CENTER FOR INFORMATION TECHNOLOGY POLICY
PRINCETON UNIVERSITY

THE UNIVERSITY OF CHICAGO 1890

June 22, 2020

Before the
Federal Trade Commission
Washington, D.C.

Endorsement Guides, P204500

Thank you for the opportunity to provide comments on the effectiveness of the FTC's Endorsement Guidelines (Guides). We write to respond to the questions asked by the FTC in the Request for Comment (Request) in connection with its systematic review of the Guides.

We are academic researchers associated with the Center for Information Technology Policy (CITP) at Princeton University and the Department of Computer Science at the University of Chicago. We draw on our collective experience in computer science and law to encourage the Commission to take specific steps to protect consumers and to conduct further studies about how online disclosures are used across a range of platforms. In particular, our research about user perceptions of online endorsements highlights the importance of influencers being fully transparent about brand relationships. We look forward to further opportunities to engage with the Commission's staff to answer any questions.

1. **Empirical research shows that influencers are not following the Guides.**

Our research speaks directly to Questions 3, 4 and 6 in the Request that seek comment on the prevalence of undisclosed endorsements online and the effectiveness of the Guides. In 2018, we conducted a large-scale empirical study of affiliate marketing on YouTube and Pinterest.[1] We gathered a dataset of randomly sampled YouTube videos (0.5 million) and Pinterest pins (2 million) and examined whether they contained affiliate marketing links. We found that just under 1% of the content we collected had

[1] Mathur, A., Narayanan, A., and Chetty, M. (2018) Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest. CSCW 2018.

such links. We then examined whether the links were disclosed to consumers. We found that about 11% of the YouTube content that had links made disclosures. On Pinterest, about 7% of that content had accompanying disclosures. This suggests that many influencers are not abiding by the Guides request to clearly and conspicuously disclose affiliate marketing content to encourage truth in advertising.

Even in the minority of instances when content was disclosed as containing affiliate marketing, those disclosures did not follow the Guides best practices on how to inform consumers. Our study divided the disclosures into three categories: "Affiliate Link," "Channel Support," and "Explanation" type disclosures. The Affiliate Link disclosures were typically the simplest form of disclosure. These included minimal statements such as "Affiliate links may be present above" or #affiliatelink. In the online supplement to the Guides ("What People Are Asking") the FTC correctly discourages the use of this kind of disclosure because "consumers might not understand that `affiliate link' means that the person placing the link is getting paid for purchases through the link."[2]

Our study found that on YouTube the majority of videos (~70%) that had some form of affiliate marketing disclosure, relied on minimal Affiliate Link disclosures to identify such content. Similarly, on Pinterest, the majority of pins (~65%) with affiliate marketing disclosures relied on this disclosure technique. Content on YouTube also used Channel Support disclosures. These disclosures explain that influencers receive financial support for their channel via the affiliate links e.g., "Shop using these links to support the channel." We found that ~20% of videos with disclosures used this technique.

Finally, a minority of influencers who disclosed affiliate marketing (~16% of YouTube videos and ~35% of Pinterest pins) used Explanation disclosures to be more direct about the financial link between influencers and brands they promoted. These disclosures used language such as "This video contains affiliate links, which means that if you click on one of the links, I'll receive a small commission." In the online supplement to the Guides, the FTC explicitly recommends using an Explanation disclosure for affiliate links to make a relationship to an advertiser clear.

[2] https://www.ftc.gov/tips-advice/business-center/guidance/ftcs-endorsement-guides-what-people-are-asking

Our research shows that there is a wide variation in how users comprehend these different disclosures. For example, we found that most users do not notice "affiliate link" disclosures. We surveyed a sample of 1791 users and found that they did not notice the Affiliate link disclosures regardless of where they were placed on the site. When we specifically alerted users to the disclosures (by calling it to their attention and not requiring them to click through), we still found that most of them did not understand that the Affiliate Link disclosure conveyed that the influencer was paid to promote the product.

By contrast, users were able to notice Explanation disclosures without intervention from the experiment. We also observed that Explanation disclosures were generally more effective in communicating to users. On Pinterest, this type of disclosure increased users' ability of identifying the underlying advertisement by 2.3–2.8 times as compared to the Affiliate Link type disclosure. It may be that Explanation disclosures were more effective on Pinterest because of the platform design: Pinterest has an easy to find pin description that is available to the user with no additional clicks, meaning a disclosure might be more easily noticed and read. On YouTube, users have to click "Show More" to read a video description and then find the corresponding disclosure if present. This result suggests that norms and design of the user interface in social media platforms might influence the efficacy of disclosures. Further research is required to investigate how user interface design affects disclosure efficacy.

Based on our research, we recommend that the FTC conduct further studies about how influencers disclose affiliate marketing online and about how users comprehend those disclosures. Aside from the wording of disclosures, we have observed that visual indicators play a major role in aiding user comprehension. A careful evaluation of these factors can help aid the FTC in developing more effective guidelines.

2. **Standardized disclosures may be more effective tools to educate consumers.**

In response to Question 7 -- how disclosures can be improved -- we have some specific suggestions based on our research. Specifically, in a separate study, we analyzed whether it is possible to automatically detect affiliate marketing links and alert users to their presence in a standardized manner. To do this research, we built a browser extension called AdIntuition (released for Chrome[3] and Firefox[4]) to automatically detect affiliate marketing links in YouTube videos and flag these videos

[3] https://chrome.google.com/webstore/detail/adintuition/pjpiddgaambjenhikcpbcbgjckidgpce?hl=en
[4] https://addons.mozilla.org/en-US/firefox/addon/adintuition/

using a banner across the video. Our preliminary findings are that it is possible to detect affiliate marketing links automatically, but that there is more advertisers and influencers can do to encourage greater transparency. We also found that standardized disclosures about affiliate marketing links could be helpful in more effectively communicating with users.

We studied 472 AdIntuition users (Chrome and Firefox) over 82 days, we found that the median user watched between 2 to 3 videos a day, and that 7.5% of the videos watched had affiliate marketing content.[5] We compiled a data set of 60,000 videos that the study's users watched and then examined who were the 15 most commonly observed influencers. In the majority of cases, a single user was responsible for watching all videos in the data set from one of these influencers. This suggests that certain users may disproportionately encounter affiliate marketing content based on the type of videos they consume.

This finding is supported by data from our 2018 study where we saw that affiliate marketing links were most commonly found across content within technology, fashion, and beauty related categories on YouTube. In these cases, it is possible that some users may be more attuned to and more able to recognize affiliate marketing than others. But, users may be more susceptible to undisclosed affiliate marketing from unfamiliar influencers when, for example, they are searching for product reviews prior to purchasing a product. The Guides could make targeted suggestions for influencers in common categories such as fashion or technology where affiliate marketing content is more pervasive. Further research is warranted to ascertain best practice disclosures for various categories of sellers.

In our 2020 study, we conducted a preliminary evaluation of the automated AdIntution disclosures about affiliate marketing content with 300 Amazon Mechanical Turk users. We observed that 71% of those who saw the automatic disclosure banner were able to identify an advertising relationship in an affiliate marketing video as compared to 56% of users who did not see an automatic disclosure banner.

We subsequently conducted a qualitative study of 11 users over the period of 2 weeks where users installed AdIntuition and then logged details in an electronic journal about every video they encountered with an AdIntuition advertising disclosure. Users also participated in pre- and post-study interviews. Although this is a very small initial sample of users, it still yielded preliminary data on how users feel when they encounter

---

[5] Swart, M., Lopez, Y., Mathur, A., and Chetty, M. (2020) Is This An Ad?: Automatically Disclosing Online Endorsements On YouTube With AdIntuition. CHI 2020.

automatic disclosures. In their electronic journals, a number of users reported being surprised that the influencer was doing affiliate marketing when they saw an automatic disclosure. Often their surprise was because the product did not align with the type of video or influencer. Others were not surprised to see an automatic disclosure and usually this was because they expected the influencer to be promoting merchandise. We also observed that while disclosures raised awareness of affiliate marketing with many users, they also made some users have negative reactions to an influencer. These negative sentiments included a decrease in trust, questioning the influencer's true intentions behind creating specific content, and feeling disappointed in the influencer. As for the products being promoted, a few participants expressed an increase in skepticism of the legitimacy of the endorsements and true quality of the product in videos that had the AdIntuition banner.

Our research suggests that many advertising practices commonly encountered online are opaque to consumers and difficult for third parties to detect systematically (e.g., paid reviews to promote products). We suggest that in addition to recommendations directed at influencers, the Guides also develop recommendations for advertisers to disclose who are the influencers they have paid for a review. Such disclosures would also aid with automatic detection and disclosure. For instance, affiliate marketing companies could disclose all the influencers with which they have a relationship on their websites. This listing would provide a source of ground truth for whether an influencer has a brand-relationship and allow third parties to verify compliance.

To implement this proposal in practice, there could be a threshold amount should be set for which companies need to automatically disclose advertising relationships. Having this information be available in a centralized place for each affiliate marketing company would enable developers to build tools like AdIntuition or even for browsers to integrate this automatic detection and disclosure functionality across multiple platforms. This would also give consumers multiple ways to access the influencer's disclosure of affiliate marketing links.

We also recommend that the FTC encourage affiliate marketing companies to provide influencers with standardized disclosures to increase consumer awareness of financial ties and minimize the burden on influencers. In the same vein, we recommend that the FTC encourage platforms to provide standardized disclosures that are easy for influencers to add. Indeed, where possible, platforms could automatically detect advertising relationships and encourage influencers to include a standardised disclosure. This would more efficiently spread the task of disclosing advertising relationships in a way that users can readily comprehend among advertisers, platforms, and influencers.

The Guides also need to be clear about which disclosures are ineffective and which ones epitomize best practice disclosures. In addition, the Guides should provide template disclosures for a variety of platforms. The Guides are currently vague in this regard. Providing template wording would make it consistent for users and enable them to more easily recognize what the disclosure means. In addition, this would more easily allow influencers to be compliant. Coupled with the current text descriptions, the Guides should include visual examples of effective disclosures on a variety of platforms.

* * *

As brands rely on influencer marketing, the FTC Guides can serve as an important source of protection for consumers. Our research shows that the FTC can aid consumers by providing clear, unambiguous suggestions for how to disclose brand relationships. We are available to assist the FTC towards the goal of helping users better identify and comprehend advertising disclosures and to encourage influencers, advertisers and platforms to better disclose advertising relationships.

Respectfully submitted,

Marshini Chetty*
Assistant Professor, Department of Computer Science, University of Chicago

Mihir Kshirsagar*
Technology Policy Clinic Lead, Center for Information Technology Policy, Princeton University

Arunesh Mathur
Graduate Student, Department of Computer Science, Princeton University

Jonathan Mayer
Assistant Professor of Computer Science and Public Affairs, Princeton University

Arvind Narayanan
Associate Professor, Department of Computer Science, Princeton University

*denotes principal comment authors

Contact:

Website: https://citp.princeton.edu
Phone: 609-258-5306
Email: mihir@princeton.edu

Website: https://www.cs.uchicago.edu
Phone: 773-702-2323
Email: marshini@uchicago.edu

Appendix A

Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest

ARUNESH MATHUR, Princeton University, USA
ARVIND NARAYANAN, Princeton University, USA
MARSHINI CHETTY, Princeton University, USA

Online advertisements that masquerade as non-advertising content pose numerous risks to users. Such hidden advertisements appear on social media platforms when content creators or "influencers" endorse products and brands in their content. While the Federal Trade Commission (FTC) requires content creators to disclose their endorsements in order to prevent deception and harm to users, we do not know whether and how content creators comply with the FTC's guidelines. In this paper, we studied disclosures within affiliate marketing, an endorsement-based advertising strategy used by social media content creators. We examined whether content creators follow the FTC's disclosure guidelines, how they word the disclosures, and whether these disclosures help users identify affiliate marketing content as advertisements. To do so, we first measured the prevalence of and identified the types of disclosures in over 500,000 YouTube videos and 2.1 million Pinterest pins. We then conducted a user study with 1,791 participants to test the efficacy of these disclosures. Our findings reveal that only about 10% of affiliate marketing content on both platforms contains any disclosures at all. Further, users fail to understand shorter, non-explanatory disclosures. Based on our findings, we make various design and policy suggestions to help improve advertising disclosure practices on social media platforms.

CCS Concepts: • **Information systems** → **Sponsored search advertising**; **Social advertising**; *Social networks*; • **Human-centered computing** → **Empirical studies in HCI**; • **Social and professional topics** → *Governmental regulations*;

Additional Key Words and Phrases: Online Advertising; Advertising Transparency; FTC Endorsement Guidelines; Social Media; Sponsored Content

ACM Reference Format:
Arunesh Mathur, Arvind Narayanan, and Marshini Chetty. 2018. Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest. In *Proceedings of the ACM on Human-Computer Interaction*, Vol. 2, CSCW, Article 119 (November 2018). ACM, New York, NY. 26 pages. https://doi.org/10.1145/3274388

1 INTRODUCTION

Endorsement-based advertising is one of many advertising strategies that allows Internet content creators—sometimes called *influencers* in marketing discourse—to monetize their content on social media platforms and blogs. Because such advertising often appears in conjunction—and is merged—with content creators' non-advertising content, Internet users encountering these advertisements may not recognize them as such, and may be misled or deceived [48].

Authors' addresses: Arunesh Mathur, Princeton University, 35 Olden Street, Princeton, NJ, 08540, USA, amathur@cs.princeton.edu; Arvind Narayanan, Princeton University, 308 Sherrerd Hall, Princeton, NJ, 08540, USA, arvindn@cs.princeton.edu; Marshini Chetty, Princeton University, 35 Olden Street, Princeton, NJ, 08540, USA, marshini@princeton.edu.

2573-0142/2018/11-ART119 $15.00
https://doi.org/10.1145/3274388

To avoid any such deception and harm to users, several guidelines and regulations around the world—including from Federal Trade Commission (FTC) [18] in the United States (US)—require content creators to disclose and describe their relationships with advertisers to users. Having adequate disclosures can better inform users when a piece of content is an advertisement, and ensure that users appropriately weigh up content creators' endorsements. In a recent case, the FTC blocked an endorsement-based advertisement where bloggers were misleading consumers by not disclosing hidden charges for products [9]. The FTC has also warned numerous other content creators—including social media celebrities—who failed to disclose their relationships with brands and companies in their social media posts [1, 13, 16, 17].

Unlike traditional online advertising, where disclosures are communicated using standardized icons (e.g., AdChoices [2] for Online Behavioral Advertising or OBA) or text (e.g., Twitter's *promoted* tweets [41]) endorsement-based advertising disclosures are relatively unstructured, open-ended in nature, and written primarily by individual content creators. As a result, we know little about how many content creators actually disclose their relationship with advertisers in the first place, and whether users notice and understand the disclosures' underlying message.

In this paper, we examine disclosures accompanying *affiliate marketing*—an endorsement-based advertising strategy that earns content creators money when users click on their customized URLs—on social media platforms. We tackle three primary research questions:

- First, how prevalent are disclosures in affiliate marketing content on social media platforms?
- Second, how do content creators word and frame disclosures in affiliate marketing content?
- Third, and finally, how effective are these current forms of affiliate marketing disclosures from a user standpoint?

Through our study, we aim to evaluate whether content creators on social media platforms follow affiliate marketing disclosure guidelines set by various regulatory bodies such as the FTC, whether their current disclosures effectively inform users about the underlying endorsement-based advertisement, and if not, determine what steps various stakeholders in the affiliate marketing industry, in policy and in design can take to improve disclosure practices.

To answer our research questions, we first measured the prevalence of affiliate marketing content on two popular social media platforms that serve and enable user-generated content: YouTube and Pinterest. We sampled ~500,000 unique YouTube videos and ~2.1 million unique Pinterest pins, filtered content that contained affiliate URLs, and then examined this content for disclosures. Following this analysis, we conducted a user study with 1,791 Amazon Mechanical Turk (MTurk) users to determine how effective the disclosures we discovered are in practice.

We have four main findings:

- First, we find that only ~10% of all affiliate marketing content on both YouTube and Pinterest contains accompanying disclosures.
- Second, we find that the affiliate marketing disclosures on these social media platforms can be broadly grouped into three distinct formats, which we term *Affiliate Link* disclosures, *Explanation* disclosures, and *Channel Support* disclosures. We find that *Explanation* disclosures—which the FTC explicitly advocates for using in its guidelines—occur least frequently: only 1.82% of all affiliate YouTube videos and 2.43% of all affiliate Pinterest pins contain *Explanation* disclosures.
- Third, we find that users are able to—by themselves—notice and understand *Explanation* disclosures but not the other types. Disclosures of this type decreased users' perceptions of the content creator favoring the endorsed product by 0.46–0.53 times, and increased users' ability of identifying the underlying advertisement by 2.3–2.8 times. However, we only observed

this effect on Pinterest, suggesting that norms and design of social media platforms might influence the efficacy of these disclosures.
- Fourth, and finally, we find that *Affiliate Link* disclosures are only half as effective as *Explanation* disclosures in communicating the underlying advertisement to users *even when these disclosures are presented to users without them having to look for it.*

We make the following contributions through our study:

- To the best of our knowledge, we are the first to empirically measure the prevalence of affiliate marketing on social media platforms, and the compliance of content creators' accompanying disclosures in affiliate marketing content with the FTC's endorsement guidelines.
- We describe a method to detect affiliate marketing content that is broadly applicable to other social media platforms and the Web. As a result of our analysis, we compiled the most comprehensive publicly available list of affiliate marketing companies and their URLs. We make this list available in the Appendix.
- We describe a method to retrieve affiliate marketing disclosures from text data. Our method can be easily extended to retrieve disclosures accompanying other forms of endorsement-based marketing techniques such as sponsored content.
- We provide new evidence that the wording and framing of disclosures impact user understanding of affiliate marketing advertising on social media platforms.

Based on our findings, we make design and policy suggestions aimed at various stakeholders such as social media platforms and affiliate marketing companies to help enable content creators to easily and clearly disclose advertising relationships. We also outline directions for future work in this area to help users notice and understand affiliate marketing disclosures, and more generally, help enable improved advertising disclosure practices on social media.

2 RELATED WORK

In this section, we describe endorsement-based advertising practices including relevant research on advertising disclosures, and place our own study in context.

