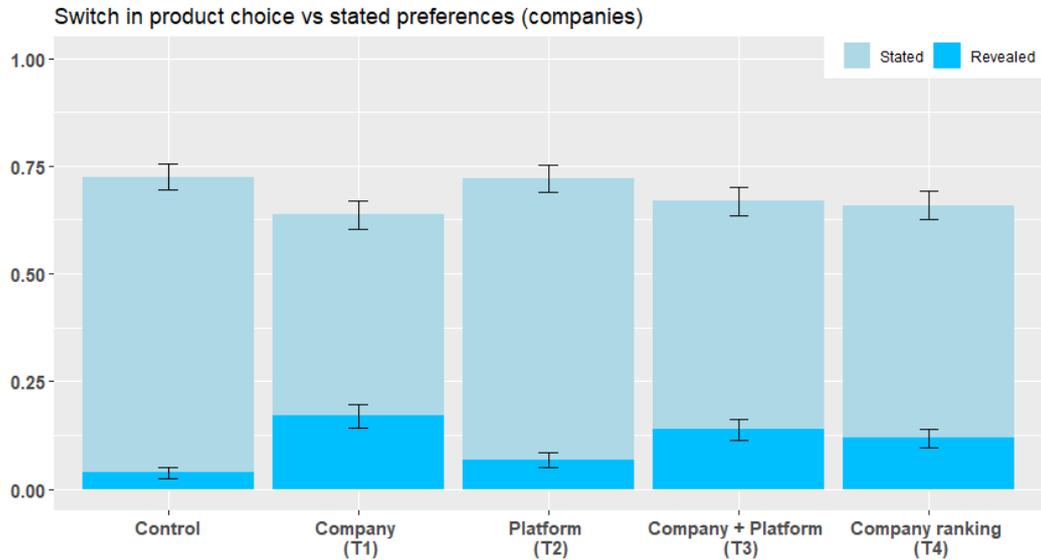
Table A6: Sample text responses from participants explaining their choice of gift card.

Sample	Text classification	Text response
1.	Misinformation	I will use a food delivery service more than using a driver service. I changed to grub hub because door dash allows their ads on websites with incorrect information.
2.	Misinformation	I first chose Uber as my choice because it is the only one that I use from the choices. However, I would happily switch to Lyft if their practices are more ethical.
3.	Misinformation	Subway was not a company that advertised on misinformation websites.
4.	Misinformation	I feel guilty about taking the burger king card if it is being used to further false information.
5.	Misinformation	I don't want to support the spread of misinformation.
6.	Misinformation	It was not listed among the sites that were linked to a misinformation site.
7.	Misinformation	I equally like door dash and grub hub but don't want to support a business associated with misinformation.
8.	Product usage	I can use this to go to work.
9.	Product usage	Doordash is the only company out of these choices that I use on a regular basis.
10.	Product usage	I chose this card because over the past two years I have bought more subs than other food places.
11.	Product usage	Because i would most likely use this gift card on my next visit to Burger King and it is less likely that i would use the others.
12.	Product usage	I chose Burger King because it's the only restaurant and service I actually use from the above list.
13.	Product usage	I frequent this restaurant quite a bit, so it would be a good fit for me.
14.	Product usage	I chose the above gift card because it's the one that I'd get the most utility from.
15.	Product preference	This is one of my favorite fast food restaurants.
16.	Product preference	I love Burger King. There plenty of items on menu that are worth getting excited about. Yummy food.
17.	Product preference	I eat at Subway and I like the food.
18.	Product preference	The have a selection that I like with fast delivery.
19.	Product preference	I would like Doordash because it is my go to food app. I love that I get to choose from a variety of food restaurants and even for beverages. My children love it as well and that gift card is going to go to them.
20.	Product preference	This gift card is the one that will be most beneficial for my family.
21.	Product preference	Subway is mine and my children's favorite local restaurant. We love to "eat fresh" and at subway everything is always fresh and delicious!
22.	Product features	subway is good to eat because of the calories that are in the food.
23.	Product features	I personally use door dash quite a bit and it fits into the convenience of my life.
24.	Product features	Health choice and trying to be healthy.
25.	Product features	Subway has convinient locations and great food at good prices.
26.	Product features	I chose this one because it is a lot closer and there is a person at burger king i am trying to become friends with.
27.	Product features	I am in a rural area now where food delivery is non existent so I would like it only to take my family out.
28.	Product features	I chose this gift card because there is a Subway close enough that i can walk to. I dont have a vehicle to drive to burger king and I dont believe lyft and uber are offered here.

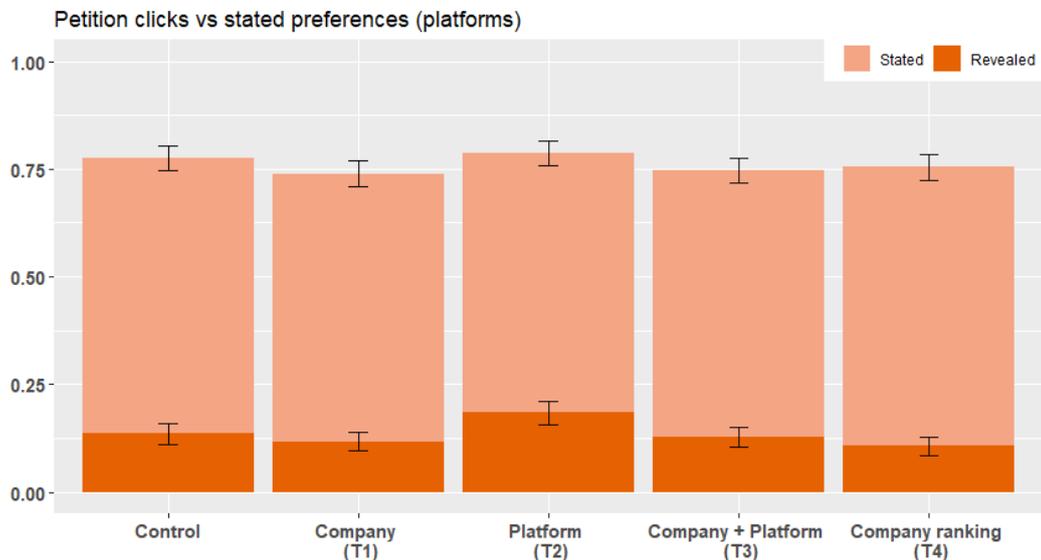
Table A7: Sample text responses from participants explaining their choice to sign an online petition.

