

An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest

Arunesh Mathur, Arvind Narayanan, Marshini Chetty
Princeton University
{amathur, arvindn, marshini}@cs.princeton.edu

Abstract—While disclosures relating to various forms of Internet advertising are well established and follow specific formats, endorsement marketing disclosures are often open-ended in nature and written by individual publishers. Because such marketing often appears as part of publishers’ actual content, ensuring that it is adequately disclosed is critical so that end-users can identify it as such. In this paper, we characterize disclosures relating to affiliate marketing—a type of endorsement based marketing—on two popular social media platforms: YouTube & Pinterest. We find that only roughly one-tenth of affiliate content on both platforms contains disclosures. Based on our findings, we make policy recommendations geared towards various stakeholders in the affiliate marketing industry, highlighting how both social media platforms and affiliate companies can enable better disclosure practices.

I. INTRODUCTION

Advertising on the Internet appears in various forms: as embedded content on websites and emails (e.g. Google AdSense served on The New York Times website), as native advertising on websites (e.g. Twitter’s *promoted* tweets), and as products endorsed by individual users—often termed *content publishers* or *influencers*—on user-generated content websites such as social media and blogs.

To ensure that users can identify advertisements as such, several regulations require advertisers to place disclosures on or around the advertisements. For instance, AdChoices [1] is a self-regulatory program followed by advertisers in the United States (US), Canada and Europe, which requires them to inform users that data about them and their activities is used to serve tailored advertisements. Similarly, social networks such like Facebook and Twitter tag native advertisements as *sponsored* and *promoted* respectively to distinguish them from regular user-generated content.

While both these kinds of disclosures are largely standardized by each advertising firm or body—whether that be that in the form of text or icons—endorsement advertising disclosures by individual content publishers are largely un-standardized and open-ended in nature. In the United States, these disclosures are guided by the Federal Trade Commission (FTC), which requires that publishers clearly and unambiguously disclose their relationship with merchants and brands. In recent times, the FTC has lodged several cases against publishers who have failed to adequately disclose their relationships with brands and merchants [2], [3], [4].

However, while such regulatory authorities provide guidelines for publishers to follow, it is unclear how many do,

and for the ones that do disclose, whether they follow the guidelines put forward. Ensuring that publishers follow these disclosure guidelines is particularly important, since content that they advertise often appears together with their original content, and may therefore be harder for users to identify.

In this paper, we specifically examine disclosures accompanying *affiliate marketing*—an endorsement marketing strategy that pays affiliates (the content publishers) money when users click on their customized URLs—on YouTube and Pinterest, two popular social media platforms. We investigate affiliate marketing’s growth over the years on these platforms, the kind of affiliates that engage in affiliate marketing, the companies that enable affiliate marketing, and finally, how affiliates disclose their relationship with these companies to end-users.

We have three main findings. First, we find that—on both YouTube and Pinterest—content with affiliate URLs have significantly higher user engagement compared to content that does not contain affiliate URLs. Second, we find that affiliate marketing disclosures appear in three distinct formats—which we term *Affiliate Link* disclosures, *Explanation* disclosures, and *Support Channel* disclosures—as opposed to one standardized format. Third, we find that the overall prevalence of these disclosures is low: only $\sim 10\%$ of all affiliate content on each platform has accompanying disclosures.

Based on our findings, we analyze how the disclosures we discovered comply with existing FTC disclosure guidelines. We forge policy recommendations for social media platforms and affiliate marketing firms to consider in order to enable affiliates to disclose such marketing relationships easily and clearly. We also outline directions for future work.

