Exploring Malvertising Driven by Brand Impersonation in Search Engine Ads

Hans Dekker

Exploring Malvertising Driven by Brand Impersonation in Search Engine Ads

THESIS

submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE

in

COMPUTER SCIENCE

by

Hans Dekker
born in Den Haag, the Netherlands

Cybersecurity Group
Intelligent Systems Department
Faculty EEMCS, Delft University of Technology
Delft, the Netherlands
www.ewi.tudelft.nl

BforeAI
108 W 13TH ST, STE 100
Wilmington, DE, United States
www.bfore.ai

Exploring Malvertising Driven by Brand Impersonation in Search Engine Ads

Author: Hans Dekker
Student id: 4692357

Abstract

Malvertising is a significant threat, in which attackers leverage online advertisements to deceive users and distribute scams, phishing pages, and malware. While prior research has largely focused on low-tier ad networks and high-risk websites, this study examines brand impersonation in mainstream search advertising platforms, specifically Google Ads.

We queried Google with brand-related search terms, capturing and analyzing the advertisements displayed to assess the scale and nature of impersonation. Over a 24-day period, our scraper collected a dataset of 52k ads across 605 brands, extracting key features such as advertiser identity, redirection chains, and landing page content.

Using a combination of manual inspection and six brand-agnostic heuristics, we identify various forms of abuse, including phishing pages, tech support scams, and a previously undocumented category, affiliate brand bidding. This last technique, in which affiliates place search ads to divert users through affiliate links, affects at least 189 brands in our dataset.

In total, 4,160 ads (7.9%) were flagged as abusive, 3781 of which involved affiliate brand bidding. Our results further reveal that verified Google Ads accounts are being rented or resold, enabling systematic evasion of identity checks. These findings expose enforcement gaps in Google's ad review and verification systems.

Thesis Committee:

Chair:	Dr. H.J. Griffioen, Faculty EEMCS, TU Delft
University supervisor:	Dr. H.J. Griffioen, Faculty EEMCS, TU Delft
Company supervisor:	Dr. P. Kotzias, BforeAI
Committee Member:	Dr. M.J.G. Olsthoorn, Faculty EEMCS, TU Delft

Preface

This thesis has been an exciting project, combining technical challenges with real-world relevance. I really enjoyed the opportunity to explore this topic in depth, and I'm grateful to all my family and friends who supported me along the way. I'd especially like to thank a few people who played an important role in making this journey both possible and enjoyable:

First, a huge thanks to my supervisor, Platon, for his valuable support, insightful feedback, and steady mentorship throughout this process. Your guidance really helped shape this project.

To Luigi and Sebastian, thank you for giving me the opportunity to explore this topic at BforeAI. It's been an amazing experience, and I've learned a lot.

To my amazing girlfriend Elise, thank you for designing the cover for this report, it's by far the most stylish part of this thesis.

To Jip, Mette, Marijn, and Wessel, thanks for being great company at the library these past few months. The coffee breaks, lunches, and (very importantly) the rides to the library, all made the process much more enjoyable.

To Professor Griffioen, thank you for being open-minded and supportive of doing a company-based thesis. That flexibility made this whole project possible.

Hans Dekker
Delft, the Netherlands
April 17, 2025

Contents

List of Figures

Chapter 1

Introduction

The ad ecosystem is a central pillar of the modern Internet, that is used to fund free content, services, and applications that users rely on daily. Digital advertising now drives a multibillion-dollar industry, with global ad spending projected to exceed $790 billion annually in 2025 [7], as advertisers seek to reach audiences across platforms and devices through targeted, data-driven approaches. Various malicious actors attempt to exploit this ecosystem in different forms including click-fraud, phishing and scam distribution, affiliate marketing fraud, ad-injection, etc.

Most previous research on the usage of advertisements as a vector to spread malware or scams has focused on display ads: graphical or multimedia-based advertisements embedded within webpages, typically showing up as banners, pop-ups, or video ads [36, 55, 52, 54]. However, recent reports indicate a growing shift toward search engine ads as a tool for delivering attacks [43, 49, 50].