Sample	Text classification	Text response
1.	Company responsibility	Companies like Subway absolutely should do this. The war on disinformation requires private and government action.
2.	Company responsibility	I think they should block their ads because of these misinformation sites causing their reputation harm.
3.	Company responsibility	Because it gives a bad reflection on the company and their brand if their ads are on websites that share misinformation.
4.	Company responsibility	All companies should be mindful of how they gain revenue and operate in society. Being ethical should always be at the forefront of their mission.
5.	Company responsibility	It can taint a company's image to be seen on misinformation websites.
6.	Platform responsibility	Because companies like subway depend on digital ad platforms to place their ads the responsibility lies with the ad platforms.
7.	Platform responsibility	Digital ad platforms should accept responsibility for placing ads on inappropriate and misleading sites.
8.	Platform responsibility	I feel like if we stop the use of ad platforms on misinformation sites in the first place then it would help out more in the long run.
9.	Platform responsibility	Digital ad platforms seem to make it easier to allow ads on misinformation websites.
10.	Platform responsibility	I feel that the onus is on digital ad platforms.
11.	Misinformation concerns	Supporting misinformation websites is horrible.
12.	Misinformation concerns	I do not want any misinformation sites to show ads.
13.	Misinformation concerns	Ads shouldn't help pay for misinformation.
14.	Misinformation concerns	I've always gotten misleading information on multiple occasions and needs to stop.
15.	Misinformation concerns	No one should be supporting misinformation.
16.	Best option	It eliminates more of the problem than the others.
17.	Best option	sounded like the most plausible choice.
18.	Best option	It is the best way to cancel out their problem
19.	Best option	It is the right thing to do.
20.	Best option	This statement seems to address the problem on a more widespread basis.
21.	No interest	I have not seen any of these ads we are taking the survey about.
22.	No interest	Freedom of speech. Up to consumers to educate themselves via various platforms.
23.	No interest	Who decides what is misinformation. Today these claims may be true, but if legislation is enacted and it becomes what corporations or government disagree with, this subverts the first amendment.
24.	No interest	I am not interested in governing what people or companies advertise or report as news. They are within their right to do so. This is America and in America people have the right to be wrong. If they don't want to do the research to find if the information they are getting is false than that's also people's right to be lazy. It's unfortunate but true.
25.	No interest	I don't want to sign the petition because its not for me to tell a company how or who to run their company ads whether i agree with it or not.

Figure A1: Participants' stated and revealed responses in terms of (a) exit and (b) voice.



(a) This figure shows participants' revealed preferences against their stated preferences regarding the role of advertising companies in financing misinformation. Revealed preferences are measured by the proportion of participants in each group who switch their gift card choice (i.e. "exit") after receiving the information treatment. Stated responses show the proportion of participants' who agree or strongly agree with the statement "Companies have an important role to play in reducing the spread of misinformation through their advertising practices". The vertical bars represent 95% confidence intervals.

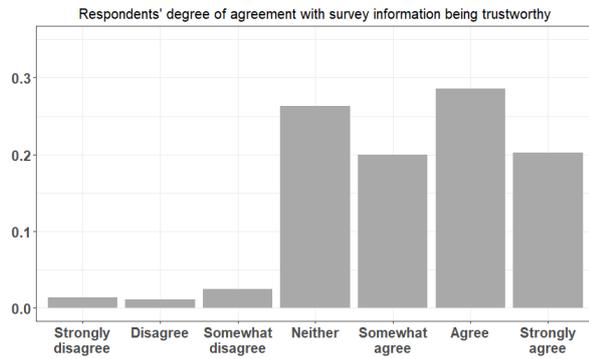


(b) This figure shows participants' revealed preferences against their stated preferences regarding the role of digital ad platforms in financing misinformation. Revealed preferences are measured by the proportion of participants in each group who click on a link to sign a petition suggesting that digital ad platforms should block ads on misinformation websites. Stated responses show the proportion of participants' who agree or strongly agree with the statement "Digital platforms should give companies the option to avoid advertising on misinformation websites." The vertical bars represent 95% confidence intervals.