2.1 Endorsement-based Advertising on Social Media Platforms

Native advertising is a form of advertising where advertisements take the form, and structure of non-advertising content, making it harder for users to identify it as such [3, 27, 29, 46]. For example, an advertisement might masquerade as a news article along with legitimate news articles on a news website. Endorsement-based advertising on social media platforms is a recent trend within native advertising, where the advertisement is created by a user—termed the content creator—of the platform to generate revenue for themselves rather than the platform itself. Broadly, such advertising manifests in one the following three ways [47]:

- *Sponsored Content*, in which a content creator partners with an advertiser or merchant to promote a product
- *Affiliate Marketing*, in which a content creator posts affiliate URLs to merchants in their content to earn money from the resulting sales
- *Product Giveaways*, in which a content creator receives product samples from an advertiser or merchant to promote and review

In this paper, we consider affiliate marketing advertising on social media platforms in depth. We leave the other advertising strategies—*Sponsored Content* and *Product Giveaways*—for future work.

Affiliate Marketing. Affiliate marketing primarily involves three entities: a content creator, a merchant and an affiliate marketing company. It comprises of two relationships: one between the

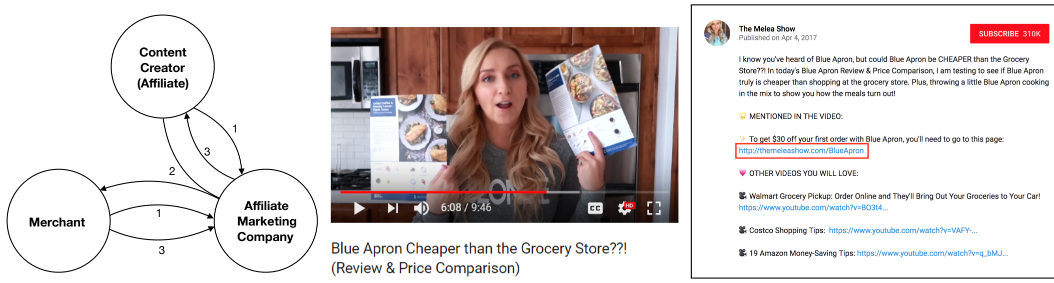

(a) (b) (c)

Fig. 1. Overview of Affiliate Marketing and Its Use on YouTube. (a) Affiliate Marketing Ecosystem. (b) YouTube Content Creator Endorsing a Product. (c) Affiliate URL in the Description of the Video (Red Box).

content creator and the affiliate marketing company, and another between the affiliate marketing company and the merchant. Illustrated in Figure 1a, affiliate marketing typically works as follows:

- Merchants and content creators sign up with an affiliate marketing company (1)
- Content creators drive sales to the merchant through the affiliate marketing company; the sales are tracked by means of customized URLs—called *affiliate URLs*—that are published by the affiliate marketing company for the content creator to distribute (2)
- Each time a sale is made using the affiliate URLs, the merchant pays the content creator a cut of the sale through the affiliate marketing company (3)

Figure 1b and Figure 1c illustrate how this works in practice on YouTube. A YouTube content creator endorses a specific brand product (Figure 1b) in their video, and includes an affiliate URL in the video's description (Figure 1c). Any purchases through this URL provides a share of the sales—a commission—to the YouTube content creator.

2.2 Advertising Disclosures

In the US, the FTC holds advertisers—regardless of the medium where they place their advertisements (e.g., radio, television, print media, or the Internet)—to the *truth-in-advertising* standard, meaning that the advertisements they publish should not be misleading in nature. The FTC requires that if additional information is required to prevent an advertisement from being misleading, then that information must be disclosed to users by adhering to the *clear and conspicuous* standard [10]; it investigates any unfair and deceptive practices under Section 5 of the FTC Act [15].

The FTC categorizes an advertisement as misleading or deceptive if it conceals its commercial nature from consumers [11]. Calling these advertisements—e.g., native advertisements, endorsements, sponsored content—*Deceptively Formatted Advertisements*, the FTC requires that they be disclosed to consumers as advertisements so that consumers can identify them as such.

In the following sub-sections, we touch upon the landscape of such advertising disclosures.

2.2.1 Native Advertising Disclosures. Given that users find it hard to identify and distinguish native advertisements (e.g., search engine advertisements, advertorials) from non-advertising content [3, 27, 29, 46], several studies have examined ways in which the disclosures present on these native advertisements can be designed for better recognition by users. These factors include changing the position, visual appearance and prominence of the disclosure, and the branding of the advertisement content itself. For instance, in a series of studies, Wojdynski *et al.* [44–46] found that greater visual prominence of the disclosure such as by increasing text size, color and contrast helped users better recognize the advertisement. Further, even varying the position of the disclosure to the middle and bottom of the article, and using the word *sponsored* was more effective than using a phrase such as *presented by [sponsor]* in helping users recognize the sponsored content.

In a recent report [38], the FTC also discovered that various aesthetic changes to disclosures, such as including borders to surround the advertisement or contrasting the color of the advertisement from other content can significantly increase users' ability to identify native advertisements. Collectively, these studies suggest that effective disclosures need to be visually distinguishable from other content.

2.2.2 Endorsement-based Advertising Disclosures. Disclosures requirements for endorsement-based advertising are described in the FTC's endorsement guidelines [18]. The guidelines state that *if there exists a connection between an endorser and the marketer that consumers would not expect and it would affect how consumers evaluate the endorsement, that connection should be disclosed.* The FTC also requires these disclosures to abide by the *clear and conspicuous* requirement so that consumers can identify them and subsequently decide how much weight to provide to the content creator's endorsement. Further, the disclosures should also not be buried in an *About* or *Terms of Use* page.

For affiliate marketing, the FTC's guidelines state that disclosures need to be placed close to any recommendations and URLs included by the content creator. The guidelines highlight that using *affiliate link* as a disclosure is insufficient, as users may not understand what an affiliate link is. Instead, the FTC recommends using a short phrase such as *I get commissions for purchases made through links in this post* close to the recommendation.

Only recently have studies begun examining how and whether users notice and understand the disclosures made by social media content creators. For instance, Evans *et al.* [12] found that stating *paid ad* was more effective than other disclosures in helping users identify advertisements on Instagram. In another study [43], van Reijmersdal *et al.* studied how disclosing sponsored content in blogs affected users. They found that after viewing the disclosure, users became more resistant to the endorsement of the blogger.

2.2.3 Other Types of Advertising Disclosures. Numerous studies have examined advertising disclosure practices in other forms of media such as television content [5, 6, 23, 28, 34]. One set of studies examined differences in television advertisements over time [23] and across demographics [34], or content [28]. These studies collectively discovered that while the presence of disclosures in advertisements have increased over the years in television advertisements, they still fail to adhere to the *clear and conspicuous* requirement.

Other more recent studies [5, 6] have examined how the length and position of the disclosure affects users' perceptions of sponsored television advertisements. They found that longer disclosures and disclosures made before the sponsored content is presented to users improves users' understanding of the nature of the content. These studies suggest that the timing of when a user sees a disclosure is important for the understanding the meaning of the disclosure.

In OBA—one of the most common forms of online advertising today—users see targeted advertisements based on their interests, demographics, and browsing histories. These online advertisements and the ability to opt-out of them are largely disclosed to users by means of the AdChoices icon—the current advertising industry self-regulation standard. This icon appears in nearly 60% of online advertisements from just 20 of the top 500 Alexa news websites [40].

Several recent studies [19, 26, 33, 42] have examined what various OBA disclosures communicate to users in practice, collectively revealing several shortcomings of the AdChoices disclosure. Hastak and Culnan [26] were the first to examine this in a survey of nearly 2,600 US adults users. They studied the efficacy of various OBA disclosures in communicating the underlying notice (recognizing the advertisement) and choice (being able to opt-out of tracking) to users. They found that the

AdChoices tagline did not work as well other tested taglines in aiding users' comprehension of the disclosure, suggesting that the wording of disclosures is important.

More recently, Ur *et al.* [42] and Leon *et al.* [33] found that non-expert users have several misconceptions about the AdChoices icon, often assuming that clicking on the icon would lead to more information about the advertised product or lead to more advertisements being shown. Garlach and Suthers [19] tested users' understanding of the AdChoices icon on mobile devices and found that participants were unable to locate and click on the icon on the small screens. These studies suggest that users may misunderstand visual indicators for disclosures altogether if they are unfamiliar with what a particular visual is supposed to represent.

2.3 Summary

Unlike the above studies, our study is the first to examine affiliate marketing disclosures in the context of endorsement-based advertising on social media. We perform a large scale measurement of social media content that contains affiliate URLs which drive revenue to the creator of the social media content. Further, rather than developing and testing new disclosures, we examine how effective the disclosures content creators currently make are in communicating the underlying endorsement-based advertisement to users. Based on prior work, we test how the placement and wording of disclosures affect user comprehension of disclosures.

3 CHARACTERIZING AFFILIATE MARKETING DISCLOSURES

In this section, we describe how we gathered content from two social media platforms to find affiliate marketing content and disclosures. Through this analysis, we answer our first and second primary research questions: how prevalent are disclosures on affiliate marketing content, and how do content creators disclose their affiliate marketing relationships with merchants to users? Figure 2 illustrates our data collection and analysis process.

We chose to specifically study YouTube and Pinterest to answer our research questions since these are two popular social media platforms that are designed to share content and reviews. However, our methods are platform agnostic and can be extended to other social media platforms.

3.1 Method

3.1.1 Data Collection. To answer our research questions, we first gathered data from YouTube and Pinterest. Our goal in this process was to gather the least possibly biased sample of YouTube videos and Pinterest pins in order to accurately model the prevalence of affiliate marketing content and disclosures on these platforms.

We considered three sampling procedures—all of which have been used in past studies—for collecting our data: sampling from large graphs, sampling using keywords, and prefix sampling. Sampling from large graphs involves sampling *Related* content graphs or social network graphs—both of which appear on YouTube and Pinterest as the *Related Videos* and *Related Pins* graph respectively. However, because *Related* content graphs have non-randomly selected edges, which are often biased towards content with high engagement—such as videos with higher view counts in the case of YouTube [49]—they result in non-uniform samples. Sampling using keywords (e.g., [4]) involves gathering samples by querying the search facilities of platforms using a set of compiled keywords. However, samples resulting from this procedure are likely to be biased towards the set of keywords, and in most cases, compiling a set of representative keywords is challenging in itself.

We therefore employed prefix sampling which has previously been used for sampling from YouTube [49] and Pinterest [20]. Prefix sampling works by generating—ahead of time—part of the identifier of some record(s) in the population, a *prefix*, which is then used to sample records beginning with that prefix. If the prefixes are uniformly generated, then the resulting samples will

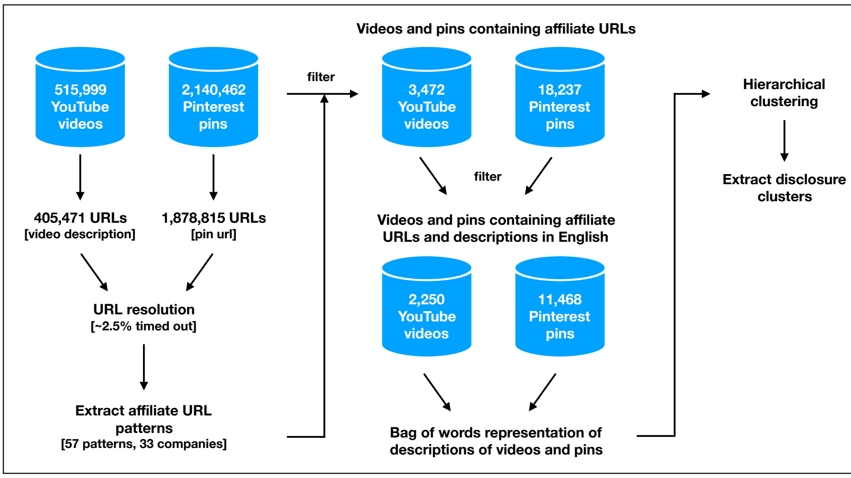

Fig. 2. Overview of the Data Collection, Data Cleaning, and Disclosure Extraction Stages.

be uniform too. This sampling procedure is particularly useful when the search space of identifiers is significantly larger than the number of already issued identifiers, and it is impractical to randomly generate the issued identifiers.

To apply prefix sampling on YouTube, we first observed that YouTube videos are assigned an eleven character long identifier. In order to gather our uniform sample, we first randomly generated video identifiers of length five, a technique used previously [49]. We then searched for those five character long identifiers using the YouTube Search API [21] which returned a list of videos beginning with that prefix. Similarly on Pinterest, pins are assigned a varying length identifier consisting wholly of numbers. The last five digits of the identifier however, represent a timestamp [20]. To gather a set of random prefixes, we first retrieved 25 pins—the number of pins returned on average—from the Pinterest *Categories* page [35] for each category. However, these pins were always the most recent pins in that category, and the prefixes they contained were non-random. Therefore, we then fetched all the related pins of these pins and sampled them randomly to create a set of *seed* pins. We then varied the last five digits of these *seed* prefixes from 00000–01500 (arbitrarily) to construct the final sample of pins. In total, we retrieved 515,999 unique YouTube videos and 2,140,462 unique Pinterest pins. While retrieving the videos and pins, we also recorded their characteristics such as what category of content they belonged to, how many times they had been viewed, how many comments had been made on the content, and details about their creators. We gathered this data between August and September 2017.

3.1.2 Discovering Affiliate Marketing Content. After gathering the samples of YouTube videos and Pinterest pins, we began the process of identifying affiliate marketing content in the samples. We first compiled a database of all the URLs in the videos and pins, specifically looking for URLs in the description of the YouTube videos and in the Pinterest pins. In total, we compiled a list of 405,471 and 1,878,815 URLs from the videos and pins respectively. To identify affiliate URLs from this set of URLs, we relied on the observation that affiliate URLs contain specific patterns [8]. For instance, Amazon's affiliate URL contains a *tag* parameter which indicates the identifier of the affiliate who stands to gain money from the purchase. However, while Amazon's affiliate URL appears directly on the destination website—the product page on Amazon.com—others may also appear during the intermediate redirects to the destination. Further, unlike Amazon, affiliate URL patterns may not necessarily only emerge as URL parameters; they may also be present in other parts of the URL including its path and sub-domain.

To ensure that our analysis accounted for all such cases, we first resolved all the URLs in our database, following both server-side and client-side redirects (HTTP 3XX, Meta refresh), and recorded the resulting intermediate URLs and all HTTP response codes. A total of ~2.5% of the URLs across both datasets failed to resolve either due to timeouts or HTTP 404 error codes; we ignored these URLs. Next, we performed a frequency analysis using each resolved URL's domain, sub-domain, path, and parameters, creating a list of commonly occurring patterns sorted by decreasing order of appearance (counts). We reasoned that if there existed any patterns across the URLs we resolved and visited, our frequency analysis would capture and bubble these to the top of the list. Starting with the sub-domain and path, we first recorded how many sub-domains (paths) each domain appeared with. A high number of sub-domains (paths) would signal that an affiliate marketing company likely caters to different merchants through unique sub-domains (paths), or that a constant sub-domain (path) appears as part of the affiliate URL. Next, we turned our attention to domains and their URL parameters. Rather than recording the number of parameters associated with each domain, we recorded the number of times a domain appeared with a URL parameter. A frequent co-occurrence of domains and URL parameters would signal that the parameter conveys some information about the content creator to the affiliate marketing company.

Because these lists contained a high number of false positives, we manually scanned each list, examining which of the domains, sub-domains, paths, parameters corresponded to affiliate marketing companies. To aid our examination, we used a combination of Google Search results, the WHOIS database, and the FMTC affiliate database [14], which provides a mapping from merchants to affiliate programs. Where possible, we also signed up on these programs as content creators to validate our findings about the affiliate programs. To limit the effort required to manually examine these lists, we only examined those combinations of domains and sub-domains/paths/parameters that appeared at least 15 times.

3.1.3 Discovering Affiliate Marketing Disclosures. After finalizing the affiliate URL patterns, we first filtered the list of resolved URLs to only retain those corresponding to these patterns. We then filtered the YouTube videos and Pinterest pins datasets to only those containing the affiliate URLs. Following this step, we began the process of extracting the disclosures from these videos and pins. We could have expected to find the disclosures—in theory—either during the course of the videos, the image of the pins, or the in the description boxes surrounding the videos and pins.

As a first step, we randomly sampled 20 videos and pins each from our filtered dataset of affiliate marketing videos and pins. We then examined these videos and pins to look for any affiliate marketing disclosures during the course of the videos and the pins' images. Because we found no disclosures in these random samples, we resorted to extracting disclosures in the descriptions of the videos and pins.

To extract the disclosures present in these videos' and pins' descriptions, we first split each description by its newlines and then by the sentences contained in each newline. We then tokenized each resulting sentence into a bag-of-words representation, and clustered the sentences using hierarchical clustering [36] with the euclidean distance metric. We chose a fairly low cut-off for the clusters based on the idea that the relevant smaller clusters containing the affiliate disclosures may already have been formed at that cut-off. We then manually examined these clusters one after the other, and recorded ones that contained disclosures pertaining to affiliate marketing. For this analysis, we only considered those descriptions that were written in English; doing so retained 64.8% and 62.8% of all affiliate videos and pins respectively.

3.1.4 Limitations. Our analyses of affiliate marketing content and disclosures has limitations. First, our method to discover affiliate URLs should be considered a lower bound on the number of affiliate URLs and consequently on affiliate marketing companies since we ignored those URL co-occurrences that appeared less than 15 times. Though we may have missed less prevalent companies owing to our frequency analysis, we are confident that our analysis presents a close approximation of affiliate URL prevalence. We also did not consider those affiliate programs that use coupon codes to track and attribute sales. Second, our method to discover affiliate marketing disclosures was limited to descriptions written in English. As a consequence, our findings may not generalize to other languages. In our data set, the descriptions of ~40% of all affiliate marketing videos and pins were written in a non-English language.

3.2 Findings

In the following sections, we present our findings, describing the characteristics of affiliate marketing content on both platforms, the types of disclosures made by content creators, and the characteristics of these disclosures.

3.2.1 Examining the Affiliate Marketing Landscape. Through our analysis, we compiled a total of 57 unique affiliate URL patterns from 33 different affiliate marketing companies—the most comprehensive publicly available list of its kind. This list facilitated automated detection of affiliate marketing content on both the social media platforms in our work and can be used to detect similar content on other platforms in future studies. We found at least one affiliate URL in a total of 0.67% or 3,472 videos on YouTube, and a total of 0.85% or 18,237 pins on Pinterest.

Table 5 in the Appendix lists the affiliate marketing companies we discovered along with their URL patterns, and the number of times we observed their presence across all URLs resolutions in our entire dataset. Note that a YouTube video may contain multiple affiliate URLs in its description unlike a Pinterest pin, which only contains one URL (that is, the pinned URL).

