Work in progress: Exploring Deceptive Affiliate Marketing

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

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

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

***Abstract*—Internet users are often exposed to advertisements that promote deceptive products and services, risking to lose their money or personal data. In our ongoing work, we explore how these malicious practices are often supported by the 'deceptive affiliate marketing' model, where deceptive merchants can shift liability to independently operating and often abusive advertisers ('affiliates'). We first develop a taxonomy of deceptive products and services advertised through affiliate marketing. We then present preliminary findings using a novel data set collected from the vantage point of the affiliate, highlighting how different product categories, countries, and user agents are valued differently. This emphasizes the need to obtain sufficiently diverse coverage when studying or defending against malicious advertising. We conclude with our plans for future work identifying main actors and intervention points in the ecosystem.**

1. Introduction

While browsing the web, using social media, or reading their email, Internet users are often exposed to advertisements that promote deceptive products, such as dietary supplements and cryptocurrency scams, or tactics that seek to collect personal information, such as through fake contests. These users then run the risk of losing their money or personal data.

In our ongoing work, we explore how these malicious practices are often supported by the 'deceptive affiliate marketing' model [6, 8, 72]. Here, individual *affiliates* exploit advertising channels to promote deceptive products or services ('offers') created by a variety of untrustworthy *merchants*, in return for a commission on each sale made, supported by intermediary *affiliate networks* [60, 65]. This marketing model is attractive to malicious entities. Merchants can shift liability to affiliates when the latter use deceptive or abusive tactics to promote products, such as fake celebrity endorsements [8, 15, 20, 72]. Conversely, affiliates do not need to consider the quality or even legality of the merchants' products they promote, as they play no part in the actual production and distribution [70]. In between, affiliate networks appear to be aware of or even encourage the deceptive practices of both affiliates and merchants [9, 41, 46]. This ecosystem operates at a very large monetary scale: one merchant earned \$179 million over five years [24], one affiliate network was estimated to have \$100 million in yearly revenue [16], and affiliates using one of the largest tracking platforms were estimated to purchase \$1.7 billion in advertising a year [20].

Our preliminary findings indicate that large parts of the deceptive and malicious content that end users encounter on the web are connected to this ecosystem. While previous research has already investigated parts of this ecosystem [7, 13, 34, 37, 39, 45, 48, 53, 71], this was done in isolation and without connecting it to the actors behind them. For our study, we retrieve data from the vantage point of the affiliate. We build custom scrapers for 21 aggregators that list *offers* (i.e., products and services) that merchants publish to affiliate networks and make available for affiliates to promote. We collect data on a daily basis and for already more than one year to obtain longitudinal coverage. The advantage of our vantage point is that we gather ground truth on the breadth of deceptive products on offer, often with detailed metadata that includes commission amounts, advertising channels, and content previews. Our data is also more comprehensive in terms of global coverage, often ignored by previous research.

In our research, we seek to understand the extent of the ecosystem through the following research questions:

- What is the prevalence and breadth of deceptive products and services on offer?
- What is the prevalence and breadth of deceptive strategies used to promote offers?
- Who are the main (identifiable) ecosystem players?
- Which categories of deceptive products/services and which countries are more valuable to affiliates?
- Which infrastructural and financial providers are used by players in the ecosystem?
- Can we identify and implement intervention points?

Using a subset of our collected data, we present preliminary insights. We see that sweepstakes, dating, and health-related offers are the most prevalent, and identify specialized affiliate networks with over ten thousand offers each. Commissions vary by vertical and country, with cryptocurrency offers being the most lucrative, routinely yielding commissions over \$100. Finally, we discover common domains that relate to affiliate networks, providing potential venues for intervention by disabling one link in the redirection chain from advertisement to product. Ultimately, such interventions could prevent users from being exposed to deceptive affiliate marketing and subsequently losing money or personal data.

After an example of deceptive affiliate marketing (section 2) and an introduction to the ecosystem (section 3), we develop a taxonomy of deceptive products and services advertised through affiliate marketing (section 4). We then discuss our data collection (section 5) and present our preliminary findings (section 6). After describing related work (section 7), we conclude with avenues for future work (section 8).

2. A real-world example

We illustrate how a product is promoted through affiliate marketing with a real-world example (Figure 1).

A user is browsing a news website, and in the margin of an article sees an ad claiming that 'Musk announces departure from Tesla' (a), bought by an affiliate from an ad exchange (the traffic source). Intrigued by the headline, they click the ad and are redirected to a pre-landing page hosted on `yourtopstories.com` (b). This page contains an elaborate article about Elon Musk's investment in a company called 'QuantumAI', advertising how its quantum computer-based stock trading algorithm can make investors wealthy. The page attempts to seem trustworthy by displaying the logo and mimicking the layout of both the UK newspaper *The Guardian* and US news network *CNN*, and contains a comment section with fake profiles further acclaiming the product.

All links on the pre-lander point to a URL hosted on `holdon1sec.com`, an intermediate page that checks whether the correct `Referer` header is set (else it displays a `Bad Request` page) and then redirects through an affiliate network tracking link to a landing page on `quantum-ai-technology.com` (c). This page allows users to sign up for the 'QuantumAI' product, claiming it will "cure their poverty". Users can view a video showing images from a presentation given by Musk with a voice actor mimicking Musk's voice advertising the product, followed by testimonials by the 'owner of the company' and several 'customers'; however, these actors are freelancers recording a predetermined message [49]. The page also displays testimonials by Jeff Bezos and Bill Gates (as 'advisors') and claims IBM, Microsoft and OpenAI are 'partners'.

The Elon Musk prelander and QuantumAI service are part of an offer created by the affiliate network 'Affiliate Interactive'. The affiliate with ID `1247` has signed up for this network and picked this offer named 'Elon Musk - Prelander - AU, DK, FL, IS, IR, NL, NZ, NO, SG, SE, UK' (ID `166`) from the 2,540 offers made available by the network. The merchant previously separately submitted this offer to the affiliate network, as part of their contract where the affiliate network will search affiliates to promote the merchant's service. The offer is listed as part of the 'cryptocurrency' vertical, i.e., the offer's category. In the offer description, the network stipulates that "all traffic types [are] allowed except incentives", meaning that other than promising users a reward for completing the offer (e.g. cash or access to content), the affiliate is free to employ any traffic source to advertise the product, such as search engines, social media posts, or email spam. The Elon Musk pre-landing page is part of the creatives (marketing material) provided by the merchant as part of the offer.

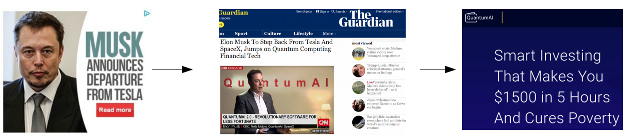

(a) *bloomberg.com* (b) *yourtopstories.com* (c) *quantum-ai-technology.com*

Figure 1. Example promotional chain for an affiliate marketing offer.

In this case, the affiliate has set up an ad campaign buying advertisement space on a news website using an ad exchange (the traffic source), using the tracking URL from the affiliate network that contains both the affiliate ID and offer ID. As indicated by the countries or GEOs listed in the offer name, the offer may only be advertised to users in Australia, Denmark, Finland, etc.; this is explicitly checked by the intermediate page before redirecting to QuantumAI. Finally, the offer sets out the requirements for the affiliate to get paid (the conversion): in this cost per sale model, a user from the allowed set of countries who was redirected to the QuantumAI website by the affiliate is required to deposit at least 250 dollars to trade in Bitcoin (the 'sale'). If this is the case, the affiliate receives a commission of up to 570 dollars from the affiliate network; the network charges at least this amount to the merchant operating the QuantumAI site.

3. Deceptive affiliate marketing ecosystem

3.1. Key terminology

3.1.1. Ecosystem players. There are three main types of players in the affiliate marketing ecosystem. **Merchants** (advertisers) have a product or service that they want to promote and sell. They seek **affiliates** (publishers, marketers, partners, advertisers[1]) who will promote the merchant's product. Merchants and affiliates find each other through **affiliate networks** (advertiser network) that act as intermediaries. These networks also provide the technical and financial infrastructure that ultimately pays the affiliate when they successfully promoted the merchant's product.

3.1.2. Monetization. Merchants post **offers**[2] to affiliate networks, who in turn make these offers available to affiliates for promotion. An offer is usually for one specific product, belonging to a certain **vertical** (niche, category). An offer will also include restrictions on who the product may be advertised to. Offers are usually targeted at specific countries or **GEO**s, which are divided into **tier**s to reflect their perceived wealth and therefore attractiveness. Certain advertising channels may also be (dis)allowed. Additionally, an offer stipulates the conditions and amounts (commissions, payouts) for a successful **conversion**, i.e. when a customer completes the offer and the affiliate is paid out. Usually, the affiliate receives a one-time fixed-amount **commission** upon conversion. Alternatively, an affiliate can be paid through revenue sharing (RevShare), where they receive a percentage of all sales made to the customer over some period of time. Two models prevail for awarding a commission: **cost per sale** (CPS, pay per sale, PPS) where a customer must purchase the product or service, and **cost per lead** (CPL, pay per lead, PPL) where a customer only needs to provide their contact information or personal data ('lead generation'). Within the ecosystem, the term **CPA marketing** (cost per action, cost per acquisition) is often used as a synonym for affiliate marketing. However, the 'action' may refer either to only a lead, or to both a sale or a lead.

