Poster: Characterizing Deceptive Affiliate Marketing Offers

Victor Le Pochat
imec-DistriNet, KU Leuven
victor.lepochat@kuleuven.be

Cameron Ballard
New York University
clb478@nyu.edu

Tom Van Goethem
imec-DistriNet, KU Leuven
tom.vangoethem@kuleuven.be

Wouter Joosen
imec-DistriNet, KU Leuven
wouter.joosen@kuleuven.be

Damon McCoy
New York University
mccoy@nyu.edu

Tobias Lauinger
New York University
lauinger@nyu.edu

While browsing the web, using social media, or reading their email, Internet users are regularly exposed to deceptive advertisements [5, 10]. These promote low-quality products and services, such as questionable dietary supplements [1] or cryptocurrency investment platforms [7], and mislead consumers with, e.g., fake celebrity endorsements [1]. Others trick users into installing potentially unwanted software [9] or into disclosing personal data, such as through fake contests [2].

These seemingly unconnected scams are all enabled by the **deceptive affiliate marketing** model [6, 11]. The deceptive ads (Fig. 1) are not run directly by the *merchants* who sell the corresponding product or service (*offer*). Instead, merchants "outsource" the marketing to independent *affiliates*, who exploit advertising channels to promote these offers in return for a commission on each successful conversion. *Affiliate networks* mediate this interaction between merchants and their affiliates [6, 8]. Unclear liability for the activities of the other parties may make abuse more attractive, with some ecosystem players even openly acknowledging using deceptive tactics [3].

We present a novel method to holistically measure the deceptive affiliate marketing ecosystem, using the point of view of affiliate marketers. We discover deceptive products and services on so-called *aggregators*, where merchants and affiliate networks advertise marketing opportunities to prospective affiliates. They provide us with ground truth on the breadth of deceptive products on offer, across all verticals in the ecosystem. Metadata often includes commission rates (payouts), offer conditions (conversion criteria), and content previews, which enables us to reliably understand the dynamics of the ecosystem, including monetary incentives. We develop web page scrapers that extract available offers from aggregators discovered through an iterative search. Afterwards, we merge duplicate offers (identical offers of the same network that are published on multiple aggregators; slightly differing offers of the same network (e.g., a different targeted country but the same product/service); and offers published by multiple networks). We also remove offers from reputable businesses based on 4 other aggregators of legitimate affiliate programs. To characterize the current state of the deceptive affiliate marketing ecosystem, we collect an extensive data set from 23 English- and Russian-language aggregators over 21 months. So far, we have gathered nearly 400,000 deceptive offers (i.e., products and services available for promotion) across 842 affiliate networks.

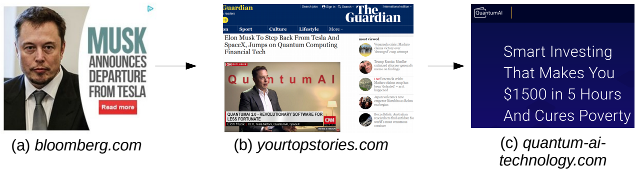
MUSK
ANNOUNCES
DEPARTURE
FROM TESLA
Read more
The Guardian
Elon Musk To Step Back From Tesla And SpaceX, Jumps on Quantum Computing Financial Tech
QuantumAI
QUANTUMAI 2.0 - REVOLUTIONARY SOFTWARE FOR LESS FORTUNATE
QuantumAI
Smart Investing That Makes You $1500 in 5 Hours And Cures Poverty

(a) *bloomberg.com*

(b) *yourtopstories.com*

(c) *quantum-ai-technology.com*

Fig. 1. An example of deceptive affiliate marketing.