II. AFFILIATE MARKETING

Affiliate marketing primarily involves three entities: the affiliate (sometimes called publisher), the merchant and the affiliate network, and consists of an agreement between the affiliate and the affiliate network, and the affiliate network and the merchant. As illustrated in Figure 1, affiliate marketing typically works as follows:

- 1) Merchants, affiliates sign up with an affiliate network
- 2) Affiliates associate with merchants they wish to promote through the affiliate network
- 3) Affiliate networks publish customized URLs for affiliates to distribute along with their content

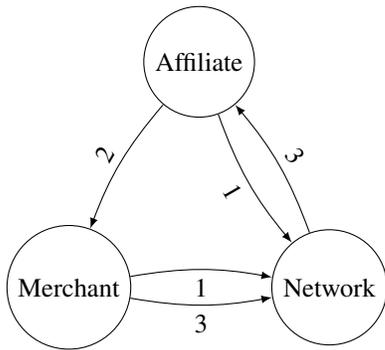


Fig. 1. An Overview Of Affiliate Marketing. (1) Affiliates and Merchants Register With an Affiliate Network. (2) Affiliates Drive Traffic to Merchants through the Affiliate Network. (3) Merchants Pay the Affiliate Network For Each Sale, Who In Turn Pay the Affiliate.

- 4) Every time a sale is made through the URLs, the merchant pays the affiliate a cut of the sale through the affiliate network

In the United States, the FTC’s endorsement guidelines [5] have largely described the disclosure standards for affiliates to maintain. The FTC states 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*. Specifically towards affiliate marketing, the guidelines state that disclosures need to be placed close to the recommendation, and not buried inside an *About Us* or *Terms of Service* pages.

While previous research has examined fraud in affiliate marketing [6], and how affiliate programs drive spam revenue and sale of unauthorized items [7], [8], we know very little about affiliate marketing disclosures. In this work, we examine the prevalence of affiliate marketing disclosures, and whether they are in compliance with the FTC’s guidelines.

III. METHOD

In this section, we describe our methodology: our data collection process, and the techniques we employed to identify the affiliate URLs and disclosures.

A. Data Collection

We examined affiliate marketing on YouTube and Pinterest, two social media platforms that are designed to share reviews and content. To ensure that our investigation uncovered the most accurate state of affiliate marketing and its disclosures on these platforms, we wanted to gather the least possibly biased sample of videos and pins. One strategy to collect videos and pins is to sample from the *related* videos and pins graphs on the network. Sampling from such large graphs has been extensively studied in the literature [9], with each approach having its benefits and drawbacks. For instance, random walks have been shown to produce less biased samples compared to breadth-first and depth-first searches.

However, because *related* 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 [10]—they result in non-uniform samples. Another strategy is using keyword searches, where videos are sampled by selecting results arising from specific search terms. However, such samples are also likely to be biased towards the search terms and the platforms’ search algorithms.

To collect a more representative and less biased sample, we employed prefix sampling which has previously been used for sampling from YouTube [10] and Pinterest [11]. Prefix sampling works by gathering, ahead of time, part of the identifier of each record called the *prefix*, which is then used to sample records beginning with that prefix. If the prefixes are uniformly gathered or generated, then the resulting samples will be uniform too. This sampling methodology is particularly useful when the search space of identifiers is significantly larger than the number of already issued identifiers, and it is not possible to generate issued identifiers randomly.

On YouTube, videos are assigned an eleven character long identifier. In order to gather a uniform sample, we first randomly generated video identifiers of length five. We then searched for those five character long identifiers using the YouTube Search API [12]. This, by way of a feature of the API, 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. To gather the prefixes, we retrieved twenty five pins from the Pinterest *Categories* page [13], and sampled their related pins randomly. We then varied the last five digits from zero to ten thousand to retrieve the 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 their categories, view counts, comment counts, and details about their creators. We collected this data between August and September 2017.

B. Data Analysis

1) *Discovering Affiliate URLs*: After sampling the videos and pins, we began identifying affiliate URLs from our corpus. First, we gathered a list of all URLs from the descriptions of the YouTube videos and the URLs of the Pinterest pins. Next, we resolved each URL following both server-side and client-side redirects (HTTP 3XX, Meta refresh), and recorded the resulting URLs and their HTTP response codes.

To identify affiliate URLs from this set of resolved URLs, we relied on the observation that affiliate URLs contain predictable patterns. 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. While Amazon’s affiliate URL appears at the end of a redirect—to Amazon’s website—affiliate URLs may also appear during the intermediate redirects. 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 discovered all such cases, 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 those 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 merchant/affiliate 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 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. We also queried the FMTC affiliate database [14], and where possible, signed up on these programs as affiliates to validate our findings.

