Partnërka in Crime: Characterizing Deceptive Affiliate Marketing Offers

Victor Le Pochat[1][0000−0003−2297−8328], Cameron Ballard[2], Lieven Desmet[1][0000−0001−5155−7472], Wouter Joosen[1][0000−0002−7710−5092], Damon McCoy[2][0000−0001−7386−7260], and Tobias Lauinger[2][0000−0002−5779−0643]

[1] DistriNet, KU Leuven
[2] New York University

Abstract. The deceptive affiliate marketing ecosystem enables a variety of online scams causing consumers to lose money or personal data. In this model, affiliates promote deceptive products and services on behalf of merchants in exchange for a commission, mediated by affiliate networks. We monitor the ecosystem holistically by taking the vantage point of affiliates and collecting ground truth from 23 aggregators that list deceptive products and services available for promotion across scam types and affiliate networks. Using our novel longitudinal data set, we characterize the ecosystem by taxonomizing the 9 main categories of deceptive products and services composing the ecosystem, and describing the main tactics used to mislead consumers. We quantify the extent of the nearly 450,000 offers in the ecosystem and the differences in the value that is attached to different types of scams, monetization models, and countries. Finally, we identify core affiliate networks and analyze longitudinal trends to track the dynamics of the ecosystem over time. The more complete coverage provided by our novel data set enables not only a broader understanding of the ecosystem, but also adds insights and metadata for developing earlier, data-driven interventions to protect consumers.

1 Introduction

Internet users are regularly exposed to deceptive products and services online [85, 105]. Scammers seek to defraud users through low-quality products and services, such as questionable dietary supplements [11] or cryptocurrency investment platforms [93], or seek to trick users into installing potentially unwanted software [64, 85, 97, 102] or into disclosing personal data, such as through fake contests [18]. Simultaneously, these products and services are promoted using misleading advertisements, employing tactics such as fake celebrity endorsements [11, 23, 28, 114], manipulative "clickbait" headlines [121], and other "dark patterns" [76]. Despite long-term awareness of these harms (since at least the late 2000s [91, 99]), numerous regulatory interventions [29–31, 33–41], and media investigations [14, 24, 61, 65, 80, 81, 94, 106, 119], this wide range of deceptive operations continues to defraud consumers worldwide at a massive scale. Individual operations reach revenues of hundreds of millions of U.S. dollars [24, 28, 40, 78], and total spend on deceptive ads through one large tracker was estimated at $1.7 billion a year [28].

While the different scams may seemingly be unconnected, we observe that many are embedded in the **deceptive affiliate marketing** ecosystem [11, 30, 101, 114].[3] In this model, *merchants* outsource the marketing of their low-quality products and services ('offers') to independent *affiliates*, who exploit advertising channels to run deceptive ads that promote these offers, in return for a commission on each successful conversion. While legitimate affiliate marketing programs are common, prior work documented how this advertising model is also central to many cybercriminal activities [91, 101], for example illegal pharmaceuticals [58, 72, 78, 91], counterfeit luxury goods [59, 107], and ransomware [53]. We reveal how it is also used in promoting and propagating other online scams.

In this paper, we comprehensively measure *how* and by *whom* the deceptive affiliate marketing ecosystem is operated. We use a novel ground-truth vantage point on the ecosystem to create a large-scale longitudinal data set that enables us to identify the relations between merchants and affiliates. We take the point of view of affiliates themselves, discovering deceptive products and services on so-called *aggregators*, which the affiliates use to search offers across many merchants at once. The advantage of this novel data set is that it provides us with complete 'insider' data on the breadth of deceptive products on offer, across all verticals (categories), countries, device configurations, and targeted populations, making our work global and longitudinal. Moreover, these aggregators publish metadata that is unavailable in traffic-based measurements, including the commission paid out to affiliates, which enables us to quantify the monetary dynamics of the ecosystem and how the different scams within it compare, which was previously impossible to achieve.

Our novel vantage point allows us to overcome inherent limitations in coverage from prior research that measures online scams from a user's perspective through the collection of traffic on deceptive ads or websites [43, 72, 90, 98, 105], which leads to only observing a subset of traffic sources, ads and scams. Instead, we provide a broad view on the connection between these scams and the deceptive affiliate marketing ecosystem that supports them. In addition, in contrast to prior work on affiliate models in cybercrime, we identify how a third, intermediary actor emerges. *Affiliate networks* serve as the primary point of contact for both merchants and their affiliates, and mediate their interactions by supporting them in attracting affiliates or discovering merchants respectively, and subsequently tracking sales and paying commissions. We describe how these affiliate networks run as large operations that have a central role in the ecosystem [19], which makes them a valuable target to prioritize and maximize the impact of technical, financial, and legal interventions [62].

To characterize the current state of the deceptive affiliate marketing ecosystem, we collect an extensive data set from 23 aggregators over more than 4 years. We analyze nearly 450,000 deceptive offers (i.e., products and services available for promotion) across 1,165 affiliate networks, which we derive after normalizing offers across aggregators and filtering out legitimate offers. We develop a new taxonomy with 9 verticals (i.e., categories) of products and services that constitute the

[3] Affiliate programs are also known as "partnërka" in Russian [91].

deceptive affiliate marketing ecosystem. We observe that merchants use a wide variety of deceptive tactics to mislead consumers across these verticals: making exaggerated claims, misrepresenting endorsements, using dark patterns [76] to pressure users, and tricking unsuspecting customers into authorizing (recurring) payments. In general, products and services are of low quality and overpriced. Using our taxonomy, we then quantify how the payouts to affiliates depend on the type of product and what a customer has to do or pay to "complete" the offer, to compare the monetary incentives from the affiliate's point of view across the range of scams in the ecosystem. The most lucrative offers relate to investment scams, capitalizing on trends such as cryptocurrency, with payouts in the hundreds of U.S. dollars once customers make a deposit to a fraudulent trading platform. Physical goods (e.g., useless health and beauty products) and subscription services (e.g., fake dating sites) also command relatively higher payouts, while virtual goods, app installs, or providing personal data pay single-digit figures, or less. At the country level, offer selections and payouts differ based on income levels, monetization methods, and regulations. Our data also allows us to identify the largest affiliate networks that mediate the ecosystem, and quantify their scale. Finally, we observe longitudinal trends in the emergence of offers, verticals, and networks, showing the dynamism of the ecosystem.

In summary, our contributions are:

- We study the deceptive affiliate marketing ecosystem from the previously unexplored vantage point of the affiliates themselves, and present our ongoing data collection of a novel data source: 23 'aggregators' that allow us to discover offers across all verticals and countries, allowing for early intervention (section 3).
- We develop a taxonomy of 9 verticals (section 4), and label offers to establish their deceptive nature. We find a broad spectrum of tactics designed to mislead users and coax them into handing over money or personal data.
- We collect 449,891 offers across 1,165 affiliate networks to quantify the breadth of the ecosystem (section 5). We find different valuations across verticals, monetization models (e.g., physical vs. virtual goods), and countries.
- We discuss how specialization drives access to the ecosystem, contributing to its growth, and how intervention strategies can prevent users from being harmed by deceptive affiliate marketing offers (section 6).

2 Key concepts

We introduce key concepts and terms, as they are used in the ecosystem itself, based on ecosystem guides and resources [7, 8, 22, 52]. Figure 1 summarizes the main ecosystem players, and the flow of payments and traffic between them.

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) who will promote the merchant's product. Merchants and affiliates find each other through **affiliate networks** (advertiser network) that act as intermediaries. These

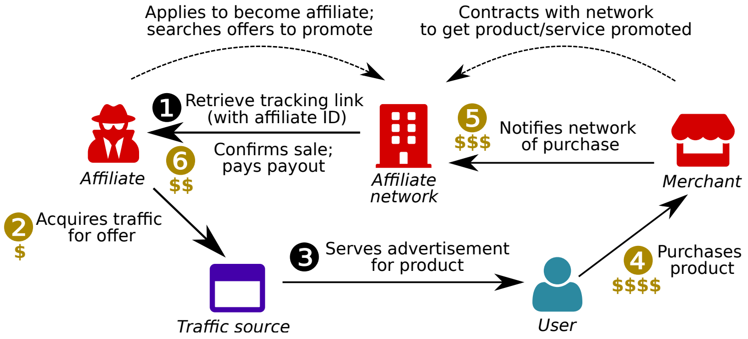

Fig. 1: A typical flow for a cost-per-sale offer shows how the three main types of ecosystem players engage to promote an offer and collect payment.

networks also provide the technical and financial infrastructure that ultimately pays the affiliate when they successfully promoted the merchant's product. There is a many-to-many-to-many relation in the ecosystem: affiliates can promote offers of any number of merchants, a merchant's offers can be promoted by any number of affiliates. a network can accept any number of affiliates and merchants, and affiliates and merchants can join any number of networks. Next to these three players, third parties provide supporting services, such as ad tracking tools, landing page builders, forums, or the offer aggregators that we use (Section 3).

Monetization Merchants post **offers** 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**, which are divided into **tiers** to reflect their perceived wealth, and therefore attractiveness. Certain advertising channels may also be (dis)allowed. Additionally, an offer stipulates the conditions and payment amounts for a successful **conversion**, i.e., the action that a customer has to take for the affiliate to be paid out (e.g., a payment or submission of an email address). Usually, the affiliate receives a one-time fixed-amount **payout** (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. The merchant will pay these payouts plus a service fee to the affiliate network, who passes the payout onto the affiliate. Two models prevail for awarding a payout: **cost per sale** (CPS, pay per sale, PPS) where a customer must pay for 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. In the CPA model, a consumer action is required for the merchant to pay the affiliate, distinguishing it from advertising models where the merchant pays per ad click or impression.

Table 1: An example of metadata available for an offer, upon which we base our analysis. Fields in *italics* are computed during data postprocessing. Note that not all metadata is always available, as quantified in subsection 5.1.

Field	Value	Field	Value
Name	QuantumAi (AU) (CPS)	Payout	560 USD
Normalized name	*quantumai*	Conversion type	CPS (cost-per-sale)
Affiliate network	AdsMain	Countries	{Australia}
Verticals	{Crypto offers, Finance}	Aggregators	{Affplus, Affscanner, OfferVault}
Normalized vertical	*Finance*	Observed on	{2020-05-06, 2020-05-07, ...}

Redirect chain The affiliate receives a **tracking link** (affiliate link) from the affiliate network, which contains a unique affiliate ID. The affiliate then selects a **traffic source** through which they promote the offer: e.g., their own website, email, or advertisements on platforms such as search engines or social networks. The customer is then redirected from the affiliate link to the **landing page** (offer page, lander) maintained by the merchant, where the customer can complete the offer. The merchant reports the affiliate ID that triggered the conversion to the network, which will pay the offer's payout to the affiliate.