We develop a taxonomy of 9 verticals of products and services that constitute the deceptive affiliate marketing ecosystem: *Dating/adult* (the most popular vertical overall), *Entertainment*, *Finance*, *Gambling*, *Games*, *Health/beauty*, *Software*, *Sweepstakes*, and *E-commerce*. In order to analyze the primary deceptive strategies used by marketers, four authors manually coded a sample of 750 offers, 75 for each of the nine verticals and 75 uncategorized offers. The annotators labeled offers based on the preview landing page, a preview screenshot from the aggregator, the offer's description and name. Annotators were asked to only label an offer if they had sufficient data to reliably assess its contents, and to verify that the metadata was consistent. Annotators were encouraged to explore the full landing page and links on it (e.g., terms of service), but were not required to submit any forms. We therefore do not know what happens to personal data or credit card details when they are submitted. We used deductive coding, starting from a codebook developed based on prior domain knowledge, as well as a codebook used in prior research on deceptive advertising [12, 13] and dark patterns [4]. Based on this labeling, we observe the use of a wide variety of deceptive tactics: making exaggerated claims, misrepresenting endorsements, using dark patterns [4] to pressure users, and tricking unsuspecting customers into authorizing (recurring) payments. In general, products and services tend to be of low quality and overpriced.

Quantitatively, payouts to affiliates depend on the type of

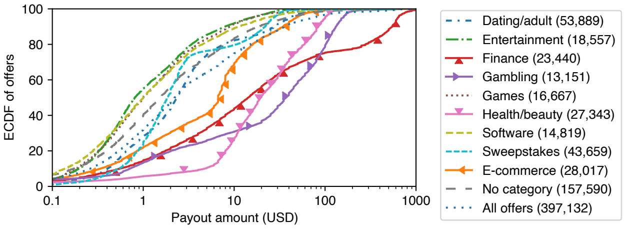

Fig. 2. CDF of payouts per vertical. (Legend lists number of offers.)

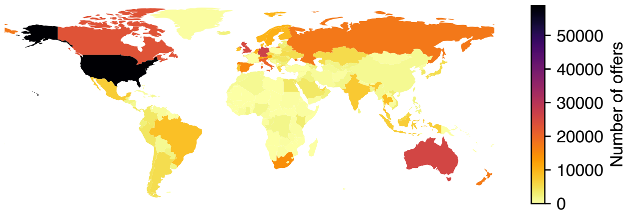

Fig. 3. Number of offers available in each country.

product and what a customer has to do or pay to "complete" the offer from the affiliate's point of view (Fig. 2). The most lucrative offers relate to investment scams, capitalizing on new trends such as cryptocurrency, with payouts in the hundreds of U.S. dollars once customers make a deposit. Physical goods (e.g., health and beauty products) and subscription services (e.g., fake dating sites) also command relatively higher payouts, while virtual goods, app installs, or personal data pay single-digit figures, or less.

At the country level, offer availability differs based on income levels, monetization methods, and regulations. The ecosystem usually separates countries into three tiers, which we describe here based on definitions from major ecosystem players. *Tier-1* countries comprise English-speaking Western countries, and usually also the rest of Western Europe and wealthy Asian countries. They are considered the most desirable, as consumers have high incomes. Investment-related *Finance*, *Sweepstakes* offers are more prevalent in Tier-1 countries. *Tier-2* countries tend to comprise Latin America, the Middle East, Eastern Europe and Russia, and some Southeast Asian countries. Consumers in these countries have moderate incomes. Credit or insurance-related *Finance*, *Gambling*, *Health/beauty* offers are more prevalent in Tier-2 countries. *Tier-3* countries comprise most African countries, and the remaining countries in South America and Asia. Digital goods in the *Entertainment*, *Games*, and *Software* verticals are most common in tier-2 and 3 countries. In absolute numbers, the United States has the most offers, at 58,793 offers (Fig. 3). Germany comes a distant second at 26,469 offers. Overall, more offers are available in North America, Europe, Russia, Australia, South Africa, and to a lesser extent Brazil and India. This correlates with the higher 'tiers' assigned to these countries, suggesting merchants prefer selling products and services to higher-income audiences.