2) *Characterizing Disclosures*: Once we finalized 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. 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 [15] 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.

IV. FINDINGS

Across YouTube, we discovered a total of 0.67% or 3,472 of 515,999 videos, and across Pinterest, a total 0.85% or 18,237 of 2,140,462 pins contained at least one affiliate URL. In this section, we present our findings, describing the characteristics of affiliate content on both platforms, the types of disclosures we discovered, and the characteristics of these disclosures.

TABLE I
PERCENTAGE OF AFFILIATE CONTENT BY CATEGORY ON YOUTUBE AND PINTEREST. ONLY TOP 10 CATEGORIES ARE LISTED.

YouTube Category	Perc.	Pinterest Category	Perc.
Science & Technology	3.61	Women’s Fashion	4.62
Howto & Style	3.49	Products	2.21
Travel & Events	1.93	Hair & Beauty	2.04
Film & Animation	1.59	Sports	1.63
Shows	1.36	Design	1.53
Music	0.94	Outdoors	1.21
Entertainment	0.68	Technology	1.16
Education	0.64	Men’s Fashion	1.12
Gaming	0.63	Animals	1.00
People & Blogs	0.39	<i>No Category</i>	0.92

A. Affiliate Marketing Companies

Table VI 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 URL resolutions. Within each platform, Amazon’s Amazon Associate Program¹ had the largest presence (YouTube = 7,308, Pinterest = 7,368), closely followed by AliExpress’ Affiliate Program² (YouTube = 2,167, Pinterest = 785). We also discovered certain merchants hosted in-house affiliate programs, as opposed to explicitly redirecting through an affiliate marketing company. For instance, Booking.com³ and Apple⁴ marketed products through their own affiliate programs.

B. Characteristics of Affiliate Content

1) *Affiliate Content Categories*: We found that across YouTube and Pinterest there existed several similarities in the categories with highest affiliate content. Overall, affiliate content on both platforms was dominated by affiliates working with fashion, beauty, and style merchants. Prevalence across YouTube’s *Howto & Style* category stood at 3.49%, and Pinterest’s *Women’s Fashion* and *Hair & Beauty* categories stood at 4.62% and 2.04% respectively. Similar patterns existed in the travel category, with YouTube’s outdoors *Travel & Events* at 1.93%, and Pinterest’s *Outdoors* at 1.21%.

There also existed dissimilarities. For instance on YouTube, affiliate content was most popular in the *Science & Technology* category (3.61%) whereas Pinterest’s *Technology* category stood at 1.16%. The list of categories ordered by their affiliate content is presented in Table I.

2) *Affiliate Content Engagement Metrics*: In addition to examining the prevalence of affiliate content by content category, we examined how affiliate and non-affiliate content (content without affiliate URLs) correlated with engagement metrics such as view, like, and comment counts. We conducted Mann–Whitney *U* tests to assess statistical significance, correcting for multiple testing using the Bonferroni method.

Across both YouTube and Pinterest, we noted a common thread: affiliate content correlated with higher engagement

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/>

metrics. On YouTube, affiliate content was longer in duration ($U \sim 7.95 \times 10^8$, $p < 0.0001$), had higher view counts ($U \sim 7.72 \times 10^8$, $p < 0.0001$), had higher like counts ($U \sim 6.96 \times 10^8$, $p < 0.0001$), and had higher dislike counts ($U \sim 6.52 \times 10^8$, $p < 0.0001$). Similarly, on Pinterest, affiliate content had 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 the pins in our dataset.

C. Types of Affiliate Disclosures

Across YouTube and Pinterest, we discovered that 10.49% of all affiliate videos and 7.03% of all affiliate pins contained an affiliate marketing disclosure. In total, we discovered three distinct types of disclosures; Table II summarizes our findings.