Across both YouTube and Pinterest, Amazon's Associate Program[1] had the largest presence (YouTube = 7,308, Pinterest = 7,368), closely followed by AliExpress' Affiliate Program[2] (YouTube = 2,167, Pinterest = 785). In addition to Amazon, we discovered other merchants that hosted in-house affiliate programs, as opposed to explicitly redirecting through an affiliate marketing company. For instance, Booking.com[3] and Apple[4] marketed products through their own affiliate programs. Some of the companies in our list are easily recognizable but many are not likely to be familiar to users.

3.2.2 Fashion, Beauty, and Technology Content Contain Highest Affiliate Marketing. Next, we analyzed how much content within each category—as defined by the platform—on YouTube and Pinterest contained affiliate URLs, and ranked the resulting categories. For this analysis, we only ranked those categories that had at least 100 affiliate marketing videos or pins.

The blue bar charts in Figure 3 illustrate our findings. Across both platforms, we found that affiliate marketing was most commonly found across content within technology, fashion, and beauty related categories. Prevalence of affiliate marketing across YouTube's *Science & Technology* and *Howto & Style* categories stood at 3.61% and 3.49% respectively. Similarly, prevalence of affiliate marketing across Pinterest's *Women's Fashion* and *Hair & Beauty* categories stood at 4.62% and 2.04% respectively.

We suspect the presence of affiliate marketing in these categories may be attributed to at least two different reasons. First, technology products, apparel, and beauty/cosmetics items are some of

[1]https://affiliate-program.amazon.com/
[2]https://portals.aliexpress.com
[3]https://www.booking.com/affiliate-program/v2/index.html
[4]https://www.apple.com/itunes/affiliates/

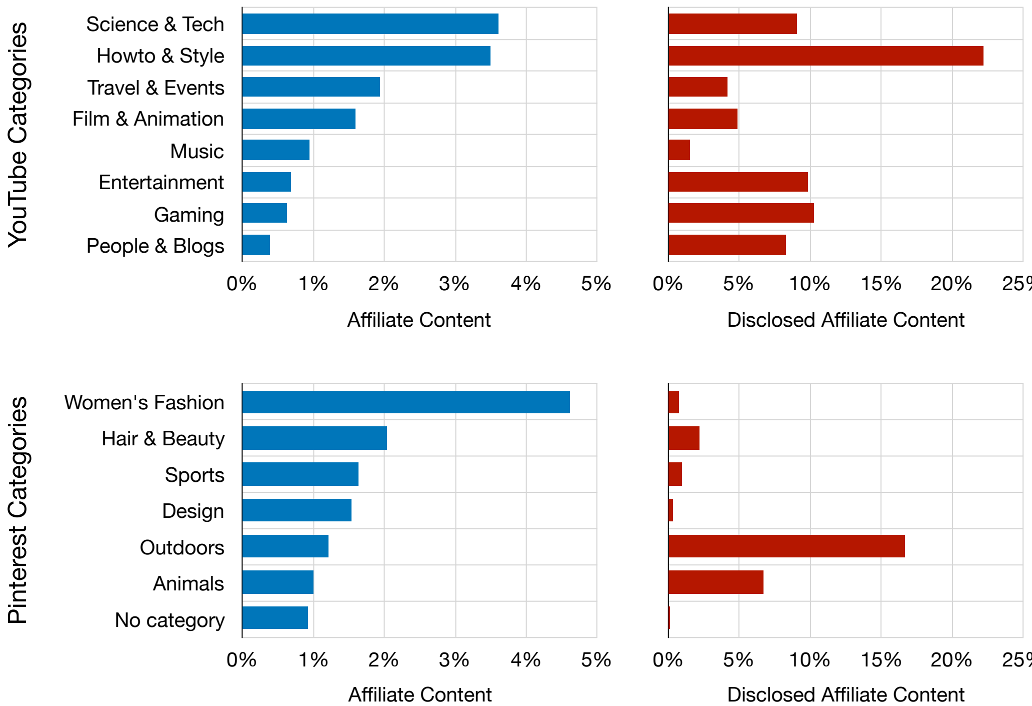

Fig. 3. Percentage of Affiliate Marketing Content Across All Content (Blue Bar Charts) and Percentage of Disclosed Affiliate Content (Red Bar Charts) Within Each YouTube and Pinterest Category. Only Categories With More Than 100 Affiliate Marketing Videos and Pins Are Shown.

the most common products consumers shop for online [25, 31, 39], so their presence in affiliate marketing might be indicative of this larger trend. Second, recent surveys have shown that users proactively seek out reviews before shopping for technology products online, and therefore this may be reflected in affiliate marketing content on these platforms [30].

3.2.3 Affiliate Marketing Content Has Higher User Engagement. Next, we examined how affiliate and non-affiliate marketing content correlated with user engagement metrics such as view, like, and comment counts. We conducted Mann–Whitney U tests to assess statistical significance. Accounting for multiple testing using Bonferroni correction, we tested for significance at the 0.01 level.

Across both YouTube and Pinterest, we noted a common thread: videos and pins with affiliate marketing content correlated with higher user engagement metrics. That is, these videos and pins had more comments, were liked more, and viewed more often than other non-affiliate marketing content. On YouTube, affiliate marketing videos correlated with longer duration length ($U \sim 7.95 \times 10^8$, $p < 0.0001$), higher view counts ($U \sim 7.72 \times 10^8$, $p < 0.0001$), higher like counts ($U \sim 6.96 \times 10^8$, $p < 0.0001$), and higher dislike counts ($U \sim 6.52 \times 10^8$, $p < 0.0001$). Similarly, on Pinterest, affiliate marketing pins correlated with higher repin counts ($U \sim 1.93 \times 10^{10}$, $p < 0.0001$). We could not directly compare the like and comment counts on Pinterest since the Pinterest API returned both as zero for all the pins in our dataset. Given the higher user engagement, affiliate marketing content is likely to surface through recommendations algorithms, and ensuring content creators disclose these endorsements becomes an even more pressing issue.

Table 1. Prevalence of Various Affiliate Disclosure Types On YouTube and Pinterest. Prevalence is Computed Across All Affiliate Content.

Disclosure	Platform	Prevalence (%)	Example
Affiliate Link	YouTube	7.02	Affiliate links may be present above
	Pinterest	4.60	(aff link)
Explanation	YouTube	1.82	This video contains affiliate links, which means that if you click on one of the links, I'll receive a small commission
	Pinterest	2.43	(This is an affiliate link and I receive a commission for the sales)
Channel Support	YouTube	2.44	AMAZON LINK: (Bookmark this link to support the show for free!!!)

3.2.4 Three Types of Affiliate Marketing Disclosures. Across YouTube and Pinterest, we discovered that 10.49% and 7.03% of affiliate marketing videos and pins respectively were disclosing the affiliate URLs in their content descriptions as such to users. We found three distinct types of disclosure clusters which we describe below; Table 1 summarizes our findings.

***Affiliate Link* Disclosures**: The first type of disclosures communicated to users that affiliate URLs were present in the content. On YouTube, these disclosures appeared in the video description either as blanket disclosures—a single disclosure across the entire description—or as disclosures highlighting the individual affiliate URLs, or both. Exactly 7.02% of affiliate marketing videos were disclosing in this manner. The following statements describe how YouTube content creators made these types of disclosures:

- *Affiliate links may be present above*
- *Some of the links may be affiliate links*
- *(Disclosure: These are affiliate links)*
- **Amazon link(s) are affiliate links*

On Pinterest, *Affiliate Link* disclosures appeared in the description of 4.60% of all affiliate marketing pins. Unlike YouTube, content creators' disclosures did not point to specific URLs, since the Pinterest pins only contained the pinned URL. The following statements describe how Pinterest content creators made these types of disclosures:

- *(aff link)*
- *(affiliate)*
- *#affiliatelink*
- *This is an Amazon Affiliate link*

***Explanation* Disclosures**: The second type of disclosures content creators made offered users a verbose explanation about affiliate marketing and affiliate URLs, and described how clicking on the URLs impacts users viewing the content. In comparison to *Affiliate Link* disclosures, content creators who used *Explanation* disclosures included more details, and often quoted specific merchants or affiliate marketing companies. On YouTube, these disclosures appeared in the descriptions of 1.82% all affiliate marketing videos. The following statements describe how YouTube content creators made these types of disclosures:

- *This video contains affiliate links, which means that if you click on one of the links, I'll receive a small commission*

- *I am an affiliate with eBay, Amazon, B&H and Adorama, which means I get a small commission when you buy through my links*
- ***Links that start with http://rstyle, Beautylish & MUG links are affiliate links, I do earn a small commission when you purchase through them, which helps me purchase products for review & improve my channel*

On Pinterest, *Explanation* disclosures appeared in the descriptions of 2.43% of all affiliate marketing pins. The following statements describe how Pinterest content creators made these disclosures:

- *(This is an affiliate link and I receive a commission for the sales)*

***Channel Support* Disclosures**: The third type of disclosures content creators made communicated to users that they would be supporting the channel by clicking on the affiliate URLs, without clearly explaining how. Only content creators on YouTube made these type of disclosures, which appeared in the descriptions of 2.44% of all affiliate marketing videos. The following statements describe how YouTube content creators made these types of disclosures:

- *AMAZON LINK: (Bookmark this link to support the show for free!!!)*
- *Support HWC while shopping at NCIX and Amazon*
- *Purchase RP here and help support this channel via the amazon affiliate program*
- *Shop using these links to support the channel*

To summarize, we found that only ~10% and ~7% of videos and pins respectively with affiliate marketing URLs contained disclosures by content creators. When present, these disclosures were of three types: *Affiliate Link* disclosures, *Explanation* disclosures, and *Channel Support* disclosures. Our results show that *Explanation* disclosures, which the FTC explicitly advocates content creators to use in its endorsement guidelines—and are also the most effective as we show our in user study—appeared least frequently.

3.2.5 Affiliate Marketing Disclosure Prevalence by Content Category. We also examined how the prevalence of content creators' affiliate marketing disclosures varied across the videos' and pins' categories (as we described in Section 3.2.2). To account for the number of content creators contributing towards the disclosures in each category, we scaled the overall disclosure percentage by the ratio of number of unique content creators accounting for those disclosures and the number of videos and pins containing disclosures. If the all the videos and pins containing disclosures were accounted for by unique creators, then the percentage remained unchanged.

The red bar charts in Figure 3 illustrate our findings. We observed some variance in the percentage of affiliate marketing content on both YouTube and Pinterest that contained any form of disclosure, with some outliers on both platforms. At about 22.5%, affiliate marketing videos in the *Howto & Style* category on YouTube contained more disclosures than affiliate marketing videos in the other categories. In stark contrast, disclosure rates in the corresponding *Hair & Beauty* and *Women's Fashion* categories on Pinterest was relatively low (0.5%–2.5%). Similarly, affiliate marketing pins in the *Outdoors* category on Pinterest had an unusually high overall disclosure percentage, at about 15%. We suspect the relatively high rate of disclosures on YouTube's fashion category may partially be explained by recent awareness of the FTC's guidelines among fashion and beauty vloggers on YouTube [1, 32]. However, as stated before, our analysis only examined disclosures written in English; the distribution of disclosures in affiliate content for other languages remains an area for future exploration.

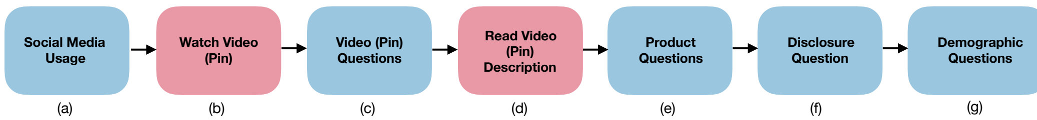

Fig. 4. The Overall Structure and Flow of the YouTube and Pinterest Experiments. Blue Boxes Represent Sections of the Experiment Where Participants Only Answered Questions; Red Boxes Represent Sections of the Experiment Where Participants Only Viewed Content.

4 USER STUDY OF AFFILIATE MARKETING DISCLOSURES

Having discovered the different types of affiliate marketing disclosures on YouTube and Pinterest, we turned to our third primary research question: How well do users notice and understand current forms of these affiliate marketing disclosures? More specifically, we asked:

- Do users notice the disclosures present in the description of the affiliate marketing content?
- Does the position of the disclosure affect whether users notice the disclosure?
- Regardless of whether users notice the disclosures, what do these disclosures communicate to users?
- How much do users think of the product embedded in affiliate marketing content (e.g., Figure 1b) as an endorsement by the content creator?
- How does the presence of a disclosure change users' perception of the endorsement?

To answer these questions, we conducted online experiments on the MTurk platform. We describe the logistics of the experiment, our subsequent data analysis, and findings in the following sections.

4.1 Method

4.1.1 Experiment Instrument. We conducted two online randomized controlled experiments on MTurk—one each for YouTube and Pinterest—employing a between-subjects full-factorial design in each experiment. We examined the effect of two independent variables—*disclosure type* and *video/pin*—on several dependent variables, which we describe in the following paragraphs. The first independent variable, *disclosure type* was a fixed factor and its levels consisted of the disclosure types that we found in the previous analysis. We further varied the *disclosure type* by position to appear either at the top or the bottom of the content description. In the YouTube experiment, this resulted in six levels with three disclosure types—(*Affiliate Link*, *Explanation*, and *Channel Support*)—spread across two disclosure positions (top, bottom). We chose the following disclosure statements for each disclosure type:

- Control: [No Disclosure]
- Affiliate Link Disclosure: *Affiliate links may be present*
- Explanation Disclosure: *This video contains affiliate links, which means that if you click on one of the links, I receive a commission for the sales*
- Channel Support Disclosure: *Shop using these links to support the channel*

Similarly, in the Pinterest experiment, the *disclosure type* independent variable contained four levels with two disclosure types—(*Affiliate Link* and *Explanation*)—spread across two disclosure positions (top, bottom). We chose the following disclosure statements for each disclosure type (Recall, Pinterest had no *Channel Support* type disclosures):

- Control: [No Disclosure]
- Affiliate Link Disclosure: *Affiliate link*
- Explanation Disclosure: *This is an affiliate link and I receive a commission for the sales*

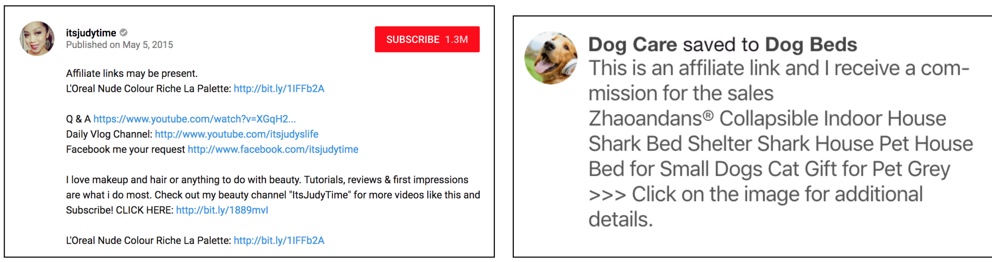

(a) *Affiliate Link* [Top] Condition (cropped) (b) *Explanation* [Top] Condition

Fig. 5. Examples of the Treatment Conditions in the (a) YouTube Experiment and (b) Pinterest Experiment. The Disclosure Statements Were Added to the Top of the Descriptions in these Treatment Conditions.

Our choice of the disclosure statements warrants a justification. Our goal with the user study was to determine the efficacy of the types of disclosures currently being made; we did not set out to design and test new hypothetical affiliate marketing disclosures, and we leave doing so for future work. This reasoning is reflected in our choices above: for each level in the *disclosure type* factor—whether in the YouTube or Pinterest experiment—we picked an instance of the disclosure that we thought would best represent disclosures in that category on that platform, but at the same time would generalize across videos and across pins.

The second independent variable *video/pin*, reflected the videos and pins participants watched during the experiments. This independent variable was a random factor in that we created its levels from the affiliate marketing videos and pins in our dataset. We considered *video/pin* as a random factor since we wished to examine how our dependent variables varied considering the differences across videos and pins, rather than examining the effect of video/pin on the dependent variables per se. To construct the levels of the *video/pin* random factor we sampled five videos and five pins from the list of affiliate marketing videos and pins our dataset. However, in order to limit the duration of the YouTube experiment, we only sampled the five videos from around the overall median affiliate marketing video length (~210 seconds); we added no such constraints to the five randomly selected affiliate marketing pins in the Pinterest experiment. We specifically chose five videos and pins to balance experimental overhead; a larger number would require more participants for a between-subjects design. None of these videos and pins contained any disclosures during the video or on the pin's image.

With these two independent variables and including the control, we arrived at a (5 X 7) and a (5 X 5) full-factorial design in the YouTube and Pinterest experiments respectively. In each experiment, we randomly assigned participants to one of the *cells* of the experiments.

Figure 4 illustrates the overall structure of the YouTube and Pinterest experiments. Each experiment contained five parts.

Part One. In the first part of the experiment (Figure 4a), we gathered details about participants' social media use. Specifically, we asked participants whether they had an account on a list of popular social media platforms including YouTube and Pinterest, how often they visited those platforms, and how often they posted content on those platforms.

Part Two. In the second part of the experiment, participants either watched a YouTube video or looked at a Pinterest pin (Figure 4b) which was embedded in the experiment page without its textual description. Immediately after this step, we asked participants' opinions about the content they watched (*video, pin impression*) on a Likert scale (extremely negative–extremely positive; 5 point), and by means of an open-ended question (Figure 4c). Based on how accurately they described the content, the responses to this question helped us ascertain whether participants actually viewed the content. In the YouTube experiment, we also logged participants' interactions with the embedded video player (play, pause, stop) to determine whether they watched the video.

Part Three. In the third part of the experiment (Figure 4d) we evaluated users' perceptions of the product embedded in the video or pin as an advertising endorsement by the content creator. In this part, we presented participants with a screenshot of the description of the video or pin they had just viewed. Depending on the condition, this was—including the control—either one of the seven conditions in YouTube experiment or one of the five conditions in the Pinterest experiment. Figure 5 shows one of treatments conditions for one video and one pin from both the YouTube and Pinterest experiments.

In the following page (Figure 4e), we presented participants with an image of the merchant's website that contained the product that was linked using an affiliate URL in the video or pin's description. We then asked participants to report their impression of the product (*product impression*) on a Likert scale (extremely negative–extremely positive; 5 point), and how much they thought the content creator of the video/pin favored the product (*content creator favors*) on a Likert scale (does not favor at all–strongly favors; 5 point). We used the latter as a measure of users' perceptions of the underlying endorsement by content creators.

To understand whether users notice the disclosures present in the descriptions of affiliate marketing content, we then asked participants how likely they thought there was a relationship between the content creator and the organization selling the product (*content creator relationship*) on a Likert scale (extremely unlikely–extremely likely; 5 point). We also included an open-ended question asking participants to explain their reasoning behind their answer to the previous question.