A specific tactic known as brand impersonation involves malicious actors creating ads that mimic legitimate search results of well-known brands, deceiving users into believing that they are clicking on an official site [44]. Since more than 44% of Google searches are for branded search terms [31], this attack vector has the potential to harm many users. In fact, security vendor Netskope reports that in their measurements, the majority of clicks that lead to phishing sites originated from search engines, either through advertisements or SEO poisoning [6].

Although Google Search Ads represent a significant percentage of the digital advertising market, there has been little research on these threats, leaving significant gaps in understanding their characteristics, prevalence, and potential detection strategies.

Prior work has shown that brand-focused search queries are an effective way to surface abuse targeting popular brands [27]. Inspired by this, we design a search-driven methodology that collects Google Search ads shown for such brand queries.

1.1 Research Question

Since search engine malvertising remains an understudied threat, and previous research has primarily focused on ads served on publisher sites, our goal is to analyze how brand

impersonation-based malvertising operates in Google Search Ads.

This leads to the central research question:

> **How does malvertising driven by brand impersonation manifest in Google Search Ads, and what can be learned from its observable types, prevalence, and advertiser behavior in scraped ad data?**

To answer this, we break the problem down into three sub-questions:

- **Q1: Can we systematically observe brand impersonation malvertising by collecting Google Search ads targeted at a broad selection of brands?**
  - How can we design a scraping setup to collect real-world instances of search-based malvertising?
- **Q2: What types of malvertising occur in Google Search ads?**
  - What categories of deceptive ads can be observed and what types of landing pages do these ads direct users to?
  - How prevalent are these different malvertising types in our dataset?
- **Q3: What patterns characterize malicious advertisers in Google Search?**
  - Can we identify recurring patterns in advertiser behavior?
  - What techniques allow these advertisers to continue operating despite Google's security mechanisms?

By addressing these questions, this thesis aims to provide a **structured understanding of search-based malvertising**, uncovering how these fraudulent ad campaigns are executed, how they persist despite enforcement measures, and what implications this has for online security.

1.2 Approach and Contributions

This thesis investigates search-based malvertising through a pipeline that combines automated data collection, heuristic-based detection, and in-depth analysis of advertiser behaviors.

Approach. We developed an automated scraping system to observe malvertising on Google Search by simulating behavior reflecting user search intent. The pipeline takes a list of brands across diverse industries, such as fintech, e-commerce, marketing tools, including both Fortune 500 companies and startups, and systematically combines each with relevant user intent keywords like "login," "support," or "pricing". For each brand–keyword combination, our system performs a Google Search query, captures all sponsored results, and follows ad clicks to their landing pages. We extract and store the ad's metadata (title, display

URL, advertiser identity), screenshots and HTML of the search results and landing pages, and network request logs (HAR). The collected data is stored in a structured database that allows us to track ad reoccurrence, compare different advertisers, and inspect landing-pages. After manual inspection of the data we developed several heuristic checks (e.g., domain mismatch, VirusTotal lookups, suspicious URL parameters), to flag abusive ads and categorize the tactics used by malicious advertisers into common malvertising types, including phishing pages, tech support scams, and affiliate brand bidding.

Using OSINT techniques, we then systematically investigated the advertisers most frequently flagged for abuse, evaluating whether verified accounts were being reused or resold for abuse.

Key Contributions.

1. **Automated Data Collection Pipeline.** We designed and implemented a scraper that systematically gathered ads and landing page data by executing brand-focused queries on Google Search. This pipeline produced a dataset of **52K ads** targeted at **605 brands** across multiple industries before scraping restrictions made large-scale collection infeasible.