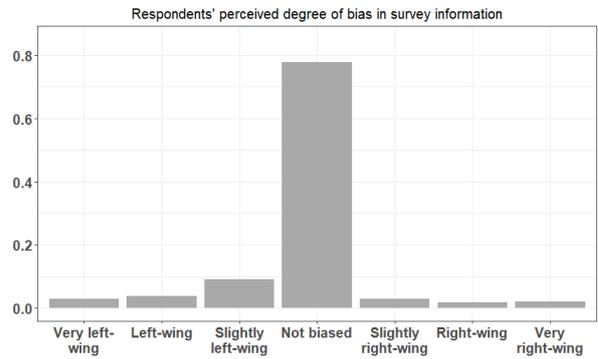
Table A8: Treatment prediction confusion matrices for the consumer experiment

	Predicted Control	Predicted Treated
Panel B: Control vs. T1		
True Control	92	107
True Treated	90	115
<i>Overall accuracy: 51.2%</i>		
Panel B: Control vs. T2		
True Control	106	92
True Treated	104	100
<i>Overall accuracy: 51.2%</i>		
Panel C: Control vs. T3		
True Control	106	91
True Treated	101	106
<i>Overall accuracy: 52.5%</i>		
Panel D: Control vs. T4		
True Control	87	123
True Treated	82	113
<i>Overall accuracy: 49.4%</i>		

Notes: This table presents the confusion matrices for the study purpose responses by participants in our consumer experiment. Each cell counts the number of participants assigned to the randomized group in the row and classified by the Support Vector Machine to be in the randomized group in the column.

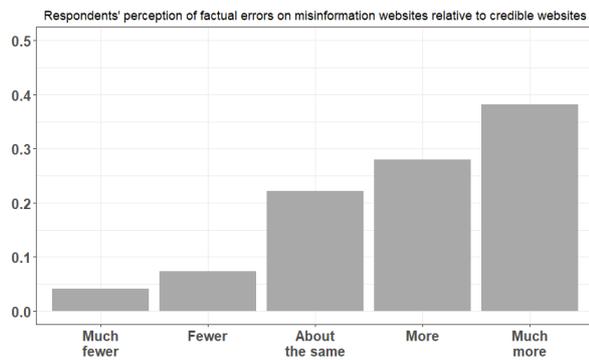


(a) Distribution of participants' responses to the question "The information provided in this survey is trustworthy."

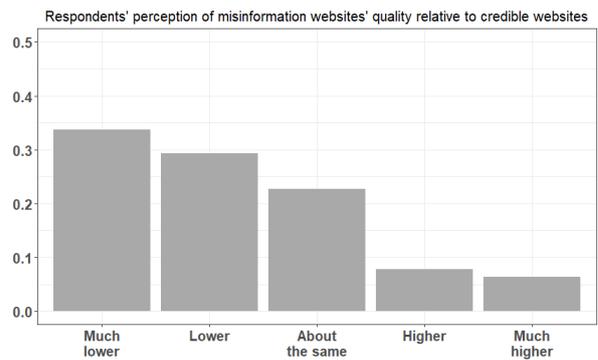


(b) Distribution of participants' responses to the question "Do you think that this survey was biased?"

Figure A2: Participants' perception of survey information.

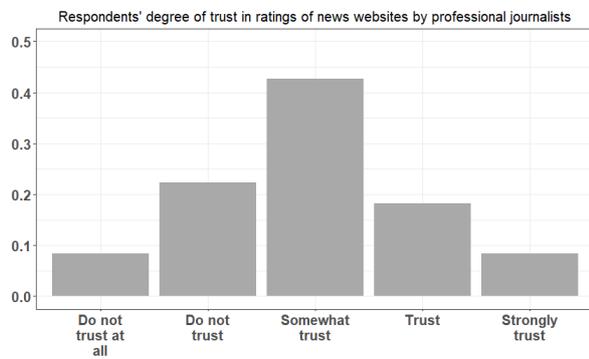


(a) Distribution of participants' responses to the question "Relative to credible websites, how many factual errors do you expect misinformation websites to have?"

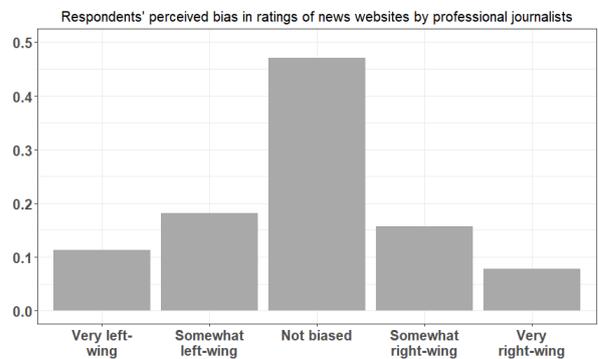


(b) Distribution of participants' responses to the question "Relative to credible websites, what quality do you expect misinformation websites to have?"

Figure A3: Participants' perception of misinformation websites.



(a) Distribution of participants' responses to the question "How much do you trust the ability of independent third party professional journalists to rate news websites?"



(b) Distribution of participants' responses to the question "What kind of political bias do you expect third party ratings of news websites by professional journalists to have?"

Figure A4: Participants' perception of website ratings by journalists.

Decision-maker study results

Table A9: Summary Statistics for Decision-maker Study

		Full sample	Treated	Control	P-value
Characteristics	Full-time employee	0.72	0.71	0.74	0.473
	Top executive role	0.56	0.58	0.54	0.467
	Marketing role	0.10	0.08	0.11	0.252
	Gender (Female)	0.21	0.21	0.20	0.808
	Duration in role (> 5 years)	0.58	0.60	0.57	0.512
	Number of employees (> 100)	0.59	0.55	0.63	0.080
	Headquartered in the U.S.	0.43	0.45	0.42	0.483
Beliefs	Estimated consumer backlash	0.41	0.40	0.42	0.708
	Company beliefs	64.3	65.8	62.9	0.308
	Prior platform beliefs	53.0	53.5	52.5	0.773
	Own company belief (yes)	0.20	0.19	0.21	0.634
	Own company belief uncertainty	0.21	0.23	0.19	0.288
Stated preferences	Misinformation control	0.88	0.87	0.88	0.641
	Company responsibility	0.76	0.80	0.73	0.062
	Platform responsibility	0.86	0.88	0.84	0.239
	Stated ad check demand	0.76	0.88	0.84	0.133
	Stated solution demand	0.71	0.72	0.69	0.598
Observations		442	216	226	