Part Four. In the fourth part of the experiment (Figure 4f), we presented the disclosure statement—which again varied depending on the treatment condition—to participants and asked them to explain what the statement meant in their own words by means of an open-ended question (*explain*). Using this question, we evaluated what these disclosures communicate to participants when they are presented—and are paying attention—to the disclosure statement.

Part Five. Finally, the fifth part of the experiment (Figure 4g) contained demographic questions.

Throughout the experiment(s), we referred to the product that was linked using the affiliate URL and the merchant as an *item* and *organization selling the item* to ensure that our questions appeared neutral and not leading.

4.1.2 Experiment Pilot and Deployment. We piloted the experiment(s) with 10 participants on UserBob[5], a usability testing website that employs MTurk workers for think-aloud screen-capture sessions. During the sessions, we asked participants to describe—in their own words—what the purpose of each question was, and whether they experienced any difficulty in answering them. Using this feedback, we restructured and refined the questions.

We then deployed the experiment(s) on MTurk as a *Give us your opinion about this YouTube Video (or Pinterest Pin)* task. We used this neutral phrase to avoid leaking any information about our motives behind the experiment. We required that participants completing the experiment be from the US, and have a MTurk approval rating of 95% or higher. Accounting for the minimum federal wage in the US ($7.25/hour), we paid participants $1.25 for completing the YouTube experiment and $0.75 for completing the Pinterest experiment since these took no longer than 10 and 5 minutes to complete respectively. We gathered all the data in March 2018. The study was approved by the Institutional Review Board of our university.

[5]http://userbob.com/

4.1.3 Participants. We recruited a total of 1,052 participants in the YouTube experiment and 753 participants in the Pinterest experiment, resulting in ∼30 participants in each condition for each video and pin. Of the 1,052 participants in the YouTube experiment, 14 (∼1.3%) skipped watching the video—as revealed by our measurements—or were not paying attention—as revealed by their open-ended responses—and were subsequently filtered out.

The mean age of participants in the YouTube experiment was 37.78, with a standard deviation of 12.11. Over half (52.60%) of all these participants were male, and nearly two-thirds (64.07%) had either a college or bachelor's degree. Close to nine-tenths (91.62%) of all participants had a YouTube account. Similarly, the mean age of participants in the Pinterest experiment was 38.29, with a standard deviation of 12.49. A little over two-fifths (41.57%) of all participants were male, and nearly two-thirds (63.21%) had either a college or bachelor's degree. Close to fourth-fifths (82.50%) of all participants had a Pinterest account.

4.1.4 Data Analysis. We built ordinal logistic regression mixed-models to analyze participant responses. Ordinal regression is used when the dependent variable is ordinal in nature (e.g., a Likert scale response). We used mixed-models to account for the *video/pin* independent variable, which we considered as a random factor. We built separate models for both YouTube and Pinterest using the ordinal package in R[6]. For each model, we also verified that the proportional odds assumption was not violated.

To analyze the open-ended response in which we asked participants to explain what the disclosure statement meant to them (*explain*), we adopted a qualitative data analysis approach and performed deductive coding [37]. Two researchers created and agreed upon a codebook after an initial exploration of all the responses. Using this codebook, both researchers independently coded the responses, blind to the condition of the responses to avoid biasing the assigned codes. We then calculated inter-rater reliability using Krippendorff's alpha (α), and achieved an $\alpha = 0.82$ and $\alpha = 0.84$ in the YouTube and Pinterest experiment responses, indicating high agreement. We then resolved the disagreements, and arrived at the final codes for each response. We make our codebook available in the Appendix (Table 6).

4.1.5 Limitations. Our user study has limitations. First, we only tested one disclosure statement within each disclosure type on both YouTube and Pinterest. We picked statements that were fairly generic enough to generalize across all the videos and pins. Second, our studies were limited in generalizability to the broader MTurk population. While the demographics of our participants, including age, gender, and education, were broadly comparable to the Internet users population [7], MTurk users may have differed in terms of Internet literacy, proficiency, and experience. Third, we randomly sampled the five YouTube videos around the median length of affiliate videos, and as a result our results may not generalize to videos of relatively longer and shorter duration.

4.2 Findings

In the following sections, we present our findings, highlighting the efficacy of the various disclosure types on both platforms. The output of the various ordinal logistic regression models for the YouTube and Pinterest experiments is summarized in Table 2 and Table 3 respectively. For each dependent variable across both experiments, the regression models compared the treatment conditions against the control.

4.2.1 Users Notice Explanation Disclosures, But Only On Pinterest. We asked participants how likely it was that there was a relationship between the content creator and the merchant selling the product. As seen in Table 2 we found no evidence that suggested any of the disclosure types

[6] https://cran.r-project.org/web/packages/ordinal/index.html

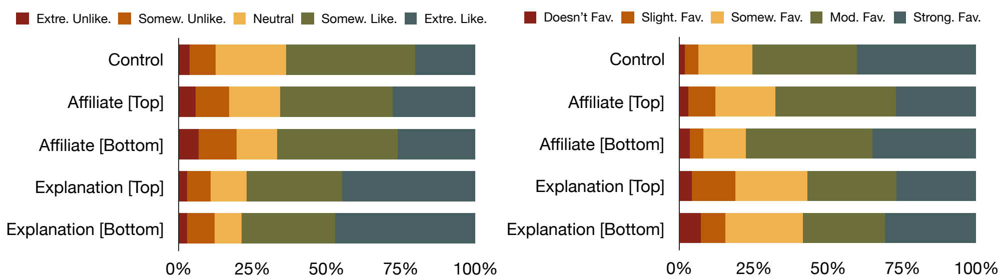

(a) Content Creator Relationship

(b) Content Creator Favors

Fig. 6. Distribution of the *Creator Relationship* and *Creator Favors* Variables from the Pinterest Experiment.

in the YouTube experiment had a statistically significant effect on participants' awareness of this relationship when compared to the control condition. This suggests that the participants either failed to notice the disclosures or interpreted them incorrectly, regardless of their type and position.

However, in the Pinterest experiment, and as summarized in Table 3, we found that the presence of a *Explanation* disclosure either at the top ($O.R. = 2.37$, $p < 0.0001$) or bottom ($O.R. = 2.82$, $p < 0.0001$) of the description multiplicatively increased the odds of participants' awareness of the relationship between the content creator and the merchant selling the product by 2.37 and 2.82 times respectively when compared to the control condition. Figure 6a shows the distribution of the *content creator relationship* dependent variable aggregated across the pins. While about 60% of participants in the control condition stated that it was either somewhat or extremely likely that there was a relationship between the content creator and the merchant, nearly 75% did so in the *Explanation* disclosure (Top & Bottom) conditions. This suggests that participants noticed and correctly interpreted the *Explanation* disclosures regardless of their position in the pins' description.

Across both platforms, we found no evidence to support that *Affiliate Link* disclosures were effective. As we show in a later result, this is likely not just because users may have not noticed these disclosure, but also due to unclear nature of the disclosure wording itself.

4.2.2 Similar Product Impressions Across All Conditions. We asked participants about their impression of the videos and pins (*video, pin impression*) before they were exposed to any treatment. As shown in Table 2 and Table 3, we found no evidence that suggested any of the disclosure types had a statistically significant effect on participants' impressions of the videos and pins when compared to the control condition in both experiments. We expected this result because the participants had not yet been exposed to the treatments when they answered this question. Therefore, any impressions of the videos and pins were likely averaged out due to randomization.

After participants read through the description of the video and pin—which varied depending on the condition of the experiments—we asked them about their impression of the product (*product impression*) that was linked by the content creator using an affiliate URL. As shown in Table 2 and Table 3, we found no evidence that suggested any of the disclosure types had a statistically significant effect on participants' impressions of the product when compared to the control condition in both the YouTube or Pinterest experiments. This suggests that participants' impression of the products did not vary due to the presence of any of the kinds of disclosures.

4.2.3 Explanation Disclosures Decrease Participants' Perceptions of Content Creators' Endorsement On Pinterest. After they had read through the videos' and pins' descriptions, we asked participants how much they thought the content creator favored the product that was linked using an affiliate URL (*content creator favors*). Only 3.13% and 3.71% of participants—across all the conditions—in the

Table 2. Ordinal Logistic Regression Models for the YouTube Experiment. Each Column Represents an Ordinal Regression Model. The Name of the Column Indicates the Dependent Variable. Values Within Brackets Represent the 95% Confidence Interval.

$*$ $p < 0.05$ $**$ $p < 0.01$ $***$ $p < 0.001$ $****$ $p < 0.0001$

Experimental Condition(s) [Control: No Disclosure]	**Video Impression**	**Product Impression**	**Creator Favors**	**Creator Relationship**
	Odds Estimate [95% C.I.]			
Affiliate Link [Top]	0.90 [0.60, 1.36]	0.85 [0.56, 1.29]	0.78 [0.52, 1.17]	0.79 [0.53, 1.17]
Affiliate Link [Bottom]	0.81 [0.53, 1.23]	0.91 [0.59, 1.40]	1.01 [0.68, 1.55]	1.00 [0.66, 1.52]
Explanation [Top]	0.74 [0.48, 1.13]	0.89 [0.58, 1.38]	0.80 [0.52, 1.23]	1.08 [0.71, 1.66]
Explanation [Bottom]	0.86 [0.56, 1.32]	0.91 [0.63, 1.52]	0.84 [0.54, 1.30]	1.30 [0.85, 1.97]
Channel Support [Top]	0.86 [0.55, 1.33]	0.73 [0.47, 1.14]	0.89 [0.57, 1.39]	0.68 [0.44, 1.05]
Channel Support [Bottom]	0.91 [0.60, 1.39]	0.95 [0.62, 1.45]	0.93 [0.61, 1.42]	0.74 [0.49, 1.12]
Random Effect [Std. Dev.]	0.67	0.66	0.68	0.42

YouTube and Pinterest experiments respectively thought the content creator did not favor the product at all. Specifically in the control condition—having no disclosure—this number was comparable: only 2.76% and 1.75% of all participants in the YouTube and Pinterest experiments respectively thought the content creator did not favor the product at all. This suggests that participants do think of the products featured through affiliate URLs as endorsements by the creators.

As seen in Table 2, we found no evidence that suggested participants' perceptions of the creators' endorsement of the affiliate marketing product varied across the disclosure types in the YouTube experiment compared to the control condition. However, in the Pinterest experiment, and as summarized in Table 3, we found that the presence of an *Affiliate Link* disclosure at the top of the description ($O.R. = 0.62$, $p < 0.05$), or an *Explanation* disclosure both at the top ($O.R. = 0.46$, $p < 0.001$) or the bottom ($O.R. = 0.53$, $p < 0.01$) of the description, all multiplicatively decreased the odds of participants' perceptions of content creators' endorsement of the product by nearly half when compared to the control condition. Figure 6b shows the distribution of the *content creator favors* dependent variable aggregated across the pins. While about 75% of participants in the control condition stated that the content creator either moderately or strongly favored the product, only 50% of participants did so in the *Explanation* (Top & Bottom) conditions. This result from the Pinterest experiment suggests that participants noticed and correctly interpreted the *Explanation* disclosure, and as a result perceived the content creators' disposition to be more neutral relative to the control condition. As before, we found no evidence to support that *Affiliate Link* disclosures were effective.

4.2.4 Participants Accurately Interpret Explanation Disclosures. We asked participants to explain what the disclosure statement meant to them in their own words by means of an open-ended

Table 3. Ordinal Logistic Regression Models for the Pinterest Experiment. Each Column Represents an Ordinal Regression Model. The Name of the Column Indicates the Dependent Variable. Values Within Brackets Represent the 95% Confidence Interval.

∗ $p < 0.05$ ∗∗ $p < 0.01$ ∗ ∗ ∗ $p < 0.001$ ∗ ∗ ∗∗ $p < 0.0001$

Experimental Condition(s) [Control: No Disclosure]	Pin Impression	Product Impression	Creator Favors	Creator Relationship
	Odds Estimate [95% C.I.]			
Affiliate Link [Top]	0.98 [0.65, 1.47]	1.49 [0.99, 2.24]	0.62^{*} [0.42, 0.93]	1.21 [0.81, 1.81]
Affiliate Link [Bottom]	0.92 [0.61, 1.38]	1.44 [0.96, 2.16]	0.90 [0.61, 1.35]	1.13 [0.76, 1.68]
Explanation [Top]	1.01 [0.67, 1.52]	0.93 [0.62, 1.40]	0.46^{***} [0.30, 0.69]	2.37^{****} [1.58, 3.56]
Explanation [Bottom]	1.29 [0.85, 1.87]	1.24 [0.82, 1.87]	0.53^{**} [0.35, 0.80]	2.82^{****} [1.86, 4.29]
Random Effect [Std. Dev.]	0.49	0.25	0.12	0.44

Table 4. Participants' Interpretation of the Disclosure Statement Broken Down by Disclosure Type.

Disclosure Type	YouTube		Pinterest	
	Incorrect	Correct	Incorrect	Correct
Affiliate Link	115	217	145	146
Explanation	14	270	15	276
Channel Support	40	236	NA	NA

question. We classified the resulting codes from the qualitative data analysis into two groups: *Correct* and *Incorrect*. The first group consisted of all those codes that suggested a relationship—not just financial—between the content creator and the merchant, or that the content was an advertisement, was sponsored or was a promotion. The second group consisted of all other codes; these represented either incorrect interpretations of the disclosure statements or null responses. Our goal with this division of codes was not motivated by whether participants understood the specifics of affiliate marketing, but more generally, that they understood that the content creator benefited in some way from a partnership with the merchant. Table 6 in the Appendix indicates the codes we marked *Correct* and *Incorrect*.

Table 4 summarizes our results. In the YouTube experiment, nearly 95%, 85%, and only 65% of participants in the *Explanation*, *Channel Support*, and *Affiliate Link* conditions respectively correctly interpreted the disclosure statement. A Chi-sq test of independence was significant ($\chi^2 = 93.14$, $p < 0.0001$), indicating indicating that the disclosure types and their interpretation—as defined by our groups—were not independent of each other. Post-hoc comparison tests revealed that there was a statically significant difference between the *Affiliate Link* and *Explanation* conditions ($\chi^2 = 79.82$, $p < 0.0001$), the *Affiliate Link* and *Channel Support* conditions ($\chi^2 = 31.15$, $p < 0.0001$), and the *Explanation* and *Channel Support* conditions ($\chi^2 = 13.62$, $p < 0.001$). Similarly, in the Pinterest experiment, nearly 95% and only 50% of participants in the *Explanation* and *Affiliate Link* conditions respectively correctly interpreted the disclosure statement. A Chi-square test of independence was—as before—significant ($\chi^2 = 183.85, p < 0.0001$).

Overall, when presented with the disclosure statement, the majority of participants correctly interpreted the *Explanation* and *Channel Support* disclosures, understanding that posting the affiliate marketing content benefited the content creator in some way. In contrast, fewer participants found *Affiliate Link* disclosures interpretable.

In summary, we found that *Explanation* disclosures are more effective than *Affiliate Link* disclosures when users can find them (e.g., on Pinterest); users likely fail to either notice or correctly interpret *Affiliate Link* disclosures. Further, we think the *Explanation* disclosures likely had an effect in the Pinterest experiment but not the YouTube experiment because of the norms and design of social media platforms. In the five affiliate pins we used in the experiment—and this is true for affiliate Pinterest pins more generally (*Median = 1 line*)—the descriptions were short and no more than one or two lines. On the other hand, because YouTube descriptions in the five videos were longer—with this being true for affiliate YouTube videos more generally (*Median = 14 lines*)—this may have limited the ability of users to identify the disclosures.

5 DISCUSSION

In this section, we discuss the implications of our findings and suggest directions for future work.

5.1 Detecting Affiliate Marketing Content and Disclosures

Through our work, we demonstrated a new method for discovering affiliate marketing companies' URL patterns. Using this method, we were able to detect affiliate marketing content on two social media platforms and generate a publicly available list for researchers to contribute to and use. Similarly, our method for discovering disclosures within content creators' textual descriptions can aid in building automated disclosure detection tools; we think this can be of broad use for regulators, platforms, and researchers to ultimately help users better identify endorsement-based advertisements. Lastly, both methods are extensible and can be used to replicate our study on other social media platforms.

5.2 Effectiveness of the FTC's Endorsement Guidelines

5.2.1 Low Prevalence of Disclosures. Our results showed that only a small percentage of affiliate marketing content from YouTube and Pinterest contained any disclosures at all. This suggests that the FTC's endorsement guidelines have only had a limited effect in ensuring that content creators disclose their affiliate marketing relationship to consumers. Even within the small set of disclosures that are made, the majority are of the kind the FTC specifically discourages using: the short *Affiliate Link* disclosures.

This presents an opportunity for further investigation and study: why do content creators fail to follow the guidelines? We think there might be at least two contributing reasons. First, it is not clear whether content creators are even aware of the guidelines to begin with. Future work could—by means of surveys and interviews—examine the reasons behind why content creators on social media choose to or fail to disclose their affiliate marketing content. Second, the FTC has only recently begun enforcing its endorsement guidelines, setting precedents for violations [1, 13, 16, 17]. We expect that endorsement disclosures—including those related to affiliate marketing—will appear more frequently as awareness around the guidelines grows.

5.2.2 Explanatory Disclosures Help But Only When Users Can Find Them. Our study validates the FTC's endorsement guidelines that explanatory disclosures are indeed effective. Users are able to better understand the underlying relationship between the content creator and the merchants through the *Explanation* disclosure, but our results also show platform-specific effects: users may not be able to notice and interpret disclosures that are buried in long texts, regardless of their clarity, such as on YouTube. Therefore, we recommend that content creators ensure that they create disclosures in a manner that facilitates discovery (e.g., by surrounding the disclosure by asterisks and underscores).

5.3 Shifting the Burden to Affiliate Marketing Companies

Content creators could be held accountable to better disclosure practices by the affiliate marketing companies they associate with. We examined the affiliate marketing terms and conditions, where publicly available, of eight of the most prevalent affiliate marketing companies from our dataset: Amazon, AliExpress, Commission Junction, Rakuten Marketing, Impact Radius, RewardStyle, ShopStyle and ShareASale. We could not find any publicly available terms and conditions from Impact Radius or RewardStyle. Only Amazon[7] and ShopStyle[8] explicitly referenced the FTC's guidelines in their affiliate terms. While we did not find any references in the Rakuten Marketing, ShareASale, Commission Junction affiliate terms, we noted all three companies blogging about the guidelines on their company blogs[9,10,11].

However, its unclear how many other affiliate marketing companies follow this practice, and whether they explicitly point the content creators that register with them to the FTC's guidelines. In addition to focusing on the content creators, we suggest that the FTC could work with these companies to revise their disclosure requirements and standards.