2. **Heuristic Detection & Brand-Specific Abuse Prevalence** After manual inspection, we developed six automated checks to identify diverse forms of malvertising. Applied to our dataset of over 52K ads, these heuristics collectively flagged 4,160 ads (7.9%) as abusive. The most frequent flags included affiliate brand bidding (3,781 ads, 7.1%), discount/referral site abuse (153 ads, 0.29%), and malicious landing domains as detected by VirusTotal (84 ads, 0.16%). For some individual brands, a significant percentage of the ads targeted at their brand name were flagged as malvertising, for example, over 10% of all ads for McAfee and Facebook were flagged.

3. **Identification of Affiliate Brand Bidding.** We uncover a new form of brand-targeted abuse, *affiliate brand bidding*, in which affiliates bid on brand keywords, redirecting users through hidden affiliate links and imposing costs on both brands and legitimate affiliates. To our knowledge, this phenomenon has not been previously described in academic literature. It appears in 3781 (7.1%) of our collected ads and targets a wide range of brands, with 189 brands (31% of all brands in our study) affected by affiliate brand bidding at least once.

4. **Discovery of a Large-Scale Advertiser Verification Vulnerability.** Our OSINT investigation of advertisers repeatedly appearing in flagged abusive ads, reveals how privileged Google Ads accounts, with verified advertiser status, that often run tens of thousands of ads according to the Google Ads Transparency Center, are being rented or resold to malicious actors, allowing them to evade Google's identity verification requirements and continue running abusive ads at scale.

5. **Policy Recommendations for Google Ads.** Based on these findings, we propose improvements to policy enforcement, such as strengthening advertiser verification measures and refining Display URL vs Final URL checks.

Overall, our work reveals how malicious advertisers operate on Google Search, from basic brand impersonation to more intricate techniques, and demonstrates a critical need for stronger oversight and enforcement within the platform.

Comparison with Closest Existing Work

While prior research on malvertising has largely focused on display ads, a recent study by Musteata et al. (2024) [39] is most methodologically similar to ours. They collect Google Search ads using a Playwright-based crawler and identify impersonation based on an allowlist of legitimate landing domains for 50 well-known brands. Their method successfully identifies 513 malicious ads in a dataset of 10K sponsored results. However, this detection strategy would require significant manual tuning for larger brand sets and requires brand-specific rules that limit generalisation.

In contrast, our study collects 52K ads across 605 brands using a scalable detection approach based on six brand-agnostic heuristics, which flag over 4K ads as malvertising. We also analyze abuse at the advertiser identity level, uncovering systemic enforcement gaps in Google's verification process, particularly the use of rented, verified accounts by malicious actors. Additionally, we document a previously undescribed, yet prevalent, category of abuse, affiliate brand bidding, in which affiliates exploit brand-targeted search ads to redirect users through affiliate links. Such a form of abuse would not be surfaced with a methodology that starts from a brand domain allowlist, as it involves legitimate domains used in unauthorised ways.

Our work expands the scope of search-based malvertising research and introduces techniques for more generalisable and scalable detection.

1.3 Outline

In Chapter 2 background information is provided on the ad ecosystem and specifically search engine ads, malvertising and brand impersonation, as well as an introduction to affiliate marketing, to contextualise the new form of abuse we identified: affiliate brand bidding. Next, Chapter 3 gives an overview of the related work, and shows the research gap this thesis aims to bridge. In Chapter 4 we explain the methodology used to address the research questions and obtain our results. Then, Chapter 5 describes the results obtained, analyzing the ads flagged by our heuristics and detailing the findings of our OSINT investigation into advertisers repeatedly engaged in abusive practices. Finally, Chapter 6 provides a brief discussion and reflection on this work. It answers the research questions stated earlier and outlines the limitations of this study. Additionally, it presents recommendations for future research, as well as practical recommendations for both Google and brands to mitigate the identified issues.

Chapter 2

Background

2.1 The Digital Advertising Ecosystem

Online advertising supports a large portion of the Internet, funding free content, services, and applications. Advertisers bid on keywords or user demographics to reach potential customers, while publishers earn revenue by displaying these ads. This complex ecosystem has many different layers (such as *ad networks*, *exchanges*, and *programmatic ad platforms*) that allow advertisers to target users across different websites and platforms.