Notes: *Estimated consumer backlash* is a binary variable that takes the value 1 if a participant agrees or strongly agrees that “Consumers react against companies whose ads appear on misinformation websites” and zero otherwise. *Company beliefs* is the estimated proportion of companies whose ads appear on misinformation websites. *Prior platform beliefs* is the estimated proportion of companies that use digital ad platforms and whose ads on appear on misinformation websites. *Own company belief (yes)* is the proportion of respondents who report “yes” when asked whether they believe their own company’s ads appeared on misinformation websites. *Own company belief uncertainty* is the proportion of respondents who report being uncertain (unsure or very unsure) about their response to the aforementioned company belief question. *Misinformation control* is a binary variable that takes the value 1 if a participant agrees or strongly agrees that “It is important to control the spread of misinformation in society” and zero otherwise. *Company responsibility* is a binary variable that takes the value 1 if a participant agrees or strongly agrees that “Companies have an important role to play in reducing the spread of misinformation through their advertising practices” and zero otherwise. *Platform responsibility* is a binary variable that takes the value 1 if a participant agrees or strongly agrees that “Digital platforms should give companies the option to avoid advertising on misinformation websites” and zero otherwise. *Stated ad check demand* is a binary variable that takes the value 1 if a participant agrees or strongly agrees to the statement “I would like to find out whether my company’s ads are appearing on misinformation websites” and zero otherwise. *Stated solution demand* is a binary variable that takes the value 1 if a participant agrees or strongly agrees to the statement “I would recommend that my company adopt a product to avoid advertising on misinformation websites” and zero otherwise.

Figure A5: Distribution of Beliefs About Advertising Companies

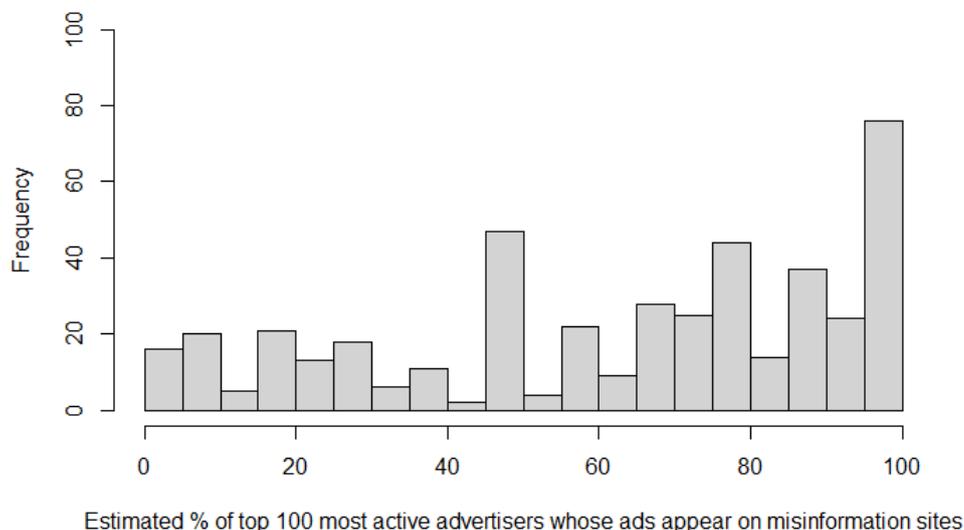


Figure A6: Distribution of Beliefs About Digital Ad Platforms

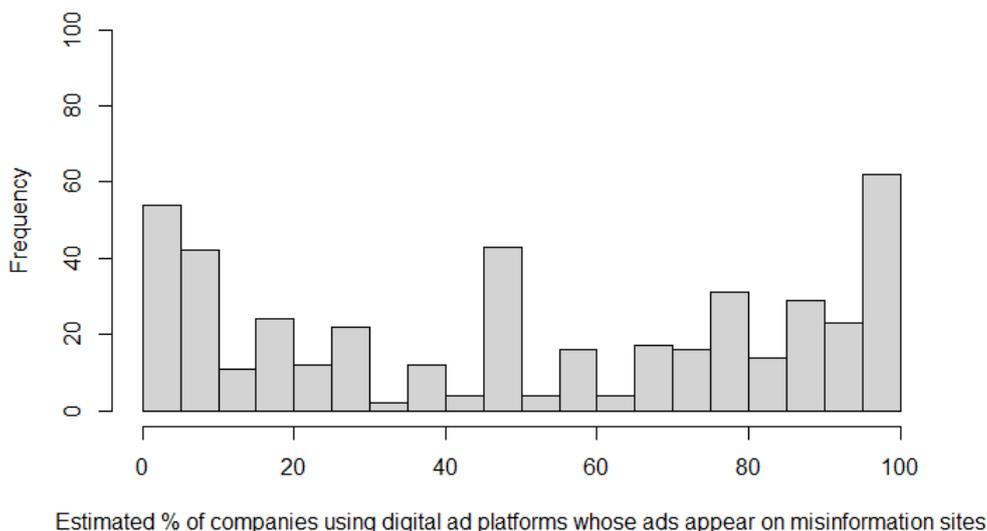


Table A10: Treatments Effects on Donation to the Global Disinformation Index (GDI)