3 Data collection and processing

We scrape offer data from *"offer aggregators"* that operate as search engines, who collect offers from multiple affiliate networks (and therefore across multiple merchants) and publish them in one interface where affiliates can search and filter them. Affiliates can find interesting offers and the affiliate networks that manage them through the aggregator, but will still need to separately join those networks. Aggregators provide broader coverage than individual affiliate networks, and unlike the latter, make their data public without requiring registration with an affiliate network, which may entail a vetting procedure. In contrast to traffic-based measurements, aggregators index current and past offers, and provide key metadata such as the network, vertical, countries, and conversion terms. Table 1 shows an example of the offer metadata that aggregators make available.

3.1 Aggregator discovery

We employ a multi-tiered approach to discover the most common 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 (AffiliateFix, affLift, BlackHatWorld) and sponsor lists for major affiliate conferences (Affiliate Summit, Affiliate World) to find additional aggregators. We continue searching aggregators until we reach saturation. Overall, we cover 16 English- and 7 Russian-language aggregators of deceptive offers (Table 2). In addition, we collect data from 4 aggregators for reputable mainstream brands, in order to filter reputable offers from our data set (subsection 3.4).

Table 2: Overview of the offer aggregators for which we collected data. We list the number of records retrieved (including duplicates of the same network offer retrieved every day), the number of network offers observed (merged on normalized offer name), and the number of networks observed.

Aggregator	# networks	# network offers	# records
English-language aggregators of deceptive offers (16)			
AdNetworksHub	2	386	85,273
Affbank	144	257,464	62,041,167
Affhomes	49	14,542	11,413,418
affNext	9	3,097	2,494,100
Affplus	229	338,132	41,053,017
Affpub	75	20,945	3,306,317
Affscanner	62	69,769	14,826,290
BestAffiliatePrograms	56	50,444	13,757,137
BigFishOffers	7	56,869	39,504,203
Click4ads	147	72,164	7,706,208
Horje	–	24,662	43,523,554
oDigger	71	46,938	8,068,918
OfferLibrary	24	44,020	16,947,738
OfferVault	240	190,609	25,357,080
WOW TRK	33	47,638	10,146,934
XOffers	126	6,083	15,183
Russian-language aggregators of deceptive offers (7)			
ActualTraffic	74	22,334	4,341,738
Admakler	8	2,656	2,721,448
atlasio	36	8,569	3,678,416
AVF	32	3,044	984,910
CPA Daily	116	82,882	5,592,100
CPA Inform	38	34,684	17,541,949
Partnerkin	90	114,023	16,784,124
English-language aggregators of reputable offers (4)			
Affi.io	125	200,634	5,673,271
FMTC	45	52,993	3,216,827
LinkPizza	11	13,927	2,689,686
Publisher Rest	59	19,403	8,988,989

3.2 Retrieval

We extract available offers through web page scrapers custom-built for every aggregator. Unless a better parsable format is available (e.g., JSON), we retrieve raw HTML through simple HTTP requests and then extract offer data from relevant elements. Most aggregators present a paginated overview listing all offers, which we retrieve and store on a daily basis. Where necessary, we implement additional logic that addresses frequently occurring limits on the number of offers that the pagination supports, e.g., by collecting offers network by network through additional filters. Afterwards, for each offer, we request more detailed data by visiting and scraping the individual offer page that is linked from the overview. We retrieve and store 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 optimize scraping wherever possible, e.g., using internal APIs and maximizing the number of offers per page, which also reduces strain on the aggregators.

3.3 Normalization

The records in our offer data contain duplicates of four kinds: identical offers of the same network on one aggregator that we retrieve on multiple days; 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. Not all aggregators merge offers across all these dimensions, so we must first discover and merge duplicate offers ourselves. Since not all aggregators provide the same metadata, merging offers also improves data completeness.

We opt to merge offers if they share their (normalized) name. Our intuition is that the offer name lists the product or service name, and that this name is relatively unique. Given the variation in available metadata across aggregators, other fields, including an aggregator's internal identifier, are unreliable for ensuring that similar offers will be merged. Our name normalization algorithm extracts and removes common keywords that are unrelated to the product or service but rather describe the offer's terms (e.g., 'CPL'). This ensures that we merge offers for the same product/service across countries or conversion requirements. In addition, we strip punctuation and convert to lowercase. We use 449,891 *offers* to refer to distinctly named products and services, which derive from *network offers*, which are distinct pairs of product/service name and affiliate network.

Aggregators may use different labels for the same offer verticals. We relabel them according to our taxonomy of verticals (section 4), manually mapping original vertical labels used at least 1,000 times. Based on this labeling, we map 412,519 offers (91.7%) to one of our nine verticals. 11.3% of offers mapped to multiple verticals, for which we applied majority voting to select the most common vertical, after stripping more generic verticals (i.e., e-commerce and software). We opt to use only the categorization provided by the aggregators instead of classifying offers through, e.g., a machine learning model, as the metadata does not necessarily allow a reliable automated classification. The offer's name, for instance, may reference a unique product or service name without a clear indication of its vertical. A description may not always be available, or contain generic keywords related to an offer's terms that are not indicative of the vertical.

3.4 Filtering reputable offers

The affiliate marketing model is not inherently malicious or deceptive. Many reputable businesses use affiliate marketing, with larger brands often operating their own program, [4] and smaller brands joining reputable affiliate networks. [5] Even the aggregators of deceptive offers list some reputable offers. As we aim to study deceptive offers, we identify and remove as many reputable offers as possible from our data set. For this purpose, we collect offers from 4 additional aggregators, also discovered through an online search, for which we confirm through manual inspection of an offer sample that they list reputable offers from mainstream brands. [6] We manually identify and remove 97 affiliate networks that appear on both types of aggregators and that promote reputable products and services. We further remove offers if their landing page URL has the same domain as an affiliate program listed on an aggregator of reputable offers under a 'bare' root domain (i.e., with an empty path). These measures resulted in the removal of 298,947 legitimate offers from our data set. Our approach errs on the side of retaining reputable offers rather than removing deceptive offers, to ensure we capture deception in the ecosystem as broadly as possible.

[4] For example Amazon (`https://affiliate-program.amazon.com/`) or eBay (`https://partnernetwork.ebay.com/`).

[5] For example Awin `https://www.awin.com/` or CJ (`https://www.cj.com/`.

[6] Note that 'reputable' or 'mainstream' does not imply that the brand is large.

4 Analysis of offer verticals & deceptive tactics

In the first part of our analysis, we study the different scams that are all prevalent in the deceptive affiliate marketing ecosystem. While prior work has studied individual scam categories, it did not situate them within the broader ecosystem, thus to date we lack a holistic overview of the ecosystem. To establish this holistic overview and show how the ecosystem underlies a variety of deceptive practices on the web, we build a common taxonomy covering the nine main verticals (categories) of deceptive products and services provided by merchants and promoted by affiliates. We develop this taxonomy starting from guides from major affiliate marketing ecosystem players [1–3, 22, 44, 67, 84, 103] as well as the categories listed in the offer aggregators, and further refine it through a manual review of a sample of 750 offers by four authors as well as an analysis of common phrases present in the offer names. For each of the nine verticals in our taxonomy, we qualitatively describe the primary tactics used to mislead consumers, also referencing prior observations of these tactics, to understand in more depth how they are shared or differ across verticals.

1. **Dating/adult** Dating services promoted through affiliate marketing usually operate on a subscription basis, often through short-term trials and recurring billing that may be hard to cancel [13, 61]. Sites further deceive users by listing fake profiles that sites themselves admit (in their terms) to be only for "entertainment" purposes [13]. These profiles are meant to elicit chat messages, for which customers have to pay [65]. These sites also make unsubstantiated claims about the speed and ease of finding a match on the site, supported by fake testimonials. Separately, some sites show adult content, which may be undesirable or inappropriate for users to see, especially if it is shown in a non-adult context or to minors.
2. **Entertainment** Entertainment offers largely cover two types: content distribution platforms, such as for videos or music (also visible in offer names), and mobile content portals, with offer names referencing mobile carriers (e.g., Claro, Movistar, Vodafone). As we did not sign up for these sites, we are unaware whether they (legally) host any (interesting) content. Both are prone to hidden subscriptions, and also collect personal data such as email addresses or phone numbers, potentially for abuse through resale.
3. **Finance** Trading platforms for cryptocurrencies (with Bitcoin being the most common phrase in offer names), binary options, foreign exchange, or commodities (e.g., oil) make exaggerated 'get-rich-quick' claims that users can earn vast amounts of money in a short time, supported by fake user testimonials and fake celebrity endorsements. However, users may instead lose their investment or are unable to withdraw funds, after having been pushed to make large deposits [14, 106]. These trading platforms have been linked to fraud in the past and are therefore often heavily regulated or even prohibited [110]. Some offers concern 'business opportunities' ('BizOpp') for building one's own business, often in the form of self-help guides. Finally, some sites collect personal data for lead generation on credit, skewed towards payday loans with excessive interest rates;

insurance, often with low coverage [80]; or banks, e.g., account sign-ups. Some offers claimed to help obtain Social Security disability benefits, possibly targeting vulnerable populations.