1. The term 'advertiser' is sometimes used for affiliates, as they are the ones who will advertise the product to potential customers.

2. Offers are sometimes also called 'affiliate programs', although this can also refer to more legitimate businesses, see subsection 3.2.

3.1.3. Redirect chain. Once an affiliate has selected an offer, they will set up a specific **campaign** to promote it. The affiliate receives a **tracking link** (affiliate link) from the affiliate network: the affiliate IDs in this link will ultimately allow the network to determine which affiliate to pay if the offer converts. The affiliate then selects the **traffic source** through which to promote the offer: e.g., their own websites, advertisements or email. If a customer follows the link in this traffic source, they will be led to the advertised product through a chain of redirects. The traffic source may point to this link directly or redirect to it from e.g. URL shorteners and/or from a custom tracking URL that allows to monitor the campaign with (often specialized) tracking software.

The potential customer is then led to and through a 'funnel', i.e. a number of **creatives** (content pages) promoting the offer. The tracking URL may then first redirect to a **pre-landing page** (pre-lander), a page that entices customers to proceed with the offer. Example pre-landing pages are blog posts with fake celebrity endorsements of a product, surveys that announce that the user has won a prize or warning pages that may claim a technical issue with the user's computer. The customer is then led either directly from the tracking URL or from the pre-landing page (potentially through intermediate pages) to the **landing page** (offer page, lander), where a customer is invited to complete the offer, e.g. by purchasing the product or service (cost per sale), or by entering their personal details (cost per lead). In the former case, this may cause a redirect to a payment page.

3.1.4. Responsibilities. The responsibilities for creating and/or accepting the contracts between affiliates, merchants and affiliate networks as well as the contents of pages in the redirection chain depend on which services the different ecosystem players provide. These distributed responsibilities allow to shift liabilities to the other parties in the ecosystem, but also make it unlikely that any party is not at least partially aware of the deceptive practices.

Before an affiliate can start promoting offers from a specific affiliate network, they must first get accepted into that network. Different networks set different requirements based on how restrictive and exclusive they want to be. They may request a face-to-face interview (over videoconferencing), and may ask which traffic sources, verticals and countries the affiliate plans to target. Networks may ask what previous experience the affiliate has in the ecosystem: more restrictive networks will only accept affiliates with a proven track record and sufficient traffic and revenue. Some networks may even proactively approach affiliates to join them. If the affiliate gets accepted, they are assigned an affiliate manager, who will be the primary point of contact as well as the person responsible for approving requests by the affiliate to promote a specific offer. Similarly, an account manager will maintain the relationship between a merchant and the affiliate network.

In terms of providing 'creatives', the affiliate is usually responsible for the promotional material on the traffic source. However, landing pages can be created by either the merchant or the affiliate network. Resources for pre-landing pages can be provided by any of the three parties. These different parties also set different constraints on which traffic sources and creatives are acceptable. The affiliate manager is usually responsible for approving the promotional material created by the publisher and confirm that it falls within the restrictions put forward by the network and/or merchant. The traffic source may also put (different) restrictions on this material, and check whether it is acceptable and leads to legitimate (pre-)landing pages.

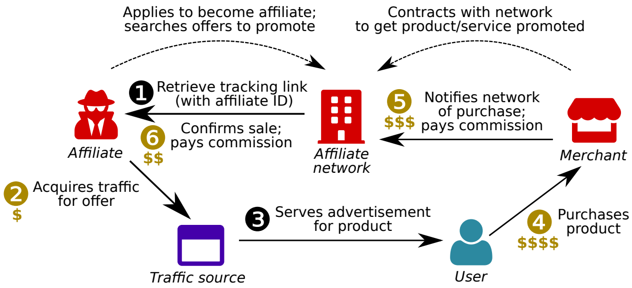

Figure 2. A typical payment flow for a cost per sale offer.

3.1.5. Payment flows. Throughout the process of affiliates attracting customers, money changes hands multiple times, as shown in Figure 2. Once an affiliate has selected an offer and retrieved a tracking link (1), they will acquire (paid) traffic promoting the offer from a traffic source (2): e.g. purchasing advertising space but also hiring botnets for email spam. The affiliate can indicate which user demographics the traffic source should target, and the source will display the promotional material (leading to the tracking link) to a user (3). The user is then led to the merchant's landing page. In the case of a cost per sale offer, if the user is persuaded to purchase the promoted product, they will make a payment to the merchant (4). In the case of cost per lead, the user only provides personal information directly to the advertiser; this advertiser monetizes the lead in another way (such as contacting the user later on to sell a product, or reselling the personal data). Once the merchant is satisfied that the offer has been successfully completed, they will notify the affiliate network of the purchase and pay their commission to the network (5). The network then pays the agreed commission to the affiliate (6), taking a service fee from the merchant's commission.

The network/merchant may apply a hold (locking period) between the purchase and awarding the commission, e.g. to verify that the customer's payment succeeds and that the traffic complies with the offer's restrictions. Affiliate networks further differentiate themselves with regard to the supported payment providers, payment frequency and minimum payment, striking a balance between having attractive payment conditions for affiliates (with easier and faster payments) and safeguarding against fraudulent behavior. More successful affiliates can also negotiate better payment terms (e.g. higher commissions or more frequent payments) with the network.

3.2. Legitimate, deceptive, malicious or illegal?

We use the term 'deceptive affiliate marketing' to describe the ecosystem that we study. The term 'malicious (advertising)' ('malvertising') appears to be usually reserved for abuse that ultimately breaches the security of a user's system, e.g., by propagating malware that may include the system in a botnet. Instead, the 'deceptive' practices that we observe more directly harm the consumer: such as by making them purchase low-quality products or services, by having them install unwanted software, or by tricking them into disclosing personal data. In this respect, the practice has also been called 'social engineering' [71].

Deceptive affiliate marketing operates in somewhat of a 'grey zone': the practice is usually not illegal in and of itself, but ecosystem players may engage in particular behavior that violates certain legislation or otherwise deceives consumers. For example, affiliates may advertise products using deceptive claims [50], or create fake celebrity endorsements [8, 15, 20, 72]. Meanwhile, merchants may hide crucial information about their products and services, although this deceptive nature is sometimes misunderstood. Physical goods may actually be shipped to consumers (instead of an outright scam where the user receives nothing), but the scam lies in hiding repeat billing [8, 50] or sending products that are of very low utility [20, 50]. Online contests are also not necessarily 'fake', i.e., the giveaway may actually happen [20], but their terms are often very limited with only a handful items being given away per year, or trick users into expensive subscriptions before becoming eligible for the contest [20].

Legally, players in the affiliate marketing ecosystem would have to adhere to consumer protection regulation. Internationally, laws in major jurisdictions enforce that merchants and advertisers must not engage in unfair commercial practices, prohibiting them from using misleading and deceptive advertising [17]. For example, in the United States, the Federal Trade Commission has the authority to regulate on such practices, and has successfully targeted deceptive affiliate marketing actors in the past [4, 6, 21–27, 29, 50]. However, such rules are often broadly defined and may require clarification and interpretation by regulators or courts [17], making it less obvious to determine when these laws are violated. Finally, merchants may be legally liable for any damage caused by their products [59].

In our deceptive affiliate marketing model, usually all ecosystem players are complicit in the deceptive or malicious behavior: merchants sell low-value products or levy hidden charges, affiliates use deceptive advertisements to persuade consumers, and affiliate networks condone both practices by accepting these – obviously deceptive – merchants and affiliates into the network. Ultimately, consumers are the main victim of these practices, as they lose money to these actors or are otherwise deceived. Sometimes, traffic sources are also victims, as their platforms are abused for deceptive advertising, but others appear to knowingly attract and permit abusive practices. For example, Vadrevu et al. [71] found three ad networks where over half of all advertisements were deceptive. Such traffic sources as well as affiliate networks sometimes openly guide affiliates on how to circumvent restrictions on abusive content [9, 41, 46], indicating their awareness.

The affiliate marketing model is not inherently malicious or deceptive; in fact, many legitimate businesses (e.g., Amazon[3], eBay[4], AliExpress[5]) use the model to award commissions to affiliates who successfully promote their products. We do observe legitimate products being listed on the offer aggregators that we track (and will seek to identify them and process them separately), but we argue that their listing on these aggregators can still 'attract' abuse. Legitimate and deceptive offers are intertwined on these platforms, so affiliates browsing the aggregators may not (be able to) distinguish between these two kinds of offers. Moreover, legitimate brands may be harmed when affiliates use abusive tactics to promote their products. For example, Farooqi et al. [19] found mainstream developers among incentivized mobile app install campaigns who were unaware of being part of such campaigns.