2.2 Search Engine Advertising

One of the most prominent forms of digital advertising is search engine advertising, where sponsored listings appear above or alongside organic search results. Platforms such as Google Ads use an auction-based model: advertisers bid on keywords (e.g., “laptop deals,” ”bank login”), and the search engine evaluates Cost-Per-Click (CPC) bids, relevance, and other factors to determine which ads to show [10].

A typical search ad includes:

- A **Title** and **ad copy**,
- A **Display URL**, which the user sees (e.g., `www.example.com`),
- A **Final URL**, which is the actual landing page the user reaches after clicking [8].

To prevent advertisers from misleading users, Google’s policy requires that the Display URL matches the second-level domain of the Final URL [8]. For example, if the Display URL is `brand.com`, then a Final URL like `brand.com/product` is valid (matching the `brand.com` domain), while `brandhelpdesk.net/support` would be invalid because it has a different second-level domain.

Figure 2.1 shows an example of a Google Search results page with a sponsored ad.

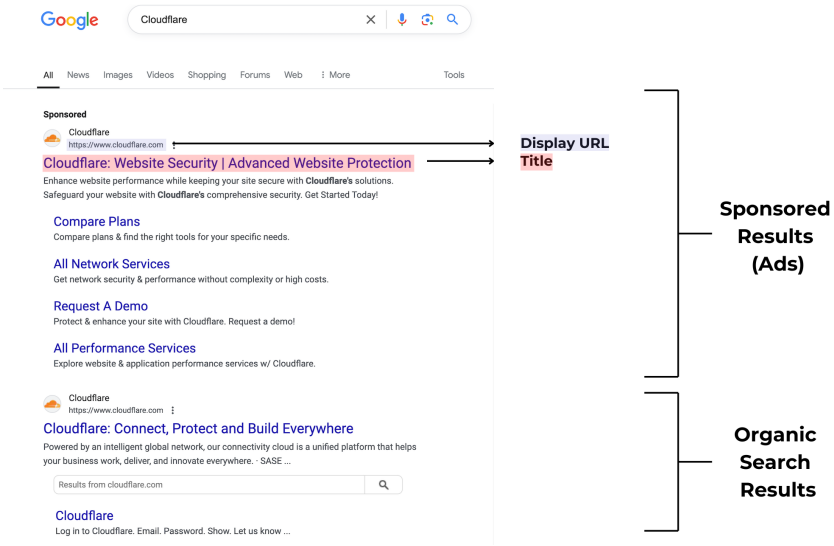

Figure 2.1: Example Google Search Results Page with Advertisement

2.2.1 Ads Transparency Center and "About This Advertiser"

When a user clicks the three dots next to an ad's Display URL (Figure 2.1), a pop-up shows advertiser details such as the legal name and location, as seen in Figure 2.2. This disclosure helps users identify the advertiser behind the ad. Google also offers an *Ads Transparency Center*, a searchable repository that shows previous ads served by a given advertiser on Google platforms (Figure 2.3), this overview only exists for advertisers that are verified. Google requires identity verification to get access to all features, and in case of suspicious behavior [3], the exact details of when Google will force an advertiser to verify themselves, are not clearly specified. In some regions, regulatory requirements mandate Google publicly disclose additional details, including targeting criteria and audience size [2].