5.4 Design Suggestions for Effective Disclosures

Our findings also unlock several design suggestions that various stakeholders in the social media advertising ecosystem—including social media platforms and web browsers—can incorporate to improve disclosure practices, and help users identify these disclosures.

5.4.1 Affordances Through Social Media Platforms. Along with content creators, social media platforms play a critical role in shaping the disclosures that content creators make. The disclosures, generally speaking, are limited by the character space available to them. For instance, the description length can be as long as 5,000 characters on YouTube, but on Pinterest it is capped to 500 characters; similarly, tweets can only be as long as 280 characters on Twitter. Because some of these interfaces may be more more conducive for content creators to add disclosures than others, social media platforms can help design their interfaces to make it easier for content creators to disclose without crowding their other promotion text.

In fact, Instagram recently added an option for sponsored content to be disclosed by using the *Paid partnership* tool, which enables disclosures outside of the traditional image description [24]. Similarly, YouTube recently added the ability for create a *Contains Product Placement* overlay to their videos [22]. Such disclosure tools are a step in the right direction, however it is unlikely that any one blanket disclosure will cover all marketing practices. As the FTC's endorsement guidelines evolve, social media platforms can also help design tools that enable content creators to retroactively edit disclosures from their past posts. Future research could focus on examining in depth the kinds of affordances that can be built into social media platforms to enable content creators to disclose their advertisements clearly and conspicuously.

[7]https://affiliate-program.amazon.com/help/operating/agreement
[8]https://www.shopstylecollective.com/terms
[9]http://blog.shareasale.com/2017/09/08/ftc-updates-and-faq-s/
[10]https://blog.marketing.rakuten.com/topic/ftc-disclosure-guidelines
[11]http://junction.cj.com/article/disclosures-what-you-need-know

5.4.2 Protection Through Web Browsers. Web browsers could also help inform users that the content they are watching has affiliate URLs by means of in-built support or through add-ons and extension. Such tools could function like current Ad-blockers, but rather than blocking advertisements, they could either highlight when a piece of content should contain disclosures, or highlight the actual disclosures when they are present. These tools could leverage *Explanation* disclosures' high interpretability rate in presenting the disclosure information to users.

Web browsers can achieve this in practice using machine learning and natural language processing based approaches in which algorithms can be trained on large datasets of labelled social media content. Our findings provide a starting point, showing that the three categories of affiliate disclosures often use wording that exhibit a common pattern. In future work, we aim to build such tools, and conduct user studies to evaluate how well they work in practice.

6 CONCLUSION

We examined advertising disclosures accompanying affiliate marketing content on two social media platforms: YouTube and Pinterest. Our study reveals that only about 10% of affiliate marketing content on both platforms contains *any* disclosures at all. Further, the explanatory disclosures which the FTC recommends using are indeed the ones we found most effective; however, these were also the least prevalent of all disclosure types. We offer practical recommendations from both a design and policy perspective to help enable better disclosures for users. Doing so can ultimately help increase transparency surrounding social media endorsements.

ACKNOWLEDGMENTS

We thank our participants for taking part in our study. We also thank Nathan Malkin and the anonymous reviewers for their feedback on our paper. This research was funded by NSF grant CNS-1664786.

REFERENCES

[1] AdAge. 2016. FTC Cracking Down on Social Influencers' Labeling of Paid Promotions. http://adage.com/article/digital/ftc-cracking-social-influencers-labeling-promotions/305345/.

[2] AdChoices. 2018. AdChoices. http://youradchoices.com/.

[3] Michelle A Amazeen and Ashley R Muddiman. 2018. Saving Media or Trading on Trust? The effects of native advertising on audience perceptions of legacy and online news publishers. *Digital Journalism* 6, 2 (2018), 176–195.

[4] Lisa Anthony, YooJin Kim, and Leah Findlater. 2013. Analyzing User-generated Youtube Videos to Understand Touchscreen Use by People with Motor Impairments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1223–1232. https://doi.org/10.1145/2470654.2466158

[5] Sophie C Boerman, Eva A Reijmersdal, and Peter C Neijens. 2012. Sponsorship disclosure: Effects of duration on persuasion knowledge and brand responses. *Journal of Communication* 62, 6 (2012), 1047–1064.

[6] Sophie C Boerman, Eva A Reijmersdal, and Peter C Neijens. 2014. Effects of sponsorship disclosure timing on the processing of sponsored content: A study on the effectiveness of European disclosure regulations. *Psychology & Marketing* 31, 3 (2014), 214–224.

[7] Pew Research Center. 2018. Internet/Broadband Fact Sheet. http://www.pewinternet.org/fact-sheet/internet-broadband/.

[8] Neha Chachra, Stefan Savage, and Geoffrey M Voelker. 2015. Affiliate crookies: characterizing affiliate marketing abuse. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*. ACM, 41–47.

[9] Federal Trade Commission. 2017. What's affiliate marketing? Should I care? https://www.consumer.ftc.gov/blog/2017/09/whats-affiliate-marketing-should-i-care.

[10] Federal Trade Commission. Accessed 2018. Dot Com Disclosures: Information about Online Advertising. https://www.ftc.gov/sites/default/files/attachments/press-releases/ftc-staff-issues-guidelines-internet-advertising/0005dotcomstaffreport.pdf.

[11] Federal Trade Commission. Accessed 2018. Enforcement Policy Statement on Deceptively Formatted Advertisements. https://www.ftc.gov/system/files/documents/public_statements/896923/151222deceptiveenforcement.pdf.

[12] Nathaniel J. Evans, Joe Phua, Jay Lim, and Hyoyeun Jun. 2017. Disclosing Instagram Influencer Advertising: The Effects of Disclosure Language on Advertising Recognition, Attitudes, and Behavioral Intent. *Journal of Interactive Advertising* 17, 2 (2017), 138–149. https://doi.org/10.1080/15252019.2017.1366885 arXiv:https://doi.org/10.1080/15252019.2017.1366885

[13] Lesley Fair. 2016. FTC challenges influencer campaign for Warner Bros.' Shadow of Mordor game. https://www.ftc.gov/news-events/blogs/business-blog/2016/07/ftc-challenges-influencer-campaign-warner-bros-shadow-mordor.

[14] FMTC. 2018. FMTC Merchant Hub. http://bit.ly/2dqgQTz.

[15] FTC. 1914. Federal Trade Comission Act. https://www.ftc.gov/enforcement/statutes/federal-trade-commission-act.

[16] FTC. 2016. Lord & Taylor Settles FTC Charges It Deceived Consumers Through Paid Article in an Online Fashion Magazine and Paid Instagram Posts by 50 "Fashion Influencers". https://www.ftc.gov/news-events/press-releases/2016/03/lord-taylor-settles-ftc-charges-it-deceived-consumers-through.

[17] FTC. 2017. CSGO Lotto Owners Settle FTC's First-Ever Complaint Against Individual Social Media Influencers. https://www.ftc.gov/news-events/press-releases/2017/09/csgo-lotto-owners-settle-ftcs-first-ever-complaint-against.

[18] FTC. 2018. The FTC's Endorsement Guides: What People Are Asking. https://www.ftc.gov/tips-advice/business-center/guidance/ftcs-endorsement-guides-what-people-are-asking.

[19] Stacia Garlach and Daniel Suthers. 2018. "I'm supposed to see that?" AdChoices Usability in the Mobile Environment. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.

[20] Eric Gilbert, Saeideh Bakhshi, Shuo Chang, and Loren Terveen. 2013. "I Need to Try This"?: A Statistical Overview of Pinterest. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2427–2436. https://doi.org/10.1145/2470654.2481336

[21] Google. 2018. Add YouTube functionality to your app. https://developers.google.com/youtube/v3/.

[22] Google. 2018. Paid product placements and endorsements. https://support.google.com/youtube/answer/154235?hl=en.

[23] Mariea Grubbs Hoy and J Craig Andrews. 2004. Adherence of prime-time televised advertising disclosures to the "clear and conspicuous" standard: 1990 versus 2002. *Journal of Public Policy and Marketing* 23, 2 (2004), 170–182.

[24] Anthony Ha. 2017. Instagram is testing a new way for celebrities and influencers to identify their sponsored posts. https://techcrunch.com/2017/06/14/instagram-sponsored-posts/.

[25] Sarah Halzack. 2015. The surprising thing that got the biggest share of online shopping dollars in 2015. https://goo.gl/ZTrJ8d.

[26] Manoj Hastak and C Culnan. 2010. Future of privacy forum online behavioral advertising "icon" study.

[27] Chris Hoofnagle and Eduard Meleshinsky. 2015. Native advertising and endorsement: Schema, source-based misleadingness, and omission of material facts. (2015).

[28] Mariea Grubbs Hoy and Michael J. Stankey. 1993. Structural Characteristics of Televised Advertising Disclosures: A Comparison with the FTC Clear and Conspicuous Standard. *Journal of Advertising* 22, 2 (1993), 47–58. http://www.jstor.org/stable/4188880

[29] Mengtian Jiang, Brigitte A. McKay, Jef I. Richards, and Wally Snyder. 2017. Now You See Me, But You Don't Know: Consumer Processing of Native Advertisements in Online News Sites. *Journal of Interactive Advertising* 17, 2 (2017), 92–108. https://doi.org/10.1080/15252019.2017.1399839 arXiv:https://doi.org/10.1080/15252019.2017.1399839

[30] KPMG. 2015. Survey Confirms the Value of Reviews, Provides New Insights. https://www.powerreviews.com/blog/survey-confirms-the-value-of-reviews/.

[31] KPMG. 2017. The truth about online consumers. https://assets.kpmg.com/content/dam/kpmg/xx/pdf/2017/01/the-truth-about-online-consumers.pdf.

[32] The Fashion Law. 2016. Beauty Vloggers Are Probably Violating FTC Guidelines. http://www.thefashionlaw.com/home/beauty-vloggers-are-probably-violating-ftc-guidelines.

[33] Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, and Guzi Xu. 2012. What Do Online Behavioral Advertising Privacy Disclosures Communicate to Users?. In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society (WPES '12)*. ACM, New York, NY, USA, 19–30. https://doi.org/10.1145/2381966.2381970

[34] Darrel D. Muehling and Richard H. Kolbe. 1998. A Comparison of Children's and Prime-Time Fine-Print Advertising Disclosure Practices. *Journal of Advertising* 27, 3 (1998), 37–48. https://doi.org/10.1080/00913367.1998.10673561 arXiv:https://doi.org/10.1080/00913367.1998.10673561

[35] Pinterest. 2018. Categories. https://www.pinterest.com/categories/.

[36] Lior Rokach and Oded Maimon. 2005. *Clustering Methods*. Springer US, Boston, MA, 321–352. https://doi.org/10.1007/0-387-25465-X_15

[37] Irving Seidman. 2013. *Interviewing as qualitative research: A guide for researchers in education and the social sciences*. Teachers college press.

[38] FTC Staff. 2017. An Exploration of Consumers' Advertising Recognition in the Contexts of Search Engines and Native Advertising. (2017).

[39] Statista. 2016. Share of internet users who have ever purchased products online as of November 2016, by category. https://www.statista.com/statistics/276846/reach-of-top-online-retail-categories-worldwide/.

[40] Grant Storey, Dillon Reisman, Jonathan Mayer, and Arvind Narayanan. 2017. The future of ad blocking: An analytical framework and new techniques. *arXiv preprint arXiv:1705.08568* (2017).

[41] Twitter. 2018. What are Promoted Tweets? https://business.twitter.com/en/help/overview/what-are-promoted-tweets.html.

[42] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. ACM, 4.

[43] Eva A van Reijmersdal, Marieke L Fransen, Guda van Noort, Suzanna J Opree, Lisa Vandeberg, Sanne Reusch, Floor van Lieshout, and Sophie C Boerman. 2016. Effects of disclosing sponsored content in blogs: How the use of resistance strategies mediates effects on persuasion. *American Behavioral Scientist* 60, 12 (2016), 1458–1474.

[44] Bartosz W Wojdynski. 2016. The deceptiveness of sponsored news articles: How readers recognize and perceive native advertising. *American Behavioral Scientist* 60, 12 (2016), 1475–1491.

[45] Bartosz W Wojdynski, Hyejin Bang, Kate Keib, Brittany N Jefferson, Dongwon Choi, and Jennifer L Malson. 2017. Building a Better Native Advertising Disclosure. *Journal of Interactive Advertising* (2017), 1–12.

[46] Bartosz W Wojdynski and Nathaniel J Evans. 2016. Going native: Effects of disclosure position and language on the recognition and evaluation of online native advertising. *Journal of Advertising* 45, 2 (2016), 157–168.

[47] Katrina Wu. 2016. YouTube marketing: Legality of sponsorship and endorsements in advertising. *JL Bus. & Ethics* 22 (2016), 59.

[48] David Zax. 2017. The War To Sell You A Mattress Is An Internet Nightmare. https://www.fastcompany.com/3065928/sleepopolis-casper-bloggers-lawsuits-underside-of-the-mattress-wars.

[49] Jia Zhou, Yanhua Li, Vijay Kumar Adhikari, and Zhi-Li Zhang. 2011. Counting YouTube Videos via Random Prefix Sampling. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (IMC '11)*. ACM, New York, NY, USA, 371–380. https://doi.org/10.1145/2068816.2068851

Received April 2018; revised July 2018; accepted September 2018

Table 5. List of Affiliate Marketing Companies Discovered by Our Analysis. Count Indicates the Number of times the URL Pattern Appeared When We Resolved the Retrieved URLs.

Company Name	Domain	URL Pattern	YouTube Count	Pinterest Count
admitad	admitad	https://ad.admitad.com/g/… https://ad.admitad.com/goto/…	245	1
affiliaXe	affiliaxe	http://performance.affiliaxe.com/…&aff_id=…	151	0
AliExpress	aliexpress	https://s.aliexpress.com/…&af=…	2167	785
Amazon	amazon	http://www.amazon.(com,de,fr,in,it)/…&tag=…	7308	7368
Apple	apple	https://itunes.apple.com/…&at=…	669	61
Audiobooks	audiobooks	https://affiliates.audiobooks.com/…&a_aid=…&a_bid=…	129	0
AvantLink	avantlink	http://www.avantlink.com/…&pw=…	34	12
Avangate	avangate	https://secure.avangate.com/…&AFFILIATE=…	12	0
Awin	awin1 zanox zenaps	http://www.awin1.com/…&awinaffid=… http://ad.zanox.com/ppc/?… http://www.zenaps.com/rclick.php?…	129	211
Banggood	banggood	http://www.banggood.com/…&p=…	88	13
Book Depository	bookdepository	https://www.bookdepository.com/…&a_aid=…	103	0
Booking.com	booking	https://www.booking.com/…&aid=…	257	7
Clickbank	clickbank	http://….hop.clickbank.net/…	678	262
CJ Affiliate	anrdoezrs dotomi dpbolvw emjcd jdoqocy kqzyfj qksrv tkqlhce	http://www.anrdoezrs.net/click-[0-9]+-[0-9]+… http://cj.dotomi.com/… http://www.dpbolvw.net/click-[0-9]+-[0-9]+… http://www.emjcd.com/… http://www.jdoqocy.com/click-[0-9]+-[0-9]+… http://www.kqzyfj.com/click-[0-9]+-[0-9]+… http://qksrv.net/… http://www.tkqlhce.com/click-[0-9]+-[0-9]+…	341	2413
Ebay	ebay	http://rover.ebay.com…&campid=…	99	1963
Envato	audiojungle codecanyon envato graphicriver themeforest videohive	https://audiojungle.net/…&ref=… https://codecanyon.net/…&ref=… https://marketplace.envato.com/…&ref=… https://graphicriver.net/…&ref=… https://themeforest.net/…&ref=… https://videohive.net/…&ref=…	108 14 175 15 19 578	0 76 262 1465 200 33
e-Commerce Partners Network	buyeasy	http://buyeasy.by/cashback/… http://buyeasy.by/redirect/…	741	7
Flipkart	flipkart	https://www.flipkart.com/…&affid=…	81	20
GT Omega Racing	gtomegaracing	http://www.gtomegaracing.com/…&tracking=…	56	0
Hotellook	hotellook	https://search.hotellook.com/…&marker=…	165	5
Hotmart	hotmart	https://www.hotmart.net.br/…&a=…	211	8
Impact Radius	7eer evyy	http://….7eer.net/c/[0-9]+/[0-9]+/[0-9]+… http://….evyy.net/c/[0-9]+/[0-9]+/[0-9]+…	180	529
KontrolFreek	kontrolfreek	https://www.kontrolfreek.com/…&a_aid=…	117	0
Makeup Geek	makeupgeek	http://www.makeupgeek.com/…&acc=…	57	0
Pepperjam Network	gopjn pjatr pjtra pntra pntrac pntrs	http://www.gopjn.com/t/[0-9]-[0-9]+-[0-9]+-[0-9]+… http://www.pjatr.com/t/[0-9]-[0-9]+-[0-9]+-[0-9]+… http://www.pjtra.com/t/[0-9]-[0-9]+-[0-9]+-[0-9]+… http://www.pntra.com/t/[0-9]-[0-9]+-[0-9]+-[0-9]+… http://www.pntrac.com/t/[0-9]-[0-9]+-[0-9]+-[0-9]+… http://www.pntrs.com/t/[0-9]-[0-9]+-[0-9]+-[0-9]+…	2	79
Rakuten Marketing	linksynergy	http://click.linksynergy.com/…&id=…	189	1877
Skimlinks	redirectingat	http://go.redirectingat.com/…&id=…	43	155
Smartex	olymptrade	https://olymptrade.com/…&affiliate_id=…	65	0
RewardStyle	rstyle	http://rstyle.me/…	402	2711
ShopStyle	shopstyle	http://shopstyle.it/…	111	9239
ShareASale	shareasale	http://www.shareasale.com/r.cfm… http://www.shareasale.com/m-pr.cfm… http://www.shareasale.com/u.cfm…	199	616
Studybay	apessay	https://apessay.com/…&rid=…	141	0
Zaful	zaful	http://zaful.com/…&lkid=…	32	786

Table 6. Codebook for the Open-Ended Response (*Explain*) from the YouTube and Pinterest Experiments. *Correct* Indicates Whether the Code was Marked to be Interpreted Correctly in the Participant Reponses.