1) *“Affiliate Link” Disclosures*: The first type of affiliate disclosure we discovered 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. Disclosures of this type were present in 7.02% of all affiliate videos. The following statements describe them:

- *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, these disclosures appeared in similar formats in the description of the pins, and were present on 4.60% of all affiliate pins. Unlike YouTube, these disclosures did not point to specific URLs, since the pins only contained one URL: the URL actually pinned. The following statements describe them:

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

2) *Explanation Disclosures*: The second kind of affiliate disclosure we discovered offered users a verbose explanation about affiliate marketing and affiliate URLs, and how clicking on affiliate URLs impacts users. Relative to *Affiliate link* disclosures, these types of disclosures were more detailed, and often quoted specific merchants or affiliate marketing companies. On YouTube, these disclosures were present in 1.82% of all affiliate videos. The following statements describe them:

- *This video contains affiliate links, which means that if you click on one of the product 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](http://beautylish) & [MUG](http://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*

TABLE II
PREVALENCE OF AFFILIATE DISCLOSURES ON YOUTUBE AND PINTEREST.

Disclosure Type	YouTube Perc.	Pinterest Perc.
<i>Affiliate Link</i> Disclosures	7.02	4.60
<i>Explanation</i> Disclosures	1.82	2.43
<i>Channel Support</i> Disclosures	2.44	NA

TABLE III
PREVALENCE OF DISCLOSURES BY CATEGORY ON YOUTUBE. ONLY CATEGORIES WITH MORE THAN 100 SAMPLES ARE LISTED. ONLY THE TOP 10 CATEGORIES ARE LISTED.

YouTube Category	Prevalence (Perc.)	Affiliates Disclosing
Howto & Style	23.68	59 of 246
Gaming	11.22	21 of 152
Entertainment	11.21	21 of 207
Science & Technology	9.72	13 of 133
People & Blogs	8.87	40 of 443
Film & Animation	4.86	9 of 143
Travel & Events	4.12	7 of 150
Music	1.77	6 of 366

On Pinterest, these disclosures appeared in similar formats in the description of the pins, and were present in 2.43% of all affiliate pins. The following statement describe them:

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

3) *Channel Support Disclosures*: The third kind of affiliate disclosures we discovered communicated to users they would be supporting the channel by clicking on the affiliate URLs, without explaining how exactly. These disclosures appeared exclusively on YouTube, and were present in 2.44% of all affiliate videos. The following statements describe them:

- *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 (NA): <http://amzn.to>*

D. Disclosure Characteristics

1) *Prevalence by Content Category*: We examined how the disclosures varied by content category on both YouTube and Pinterest to investigate whether certain categories of content—and therefore, merchants and affiliate marketing companies—are more or less likely to disclose than others. Rather than examining the prevalence of each type of disclosure, we considered all content that had at least one of the three types of disclosure. We also limited our analysis to those content categories for which we had at least 100 samples in our dataset. Table III and Table IV summarize our findings.

Overall, we found that the coverage of affiliate disclosures across content categories varied significantly, with certain categories like Pinterest’s *Animals* as high as 60%, and others such as YouTube’s *Music* as low as 1.77%. On YouTube, the highest disclosure coverage occurred in the *Howto & Style* category with 23.68%.

TABLE IV
PREVALENCE OF DISCLOSURES BY CATEGORY ON PINTEREST. ONLY CATEGORIES WITH MORE THAN 100 SAMPLES ARE LISTED. ONLY THE TOP 10 CATEGORIES ARE LISTED.

Pinterest Category	Prevalence (Perc.)	Affiliates Disclosing
Animals	60.67	26 of 71
Food & Drink	14.08	7 of 68
Health & Fitness	12.76	5 of 45
Hair & Beauty	11.37	20 of 238
Kids	8.33	6 of 74
Home Decor	7.32	18 of 248
DIY Crafts	5.23	14 of 162
Women’s Fashion	4.21	39 of 960

TABLE V
PREVALENCE OF DISCLOSURES BY YEAR ON YOUTUBE AND PINTEREST.