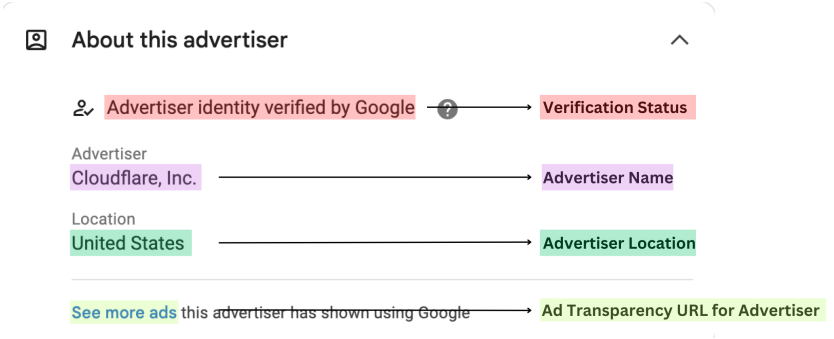

Figure 2.2: Google "About this advertiser" Information

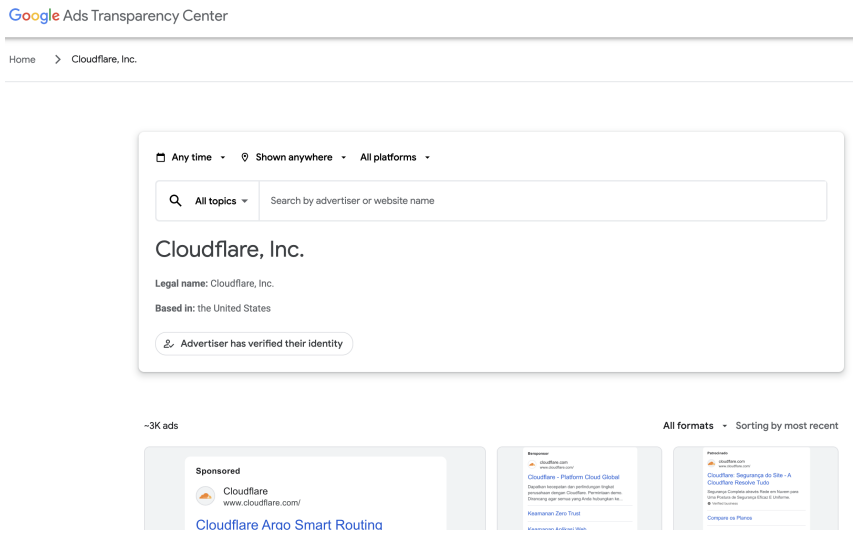

Figure 2.3: Example Google Ads Transparency Center Page

While these transparency features allow users to verify the advertisers behind an ad, they rely on users actively clicking to check the advertiser details. In practice, many users may not look at this information and instead assume that a prominently placed ad belongs to the legitimate brand it appears to represent.

2.3 Malvertising and Brand Impersonation

2.3.1 Malvertising Overview

Malvertising combines "malicious" and "advertising" to describe scenarios where ads are used to distribute malware, scams, or social engineering attacks [42]. While the majority of malvertising research has focused on *display ads* embedded in webpages [36, 55, 52, 54], recent reports suggest that *search ads* are increasingly used to serve the same malicious purpose [49, 50].

2.3.2 Brand Impersonation in Search Ads

A common tactic in search engine malvertising is brand impersonation, where attackers use a trusted brand's name, logo, or layout to make their ad appear as if it originates from that brand [44]. Because 44% of Google searches are estimated to be brand-centric [31], the impersonation of popular brands has the potential to harm many users, who are likely to trust the top result of the search engine results page (SERP). Figure 2.4 provides an example of brand impersonation in an ad our scraper encountered, where the ad appears to be from Zoom, but actually leads to the malicious page shown in 2.5.