	Full sample	Prior belief: No		
	(1)	All (2)	Certain (3)	Uncertain (4)
Treatment	-0.01 (0.05)	-0.00 (0.05)	-0.02 (0.06)	0.07 (0.17)
Observations	442	354	286	68
Controls	Yes	Yes	Yes	Yes

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows OLS regression results where the dependent variable is donation to the Global Disinformation Index or GDI, a binary variable that takes a value of one when a participant chooses to donate to GDI and zero when a participant chooses to donate to DataKind, the alternative charity option provided (see Figure A28 for the corresponding question). Column (1) shows results for the full sample of participants and columns 2-4 show results for the sub-sample of participants who reported “No” to the question “Do you think your company or organization had its ads appear on misinformation websites during the past three years (2019-2021)?”. In column 2, we report results for all participants who reported “No” to the aforementioned question. Column 3 shows results for participants who report being certain about their response to the aforementioned question (choosing “Somewhat sure”, “Sure” or “Very sure”). Column 4 shows results for participants who report being uncertain about their response to the aforementioned question (choosing “Unsure” or “Very unsure”). We control for decision-makers’ characteristics and prior beliefs in all specifications. Robust standard errors in parentheses.

Appendix B: Experiment Design and Survey Instruments

Consumer experiment

Design of the survey experiment

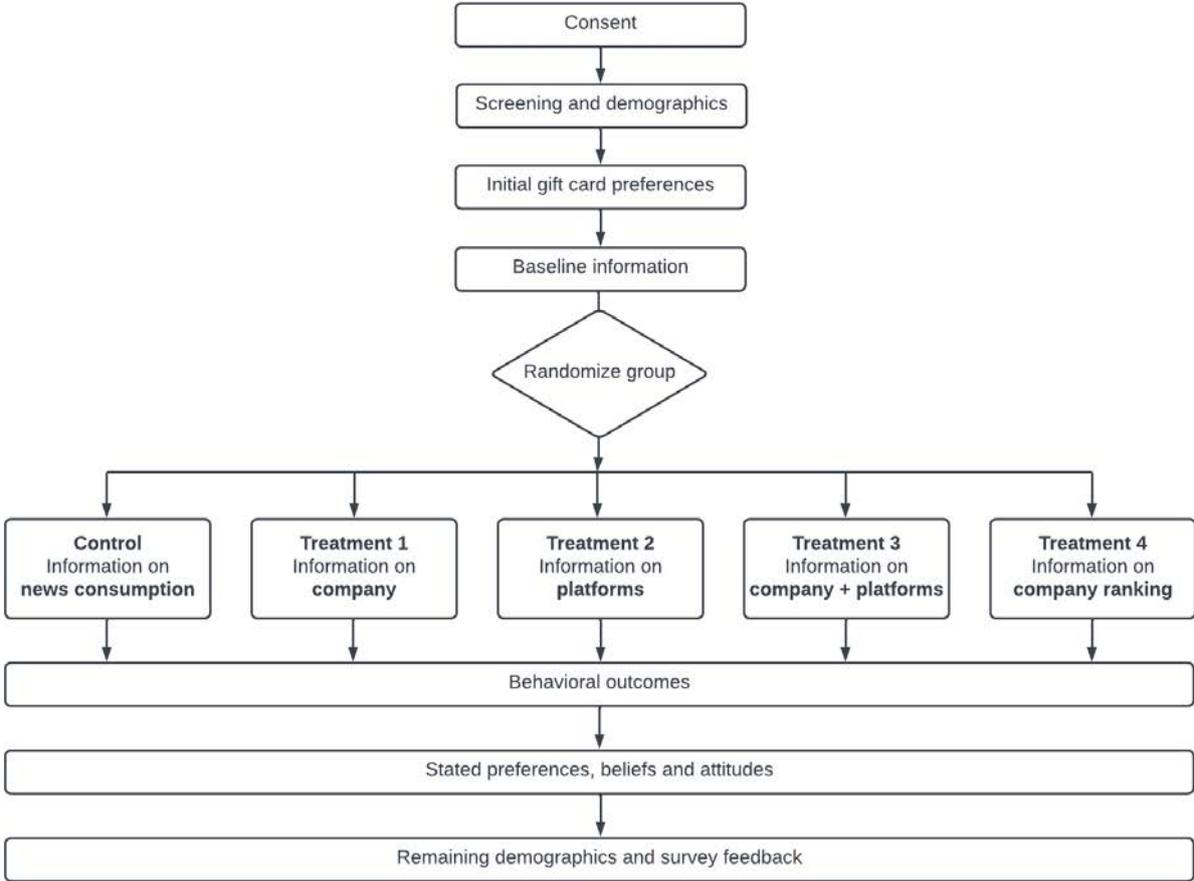


Figure A7: Design of the consumer survey experiment.

Measuring initial gift card preferences

Which gift card are you most interested in receiving?

Rank your choices on a scale from **your 1st choice (most preferred)** to **your 6th choice (least preferred)**.

	1st choice (MOST preferred)	2nd choice	3rd choice	4th choice	5th choice	6th choice (LEAST preferred)
Subway	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Grubhub	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Uber	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
DoorDash	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Burger King	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Lyft	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure A8: Measuring participants' initial preferences using a ranking of gift card options

How much weight would you assign to each of the following gift card options?

Higher weights mean greater interest. For example, if you are most interested in an Uber gift card, assign it a greater weight than all of the available options. If you are equally interested in Uber and Lyft gift cards, assign both options with equal weights.

These weights will help us understand your interests better and enable us to offer you relevant information depending on your interest.