Year	YouTube Perc.	Pinterest Perc.
2013	9.80	0.00
2014	8.74	0.00
2015	12.61	0.00
2016	12.36	1.93
2017	8.99	11.20

2) *Prevalence by Year*: In addition to examining the distribution of affiliate disclosures by content category, we also investigated how the affiliate disclosure coverage varied over the past years. Again as before, we only restricted our analyses to those years—across both YouTube and Pinterest—that contained a minimum of 100 samples. We examined the overall rate of disclosures and did not control for possible confounders such as the types of disclosure, the category of content, or the affiliate’s country since the prevalence of disclosures was low. Table V lists our findings.

Overall, across both platforms we found that the rate of disclosures by year remained mostly steady. Unlike Pinterest, where we noticed no disclosure coverage until 2016, YouTube had disclosures first appear in 2013 in our dataset.

V. DISCUSSION

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

A. Have the FTC’s Endorsement Guidelines Been Effective?

In the United States, the FTC first updated its affiliate marketing guidelines in 2013, stating that affiliate links embedded inside reviews must be disclosed to end-users so they can decide how much weight to provide to the publisher’s endorsement. In fact, in its current version of the guidelines, the FTC highlights that simply stating *affiliate link* is not enough, as users may not understand what that means. Instead, the FTC recommends using a short phrase such as *I get commissions for purchases made through links in this post*—similar to the *Explanation* disclosure—close to the recommendation.

Concerningly, our results show that the overall prevalence of affiliate disclosures is low, and that the disclosures are largely of the variety the FTC specifically advocates against: the *Affiliate Link* disclosures. In fact, *Explanation* disclosures—which

the FTC recommends—only appear in 1.82% and 2.43% of affiliate content on YouTube and Pinterest respectively.

This presents an opportunity for further examination: why is it that despite the FTC’s request for specific disclosure requests in affiliate marketing, affiliates fail to follow the guidelines? Is it because the affiliates are unaware that they need to disclose, or is it because they unaware of the FTC’s specific guidelines? Future work could examine this by the means of surveys and detailed interviews with affiliates.

Our work further highlights why having disclosures in affiliate content is important: affiliate content tends to have higher user engagement metrics, which means they are likely to be picked up by generic recommendations algorithms, and shown to users via search or otherwise.

B. Role of Stakeholders in Enabling Disclosures

1) *Role of Social Media Platforms*: Along with affiliates, social media platforms play a critical role in shaping the disclosures. Affiliates’ disclosures, generally speaking, are limited by the character space available to them. For instance, the description length can be as long as 5000 characters on YouTube, but on Pinterest it is capped to 500 characters. Similarly, tweets can only be as long as 280 characters on Twitter. Therefore, social media platforms can help design their interfaces to make it easier for affiliates to disclose without crowding their promotion text. For instance, 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 [16]. Similarly, YouTube also added the ability for create a *Contains Product Placement* overlay to their videos [17].

Such disclosure tools are a step in the right direction, however it is unlikely that such blanket disclosures will cover all marketing strategies. Future work could investigate what kind of affordances should be designed into social media platforms to enable affiliates to disclose clearly, and such that users fail to miss the disclosures.

2) *Role of Affiliate Marketing Companies*: Affiliates could also be held accountable to better disclosure practices by the affiliate marketing companies they sign up for. 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 Amazon⁵ and ShopStyle⁶ explicitly referenced the FTC’s guidelines in their affiliate terms. While we did not find any references in the Rakuten Marketing and ShareASale affiliate terms, we noted both companies blogging about the guidelines on their company blogs^{7,8}. However, its unclear how many other programs follow this practice,

⁵<https://affiliate-program.amazon.com/help/operating/agreement>

⁶<https://www.shopstylecollective.com/terms>

⁷<http://blog.shareasale.com/2017/09/08/ftc-updates-and-faq-s/>

⁸<https://blog.marketing.rakuten.com/topic/ftc-disclosure-guidelines>

what standards they hold their affiliates accountable to, and whether they explicitly point their affiliates to the FTC’s guidelines. Future work could also examine ways in which affiliate marketing companies can reward or penalize those affiliates who fail to adequately disclose.