Code	**Interpretation**	**Marked *Correct?***
creator_makes_money_url_purchase	The creator makes money when a product is purchased through the URL	Yes
creator_makes_money_url_click	The creator makes money when the URL is clicked	Yes
creator_makes_money_sponsored	The creator makes money through sponsorships	Yes
creator_makes_money_url_owner	The creator makes money thorough the owner of the URL	Yes
creator_makes_money_through_ads	The creator makes money through ads on the URL's page	Yes
creator_makes_money_from_company	The creator makes money from the company selling the product	Yes
creator_receives_benefit	The creator receives "credit", or "points", or "discounts", or "incentives"	Yes
creator_makes_money_from_youtube	The creator makes money from YouTube for clicks on the URL	Yes
creator_makes_money_from_showcasing	The creator makes money from showcasing the product	Yes
creator_affiliated_company	The creator is connected or affiliated with a company (e.g., Amazon)	Yes
creator_content_sponsored	The creator's content is sponsored	Yes
url_is_advertising	The URL itself is an advertising	Yes
content_paid_ad	The content is a paid advertisement, or is paid publicity	Yes
third_party_promotion	The video is a a third-party promotion of a product	Yes
url_to_page	The URL is a URL to a product page for convenience	No
url_track_traffic	The URL is to track traffic	No
do_not_know	Explicit statement stating "no idea" or similar	No
not_clear	Respondent found the statement unclear	No
creator_paying_for_ads	The creator is paying for ads	No

Appendix B

Is This An Ad?: Automatically Disclosing Online Endorsements On YouTube With AdIntuition

Michael Swart
Princeton University

Ylana Lopez
Princeton University

Arunesh Mathur
Princeton University

Marshini Chetty
University of Chicago

ABSTRACT

Undisclosed online endorsements on social media can be misleading to users who may not know when viewed content contains advertisements. Despite federal regulations requiring content creators to disclose online endorsements, studies suggest that less than 10% do so in practice. To overcome this issue, we need knowledge of how to best detect online endorsements, knowledge about how prevalent online endorsements are in the wild, and ways to design systems to automatically disclose advertising content to viewers. To that end, we designed, implemented, and evaluated a tool called AdIntuition which automatically discloses when YouTube videos contain affiliate marketing, a type of social media endorsement. We evaluated AdIntuition with 783 users using a survey, field deployment, and diary study. We discuss our findings and recommendations for future measurements of and tools to detect and alert users about affiliate marketing content.

Author Keywords

social media, browser extension, advertisements, influencer

CCS Concepts

•Human-centered computing → Human computer interaction (HCI);

INTRODUCTION

Online endorsements are a form of advertising that help social media influencers to monetize their content [16]. These influencers are paid because of the perception that they are able to shape the opinion of their followers on a daily basis. For instance, Kylie Jenner, a social media influencer, reportedly made \$26.5 Million from just 53 Instagram advertisements [6]. The high fees that brands are willing to pay for online endorsements show how lucrative they can be for content creators [39]. When content creators do form a connection to a brand, they are required by the Federal Trade Commission (FTC) to disclose this relationship on the platforms they use for endorsements [11]. However, a recent study suggests less than 10% actually do so on YouTube and Pinterest [21] and this is problematic because viewers may be misled by these disguised advertisements.

The consequence of undisclosed endorsements may be benign: a viewer may not realize that an influencer's endorsement of a product is inauthentic. More extreme consequences include financial loss such as when the content creator is being deceptive [23] or the product being marketed has deceptive practices [8]. Arguably, influencers 'disguised advertisements' are a form of 'dark pattern' [4], a choice to leave out information in order to lead a viewer down a decision-making path for the benefit of the influencer. Undisclosed advertising content could even be classified as misinformation, causing viewers to falsely believe the content is unbiased, which can be particularly egregious if in aid of a political agenda [30].

Although some platforms such as YouTube do allow content creators to self-indicate a sponsorship, for instance, checking a box to indicate 'Paid Promotion' [13], it is unclear how often content creators use these features. Yet, studies have shown having these types of disclosures can help users identify advertisements and form more critical attitudes towards the brand being promoted [9, 34, 3, 37]. The question then is how can we automatically detect and communicate to a user when content contains online endorsements so that users are informed about the content they are consuming? We set out to address this issue by focusing on one type of online endorsement, affiliate marketing, on the YouTube platform. In affiliate marketing, influencers with a brand relationship are paid for sales or referrals generated from their content consumers. For example, a YouTuber may earn a commission if a video viewer clicks on a brand-generated Uniform Resource Locator (URL) provided to the YouTuber that links to a brand's product or website. These links serve as a source of ground truth for automatically detecting affiliate marketing content which makes automatic disclosure feasible. In our work, we posed the following research questions:

- How can we best detect and measure affiliate marketing content automatically and in real-time?
- How can we design and implement automatic ad disclosures for affiliate marketing content?
- How do users react to real-time automatic ad disclosures ?

To answer these questions, we used the only publicly available existing data set of known affiliate marketing link patterns [19] and a set of 0.5 million YouTube videos compiled by Mathur et al. [21] to detect affiliate marketing content from YouTube video descriptions. We used the data set of videos to identify two additional features for detecting affiliate marketing content automatically, 1) url parameters called Uniform Tracking Modules (UTM) [14] in links in video descriptions and 2) customized coupon codes [24] used in video descriptions. We

CHI '20, April 25–30, 2020, Honolulu, HI, USA.
2020 Association of Computing Machinery.
ACM ISBN 9978-1-4503-6708-0/20/04 ...\$15.00.
http://dx.doi.org/10.1145/3313831.3376178

built a coupon code classifier and detection module for these features in YouTube videos in real-time. Next, we integrated these modules into a browser extension called *AdIntuition* that automatically detects and discloses affiliate marketing content to users using a banner displayed above a video and by highlighting links and coupon codes in video descriptions.

We evaluated the efficacy of our interface using three studies. In Study 1, we performed a survey of 300 Amazon Mechanical Turkers to evaluate as part of the AdIntuition design process to test if the banner was noticeable to users and if they understood the disclosure. In Study 2, we conducted a diary study with 11 users over two weeks to examine how users react to automatic advertising disclosures and in Study 3, we conducted a field deployment with 472 users to evaluate how well our coupon code classifier and feature detection modules performed with videos real users watch and to learn more about how real users encounter affiliate marketing content. Based on our studies, we have the following main findings:

- Users with were able to more easily identify advertising content with AdIntuition disclosures than without them.

- Users became more reflective on content creators and content when automatic disclosures were presented to them in real-time

- Affiliate marketing content on YouTube can be detected in real-time using the known affiliate link patterns from Mathur et al. [21] but also by detecting the additional features we identified, UTM parameters and coupon codes, in video descriptions.

- Users encountered affiliate marketing content very differently depending on whether the content creators they view use affiliate marketing in the first place and how frequently they view videos from the same creators

Our contributions include 1) a coupon code classifier and feature detection module to detect affiliate marketing videos in real time; 2) evidence users notice and understand AdIntuition's automatic disclosures; 3) evidence of automatic disclosures' effect on users, and 4) evidence that AdIntuition works well with real videos in the wild. We have made AdIntuition publicly available as a Chrome and Firefox extension with source code on and a list of commonly occurring UTM parameters associated with affiliate marketing on GitHub[1]. Based on our findings, we make recommendations for future work including how our tool can be extended to detect, measure, and automatically disclose advertising content on other platforms in real-time including blogs, websites, and other social media platforms. AdIntuition also serves as example of how to combat misleading online content and can inform the design of similar tools to inform consumers about dark patterns, misinformation, and disinformation.

BACKGROUND AND RELATED WORK

We describe disguised advertisements, affiliate marketing, disclosures, and related work.

[1]https://github.com/mike-swart/AdIntuition/

Disguised Advertisements

Many studies examine deceptive and manipulative practices online that steer users into making decisions for the benefit of the service provider. These practices, known as 'dark patterns' encompass disguised advertisements, which at the interface level are meant to trick the user into clicking on them to force engagement [4, 15, 20]. Arguably, advertising on social media where there is no disclosure is also a form of a disguised advertisement. In our work, we provide a tool for informing users of disguised advertisements on YouTube. Tangentially, many researchers are examining online misinformation and disinformation, where users are deliberating spreading fake or false information, sometimes in aid of a political agenda [30]. Our work is also related to these studies in that AdIntuition informs users of misleading advertising content so that users can assess the content and any bias accordingly. We also have to take care that AdIntuition itself is not misleading.

Types of Advertising on Social Media

Social media advertising takes two forms. Platform-based advertising refers to ads on social media platforms that can be purchased by merchants or sellers. Examples of such advertising include Facebook ads [10] and Twitter ads [32]. Endorsement-based advertising or online endorsement by contrast refers to ads on social media platforms in which merchants or sellers engage with specific users—influencers or endorsers—to advertise their products to other users. Online endorsements occur in three forms [38]: sponsored advertisements, product sampling, and affiliate marketing. In **sponsored advertising**, a merchant or seller pays a fee to an influencer or endorser in return for endorsing their product. In **product sampling**, merchants or sellers send free products to influencers to get them to endorse and promote the product to their followers. **Affiliate marketing**—the focus of this paper—an influencer or endorser earns fees based on the sales that they generate for the product [7].

Affiliate Marketing

In affiliate marketing, first, merchants and influencers register with an affiliate marketing company, which mediates their relationship. Next, influencers drive sales to the merchant through the affiliate marketing company. This is done by placing customized URLs or website links called affiliate URLs, or coupon codes that are published by the affiliate marketing company in influencers' content. These custom links or coupon codes e.g., *'CODE FOR 5% OFF: MELON',* enable the affiliate marketing company to track sales from customers [24]. UTM query parameters in URLs are analytics tracking parameters that are used to quantify audience characteristics and to track referrals [14]. Finally, for each tracked sale, the merchant pays the influencer a portion of the sale through the affiliate marketing company. There are few studies of affiliate marketing, notably, in the security community focused on detecting cases where content creators defraud affiliate marketing programs, e.g., through setting fraudulent cookies along the tracking chain of links [28, 29, 5] or by typosquatting [18]. In our paper, we focus on how to automatically disclose in real-time when affiliate marketing content is present.

Disclosures in Online Endorsements

According to the FTC's endorsement guidelines [11], advertisers are required to disclose their relationships with merchants to consumers so that they can recognize and assess the content in an informed way. These disclosures need to comply with the *clear and conspicuous* standard so they are identifiable to consumers. Specifically, affiliate marketing disclosures need to be placed close to the endorsement and the URLs included by the content creator. Using the text *affiliate link* as a disclosure statement is insufficient; instead, using an explanatory phrase such as *I get commissions for purchases made through links in this post* is encouraged.

Although the HCI community has studied social media for many years (e.g., [35, 12, 36]), only a few studies focus on understanding how users process and deal with advertising content. For instance, there are studies of online behavioral advertising and ad-blocking tools [33, 22]. Outside of the HCI community, researchers have examined whether influencers make advertising disclosures and whether users notice and understand these disclosures [3, 37]. These studies have found that using the text *paid ad* was effective in helping users identify sponsored Instagram post [9] and that users became more resistant to bloggers' endorsements when bloggers disclosed their sponsored content.

One prior study focuses directly on measuring the prevalence of affiliate marketing content and how users perceive affiliate marketing ad disclosures. Collecting a dataset of 500K YouTube videos and $\sim$1M Pinterest pins, Mathur et al. [21] discovered that 90% of affiliate links are undisclosed on these platforms and that those that did disclose used what the authors call *affiliate link* style disclosures. They highlighted the need for tools that alert users about endorsement-based advertising. We build directly on this work to create automatic detection techniques for disclosing affiliate marketing content to users.

Automatically Identifying Advertisements

Automatically disclosing online endorsements in real-time is an under-studied phenomenon with most tools focused on identifying and blocking malicious or privacy-infringing advertisements on the web [22]. Ad-blocking extensions like AdBlock [2] and Adblock Plus [1] maintain *filter* lists of ad servers and user interface elements that correspond to advertisements, which they use to block those advertisements. Other tools detect online behavioral ads using lightweight computer vision and image processing techniques [31]. However, no such parallel—to the best of our knowledge—exists for endorsement-based advertisements. Traditionally, it has been hard to identify these advertisements because we did not understand their identifying features until the Mathur et al. [21] study. Mathur et al. [21] note that users do not always notice or understand disclosures and that most videos do not have a disclosure. Also, Mathur et al.'s dataset consists of a random sampling of videos which may not reflect a set of videos real users may watch. We built on their work to create our AdIntuition browser extension which can immediately analyze any video in real-time for affiliate marketing content using features from this work and additional ones we identified. Our work also contributes by alerting users to this content automatically and showing how users react to automatic ad disclosures for YouTube videos. Finally, in our work we collect data on and test our tool with actual videos watched by real users.

METHODS

In this section, we describe our detection techniques for affiliate marketing content on YouTube. We then describe the AdIntuition tool design and implementation. We also describe the user studies conducted to evaluate AdIntuition, a survey of 300 users on Amazon Mechanical Turk, a diary study with 11 users, and a field deployment with 472 users.

Automatically Detecting Affiliate Marketing Content

We focus on detecting affiliate marketing since affiliate marketing links, provided by a brand, are ground truth for a brand-content creator relationship. Other types of sponsored content such as native advertising do not have a similar ground truth to detect and verify brand-relationships automatically. Our goal is to expand on the set of signals that can increase confidence in detecting affiliate marketing content, focusing specifically on UTM query parameters and coupon codes as additional signals. First, we examined which UTM query parameters and coupon code patterns are associated with affiliate marketing content to verify whether these features could serve as a ground truth in detection. We used the data set of $515,999$ YouTube videos from Mathur et al. [21] and found that 1.2% of all the videos in the original data set contained either a known affiliate link pattern, UTM query parameter in a URL, or coupon code. We performed a manual inspection of these features in the video descriptions. UTM query parameters that appeared frequently included "utm_source=", "utm_term=", "campaignid=", "utm_campaign=", "utm_content=", "aff_id=", and "utm_medium=". These often contain textual values related to affiliate marketing, such as 'aff' or 'affiliate'.

URL/UTM Parameter Detection Module In AdIntuition

For each link found in a YouTube video's description, our detection module in AdIntuition checks whether the link, including any intermediate redirected link(s), matches against the known list of affiliate marketing patterns [19] or contains the UTM parameters of interest. If the link or any link in the redirect chain is found to contain one of these two features, then the relevant URL is highlighted in the video description to show that it is an affiliate marketing link.

Coupon Code Clustering and Classification

We also built a classifier to detect coupon codes in video textual descriptions.

Finding Coupon Codes In Existing Data Set: Starting with the $515,999$ videos in the Mathur et al. [21] dataset, we found that $174,885$ had descriptions that were in English. We tokenized these descriptions into $1,139,880$ individual sentences. To shrink the number of inputs to our clustering algorithm, we generated a unique set of potential coupon codes for each video based off of the content creator's channel name. For example, Casey Neistat is a famous YouTuber whose channel name is *'CaseyNeistat'*. Our dictionary generator checked the description of his videos for his full channel title, *'CaseyNeistat'*, individual components of his channel title split on lowercase to uppercase transitions or spaces, "Casey" and

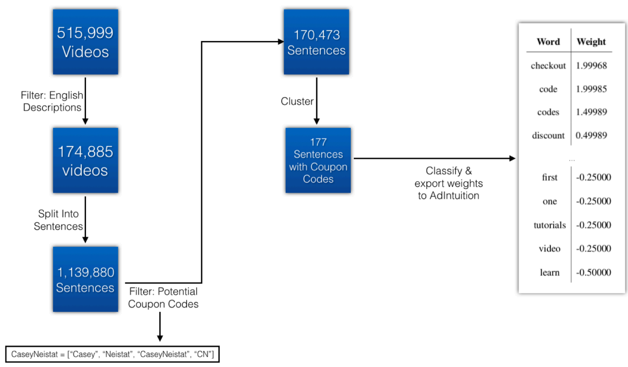

Figure 1. Clustering and Classification Funnel

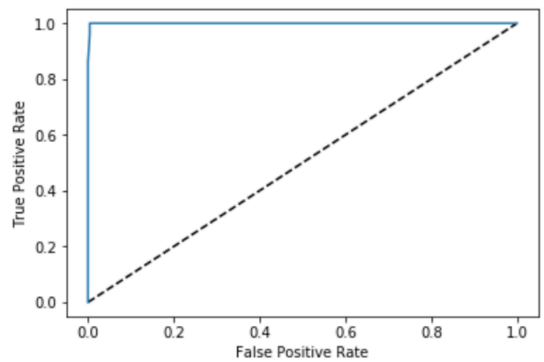

Figure 2. ROC Curve for the Coupon Code Classifier

"Neistat", and then the initials of his channel title components, "CN".

Using the dictionary generating technique, we found that 170,473 sentences contained a match. Upon closer inspection, we noted that many of these sentences did not contain coupon codes because content creators tend to use the same username across social media platforms. For example, Casey Neistat's twitter page is @CaseyNeistat and his Instagram name is @caseyneistat, both of which correspond to a case-insensitive search term that was generated in the list of possible coupon codes. We manually labelled 129 sentences as containing coupon codes but only 31.8% of these coupon codes were related to the channel name.

Clustering Sentences With Coupon Codes: We filtered the data set to sentences that might contain a coupon code using our coupon code generator. Using a bag-of-words representation of all of the sentences, we computed the cosine distance between each representation and all of the other sentences' representations. We used the cosine distance so that long sentences and short sentences would not necessarily be farther apart from each other, as can be the case with Euclidean distance calculations. The clustering algorithm automatically chooses a threshold to cluster the data into distinct sets. Using this algorithm, we clustered the sentences into 253 clusters. The first author manually inspected each cluster to note which contained most of the sentences with coupon codes which we could use to train the classifier. Figure 1 shows the funnel for clustering and classification of videos descriptions into sentences and then sentences selection. The words and weights shown in this figure are most (positive valence) or least predictive (negative valence) of a sentence containing a coupon code e.g., "checkout".

Building A Classifier To Find Coupon Codes In New Data: With a cluster of 177 sentences containing coupon codes, we then built a classifier that we could use to determine new coupon codes. We sampled 1,000 other random sentences in addition to the 177 that were in the coupon code cluster to add negative examples to our training dataset. Next, we used a bag-of-words representation for each sentence and assigned each sentence a value of 1 if it contained a coupon code and 0 otherwise. We performed three-fold cross validation to train, validate, and test the classifier [27] as described in the next subsection. For each round of training, validating, and testing, we created a vocabulary based off of the sentences in the training set and then fit the classifier to the bag-of-words representations of those sentences. We then evaluated and tuned the parameters that the classifier used with the validation set. Finally, we quantified the performance of the classifier using the testing set. The performance was evaluated using r-squared score, F1 Score, and ROC curve shown below.