4. **Gambling** Online gambling sites provide e.g., casino games (the most common offer name phrase) or sports betting. These sites claim large potential earnings for users, and offer large "welcome bonuses" that may actually be difficult to earn. The sites often operate with valid licenses from permissive countries, e.g., Curaçao and Malta [75], but may be illegal to visit or advertise in some countries [116].
5. **Games** Games or game guides on mobile (with carrier references in offer names) often entail hidden recurring subscriptions, to an extent that Chrome has introduced warnings for them [92]. For browser and other mobile games, developers may simply require that users install their game, and monetize users separately, e.g., using microtransactions.
6. **Health/beauty** Health/beauty offers cover products for, a.o., weight loss (e.g., keto diet), 'nutraceuticals' (foods and supplements with claimed health benefits), CBD, male enhancement, or skincare, as are visible in offer names. Affiliates were also found to promote price-gouged masks [95], and go so far as promoting quack cures for diseases such as diabetes. These physical goods may actually be shipped to consumers [72], but they may unknowingly start a subscription for regular deliveries [11, 41]. Products may have no actual utility, in particular compared to the claims made, and users consider these supplements particularly "scammy" [121]. Likely fake testimonials from doctors, regular consumers, and celebrities [11, 23, 28, 114] further acclaim the products. The low value of these products may also be apparent from a need to add disclaimers such as "not FDA approved".
7. **Software** Promoted software applications are often 'utilities': software that purports to protect or improve devices, visible in offer names through common phrases such as antivirus software, VPNs, media but also (fake) Flash players, WhatsApp add-ons, ad-blocking browser extensions, or "phone cleanup" tools, both on desktop and mobile [102]. On the one hand, these may be mainstream apps whose affiliate programs are 'arbitraged', i.e., deceptive affiliate networks 'repackage' affiliate programs from reputable merchants [27], a.o., to make them more accessible to affiliates [6]. On the other hand, these may be low-quality apps that are monetized directly through recurring mobile subscriptions or intrusive advertisements. Users find ads for these software downloads deceptive [121]. Common offer name keywords show these offers are often targeted at specific platforms, e.g., Chrome or iPhones. Outright malware is also sometimes spread, such as the Shlayer trojan through fake Flash Player software [56]. Next to utilities, offers may concern purchased incentivized app installs, where users are rewarded for installing the app [27, 115]. As the practice artificially boosts app installs, the practice is discouraged or prohibited by major app stores [27].
8. **Sweepstakes** Free products (with iPhone and Samsung Galaxy being common in offer names) or vouchers/gift cards (Amazon being common) are given away in exchange for a user action [18, 20]. A user may need to leave contact details, to be sold to third parties ('lead generation'). A user may alternatively

be required to submit their credit card details, (unknowingly) starting subscriptions [28], for e.g., low-quality streaming sites, recipe catalogues, discount shops, or website builders. These giveaways often suggest affiliation with major brands, who have warned users not to participate [46, 57]. Users may also be prematurely congratulated, addressed as a 'winner' or told they can claim the prize even though they might not win the contest. The giveaways may actually happen, although a product may be substituted for a gift card of the same value, and their terms are often very limited with only a handful items being given away per year [28]. Instead of contests, products/gift cards may also be given as a reward for completing other offers.

9. E-commerce This vertical includes a variety of deceptive products and services for sale that do not fit the other categories. These can be physical goods, ranging from electronics ("electricity saving boxes") to lucky charms (money amulets), with exaggerated claims about their utility and likely fake testimonials from customers. Arbitrage of existing affiliate programs for online shops, travel sites, etc. also exists, as does lead generation for, e.g., home improvement (solar panel incentive programs).

In general, deception lies in a mismatch between price and value, with products and services being of dubious quality or egregiously priced. In addition, merchants deceive consumers by persuading them with exaggerated claims, hiding important terms and conditions in small print, or hiding the implications of the offers, such as recurring payments or the sale of personal data to third parties. Abuse of "dark patterns" to pressure users [76], such as timers and indicators that the offer will expire soon or fake user activity notifications ("X has just bought Y"), as well as fake endorsements by experts, celebrities, purported previous buyers, or unaffiliated brands to gain trust are all common. In the US, the Federal Trade Commission has alleged that these practices violate various consumer protection laws, such as the FTC Act [30, 32, 40]. In the European Union, the unfair commercial practices directive is similarly interpreted to ban such practices [21]. Outside actors also recognize the deception within the ecosystem and the harms that stem from it. Many verticals within the deceptive affiliate marketing ecosystem concern industries that are considered high-risk by financial institutions, with "more operational, regulatory, and reputational risk" [26] including higher levels of fraud [48], forcing merchants into transacting with a small set of banks who are willing to take the risk [77]. The ecosystem appears to be diverse and ever-evolving in the scams that are perpetrated, warranting the broad and continuous monitoring provided by our data collection.

5 Quantitative analysis of offers

In the second part of our analysis, we quantitatively characterize the dynamics and differentiation within the deceptive affiliate marketing ecosystem. Using the metadata unique to our data set, we seek to understand how the monetary incentives for affiliates depend on the type of scam and how it is monetized. This provides context for the strategies that ecosystem players might use and the risks they might be willing to take to propagate scams. In particular, we want to leverage our more comprehensive coverage to see whether users will see different abuse, e.g., depending on where they live, as this insight is necessary to develop comprehensive interventions that globally protect users. Finally, we also want to observe whether there is concentration in the ecosystem, as identifying the large-scale ecosystem players that enable a large number and variety of scams is useful to prioritize interventions. Note that our view on the ecosystem only allows us to analyze how many offers merchants publish and make available to affiliates. We cannot know which offers affiliates promote more or less, more successfully convert, and therefore cause the most harm to consumers.

5.1 Summary statistics

For our analysis, we use data collected between March 24, 2020 and September 1, 2024, providing us longitudinal coverage[7]. We collect 449,891 *offers*, i.e., distinctly named products and services. These derive from 625,106 *network offers* across 1,165 affiliate networks, i.e., distinct pairs of product/service name and affiliate network. Between aggregators, 70.8% of networks and 62.1% of offers are only listed on 1 aggregator, showing that collecting data from many aggregators increases our coverage. We specifically find 150 networks that are only listed on the Russian-language aggregators, supporting our decision to include them. Orthogonally, combining metadata across aggregators improves data completeness: for 36.3% of offers, one aggregator had metadata that was missing at another aggregator. Through the aggregators, we capture data on a broad set of affiliate networks, thus we observe a larger share of the ecosystem than if we were to focus on specific affiliate networks.

5.2 Payouts and conversion criteria

The merchant who makes an offer available will pay the affiliate an agreed payout if the affiliate achieve that, after advertising the offer to a consumer, that consumer takes the required action to ‘convert’ the offer, i.e., fulfill the offer’s condition. The payouts to affiliates reflect how much the ecosystem values offers across certain dimensions, such as the vertical or conversion criteria, which we analyze in this section. Throughout this section, we compute a median payout across all observed payouts for offers, i.e., over countries, affiliate networks, time, and aggregators, and convert all payouts to U.S. dollars using exchange rates on September 1, 2024. 444,745 (98.9%) offers have payout metadata.

We first plot the cumulative distribution of payouts per normalized vertical (subsection 3.3) in Figure 2, giving us insight into how much merchants are willing to pay affiliates depending on the type of product or service and therefore the type

[7] Gaps in coverage do occur: aggregators became unavailable while others were added later on; aggregators blocked our scrapers or changed their site layout which broke our scrapers, or our scraping infrastructure was down.

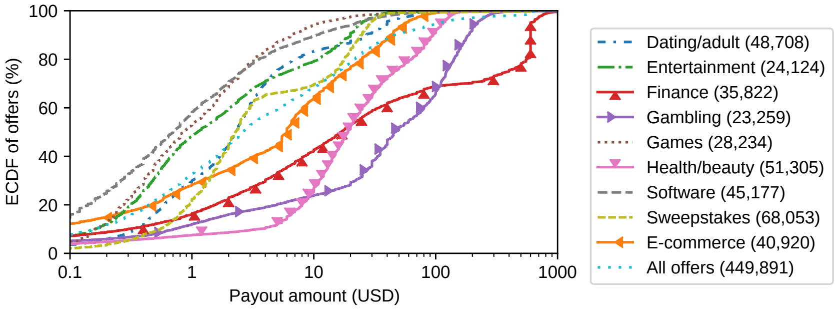

Fig. 2: The cumulative distribution of median payouts for offers per vertical shows that scam types significantly differ in their valuation within the ecosystem. (Legend lists number of offers.)

of scam. *Finance*, *Gambling*, and *Health/beauty* command the highest payouts, with over half of offers having a median payout of more than $10. Conversely, over 60% of offers in the remaining verticals (other than *E-commerce*) have a payout of less than $5. Across all offers, the median payout is $2.60. In order to better understand how much different types of deceptive offers and the associated conversion types are worth to affiliates, we extract common phrases from offer titles and compute the median payout for the respective offers.

In general, high payouts reflect that a consumer needs to make a (higher) payment to the merchant. We see payouts of over $100 for 31.4% of *Finance* offers, specifically trading and investment platforms, with 'Bitcoin' and 'crypto' offers having a median payout of $582.5 and $550.0, respectively. A successful conversion requires affiliates to convince consumers to make a minimum deposit, usually $250. (Interestingly, the payouts to affiliates sometimes exceed this deposit.) This suggests that the merchants of these platforms believe that users will invest large amounts of money over time, and are therefore willing to incentivize affiliates to promote the platforms through the promise of large rewards [50]. On a similar note, *Gambling* offers command higher payouts when they require players make a deposit to a gambling website, with 'casino' offers having a median payout of $87.5. For *Health/beauty* and *E-commerce* offers, the delivery of physical goods inherently solicits larger payments and thus higher affiliate payouts. For example, the median payout for 'keto' diet offers is $80.0, and is $41.7 for 'gadgets.' However, physical goods have a higher cost to produce and ship than for virtual trading or gambling platforms, so payouts are not as high as they are for the latter platforms. For *Dating/adult* and *Sweepstakes* offers, around 15% and 35% of offers respectively have around 10 times higher payouts than the rest of the vertical, which usually coincides with the offer converting on a consumer paying for a (trial/premium) subscription. *Entertainment*, *Games*, and *Software* offers have the lowest payouts, even when they require consumers to make a payment. *Games* 'subscription' offers, for example, have a median payout of $4.2. Presumably, this is due to the typically lower sales prices of the respective digital goods.

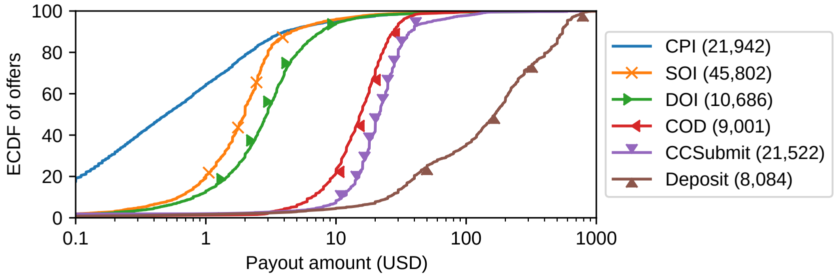

Fig. 3: The cumulative distribution of median payouts for offers per conversion type shows that higher payments or sharing more personal data by consumers correlates with the valuation of an offer. (Legend lists number of offers.)