3) *Role of Web Browsers*: Finally, Web browsers could help increase transparency into affiliate marketing advertising practices by means of in-built support or through add-ons and extensions. 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 the disclosures when they are present. Machine learning based approaches naturally lend themselves to developing such tools as various models could be trained on large datasets of social media content.

C. Future Work

In future work, we hope to take three specific directions. First, we hope to validate the FTC’s guidelines on affiliate marketing, and verify that the suggestions they put forward are actually effective in practice. That is, do end-users interpret the disclosures expected by the FTC correctly, and if not, how can we better improve them? Second, our findings show that the three categories of affiliate disclosures often use wordings that exhibit a common pattern. We hope to build a tool—such as a browser extension—that uses machine learning to detect these patterns automatically and highlights them to end-users. Such a tool will be able to help end-users better identify the disclosures that are often buried in the content they view. Third, we hope to examine how the disclosures we discovered align with the disclosure guidelines put forward by agencies in other countries such as the Advertising Standards Authority (ASA) in the United Kingdom.

VI. LIMITATIONS

Our study has two primary limitations. First, we only examined for affiliate marketing disclosures in the description boxes of the YouTube videos and Pinterest pins. Affiliates may also have disclosed their affiliate relationships during the course of the YouTube video, or on the image containing the Pinterest pin. In an initial analysis of 20 randomly selected affiliate videos and pins, we found no disclosures in these locations, and instead focused our attention to the descriptions. In future work, we will expand our analysis to cover both these cases more thoroughly. Second, our method to discover affiliate links should be considered as a lower bound on the number of affiliate marketing programs, since we may have missed the less prevalent companies owing to our frequency analysis. We also did not consider those affiliate programs that use coupon codes to track sales. In future work, we will expand our analysis to encompass a greater set of affiliate programs.

VII. CONCLUSION

In this paper, we analyzed affiliate marketing based disclosures on two social media platforms: YouTube and Pinterest. We found that disclosures on these platforms fall into three

categories, and the overall rate of disclosure is low. Our findings provide a starting point for several policy based discussions.

VIII. ACKNOWLEDGEMENTS

Our research was funded by Princeton University. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Princeton University.

REFERENCES

- [1] AdChoices, “Adchoices,” <http://youradchoices.com/>, 2018.
- [2] AdAge, “Ftc cracking down on social influencers’ labeling of paid promotions,” <http://adage.com/article/digital/ftc-cracking-social-influencers-labeling-promotions/305345/>, 2016.
- [3] FTC, “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>, 2016.
- [4] L. Fair, “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>, 2016.
- [5] FTC, “The FTC’s Endorsement Guides: What People Are Asking,” <https://www.ftc.gov/tips-advice/business-center/guidance/ftc-endorsement-guides-what-people-are-asking>, 2018.
- [6] N. Chachra, S. Savage, and G. M. Voelker, “Affiliate crookies: characterizing affiliate marketing abuse,” in *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*. ACM, 2015, pp. 41–47.
- [7] D. McCoy, A. Pitsillidis, J. Grant, N. Weaver, C. Kreibich, B. Krebs, G. Voelker, S. Savage, and K. Levchenko, “Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs,” in *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)*. Bellevue, WA: USENIX, 2012, pp. 1–16. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity12/technical-sessions/presentation/mccoy>
- [8] C. Kanich, N. Weaver, D. McCoy, T. Halvorson, C. Kreibich, K. Levchenko, V. Paxson, G. M. Voelker, and S. Savage, “Show me the money: Characterizing spam-advertised revenue.” in *USENIX Security Symposium*, 2011, pp. 15–15.
- [9] J. Leskovec and C. Faloutsos, “Sampling from large graphs,” in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2006, pp. 631–636.
- [10] J. Zhou, Y. Li, V. K. Adhikari, and Z.-L. Zhang, “Counting youtube videos via random prefix sampling,” in *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, ser. IMC ’11. New York, NY, USA: ACM, 2011, pp. 371–380. [Online]. Available: <http://doi.acm.org/10.1145/2068816.2068851>
- [11] E. Gilbert, S. Bakhshi, S. Chang, and L. Terveen, ““i need to try this”?: A statistical overview of pinterest,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI ’13. New York, NY, USA: ACM, 2013, pp. 2427–2436. [Online]. Available: <http://doi.acm.org/10.1145/2470654.2481336>
- [12] Google, “Add YouTube functionality to your app,” <https://developers.google.com/youtube/v3/>, 2018.
- [13] Pinterest, “Categories,” <https://www.pinterest.com/categories/>, 2018.
- [14] FMTC, “Fmte merchant hub,” <http://bit.ly/2dqgQTz>, 2018.
- [15] L. Rokach and O. Maimon, *Clustering Methods*. Boston, MA: Springer US, 2005, pp. 321–352. [Online]. Available: https://doi.org/10.1007/0-387-25465-X_15
- [16] A. Ha, “Instagram is testing a new way for celebrities and influencers to identify their sponsored posts,” <https://techcrunch.com/2017/06/14/instagram-sponsored-posts/>, 2017.
- [17] Google, “Paid product placements and endorsements,” <https://support.google.com/youtube/answer/154235?hl=en>, 2018.