Evaluating The Classifier With Existing Data: Testing different classifiers, we found that Support Vector Classifier (SVC) worked best instead of a Random Forest Classifier or a Decision Tree Classifier, both of which were also tested. We used the F1 score metric to evaluate two important aspects of the classifier: precision and recall. Precision refers to how well the classifier performs at finding sentences that actually contain a coupon code and how well it avoids false positives. Recall refers to how well the classifier is at identifying a large number of the sentences that are known to have a coupon code in them. The classifier had an F1 score of 0.992, which meant that it performed reasonably well given that the maximum score possible is 1. The r-squared value characterizes how well the model's division of the sets matches up with the data. The SVC model's r-squared value was 0.976, and the ROC Curve is shown in Figure 2.

Since the classifier used a bag-of-words representation, words were given weights for how important they were in finding coupon codes. Words such as "first" and "video" do not correlate with the presence of a coupon code, whereas "code" correlates strongly. Though trained on sentences that contained coupon codes that were related to channel names, the classification focuses on the words that surround the coupon code, not the coupon code itself. Out of 129 sentences that we manually labeled to have coupon codes, only 31.8% of them corresponded to the channel name. Meanwhile, the classifier was able to correctly identify 93.8% of the sentences with coupon code (true positive rate), many of which were not related to the channel name.

Integrating The Coupon Code Classifier Into AdIntuition: Once the classifier was trained, the vocabulary words and associated weights were exported into JavaScript code that is used in the browser extension. In the AdIntuition extension, when a YouTube video loads, the description is parsed into sentences which are transformed into a bag-of-words vector in the same way as before. This vector is then multiplied by the weights vector to get a score, which can be used to determine if a sentence has a coupon code. We used the GridSearchCV library to optimize parameters for the SVC with their suggested threshold value of 1.0 [26].

Designing AdIntuition Automatic Disclosures

We integrated the detection techniques above into our AdIntuition browser extension for Chrome and Firefox, to automatically detect and disclose affiliate marketing content on YouTube videos. We had the following design goals building on findings from prior work [21]:

1. Display a noticeable automatic disclosure (since users often do not notice ones in the video description)
2. Use clear explanatory disclosures (since these are most understood by users)
3. Highlight affiliate marketing links and coupon codes in video descriptions (to aid with alerting users)
4. Distinguish between known affiliate marketing links and ‘suspected’ affiliate marketing content based on UTM parameters and coupon codes (so users are not misled by false positives in our detection techniques)

To meet these goals, AdIntuition displays an unobtrusive but detectable banner directly above a video when it loads if the extension detects any features of affiliate marketing content. We used three colors to distinguish between clear versus known sponsorships to avoid being misleading. Users can click the AdIntuition extension for the color key. For known affiliate marketing links, the banner is pink and states: *‘This video contains affiliate marketing content. The creator may make a commission if you click on the highlighted portions of the description’*. For content AdIntuition flags based on UTM parameters or coupon codes alone, we display a yellow or orange banner respectively and change the wording to ‘contains suspected’ affiliate marketing content to reflect our confidence in the classification. Finally, AdIntuition highlights affiliate marketing links, links with UTM parameters we have flagged, and coupon codes in the video description, also using color to distinguish which flag was set.

Instrumenting AdIntuition To Log Data

We instrumented AdIntuition to log certain events of interest to an Amazon Web Services (AWS) DynamoDB linked to an API Gateway to help improve our detection techniques and measure real world usage. Using the API Gateway, AdIntuition logs data from each user to the central database. For each user, we generate a random user id that is used to distinguish between different AdIntuition extension users. This random user id allows us to promote anonymity in the data and minimize risk to individual users. To preserve privacy we also do not log any video watched that is not flagged as affiliate marketing content.

AdIntuition collects information for the following events:

- add_username: this event is logged when the user first downloads the extension and when the user id was created. User ids are generated once and then used for the rest of the time that the user uses the extension.
- vid_watch: the user watched a video. Video id is logged only if the video was flagged as containing affiliate marketing content.
- utm: a URL with UTM parameters of interest was found in the video description. The video id and matched URL are logged.
- aff: a URL with a known affiliate link from this list was found in the video description. The video id and matched URL are logged.
- coupon_code: a sentence containing a coupon code was found in the video description. The video id and sentence are logged.

AdIntuition Evaluation

We conducted three user studies, a survey, a diary study, and a field deployment to evaluate how users interact with AdIntuition, all of which were approved by our institution’s Institutional Review Board (IRB). In this paper, we only report on findings relevant to AdIntuition usage.

Study 1: Evaluating AdIntuition Interface

As part of the design process, to evaluate if the AdIntuition banner was noticeable and understandable to users, we designed an online survey. This survey contained 17 questions based on Mathur et al.’s user study of how users interpret disclosures on YouTube [21]. We recruited 300 participants on Amazon’s Mechanical Turk (MTurk) service to participate in the survey between April 24^{th} and May 1^{st}, 2019. Participants were required to be at least 18 years old, be in the United States, and have an MTurk score of 95% approval or higher. They were paid $1.25 for a maximum of 15 minutes of work, which was calculated using the minimum federal hourly wage in the U.S. of $7.25 and dividing it by 6.

Participants were asked to watch one of three YouTube videos, each of which had a known affiliate marketing link pattern and a product being marketed. For instance, one creator described a new cinnamon donut flavored cookie and provided a link to buy the cookies in the video description. All participants were split randomly into two groups. Users in the control group were shown one of the YouTube videos with no AdIntuition banner. The other group was shown one of the YouTube videos with an automatic pink disclosure and highlighted link as it would appear in AdIntuition. Note, we only tested the pink banner to ensure users noticed and understood our design in a controlled setting. We had 6 total conditions with a control and treatment group for each of three videos. The survey first asked about users social media usage. Next, participants were asked to watch their randomly assigned YouTube video and to read the video’s description. Participants were unable to continue the survey until they had stayed in this section for at least 3 minutes to ensure they watched the video. Participants

Group	Video 1	Video 2	Video 3	Total
Control	62	39	48	149
Treatment	50	50	52	151

Table 1. Number Of Participants In Each Condition (Study 2)

were asked to describe the video that they watched, rate their opinion of video content, and to provide reasons for their answers. After this section, participants were shown the product that was described in the video and asked their impression of the product. They were then asked to rate how likely they felt a relationship between the content creator and the organization selling the item existed and why. Following this step, participants were told that the video did contain affiliate marketing content and asked if a banner or a highlighted link would or did assist them in determining the relationship. Finally, we collected demographic information.

Analysis: We conducted a descriptive analysis on the survey data. In addition, we performed inductive thematic analysis [25] on the open-ended answers using a codebook of 8 codes that we created after reading participant responses and team meetings. We met regularly to discuss themes to reach consensus. Examples of codes include 'influence of the AdIntuition banner' and 'knowledge of YouTube sponsorship trends'. Each coded survey response was reviewed by at least two team members and then we wrote summaries of emerging themes for each code. After multiple team meetings, we reached consensus on the main themes from the study. We denote participants in Study 2 with the identifier *AMT* and a participant id.

Participant Demographics: 300 participants were randomly assigned to a group and had valid responses. The number of participants in aach video and group combination is shown in Table 1. The median age for all groups was 37 except for the control group for video 1 which was 40. 94% of users in each group reported using YouTube several times a week, 75% also said at least once a day. Across all 300 participants, 45% of participants were female, 54% were male, and 1% did not disclose gender. Ages ranged from 23 to 73 years old with a median of 37 years old. 97.3% of participants used YouTube at least once a week, with 57.7% reporting using it several times a day.

Study 2: Evaluating AdIntuition User Experience With Two Week Diary Study

To evaluate the user experience of AdIntuition, we conducted a diary study [17] in July-August 2019. This study allowed us to interview users about their real-world encounters with and perceptions of AdIntuition banners with all three colors and their effects on user opinions. We recruited 12 participants through online mailing lists and social media postings. We recruited participants who were regular YouTube users over 18 who lived in the United States and had access to a device and browser compatible with the extension (Google Chrome or Firefox). We defined a 'regular' YouTube user as anyone who watches videos on YouTube multiple times a week.

The diary study was divided into 3 parts: pre-study interview with installation of the tool, diary log period, and post-study interview. Each interview lasted up to 30 minutes. The interviews were conducted either in-person at our institution's campus or through video/audio call (Skype or Google Hangouts). All interviews were audio-taped. During the pre-study interview, we gathered baseline data: how often/what YouTube videos users watched, if they ever encountered a video they thought was sponsored before, why they thought that, and how it affected them. After the interview, the researcher helped the participants install the tool onto the participant's personal device, showed them how to use it, and verified it was working from the user_id generated by the extension. After this interview, participants were asked to complete diary log entries in a paper or electronic diary for the next 10 days. Each day, participants were asked to log information on the videos they watched on YouTube that contained an AdIntuition disclosure banner or note if they did not see any videos with these banners. Each diary log asked for video title and channel name, a rating of level of surprise at seeing a banner above the video indicating affiliate marketing content on a Likert scale, and reasons for the rating.

During this diary period, participants were instructed to watch videos as usual. To incentivize logging, we sent participants daily email reminders and compensated them $2 for each day that contained at least 1 diary entry (in addition to the $25 for participants in the pre and post study interviews for a total of $45 for 10 days maximum.) In the post-study interview, participants were asked about their experiences with using AdIntuition, the tool design, and their opinions on affiliate marketing on YouTube. Participants were also shown how to uninstall the tool but reminded that they could continue using the tool if they wanted beyond the study. If participants did not encounter any affiliate marketing content during the diary logging period, we showed them an affiliate marketing video with an AdIntuition disclosure banner before asking them the questions.

Analysis: All diary logs were converted to a digital format in an Excel spreadsheet for analysis. We transcribed the interviews and developed two codebooks based on the interview guide and diary instructions which we refined after reviewing several transcripts, diary logs, and team discussions. We used the same analysis process as Study 1 to perform inductive thematic analysis [25] on the transcripts and diary logs. Two co-authors went through and coded the transcripts, met weekly with the team, and refined the themes based on points of agreement. We had a total of 80 codes (70 for interviews, 10 for diary logs). Example codes for interviews included 'affiliate marketing knowledge', 'attitudes towards affiliate marketing disclosures', and 'thoughts on the AdIntuition design'. Example codes for the diaries included 'types of videos watched' and 'reasons for feeling surprised to see the banner above a video'. We denote interviews quotes in Study 2 with *P* and quotes from participants diary logs with *PDL* and the participant id.

Participant Demographics: Only 11 participants completed the diary study (7 female, 4 male) for a minimum of 10 days. The age range was 19 to 71, with more than half the participants being under 35. Participants had a variety of occupations including undergraduate and graduate students, housewives, a librarian, an administrator, a manager, a director, and a retiree. All but one participant reported watching YouTube videos daily with 1 saying they watched videos several times a week. Over the study period, 9 participants completed 108 diary log entries with a median of 10 entries per participant. 3/9 came across the banners every day and 6/9 saw the banners only on some days. Participants who saw banners during the study period saw the following types of videos: how to/educational, entertainment, vlogging/video personalities, sports/video games, product reviews, style, and music. The remaining 2 participants did not come across any videos with AdIntuition banners during the study even though they watched multiple videos a day. Therefore these participants did not have any diary logs.

Study 3: AdIntuition Field Deployment With Real Users

To evaluate how effective our classifier's efficacy with videos real users watch in real-time, we conducted a field deployment. In this study, we were unable to interview participants since we did not collect personally identifiable information to enable a large deployment. We recruited users to test out a tool 'to identify misleading ads on YouTube' through a blog on our institutional website, Twitter, and an article in the popular media. To preserve user privacy, we created a privacy policy detailing everything the tool collects. We also built functionality for users to download a report of all of the data that AdIntuition collects for them and to delete entries in that record. Finally, we added options for users to opt into or out of data logging. Overall, we had a total of 472 downloads on the Chrome and Firefox stores for the period June 5^{th}-August 26^{th}, 2019.

Analysis: We assumed users who had just downloaded AdIntuition would be interested to see how it worked and therefore would have the potential to skew our analysis. For any metric that had to do with individual usage of the extension, we removed all users who used AdIntuition for 1 day only and considered users as an active user only if they used AdIntuition for at least 1 additional day beyond the day of download. In our analysis, we focus only on these *active* users. We also removed all AdIntuition diary study users from our analysis. We only report on videos with affiliate marketing content which were verified by manual inspection and we *exclude false positives* in all graphs and data presented.

Limitations

Our survey study evaluated a simulation of the AdIntuition interface so the results may differ if we used experience sampling with our real world deployment instead (which also presents challenges of asking users to respond in context). Further, our diary log study is limited by self-report and the sample size. Finally, to preserve privacy in the field deployment, we only collected video ids for videos AdIntuition flagged as containing affiliate marketing content. Thus, we are unable to calculate false negatives for our detection techniques.

FINDINGS

Study 1 and 2: User Experiences With AdIntuition

We report on the findings across Study 1, the survey, and Study 2, the diary study.

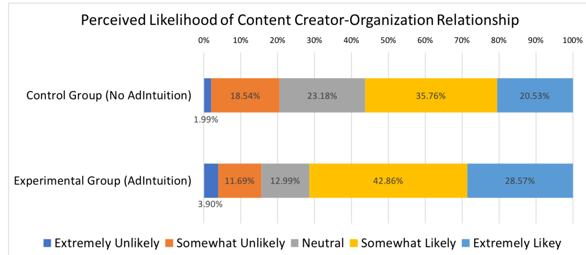

Figure 3. Perceived Likelihood Of Content-Creator-Organization Relationship In AdIntuition Survey

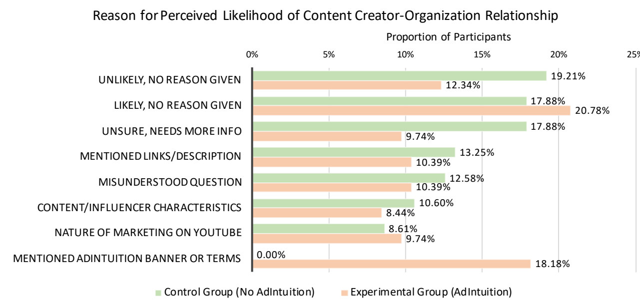

Figure 4. Reasons For Perceived Content-Creator-Organization Relationship In AdIntuition Survey

AdIntuition Helps Users Identify Affiliate Marketing Content

Our first study suggested that AdIntuition's interface helped participants to identify affiliate marketing content. As shown in Figure 3, the experimental group (with AdIntuition disclosure banner and highlighted affiliate link) was much more likely to perceive a relationship between a content creator and the organization selling merchandise in the video than in the control group which did not see an AdIntuition banner. For instance, 71.4% in the experimental group that saw the AdIntuition disclosure banner above the video and the highlighted affiliate marketing link reported that this relationship was likely or extremely likely compared to 56.3% of the participants in the control group who did not see any AdIntuition banner when they watched a video. We calculated an odds ratio of 1.941 with $p < 0.00628$ suggesting this result was not owing to random chance.

Figure 4 summarizes the reasons participants provided for their answers. Notably, 18.18% of participants in the experimental condition explicitly reported their reason for believing a relationship is likely is because of the AdIntuition banner. For instance, participant AMT196 stated *'There was a message above the video stating that it included affiliate links. More than likely this was a paid sponsorship by Chips Ahoy"* (None reported this in the control condition since they saw no banner). More participants, 17.88% , also expressed that they were unsure or needed more information to answer the question in the control condition versus 9.74% in the experimental condition. Interestingly, a similar number of participants mentioned other *implicit* indications of an affiliate marketing relationship in both conditions (32.45% control vs 28.57% experimental) such as noticing links in the video description for the products (*e.g., 'Given that a link pointed to this specific item that* *appeared to be used in the video, I'm guessing the presenter is related to DJI'* (AMT348), the behavior of the content creator, or video characteristics or expectations of marketing content on YouTube (e.g, *'It's a pretty common thing among 'influencers"* (AMT245).

Disclosures Most Unexpected If Implicit Signals Are Absent

In our second study, across all users, 26% of the total diary logs indicated that participants were surprised that a particular video had affiliate marketing, in 43% of the logs participants indicated they were not surprised to see an ad disclosure, and 31% of the entries were rated as a neutral. Often, participants logged that they were surprised to see the AdIntuition banner when the product was not aligned with the type of video or video creator. For example, one participant described being very surprised to watch a video where there was advertisement for something that contradicted the entire premise of the video: *"I didn't expect a food ad from a water fasting video, found it counterintuitive"* (P10DL6) Other reasons for being surprised included: unclear product promotion, the sponsorship not being typical of the channel or influencer, the channel having a small following, or containing little to no indication of a sponsorship in the video as summarized by Participant P3 (DL4): *"I was surprised because I didn't notice the sponsorship at all while watching the video. It was very subtle"*.

For some videos, participants did not report feeling surprised to see they contained affiliate marketing content. Their reported reasons in the interviews and diary logs largely matched the implicit indicators mentioned by participants in Study 1. In these cases, participants mentioned noticing a video was sponsored, or more implicit indicators, such as the influencer talking very favorably about a specific product in comparison to other products, talking about product very positively or for a significant amount of time of the video, the frequency of mentioning a product, a clear focus on the product, product placements disrupting the natural flow of the video, and talking 'commercial-like'. An example diary log summarizing these reasons include: *"She regularly promotes her merch that she sells (sometimes to raise money for charity). She also just released her own line of nail polish"* (P1DL26). In these logs, participants reported feeling neutral when the sponsorship made sense. Notably, the majority of participants did not understand the differences between the different color banners.

Disclosures Made Users More Reflective On Content Creators

In Study 2, we were also curious to see if participants noticed any changes in their behavior while using AdIntuition. Less than half of the participants reported no change in their YouTube watching behaviors as a result of using the extension. However, 6 of 11 participants found that using AdIntuition made them think more about sponsorships on YouTube and made them more aware and perceptive of them. Some participants even told us that they found themselves actively searching for additional videos to see which videos would have a banner, trying to see which product would be promoted in the video, and how the influencer would promote the product. As an example, participant P2 summarized: *"I did notice a difference in the way I viewed certain videos. As I was watching them I was definitely more sort of clued in, so looking for any clues as to what they were advertising. [In] some videos [it] is really obvious but some videos where it may be is more subtle then quite often, I found myself trying to watch to figure out where they were advertising stuff"*.

