

Adapting Security Warnings to Counter Online Disinformation

Ben Kaiser
Princeton University

Jerry Wei
Princeton University

Eli Lucherini
Princeton University

Kevin Lee
Princeton University

J. Nathan Matias
Cornell University

Jonathan Mayer
Princeton University

Abstract

Disinformation is proliferating on the internet, and platforms are responding by attaching warnings to content. There is little evidence, however, that these warnings help users identify or avoid disinformation. In this work, we adapt methods and results from the information security warning literature in order to design and evaluate effective disinformation warnings.

In an initial laboratory study, we used a simulated search task to examine contextual and interstitial disinformation warning designs. We found that users routinely ignore contextual warnings, but users notice interstitial warnings—and respond by seeking information from alternative sources.

We then conducted a follow-on crowdworker study with eight interstitial warning designs. We confirmed a significant impact on user information-seeking behavior, and we found that a warning’s design could effectively inform users or convey a risk of harm. We also found, however, that neither user comprehension nor fear of harm moderated behavioral effects.

Our work provides evidence that disinformation warnings can—when designed well—help users identify and avoid disinformation. We show a path forward for designing effective warnings, and we contribute repeatable methods for evaluating behavioral effects. We also surface a possible dilemma: disinformation warnings might be able to inform users and guide behavior, but the behavioral effects might result from user experience friction, not informed decision making.

1 Introduction

Disinformation is spreading widely on the internet, often propelled by political motives [1, 2]. Platforms are responding by attaching warnings to disinformation content, in order to inform users and guide their actions. Facebook implemented disinformation warnings as early as December 2016 [3], and Google [4], Bing [5], and Twitter [6] have adopted similar content notices. There has been substantial public debate about the propriety of disinformation warnings, especially after Twitter began labeling tweets by U.S. President Donald Trump in May 2020 [7]. But recent studies provide scant evidence that these warnings can meaningfully influence user beliefs or behaviors, and it is an open question whether warnings are promising or futile for combating disinformation.

Security researchers faced a similar challenge over a decade ago, when studies showed that warnings for malware, phishing, and other online threats broadly failed to protect users [8,

9]. After a series of iterative, multi-method studies [10–21], security warnings now reliably inform user security decisions and help users avoid harmful and inauthentic content [10, 17]. In this work, we adapt methods and results from the information security warning literature in order to design and evaluate effective disinformation warnings.

A key finding from security research that we adapt to disinformation is that *contextual* warnings, which do not interrupt the user or compel action, are far less effective at changing behavior than *interstitial* warnings, which interrupt the user and require interaction [8, 9, 17]. Our work is, to our knowledge, the first to evaluate interstitial disinformation warnings.

Another relevant contribution from the security literature is a set of rigorous qualitative and quantitative methods for evaluating warnings, including structured models, realistic guided tasks, user interviews, and field studies (e.g., [11, 13, 15–18]). Our work adapts these methods to empirically examine contextual and interstitial disinformation warnings.

Across two studies, we use qualitative approaches (think-aloud exercises, interviews, and inductive coding) to understand user perceptions of disinformation warnings, as well as quantitative measures of the warnings’ effects on user behavior. We consider the following research questions:

1. After encountering contextual and interstitial disinformation warnings, how often do users change their behavior by opting for alternative sources of information?
2. Why do some users choose not to change their behaviors after encountering contextual and interstitial disinformation warnings?
3. Can interstitial warnings that are highly informative effectively change user behavior?
4. Can interstitial warnings that are highly threatening effectively change user behavior?

We first conducted a laboratory experiment ($n = 40$) in which participants searched for specific facts on Google and encountered an interstitial or contextual disinformation warning for certain search results (Section 3). The interstitial warning was substantially more effective at changing user behavior than the contextual warning, in large part because users did not notice or comprehend the more subtle contextual warning. In post-task interviews, participants described two reasons for the interstitial warning’s strong behavioral effect: the *informativeness* of the warning’s messaging and the *risk of harm* conveyed by the warning’s threatening design.

We then conducted a follow-on crowdworker study ($n = 238$), examining eight interstitial warning designs (Section 4). We confirmed the strong behavioral effects of interstitial warnings. We also found, however, that neither user comprehension nor perceived risk of harm appeared to moderate those effects.

Our results provide evidence that interstitial disinformation warnings can both inform users and guide user behavior. We demonstrate scalable and repeatable methods for measuring warning effectiveness and testing theories of effect. We also surface a possible dilemma: the behavioral effects of disinformation warnings may be attributable to user experience friction, rather than informed decision making. Our work highlights a path forward for designing effective warnings, and we close by encouraging iterative research and improvement for disinformation warnings—just like the information security community has successfully done for security warnings.

2 Background and Related Work

Disinformation research is dispersed across academic disciplines.¹ Recent work has predominantly focused on measurement (e.g., of content, campaigns, or user interactions) [22–40], or on developing automated detection methods [41–55].

In this section, we begin with background on disinformation websites, which are the targets for the warnings in our studies. We then discuss related work on the effects of and responses to disinformation. Finally, we describe the security warnings literature, which is the inspiration for this work.

2.1 Disinformation Websites

Disinformation campaigns are often multimodal, exploiting many different social and media channels at once [56]. These campaigns use websites as an important tool to host content for distribution across platforms and generate ad revenue [25, 28, 57–60]. Disinformation websites are designed to intentionally deceive users into believing that they are legitimate news outlets.² Our work examines whether warnings can counter this deception and help users contextualize or avoid disinformation websites.

2.2 Effects of Disinformation

Disinformation campaigns hijack the heuristics that users rely on to make accurate judgments about the truthfulness of information [63]. For example, disinformation campaigns often mimic credibility indicators from real news sources [64] or use social media bots to create the appearance of support [29].

Misperceptions that individuals hold after consuming disinformation are difficult to dispel [63]. Collectively, a misin-

formed populace may make social and political decisions that are not in the society’s best interests [65] (e.g., failing to mitigate climate change [66]). Influencing policy in this way—by shaping public perception and creating division—is a goal of many campaigns, especially by state-level actors [67].

Presenting a warning before exposure to disinformation can prevent harmful effects in several ways. Warnings can induce skepticism, so that users are less likely to take disinformation at face value [68]. Warnings can also make users more susceptible to corrections [63, 69]. Finally, warnings may cause users to not read the disinformation at all.

2.3 Responses to Disinformation

There are three main types of responses to disinformation that platforms and researchers have considered [70]. The first is deranking disinformation by changing recommendation algorithms [71]. Academics have studied this approach in simulated models of social networks [72–77], although not in realistic settings or with real users. Second, platforms have repeatedly removed disinformation content and banned accounts that promote disinformation [78–80]. Neither platforms nor researchers have established evidence on the effects of these takedowns. Finally, platforms have added warnings and similar forms of context to posts [4–6, 81, 82].

We note that there are important speech distinctions between these responses. When a platform removes content, it unilaterally makes speech less accessible to users. When a platform deranks content, it leaves the content available, but it unilaterally curtails speech distribution and discoverability. The potential promise of disinformation warnings is that they respond to problematic speech with counterspeech: platforms inform and protect users, without making unilateral decisions about content availability, distribution, or discoverability. As we discuss in Section 4.8, our work poses a possible dilemma for disinformation warnings as speech regulation: warnings can inform users and guide user behavior, but the behavioral effects may not be attributable to informed decision making.

Fact Check Warnings The most well-studied disinformation warnings are contextual labels indicating a story has been “disputed” or “rated false” by fact checkers [83–88]. These labels constituted Facebook’s first major effort to counter disinformation [81], and Google [4], Bing [5], and Twitter [6] have taken similar approaches. Facebook eventually discontinued use of “disputed” warnings after determining based on internal studies that the warnings were of limited utility [81]. More recently Facebook, Instagram, and Twitter all deployed new warning formats, including interstitials [89, 90].

Some studies of fact check warnings reported no significant effects on participant perceptions of disinformation [84, 85], while others found moderate effects under certain conditions [83, 87, 88, 91, 92]. Pennycook et al. found that fact check warnings caused participants to rate disinformation as less accurate after repeated warning exposure, but not with a single exposure [83]. Another study by Pennycook et al.

¹We use the term “disinformation” here and throughout this work, because our studies focus on warning users about intentionally misleading websites. We believe that our results generalize to misleading content regardless of intent (i.e., “misinformation”), because laboratory participants did not identify a website’s motive as a salient concern and we find that interstitial warning behavioral effects are not significantly related to warning content.

²These websites are sometimes termed “fake news” or “junk news” in related work (e.g., [57, 61, 62]).

identified a counterproductive *implied truth effect*: attaching a fact check warning to some headlines caused participants to perceive other headlines as more accurate [87]. Seo et al. found that fact check warnings caused participants to perceive stories as less accurate, but the effect did not persist when participants encountered the same stories later [88]. Mena found that fact check warnings had small negative effects on perceived credibility of news content on social media and self-reported likelihood of sharing [91]. Finally, Moravec et al. examined how fact check warnings can induce instinctual cognitive responses in users and cause users to thoughtfully incorporate new information into their decision making; a warning design that combined both mechanisms showed a moderate effect on social media post believability [92].

Related Links Bode and Vraga examined the effects of providing related links to alternative, credible sources of information alongside misinformation [93]. The study found that when related links debunked misinformation, participants who initially believed the disinformation showed a limited tendency toward corrected beliefs. Facebook, Google Search, and Bing all currently use related links warning designs.

Highlighting Falsehoods Garrett et al. tested a two-part warning, where participants were first informed that a fact-checker had identified factual errors in a story, then those errors were highlighted in the body of the story [94]. Among users already predisposed to reject the misinformation, this treatment significantly increased accuracy of beliefs, but it had no effect among users inclined to believe the misinformation.

Methods of Prior Work In all of these studies, participants were presented with screenshots of simulated social media posts, then were posed survey questions such as how truthful they thought the posts were and whether they would consider sharing the posts on social media. These methods can inform theories about how users will respond in real-world settings, but generalizations are tenuous because the methods involve highly artificial tasks and self-reported predictions about behavior. As we discuss below, security research has found that in order to measure realistic responses to warnings, it is important to design experimental tasks that involve realistic systems, realistic risks, and measurement of actual participant behavior [9, 11, 95, 96].

2.4 Security Warnings

Effective warnings are essential for security, because there are certain security decisions that systems cannot consistently make on behalf of users. Adversaries deliberately exploit judgment errors associated with these human-in-the-loop security decisions [97]. Early studies of security warnings found that the warning formats that were currently in use generally failed to protect users from online risks [8, 95]. Modern warnings, by contrast, are extremely effective: a recent study of over 25 million browser warning impressions showed that the warnings protected users from malware and phishing websites

around 75-90% of the time [10]. The immense progress in security warning effectiveness is due to numerous, rigorous studies that for over a decade have tested varied warning designs using diverse experimental methods and analytic lenses.

The primary methods of early security warning studies were laboratory experiments involving supervised tasks, user interviews, and surveys [8, 9, 11, 18, 20, 95, 96]. These studies examined users' beliefs and decision-making processes, in part by using structured models from warning science literature to identify reasons that warnings failed to change user behaviors. Security researchers typically used the Communication–Human Information Processing (C-HIP) model, which describes five information processing stages that users undergo when receiving a warning. Users must *notice* the warning, *comprehend* the warning's meaning, *believe* the warning, be *motivated* to heed the warning, and finally, *behave* in the manner intended by the warning issuer [98]. By determining the stage of the C-HIP model at which information processing was failing, researchers learned how to modify warning designs to increase the strength of the desired effect [9, 11].