3. https://affiliate-program.amazon.com/
4. https://partnernetwork.ebay.com/
5. https://portals.aliexpress.com/

4. Verticals

We taxonomize the verticals (categories) targeted by affiliate marketers, to understand which products and services are prevalent in the deceptive affiliate marketing ecosystem and highlight how they deceive consumers. We base our taxonomy on the verticals listed in the offer aggregators, complemented by guides from major affiliate marketing ecosystem players [1–3, 28, 42, 55, 58].

Sweepstakes Merchants hold sweepstakes where they promise to give away free products (e.g., iPhones [14]) or vouchers (e.g., for shops). Major brands have warned users that they are being misrepresented as supporting these sweepstakes, and that users should not go along with the offer [30, 33]. Sweepstakes break down into two major types: those designed for 'lead generation', where the primary goal is collecting (private) contact details from users to sell it onto third parties, and those that lure users into submitting their credit card details and (unknowingly) starting subscriptions (e.g., for fake entertainment platforms).

Health and beauty Merchants offer various physical health and beauty-related products for sale. These cover categories such as weight loss, 'nutraceuticals' (foods and supplements with claimed health benefits), CBD, or male enhancement. Merchants may legitimately ship products [45], but users may unknowingly start a subscription for regular deliveries, or receive products without any actual utility.

Financial products and services Merchants offer various (mostly virtual) financial products and services. These cover categories such as trading platforms for cryptocurrency, foreign exchange, or binary options; credit; insurance; or 'business opportunities' ('BizOpp') for building one's own business (akin to 'get-rich-quick schemes'). Financial products and services are often heavily regulated, and they may be illegal to sell or promote depending on the jurisdiction, in the least if necessary details to understand the financial impact of the scheme (such as the risk involved or the credit rate applied) are omitted.

Dating and adult content Merchants promote sites hosting dating services or adult content. These services usually expect users to start a subscription. Services sometimes deceive users by matching them with fake profiles. It may be illegal to operate or promote these services in certain jurisdictions.

Gambling Merchants promote sites offering online gambling services such as sports betting or casinos. Gambling sites may operate illegally in certain countries, and their promotion may be prohibited or restricted.

Gaming Merchants promote games that either run in a browser or as an app. These games may cost money to install, require a subscription or monetize users through micro-transactions.

Software Merchants offer software for installation on desktop and mobile, ranging from legitimate apps (for which existing affiliate programs are often 're-published' as an offer [19], for example in the case of VPN apps) over low-quality or potentially unwanted software (such as purported anti-virus software or browser extensions that add a toolbar and inject ads; sometimes called 'utilities') to outright malware (such as fake Flash Player software).

E-commerce Products and services on offer are not restricted to the previous categories. Merchants may promote various types of physical goods (sometimes with low utility) or services (sometimes reusing existing affiliate programs, e.g., for travel sites).

5. Data collection

5.1. Aggregator discovery

In our data collection, we cover 21 "offer aggregators", i.e., search engines for offers from multiple affiliate networks (Table 1). As affiliate networks often originate in Russia [53, 60], we cover 14 English- and 7 Russian-language aggregators. We find networks that are only listed on the Russian-language aggregators, confirming our decision to search and include them. Overall, we observe that offers listed on these platforms cover all countries worldwide.

We employ a multi-tiered approach to discover the most popular aggregators. We use the Google search engine with generic keywords (such as "affiliate offers", "CPA offers") and with the names of major networks listed on previously discovered aggregators (such as "AdCombo", "MaxBounty"). We also consult specialized forums (such as AffiliateFix, BlackHatWorld, affLift) and sponsor lists of major affiliate conferences (such as Affiliate Summit, Affiliate World) to find additional aggregators. We conduct these searches until we reach saturation in the list of discovered aggregators.

While we cannot independently confirm the reliability of offer data, we suspect that aggregators obtain it directly from affiliate networks. Metadata available at some aggregators suggests that they integrate directly with the offer management platforms of affiliate networks. Aggregators then regularly retrieve the most current set of offers, through either an API or scraping (using a provided account at the network). We plan to verify the accuracy and timeliness of aggregator data by comparing between aggregators as well as with offer data made available publicly by affiliate networks (i.e., without registration).

In addition, we argue that there is an incentive for networks and aggregators to provide accurate data to affiliates. Underground activities operate on a reputation system, where breaches of trust result in negative feedback on e.g., underground forums [31]; similarly, we can expect dishonest aggregators to be called out. If inaccuracies in the data are present, we therefore expect this data to be outdated rather than purposefully wrong.

TABLE 1. OVERVIEW OF OFFER AGGREGATORS.

		Number of		
Aggregator	Lang.	networks	offers	observations
ActualTraffic	RU	78	17,909	5,012,578
AdMakler	RU	–	1,378	88,908
AdNetworksHub	EN	2	446	69,424
AffBank	EN	117	173,883	4,735,138
AffHomes	EN	48	20,070	2,718,869
AffNext	EN	7	3,129	2,538,928
AffPlus	EN	248	545,502	38,162,212
AffPub	EN	75	30,239	7,171,630
AffScanner	EN	63	101,258	12,229,963
Atlas.io	RU	–	8,267	23,694,058
AVF	RU	32	–	1,107,777
BestAffiliatePrograms	EN	73	60,954	12,591,277
BigFishOffers	EN	8	475,365	95,564,754
Click4Ads	EN	149	279,904	5,675,809
CPADaily	RU	–	17,286	2,248,141
CPAInform	RU	–	–	4,703,411
ODigger	EN	63	42,369	2,430,065
OfferLibrary	EN	19	30,866	4,712,839
OfferVault	EN	318	363,827	12,777,655
Partnerkin	RU	89	30,813	3,248,634
XOffers	EN	126	2,978	12,206

5.2. Data retrieval

We extract available offers through web page scrapers custom-built for every aggregator. Unless a more machine-parsable format is available (e.g., JSON), we retrieve raw HTML through simple HTTP requests using the Python `requests` library, which suffices to retrieve all necessary data. We then parse the HTML page to extract data from relevant elements using the Python `BeautifulSoup` library. Most aggregators present a paginated overview of all offers on their main page, with for each offer a link to a page with additional metadata. We first traverse the paginated overview to collect basic data for all offers (visiting N/P pages, with N the number of offers and P the number of offers per page); we automate our scrapers to retrieve the full offer overview on a daily basis. Afterwards, we individually request detailed data for each offer (visiting N pages); we retrieve detailed data for newly seen offers once a day, but recollect detailed data for all offers once a week. We believe this scraping frequency strikes a good balance between timeliness of the data and consumption of scraping resources. Moreover, we seek to optimize scraping wherever possible, e.g., leveraging internal API's or maximizing the number of offers per page, which also reduces strain on the aggregators.

Our data collection started on March 24, 2020 across the aggregators in Table 1 and is still ongoing to provide longitudinal coverage. Gaps in coverage do occur (Figure 3): aggregators have become unavailable, aggregators changed their site layout which broke our scrapers, or our scraping infrastructure was temporarily down.

5.3. Ethics

Given the often malicious nature of the players in the ecosystem, we must carefully consider how we proceed with our study and treat our findings. We believe that the goals of our study will bring about significant benefits to understanding and even combating the malicious practices

Figure 3. Availability of aggregator data.

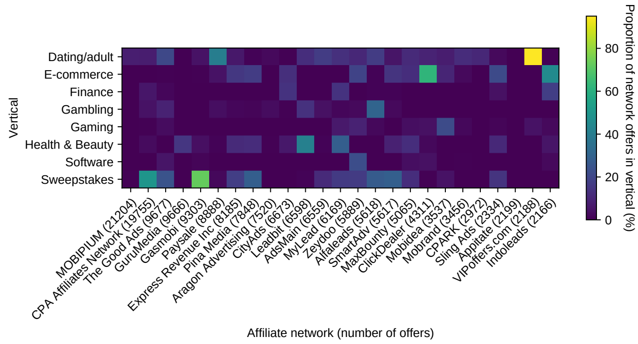

Figure 4. Top 25 networks with the most offers listed on OfferVault, and the proportion of offers per vertical per network.

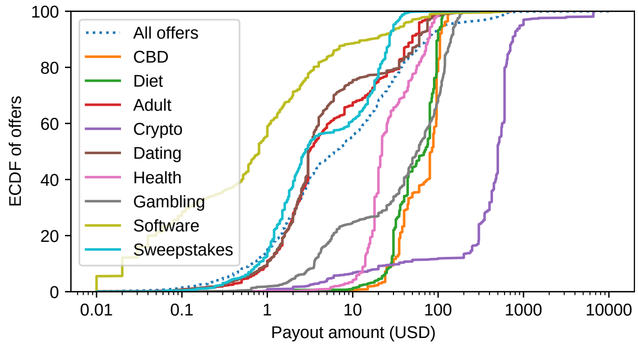

Figure 5. Distribution of payouts (USD) per offer (sub)vertical.

within the affiliate marketing ecosystems, which therefore also justifies certain experimental techniques to obtain data on and insights into the ecosystem. Ethical evaluations conducted in previous studies have lead to a consensus that given appropriate measures, the use of scraping is ethically justified especially when studying malicious ecosystems [51, 57, 62]. To the best of our knowledge, the scraped offer data does not contain personally identifiable information. Once our study has been fully developed, we plan to share our data with other researchers and parties of interest, including law enforcement when applicable.