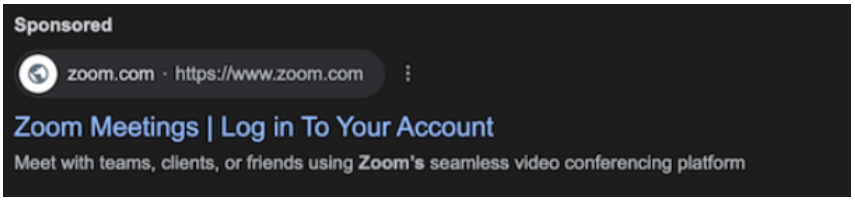

Figure 2.4: Example of Brand Impersonation targeting Zoom

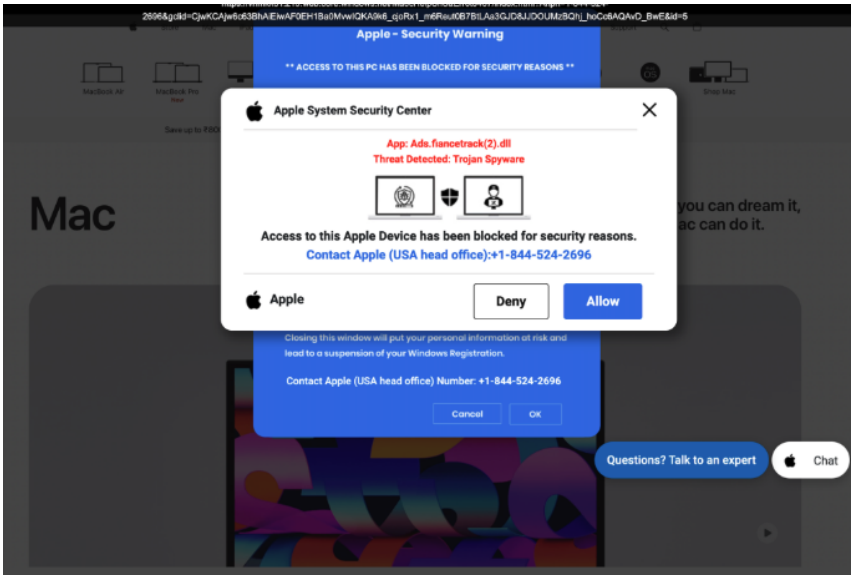

Figure 2.5: Example of Brand Impersonation targeting Zoom, Final Landing Page

2.3.3 Cloaking and Brand Impersonation Scams

Malicious advertisers often use *cloaking* to disguise harmful content. By detecting the user-agent or IP address, they serve benign pages to Google reviewers or automated crawlers, while redirecting real visitors to phishing or scam sites [12]. This allows the ad shown in 2.4 to use the Display URL ”zoom.com”, even though the ad redirects real users to a different Final URL, which is against Google policy. Cloaking enables these ads to remain active longer before detection and removal.

2.4 Affiliate Marketing and Its Abuse

Affiliate marketing is a legitimate, performance-based model where affiliates earn commissions for driving traffic or sales via tracked links or codes. Many major brands invite affiliates to advertise their products, paying a commission for each sale or conversion.

However, affiliate marketing can also be exploited by malicious actors. One form of abuse relevant to search engine ads is affiliate brand bidding, where affiliates bid on key-

words containing the brand's name (e.g., "Jira Pricing"). Instead of harming users directly, this tactic intercepts traffic already intending to visit the brand's official site, forcing the brand to pay commissions to an affiliate that did not actually participate in the promotion of the brand.

In 2.6 we show how a brand-bidding ad appears nearly identical to an organic search result. However, when a user clicks on the ad instead of the organic listing, the affiliate running the ad is credited for any resulting conversions, even though they did not contribute to the original customer acquisition.

This strategy can also have a negative impact on other affiliates. If the user, originally led to the brand by another affiliate, later searches for the brand on Google and clicks a brand-bidding affiliate's link, the first legitimate affiliate usually loses their claim at a commission due to the common last-click attribution model.