Limitations It can be difficult to cause users to perceive realistic risk in a laboratory, requiring the use of deception and thoughtful experimental design [95, 96]. Laboratory studies must also address the challenge that participants may be more likely to disregard warnings that hinder task completion [11, 96]. Later research overcame these limitations using field studies, which measure users' reactions to warnings at scale in realistic environments [10, 13–15, 17, 19].

This body of work has definitively established that actively interrupting a user's workflow with a warning is more effective at preventing risky behaviors than passively showing the warning. Wu et al. compared popup warnings to toolbar icons and found that the popups caused users to behave with significantly more caution [8]; Egelman et al. observed that 79% of participants chose not to click through interstitial warnings compared to 13% for contextual warnings [9]. As a result, interstitials and other forms of active security warnings have become standard in all major browsers [17].

Several studies compared multiple warning designs and found that clear messages and use of visual cues can improve comprehension and adherence [11–14]. Personalizing messages based on user-specific psychological factors has not, however, shown a significant effect on adherence [16].

Limits of Analogizing to Disinformation Warnings The goals of security and disinformation warnings are not identical, so to study disinformation warnings, we must adapt—not simply reuse—the findings and methods from security warning research. Security warnings protect users from harms that are typically individualized, irreversible, and otherwise difficult for users to avoid themselves. The risks of disinformation, by contrast, are usually more collective and diffuse (see Section 2.2) and reversible (e.g., by receiving accurate information). Moreover, a user who encounters disinforma-

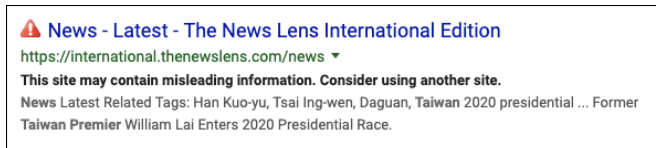


Figure 1: Search results pages displayed contextual warnings.

tion may be readily capable of protecting themselves from the risk (e.g., if they are media literate). As noted earlier, disinformation warnings also have speech implications that are distinct from security warnings. The differences between the security and disinformation problem domains motivate us to emphasize designs that inform users throughout our work.

3 Laboratory Study

We began with a laboratory study designed to examine how participants would process and react to contextual and interstitial disinformation warnings when searching for information. The search engine context is conducive to studying behavioral effects because participants have a concrete goal (finding a piece of information) and multiple pathways to achieve that goal (different search results and information sources).

We posed three research questions:

RQ1: In encounters with contextual and interstitial disinformation warnings, do users notice the warnings? Prior studies of contextual warnings note that one reason effect sizes are low or insignificant is that participants fail to notice the warnings. Effective warnings must attract user attention through conspicuous design or prominent placement.

RQ2: When users notice contextual and interstitial warnings, do they understand that the warnings have to do with disinformation? If a user misunderstands a warning, they may drop it from further cognitive consideration or respond in unintended ways that could increase risk.

RQ3: When users encounter and comprehend contextual and interstitial disinformation warnings, do they change their behaviors in the intended way by opting for alternative sources of information? This is an important outcome of warning exposure, which we aim to measure as described in Section 3.4.

3.1 Warning Designs

We adapted contextual and interstitial disinformation warnings from modern security warnings used by Google. Google’s warnings are well studied [10, 14, 15, 17, 19, 21, 99] and widely deployed, making them a useful template to design warnings that participants will believe are real.

We developed our contextual warning (Figure 1) based on a warning for compromised websites that Google displays in search results. We changed the color of the text from hyperlink blue to bold black to indicate that the text could not be clicked. We also added a red and white exclamation icon next to the search result to make the warning more noticeable.

We adapted our interstitial warning (Figure 2) from Google

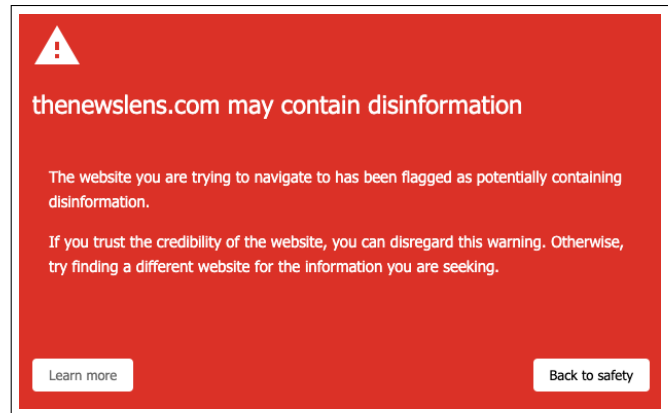


Figure 2: Participants encountered interstitial warnings after clicking search results.

Chrome’s warning page for malware. We modified the text to reference disinformation and changed the “Details” button to “Learn more.” Clicking “Learn more” revealed a message explaining that an automated system had flagged the site as disinformation and a “Continue” button that allowed the user to bypass the warning and continue to their selected page.

3.2 Study Design

In a laboratory setting, each participant completed a think-aloud role-playing task followed by an interview. By observing the participant during the task, we could tell if they noticed the warnings (RQ1) and altered their behavior in response (RQ3). Using the interviews, we could confirm whether the participant noticed the warnings, ask whether they comprehended the warnings (RQ2), and seek additional insights into how the participant processed the warnings.

Role-Playing Task The participant assumed the persona of an academic researcher trying to find answers to four questions using Google search. As our subjects were students, we believed this persona would be comfortable and aid with task immersion. For each question, we provided multiple sources of information and attached a warning to just one source so that the participant did not have to click through the warning to complete the task.³ Unknown to the participant, two questions were control rounds with no warnings and two were treatment rounds where warnings were inserted via a browser extension. We assigned participants in equal numbers to receive either contextual or interstitial warnings in both treatment rounds.

Search Tasks We designed the search tasks (shown in Table 1) to cover simple facts that could be easily found with a single Google search. For the treatment tasks, we selected facts specific to non-U.S. countries and covered by little-known, non-U.S. news sources in order to satisfy three additional design goals. First, the facts were related to current events due to the study’s focus on news disinforma-

³Studies of security warnings have shown that this “task focus effect” can bias participants’ behavior [11, 96].

tion. Second, so that participants could choose between multiple sources, each fact was publicly reported by at least two English-language news websites. Third, we aimed to select facts and websites that participants were not likely to be familiar with so as to avoid participants having preconceived biases about the information or the credibility of the sources.

3.3 Study Procedures

We explained the task and role to each participant at the beginning of their session. We asked the participant to behave as if they were using their own computer and to narrate their thoughts and actions as they performed the task. We framed the study as examining research and search engine usage behaviors to avoid priming participants to think about disinformation or warnings.

Participants began the task on the Google homepage. We informed participants that they could use either of two specific websites to find a particular piece of information, and that they should start with the first website and could switch to the second for any reason.⁴ Control rounds occurred first and third and treatment rounds occurred second and fourth, with the question order randomized within those sets.

We did not prescribe specific search queries to use, but most participants used a similar format: a set of terms relevant to the fact they needed to find combined with the name of the website on which they wanted to find the information. The participant would enter the query, navigate through results to find the requested information, and verbally provide an answer to the researcher. We would then instruct them to return to the Google homepage to begin the next round.

3.4 Data Collection

We took notes during each session to record how the participant selected search results, browsed the websites to seek information, reacted to warnings, and described their thoughts and actions. We also computed two metrics for each warning: a *clickthrough rate* (CTR) and an *alternative visit rate* (AVR).

Clickthrough Rate CTR is a commonly used metric in studies of security warnings. It measures the proportion of warning encounters where a participant dismisses the warning and proceeds, instead of going back. For contextual warnings, we recorded a clickthrough if the participant clicked a result that had an attached warning and a non-clickthrough if they chose a different result or changed the search query to use the second suggested website. For interstitial warnings, we recorded a clickthrough if the participant clicked “Learn more” and then bypassed the warning. If the participant clicked “Back to safety” or used the browser back button to go back to the search results, we recorded a non-clickthrough.

Alternative Visit Rate We also recorded whether participants visited an alternative source (i.e., the secondary website)

⁴This design directly parallels an evaluation of security warnings conducted by Sunshine et al. [11].

during a task, either because the user did not continue beyond the warning to the primary source or because the user sought to confirm the accuracy of the information from the primary source.⁵ We used this data to compute each warning’s AVR: the proportion of tasks where a participant visited an alternative source before completing the task. AVR is a new metric we devised for empirically measuring the behavioral effects of a disinformation warning.⁶ A high AVR indicates that a warning can influence users to decide to visit secondary sources of information.⁷ In some cases this will cause a user not to see the disinformation at all, and in all cases it exposes the user to alternative information.

Interview After the final round, we informed participants of the true purpose of the study, then conducted the interview. We first asked about the participant’s general understanding of disinformation: how they defined disinformation, what made them believe a website might contain disinformation, and if they had ever encountered content that they recognized as disinformation. Next, we asked the participant to describe their reactions to the warnings that they encountered during the study. We prompted participants to elaborate on these responses until we could determine whether the participant had noticed and comprehended the warnings (RQ1 and RQ2).

Before the next round of questions, we showed the participant printouts of the contextual and interstitial warnings used in the study. We then asked whether the participant believed the warnings would be effective in use, if they felt that one format would be more effective than the other, and if they had recommendations for how disinformation warnings in general could be made more effective.

Finally, we asked about the participant’s demographics, academic background, and level of news consumption.⁸

Coding We combined interview transcripts with our notes to form a single report for each session, then open coded the reports using Dedoose. One author performed the initial coding, producing 253 unique codes, then condensed the codes into themes. A second author validated this work, ensuring that the codes accurately reflected the study data and that the proposed themes were justified by the codes.

3.5 Participant Recruiting

We recruited participants through the Princeton University Survey Research Center, which advertises to randomly

⁵CTR and AVR are closely related: a non-clickthrough is a type of alternative visit. As a result, $AVR \geq 1 - CTR$.

⁶Another advantage of the AVR metric is that it is available in control conditions, not just treatment conditions. We did not record alternative visits for control rounds in the laboratory study, but we make extensive use of control round AVR in the crowdworker study.

⁷AVR does not capture user perceptions of warnings or the accuracy of user beliefs, which is why we pair this approach with qualitative methods. It is an important open question whether encounters with high AVR warnings are associated with more accurate beliefs, easier correction of misperceptions, or better ability to distinguish disinformation from real news.

⁸We list all survey questions in supporting materials [100].

Table 1: We measured clickthrough rates (CTR) and alternative visit rates (AVR) in treatment rounds of the laboratory study.