By scraping offer aggregators, we avoid the need to register for individual affiliate networks. We observe that this registration process ranges from basic username/password registration, over providing contact details (email address, phone number, instant messaging accounts), to interviews with those managing the affiliate network. Next to reducing the effort in collecting data, we do not expose ourselves to the players in the ecosystem, nor do we have to resort to deception when describing our goals or contact details.

6. Preliminary results

We now present preliminary findings from our data analysis that show how the deceptive affiliate marketing values fraud types and countries differently. This analysis is currently limited to data from the OfferVault aggregator until October 20, 2020, as the effort to post-process, merge and verify data from all aggregators, which requires identifying identical offers and normalizing metadata across the aggregators, is still ongoing. Our results are therefore only indicative of trends across a sample, but do not yet fully describe the ecosystem.

Our data sample contains 231,422 offers with a distinct identifier. The top verticals are sweepstakes (at least 27,745 offers), dating/adult (at least 18,512), and health/beauty (at least 6,373). Figure 4 shows the top 25 affiliate networks according to the number of offers. The top network, MOBIPIUM, has 21,204 distinct listed offers, and is geared towards a variety of offer types for mobile devices. Based on Figure 4, we see that different networks tend to specialize in certain verticals. As we capture data on a broad set of affiliate networks through the aggregators, we observe a larger share of the ecosystem than if we were to focus on specific affiliate networks.

Figure 5 shows that merchants award the highest commissions for cryptocurrency offers, often running into the hundreds of US dollars. Conversely, offers for software are often worth around 1 dollar or less. Overall, the median commission is 6 US dollars, with health (including diet and CBD) and gambling commanding higher commissions, while sweepstakes, adult and dating offers are worth less. section A displays examples of offers, together with the listed vertical and payout, which correlate with the trends observed per vertical.

Figure 6 shows that countries such as United Arab Emirates, Sweden, Russia, Japan, and Singapore see higher payouts than average; this may be due to higher valuations for consumers there, but also due to more lucrative offer types being preferred there. Table 2 lists average payout grouped by category and country. We see larger variations that are not only explained by an overall higher valuation for a certain country: for example, in the U.S., cryptocurrency or dating offers are worth less than in other countries, while health and gambling are among the more lucrative. Given the different valuations as well as varying availability of certain offers by country, defenders must ensure that they discover malicious and deceptive websites in a sufficient number of countries, in order to equally and comprehensively protect all users worldwide.

Certain offers indicate in their name that they are only valid for certain user agents. Table 3 lists summary counts for offers with names of OSes or browsers. We find more Mac-related offers, although they are worth less on average than Windows offers. Chrome offers are much more prevalent than other browsers, although Safari offers are most valued. We anecdotally observe that such offers often relate to fake software or deceptive browser extensions, e.g., a ‘Mac Flash Player’ (Figure 7), revealing their malicious nature. Moreover, we find that such offers

TABLE 2. AVERAGE PAYOUTS (USD) PER OFFER CATEGORY AND COUNTRY. WE OMIT THE AMOUNT IF THERE ARE FEWER THAN 10 OFFERS IN THE GIVEN CATEGORY AND COUNTRY.

	BR	IN	JP	RU	SG	ZA	SE	TH	AE	US
Adult	15.3	7.9	23.0	6.4	23.7	4.1	22.5	9.1	50.5	9.0
CBD	–	58.1	54.7	–	–	128.8	55.2	55.3	133.2	74.4
Crypto	443.6	644.7	–	375.2	550.5	497.8	528.8	570.0	544.5	137.4
Dating	30.7	24.4	32.8	13.2	26.7	15.9	27.4	3.2	52.0	13.0
Diet	77.4	32.0	53.7	–	78.0	43.2	43.9	–	67.4	69.9
Gambling	46.9	35.1	116.7	47.3	112.9	71.2	81.8	76.4	75.9	91.0
Health	53.9	51.0	87.5	–	–	–	47.8	25.4	32.0	61.2
Software	10.5	13.0	4.7	20.0	33.0	15.1	8.6	1.5	24.7	9.3
Sweepstakes	2.6	4.3	23.3	–	7.8	3.8	18.2	0.5	15.4	9.1

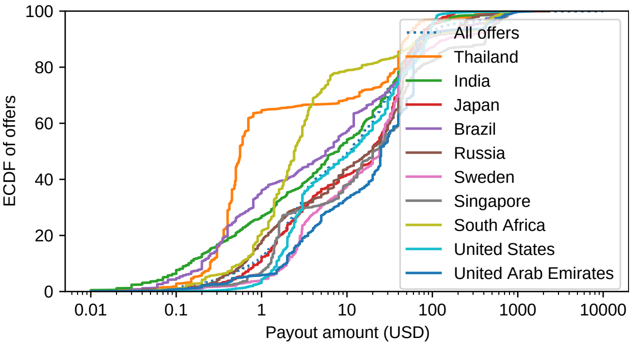

Figure 6. Distribution of payouts (USD) per country.

TABLE 3. COUNT OF DISTINCT OFFERS AND AVERAGE PAYOUTS (USD) PER USER AGENT FOUND IN AN OFFER TITLE.

User agent	Count	Average payout ($)
Windows	228	7.4
Mac	401	5.8
Chrome	677	3.5
Firefox	176	2.8
Safari	251	3.9
Internet Explorer	34	0.7

are often 'cloaked' [32]: the real contents of the landing page are only available if it is visited through the correct user agent (or also country), otherwise it redirects to another benign site (e.g., a search engine). This highlights that measurements limited to only one user agent(/country) will not fully capture the breadth of malicious web content.

Table 4 lists a selection of common keywords (n-grams) in the titles of offers, together with the average payout. Certain verticals are very common (sweepstakes, gambling, health), while others are more highly valued (in particular crypto trading at an average payout of $664). Affiliates may be more attracted by offers with such high commissions, and may consider more abusive advertising tactics given the high return if they succeed. In turn, end users may therefore be exposed more often and more aggressively to advertising for e.g., cryptocurrency trading platforms.

Table 5 lists the most common hosts found in the 'preview links' provided for some offers. By selecting the most prevalent hosts, we obtain a skewed view of the utility of these links: many websites provide services to share screenshots, which affiliate networks use to show the contents of the offer while not disclosing the actual website where it is hosted. Still, such images could be used to cluster similar offers, match recognized texts with landing

TABLE 4. SELECTED COMMON KEYWORDS IN OFFER TITLES.

Keyword	Count	Payout (avg. $)	Vertical
`gift card`	1130	3.7	Sweepstakes/surveys
`casino`	947	77.1	Gambling
`cbd`	905	64.9	Health (CBD)
`iphone 11 pro`	770	17.9	Sweepstakes/surveys
`bitcoin`	460	664.1	Crypto trading
`galaxy s20`	403	20.3	Sweepstakes/surveys
`weight loss`	391	56.2	Health (Diet)
`male enhancement`	290	53.5	Health (Male enh.)
`playstation 5`	100	8.1	Sweepstakes/surveys
`keto diet`	83	61.3	Health (Diet)

TABLE 5. MOST COMMON HOSTS IN PREVIEW LINKS.

Hostname	Count	Domain purpose
`bit.ly`	16110	URL shortener
`snipboard.io`	9741	Screenshot
`gyazo.com`	8998	Screenshot
`gurumedia.info`	8435	Affiliate network
`prnt.sc`	7615	Screenshot
`img.clickdealer.com`	3861	Affiliate network
`integration.alfashops.ru`	3378	Affiliate network
`img.adplato.com`	3024	Affiliate network
`app.zeydoo.com`	2865	Affiliate network
`prntscr.com`	2834	Screenshot
`apps.apple.com`	2280	App store
`play.google.com`	2146	App store
`i.gyazo.com`	1913	Screenshot
`monosnap.com`	1799	Screenshot
`www.screencast.com`	1693	Screenshot
`hyperstech.com`	1521	Deceptive products
`snag.gy`	1439	Screenshot
`affiliate.cpamatica.io`	1360	Affiliate network
`terraleads.com`	1271	Affiliate network
`www.nutaku.net`	1147	Adult content

pages, and understand the nature of the offer in general. Moreover, we find hosts directly related to the affiliate networks, which could indicate the tracking URLs used. As we will discuss in section 8, these tracking websites may prove to be effective intervention targets.

7. Related work

Prior work discussed the inner workings of the affiliate marketing model in the context of cybercrime. Samosseiko [60] first outlined the role of affiliate networks in spam-advertised pharmacies and counterfeit software, focusing on Russian 'partnerka' networks. For a few examples of popular networks (at the time), he explains the mechanisms behind affiliate marketing through tools of the trade, infrastructural analyses, internal data on available offers and payouts, and revenue estimates. Kanich et al. [34] and McCoy et al. [53] executed detailed analyses of the economics behind major pharmaceutical and counterfeit software affiliate networks, studying customer purchasing behavior, as well as revenue for the networks and their affiliates. The latter study benefits from leaked ground truth of the networks themselves. Levchenko et al. [45] linked products advertised in spam to their respective affiliate networks, and studied to what extent they relied on shared network and payment infrastructure. As part of their systematization of the underground economy, Thomas et al. [70] describe how the affiliate marketing model is central to many organized cybercrime operations.