When offers convert on a consumer action other than a payment, affiliate payouts tend to be lower. In these cases, merchants need to monetize their offers indirectly, such as by leveraging the consumer's personal data. Lead generation offers in the *Finance* vertical, for instance, yield relatively high median payouts (e.g., $11.0 per consumer submitting their personal data in 'loans' offers) compared to lead generation offers in other verticals, but they are an order of magnitude smaller than for investment platform offers requiring a deposit. Other common types of offer conversion are registration on a website, which exists in all verticals but is most prevalent for *Dating/adult*, *Entertainment*, and *Gambling*, or the installation of an app (*Games* and *Software*).

The distribution of payouts per conversion type [22] across all verticals (Figure 3) confirms this relationship. The lowest payouts are a median of $0.5 for app installs (CPI, cost per install). This is followed by two forms of lead generation. Median payouts are slightly higher at $2.95 when a consumer confirms the validity of their contact details (DOI, double opt-in – clicking a confirmation link in an email, for instance) as opposed to $2.0 for only providing the data (SOI, single opt-in). Once consumers make payments, payouts increase around tenfold compared to the value of personal data for leads. The payout depends on the payment method: at $15.35, the median is slightly lower for cash-on-delivery (COD – payment upon receipt of the item) than the $21.2 for 'credit card submit' (CCSubmit). This may cover the risk of non-payment on delivery for a product that is already shipped, compared to a credit card that can be charged ahead of shipment, but could also reflect the popularity of COD in emerging markets [47], where offer payouts are lower overall (subsection 5.3). Finally, at $165.7, median payouts are the highest for offers where the consumer makes a deposit, which correlates with offers where no goods need to be shipped.

Different conversion types and payouts likely drive affiliates to pursue different strategies to attract audiences that might result in a successful conversion. Knowledge of what an affiliate needs to achieve for a conversion, and how much they can afford to spend on promotional efforts to acquire conversions, can help researchers devise more targeted measurements or interventions.

5.3 Targeted countries

Another interesting aspect of affiliate offers is how merchants target their deceptive products or services at specific countries. The ecosystem usually separates countries into three tiers, depending on purchasing power, language skills, and regulatory frameworks. We describe these tiers based on the definitions from major ecosystem players [51, 70, 79, 104, 111–113]. *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. This means that many traffic sources and offers with high payouts are available to affiliates. However, competition between affiliates might be higher and markets tend to be more regulated, restricting which (deceptive) products, services, and promotional tactics are allowed. *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. The tier is seen as a starting point for affiliates, with fewer legal restrictions and cheaper traffic, but also lower payouts and conversions. *Tier-3* countries comprise most African countries, and the remaining countries in South America and Asia. Consumers in this tier have the lowest purchasing power, and are less attractive to affiliates as payouts are low. Local languages and customs may make consumers less accessible, but this also means there is less competition for cheaper traffic in a less regulated market.

Out of the 406,705 offers with associated country data (90.4% of all offers), 77,069 (18.9%) target a worldwide audience.[8] Conversely, 229,272 (56.4%) of offers with country data are targeted at only one country. In absolute numbers, the United States has the most offers, at 87,550 offers (Figure 4). Germany comes a distant second at 35,587 offers. Overall, more offers are targeted to countries 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.

In terms of verticals, a country tends to have a higher or lower share of its targeted offers in higher or lower-paying verticals in accordance with the country's tier. Due to space constraints, we report on the most significant findings from our analysis of each vertical's share of a country's total targeted offers. *Dating/adult* offers have a relatively high share of offers in most countries, no matter the tier; it is also the most common vertical overall. *Finance* offers tend to be more frequent in tier-2 locations, such as Latin America (e.g., 32.4% in Bolivia), the Middle East (e.g., 27.2% in Oman), and the Baltics (e.g., 24.3% in Latvia). *Gambling* offers are more frequent in CIS countries, led by Turkmenistan at 45.5%. *Health/beauty* offers are more frequent in Eastern Europe, e.g., 65.8% in Bosnia and Herzegovina and 50.6% in Montenegro. while other *E-commerce* product offers are particularly frequent in Russia (28.3%) and CIS countries (e.g.,

[8] They list (an equivalent of) 'Worldwide' as their country or in their name, or cover over 100 countries. For reference, the United Nations has 193 member states, and 249 territories have been assigned an ISO 3166-1 alpha-2 code.

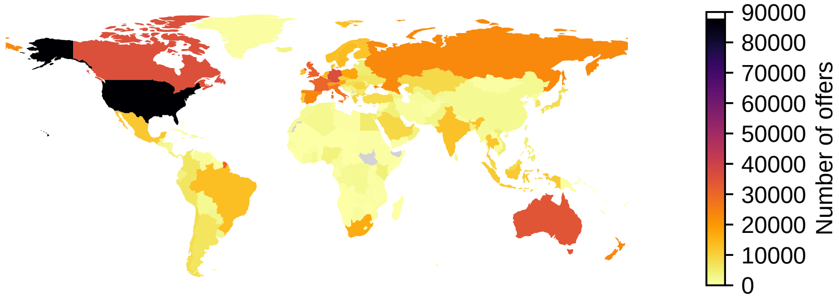

Fig. 4: There is an uneven distribution in the number of offers per country where merchants target that country. This distribution correlates with the 'tiers' in which the ecosystem separates countries, which itself depends on purchasing power, language skills, and regulatory frameworks.

24.0% in Kazakhstan). *Sweepstakes* are common mostly in tier 1: Europe (e.g., 32.1% in France, 27.0% in Germany), the United States (22.0%), and Australia (22.0%), but also Indonesia (23.8%) and Singapore (23.2%). Digital goods in the *Entertainment*, *Games*, and *Software* verticals are most common in tier-2 and 3 countries, i.e., Latin America, Africa, the Middle East, and parts of Asia. These offers are mostly mobile-based, and may therefore reflect the wider availability of mobile payments compared to other payment methods in these countries. The above analysis shows that specialization also occurs on the country axis. Differing offer market shares of verticals in countries primarily depends on the willingness and ability of merchants to do business in a country, which in turn relates to the (perceived) profitability, as gauged by intrinsic consumer interest but also factors such as income, but may also depend on external factors such as the availability of payment methods, ease of fulfillment, legal restrictions, etc.

5.4 Affiliate networks

The offers in our data set originate from 1,165 affiliate networks. We see a concentration of offers with large networks: 50% of all offers stem from only 47 networks. Between networks, 17.6% of offers with the same name are listed on more than 1 network, suggesting that certain merchants seek to increase their reach of potential affiliates through multiple networks. Our analysis confirms the tendency for specialization in this cybercriminal ecosystem [101, 109], as affiliate networks tend to specialize in one or a handful of verticals. 50.4% of networks have only one vertical with at least 5% of their offers. Similar to the distribution for offers, *Dating/adult* is the most common vertical, with 403 networks having more than 5% of their offers in this vertical (Table 4). Figure 6 shows the top 20 affiliate networks according to the number of offers, with the three largest affiliate networks in our data, Mobusi, Golden Goose and Algo Affiliates, listing over 27,000 products and services each. Of note is that some of these networks have previously been linked to scams: ClickDealer was linked to Bitcoin investment

Table 3: The top 5 affiliate networks specialize in the countries they target, as shown by their top 5 countries in terms of percentage of offers. (Offers can target multiple countries, so percentages per network do not sum to 100%.)

Network	Country 1		Country 2		Country 3		Country 4		Country 5	
mobusi	ZA	9.05%	TH	7.77%	IT	6.61%	BR	6.18%	IN	5.84%
goldengoose	SA	7.27%	CH	4.73%	ZA	4.49%	RO	4.32%	TR	4.27%
algoaffiliates	US	96.24%	FR	93.79%	DE	93.75%	CH	93.74%	CA	93.74%
mylead	PL	24.68%	DE	22.16%	US	17.27%	AU	17.10%	CA	14.94%
trafficlight	RU	50.21%	KZ	15.20%	BY	8.87%	UA	8.34%	IT	8.13%

scams [14, 106] Adscend Media was sued for spreading spam [108], and CPALead was linked to survey scams [18]. The latter two accusations date from 2012 and 2013, respectively, highlighting the longevity of these affiliate networks and their deceptive tactics.

Among the top 3, we again see a different focus in terms of targeted verticals. Mobusi and Golden Goose target mobile devices with *Dating/adult*, *Entertainment*, *Games*, and *Software* offers, while Algo Affiliates has a strong focus on *Finance* offers. According to their websites, these three networks have over 20,000, 10,000 and 100,000 affiliates, respectively, reinforcing the vast scale of the deceptive ecosystem. In the same spirit, specialization happens at a country level: among the 5 affiliate networks with the most offers, offers target a very varied set of countries (Table 3). For example, the countries that Mobusi targets most are mostly tier-2 countries, while Golden Goose, Algo Affiliates and MyLead target mostly tier-1 countries, and Traffic Light targets CIS countries.

Finally, the longitudinal aspect of our data set gives us a view on affiliate networks entering or leaving the market. Omitting the start and end of our measurement period (to account for left/right censoring), we saw 639 networks appear and 544 networks disappear over our 4 years measuring offer aggregators (Figure 5), which may reflect networks being created or disbanded as well as networks deciding to start or stop advertising their offers on aggregators. This shows the dynamism of the ecosystem, and highlights the need to monitor it longitudinally to discover and target new deceptive actors.

5.5 Longitudinal popularity

Through the longitudinal view that our data set provides on the ecosystem, we can observe the emergence of new topics in deceptive offers. After the initialization of our data set in March 2020, the count of newly observed offers has been fairly stable over time (Figure 7), at an average rate of around 4,000 offers per week. Compared to this global trend, we see that the release of new and popular products coincides with merchants creating new offers that mention them. As one case study (Figure 7), we see that *Sweepstakes* offers where different iPhone models are promised as a prize start appearing on or slightly before the products' release dates. This suggests that ecosystem players are capitalizing on the novelty

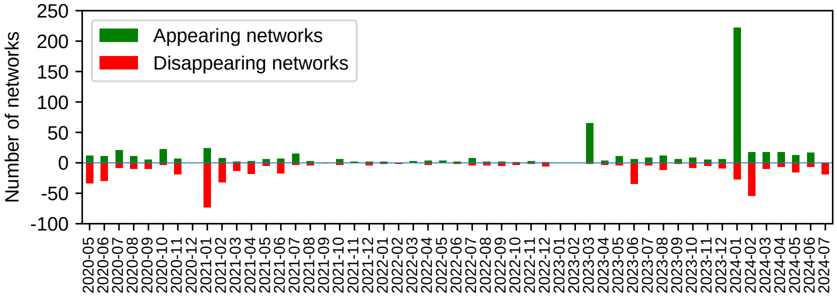

Fig. 5: Affiliate networks (dis)appear over our 4-year measurement of aggregators.