Round	Participant Instructions	Contextual Warning		Interstitial Warning	
		CTR	AVR	CTR	AVR
Control 1	Find the total area of Italy in square kilometers on Wikipedia or WorldAtlas.	–	–	–	–
Control 2	Report the price of a pair of men’s New Balance 574 on JoesNewBalanceOutlet or 6pm.com.	–	–	–	–
Treatment 1	Find the political party of Taiwan’s Premier on TheNewsLens or FocusTaiwan.	15/20	7/20	7/20	13/20
Treatment 2	Find the name of the girl reported missing in Barbados on Feb 11, 2019 on BarbadosToday or LoopNewsBarbados.	18/20	4/20	11/20	10/20

selected students. We also sent recruiting emails to distribution lists of various student organizations. We received 76 responses and selected 40 participants. Our participant group consisted of 16 women and 24 men aged 18-28 years old, studying across 17 disciplines.

Clearly this sample is biased in several respects, including age, education level, and social group. Later in this work, we evaluate a significantly more diverse sample recruited online (see RQ1 in Section 4). In the context of security warning studies, student populations have been shown to provide similar results to more representative samples [96].

The recruiting and consent materials provided to participants indicated that the study would take 30-45 minutes and focus on the user experience of search engines, with no mention of disinformation or browser warnings. Participants signed consent forms before beginning the study and were paid \$15. The study was approved by the Princeton IRB.

3.6 Results

We present quantitative results for the warnings’ behavioral effects (Table 1). We also discuss how notice and comprehension related to participant behavior and present qualitative results on user opinions and perceptions of the warnings.

3.6.1 Behavioral Effects

Contextual The CTR for the contextual warning was very high: 33/40. There were a total of 11/40 alternative visits: 7 non-clickthroughs and 4 occasions where a participant who clicked through a warning went back to search again using the secondary source.

Interstitial The CTR for the interstitial warning was much lower: 18/40. We observed 1 alternative visit after a click-through and 22 alternative visits after non-clickthroughs, for a total AVR of 23/40.

3.6.2 Notice and Comprehension

Contextual In the contextual treatment group, 13 out of 20 participants stated during interviews that they were not aware they had been shown disinformation warnings. All of these participants clicked through both warnings. 4 reported that they did not notice the warnings at all. Among the 16 that did notice the warnings, 9 noticed the icons but not the text.

Interstitial All 20 participants noticed the interstitial warnings. 12 understood that the warnings were about disinformation. 7 believed the warnings communicated a risk of “harm,” a “virus,” or another “security threat” and quickly clicked to go back without reading the text. The remaining participant clicked through both warnings; when asked why, he explained that he was focused on completing the study and “probably would have reacted differently” outside of the study.

3.6.3 Opinions on Disinformation Warnings

As part of the interview, we displayed printouts of both warning designs and asked for the participant’s opinions about the warnings’ relative merits and general effectiveness.

Contextual When asked which warning design they believed would be more effective in general, a small minority (6/40) chose the contextual warning. 5 of these participants were in the interstitial treatment group.

5 participants noted that the contextual warning could be seen before a user “commits” by clicking a link. 1 participant explained, “*you’re immediately presented with alternatives, whereas for the interstitial I’m already there and committed a click, so I want to go forward.*” Another preferred the contextual warning because it was easier to bypass: “*[I] just wanted to find a link to click on very quickly, it doesn’t take as much effort to avoid compared to the interstitial.*”

5 other participants emphasized the “always-on” nature of the contextual warning. 1 participant liked how they could “*always see the warning when browsing Google search results without having to click around.*” Another felt that the contextual warning was paternalistic because it tilted the otherwise level playing field among search results, “*direct[ing] you to which [results] you should visit.*”

15 participants said that the contextual warning was not noticeable. 1 specified that “*the exclamation point is very subtle... you’re not going to notice it.*”

Interstitial 34 participants—an overwhelming majority—believed that the interstitial warning would be more effective in general. When asked why, 32 mentioned that it was more noticeable. 17 mentioned the color red, with 1 participant noting that “*everybody knows red means stop.*”

19 participants remarked on how the warning requires user input to proceed. 1 participant observed that *“it stops the flow of the user and forces them to take an action.”* Other responses suggest that design cues contributed to the warning’s effectiveness; participants mentioned that the red color and text *“implied that the user is in danger”* and that the text was *“more instructive than the text on the contextual warning.”*

When asked about drawbacks to the interstitial warning, 5 participants focused on the inconvenience and the potential for warning fatigue. 1 participant noted that they would *“probably turn it off in the settings”* if the warning showed up frequently.

Improving Warning Designs Many participants (17) suggested that more informative disinformation warnings would be more effective. Recommendations included adding *“more about why this particular site was flagged,”* definitions of terms, and more explicit descriptions of potential harms. Conversely, 7 others urged *“short and concise”* messages that *“[get] the point across quickly.”*

5 participants suggested using different warnings depending on severity and whether the user had visited the website before. Another 5 recommended that warnings persist even if a user had visited a website before, arguing that warnings would be more effective if users were *“consistently reminded that [the page] may not be completely safe or factual.”*

Participants also suggested alternate warning forms: pop-ups, banner messages, or highlighting false claims.

Trust The source of the warning was important to many participants. 8 indicated that they were more likely to be deterred by a browser warning if they knew that it was triggered by Google, since they trusted the company’s judgment. 1 participant explained that they clicked through the interstitial warning because they understood that Princeton University had flagged the website, and they felt that the university was not a credible source of judgment about online disinformation. Another theme underlying trust judgments was previous experiences with browser warnings. 7 participants expressed that they distrusted the warnings due to previous encounters with false positive warnings or overly restrictive site blockers on institutional networks (e.g., in high school).

Risk 7 participants expressed the opinion that disinformation is not a serious threat or that it is not as harmful as malware. One participant explained that they typically comply with browser warnings but reacted differently to the disinformation warning: *“I don’t like the idea of someone telling me where or what I am allowed to access. You can give me suggestions. It was because I realized it was a disinformation warning and not a malware warning that I went back to try and get to the website.”* Another participant characterized this sentiment sharply, saying *“[d]isinformation warnings should not make it harder to access the site.”*

3.7 Discussion

Contextual vs. Interstitial Warnings The interstitial warning was distinctly more noticeable and comprehensible than the contextual warning, and also far more effective at inducing behavioral changes. Similar findings in security warning research prompted the field to shift from contextual to interstitial formats. Platforms are only just beginning to make this shift for disinformation; contextual warnings are currently the dominant format in both research and deployment. While contextual warnings may still have a role in countering disinformation, interstitial warnings and other forms of active, interruptive warnings clearly merit further consideration.

Impact of Visual Design Choices The iconography, colors, and text styles in our warnings impacted participant attention, comprehension, and behavior. The icon we added to the contextual warning made the warning more noticeable but did not necessarily aid with comprehension, as many participants who noticed the icon still failed to notice the text. The red background of the interstitial warning contributed to its effectiveness, but may also have reduced comprehension as participants seemed to react quickly upon seeing the red color without taking the time to read or understand the warning. Again drawing from security warning research, future work should use comparative effectiveness studies to isolate the effects of individual visual design choices.

3.7.1 Mechanisms of Effect

So few participants complied with the contextual warning that it is difficult to draw conclusions about what caused the behavioral effect. For the interstitial warning, however, we found evidence for three different mechanisms by which the warning induced behavioral changes.

Informativeness Warning science literature focuses on educating the user and enabling them to make an informed decision about how to proceed [101, 102]. Across both warning designs we tested, participants who understood the warning visited an alternative website in over half of cases (21/38), while participants who did not understand the warning did so in only a third of cases (13/42). Moreover, nearly half of participants recommended making the warnings more informative to improve effectiveness. These results reinforce that informing users is a possible mechanism of effect for interstitial disinformation warnings.

Fear of Harm The interstitial warnings had a threatening appearance, which many participants identified as a factor in why they did not continue. Some participants visited an alternative site because they perceived a non-specific risk of personal harm, without comprehending the warning. Other participants misinterpreted the warning and believed that it described a risk of receiving a computer virus or other security threat. If a warning conveys a risk of harm, it should be specific and narrowly scoped; otherwise users may perceive the warning as irrelevant or a false positive, which could reduce

the behavioral effect [10, 15, 18, 20, 103]. As long as the specific harm is clear, though, using design cues to further convey a general risk of harm may improve warning effectiveness.

Friction The interstitial warning’s strong effect was due, in part, to the friction it introduced into the task workflow. Some participants preferred to choose another source rather than read the warning, decide whether to believe and comply with it, and click through it. As with the “fear of harm” mechanism, friction must be carefully calibrated to avoid inducing warning fatigue or habituating users to ignore warnings. Friction also has serious drawbacks as a causal mechanism: it degrades the user experience, makes platforms more difficult to use, and does not rely on an informed decision about disinformation.

3.7.2 Limitations

Security warning research has observed challenges in studying behavioral responses to risks in laboratory settings, particularly with respect to ecological validity [96]. We encountered similar challenges in this study.

Our sample was small in size and biased in several ways; our findings should be understood as illustrative but not representative. It is important to identify if our findings can be replicated by larger, more diverse populations.

Because participants used our computer and we were watching during the task, some participants reported that they behaved differently in the study than they might have in real life. Others may not have reported this effect because they were reluctant to admit that they behaved with bias or because the effect was unconscious. An experimental design that allows participants to use their own computers in their own environments (i.e., not in a laboratory) could offer more realistic observations of how participants assess risk.

Although participants appeared to be driven to complete the research tasks, they were not personally invested in completing the tasks or finding correct answers to the queries. The role-playing aspect of the study may not have been strongly immersive, and there were no extrinsic rewards or penalties that incentivized correct answers.

Finally, because our search tasks did not pertain to participants’ social or political contexts, participants had little reason to engage in motivated reasoning. Motivated reasoning can strongly influence a user’s perceptions of information relevant to their social or political beliefs [104], so in those contexts, the warning effects that we demonstrate may be weaker.

4 Crowdsworker Study

In our second study, we aimed to verify the behavioral effects of interstitial disinformation warnings. We also examined the mechanisms for those effects, so that we could reason about the utility and limitations of deploying the warnings.

Our research questions were:

RQ1: Do interstitial disinformation warnings cause users to choose alternative sources of information? We investigated whether population sample bias or task design

significantly affected the results of our laboratory study. We recruited a larger, more diverse participant pool from Amazon Mechanical Turk (Section 4.6) and adjusted the task to account for limitations in the laboratory study (Section 4.2).

RQ2: Do interstitial warnings that effectively inform users about the risks of disinformation cause users to choose alternative sources of information? We tested whether participants understood the warnings, then compared the behavioral effects of informative and uninformative warnings to isolate the impact of informativeness on behavior.

RQ3: Do interstitial warnings that communicate a risk of personal harm cause users to choose alternative sources of information? We tested whether warnings caused participants to fear harm, then compared the behavioral effects of warnings that did and did not evoke a fear of harm.

RQ4: Does user partisanship (with respect to U.S. politics) moderate behavioral effects or perceptions of interstitial warnings? Research in political science indicates that political orientation affects judgments of information credibility [105] and efficacy of misinformation warnings [87]. We included this research question to detect if partisan alignment created a bimodal distribution in responses to warnings.

The task structure and key behavioral outcomes remained the same as in the laboratory study. We informed participants that they were joining a study of search engine usage and research behaviors. We then guided participants through four research tasks using a search engine, alternating between control and treatment rounds. In each treatment round, the participant encountered one of eight candidate interstitial disinformation warnings after clicking certain search results. We measured whether the participant clicked through the warning and whether they visited an alternative website. We examine the CTR and AVR across all observations to answer RQ1.