Further work identified affiliate marketing in detailed studies of specific malicious ecosystems. Caballero et al. [11] analyzed the affiliate structure behind 'pay-per-install' malware, identifying popular programs and the most prevalent affiliates. Kotzias et al. [40] and Thomas et al. [69] discussed popular 'pay-per-install' affiliate programs, including their offers and payouts. Stone-Gross et al. [64] studied the economics of fake antivirus software, quantifying revenue and detecting major actors in the ecosystem. Karami [36] analyzed 'Tower of Power', an affiliate program for herbal supplements and replica luxury goods. Through a database dump, they analyzed products on offer and their prices, customer and affiliate characteristics including revenue, and the domain name infrastructure. Clark and McCoy [13] analyzed the affiliate networks behind survey scams distributed through Facebook ads. They discuss affiliate network prevalence and tracking URL formats, and estimate revenue through affiliate account age and offer payouts. Liao et al. [47] found reputable affiliate networks to be implicated in spam campaigns hosted at cloud providers for 'long-tail' search engine optimization (i.e. targeting longer, less popular phrases). They study the most targeted affiliated programs and most active affiliates. A recent blog post by Palo Alto Networks' "Unit 42" [72] describes the identification and subsequent takedown of one affiliate marketing campaign abusing celebrity endorsements to advertise nutraceuticals. We seek to harmonize the 'common denominator' within these studies, showing how the affiliate marketing model underpins a diverse set of malicious and deceptive practices.

The affiliate marketing model plays host to other abuse types. In affiliate fraud or affiliate abuse, the affiliate tricks the user into unknowingly opening their affiliate link, after which the merchant undeservedly must pay commissions to the affiliate on future purchases. Only the affiliate therefore has a malicious intent while the (mostly legitimate) merchant is the victim and the end user is an unwitting participant. Edelman and Brandi [18] describe the economics of this abuse, while previous work has studied the technical means: loading malicious web pages opening the affiliate link in hidden elements [12, 61], injecting affiliate identifiers through browser extensions or binaries [35, 68], or redirecting users through the affiliate link notably using cybersquatting domains [5, 38, 43, 54, 56, 63, 73]. Finally, affiliates may also be in breach of advertising regulations due to insufficient disclosure of their affiliate status with legitimate merchants [10, 52, 66].

8. Conclusion and future work

We provide a first look at our ongoing data collection and analysis for the 'deceptive affiliate marketing' ecosystem. We show how the ecosystem covers many verticals that are implicated in deceptive practices. Based on data from 21 aggregators over the past year, we already can see that different verticals are valued differently, with for example cryptocurrency trading platforms yielding high commissions into the hundreds of dollars. This may imply that affiliates resort to more nefarious strategies if the payout could be higher. Countries and user agents are also differentiated, showing how research into malicious ecosystems must have sufficient coverage in order to be comprehensive and representative.

We plan to complement our existing data collection with additional crawling of the preview links, and with client-side data from the advertising channels that are used to reach consumers to determine the deceptive advertising strategies used by affiliates and confirm the prevalence of the ecosystem. We would then gain an end-to-end understanding of the deceptive affiliate marketing ecosystem. This could prove particularly useful for developing countermeasures to protect end users against these scams. Through the redirect chains between an ad and the final landing page, we may identify the tracking domains that affiliate networks use. These are a prime target for effective takedowns [44, 67]: these URLs cannot change without causing ad campaign and revenue interruptions, as otherwise URLs in ads from affiliates would no longer redirect to the offer. This data could also be used for client-side intervention, by enhancing blocklists with offer preview and tracking URLs, or even to provide users with transparency, for example with a browser extension that displays offer metadata for a particular ad. We also plan to collect metadata on the affiliate networks themselves (once again through aggregators), including provided payment methods and additional tracking domains. This would allow us to understand the financial and infrastructural elements supporting these networks, which are again potential targets for interventions.

Acknowledgments

We thank the anonymous reviewers and our shepherd Merve Sahin for their valuable feedback. This research is partially funded by the Research Fund KU Leuven, and by the Flemish Research Programme Cybersecurity. Victor Le Pochat holds a PhD Fellowship of the Research Foundation - Flanders (FWO 11A3421N). This work was supported by research credits from Google Cloud.

References

[1] *A Complete Overview of the Health & Beauty (Nutra) Vertical*. Advidi. Apr. 28, 2017. URL: https://advidi.com/overview-health-beauty-nutra-vertical/ (visited on 12/29/2020).

[2] *A Complete Overview of the Mainstream Vertical*. Advidi. Oct. 10, 2017. URL: https://advidi.com/complete-overview-mainstream-vertical/ (visited on 12/29/2020).

[3] *A Guide to the Finance Vertical*. Advidi. May 19, 2020. URL: https://advidi.com/a-guide-to-the-finance-vertical/ (visited on 12/29/2020).

[4] *Affiliate Marketers to Pay More Than $4 Million to Settle Charges That They Promoted a Fraudulent Business Coaching and Investment Scheme*. Federal Trade Commission. Mar. 5, 2020. URL: https://www.ftc.gov/news-events/press-releases/2020/03/affiliate-marketers-pay-more-4-million-settle-charges-they (visited on 11/23/2020).

[5] Pieter Agten, Wouter Joosen, Frank Piessens, and Nick Nikiforakis. "Seven Months' Worth of Mistakes: A Longitudinal Study of Typosquatting Abuse". In: *2015 Network and Distributed System Security Symposium*. Network and Distributed System Security Symposium. San Diego, CA: Internet Society, 2015. ISBN: 978-1-891562-38-9. DOI: 10.14722/ndss.2015.23058. URL: https://www.ndss-symposium.org/ndss2015/ndss-2015-programme/seven-months-worth-mistakes-longitudinal-study-typosquatting-abuse/ (visited on 10/25/2019).

[6] *Another Group of Marketers Behind Phony 'Gift Card' Text Spam Settles FTC Complaint*. Federal Trade Commission. Feb. 28, 2014. URL: https://www.ftc.gov/news-events/press-releases/2014/02/another-group-marketers-behind-phony-gift-card-text-spam-settles (visited on 07/09/2021).

[7] Emad Badawi, Guy-Vincent Jourdan, Gregor Bochmann, Iosif-Viorel Onut, and Jason Flood. "The "Game Hack" Scam". In: *19th International Conference on Web Engineering*. Ed. by Maxim Bakaev, Flavius Frasincar, and In-Young Ko. ICWE '19. Cham: Springer International Publishing, 2019, pp. 280–295. ISBN: 978-3-030-19274-7. DOI: 10.1007/978-3-030-19274-7_21.

[8] 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.

[9] *Ban Your Stereotypes about FB*. AdCombo. Apr. 26, 2018. URL: https://blog.adcombo.com/ban-your-stereotypes-about-fb/ (visited on 07/06/2021).

[10] Laura E. Bladow. "Worth the Click: Why Greater FTC Enforcement Is Needed to Curtail Deceptive Practices in Influencer Marketing Notes". In: *William & Mary Law Review* 59.3 (2017–2018), [i]–1164. URL: https://heinonline.org/HOL/P?h=hein.journals/wmlr59&i=1157 (visited on 07/20/2020).

[11] Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. "Measuring Pay-per-Install: The Commoditization of Malware Distribution". In: *20th USENIX Security Symposium*. 20th USENIX Security Symposium. USENIX Security '11. 2011.

[12] Neha Chachra, Stefan Savage, and Geoffrey M. Voelker. "Affiliate Crookies: Characterizing Affiliate Marketing Abuse". In: *2015 ACM Conference on Internet Measurement Conference*. IMC '15. 2015, pp. 41–47. ISBN: 978-1-4503-3848-6. DOI: 10.1145/2815675.2815720. URL: http://dl.acm.org/citation.cfm?doid=2815675.2815720 (visited on 09/15/2019).

[13] Jason W. Clark and Damon McCoy. "There Are No Free iPads: An Analysis of Survey Scams as a Business". In: *6th USENIX Workshop on Large-Scale Exploits and Emergent Threats*. 6th USENIX Workshop on Large-Scale Exploits and Emergent Threats. LEET '13. 2013.

[14] Graham Cluley. *Beware! Free Apple Products Used as Lure in Text Scams*. Naked Security. Aug. 8, 2012. URL: https://nakedsecurity.sophos.com/2012/08/08/free-apple-products-text-scam/ (visited on 01/13/2020).

[15] Jerome Dangu. *Fake Celebrity-Endorsed Bitcoin Scam Abuses Ad Tech to Net $1M in 1 Day*. Medium. Jan. 27, 2020. URL: https://blog.confiant.com/fake-celebrity-endorsed-scam-abuses-ad-tech-to-net-1m-in-one-day-ffe330258e3c (visited on 01/30/2020).