Table 4: Networks differ in the verticals they specialize in, as shown by the number of networks for whom the vertical covers at least a given percentage (1, 5, or 10%) of their offers.

Vertical	# networks		
	1%	5%	10%
Dating/adult	519	403	321
Finance	520	376	291
E-commerce	499	360	278
Sweepstakes	361	280	218
Health/beauty	432	271	206

Vertical	# networks		
	1%	5%	10%
Gambling	327	205	170
Software	320	177	122
Other	172	105	87
Games	208	97	65
Entertainment	217	74	46

of these products to make their offers more attractive, and that aggregator data can be useful for discovering new offers referencing these products early on. On the level of verticals, we also measure the longitudinal trends within the relative share of new offers for each vertical (Figure 8). Among the most significant trends, *Dating/adult* offers were more common in 2020 and 2021, perhaps owing to the COVID-19 pandemic and lockdowns making virtual dating sites more attractive. *Finance* offers spiked mid-2022, before receding to a much lower level, although these offers seem to resurge in 2024. *Software* offers see an increase over time, achieving a peak share of offers in 2024 that was four times higher than the lowest shares in 2020 or 2021. Ecosystem players are therefore also adapting to long-term trends, causing certain scams to fall out of favor or in contrast are published by merchants more often, which reinforces the need to monitor the ecosystem and its evolution over time.

5.6 Summary

The deceptive affiliate marketing ecosystem enables a broad variety of online scams – even though they might seem unconnected at first – with both more well-known scams such as unwanted software, and previously less studied scams such as finance or entertainment scams. Through the payout metadata unique to our data set, we see that some offers yield very high rewards, up to hundreds of

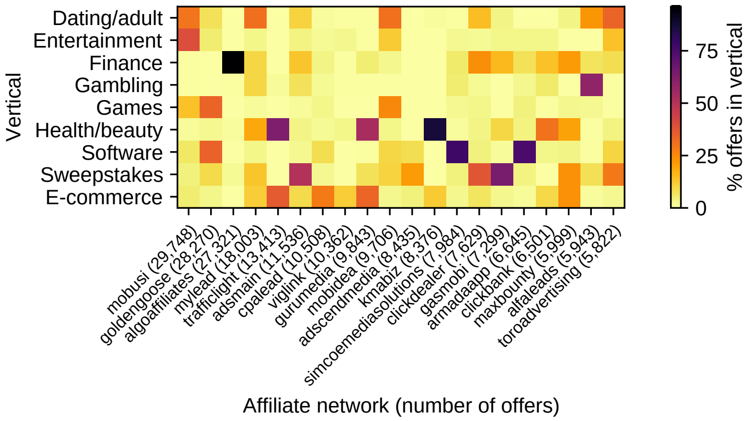

Fig. 6: For the 20 networks with the most offers, the proportion of offers per vertical shows that networks specialize in the scam types they manage.

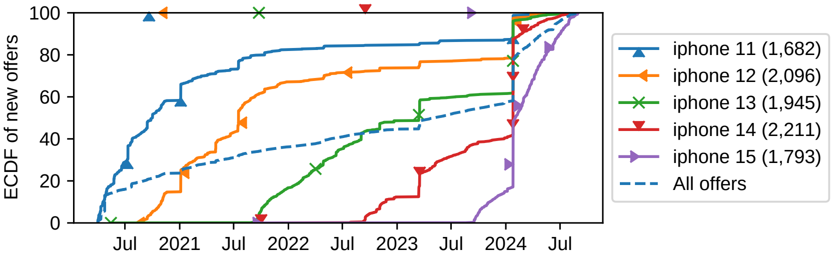

Fig. 7: The longitudinal cumulative distribution of newly observed offers overall shows that the ecosystem grows steadily. The distribution of offers containing keywords for iPhone product launches is related to the release dates (shown by markers along the top), which shows that merchants capitalize on new trends.

dollars for investment scams. Payouts correlate with what the consumer provides in terms of money or personal data. This can inform which affiliates are attracted to the offers and the tactics they are willing to use to achieve a certain revenue. The deceptive affiliate marketing ecosystem operates on a global scale, but the selection of scams is adapted to the country, reflecting, a.o., incomes. In general, offers can be targeted at very specific audiences. Through the longitudinal view that our data set provides, we also discover that affiliate marketers capitalize on new trends. Our data set can be used to achieve broad coverage of the ecosystem to develop more comprehensive interventions that protect consumers across countries, verticals, and time. In addition, we can use it to identify the large affiliate networks, against which interventions may be the most effective.

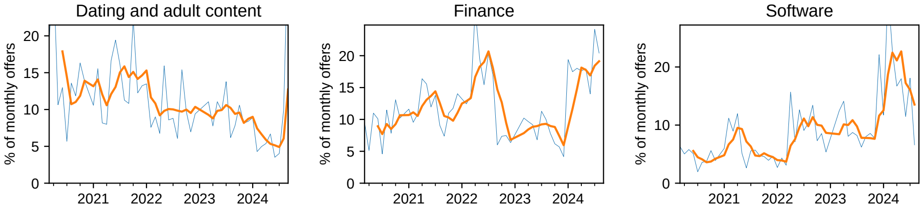

Fig. 8: The longitudinal share of new offers (i.e., observed for the first time in our data set) per month for three selected categories shows that the publication of offers from different scam types is affected by long-term trends. The thick orange line is a four-month moving average, the thin blue line is the monthly share.

6 Discussion

Specialization. The deceptive affiliate marketing ecosystem embeds specialization across multiple dimensions, a trend that has been observed in other cybercriminal ecosystems [54, 87, 88, 101, 109] and that we observe as well. Affiliate marketing is itself a form of specialization, where merchants no longer need to manage advertising themselves and can focus on fulfillment, and instead contract promotion out to individual affiliates. These affiliates can focus on marketing the offers to specific audiences that they believe to be the most susceptible to the deceptive scams, through the most effective advertising channels. We find that our third main actor, the affiliate network, extends this division of labor, managing affiliate recruitment, offer tracking, and commission transfers. These affiliate networks also introduce centralization through their 'marketplace' for online scams, where affiliates can easily discover a large offer inventory and merchants can easily access a large pool of affiliates, increasing the ease with which ecosystem players can encounter each other and trade services.

With the affiliate network acting as a central party to all transactions, affiliates become disconnected from the concrete merchants whose products they promote, and see offers as a commodity; the same holds for merchants towards the affiliates promoting their products. This ties in with risk distribution: merchants only need to pay affiliates when the product or service actually is sold, while affiliates can be agile in switching merchants if the product or service does not perform well. Similarly, these distributed responsibilities could be seen as liability distribution: affiliates, affiliate networks, and merchants could claim they are unaware of the deceptive practices of the other parties, or at least not responsible for them. Nevertheless, we see that some ecosystem players openly acknowledge, discuss, or even encourage deceptive tactics [4, 12, 28, 66, 73], suggesting complicity in the scams. This relates to our broader observation that the deceptive affiliate marketing ecosystem is "hiding in plain sight" [88]: the offer aggregators from which we obtain our data operate in the open, as do most of the actors and supporting services, who seem to consider their individual participation in the ecosystem above board. In this light, we also consider that the listing of reputable products and services on deceptive offer aggregators (Section 3.4) can still attract abuse. Reputable and deceptive offers are intertwined on these aggregators, so affiliates browsing the aggregators may not (be able to) distinguish between these two kinds of offers. Reputable brands may then be harmed when affiliates use abusive tactics to promote their products [27].

Throughout our analysis in Sections 4 and 5, we also observe specialization within the set of offers that merchants make available to affiliates. We taxonomized offers into nine verticals for which we see greatly varying payouts, which correlate with the varying valuation for the countries where the offer is available. This means that merchants and affiliates can select their level of expertise but also comfort and exposure to intervention with regards to the scams they want to engage in. Indeed, new affiliates are recommended to promote lower payout offers [15, 22], i.e., those from specific verticals and (tiers of) countries, in part as they are said to more easily convert, have less competition, and are less regulated. These new affiliates (and merchants) may start off with 'small-scale' scams, which nevertheless deceive users into losing money or personal data. Then, through the experience they gain, affiliates and merchants may eventually "graduate" to more high-profile, high-payout scams, such as harmful scams for nutraceuticals, gambling sites, or financial products, increasing the harm to online consumers. All these elements of specialization combined – distribution of tasks, risk and liability; networks centralizing offers; and diversification of offers – therefore increase the ease of joining the deceptive affiliate marketing ecosystem, and contribute to its scale and growth, as we observed in our longitudinal analysis (Section 5.5). Having our ground-truth vantage point helps in understanding this specialization and maintaining wide coverage of the ecosystem despite it, which is crucial for broadly effective interventions.

Intervention. Ultimately, consumers should be prevented from viewing deceptive offers, while merchants, affiliate networks, and affiliates should face consequences for publishing deceptive offers. The division of responsibilities between the deceptive affiliate marketing ecosystem players can both help and hinder such interventions. It provides more venues for intervention: as long as one of the links in the chain from traffic source to landing page is broken, a consumer is no longer able to complete the offer. Conversely, any one intervention will only address a small part of the ecosystem, as all other players can remain active. However, any intervention introduces an additional burden for ecosystem players, and may prevent or deter them from participating in the ecosystem.

On a technical level, interventions can start on the consumer's side: through client-side blocks of deceptive offers, e.g., through (DNS-based) blocklists, browser interstitials or extensions blocking affiliate marketing-related URLs or detecting deceptive affiliate marketing content. These interventions have relatively low overhead and can therefore be deployed quickly, but only protect the users who install them. The early look provided by our data set, tracking new offers as they emerge on aggregators, can lead to faster ingestion of offers and landing pages into blocklists, potentially even preventing users from ever seeing the offer. Next, the traffic source could detect and take down ads for or links to known deceptive offers, preventing all users of that traffic source from being redirected to those offers. Finally, along the redirect chain from the traffic source to the landing page, the supporting infrastructure, such as servers or domain names, can be taken down [107]. Such a takedown may be requested by law enforcement or by affected brands [10, 55, 69] and implemented by registries, registrars and/or hosting providers. This process protects any Internet user from seeing the deceptive offer, as (in the end) the landing page becomes unreachable. However, takedowns involve more stakeholders and have more overhead [55], and may therefore be slower to take effect. To maximize their impact and effectiveness, they should target the "bottleneck" infrastructure resources in the redirect chain [71, 100], such as affiliate network tracking domains. These can be identified by combining our data set with end-to-end measurements starting from the traffic source.