We used surveys after each warning encounter to measure how informative the warning was and how strongly the participant perceived the warning to convey a risk of harm. RQ2 and RQ3 concern the relationship between AVR and these survey responses. A standard analysis approach would have been to randomly assign participants to warnings, then compute statistical tests across the conditions. Unless the differences in effect between warnings were dramatic, however, this approach would have required a massive number of observations on each warning to establish statistical significance.

We instead employed a multi-armed bandit algorithm, which allows efficient exploration of a larger design space than is traditionally possible. As we received successive observations, the bandit increased the odds that participants encountered the warnings proving to be most and least informative and most and least effective at conveying fear of harm. After all observations were completed, significantly more participants had encountered these top- and bottom-performing warnings, providing us with the statistical power needed to test our hypotheses. Section 4.5 discusses the design and implementation of the multi-armed bandit algorithm.

4.1 Warning Designs

We created eight candidate interstitial disinformation warnings: four designed for informativeness and four designed to evoke fear of harm (Table 2). The warnings shared a common layout, consisting of an icon, a title, a primary message, an optional detailed message, and two buttons. This layout differed from the laboratory interstitial warning in two ways.

First, in the laboratory warning design, the detailed message (and the “Continue” button) were hidden at first and would only be revealed after the participant clicked “Learn more.” In the crowdworker study, we wanted to ensure that the full warning message was always displayed, because part of what we were measuring was the effect of different messages. We eliminated the “Learn more” button and instead displayed the detailed message and the “Continue” button on all warnings.

The second change addressed the “Back to safety” button. This button text implied that the user was in danger, which was inappropriate for the informative warnings. We changed the button to read “Go back” and applied this change to both informative and harm-based warnings in order to maintain a common interaction design across all warnings.

For both groups of warnings, we generated several options for icons, titles, and messages. We then created candidate designs by choosing combinations of these elements and inserting them into the layout template.

Informative Warnings We designed the informative warnings to be visually nonthreatening and clearly explanatory in their messages (see Figure 3a). The warnings included one of two icons—either a generic exclamation point icon or a policeman silhouette—and displayed black text against a white background. The warning messages explained the nature and consequences of disinformation in varying detail: some explicitly defined the term “disinformation,” some asserted that “experts” had labeled the site as containing “false or misleading” information, and others provided clear guidance on how to behave (“finding alternative sources of information”).

Harm Warnings The harm warnings contained less text and used forceful language, colors, and icons to suggest a serious threat (see Figure 3b). The warnings used either a skull-and-crossbones icon or a policeman icon, and the content was colored white against a red background. The most extreme warning design simply said: “WARNING: This website is dangerous.” The other three warning designs were titled “Security Alert” and indicated in their messages that the threat had to do with information quality.

4.2 Task Design

We aimed to keep procedures for the crowdworker study as similar to the laboratory study as possible, but the different setting and research questions necessitated three changes.

First, because crowdworkers participate remotely and use their own computers, we could not easily measure their browsing, insert warnings, or control search queries or results. In-

stead of using live Google searches, we developed a web application to guide participants through the experiment and realistically simulate a search engine (Figure 4). We populated results for each query from real Google results for the query, including the snippet for each search result. In order to simulate the story page after clicking a search result—and to ensure that participants saw the same content—we used full-page screenshots of story pages. Participants could browse these screenshots similar to real webpages (Figure 5). Unlike in the laboratory study, we specified the queries to use and did not direct participants to specific sources.

Second, crowdworkers participated in our study to earn a wage. This is a very different motivation than that of our laboratory subjects. Crowdworkers may be more focused on completing the task quickly (in order to earn the task fee) than on engaging meaningfully with the study and behaving as “good” research participants. One way we addressed the risk was to ensure that only workers with track records of submitting quality work participated in the study (see Section 4.6). We also offered a bonus of \$1 (43% of the base fee) to participants who correctly answered all four search questions. The bonus incentivized crowdworkers to engage with the tasks, read instructions carefully, seek accurate information, and take disinformation warnings seriously.

Finally, we used a series of surveys in lieu of directly observing participants and conducting interviews. Each participant completed a pre-task survey about their partisan alignment, surveys in each round about their behavior and perception of the warning, and a post-task demographic survey.⁹









Search Tasks As in the laboratory study, we selected facts for participants to retrieve that were reported by multiple sources and were obscure enough that participants would likely be unfamiliar with the topic or sources. We also ensured that all search results came from news outlets and that no two results came from the same outlet, giving participants a greater variety of choices for sources of information. All four tasks pertained to events in the U.S. to make the topics more relevant to our U.S.-based participant population. We also designed the treatment tasks to cover political scandals, so that participants would find it plausible that news outlets might publish disinformation about these topics. Table 3 presents the queries for control and treatment rounds in the study.


Procedures After accepting our job on Amazon Mechanical Turk, participants navigated to our study web application. The landing page displayed instructions and a visual guide for the study user interface, then directed participants to begin the first search round.

Each round consisted of a research task where the participant used our simulated search engine to find a particular fact. The participant began on a generic search query page, which specified the fact to search for and the query to use (Figure 4).


⁹We provide survey details in supporting materials [100].

Table 2: We developed eight interstitial warning designs for the crowdworker study. Figure 3 shows sample designs.

	Harm (white on red background)				Informativeness (black on white background)			
ID	h1	h2	h3	h4	i1	i2	i3	i4
Icon	Skull 	Skull 	Policeman 	Policeman 	Exclamation 	Policeman 	Policeman 	Exclamation 
Title	WARNING	Security Alert	Security Alert	Security Alert	False or Misleading Content Warning	Fake News Warning	False or Misleading Content Warning	Fake News Warning
Primary message	This website is dangerous.	This website contains misleading or false information.	This website is dangerous.	This website contains misleading or false information.	This website presents itself as news, but it contains information that experts have identified to be false or misleading	This website contains misleading or false information.	This website contains misleading or false information.	This website presents itself as news, but it contains information that experts have identified to be false or misleading
Details	None	None	Consider finding alternative sources of information.	None	This website spreads disinformation: lies, half-truths, and non-rational arguments intended to manipulate public opinion. It can be difficult to tell the difference between real news and disinformation, but it poses a serious threat to national security, election integrity, and democracy.	This website spreads disinformation: lies, half-truths, and non-rational arguments intended to manipulate public opinion. It can be difficult to tell the difference between real news and disinformation, but it poses a serious threat to national security, election integrity, and democracy.	Consider finding alternative sources of information.	Consider finding alternative sources of information.



(a) Informative warning design i2



(b) Harm warning design h3

Figure 3: Examples of warning designs in the crowdworker study. Table 2 describes all eight warning designs.

When the participant submitted the query, our study application presented a search results page populated with eight results. Clicking on a result led to a story page containing the news article that the search result snippet described (Figure 5). These story pages were full-size screenshots of real news article webpages, allowing participants to scroll through and read the articles as if they were browsing the real webpages.

Three of the results in each round were *target results*: the search result snippets clearly pertained to the query topic, and the story pages contained articles that clearly provided the answer to the task question. The other five results were *nontarget results*, which linked to story pages that did not contain the answer. Some nontarget results could be readily identified as irrelevant from the results page snippet (i.e., they

Table 3: The crowdworker study involved two control rounds and two treatment rounds. We report outcomes in Table 4.

Round	Participant Instructions
Control 1	A train from the Massachusetts Bay Transportation Authority derailed in Boston on June 8. Please find the color of the line that the train was on using the query “MBTA derailment.”
Control 2	In Alabama, the Muscogee tribe filed a lawsuit over a casino that was built on sacred ground. Please find the name of the casino using the query “muscogee nation casino lawsuit.”
Treatment 1	Please find the names of the two Clatskanie School Board members who were recalled in May 2018 using the query “clatskanie school board recall.”
Treatment 2	Mike Sanders, the Executive of Jackson County, Kansas, was sentenced to prison in 2018 for embezzlement. Please find the name of the political party to which he belonged using the query “mike sanders jackson county.”

contained some search terms but in a different context), while other nontarget results were germane to the query topic but the story page did not contain the answer to the task question. We ordered search results so that the target results would be easy to find: the top result was always a target result, and the other two target results appeared randomly within the top five results. The rest of the results appeared in random order.

On each story page, an instruction box repeated the question and allowed submitting an answer or returning to the search results. If the participant returned to the results, each result they previously clicked was grayed out and disabled.

In treatment rounds, the participant saw an interstitial warning the first time they clicked a target result. Our multi-armed bandit algorithm (Section 4.5) selected the warning design that the participant encountered. All warnings included two buttons: “Go back” or “Dismiss and continue.” If the user clicked “Go back” or the browser back button, they returned to the search results. The user did not encounter a warning when they clicked on any other result. If the user clicked “Dismiss and continue,” they were taken to the story page, where they could submit an answer or return to the search results.

When the participant submitted an answer, they were presented with a survey about why they chose particular search results. This survey was a misdirection to maintain the false premise that the experiment was studying search engine usage. In control rounds, submitting the survey led to the next round. In treatment rounds, the next page was a second survey, designed to capture whether the participant comprehended the purpose of the warning and whether the participant perceived a risk of harm (see Section 4.3). This survey also included an attention check question so that we could discard responses from participants who did not carefully read the instructions.

After completing all four rounds, participants were navigated to a final demographic survey, then compensated.

4.3 Measuring Participant Perceptions

After each treatment round, we presented a survey to measure whether the participant comprehended the warning or perceived a risk of harm. We developed the survey questions based on our laboratory results and small-scale pilot studies.

Informativeness We designed three survey questions to measure whether participants comprehended the purpose of a warning. Recall that in our laboratory study, participants who misunderstood a warning typically believed the warning was related to malware or another security threat. We grounded our informativeness questions in this observation, asking participants to indicate whether the warning was about three topics: malware (incorrect), information theft (incorrect), and disinformation (correct). The survey presented the following statements to participants and asked about agreement on a 5-point Likert scale (“Strongly disagree” to “Strongly agree”).

- $info_1$: The warning said that this site is trying to install malware.
- $info_2$: The warning said that this site is trying to steal my information.
- $info_3$: The warning said that this site contains false or misleading information.

We used the survey responses to compute an informativeness score $i_{p,w}$ in the range $[-2, 2]$, which captured participant p ’s certainty that warning w was about disinformation. $i_{p,w} = 2$ if p “strongly agreed” that w was about false or misleading information and “strongly disagreed” that w was about malware and stealing information. For each point deviation from these “correct” responses on the Likert scales, we reduced $i_{p,w}$ by 1, resulting in a lower score when the participant was uncertain in their answer or had an incorrect understanding of the warning. The scoring formula was:

$$i_{p,w} = \max(-2, info_3 - info_2 - info_1 - 1)$$

Harm We designed two survey questions to measure whether a warning caused a participant to perceive a risk of harm. The survey asked about agreement with the following statements, using the same 5-point Likert scale as above.

- $harm_1$: After seeing the warning, I believe the website may harm me if I visit it.
- $harm_2$: After seeing the warning, I believe the website may harm other people who visit it.

The two questions distinguish between personal harm (the first question) and societal harm (the second question). Recall that in our laboratory study, we identified personal harm as a possible mechanism of effect for disinformation warnings.