[16] Nicholas De Rosa, Jeff Yates, and Brigitte Noël. *Un empire montréalais de l'arnaque en ligne*. Radio-Canada.ca. June 21, 2021. URL: https://ici.radio-canada.ca/recit-numerique/2140/adcenter-hyuna-philip-keezer-streaming-concours (visited on 06/14/2021).

[17] Mateja Durovic and Hans W. Micklitz. "International Law on (Un)Fair Commercial Practices". In: *Internationalization of Consumer Law: A Game Changer*. Ed. by Mateja Durovic and Hans W. Micklitz. SpringerBriefs in Political Science. Cham: Springer International Publishing, 2017, pp. 25–48. ISBN: 978-3-319-45312-5. DOI: 10.1007/978-3-319-45312-5_3. URL: https://doi.org/10.1007/978-3-319-45312-5_3 (visited on 07/06/2021).

[18] Benjamin Edelman and Wesley Brandi. "Risk, Information, and Incentives in Online Affiliate Marketing". In: *Journal of Marketing Research* 52.1 (Feb. 1, 2015), pp. 1–12. ISSN: 0022-2437. DOI: 10.1509/jmr.13.0472. URL: https://doi.org/10.1509/jmr.13.0472 (visited on 09/16/2019).

[19] Shehroze Farooqi, Álvaro Feal, Tobias Lauinger, Damon McCoy, Zubair Shafiq, and Narseo Vallina-Rodriguez. "Understanding Incentivized Mobile App Installs on Google Play Store". In: *Proceedings of the ACM Internet Measurement Conference*. IMC '20: ACM Internet Measurement Conference. Virtual Event USA: ACM, Oct. 27, 2020, pp. 696–709. ISBN: 978-1-4503-8138-3. DOI: 10.1145/3419394.3423662. URL: https://dl.acm.org/doi/10.1145/3419394.3423662 (visited on 10/26/2020).

[20] Zeke Faux. "'They Go out and Find the Morons for Me'". In: *Bloomberg Businessweek* 4564 (2018), pp. 56–61. ISSN: 00077135. URL: http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=128750674&site=ehost-live&scope=site.

[21] Federal Trade Commission. *Federal Court Rules Affiliate Marketing Network and Its Parent Company Must Turn Over $11.9 Million They Received From Deceptive Marketing Scheme*. Federal Trade Commission. Apr. 6, 2015. URL: https://www.ftc.gov/news-events/press-releases/2015/04/federal-court-rules-affiliate-marketing-network-its-parent (visited on 01/20/2020).

[22] Federal Trade Commission. *FTC Seeks to Halt 10 Operators of Fake News Sites from Making Deceptive Claims About Acai Berry Weight Loss Products*. Federal Trade Commission. Apr. 19, 2011. URL: https://www.ftc.gov/news-events/press-releases/2011/04/ftc-seeks-halt-10-operators-fake-news-sites-making-deceptive (visited on 01/20/2020).

[23] Federal Trade Commission. *FTC Settlement Bars Spam Email Marketing, Baseless Weight-Loss Claims by Diet-Pill Operation*. Federal Trade Commission. Mar. 16, 2017. URL: https://www.ftc.gov/news-events/press-releases/2017/03/ftc-settlement-bars-spam-email-marketing-baseless-weight-loss (visited on 01/20/2020).

[24] Federal Trade Commission. *Internet Marketers of Dietary Supplement and Skincare Products Banned from Deceptive Advertising and Billing Practices*. Federal Trade Commission. Nov. 15, 2017. URL: https://www.ftc.gov/news-events/press-releases/2017/11/internet-marketers-dietary-supplement-skincare-products-banned (visited on 01/20/2020).

[25] *FTC Announces Crackdown on Deceptively Marketed CBD Products*. Federal Trade Commission. Dec. 16, 2020. URL: https://www.ftc.gov/news-events/press-releases/2020/12/ftc-announces-crackdown-deceptively-marketed-cbd-products (visited on 01/08/2021).

[26] *FTC Charges Marketers Used Massive Spam Campaign To Pitch Bogus Weight-Loss Products*. Federal Trade Commission. June 6, 2016. URL: https://www.ftc.gov/news-events/press-releases/2016/06/ftc-charges-marketers-used-massive-spam-campaign-pitch-bogus (visited on 11/23/2020).

[27] *FTC Charges Online Marketing Scheme with Deceiving Shoppers*. Federal Trade Commission. Aug. 4, 2017. URL: https://www.ftc.gov/news-events/press-releases/2017/08/ftc-charges-online-marketing-scheme-deceiving-shoppers (visited on 01/20/2020).

[28] Kinga Gawron. *Ranking of the Best Affiliate Marketing Niches for 2021*. Zeropark Blog. Jan. 6, 2021. URL: https://zeropark.com/blog/affiliate-marketing-best-niches-2021/ (visited on 01/22/2021).

[29] *Geniux Dietary Supplement Sellers Barred from Unsupported Cognitive Improvement Claims*. Federal Trade Commission. Apr. 10, 2019. URL: https://www.ftc.gov/news-events/press-releases/2019/04/geniux-dietary-supplement-sellers-barred-unsupported-cognitive (visited on 01/20/2020).

[30] Sarra Gray. *Lidl Shoppers Urged to Watch out for £500 Gift Card Scam - 'Don't Share Personal Details'*. Express.co.uk. May 20, 2021. URL: https://www.express.co.uk/life-style/life/

1438930/lidl-uk-scam-warning-gift-cards-email-latest-news (visited on 06/11/2021).

[31] Thomas J. Holt and Eric Lampke. “Exploring Stolen Data Markets Online: Products and Market Forces”. In: *Criminal Justice Studies* 23.1 (Mar. 1, 2010), pp. 33–50. ISSN: 1478-601X. DOI: 10.1080/14786011003634415. URL: https://doi.org/10.1080/14786011003634415 (visited on 01/12/2020).

[32] Luca Invernizzi, Kurt Thomas, Alexandros Kapravelos, Oxana Comanescu, Jean-Michel Picod, and Elie Bursztein. “Cloak of Visibility: Detecting When Machines Browse a Different Web”. In: *2016 IEEE Symposium on Security and Privacy*. 2016 IEEE Symposium on Security and Privacy (SP). S&P ’16. San Jose, CA: IEEE, May 2016, pp. 743–758. ISBN: 978-1-5090-0824-7. DOI: 10.1109/SP.2016.50. URL: http://ieeexplore.ieee.org/document/7546533/ (visited on 12/09/2019).

[33] Judy Johnson. *Aldi Scam: Supermarket Shares Warning over £250 Voucher Scam Message*. Express.co.uk. May 5, 2020. URL: https://www.express.co.uk/life-style/food/1278219/aldi-scam-message-voucher-coupon (visited on 06/11/2021).

[34] Chris Kanich, Nicholas Weaver, Damon McCoy, Tristan Halvorson, Christian Kreibich, Kirill Levchenko, Vern Paxson, Geoffrey M Voelker, and Stefan Savage. “Show Me the Money: Characterizing Spam-Advertised Revenue”. In: *20th USENIX Security Symposium*. USENIX Security ’11. 2011.

[35] Alexandros Kapravelos, Chris Grier, Neha Chachra, Christopher Kruegel, Giovanni Vigna, and Vern Paxson. “Hulk: Eliciting Malicious Behavior in Browser Extensions”. In: *23rd USENIX Security Symposium*. USENIX Security ’14. 2014, p. 15.

[36] Mohammad Karami, Shiva Ghaemi, and Damon McCoy. “Folex: An Analysis of an Herbal and Counterfeit Luxury Goods Affiliate Program”. In: *2013 APWG eCrime Researchers Summit*. 2013 APWG eCrime Researchers Summit. eCrime ’13. Sept. 2013, pp. 1–9. DOI: 10.1109/eCRS.2013.6805782.

[37] Amin Kharraz, William Robertson, and Engin Kirda. “Surveylance: Automatically Detecting Online Survey Scams”. In: *2018 IEEE Symposium on Security and Privacy*. 2018 IEEE Symposium on Security and Privacy (SP). S&P ’18. San Francisco, CA: IEEE, May 2018, pp. 70–86. ISBN: 978-1-5386-4353-2. DOI: 10.1109/SP.2018.00044. URL: https://ieeexplore.ieee.org/document/8418597/ (visited on 09/15/2019).

[38] Panagiotis Kintis, Najmeh Miramirkhani, Charles Lever, Yizheng Chen, Rosa Romero-Gómez, Nikolaos Pitropakis, Nick Nikiforakis, and Manos Antonakakis. “Hiding in Plain Sight: A Longitudinal Study of Combosquatting Abuse”. In: *2017 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’17. Dallas, Texas, USA: Association for Computing Machinery, Oct. 30, 2017, pp. 569–586. ISBN: 978-1-4503-4946-8. DOI: 10.1145/3133956.3134002. URL: https://doi.org/10.1145/3133956.3134002 (visited on 01/17/2020).