Next to the more comprehensive insights into the ecosystem that our novel data set provides, our vantage point permits to increase the breadth, speed, and impact of future interventions. The more complete coverage of our data set across countries, device configurations, or targeted populations improves the discovery of scams regardless of their audience. We can detect new offers as soon as they are created, potentially before affiliates advertise them to consumers, and the related webpages can be blocked or taken down quickly. Our data also allows to identify the large affiliate networks that mediate the ecosystem and quantify their scale, which can be used to prioritize technical, financial, and legal interventions and maximize their impact. Finally, we identify common infrastructure providers for offers from our data, both mainstream providers that are abused, and dedicated sites that are potential venues for (technical) intervention through targeted takedowns. As we continue collecting data, our data set can be leveraged as an "early warning system" to deploy timely defenses against newly emerging deceptive practices. We will therefore share our data with researchers and stakeholders in order to enrich other measurements that provide further insights into the way this ecosystem operates, and ultimately enable more powerful interventions to prevent users from being exposed to deceptive affiliate marketing and losing their money or personal data.

References

[1] C. Steven Baker. *Subscription Traps and Deceptive Free Trials Scam Millions with Misleading Ads and Fake Celebrity Endorsements*. Better Business Bureau, Dec. 12, 2018. URL: https://www.bbb.org/globalassets/local-bbbs/st-louis-mo-142/st_louis_mo_142/studies/bbb-study-free-trial-offers-and-subscription-traps.pdf.

[2] Jason W. Clark and Damon McCoy. "There Are No Free iPads: An Analysis of Survey Scams as a Business". In: *LEET*. 2013.

[3] Zeke Faux. "'They Go out and Find the Morons for Me'". In: *Bloomberg Businessweek* 4564 (2018), pp. 56–61. ISSN: 0007-7135.

[4] A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, and A. Narayanan. "Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites". In: *Proc. ACM HCI* 3 (CSCW 2019), 81:1–81:32. DOI: 10.1145/3359183.

[5] T. Nelms, R. Perdisci, M. Antonakakis, and M. Ahamad. "Towards Measuring and Mitigating Social Engineering Software Download Attacks". In: *25th USENIX Security Symposium*. 2016, pp. 773–789.

[6] Dmitry Samosseiko. "The Partnerka - What Is It, and Why Should You Care?" In: *19th Virus Bulletin International Conference*. 2009, pp. 115–120. URL: https://www.sophos.com/it-it/medialibrary/PDFs/technical%5C%20papers/samosseikovb2009paper.pdf.

[7] Craig Silverman. "Ads Inc. Shut Down, But The Tools It Used To Trick People On Facebook Have Lived On". In: *BuzzFeed News* (Dec. 1, 2020). URL: https://www.buzzfeednews.com/article/craigsilverman/ads-inc-crypto-scams-facebook.

[8] Gianluca Stringhini. "Adversarial Behaviour". In: *The Cyber Security Body of Knowledge*. 2019, pp. 223–249.

[9] Kurt Thomas et al. "Investigating Commercial Pay-Per-Install and the Distribution of Unwanted Software". In: *25th USENIX Security Symposium*. 2016, pp. 721–738.

[10] Phani Vadrevu and Roberto Perdisci. "What You See Is NOT What You Get: Discovering and Tracking Social Engineering Attack Campaigns". In: *IMC*. 2019, pp. 308–321. DOI: 10.1145/3355369.3355600.

[11] Jeff White. *Takedowns and Adventures in Deceptive Affiliate Marketing*. Unit42. Apr. 25, 2019. URL: https://unit42.paloaltonetworks.com/takedowns-and-adventures-in-deceptive-affiliate-marketing/.

[12] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. "Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites". In: *5th Workshop on Technology and Consumer Protection*. 2020.

[13] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. "What Makes a "Bad" Ad? User Perceptions of Problematic Online Advertising". In: *CHI*. 361. 2021, 361:1–361:24. DOI: 10.1145/3411764.3445459.