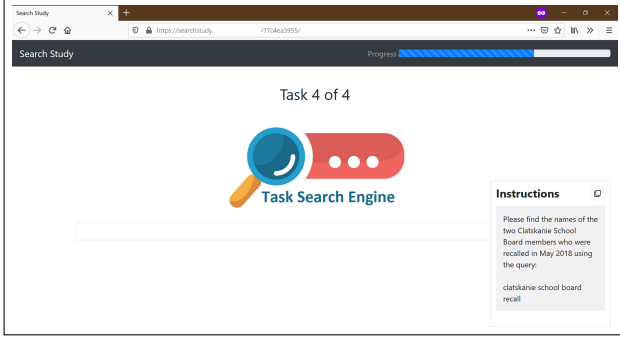


Figure 4: Each round of the crowdworker study began on this search page.

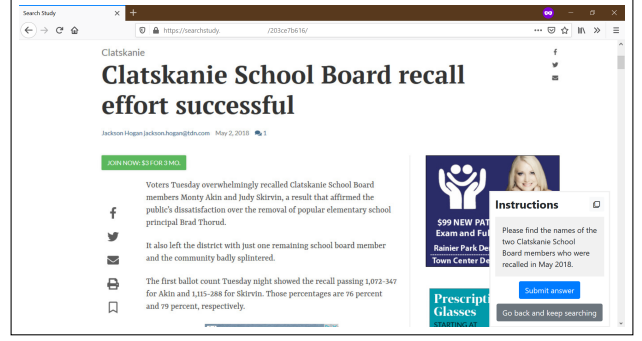


Figure 5: Clicking on a search result led the participant to a story page containing a screenshot of a real news webpage, instructions, and buttons to submit an answer or navigate back.

We found when piloting our crowdworker study that participants routinely conflated personal and societal harm when answering survey questions. We expressly asked about these two types of harm to ensure clarity for participants, and we solely used the personal harm response in our analysis. We computed the harm score $h_{p,w}$ for participant p and warning w by projecting their $harm_1$ response into the range $[-2, 2]$:

$$h_{p,w} = harm_1 - 3$$

4.4 Measuring Participant Behaviors

We measured the same behavioral outcomes as in the laboratory study (Section 3.4): clickthrough rate (CTR) and alternative visit rate (AVR). CTR represents the proportion of warning encounters where the participant clicked “Dismiss and continue.” AVR measures how often participants clicked on more than one source before submitting an answer. We recorded an alternative visit in a control round when the participant visited more than one story page, and we recorded an alternative visit in a treatment round when the participant visited a different story page after encountering a warning (regardless of whether the participant clicked through). Measuring AVR in control and treatment rounds enables us to estimate the warning’s effect with respect to a base rate.

4.5 Assigning Warnings

In order to answer our research questions about mechanisms of effect (RQ2 and RQ3), we measured for differences in behavioral effects between the warnings that achieved the highest and lowest mean scores for informativeness and harm.

The standard method for comparing effect sizes between treatments is a randomized controlled trial, in which participants are randomly assigned to treatments (often in equal numbers). The key variable determining how many observations are needed is the estimated difference in effect size between treatments. If this difference is small, the study will require a large number of observations for each treatment to achieve statistically significant confidence.

When designing our study, we observed that the difference in effect sizes between treatments could be small, meaning

that a large sample size could be necessary to evaluate our hypotheses. Our observations were expensive, however, and although we were testing 16 different conditions (8 warning treatments with 2 score outcomes each), we were only interested in comparing the effects of 4 conditions (with the top and bottom mean scores for each outcome). We therefore sought a method to efficiently assign participants to warnings so that the top- and bottom-scoring warnings achieved high confidence levels, but the other warnings (which we would not use in our hypothesis tests) would receive fewer observations and therefore consume fewer experimental resources.

In this study, we used an adaptive bandit algorithm to assign participants to warnings based on observations of previous participants. With each new observation, bandit algorithms update the probability of each condition in a study according to some *reward function* that aligns with the researchers’ scientific goals [106]. Bandits have been widely used in clinical trials, software design optimization, and political opinion research [107–109]. We discuss the full details of our multi-armed bandit implementation in supporting materials [100].

The reward function in our adaptive experiment preferred disinformation warnings that achieved high and low mean scores for informativeness and harm. For the first $n = 80$ participants, the algorithm assigned all warnings equally. For the remaining participants, the reward function prioritized warnings with the highest and lowest mean scores for informativeness and harm. As the bandit algorithm iterated, maximum and minimum scoring warnings emerged, and the algorithm improved our confidence in the mean scores for these warnings by prioritizing them for presentation to participants.

4.6 Participant Recruiting

We collected data from 250 Amazon Mechanical Turk workers who were based in the U.S. and had completed more than 5,000 jobs with an approval rate above 97%. We discarded data from 12 workers who failed an attention check question, leaving a sample population of 238. The population was roughly two-thirds male and over half of participants were between the ages of 30 and 49. The majority consumed

news media at least five days a week and paid somewhat close attention to politics and current events. We provide full population demographics in supporting materials [100].

Recruiting and consent materials described the study as related to search engine use and did not mention warnings or disinformation. We estimated the total task duration as 15-20 minutes and compensated participants \$2.33 with the opportunity to earn a \$1 bonus (43%) for retrieving the correct answer for all four queries.¹⁰ If a participant abandoned the task partway through or exceeded a 2-hour time limit, we discarded their data and recruited a replacement participant.

Our study was approved by the Princeton University IRB.

4.7 Results

We preregistered our analysis methods [100]. We computed mean ratings and 95% confidence intervals for informativeness and harm scores (Figure 6). For each political alignment and mechanism of effect, we identified the two warnings with the highest and lowest mean scores (Table 4).¹¹ We then conducted statistical tests comparing the AVR between these two warnings.¹² We treated clickthroughs and alternative visits as samples drawn from binomial distributions.¹³ For tests with large sample sizes, we used a z-test because a normal distribution approximates a binomial distribution; when the sample size was small, we used Fisher’s exact test.

Informativeness We found that *i3* had a very high mean informativeness score for liberal participants (1.41 on the scale $[-2, 2]$). *i2* and *i4* also had high, consistent informativeness scores for liberals. As for conservative participants, we found much lower mean informativeness scores for every warning we designed to be informative. *i4* had the highest mean score (0.88). The next most informative warning for conservatives was *h4*, which we had intended to convey a risk of harm.

The least informative warning was *h1*, which achieved consistent, extremely low informativeness scores from both liberals and conservatives. *h1* was the most extreme warning in the harm category; the only text it contained was “WARNING: This website is dangerous.”

Harm *h1* had the highest mean score for evoking fear of harm among liberals (1.18), with high confidence. We also found that *h1* had a high mean harm score for conservatives, but *h3* had a slightly higher score (1.15).

i4 had the lowest mean harm score for both political alignments, with a fairly low score among liberals (-0.76) and a more neutral score among conservatives (-0.2).

CTR The cumulative CTR across all treatments was 16%, which was noticeably lower than what we observed for in-

Table 4: We report alternative visit rates (AVR), clickthrough rates (CTR), and mean informativeness (\bar{i}) and harm (\bar{h}) scores with 95% confidence intervals.

	Liberal					Conservative				
	#	AVR	CTR	\bar{i}	\bar{h}	#	AVR	CTR	\bar{i}	\bar{h}
Control	318	20%	–	–	–	158	16%	–	–	–
Treatment	318	87%	16%	–	–	158	85%	17%	–	–
Selected treatments										
h1	120	85%	18%	-1.94 ± 0.06	1.18 ± 0.18	46	83%	17%	-1.91 ± 0.11	–
h3	73	84%	18%	–	–	27	81%	22%	–	1.15 ± 0.46
i3	39	87%	13%	1.41 ± 0.43	–	10	90%	10%	–	–
i4	17	82%	12%	–	-0.76 ± 0.69	25	76%	24%	0.88 ± 0.69	-0.2 ± 0.62

terstitial warnings in the laboratory (40%). No individual warning in this study demonstrated a higher CTR than the interstitial warning we tested in the laboratory.

AVR The AVR across all treatments was 86%, compared to 19% in control rounds. We used a one-sided z-test to evaluate if this difference was significant, and we found strong support for the hypothesis that the AVR in treatment rounds was greater than the AVR in control rounds ($p = 1.48e^{-111}$).

We tested whether there was a significant difference in AVR between the top- and bottom-scoring warnings for informativeness within the liberal and conservative groups. For liberals, we used a one-sided z-test and failed to reject the null hypothesis that the AVR of the top-scoring warning was less than or equal to the AVR of the bottom-scoring warning ($p = 0.27$). For conservatives, we used a one-tailed Fisher’s exact test (due to the small sample size of conservative participants) and also failed to reject the null hypothesis ($p = 0.54$).

Next, we tested for an AVR difference between the top- and bottom-scoring warnings for harm. We failed to reject both null hypotheses, with a one-sided z-test for liberals ($p = 0.41$) and a one-tailed Fisher’s test for conservatives ($p = 0.74$).

4.8 Discussion

We found that interstitial warnings have a strong effect on user behavior, confirming our laboratory study results (RQ1).

The results for warning informativeness were inconclusive (RQ2). We demonstrated that interstitial disinformation warnings can effectively inform users that a website may contain disinformation; we identified warning designs that scored well on average for informing participants. Conveying that a website may contain disinformation can prompt users to think critically about the website’s trustworthiness, and critical thinking is an important predictor of a user’s ability to correctly judge the accuracy of information [105, 110]. We did not, however, find evidence that informative warnings have a greater effect on user behavior than uninformative warnings.

The results for warnings conveying a risk of harm were similarly inconclusive (RQ3). We found that warning design can

¹⁰We expected that nearly all participants would qualify, and 81.5% did.

¹¹Our preregistered methods also included a non-overlapping 95% confidence interval criterion, but it did not affect our warning selection.

¹²We did not conduct statistical tests on warning CTRs, because our research questions focused on seeking alternative sources of information.

¹³We assumed independence between warning treatments for a participant.

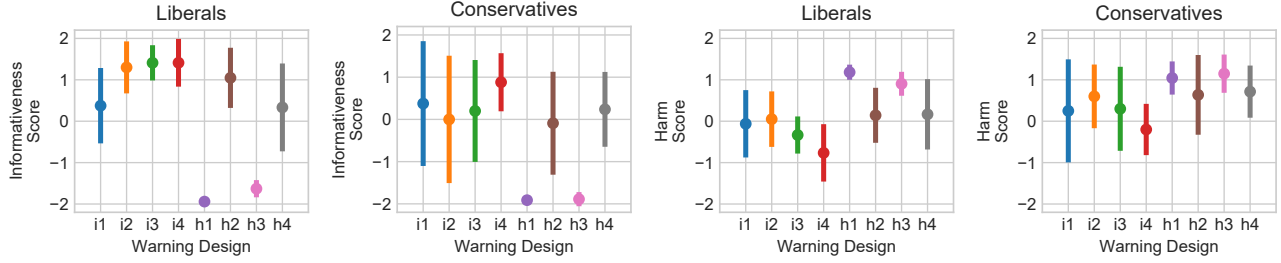


Figure 6: Mean informativeness and harm scores, with 95% confidence intervals, for liberals and conservatives.

effectively convey a risk of harm, but we did not find evidence that better conveying a risk of harm affects user behavior.