Beyond technical interventions, while deceptive affiliate marketing operates in a 'grey zone,' consumer protection regulators can take action against unfair or misleading commercial practices [25], including novel techniques that users may perceive as deceptive [121]. For example, in the United States, the Federal Trade Commission has the authority to regulate on practices such as deceptive advertising [31, 42], and has previously targeted deceptive affiliate marketing ecosystem players [29–31, 33–41]. Interventions against financial services may also be effective in disincentivizing ecosystem players by disrupting their revenue streams [60, 77, 101]. Our data set can expand the discovery by regulators of new types of scams or deceptive techniques, and help in identifying the major ecosystem players and attributing infrastructure to them, in order to further enable comprehensive interventions.

Limitations. While we saw that the 23 deceptive offer aggregators covered a wide variety of verticals and affiliate networks, they may not cover all types of online abuse that uses the affiliate marketing model. Notably missing from our data set are offers related to two trades of physical goods where researchers have previously observed the affiliate marketing model, pharmaceuticals [58, 78] and counterfeit goods [59], as well as online crimes such as bulletproof hosting providers [86] and ransomware [17, 86]. The more obvious illegality of these crimes may make it too risky to publish these offers openly on aggregators. However, the prior studies suggest that the relation between affiliates and merchants is more direct, without affiliate networks as intermediaries. The merchants may therefore rely more on word-of-mouth marketing, as to inherently select for affiliates who are more deeply involved in the ecosystem and may more likely convert offers, and therefore would not need aggregators to list their offers. In addition, affiliate networks that do publish offers on aggregators might have private offers that are only available to those affiliates who are registered with the network, or with a proven track record [22]. Again, the deceptive offers that we discover and study may serve as a stepping stone towards these even more harmful ecosystems, which confirms the need to also target the deceptive affiliate marketing ecosystem.

While we cannot independently confirm the origin and reliability of the offer data that aggregators provide, we have good reason to believe that aggregators obtain it directly from affiliate networks. Metadata from some aggregators refers to the offer management platforms of affiliate networks as data sources, suggesting that they are directly integrated. Aggregators also advertise their platform for networks to list their offers (usually for a fee) [5, 45, 83, 122], suggesting they actively collaborate to publish offer data. In addition, we argue that networks and aggregators have an incentive to provide correct data to affiliates. Next to the commercial benefits of accurately representing the available offers, underground activities operate on a reputation system, where breaches of trust result in negative feedback on e.g., underground forums [49]. Similarly, we expect dishonest aggregators or networks to be called out. If inaccuracies in the data are present, we expect them to be due to parsing errors or data staleness rather than malicious intent. We therefore consider the aggregator data reliable for our purpose.

Since we focus on gaining insights into the dynamics of the deceptive affiliate marketing offers that merchants make available, we cannot use our data set on its own to measure the ecosystem from the sides of consumers and affiliates. Future work can leverage the global overview from our data set to support end-to-end measurements with broad coverage, therefore improving the validity of future insights into the ecosystem. Such comprehensive end-to-end data can be used to determine how often consumers see and complete deceptive affiliate marketing offers. With data from traffic sources, it becomes possible to study how the advertising material created by affiliates is deceptive [120, 121], as well as which traffic sources are more commonly (ab)used for deceptive advertising [105, 120]. End-to-end traffic can also allow to identify the most active affiliates, affiliate networks, and merchants. In addition, while we quantify the monetary value of a successful conversion to the merchant, we cannot estimate the cost to the affiliate for obtaining a successful conversion, in particular the differences across offers, verticals, and countries. This also means we cannot reliably calculate total earnings across the ecosystem. Finally, in-depth investigation of the corporate entities behind main ecosystem players [24, 119] can further support interventions.

Ethical considerations. Given the often malicious nature of the players in the ecosystem, we must carefully consider how we conduct 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 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 [74, 89, 96]. To the best of our knowledge, the scraped offer data does not contain personally identifiable information, and our research was not considered eligible for ethical review from our university, due to it not being on human subjects. We will share both the data set used for this study and the data that is still continuously being collected 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 more extensive vetting including 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.

7 Related work

Prior work discussed how certain types of cybercrime use the affiliate marketing model. Samosseiko [91] first outlined the role of affiliate networks in spam-advertised pharmacies and counterfeit software, focusing on Russian 'partnerka' networks. Kanich et al. [58] and McCoy et al. [78] studied the purchases and revenues on major pharmaceutical and counterfeit software affiliate networks, using leaked ground truth of the networks. Levchenko et al. [72] 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. [101] describe how the affiliate marketing model is central to many organized cybercrime operations. For example, certain bulletproof hosting providers [86] or ransomware-as-a-service providers [17, 118] operate through affiliate programs.

Further work identified affiliate marketing in detailed studies of specific malicious ecosystems. Caballero et al. [16] analyzed the affiliate structure behind 'pay-per-install' malware. Kotzias et al. [64] and Thomas et al. [102] analyzed popular 'pay-per-install' affiliate programs. Stone-Gross et al. [97] studied the major actors and economics of fake antivirus software. Karami et al. [59] analyzed 'Tower of Power,' an affiliate program for herbal supplements and replica luxury goods. Clark and McCoy [18] analyzed the affiliate networks behind survey scams distributed through Facebook ads. White [114] describes the identification and subsequent takedown of one affiliate marketing campaign abusing celebrity endorsements to advertise nutraceuticals. Vadrevu and Perdisci [105], Koide et al. [63], and Yang et al. [117] describe detection and blocking techniques for "social engineering ads" which cover a subset of the scams that we observe, without connecting them to the deceptive affiliate marketing ecosystem. Compared to these prior works, which were done from the vantage point of a user and limited to a few verticals or traffic sources each, our vantage point and study provides a broader view on the ecosystem.

Finally, leveraging ground-truth data has proven valuable to studying cybercriminal ecosystems holistically, having been used for, a.o., online anonymous markets [68], stolen payment card marketplaces [9], and Bitcoin mixing services [82]. We are the first to use ground-truth offer metadata to study the deceptive affiliate marketing ecosystem, using it to discover deceptive products and services across all major verticals, countries, and potential traffic sources, which allows us to find new scams such as cryptocurrency investment platforms that were not yet studied in prior work, and to quantify and compare the monetary value that the ecosystem attaches to specific verticals and conversion types.

8 Conclusion

We provide an overview of the deceptive affiliate marketing ecosystem and the offers (products/services) that are promoted through it, using our longitudinal and ongoing data collection from 23 offer aggregators. We show how this ecosystem brings together a variety of scams observed on the web today. Ecosystem players deploy a wide range of tactics to mislead users, meant to suggest high quality, reputation, or demand, even though the products and services on offer tend to actually be of very low value. Depending on parameters such as the offer vertical and country, the products and services that are promoted, the affiliate networks that make them available, and the value that is assigned to them differs, highlighting the added value of the comprehensive coverage provided by our vantage points. This specialization makes the deceptive affiliate marketing ecosystem attractive and accessible to new players, resulting in its continued growth over time.

A crucial step towards combating the deceptive affiliate marketing scams is holistically understanding the inner workings and dynamics of this ecosystem. Our study provides a first global overview of this ecosystem, which becomes especially important given the specialization that is embedded in the cybercriminal deceptive affiliate marketing ecosystem, with its diverse actors and scam types. When combined with our data set, measurements on a variety of traffic sources can provide a comprehensive end-to-end view of the ecosystem, the interactions between the players, and the dynamics that steer them. Our continuous data collection pipeline provides crucial metadata that enables near real-time monitoring of newly emerging scams, which can lead to more extensive, effective and impactful interventions, to ultimately protect consumers from being scammed on the web. To support this, we will share our data with researchers and stakeholders to enrich other measurements and deploy powerful, effective and targeted defenses to prevent users from being exposed to deceptive affiliate marketing and losing their money or personal data.[9]

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. We thank the participants of the SecWeb 2021 workshop and IEEE S&P 2022 poster session for their feedback on an early version of this research. This research is partially funded by the Research Fund KU Leuven, the Cybersecurity Research Programme Flanders, and the National Science Foundation under grant 2039693. This research was supported by the Research Foundation - Flanders (FWO) through a PhD fellowship (11A3419N, 11A3421N), junior postdoctoral fellowship (1298825N) and a research visit grant (V426920N) held by Victor Le Pochat. Victor Le Pochat is currently employed by the European Commission. The views and opinions of the author expressed herein are personal and do not necessarily reflect those of the European Commission or other EU institutions.

[9] See `https://deceptive-affiliate-marketing.distrinet-research.be/`.

References

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

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

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

[4] *Ad Creatives Review.* Pushground Blog. Sept. 2, 2021. URL: https://www.pushground.com/blog/ad-creatives-review.

[5] *Add Your Network/Program.* Affplus. URL: https://www.affplus.com/add-network.

[6] *Affiliate Marketer's Guide to Antivirus and VPNs.* Pushground Blog. Apr. 1, 2021. URL: https://www.pushground.com/blog/antivirus-advertising-guide.

[7] affilinc Ltd. *AffiliateFix.* AffiliateFix. 2021. URL: https://www.affiliatefix.com/.

[8] affLIFT, LLC. *affLIFT.* affLIFT. 2021. URL: https://afflift.com.

[9] Maxwell Aliapoulios, Cameron Ballard, Rasika Bhalerao, Tobias Lauinger, and Damon McCoy. "Swiped: Analyzing Ground-truth Data of a Marketplace for Stolen Debit and Credit Cards". In: *USENIX Security.* 2021, pp. 4151–4168. URL: https://www.usenix.org/conference/usenixsecurity21/presentation/aliapoulios.

[10] Eihal Alowaisheq et al. "Cracking the Wall of Confinement: Understanding and Analyzing Malicious Domain Take-downs". In: *NDSS.* 2019. DOI: 10.14722/ndss.2019.23243.