[39] Takashi Koide, Daiki Chiba, and Mitsuaki Akiyama. “To Get Lost Is to Learn the Way: Automatically Collecting Multi-Step Social Engineering Attacks on the Web”. In: *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*. ASIA CCS ’20. New York, NY, USA: Association for Computing Machinery, Oct. 5, 2020, pp. 394–408. ISBN: 978-1-4503-6750-9. DOI: 10.1145/3320269.3384714. URL: https://doi.org/10.1145/3320269.3384714 (visited on 10/13/2020).

[40] Platon Kotzias, Leyla Bilge, Sofia Antipolis, and Juan Caballero. “Measuring PUP Prevalence and PUP Distribution through Pay-Per-Install Services”. In: *25th USENIX Security Symposium*. USENIX Security ’16. 2016.

[41] Magdalena Kukułka. *How to Run Affiliate Marketing Antivirus Campaigns with Push Traffic?* Zeropark Blog. June 22, 2021. URL: https://zeropark.com/blog/affiliate-marketing-antivirus-campaigns-push-traffic/ (visited on 07/06/2021).

[42] Magdalena Kukułka. *Top Affiliate Offers in 2021*. Zeropark Blog. Feb. 2, 2021. URL: https://zeropark.com/blog/top-affiliate-offers-in-2021/ (visited on 07/06/2021).

[43] Victor Le Pochat, Tom Van Goethem, and Wouter Joosen. “A Smörgåsbord of Typos: Exploring International Keyboard Layout Typosquatting”. In: *2019 IEEE Security and Privacy Workshops*. 4th International Workshop on Traffic Measurements for Cybersecurity. SPW ’19. San Francisco, CA, USA: IEEE, May 2019, pp. 187–192. ISBN: 978-1-72813-508-3. DOI: 10.1109/SPW.2019.00043. URL: https://ieeexplore.ieee.org/document/8844633/ (visited on 10/25/2019).

[44] Nektarios Leontiadis, Tyler Moore, and Nicolas Christin. “Measuring and Analyzing Search-Redirection Attacks in the Illicit Online Prescription Drug Trade”. In: *20th USENIX Security Symposium*. USENIX Security ’11. 2011, p. 17.

[45] Kirill Levchenko, Andreas Pitsillidis, Neha Chachra, Brandon Enright, Márk Felegyhazi, Chris Grier, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, Damon McCoy, Nicholas Weaver, Vern Paxson, Geoffrey M. Voelker, and Stefan Savage. “Click Trajectories: End-to-End Analysis of the Spam Value Chain”. In: *2011 IEEE Symposium on Security and Privacy*. 2011 IEEE Symposium on Security and Privacy. S&P ’11. May 2011, pp. 431–446. DOI: 10.1109/SP.2011.24.

[46] *Leveraging Facebook for Affiliate Marketing in 2020*. Everad. July 15, 2020. URL: https://blog.everad.com/en/leveraging-facebook-for-affiliate-marketing-in-2020/ (visited on 07/06/2021).

[47] Xiaojing Liao, Chang Liu, Damon McCoy, Elaine Shi, Shuang Hao, and Raheem Beyah. “Characterizing Long-Tail SEO Spam on Cloud Web Hosting Services”. In: *25th International Conference on World Wide Web*. WWW ’16. ACM Press, 2016, pp. 321–332. ISBN: 978-1-4503-4143-1. DOI: 10.1145/2872427.2883008. URL: http://dl.acm.org/citation.cfm?doid=2872427.2883008 (visited on 09/15/2019).

[48] Xiaojing Liao, Kan Yuan, XiaoFeng Wang, Zhongyu Pei, Hao Yang, Jianjun Chen, Haixin Duan, Kun Du, Eihal Alowaisheq, Sumayah Alrwais, Luyi Xing, and Raheem Beyah. “Seeking Nonsense, Looking for Trouble: Efficient Promotional-Infection Detection through Semantic Inconsistency Search”. In: *2016 IEEE Symposium on Security and Privacy*. 2016 IEEE Symposium on Security and Privacy (SP). S&P ’16. San Jose, CA: IEEE, May 2016, pp. 707–723. ISBN: 978-1-5090-0824-7. DOI: 10.1109/SP.2016.48. URL: http://ieeexplore.ieee.org/document/7546531/ (visited on 11/29/2019).

[49] Stephan Lindburg. *Quantum AI Review, Fake Quantum AI SCAM By Elon Musk Exposed!* Scam Crypto Robots. Nov. 14, 2019. URL: https://scamcryptorobots.com/quantum-ai-review-scam/ (visited on 01/13/2020).

[50] *Marketers Behind Fake News Sites Settle FTC Charges of Deceptive Advertising*. Federal Trade Commission. Nov. 14, 2012. URL: https://www.ftc.gov/news-events/press-releases/2012/11/marketers-behind-fake-news-sites-settle-ftc-charges-deceptive (visited on 07/05/2021).

[51] James Martin and Nicolas Christin. “Ethics in Cryptomarket Research”. In: *International Journal of Drug Policy*. Drug Cryptomarkets 35 (Sept. 1, 2016), pp. 84–91. ISSN: 0955-3959. DOI: 10.1016/j.drugpo.2016.05.006. URL: https://www.sciencedirect.com/science/article/pii/S0955395916301608 (visited on 06/11/2021).

[52] Arunesh Mathur, Arvind Narayanan, and Marshini Chetty. “Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest”. In: *Proceedings of the ACM on Human-Computer Interaction* 2 (CSCW Nov. 1, 2018), 119:1–119:26. ISSN: 2573-0142. DOI: 10.1145/3274388. arXiv: 1809.00620. URL: http://doi.acm.org/10.1145/3274388 (visited on 09/15/2019).

[53] Damon McCoy, Andreas Pitsillidis, Grant Jordan, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey M Voelker, Stefan Savage, and Kirill Levchenko. “PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs”. In: *21st USENIX Security Symposium*. USENIX Security ’12. 2012, p. 16.

[54] Tyler Moore and Benjamin Edelman. “Measuring the Perpetrators and Funders of Typosquatting”. In: *Financial Cryptography and Data Security*. Vol. 6052. FC ’10. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 175–191. ISBN: 978-3-642-14576-6 978-3-642-14577-3. DOI: 10.1007/978-3-642-14577-3_15. URL: http://link.springer.com/10.1007/978-3-642-14577-3_15 (visited on 12/20/2019).

[55] Dasha Nazarova and Irina Bystrova. *Supreme Guide to Affiliate Marketing Verticals*. RedTrack, May 19, 2020. URL: https://redtrackmarketing.s3.eu-central-1.amazonaws.com/Affiliate+Marketing+Verticals+Guide.pdf.

[56] Nick Nikiforakis, Steven Van Acker, Wannes Meert, Lieven Desmet, Frank Piessens, and Wouter Joosen. “Bitsquatting: Exploiting Bit-Flips for Fun, or Profit?” In: *22nd International Conference on World Wide Web*. The 22nd International Conference. WWW ’13. Rio de Janeiro, Brazil: ACM Press, 2013, pp. 989–998. ISBN: 978-1-4503-2035-1. DOI: 10.1145/2488388.2488474. URL: http://dl.acm.org/citation.cfm?doid=2488388.2488474 (visited on 01/12/2020).

[57] Sergio Pastrana, Daniel R. Thomas, Alice Hutchings, and Richard Clayton. “CrimeBB: Enabling Cybercrime Research

on Underground Forums at Scale”. In: *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW ’18*. The 2018 World Wide Web Conference. Lyon, France: ACM Press, 2018, pp. 1845–1854. ISBN: 978-1-4503-5639-8. DOI: 10.1145/3178876.3186178. URL: http://dl.acm.org/citation.cfm?doid=3178876.3186178 (visited on 06/11/2021).

[58] *Products and Services - Affiliates*. Advidi. URL: https://advidi.com/products-and-services-affiliates/ (visited on 06/10/2021).

[59] Mathias Reimann. “Liability for Defective Products at the Beginning of the Twenty-First Century: Emergence of a Worldwide Standard”. In: *American Journal of Comparative Law* 51.4 (2003), pp. 751–838. URL: https://heinonline.org/HOL/P?h=hein.journals/amcomp51&i=763 (visited on 07/06/2021).

[60] Dmitry Samosseiko. “The Partnerka - What Is It, and Why Should You Care?” In: *19th Virus Bulletin International Conference*. 19th Virus Bulletin Conference. Sept. 2009, pp. 115–120.

[61] Peter Snyder and Chris Kanich. “No Please, After You: Detecting Fraud in Affiliate Marketing Networks”. In: *14th Annual Workshop on the Economics of Information Security*. 14th Annual Workshop on the Economics of Information Security. WEIS ’15. 2015, p. 11.

[62] Kyle Soska and Nicolas Christin. “Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem”. In: *24th USENIX Security Symposium*. USENIX Security ’15. 2015, pp. 33–48.