We hypothesize that the user experience friction introduced by interstitial warnings may be an important causal factor for changes in user behavior.¹⁴ We found evidence for friction as a mechanism of effect in our laboratory study (Section 3.7.1), but we did not test the theory in the crowdworker study.

Our results for RQ1 and RQ2 pose a possible speech dilemma: interstitial disinformation warnings can effectively inform users, but whether a user is informed may have little relation to how they behave in response to warnings.

Finally, we did not find evidence that partisanship moderates warning perceptions or behaviors (RQ4). Figure 6 shows that warning scores were generally similar for liberal and conservative participants, and Table 4 shows that CTRs and AVRs were also close between the groups.

Limitations There may be variables we did not measure that moderate the relationships between warning designs, participant perceptions, and behavior. Detailed qualitative methods, like in our laboratory study, can surface these variables—but are challenging to implement in a crowdworker study.

We also note that while our crowdworker sample was more diverse than our laboratory sample, neither sample was representative of the U.S. population. The behavioral effects we observed were fairly consistent across demographic groups, though. Our study population only included individuals located in the U.S.; cross-cultural research is needed to understand if the effects we observed apply globally.

5 Conclusion

In this section, we provide recommendations for future evaluation and deployment of disinformation warnings.

5.1 Directions for Future Research

Future work could explore the role of user experience friction in disinformation warnings. We found limited evidence that friction is an important factor in our laboratory study. The results from our crowdworker study also suggest that friction—rather than informativeness or conveying a risk of

harm—may be the predominant cause of warning behavioral effects. We did not test friction as a mechanism of effect in our crowdworker study, and our results do not conclusively rule out the mechanisms of effect we did examine. But our results are strongly suggestive, and friction merits further study.

Future work could also evaluate other types of interstitial warnings and interstitial warnings in other contexts, especially social media. Platforms are already deploying warning popups that a user must dismiss, as well as warning overlays that obscure content until the user clicks [111, 112].

Another promising direction is evaluating how interstitial warnings interact with factors known to impact warning adherence and receptivity to disinformation. These factors include repetition of warnings [10, 15, 18, 20, 103], user age and digital literacy [113, 114], user tendency toward cognitive reflection [110, 115, 116], repeated exposure to inaccurate information [19, 83], and whether that information aligns with user political preferences [114, 117, 118].

A final direction for future study is exploring possible unintended consequences of interstitial disinformation warnings. These warnings could create an implied truth effect [87], generally undermine trust in online content [83], cause concern about the warning provider, or lead to warning fatigue [20].

5.2 Informing Platform Disinformation Warnings

Interstitial warnings can be effective tools for countering disinformation. Compared to contextual warnings, interstitial designs are much more noticeable for users and much more capable of informing users about disinformation. Platforms that use contextual warnings for disinformation should be aware that their warnings may have minimal effects.

Going forward, platforms should follow evidence-based approaches for developing and deploying disinformation warnings. By conducting internal evaluations, collaborating with independent researchers, and releasing data, platforms can significantly advance their ability to counter disinformation with warnings—just like software vendors have done for over a decade to advance security warnings [10, 13–15, 17, 19].

Acknowledgments

We thank Marshini Chetty and Elissa Redmiles for valuable early feedback on this work. Simone Fischer-Hübner provided thoughtful shepherding for our paper.

¹⁴Another possible explanation is a substitution effect in our analysis. The uninformative warnings we examine, for example, could be effective at conveying risk of harm. Future work on the role of friction in disinformation warnings could also shed light on this issue.

References

- [1] Samantha Bradshaw and Philip N. Howard. *The Global Disinformation Order: 2019 Global Inventory of Organised Social Media Manipulation*. Tech. rep. University of Oxford, Sept. 26, 2018. URL: <https://comprop.oii.ox.ac.uk/wp-content/uploads/sites/93/2019/09/CyberTroop-Report19.pdf>.
- [2] Diego A. Martin, Jacob N. Shapiro, and Michelle Nedashkovskaya. “Recent Trends in Online Foreign Influence Efforts”. In: *Journal of Information Warfare (JIW)* 18 (3 2019). URL: <https://esoc.princeton.edu/publications/trends-online-influence-efforts>.
- [3] Adam Mosseri. *Addressing Hoaxes and Fake News*. Facebook Newsroom. Dec. 15, 2016. URL: <https://about.fb.com/news/2016/12/news-feed-fyi-addressing-hoaxes-and-fake-news/> (visited on 06/03/2020).
- [4] Justin Kosslyn and Cong Yu. *Fact Check Now Available in Google Search and News Around the World*. Google Official Blog. July 2017. URL: <https://www.blog.google/products/search/fact-check-now-available-google-search-and-news-around-world/> (visited on 03/05/2020).
- [5] Microsoft Bing. *Bing Adds Fact Check Label in SERP to Support the ClaimReview Markup*. Bing Blogs. Sept. 14, 2017. URL: <https://www.blog.google/products/search/fact-check-now-available-google-search-and-news-around-world/> (visited on 03/05/2020).
- [6] Davey Alba and Kate Conger. *Twitter Moves to Target Fake Videos and Photos*. New York Times. Feb. 2020. URL: <https://www.nytimes.com/2020/02/04/technology/twitter-fake-videos-photos-disinformation.html> (visited on 04/18/2020).
- [7] Taylor Hatmaker. *Twitter Adds a Warning Label Fact-Checking Trump’s False Voting Claims*. TechCrunch. May 26, 2020. URL: <https://techcrunch.com/2020/05/26/twitter-trump-labels-fact-checking-tweet/> (visited on 06/03/2020).
- [8] Min Wu, Robert C Miller, and Simson L Garfinkel. “Do Security Toolbars Actually Prevent Phishing Attacks?” In: *Proceedings of the 2006 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Apr. 2006. DOI: 10.1145/1124772.1124863.
- [9] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. “You’ve Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings”. In: *Proceedings of the 2008 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Apr. 2008. DOI: 10.1145/1357054.1357219.
- [10] Devdatta Akhawe and Adrienne Porter Felt. “Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness”. In: *Proceedings of the 22nd USENIX Security Symposium (USENIX Security)*. Aug. 2013. URL: <https://dl.acm.org/doi/10.5555/2534766.2534789>.
- [11] Joshua Sunshine et al. “Crying Wolf: An Empirical Study of SSL Warning Effectiveness”. In: *Proceedings of the 18th USENIX Security Symposium (USENIX Security)*. Aug. 2009. URL: <https://dl.acm.org/doi/abs/10.5555/1855768.1855793>.
- [12] Marian Harbach et al. “Sorry, I Don’t Get It: An Analysis of Warning Message Texts”. In: *Proceedings of the 17th International Conference on Financial Cryptography and Data Security (FC)*. Vol. 7859. Lecture Notes in Computer Science (LNCS). Apr. 2013. DOI: 10.1007/978-3-642-41320-9_7.
- [13] Adrienne Porter Felt et al. “Experimenting at Scale with Google Chrome’s SSL Warning”. In: *Proceedings of the 2014 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Apr. 2014. DOI: 10.1145/2556288.2557292.
- [14] Adrienne Porter Felt et al. “Improving SSL Warnings: Comprehension and Adherence”. In: *Proceedings of the 2015 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Apr. 2015. DOI: 10.1145/2702123.2702442.
- [15] Joel Weinberger and Adrienne Porter Felt. “A Week to Remember: The Impact of Browser Warning Storage Policies”. In: *Proceedings of the 12th Symposium On Usable Privacy and Security (SOUPS)*. June 2016. URL: <https://dl.acm.org/doi/abs/10.5555/3235895.3235898>.
- [16] Malkin, Nathan and Mathur, Arunesh and Harbach, Marian and Egelman, Serge. “Personalized Security Messaging: Nudges for Compliance with Browser Warnings”. In: *Proceedings of the EuroUSEC 2017 – The 2nd European Workshop on Usable Security*. Apr. 2017. DOI: 10.14722/eurousec.2017.23008.
- [17] Robert W. Reeder et al. “An Experience Sampling Study of User Reactions to Browser Warnings in the Field”. In: *Proceedings of the 2018 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Apr. 2018. DOI: 10.1145/3173574.3174086.
- [18] Serge Egelman and Stuart E. Schechter. “The Importance of Being Earnest (in Security Warnings)”. In: *Proceedings of the 17th International Conference on Financial Cryptography and Data Security (FC)*. Vol. 7859. Lecture Notes in Computer Science (LNCS). Apr. 2013. DOI: 10.1007/978-3-642-39884-1_5.
- [19] Hazim Almuhammedi et al. “Your Reputation Precedes You: History, Reputation, and the Chrome Malware Warning”. In: *Proceedings of the 10th Symposium On Usable Privacy and Security (SOUPS)*. June 2014. URL: <https://dl.acm.org/doi/10.5555/3235838.3235848>.
- [20] Cristian Bravo-Lillo et al. “Harder to Ignore? Revisiting Pop-Up Fatigue and Approaches to Prevent It”. In: *Proceedings of the 10th Symposium On Usable Privacy and Security (SOUPS)*. June 2014. URL: <https://dl.acm.org/doi/abs/10.5555/3235838.3235847>.
- [21] Mustafa Emre Acer et al. “Where the Wild Warnings Are: Root Causes of Chrome HTTPS Certificate Errors”. In: *Proceedings of the 24th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. Oct. 2017. DOI: 10.1145/3133956.3134007.
- [22] Michela Del Vicaro et al. “The Spreading of Misinformation Online”. In: *Proceedings of the National Academy of Sciences (PNAS)* 113 (3 Jan. 2016). DOI: 10.1073/pnas.1517441113.
- [23] Hunt Allcott and Matthew Gentzkow. “Social Media and Fake News in the 2016 Election”. In: *Journal of Economic Perspectives (JEP)* 31.2 (2017). DOI: 10.1257/jep.31.2.211.
- [24] Hunt Allcott, Matthew Gentzkow, and Chuan Yu. “Trends in the Diffusion of Misinformation on Social Media”. In: *Research and Politics* 6 (2 2019). DOI: 10.1177/2053168019848554.
- [25] Peter Burger et al. “The Reach of Commercially Motivated Junk News on Facebook”. In: *PLoS One* 14 (8 Aug. 2019). DOI: 10.1371/journal.pone.0220446.
- [26] Andrew M. Guess, Brendan Nyhan, and Jason Reifler. “Exposure to Untrustworthy Websites in the 2016 US Election”. In: *Nature Human Behavior* (Mar. 2020). DOI: 10.1038/s41562-020-0833-x.
- [27] Arkaitz Zubiaga et al. “Analysing How People Orient to and Spread Rumours in Social Media by Looking at Conversational Threads”. In: *PLoS One* 11 (3 Mar. 2016). DOI: 10.1371/journal.pone.0150989.
- [28] Philip N. Howard et al. *Social Media, News and Political Information During the US Election: Was Polarizing Content Concentrated in Swing States?* 2018. arXiv: 1802.03573 [cs.SI].