[11] 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/article/investigations/18929-subscription-traps-and-deceptive-free-trials-scam-millions-with-misleading-ads-and-fake-celebrity-endorsements.

[12] *Ban Your Stereotypes about FB.* AdCombo. Apr. 26, 2018. URL: https://blog.adcombo.com/ban-your-stereotypes-about-fb/.

[13] *BBB Scam Alert: Looking for Love? Don't Fall for a Fake Dating Website.* Better Business Bureau. July 9, 2021. URL: https://www.bbb.org/article/news-releases/24477-bbb-scam-alert-looking-love-dont-fall-for-a-fake-dating-service.

[14] Eric van den Berg. "De verborgen industrie die Bitcoinadvertenties op Facebook zet: 'Ze weten niet eens dat we bestaan!'" In: *Brandpunt+* (Aug. 23, 2019). URL: https://www.npo3.nl/brandpuntplus/de-verborgen-industrie-die-bitcoinadvertenties-op-facebook-zet-ze-weten-niet-eens-dat-we-bestaan.

[15] Magdalena Bober. *The Ultimate Guide to Finding Profitable CPA Offers in Affiliate Marketing.* Zeropark. Sept. 25, 2020. URL: https://zeropark.com/blog/ultimate-guide-to-finding-profitable-cpa-offers/.

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

[17] Jack Cable, Ian W. Gray, and Damon McCoy. "Showing the Receipts: Understanding the Modern Ransomware Ecosystem". In: *eCrime.* 2024.

[18] 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.* 2013.

[19] Richard Clayton, Tyler Moore, and Nicolas Christin. "Concentrating Correctly on Cybercrime Concentration". In: *14th Annual Workshop on Economics and Information Security.* 2015.

[20] 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/.

[21] "Commission Notice – Guidance on the Interpretation and Application of Directive 2005/29/EC of the European Parliament and of the Council concerning unfair business-to-consumer commercial practices in the internal market". In: *Official Journal of the European Union* C 526.2021/C 526/01 (Dec. 29, 2021), pp. 1–129. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021XC1229(05).

[22] Patrick D. *How to Find CPA Offers: Choosing the Best CPA Offer to Promote.* Adsterra. Dec. 11, 2021. URL: https://adsterra.com/blog/how-to-find-the-best-cpa-offers/.

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

[24] 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.

[25] Mateja Durovic and Hans W. Micklitz. "International Law on (Un)Fair Commercial Practices". In: *Internationalization of Consumer Law: A Game Changer.* SpringerBriefs in Political Science. Springer International Publishing, 2017, pp. 25–48. DOI: 10.1007/978-3-319-45312-5_3.

[26] *Everything You Need To Know About High-Risk Industries.* LegitScript. Apr. 12, 2023. URL: https://www.legitscript.com/fraud-risk-and-prevention/high-risk-industries/.

[27] Shehroze Farooqi, Álvaro Feal, Tobias Lauinger, Damon McCoy, Zubair Shafiq, and Narseo Vallina-Rodriguez. “Understanding Incentivized Mobile App Installs on Google Play Store”. In: *IMC*. 2020, pp. 696–709. DOI: 10.1145/3419394.3423662.

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

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

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

[31] Federal Trade Commission. “Enforcement Policy Statement on Deceptively Formatted Advertisements”. In: *Federal Register* 81.74 (Apr. 18, 2016), pp. 22596–22601. URL: https://www.ftc.gov/system/files/documents/public_statements/896923/151222deceptiveenforcement.pdf.

[32] Federal Trade Commission. *Fauxmats, False Claims, Phony Celebrity Endorsements, and Unauthorized Charges*. Nov. 16, 2017. URL: https://www.ftc.gov/business-guidance/blog/2017/11/fauxmats-false-claims-phony-celebrity-endorsements-and-unauthorized-charges.

[33] 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*. Apr. 6, 2015. URL: https://www.ftc.gov/news-events/press-releases/2015/04/federal-court-rules-affiliate-marketing-network-its-parent.

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

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

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

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

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

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

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

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

[42] *FTC Policy Statement on Deception*. 1984. URL: https://www.ftc.gov/system/files/documents/public_statements/410531/831014deceptionstmt.pdf.

[43] Hongyu Gao, Jun Hu, and Christo Wilson. “Detecting and Characterizing Social Spam Campaigns”. In: *10th Internet Measurement Conference*. 2010.

[44] 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/.

[45] *Get Your Offers Listed on the Offer Engine*. WOW TRK. URL: https://www.wowtrk.com/list-offers/.

[46] 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.

[47] Sara Hamed and Sara El-Deeb. “Cash on Delivery as a Determinant of E-Commerce Growth in Emerging Markets”. In: *Journal of Global Marketing* 33.4 (Aug. 7, 2020), pp. 242–265. DOI: 10.1080/08911762.2020.1738002.

[48] *High-risk merchant accounts explained*. Stripe. June 10, 2024. URL: https://stripe.com/resources/more/high-risk-merchant-accounts-explained.

[49] 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. DOI: 10.1080/14786011003634415.

[50] Tom Hooker. *Fact Check: Are Crypto Payouts Real?- PropellerAds Blog.* PropellerAds Blog. Sept. 8, 2020. URL: https://propellerads.com/blog/adv-fact-check-are-crypto-payouts-real/.

[51] Tom Hooker. *GEO Master: Understanding the Different Country Tiers.* ActiveRevenue. Mar. 26, 2020. URL: https://activerevenue.com/blog/2020/03/26/geo-master-understanding-the-different-country-tiers/.

[52] *How to Promote Affiliate Offers.* Affise. Oct. 28, 2022. URL: https://affise.com/blog/how-to-promote-affiliate-offers/.

[53] Danny Yuxing Huang et al. "Tracking Ransomware End-to-end". In: *IEEE S&P.* 2018, pp. 618–631. DOI: 10.1109/SP.2018.00047.

[54] Keman Huang, Michael Siegel, and Stuart Madnick. "Systematically Understanding the Cyber Attack Business: A Survey". In: *ACM Computing Surveys* 51.4 (July 2018). DOI: 10.1145/3199674.

[55] Alice Hutchings, Richard Clayton, and Ross Anderson. "Taking down Websites to Prevent Crime". In: *eCrime.* 2016, pp. 1–10. DOI: 10.1109/ECRIME.2016.7487947.

[56] Anton V. Ivanov, Mikhail Kuzin, and Ilya Mogilin. *Shlayer Trojan Attacks One in Ten macOS Users.* Securelist. Jan. 23, 2020. URL: https://securelist.com/shlayer-for-macos/95724/.

[57] 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.

[58] Chris Kanich et al. "Show Me the Money: Characterizing Spam-advertised Revenue". In: *USENIX Security.* 2011.

[59] Mohammad Karami, Shiva Ghaemi, and Damon McCoy. "Folex: An Analysis of an Herbal and Counterfeit Luxury Goods Affiliate Program". In: *eCrime.* 2013, pp. 1–9. DOI: 10.1109/eCRS.2013.6805782.

[60] Mohammad Karami, Youngsam Park, and Damon McCoy. "Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services". In: *25th International Conference on World Wide Web.* 2016, pp. 1033–1043. DOI: 10.1145/2872427.2883004.

[61] Alex Kasprak. "Savage Memes and Lunar Dreams: Deceptive Dating Sites' Intimate Ties to Firefly Aerospace". In: *Snopes.com* (Feb. 12, 2020). URL: https://www.snopes.com/news/2020/02/12/savage-memes-lunar-dreams/.

[62] Aniket Kesari, Chris Hoofnagle, and Damon McCoy. "Deterring Cybercrime: Focus on Intermediaries". In: *Berkeley Technology Law Journal* 32.3 (2017), pp. 1093–1134. URL: https://heinonline.org/HOL/P?h=hein.journals/berktech32%5C&i=1137.

[63] 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: *ASIACCS.* 2020, pp. 394–408. DOI: 10.1145/3320269.3384714.

[64] Platon Kotzias, Leyla Bilge, Sofia Antipolis, and Juan Caballero. "Measuring PUP Prevalence and PUP Distribution through Pay-Per-Install Services". In: *USENIX Security.* 2016.

[65] Simon Krona. "Därför får du mystiska vänförfrågningar på sociala medier". In: *SVT Nyheter* (Mar. 7, 2020). URL: https://www.svt.se/nyheter/darfor-far-du-mystiska-vanforfragningar-pa-sociala-medier.

[66] 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/.

[67] Magdalena Kukułka. *Top Affiliate Offers in 2021.* Zeropark Blog. Feb. 2, 2021. URL: https://zeropark.com/blog/top-affiliate-offers-in-2021/.

[68] Jochem van de Laarschot and Rolf van Wegberg. "Risky Business? Investigating the Security Practices of Vendors on an Online Anonymous Market using Ground-Truth Data". In: *USENIX Security.* 2021, pp. 4079–4095. URL: https://www.usenix.org/conference/usenixsecurity21/presentation/van-de-laarschot.

[69] Victor Le Pochat et al. "A Practical Approach for Taking Down Avalanche Botnets Under Real-World Constraints". In: *NDSS.* 2020. DOI: 10.14722/ndss.2020.24161.

[70] Nick Lenihan. *Complete List of Country Tiers for Affiliate Marketing.* affLIFT. Jan. 14, 2020. URL: https://afflift.com/f/threads/complete-list-of-country-tiers.3444/.

[71] Nektarios Leontiadis, Tyler Moore, and Nicolas Christin. "Measuring and Analyzing Search-Redirection Attacks in the Illicit Online Prescription Drug Trade". In: *USENIX Security.* 2011.

[72] Kirill Levchenko et al. "Click Trajectories: End-to-End Analysis of the Spam Value Chain". In: *S&P.* 2011, pp. 431–446. DOI: 10.1109/SP.2011.24.

[73] *Leveraging Facebook for Affiliate Marketing in 2020.* Everad. July 15, 2020. URL: https://blog.everad.com/en/leveraging-facebook-for-affiliate-marketing-in-2020/.

[74] James Martin and Nicolas Christin. "Ethics in Cryptomarket Research". In: *International Journal of Drug Policy.* Drug Cryptomarkets 35 (Sept. 1, 2016), pp. 84–91. DOI: 10.1016/j.drugpo.2016.05.006.

[75] Willem Marx. "A Sunny Place for a Shady Business". In: *Bloomberg Businessweek* 4724 (Dec. 20, 2021), pp. 42–47. ISSN: 0007-7135.