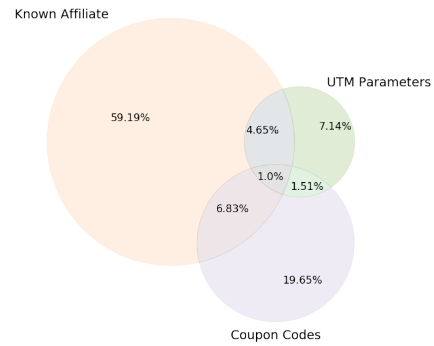

Figure 5. Total Affiliate Marketing Videos By Feature Type

All but one participant reported some change in their attitude or perception of sponsored content as a result of using AdIntuition. Some participants expressed more negative attitudes towards videos with the AdIntuition banner, such as noting that the intent of the video was unclear or that those videos were of lesser (subjective) quality in comparison to videos without banners from the same video creators. Some participants also described negative attitude changes towards the influencers, such as a decrease in trust, questioning the influencer's true intentions behind creating specific content, and feeling disappointed in the influencer, or that *'whoever was producing the video didn't have my best interests in heart'* (P11). As for the products being promoted, a few participants expressed an increase in skepticism of the legitimacy of the endorsements and true quality of the product in videos that had the AdIntuition banner. Participants also felt deceived by unexpected banners as expressed by P4: *'I felt disappointed in them even though I don't know them, but I didn't change what I watched.'*

However, other participants expressed positive attitude changes such as excitement when coming across a video with a banner, feeling more intrigued by the video, and gaining new insights into the business aspects of YouTube content creators. This is captured in a quote by Participant P2: *"The banner made me realize that it's much more of a business than just purely people just having fun. But I don't think it made me think any better or worse of the people who are making them"*. These changes contrasted with pre-study sentiments which were more neutral overall regarding sponsored content.

Study 3: Performance Of Automated Detection Methods

Next, we report findings from the field deployment, Study 3.

Multiple Features Aid Affiliate Marketing Automatic Detection

In total, 472 AdIntuition users saw a total of $60,835$ videos over the deployment period of 82 days. Of the total number of videos watched across all users, (some of which were watched more than once), 6071 were true positives and contained one of the three indicators of affiliate marketing content. Of this total

utm_source	utm_medium	utm_content
youtube (1424)	paid (326)	cta-link (94)
affiliate (360)	social (278)	[no value] (51)
open.spotify.com (75)	youtube (189)	banggoodtv (15)
blogger (51)	affiliate (188)	description (15)
youtube.com (48)	yt+main (186)	en (14)
impact_radius (40)	yt main (115)	dhalucard (14)
yt (40)	referral (113)	tw (12)
afc-linus+media+group (36)	video (97)	linustechtips (10)
vlogger (31)	description (84)	descrip (9)
refersion (31)	open (75)	youtube (9)

Table 2. Most Common UTM Parameter Values For Various UTM Query Parameters With Counts of Occurrence

number of affiliate marketing videos watched across users, there were 4494 total unique videos. Figure 5 shows the breakdown of which features these unique affiliate marketing videos were flagged on. Note, this shows cases where a video was flagged on more than one feature.[2] Of the total flagged videos (including overlap between videos with more than one feature flagged), 71.7% had at least one link with a known affiliate link pattern, 14.33% had at least one link from a domain confirmed to use UTM parameters in their affiliate marketing campaigns, and 29.02% had at least one coupon code in the video description.

UTM Parameters: In the unique affiliate marketing videos data set, we found the most commonly occurring values for each of the UTM parameters that we logged and then manually inspected this list as shown in Table 2. We then flagged the ones that appeared to relate to affiliate marketing terms such as "paid" (the most common utm_medium value), "affiliate" (the second most common utm_source and fourth most common utm_medium value), and "referral" (the seventh most common utm_medium value). We also added terms surfaced in the known affiliate marketing links list [19]. This list included the following values that we then searched for: '=aff', '=infl', 'aff_id=, 'aff=', 'sponsor','=paid', '=ref', 'ref_id=', 'promotion', 'ref=', 'referral', 'affiliate', and 'influencer'. Of the total unique links that were logged to contain UTM parameters, 33.24% contained a UTM parameter with one of these tags. These links were tagged 7386 times in total, some of them just once and some multiple times depending on the number of times a video was watched by AdIntuition users.

We filtered down the total list of 929 domains containing UTM parameters we flagged to 28 domains, each of which had at least 10 unique links present in the data set. We then manually checked whether each of these 28 domains on this list had affiliate programs by visiting their home pages to confirm if they offered such programs or in some cases, signing up on affiliate marketing company programs claiming to do affiliate marketing on behalf of these sites. 21/28 domains had affiliate marketing programs. Of the 7 remaining domains, 5 were unclear, and 2 had affiliate marketing programs that were no longer active. As a lower bound, we were thus able to verify that on these frequently occurring 28 domains, AdIntuition

[2]Only our coupon code classifier is predictive in that it can handle inputs it has not seen before based on pre-determined weights calculated in training, therefore we cannot report on predictive power for each attribute flagged.

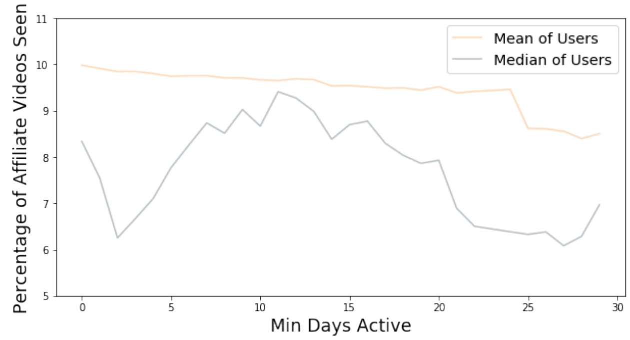

Figure 6. Affiliate Marketing Encounters Over Days Active

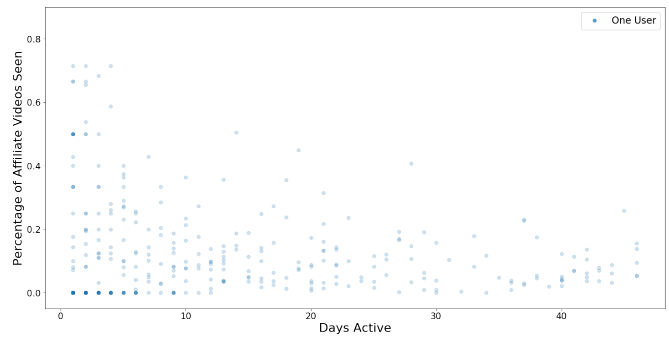

Figure 7. Individual User Encounters With Affiliate Marketing

had a true positive rate of 75% based on a UTM parameter feature alone. Less frequently occurring domains may also have affiliate marketing programs but further manual verification is required. We have made this list of commonly occurring UTM parameters and affiliate marketing domains available on GitHub to extend the known affiliate marketing patterns made available in [21].

Coupon Codes: AdIntuition flagged a total of 2539 coupon codes in the videos that had affiliate marketing content. We manually reviewed these codes and found that our classifier had falsely tagged 470 codes and that 43 codes were unclear as to whether it was a coupon code. We had a true positive rate of 81.17%, suggesting our classifier had a slightly worse true positive rate on real world data.

Affiliate Marketing Encounters Depend On Viewing Behaviors

During Study 3, the median user used the extension for 4 days and saw 2.7 videos per day. The range of total videos watched per user per day between 1 and 123. The median video views per a user was 11. Across all active users, the median percentage of affiliate marketing content seen was 7.55% of videos. As shown in Figure 6, which maps the total affiliate marketing content active users encountered, AdIntuition users who have used the extension for longer periods see slightly fewer videos with affiliate marketing than those who have only used it for a few days. However, as users interact with the extension more, the prevalence increases.

We also plotted the percentage of total affiliate marketing videos encountered by each active user over time. This is illus-

Content Creator	No Of Videos Watched Per User [User1, User2,...]
Sweet Anita	[91]
iFL TV	[43]
Tim Pool	[35]
FNG	[22]
The Alpha Male Strategies Show	[22]
Cereal Entrepreneur - Jordan Steen	[18]
LetsPlay	[18]
SAM THE COOKING GUY	[17]
Unbox Therapy	[16]
Movieclips	[15, 37]
Linus Tech Tips	[15, 14, 12, 12, 12]
Simply Nailogical	[15]
jade darmawangsa	[15]

Table 3. Most Common Content Creators And No Of Unique Videos Watched From That Creator By Various Users

trated in Figure 7, where each bubble on the graph represents one user. Noticeably, the number of users who have not seen any affiliate marketing videos (0% prevalence) stops after 11 days of AdIntuition usage. Finally, we examined the most commonly occurring content creators with more than 15 unique videos in our data set and found that the majority of them were only watched by a single AdIntuition user (As shown in Table 3, in an extreme case, 1 user watched 91 videos from 'Sweet Anita'). Our data demonstrates that users encounter affiliate marketing content very differently over time, likely owing to their viewing preferences and whether content creators they view are affiliate marketers.

DISCUSSION

We make the following recommendations based on our results.

Ethics And Automatically Disclosing Advertising Content

Creating automatic disclosure tools for flagging disguised advertisements or misleading online content in general can cross an ethical line if there are false positives which sway consumer opinions. For example, in our survey and diary study, participants sometimes viewed content more negatively when a disclosure was present. In AdIntuition, to avoid being misleading we used color and wording to convey the confidence in our classification. We need further research to understand how best to inform users about misleading content such as advertisements without having them form overly negative opinions or habituate to automatically generated disclosures. In this way, we can ensure that automated disclosure tools do not themselves perpetuate misleading information or dark patterns in cases of false positives.

Future work could examine combining our approach with crowd-sourcing to have secondary checks and balances on the information provided. Future studies could also work on incorporating a reporting tool for users to mark incorrectly flagged content. These tools should also allow influencers to view their own content with these tools and similarly report inaccurate messages. Finally, future work could also examine the effects of different types of disclosures on user opinions to mitigate negative effects. The need to present information in an unbiased informative manner is one that is common to any system that aims to inform users about misleading online content such as dark patterns, disinformation, and misinformation. Future studies can build on our work to find common guidelines for informing users about misleading content in a neutral fashion for these related domains.

Automatic And Human Aided Disclosure

AdIntuition demonstrates that automatically disclosing one type of online endorsement on social media is possible. However, our work also raises the larger discussion of who is responsible for creating these disclosures about misleading online content? Browsers could integrate automatic disclosure tools into their capabilities so that users can see advertising content more easily across the web. Platforms could similarly integrate this approach or affiliate marketing companies could more strongly require that influencers disclosure their relationships. We do see a place for third party automatic disclosure tools too for informing regulators about the prevalence of this type of content in the wild, how often disclosures are happening, and how users react to these disclosures.

An open question is how to maintain a tool for automatic disclosures? Much like an ad blocker [22], lists of known affiliate marketing programs are constantly shifting and changing so our data set needs to be continuously curated and expanded upon. Crowd-sourcing is one method that has been used successfully in AdBlockers, this technique could similarly add value to detection and automatic disclosure tools such as AdIntuition. AdIntuition can also constantly be improved as more data is gathered with which to refine its classifier and incorporate other classifiers and detection methods.

Detecting Affiliate Marketing On Other Platforms

Our results show that we can automatically detect affiliate marketing content on YouTube using known affiliate marketing link patterns, coupon codes, and UTM parameters. Future work could extend our approach to automatically detect and disclose affiliate marketing content on other social media platforms or blogs. Our user studies suggest automatic detection and disclosures are useful since they do not rely on content creators to self disclose this information. Another area for future research concerns when to show an automatic disclosure to a user. For instance, participants in the diary study suggested seeing automatic disclosures when searching for content could be useful in certain scenarios such as searching for reviews of a product one is about to buy. Future work could investigate how user experience varies when automatic disclosures are shown at different user decision making points.

CONCLUSION

In this paper, we presented new ways to detect and measure affiliate marketing content on YouTube and a tool for automatically disclosing this content to users on the platform. We also presented findings from evaluating the tool, AdIntuition which suggest that the detection techniques are performing reasonably well and that users are able to better identify advertising content with AdIntuition's automatic disclosures. Based on our findings, we recommend that future studies extend our work to build more robust online automatic ad detection tools to keep users informed about the content they are viewing.

ACKNOWLEDGMENTS

We thank our participants.

REFERENCES

[1] 2019. Adblock Plus. `https://adblockplus.org`. (2019).

[2] AdBlock. 2019. AdBlock. https://getadblock.com. (2019).

[3] Sophie C Boerman, Eva A Reijmersdal, and Peter C Neijens. 2012. Sponsorship disclosure: Effects of duration on persuasion knowledge and brand responses. *Journal of Communication* 62, 6 (2012), 1047–1064.

[4] Harry Brignull. 2018. Dark Patterns. https://darkpatterns.org/. (2018). Accessed March 12, 2019.

[5] Neha Chachra, Stefan Savage, and Geoffrey M Voelker. 2015. Affiliate crookies: characterizing affiliate marketing abuse. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*. ACM, 41–47.

[6] Amy Chozick. 2019. Keeping Up with the Kardashian Cash Flow. *New York Times* (March 2019). https://www.nytimes.com/2019/03/30/style/kardashians-interview.html?module=inline

[7] Federal Trade Commission. 2017. What's affiliate marketing? Should I care? https://www.consumer.ftc.gov/blog/2017/09/whats-affiliate-marketing-should-i-care. (2017).

[8] Federal Trade Commission. Accessed 2018. Enforcement Policy Statement on Deceptively Formatted Advertisements. https://www.ftc.gov/system/files/documents/public_statements/896923/151222deceptiveenforcement.pdf. (Accessed 2018).

[9] Nathaniel J. Evans, Joe Phua, Jay Lim, and Hyoyeun Jun. 2017. Disclosing Instagram Influencer Advertising: The Effects of Disclosure Language on Advertising Recognition, Attitudes, and Behavioral Intent. *Journal of Interactive Advertising* 17, 2 (2017), 138–149. DOI: http://dx.doi.org/10.1080/15252019.2017.1366885

[10] Facebook. 2019. Facebook Ads. https://www.facebook.com/business/ads. (2019).

[11] FTC. 2018. The FTC's Endorsement Guides: What People Are Asking. https://www.ftc.gov/tips-advice/business-center/guidance/ftcs-endorsement-guides-what-people-are-asking. (2018).

[12] Eric Gilbert and Karrie Karahalios. 2009. Predicting tie strength with social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 211–220.

[13] Google. 2018. Paid product placements and endorsements. https://support.google.com/youtube/answer/154235?hl=en. (2018).

[14] Google. 2019. Collect campaign data with custom URLs. https://support.google.com/analytics/answer/1033863. (2019).

[15] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 534, 14 pages. DOI: http://dx.doi.org/10.1145/3173574.3174108

[16] V Kumar and Rohan Mirchandani. 2012. Increasing the ROI of social media marketing. *MIT sloan management review* 54, 1 (2012), 55.

[17] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2010. *Research Methods in Human-Computer Interaction*. Wiley Publishing.

[18] Victor Le Pochat, Tom Van Goethem, and Wouter Joosen. 2019. A Smörgåsbord of Typos: Exploring International Keyboard Layout Typosquatting. In *Proceedings of the 4th International Workshop on Traffic Measurements for Cybersecurity*. IEEE Computer Society Press.

[19] Arunesh Mathur. 2018. Known Affiliate Marketing Link Patterns. https://github.com/aruneshmathur/affiliate-marketing-disclosures/blob/master/affiliate_markting_links.txt. (2018).

[20] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 81 (Nov. 2019), 32 pages. DOI:http://dx.doi.org/10.1145/3359183

[21] Arunesh Mathur, Arvind Narayanan, and Marshini Chetty. 2018a. Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 119 (Nov. 2018), 26 pages. DOI:http://dx.doi.org/10.1145/3274388

[22] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018b. Characterizing the Use of Browser-Based Blocking Extensions To Prevent Online Tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 103–116. https://www.usenix.org/conference/soups2018/presentation/mathur

[23] NYTimes. 2017. Fyre Festival, a Luxury Music Weekend, Crumbles in the Bahamas. https://www.nytimes.com/2017/04/28/arts/music/fyre-festival-ja-rule-bahamas.htmlt. (2017).

[24] Richard L Oliver and Mikhael Shor. 2003. Digital redemption of coupons: Satisfying and dissatisfying effects of promotion codes. *Journal of Product & Brand Management* 12, 2 (2003), 121–134.

[25] Johnny Saldañna. 2013. *The Coding Manual for Qualitative Researchers* (2nd ed ed.). SAGE, Los Angeles.

[26] scikit learn developers. 2007-2019a. *sklearn.model_selection.GridSearchCV*.

[27] scikit learn developers. 2007-2019b. *sklearn.model_selection.KFold*.

[28] Peter Snyder and Chris Kanich. 2015. No Please, After You: Detecting Fraud in Affiliate Marketing Networks.. In *WEIS*.

[29] Peter Snyder and Chris Kanich. 2016. Characterizing fraud and its ramifications in affiliate marketing networks. *Journal of Cybersecurity* 2, 1 (2016), 71–81.

[30] Kate Starbird, Ahmer Arif, and Tom Wilson. 2019. Disinformation As Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 127 (Nov. 2019), 26 pages. DOI: http://dx.doi.org/10.1145/3359229

[31] Grant Storey, Dillon Reisman, Jonathan Mayer, and Arvind Narayanan. 2017. The future of ad blocking: An analytical framework and new techniques. *arXiv preprint arXiv:1705.08568* (2017).

[32] Twitter. 2019. What are Promoted Tweets? https://business.twitter.com/en/help/overview/what-are-promoted-tweets.html. (2019).

[33] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. ACM, 4.

[34] Eva A van Reijmersdal, Marieke L Fransen, Guda van Noort, Suzanna J Opree, Lisa Vandeberg, Sanne Reusch, Floor van Lieshout, and Sophie C Boerman. 2016. Effects of disclosing sponsored content in blogs: How the use of resistance strategies mediates effects on persuasion. *American Behavioral Scientist* 60, 12 (2016), 1458–1474.

[35] Aditya Vashistha, Abhinav Garg, Richard Anderson, and Agha Ali Raza. 2019. Threats, Abuses, Flirting, and Blackmail: Gender Inequity in Social Media Voice Forums. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 72, 13 pages. DOI: http://dx.doi.org/10.1145/3290605.3300302

[36] Elham Vaziripour, Devon Howard, Jake Tyler, Mark O'Neill, Justin Wu, Kent Seamons, and Daniel Zappala. 2019. I Don'T Even Have to Bother Them!: Using Social Media to Automate the Authentication Ceremony in Secure Messaging. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 93, 12 pages. DOI:http://dx.doi.org/10.1145/3290605.3300323

[37] Bartosz W Wojdynski. 2016. The deceptiveness of sponsored news articles: How readers recognize and perceive native advertising. *American Behavioral Scientist* 60, 12 (2016), 1475–1491.

[38] Katrina Wu. 2016. YouTube marketing: Legality of sponsorship and endorsements in advertising. *JL Bus. & Ethics* 22 (2016), 59.

[39] David Zax. 2017. The War To Sell You A Mattress Is An Internet Nightmare. https://www.fastcompany.com/3065928/sleepopolis-casper-bloggers-lawsuits-underside-of-the-mattress-wars. (2017).