Your responses must sum up to 100.

Subway	<input type="text" value="0"/>
Grubhub	<input type="text" value="0"/>
Uber	<input type="text" value="0"/>
DoorDash	<input type="text" value="0"/>
Burger King	<input type="text" value="0"/>
Lyft	<input type="text" value="0"/>
Total	<input type="text" value="0"/>

Figure A9: Measuring participants' initial preferences using continuous weights to gift card options that sum up to 100

Baseline information

Misinformation websites repeatedly present incorrect or misleading information as fact. Examples of misinformation websites identified in this way include: thegatewaypundit.com, infowars.com, rt.com, sputniknews.com, palmerreport.com, etc. These misinformation websites are identified by trained non-partisan professionals at [independent organizations](#).

Ads from companies appear on various websites, some of which are misinformation websites. These ads are how news websites make money.

In the below example, ads from athletics company Puma appear on zerohedge.com, a website that frequently spreads misinformation (e.g. misinformation about the Russia-Ukraine war) according to trained journalists. Thus, ads such as the one shown below by Puma contribute towards financially sustaining this misinformation website.

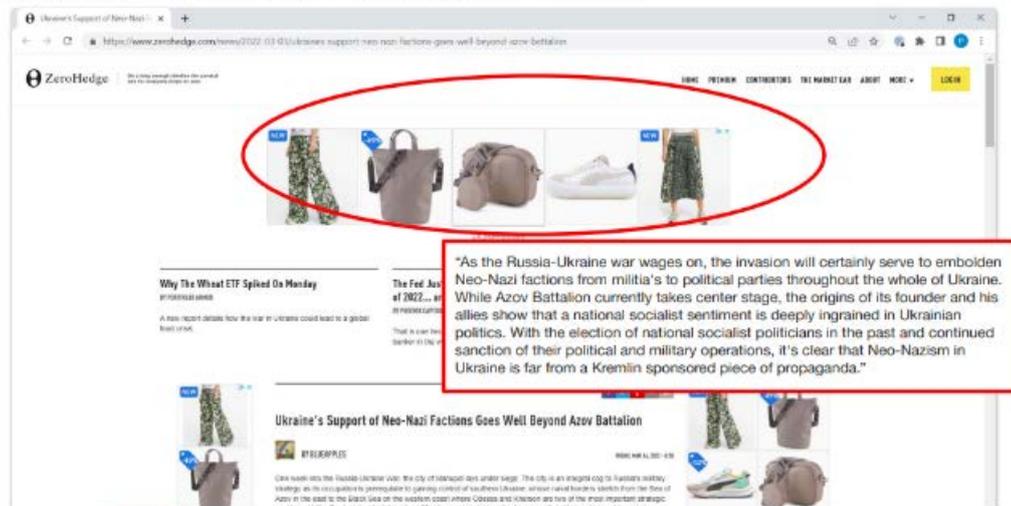


Figure A10: Participants in both experiments are given the above baseline information on misinformation and advertising prior to receiving randomized information treatments.

Randomized information treatments

A recent [study](#) found that Americans consume about five times more news from television than from online sources. Some of these news sources may include misinformation whereas others are trustworthy sources.

Figure A11: Information provided to the control group in the consumer experiment.

In the recent past, ads from X repeatedly appeared on misinformation websites. To provide more context, companies can adopt tools to control which websites their ads appear on.

Figure A12: Information provided to the “Company” treatment group (T1) in the consumer experiment. “X” is the top choice company chosen by the respondent prior to the information treatment.

Digital platforms (e.g., Google's DoubleClick platform, Microsoft's AppNexus platform, etc.) can be used to automatically place ads on different websites. In the recent past, companies that used digital ad platforms were about 10 times more likely to appear on misinformation websites than companies that did not use digital ad platforms as shown below:

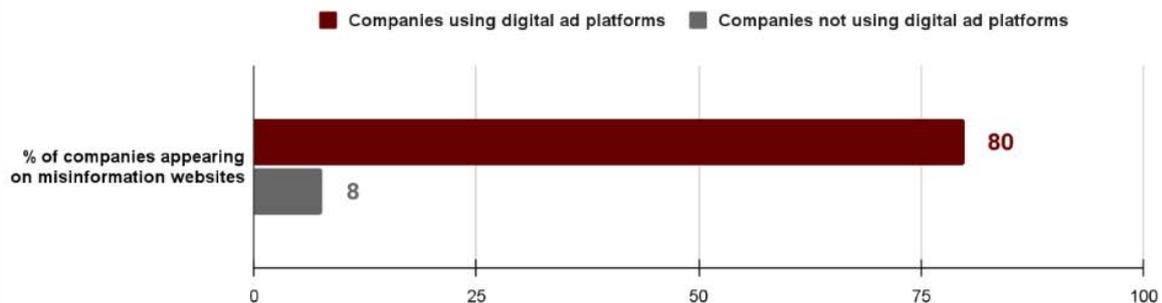


Figure A13: Information provided to the “Platform” treatment group (T2) in the consumer experiment. This information is also provided to treated participants in the decision-maker experiment.

In the recent past, ads from X repeatedly appeared on misinformation websites.