[76] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. “Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites”. In: *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW Sept. 20, 2019), 81:1–81:32. DOI: 10.1145/3359183.

[77] Damon McCoy, Hitesh Dharmdasani, Christian Kreibich, Geoffrey M. Voelker, and Stefan Savage. “Priceless: The Role of Payments in Abuse-Advertised Goods”. In: *CCS*. 2012, pp. 845–856. DOI: 10.1145/2382196.2382285.

[78] Damon McCoy et al. “PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs”. In: *21st USENIX Security Symposium*. 2012.

[79] *Media Buying Campaign: How to Choose GEO?* Udonis. Dec. 13, 2018. URL: https://www.blog.udonis.co/digital-marketing/media-buying-campaign-geos.

[80] Jeremy B. Merrill and Marshall Allen. ““Trumpcare” Does Not Exist. Nevertheless Facebook and Google Cash In on Misleading Ads for “Garbage” Health Insurance.” In: *ProPublica* (Oct. 20, 2020). URL: https://www.propublica.org/article/trumpcare-does-not-exist-nevertheless-facebook-and-google-cash-in-on-misleading-ads-for-garbage-health-insurance.

[81] Jeremy B. Merrill and Hanna Kozlowska. “How Facebook Fueled a Precious-Metal Scheme Targeting Older Conservatives”. In: *Quartz* (Nov. 19, 2019). URL: https://qz.com/1751030/facebook-ads-lured-seniors-into-giving-savings-to-metals-com/.

[82] Fieke Miedema, Kelvin Lubbertsen, Verena Schrama, and Rolf van Wegberg. “Mixed Signals: Analyzing Ground-Truth Data on the Users and Economics of a Bitcoin Mixing Service”. In: *USENIX Security*. 2023, pp. 751–768. URL: https://www.usenix.org/conference/usenixsecurity23/presentation/miedema.

[83] Uliana Moreva. *CPA Aggregators. Cheating or the Best Solution for Marketers?* Affbank. Aug. 21, 2018. URL: https://affbank.com/blog/https%5C%3A%5C%2F%5C%2Faffbank.com%5C%2Fblog%5C%2Fcpa_aggregators.

[84] 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.

[85] Terry Nelms, Roberto Perdisci, Manos Antonakakis, and Mustaque Ahamad. “Towards Measuring and Mitigating Social Engineering Software Download Attacks”. In: *USENIX Security*. 2016, pp. 773–789.

[86] Arman Noroozian, Jan Koenders, Eelco van Veldhuizen, Carlos H Ganan, Sumayah Alrwais, Damon McCoy, and Michel van Eeten. “Platforms in Everything: Analyzing Ground-Truth Data on the Anatomy and Economics of Bullet-Proof Hosting”. In: *USENIX Security*. 2019, pp. 1341–1356.

[87] Masarah Paquet-Clouston and Sebastián García. “On the motivations and challenges of affiliates involved in cybercrime”. In: *Trends in Organized Crime* (Dec. 2022). DOI: 10.1007/s12117-022-09474-x.

[88] Masarah Paquet-Clouston, Serge-Olivier Paquette, Sebastian Garcia, and Maria José Erquiaga. “Entanglement: cybercrime connections of a public forum population”. In: *Journal of Cybersecurity* 8.1 (Jan. 2022). DOI: 10.1093/cybsec/tyac010.

[89] Sergio Pastrana, Daniel R. Thomas, Alice Hutchings, and Richard Clayton. “CrimeBB: Enabling Cybercrime Research on Underground Forums at Scale”. In: *2018 World Wide Web Conference*. 2018, pp. 1845–1854. DOI: 10.1145/3178876.3186178.

[90] Elissa M. Redmiles, Neha Chachra, and Brian Waismeyer. “Examining the Demand for Spam: Who Clicks?” In: *2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–10. DOI: 10.1145/3173574.3173786.

[91] 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/en-us/medialibrary/PDFs/technical%20papers/samosseikovb2009paper.pdf.

[92] Emily Schechter, Giacomo Gnecchi Ruscone, and Badr Salmi El Idrissi. *Notifying Users of Unclear Subscription Pages.* Chromium Blog. Nov. 8, 2018. URL: https://blog.chromium.org/2018/11/notifying-users-of-unclear-subscription.html.

[93] 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.

[94] Craig Silverman. “How A Massive Facebook Scam Siphoned Millions Of Dollars From Unsuspecting Boomers”. In: *BuzzFeed News* (Oct. 16, 2019). URL: https://www.buzzfeednews.com/article/craigsilverman/facebook-subscription-trap-free-trial-scam-ads-inc.

[95] Craig Silverman and Trevor Davis. “Coronavirus Mask Ads Were Emailed To Millions Of Americans With Unsafe Claims And Inflated Prices”. In: *BuzzFeed News* (Apr. 21, 2020). URL: https://www.buzzfeednews.com/article/craigsilverman/coronavirus-mask-ads-were-emailed-to-millions-of-americans.

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

[97] 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*. 2013, pp. 55–78.

[98] Karthika Subramani, Xingzi Yuan, Omid Setayeshfar, Phani Vadrevu, Kyu Hyung Lee, and Roberto Perdisci. "When Push Comes to Ads: Measuring the Rise of (Malicious) Push Advertising". In: *IMC*. 2020, pp. 724–737. DOI: 10.1145/3419394.3423631.

[99] Sean Sullivan. *I May Never Text Again: More Facebook Spam.* News from the Lab - F-Secure Labs. Aug. 24, 2010. URL: https://archive.f-secure.com/weblog/archives/00002016.html.

[100] Janos Szurdi, Meng Luo, Brian Kondracki, Nick Nikiforakis, and Nicolas Christin. "Where Are You Taking Me? Understanding Abusive Traffic Distribution Systems". In: *Web Conference 2021*. 2021, pp. 3613–3624. DOI: 10.1145/3442381.3450071.

[101] Kurt Thomas et al. "Framing Dependencies Introduced by Underground Commoditization". In: *14th Annual Workshop on the Economics of Information Security*. 2015.

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

[103] *Top Affiliate Marketing Verticals of 2021.* Pushground Blog. Aug. 5, 2021. URL: https://www.pushground.com/blog/top-affiliate-marketing-verticals.

[104] *Types of Inventory Available.* Zeropark. Sept. 26, 2020. URL: https://web.archive.org/web/20200926093930/https://doc.zeropark.com/en/inventory_types.html.

[105] 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.

[106] Tim Verheyden, Fabienne Meijer, and Amra Dorjbayar. "Wij klikten op een valse advertentie met Philippe Geubels zodat u het niet hoeft te doen". In: *VRT NWS* (Sept. 9, 2019). URL: https://www.vrt.be/vrtnws/nl/2019/09/04/wij-klikten-op-de-nepadvertenties-met-philippe-geubbels-en-kwame/.

[107] David Y. Wang, Matthew Der, Mohammad Karami, Lawrence Saul, Damon McCoy, Stefan Savage, and Geoffrey M. Voelker. "Search + Seizure: The Effectiveness of Interventions on SEO Campaigns". In: *2014 Internet Measurement Conference*. 2014, pp. 359–372. DOI: 10.1145/2663716.2663738.

[108] *Washington State AG and Facebook Target "Clickjackers" | Washington State.* Washington State Office of the Attorney General, Jan. 26, 2012. URL: https://www.atg.wa.gov/news/news-releases/washington-state-ag-and-facebook-target-clickjackers.

[109] Rolf van Wegberg et al. "Plug and Prey? Measuring the Commoditization of Cybercrime via Online Anonymous Markets". In: *USENIX Security*. 2018, pp. 1009–1026.

[110] Simona Weinglass. "Israel Bans Binary Options Industry, Finally Closing Vast, 10-Year Fraud". In: *The Times of Israel* (Oct. 23, 2017). URL: https://www.timesofisrael.com/israel-bans-entire-binary-options-industry-finally-closing-vast-10-year-fraud/.

[111] Wewe Media. *Complete List of TIERS & Countries.* AffiliateFix. Mar. 1, 2018. URL: https://www.affiliatefix.com/threads/complete-list-of-tiers-countries.147158/.

[112] *What Is a Tier of Traffic and What Tier Should You Choose?* PropellerAds Blog. Mar. 20, 2018. URL: https://propellerads.com/blog/what-is-a-tier-of-traffic-and-what-tier-should-you-choose/.

[113] *What Tier to Choose for Ads Campaign.* RichAds Blog. Aug. 11, 2020. URL: https://richads.com/blog/what-tier-to-choose-profitable-geos-for-advertising-campaigns-in-tier-1-2-or-3/.

[114] 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/.

[115] X. Xu et al. "Dissecting Mobile Offerwall Advertisements: An Explorative Study". In: *2020 IEEE 20th International Conference on Software Quality, Reliability and Security*. 2020, pp. 518–526. DOI: 10.1109/QRS51102.2020.00072.

[116] Hao Yang et al. "Casino Royale: A Deep Exploration of Illegal Online Gambling". In: *2019 Annual Computer Security Applications Conference*. 2019, pp. 500–513. DOI: 10.1145/3359789.3359817.

[117] Zheng Yang, Joey Allen, Matthew Landen, Roberto Perdisci, and Wenke Lee. "TRIDENT: Towards Detecting and Mitigating Web-based Social Engineering Attacks". In: *USENIX Security*. 2023, pp. 6701–6718. URL: https://www.usenix.org/conference/usenixsecurity23/presentation/yang-zheng.

[118] Javier Yuste and Sergio Pastrana. "Avaddon Ransomware: An in-Depth Analysis and Decryption of Infected Systems". In: *Computers & Security* 109 (Oct. 1, 2021), p. 102388. DOI: 10.1016/j.cose.2021.102388.

[119] Conradin Zellweger and Fabian Kohler. *How scam networks use fake celebrity ads to lure online investors.* SWI swissinfo.ch. July 5, 2024. URL: https://www.swissinfo.ch/eng/life-aging/how-scam-networks-use-fake-celebrity-ads-to-lure-investors/82568794.

[120] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. "Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites". In: *ConPro*. 2020.

[121] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. "What Makes a "Bad" Ad? User Perceptions of Problematic Online Advertising". In: *CHI*. 361. 2021. URL: https://doi.org/10.1145/3411764.3445459.

[122] Zulma Corporation Limited. *Add Your Network.* OfferVault. URL: https://offervault.com/add-your-network.