APPENDIX

TABLE VI
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	http://www.awin1.com/...&awinaffid=...	129	211
	zanox	http://ad.zanox.com/ppc/?...		
	zenaps	http://www.zenaps.com/relick.php?...		
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	http://www.anrdoezrs.net/click-[0-9]+[0-9]+...	341	2413
	dotomi	http://cj.dotomi.com/...		
	dpbolvw	http://www.dpbolvw.net/click-[0-9]+[0-9]+...		
	emjcd	http://www.emjcd.com/...		
	jdoqocy	http://www.jdoqocy.com/click-[0-9]+[0-9]+...		
	kqzyfj	http://www.kqzyfj.com/click-[0-9]+[0-9]+...		
Ebay	qksrv	http://qksrv.net/...	99	1963
	tkqlhce	http://www.tkqlhce.com/click-[0-9]+[0-9]+...		
	ebay	http://rover.ebay.com...&campid=...		
Envato	audiojungle	https://audiojungle.net/...&ref=...	108	0
	codecanyon	https://codecanyon.net/...&ref=...	14	76
	envato	https://marketplace.envato.com/...&ref=...	175	262
	graphicriver	https://graphicriver.net/...&ref=...	15	1465
	themeforest	https://themeforest.net/...&ref=...	19	200
videohive	https://videohive.net/...&ref=...	578	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	http://...7eer.net/c/[0-9]+/[0-9]+/[0-9]+...	180	529
	evyy	http://...evyy.net/c/[0-9]+/[0-9]+/[0-9]+...		
KontrolFreek	kontroloffreek	https://www.kontroloffreek.com/...&a_aid=...	117	0
Makeup Geek	makeupgeek	http://www.makeupgeek.com/...&acc=...	57	0
Pepperjam Network	gopjn	http://www.gopjn.com/t/[0-9]-[0-9]+[0-9]+[0-9]+...	2	79
	pjatr	http://www.pjatr.com/t/[0-9]-[0-9]+[0-9]+[0-9]+...		
	pjtra	http://www.pjtra.com/t/[0-9]-[0-9]+[0-9]+[0-9]+...		
	pntra	http://www.pntra.com/t/[0-9]-[0-9]+[0-9]+[0-9]+...		
	pntrac	http://www.pntrac.com/t/[0-9]-[0-9]+[0-9]+[0-9]+...		
pntrs	http://www.pntrs.com/t/[0-9]-[0-9]+[0-9]+[0-9]+...			
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		http://www.shareasale.com/r.cfm...	199	616
		http://www.shareasale.com/m-pr.cfm...		
		http://www.shareasale.com/u.cfm...		
Studybay	apessay	https://apessay.com/...&rid=...	141	0
Zaful	zaful	http://zaful.com/...&lkid=...	32	786