[63] Jeffrey Spaulding, Shambhu Upadhyaya, and Aziz Mohaisen. “The Landscape of Domain Name Typosquatting: Techniques and Countermeasures”. In: *11th International Conference on Availability, Reliability and Security*. 2016 11th International Conference on Availability, Reliability and Security (ARES). ARES ’16. Aug. 2016, pp. 284–289. DOI: 10.1109/ARES.2016.84.

[64] Brett Stone-Gross, Ryan Abman, Richard A. Kemmerer, Christopher Kruegel, Douglas G. Steigerwald, and Giovanni Vigna. “The Underground Economy of Fake Antivirus Software”. In: *Economics of Information Security and Privacy III*. Ed. by Bruce Schneier. Springer New York, 2013, pp. 55–78. ISBN: 978-1-4614-1981-5.

[65] Gianluca Stringhini. “Adversarial Behaviour”. In: *The Cyber Security Body of Knowledge*. Ed. by Awais Rashid, Howard Chivers, George Danezis, Emil Lupu, and Andrew Martin. 1.0. Oct. 31, 2019, pp. 223–249.

[66] Michael Swart, Ylana Lopez, Arunesh Mathur, and Marshini Chetty. “Is This An Ad?: Automatically Disclosing Online Endorsements On YouTube With AdIntuition”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI ’20. Honolulu, HI, USA: Association for Computing Machinery, Apr. 21, 2020, pp. 1–12. ISBN: 978-1-4503-6708-0. DOI: 10.1145/3313831.3376178. URL: https://doi.org/10.1145/3313831.3376178 (visited on 04/28/2020).

[67] Janos Szurdi, Meng Luo, Brian Kondracki, Nick Nikiforakis, and Nicolas Christin. “Where Are You Taking Me? Understanding Abusive Traffic Distribution Systems”. In: WWW. 2021, p. 12.

[68] Kurt Thomas, Elie Bursztein, Chris Grier, Grant Ho, Nav Jagpal, Alexandros Kapravelos, Damon McCoy, Antonio Nappa, Vern Paxson, Paul Pearce, Niels Provos, and Moheeb Abu Rajab. “Ad Injection at Scale: Assessing Deceptive Advertisement Modifications”. In: *2015 IEEE Symposium on Security and Privacy*. 2015 IEEE Symposium on Security and Privacy. S&P ’15. May 2015, pp. 151–167. DOI: 10.1109/SP.2015.17.

[69] Kurt Thomas, Juan A Elices Crespo, Ryan Rasti, Jean-Michel Picod, Cait Phillips, Marc-Andre Decoste, Chris Sharp, Fabio Tirelo, Ali Tofigh, Marc-Antoine Courteau, Lucas Ballard, Robert Shield, Nav Jagpal, Moheeb Abu Rajab, Panayiotis Mavrommatis, Niels Provos, Elie Bursztein, and Damon McCoy. “Investigating Commercial Pay-Per-Install and the Distribution of Unwanted Software”. In: *25th USENIX Security Symposium*. USENIX Security ’16. 2016, p. 19.

[70] Kurt Thomas, Danny Yuxing Huang, David Wang, Elie Bursztein, Chris Grier, Thomas J Holt, Christopher Kruegel, Damon McCoy, Stefan Savage, and Giovanni Vigna. “Framing Dependencies Introduced by Underground Commoditization”. In: *14th Annual Workshop on the Economics of Information Security*. 14th Annual Workshop on the Economics of Information Security. WEIS ’15. 2015.

[71] Phani Vadrevu and Roberto Perdisci. “What You See Is NOT What You Get: Discovering and Tracking Social Engineering Attack Campaigns”. In: *2019 Internet Measurement Conference*. IMC ’19. 2019, p. 14. DOI: 10.1145/3355369.3355600.

[72] 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/ (visited on 10/07/2019).

[73] Yuwei Zeng, Tianning Zang, Yongzheng Zhang, Xunxun Chen, and YiPeng Wang. “A Comprehensive Measurement Study of Domain-Squatting Abuse”. In: *2019 IEEE International Conference on Communications*. ICC 2019 - 2019 IEEE International Conference on Communications (ICC). ICC ’19. Shanghai, China: IEEE, May 2019, pp. 1–6. ISBN: 978-1-5386-8088-9. DOI: 10.1109/ICC.2019.8761388. URL: https://ieeexplore.ieee.org/document/8761388/ (visited on 09/15/2019).

Appendix A.
Visual examples of affiliate offers

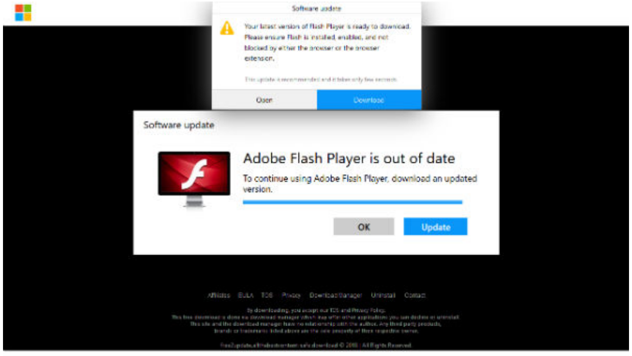

Figure 7. 'Mac Flash Player' offer; vertical: software; payout: $4.50

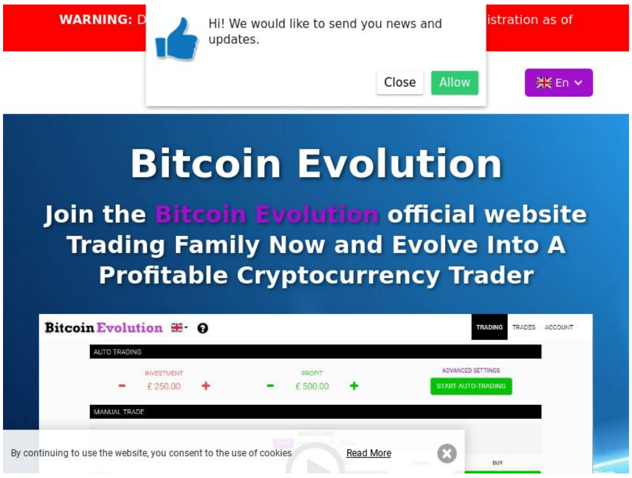

Figure 8. 'Bitcoin Evolution' offer; vertical: crypto; payout: $1,595.00

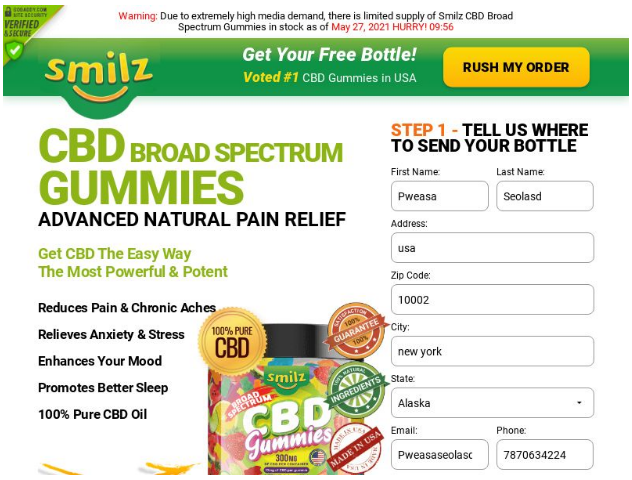

Figure 9. 'Smilz CBD Gummies' offer; vertical: CBD; payout: $130.00

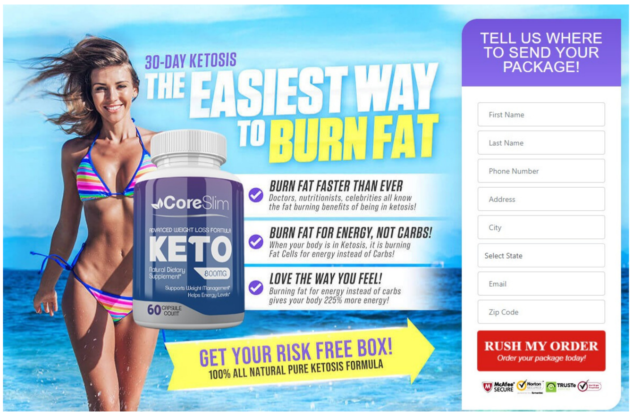

Figure 10. 'Core Slim Keto' offer; vertical: diet; payout: $110.50

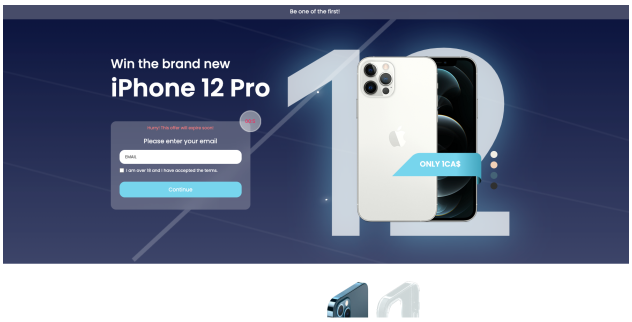

Figure 11. 'iPhone 12 Pro' offer; vertical: sweepstakes; payout: $45.00