- [29] Chengcheng Shao et al. "The Spread of Low-Credibility Content by Social Bots". In: *Nature Communications* (Nov. 2018). DOI: 10.1038/s41467-018-06930-7.
- [30] Soroush Vosoughi, Deb Roy, and Sinan Aral. "The Spread of True and False News Online". In: *Science* 359.6380 (Mar. 9, 2018). DOI: 10.1126/science.aap9559.
- [31] Srijan Kumar, Robert West, and Jure Leskovec. "Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes". In: *Proceedings of the 25th World Wide Web Conference (WWW)*. Apr. 2016. DOI: 10.1145/2872427.2883085.
- [32] Daniel Bush and Alex Zaheer. *Bing's Top Search Results Contain an Alarming Amount of Disinformation*. Stanford Internet Observatory Blog. Dec. 27, 2019. URL: <https://cyber.fsi.stanford.edu/io/news/bing-search-disinformation> (visited on 05/15/2020).
- [33] Adam Fourney et al. "Geographic and Temporal Trends in Fake News Consumption During the 2016 US Presidential Election". In: *Proceedings of the 26th ACM International Conference on Information and Knowledge Management (CIKM)*. Nov. 2017. DOI: 10.1145/3132847.3133147.
- [34] Chengcheng Shao et al. "Anatomy of an Online Misinformation Network". In: *PLoS One* 13 (4 Apr. 2018). DOI: 10.1371/journal.pone.0196087.
- [35] Adam Badawy, Kristina Lerman, and Emilio Ferrara. "Who Falls for Online Political Manipulation?" In: *Companion Proceedings of the The Web Conference (WWW) 2019*. May 2019. DOI: 10.1145/3308560.3316494.
- [36] Alexandre Bovet and Hernán A. Makse. "Influence of Fake News in Twitter During the 2016 US Presidential Election". In: *Nature Communications* (Jan. 2019). DOI: 10.1038/s41467-018-07761-2.
- [37] Nir Grinberg et al. "Fake News on Twitter During the 2016 U.S. Presidential Election". In: *Science* 363.6425 (Jan. 25, 2019). DOI: 10.1126/science.aau2706.
- [38] Emilio Ferrara. "Disinformation and Social Bot Operations in the Run Up to the 2017 French Presidential Election". In: *First Monday* 22 (8 Aug. 2017). DOI: 10.2139/ssrn.2995809.
- [39] Michele Cantarella, Nicolò Fraccaroli, and Roberto Volpe. *Does Fake News Affect Voting Behaviour?* Tech. rep. DEMB Working Paper Series n.146, June 12, 2019. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3402913.
- [40] Kate Starbird et al. "Ecosystem or Echo-System? Exploring Content Sharing Across Alternative Media Domains". In: *AAAI-ICWSM2018*. URL: <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM18/paper/view/17836>.
- [41] Andreas Vlachos and Sebastian Riedel. "Identification and Verification of Simple Claims about Statistical Properties". In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Sept. 2015. URL: <https://www.aclweb.org/anthology/D15-1312.pdf>.
- [42] James Thorne and Andreas Vlachos. "An Extensible Framework for Verification of Numerical Claims". In: *Proceedings of the Software Demonstrations of the 15th Annual Meeting of the Association for Computational Linguistics (ACL)*. Apr. 2017. URL: <https://www.aclweb.org/anthology/E17-3010.pdf>.
- [43] Naemul Hassan et al. "Toward Automated Fact-Checking: Detecting Check-worthy Factual Claims by ClaimBuster". In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery in Data Mining (KDD)*. July 2017. DOI: 10.1145/3097983.3098131.
- [44] James Thorne and Andreas Vlachos. "Automated Fact Checking: Task Formulations, Methods and Future Directions". In: *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*. Aug. 2018. URL: <https://www.aclweb.org/anthology/C18-1283>.
- [45] James Thorne et al. "FEVER: A Large-Scale Dataset for Fact Extraction and VERification". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT)*. (Long paper). June 2018. DOI: 10.18653/v1/N18-1074.
- [46] Hannah Rashkin et al. "Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking". In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Sept. 2017. DOI: 10.18653/v1/D17-1317.
- [47] Martin Potthast et al. "A Stylometric Inquiry into Hyperpartisan and Fake News". In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*. (Long paper). July 2016. DOI: 10.18653/v1/P18-1022.
- [48] Benjamin D. Horne and Siebl Adali. "This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News". In: *Proceedings of the 11th International Conference on Weblogs and Social Media (ICWSM)*. May 2017. URL: <https://aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15772>.
- [49] Austin Hounsel et al. "Identifying Disinformation Websites Using Infrastructure Features". In: *10th USENIX Workshop on Free and Open Communications on the Internet (FOCI)*. Aug. 11, 2020. URL: <https://www.usenix.org/conference/foci20/presentation/hounsel>.
- [50] Fan Yang et al. "Automatic Detection of Rumor on Sina Weibo". In: *Proceedings of the 2012 ACM Workshop on Automated Decision Making for Active Cyber Defense (SafeConfig)*. Aug. 2012. DOI: 10.1145/2350190.2350203.
- [51] Ke Wu, Song Yang, and Kenny Q Zhu. "False Rumors Detection on Sina Weibo by Propagation Structures". In: *Proceedings of the 31st International Conference on Data Engineering (ICDE)*. Apr. 2015. DOI: 10.1109/ICDE.2015.7113322.
- [52] Zhe Zhao, Paul Resnick, and Qiaozhu Mei. "Enquiring Minds: Early Detection of Rumors in Social Media from Enquiry Posts". In: *Proceedings of the 24th World Wide Web Conference (WWW)*. May 2015. DOI: 10.1145/2736277.2741637.
- [53] Jing Ma, Wei Gao, and Kam-Fai Wong. "Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning". In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*. (Long paper). Aug. 2017. DOI: 10.18653/v1/P17-1066.
- [54] Ramy Baly et al. "Predicting Factuality of Reporting and Bias of News Media Sources". In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Oct. 2018. DOI: 10.18653/v1/D18-1389.
- [55] Kashyap Popat et al. "Where the Truth Lies: Explaining the Credibility of Emerging Claims on the Web and Social Media". In: *Companion Proceedings of the The Web Conference (WWW) 2017*. Apr. 2017. DOI: 10.1145/3041021.3055133.
- [56] Claire Wardle and Hossein Derakhshan. *Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making*. Tech. rep. Council of Europe, Sept. 27, 2017. URL: <https://rm.coe.int/information-disorder-toward-an-interdisciplinary-framework-for-research/168076277c>.
- [57] Samantha Bradshaw and Phillip N. Howard. *Why Does Junk News Spread So Quickly Across Social Media?* Tech. rep. Oxford Internet Institute, Jan. 28, 2018. URL: https://kf-site-production.s3.amazonaws.com/media_elements/files/000/000/142/original/Topos_KF_White-Paper_Howard_V1_ado.pdf.

- [58] Priyanjana Bengani. *Hundreds of 'Pink Slime' Local News Outlets are Distributing Algorithmic Stories and Conservative Talking Points*. Columbia Journalism Review. Dec. 18, 2019. URL: https://www.cjr.org/tow_center_reports/hundreds-of-pink-slime-local-news-outlets-are-distributing-algorithmic-stories-conservative-talking-points.php (visited on 02/15/2020).
- [59] Tara Calishain. *Look out for Junk Sources in Google News*. ResearchBuzz. Nov. 24, 2019. URL: <https://researchbuzz.me/2019/11/24/look-out-for-junk-sources-in-google-news/> (visited on 02/15/2020).
- [60] Alex Kasprak et al. *Hiding in Plain Sight: PAC-Connected Activists Set Up 'Local News' Outlets*. Snopes. Mar. 4, 2019. URL: <https://www.snopes.com/news/2019/03/04/activists-setup-local-news-sites/> (visited on 02/15/2020).
- [61] David M. J. Lazer et al. "The Science of Fake News". In: *Science* 359.6380 (Mar. 9, 2018). DOI: 10.1126/science.aao2998.
- [62] Leon Yin et al. *Your Friendly Neighborhood Troll: The Internet Research Agency's Use of Local and Fake News in the 2016 US Presidential Campaign*. Tech. rep. SMapP Data Report, Jan. 2018. URL: https://s18798.pcdn.co/smapp/wp-content/uploads/sites/1693/2018/11/SMaPP_Data_Report_2018_01_IRA_Links.pdf.
- [63] Stephan Lewandowsky et al. "Misinformation and Its Correction: Continued Influence and Successful Debiasing". In: *Psychological Science in the Public Interest (PSPI)* 13 (3 Dec. 2012). DOI: 10.1177/1529100612451018.
- [64] Edson C. Tandoc Jr., Zheng Wei Lim, and Richard Ling. "Defining 'Fake News'". In: *Digital Journalism* 6 (2 Feb. 2018). DOI: 10.1080/21670811.2017.1360143.
- [65] James H. Kuklinski et al. "Misinformation and the Currency of Democratic Citizenship". In: *The Journal of Politics (J Polit)* 62 (3 Aug. 2000). DOI: 10.1111/0022-3816.00033.
- [66] Rakoen Maertens, Frederik Anseel, and Sander van der Linden. "Combating Climate Change Misinformation: Evidence for Longevity of Inoculation and Consensus Messaging Effects". In: *Environmental Psychology* (June 12, 2020). (Pre-print). DOI: 10.1016/j.jenvp.2020.101455.
- [67] Christina Nemr and William Gangware. *Weapons of Mass Distraction: Foreign State-Sponsored Disinformation in the Digital Age*. Tech. rep. Park Advisors, Mar. 2019. URL: <https://www.park-advisors.com/disinfo-report>.
- [68] Edith Greene, Marlene S Flynn, and Elizabeth F Loftus. "Inducing Resistance to Misleading Information". In: *The European Journal of Psychology Applied to Legal Context* 7 (1 Jan. 2015). DOI: 10.1016/S0022-5371(82)90571-0.
- [69] Ullrich KH Ecker, Stephan Lewandowsky, and David TW Tang. "Explicit Warnings Reduce but do Not Eliminate the Continued Influence of Misinformation". In: *Memory & Cognition (MC)* 38 (8 Dec. 2010). DOI: 10.3758/MC.38.8.1087.
- [70] Goldman, Eric. *Content Moderation Strategies*. Presentation at Princeton University CITP Luncheon Series. Apr. 9, 2019.
- [71] Emily Taylor, Stacie Walsh, and Samantha Bradshaw. *Industry Responses to the Malicious Use of Social Media*. NATO Stratcom CoE. Oct. 9, 2018. URL: <https://www.stratcomcoe.org/industry-responses-malicious-use-social-media> (visited on 06/09/2020).
- [72] Masahiro Kimura, Kazumi Saito, and Hiroshi Motoda. "Minimizing the Spread of Contamination by Blocking Links in a Network". In: *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*. July 2008. DOI: 10.1145/1514888.1514892.
- [73] Ceren Budak, Divyakant Agrawal, and Amr El Abbadi. "Limiting the Spread of Misinformation in Social Networks". In: *Proceedings of the 20th World Wide Web Conference (WWW)*. Mar. 2011. DOI: 10.1145/1963405.1963499.
- [74] Nam P Nguyen et al. "Containment of Misinformation Spread in Online Social Networks". In: *Proceedings of the 4th ACM Conference on Web Science (WebSci)*. June 2012. DOI: 10.1145/2380718.2380746.
- [75] Senzhang Wang et al. "Negative Influence Minimizing by Blocking Nodes in Social Networks". In: *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*. July 2013. URL: <https://dl.acm.org/doi/abs/10.5555/2908286.2908331>.
- [76] Mehrdad Farajtabar et al. "Fake News Mitigation via Point Process Based Intervention". In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*. Aug. 2017. DOI: 10.5555/3305381.3305495.
- [77] Yi Li, Ruidong Yan, and Weili Wu. "Nonlinear Combinatorial Optimization". In: 2019. Chap. Optimization on Content Spread in Social Network Studies. DOI: 10.1007/978-3-030-16194-1_13.
- [78] Twitter. *Disclosing New Data to Our Archive of Information Operations*. Twitter Company Blog. Sept. 20, 2019. URL: https://blog.twitter.com/en_us/topics/company/2019/info-ops-disclosure-data-september-2019.html (visited on 02/07/2020).
- [79] Nathaniel Gleicher and Oscar Rodriguez. *Removing Additional Inauthentic Activity from Facebook*. Facebook Newsroom. Aug. 11, 2018. URL: <https://about.fb.com/news/2018/10/removing-inauthentic-activity/> (visited on 02/07/2020).
- [80] Nick Clegg. *Combating COVID-19 Misinformation Across Our Apps*. Facebook Newsroom. Mar. 25, 2020. URL: <https://about.fb.com/news/2020/03/combating-covid-19-misinformation/> (visited on 04/18/2020).
- [81] Jeff Smith, Grace Jackson, and Seetha Raj. *Designing Against Misinformation*. Facebook Design. Dec. 20, 2017. URL: <https://medium.com/facebook-design/designing-against-misinformation-e5846b3aale2> (visited on 02/07/2020).
- [82] Manish Singh. *Whatsapp Tests New Feature to Fight Misinformation: Search the Web*. TechCrunch. Mar. 21, 2020. URL: <https://techcrunch.com/2020/03/21/whatsapp-search-web-coronavirus/> (visited on 04/18/2020).
- [83] Gordon Pennycook, Tyrone D Cannon, and David G Rand. "Prior Exposure Increases Perceived Accuracy of Fake News". In: *The Journal of Experimental Psychology: General* 147 (12 Dec. 2018). DOI: 10.1037/xge0000465.
- [84] Björn Ross et al. "Fake News on Social Media: The (In)Effectiveness of Warning Messages". In: *ICIS 2018 Proceedings*. Dec. 2018. URL: <https://aisel.aisnet.org/icis2018/social/Presentations/16/>.
- [85] Mingkun Gao et al. "To Label or Not to Label: The Effect of Stance and Credibility Labels on Readers' Selection and Perception of News Articles". In: *Proceedings of the ACM on Human-Computer Interaction (PACM HCI)*. Nov. 2018. DOI: 10.1145/3274324.
- [86] Katherine Clayton et al. "Real Solutions for Fake News? Measuring the Effectiveness of General Warnings and Fact-Check Tags in Reducing Belief in False Stories on Social Media". In: *Political Behavior* (Feb. 2019). DOI: 10.1007/s11109-019-09533-0.
- [87] Gordon Pennycook et al. "The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings". In: *Management Science* (Feb. 21, 2020). DOI: 10.1287/mnsc.2019.3478.