To provide more context, X uses digital platforms (e.g., Google's DoubleClick platform, Microsoft's AppNexus platform, etc.) for advertising. Such platforms can be used to automatically place ads on different websites. While companies can adopt tools to control which websites their ads appear on, companies that used digital ad platforms were about 10 times more likely to appear on misinformation websites than companies that did not use digital ad platforms as shown below:

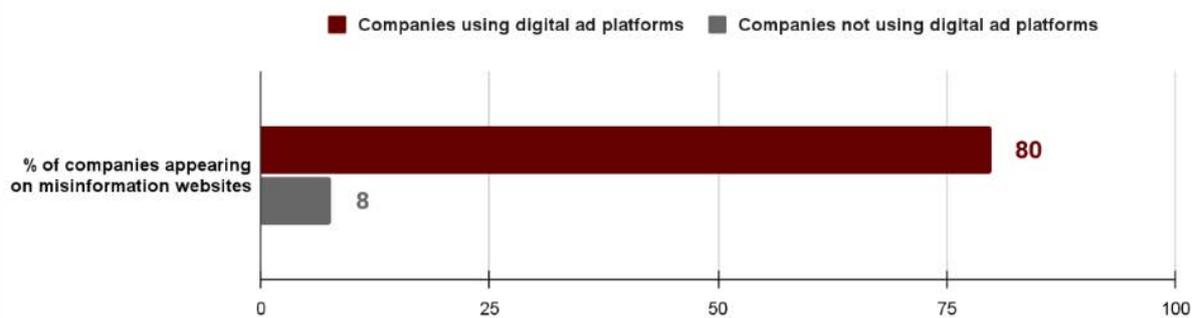
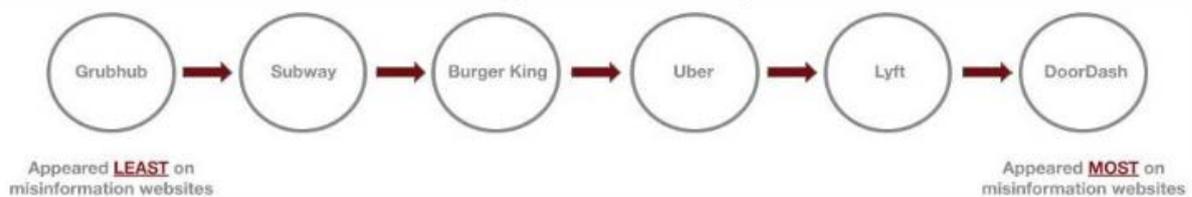


Figure A14: Information provided to the “Company and Platform” treatment group (T3) in the consumer experiment. “X” is the top choice company chosen by the respondent prior to the information treatment.

In the recent past, ads from all six companies below repeatedly appeared on misinformation websites in the following order of intensity:



To provide more context, companies can adopt tools to control which websites their ads appear on.

Figure A15: Information provided to the “Company ranking” treatment group (T4) in the consumer experiment.

Consumer experiment outcomes

Which **\$25 gift card** would you like to receive?

If you are selected, you will be given the gift card that you choose in this question.

- Subway
- Grubhub
- Uber
- DoorDash
- Burger King
- Lyft

Figure A16: Measuring exit by recording participants' final gift card choice. The gift card choices are provided in a randomized order to each participant.

Would you like to sign any of these petitions?

- Companies like Uber need to block their ads from appearing on misinformation websites.
- Companies like Uber need to allow their ads to appear on misinformation websites.
- Digital ad platforms used by companies need to allow ads to appear on misinformation websites.
- Digital ad platforms used by companies need to block ads from appearing on misinformation websites.
- I do not want to sign any petition.

Figure A17: Measuring voice by recording participants' intention to sign a petition on Change.org. With the exception of the last option above, the petition choices are provided in a randomized order. Participants who opt to sign a petition are then provided with a link to the petition of their choice.

Decision-maker experiment

Design of the survey experiment

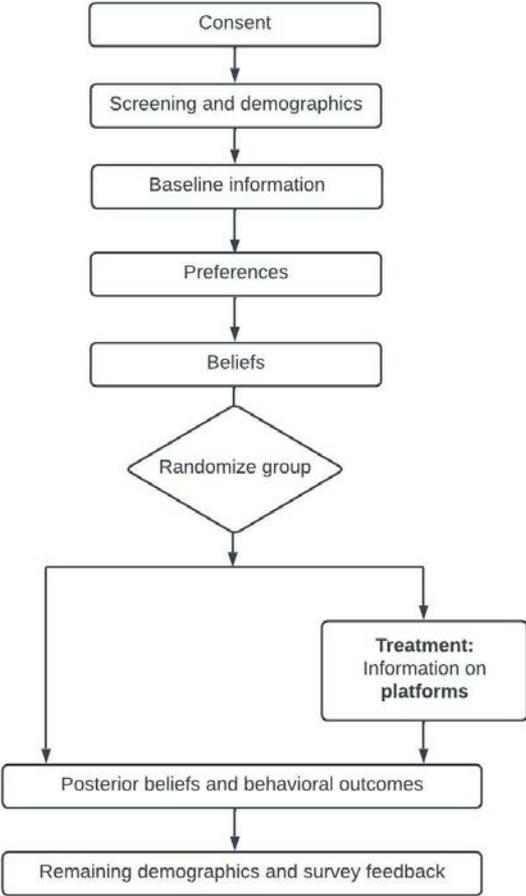


Figure A18: Design of the decision-maker survey experiment.

Measuring stated preferences and beliefs

We are interested in learning your general attitudes towards the topics discussed in this survey.

Please read each statement below carefully and rate how much you agree or disagree with it.

	Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
Companies have an important role to play in reducing the spread of misinformation through their advertising practices	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Consumers react against companies whose ads appear on misinformation websites	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I would like to find out whether my company's ads are appearing on misinformation websites	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I would recommend that my company adopt a product to avoid advertising on misinformation websites	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
It is important to control the spread of misinformation in society	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Digital platforms should give companies the option to avoid advertising on misinformation websites	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure A19: Measuring participants' stated beliefs and preferences, including the importance of controlling the spread of misinformation, the roles that companies and digital ad platforms should play in curbing misinformation, the potential for consumer backlash and the stated actions taken by the participant to get an ad check and adopt a solution to avoid advertising on misinformation outlets for their company. The rows are presented in a randomized order to participants.

Measuring revealed preferences

We conducted a research study to find out how people react to companies that advertise on misinformation websites.

In the research study, people were given information about specific companies that advertised on misinformation websites. We measured how many people responded to the information provided by changing their product choice and by opting to sign a petition advocating a change in the company's practices.

You can learn more about some of the study findings below. We will provide you with this information on the last page of the survey.

- I would like to find out how many people opted to sign a petition to voice concerns about company practices online.
- I would like to find out how many people switched away from getting the company's products or services.
- I would not like to receive this information.

Figure A20: Measuring participants' demand for information on consumer reactions

We analyzed data to examine which companies advertised on misinformation websites in the recent past. You can receive information on whether your company's ads recently appeared on such websites.

Please enter the name of your company or organization, and the email address at which we can send you this information. We will follow up to provide this information to all those interested within a month of survey completion. Please leave the fields blank if you are not interested in this information.

This information will be kept strictly confidential and will only be shared with you to inform you about where your company's ads have recently appeared.

My company's name is:

I would like to receive this information at the following email:

Figure A21: Measuring participants' demand for an ad check

Companies can adopt tools to automatically control which types of websites their ads appear on. Such tools provide regularly updated information on misinformation websites as they are identified by professional journalists.

You can sign up to learn more about how your company can manage where its ads appear via **a free 15-minute online information session** from experts in the field. You can sign-up for this information session yourself and/or invite another member of your organization to attend this session.

In order to sign up for this information session, please click on the following link: <https://forms.gle/NeRX8gPCuX4KXaa9A>. We will follow up with further details about attending this session after survey completion.

After signing up, please come back and complete this survey.

- I have signed up for this consultation.
- I would not like to sign up for this consultation.

Figure A22: Measuring participants' demand for a solution

Measuring prior beliefs

We would now like to get your best estimate on the number of companies whose ads recently appeared on misinformation websites.

Out of the top 100 most active advertisers, how many companies' ads appeared on misinformation websites during the past three years (2019-2021)?

Number of companies whose ads appeared on misinformation websites:

Figure A23: Measuring participants' beliefs about companies advertising on misinformation sites

Do you think your company or organization had its ads appear on misinformation websites during the past three years (2019-2021)?

Yes

No

How sure are you about your answer to the previous question?

Very sure

Sure

Somewhat sure

Unsure

Very unsure

Figure A24: Measuring participants' beliefs about their own company advertising on misinformation sites

Companies commonly use digital platforms (e.g., Google's DoubleClick platform, Microsoft's AppNexus platform, etc.) for advertising. Such platforms can be used to automatically place ads on different websites.

During the past three years (2019-2021), among 100 companies that **did not use** digital ad platforms, nearly **8** companies' ads appeared on misinformation websites on average.

Among 100 companies that **did use** digital ad platforms, how many companies' ads do you think appeared on misinformation websites? Please provide your best estimate below.

I think that out of 100 companies that used digital ad platforms, the number of companies whose ads appeared on misinformation websites were:

Figure A25: Measuring participants' prior beliefs (i.e. before receiving the randomized information treatment) about the role played by digital ad platforms in amplifying the financing of misinformation.

Decision-maker experiment outcomes

Misinformation websites can use digital ad platforms (e.g., Google's DoubleClick platform, Microsoft's AppNexus platform, etc.) to make money through advertising. These digital ad platforms place companies' ads on different types of websites, which can include misinformation websites.

In the recent past, misinformation websites that **did not use** digital ad platforms received ads from about 65 advertising companies per month on average.

For misinformation websites that **did use** digital ad platforms, what do you think is the average number of advertising companies per month whose ads appeared on such websites? Please provide your best estimate.

Average number of companies whose ads appeared monthly on misinformation websites that use digital ad platforms:

Figure A26: Measuring participants' posterior beliefs (i.e. after receiving the randomized information treatment) about the role played by digital ad platforms in amplifying the financing of misinformation.

We analyzed data for the most active 100 advertising companies during the past three years (2019-2021), including:

- 1) Analytics technologies that are used to improve ad performance, and
- 2) Platforms that least frequently place companies' ads on misinformation websites.

In this question, you can opt to receive this information to inform which technologies can be used by your company. We will provide you with the information you choose at the end of this study.

I would like to know which platforms least frequently place companies' ads on misinformation websites.

I would not like to receive this information.

I would like to know which types of analytics technologies are used to improve ad performance.

Figure A27: Measuring demand for information on which platforms least frequently place companies' ads on misinformation websites. The answer choices are presented in a randomized order to each participant.

The research team will donate \$100 to one of the following organizations.

1. The **Global Disinformation Index** provides ad tech platforms with trusted, non-partisan and independent ratings to identify and reduce advertising on misinformation websites.
2. **DataKind** helps mission-driven organizations increase their impact by unlocking their data science potential ethically and responsibly.

Which organization should we donate to? We will make this donation based on the decision of a randomly selected participant among the first 100 complete responses - it could be you.

DataKind

The Global Disinformation Index

Figure A28: Measuring participant's willingness to donate to the Global Disinformation Index (GDI). The answer choices are presented in a randomized order to each participant.