- [88] Haeseung Seo, Aiping Xiong, and Dongwon Lee. “Trust It or Not: Effects of Machine-Learning Warnings in Helping Individuals Mitigate Misinformation”. In: *Proceedings of the 10th ACM Conference on Web Science (WebSci)*. June 2019. DOI: 10.1145/3292522.3326012.
- [89] Guy Rosen et al. *Helping to Protect the 2020 US Elections*. Facebook Newsroom. Aug. 21, 2019. URL: <https://about.fb.com/news/2019/10/update-on-election-integrity-efforts/> (visited on 04/18/2020).
- [90] Ben Collins. *Twitter Is Testing New Ways to Fight Misinformation — Including a Community-Based Points System*. NBC News. Feb. 20, 2020. URL: <https://www.nbcnews.com/tech/tech-news/twitter-testing-new-ways-fight-misinformation-including-community-based-points-n1139931> (visited on 04/18/2020).
- [91] Paul Mena. “Cleaning Up Social Media: The Effect of Warning Labels on Likelihood of Sharing False News on Facebook”. In: *Policy & Internet (P&I)* 12 (2 July 28, 2019). DOI: 10.1002/poi3.214.
- [92] Patricia L. Moravec, Antino Kim, and Alan R. Dennis. “Appealing to Sense and Sensibility: System 1 and System 2 Interventions for Fake News on Social Media”. In: *Information Systems Research (ISR)* 31 (3 Aug. 13, 2020). DOI: 10.1287/isre.2020.0927.
- [93] Leticia Bode and Emily K Vraga. “In Related News, That Was Wrong: The Correction of Misinformation Through Related Stories Functionality in Social Media”. In: *Journal of Communication* 65 (4 Aug. 2015). DOI: 10.1111/jcom.12166.
- [94] R. Kelly Garrett and Brian E. Weeks. “The Promise and Peril of Real-Time Corrections to Political Misperceptions”. In: *Proceedings of the 2013 ACM Conference on Computer Supported Cooperative Work (CSCW)*. Feb. 2013. DOI: 10.1145/2441776.2441895.
- [95] Stuart E Schechter et al. “The Emperor’s New Security Indicators”. In: *Proceedings of the 28th IEEE Symposium on Security & Privacy (S&P)*. May 2007. DOI: 10.1109/SP.2007.35.
- [96] Andreas Sotirakopoulos, Kirsie Hawkey, and Konstantin Beznosov. “On the Challenges in Usable Security Lab Studies: Lessons Learned from Replicating a Study on SSL Warnings”. In: *Proceedings of the 7th Symposium On Usable Privacy and Security (SOUPS)*. June 2011. DOI: 10.1145/2078827.2078831.
- [97] Lorrie F Cranor. “A Framework for Reasoning About the Human in the Loop”. In: *Proceedings of the 1st Conference on Usability, Psychology, and Security (UPSEC)*. Apr. 14, 2008. URL: <https://dl.acm.org/doi/10.5555/1387649.1387650>.
- [98] Vincent C. Conzola and Michael S. Wogalter. “A Communication–Human Information Processing (C–HIP) Approach to Warning Effectiveness in the Workplace”. In: *Journal of Risk Research* 4 (4 Apr. 15, 2011). DOI: 10.1080/13669870110062712.
- [99] Adrienne Porter Felt et al. “Rethinking Connection Security Indicators”. In: *Proceedings of the 12th Symposium On Usable Privacy and Security (SOUPS)*. June 2016. URL: <https://dl.acm.org/doi/10.5555/3235895.3235897>.
- [100] Ben Kaiser et al. *Adapting Security Warnings to Counter Online Disinformation*. Supporting materials. URL: <https://osf.io/qf8e5/>.
- [101] Michael S. Wogalter. “Handbook of Warnings”. In: Jan. 20, 2006. Chap. Purposes and Scope of Warnings. DOI: 10.1201/9781482289688.
- [102] David W. Stewart and Ingrid M. Martin. “Intended and Unintended Consequences of Warning Messages: A Review and Synthesis of Empirical Research”. In: *Journal of Public Policy & Marketing (JPP&M)* 13 (1 Mar. 1994). DOI: 10.1177/074391569401300101.
- [103] Anthony Vance et al. “Tuning Out Security Warnings: A Longitudinal Examination of Habituation Through fMRI, Eye Tracking, and Field Experiments”. In: *Management Information Systems Quarterly (MISQ)* 42 (2 June 2018). DOI: 10.25300/MISQ/2018/14124.
- [104] Dan M. Kahan. *Ideology, Motivated Reasoning, and Cognitive Reflection: An Experimental Study*. Tech. rep. Yale Law School Cultural Cognition Project, Nov. 29, 2012. DOI: 10.2139/ssrn.2182588.
- [105] Gordon Pennycook and David G. Rand. “Fighting Misinformation on Social Media Using Crowdsourced Judgments of News Source Quality”. In: *Proceedings of the National Academy of Sciences (PNAS)* 116 (7 Feb. 2019). DOI: 10.1073/pnas.1806781116.
- [106] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. July 31, 2020. URL: <https://tor-lattimore.com/downloads/book/book.pdf>.
- [107] Feifang Hu and William F. Rosenberger. *The Theory of Response-Adaptive Randomization in Clinical Trials*. Apr. 10, 2006. DOI: 10.1002/047005588X.
- [108] John White. *Bandit Algorithms for Website Optimization*. Jan. 23, 2013.
- [109] Molly Offer-Westort, Alexander Coppock, and Donald P. Green. *Adaptive Experimental Design: Prospects and Applications in Political Science*. Tech. rep. Feb. 12, 2019. DOI: 10.2139/ssrn.3364402.
- [110] Gordon Pennycook and David G Rand. “Lazy, not Biased: Susceptibility to Partisan Fake News is Better Explained by Lack of Reasoning Than by Motivated Reasoning”. In: *Cognition* 188 (July 2019). DOI: 10.1016/j.cognition.2018.06.011.
- [111] Instagram. *Combatting Misinformation on Instagram*. Instagram Announcements. Dec. 16, 2019. URL: <https://about.instagram.com/blog/announcements/combatting-misinformation-on-instagram/> (visited on 05/13/2020).
- [112] Yoel Roth and Nick Pickles. *Updating our Approach to Misleading Information*. Twitter Product Blog. May 11, 2020. URL: https://blog.twitter.com/en_us/topics/product/2020/updating-our-approach-to-misleading-information.html (visited on 05/13/2020).
- [113] Kevin Munger et al. “The (Null) Effects of Clickbait Headlines on Polarization, Trust, and Learning”. In: *Public Opinion Quarterly* (Apr. 30, 2020). DOI: 10.1093/poq/nfaa008.
- [114] Andrew Guess, Jonathan Nagler, and Joshua Tucker. “Less than You Think: Prevalence and Predictors of Fake News Dissemination on Facebook”. In: *Science Advances* 5.1 (2019). DOI: 10.1126/sciadv.aau4586.
- [115] Ben M Tappin, Gordon Pennycook, and David G Rand. *Rethinking the Link Between Cognitive Sophistication and Politically Motivated Reasoning*. PsyArXiv. 2018. DOI: 10.31234/osf.io/yuzfj.
- [116] Robert M Ross, David G Rand, and Gordon Pennycook. *Beyond “Fake News”: The Role of Analytic Thinking in the Detection of Inaccuracy and Partisan Bias in News Headlines*. PsyArXiv. 2019. DOI: 10.31234/osf.io/cgsx6.
- [117] Brian F. Schaffner and Cameron Roche. “Misinformation and Motivated Reasoning: Responses to Economic News in a Politicized Environment”. In: *Public Opinion Quarterly* 81 (1 Mar. 1, 2017). DOI: 10.1093/poq/nfw043.
- [118] Dan M. Kahan. *Misconceptions, Misinformation, and the Logic of Identity-Protective Cognition*. Tech. rep. Yale Law School Cultural Cognition Project, May 24, 2017. DOI: 10.2139/ssrn.2973067.