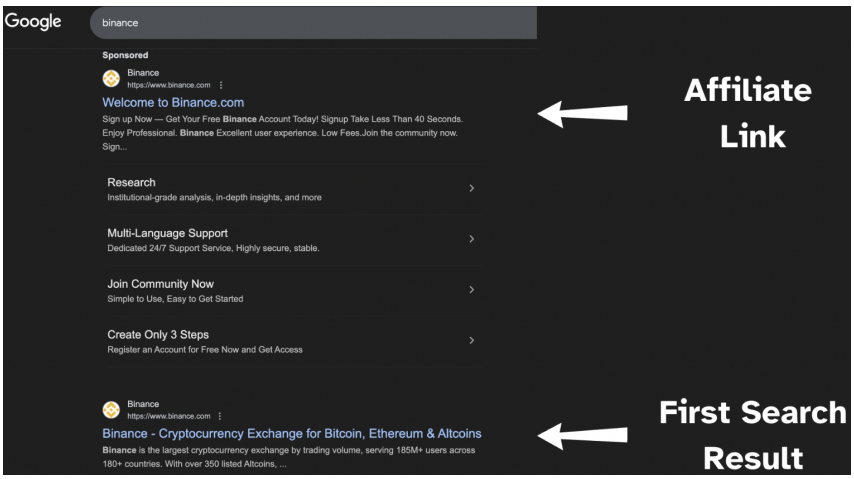

Figure 2.6: Brand Bidding Ad Targeting Binance

As we'll further discuss in Chapters 5 and 6 We found that this form of abuse is widespread on Google Search Ads, and current Google policies fail to address it.

2.5 Summary

In this chapter, we have introduced the fundamental concepts of search engine advertising, malvertising and brand impersonation necessary to understand the rest of this thesis. We also explained how legitimate affiliate marketing can be abused via affiliate brand bidding, which we will later show is a common issue in Google Search Ads.

Chapter 3

Related Work

This chapter provides an overview of existing work on malvertising, malicious activity in search engines, brand impersonation, and social engineering detection. We describe various techniques (clustering, JavaScript anomaly detection, visual classification, fingerprinting, and evasion countermeasures) used to identify malicious ads or other scam content. Finally, we show how our approach fills important gaps, focusing on brand impersonation within mainstream search engine ads.

3.1 Introduction and Terminology

Malvertising is the malicious use of digital advertisements to distribute malware, scams, or other malicious content. Attackers target low-tier ad networks, which are less regulated and serve ads on smaller publisher sites, as well as mainstream platforms, such as search engines and large social media platforms.

Social engineering (SE) refers to deceptive tactics that manipulate users into taking actions (e.g., clicking malicious links or revealing credentials).

Brand impersonation is a type of SE attack in which attackers mimic well-known brand names, domains, or appearances to gain user trust. When combined with malvertising, this results in legitimate-looking ads that lead unsuspecting users to scam or phishing sites.

3.2 Malvertising and Search Engine Ads

3.2.1 Early Studies on Malvertising on Publisher Websites

Historically large-scale malvertising research focused on ads embedded on publisher websites. Two foundational examples are:

- **Li et al. (2012)** [36] did a study over three months, collecting 25 million redirection chains linked to online ads. By identifying malicious nodes with characteristics that indicated malware, scams, or click-fraud, they highlighted that even leading ad networks were vulnerable to malvertising.

- **Zarras et al. (2014)** [55] performed one of the first large-scale analyses of malvertising. They crawled pages sourced from antivirus (AV) logs and Alexa's top-ranked sites, using Wepawet (a malware detection service), malware/phishing blacklists, and VirusTotal for ad classification. In line with [36], their work also showed that even reputable ad services could serve malvertisements.

Study	Data Collection	Detection Method
Li et al. (2012) [36]	3-month crawl, ∼25 M redirection chains	Identified malicious nodes using malware/scam signatures
Zarras et al. (2014) [55]	Crawling AV logs and Alexa top sites	Combination of Wepawet, VirusTotal and blacklists

Table 3.1: Representative early malvertising studies on publisher websites.

Despite their scale, these works examined ads embedded on publisher sites rather than malvertisements in search engine results, which both anecdotal reports [44, 30, 45] and systematic studies [6] by security vendors report to be more prevalent today. They also do not focus specifically on brand impersonation.

3.2.2 Malvertising in Free Live Streaming Sites

Other research has explored malvertising in specific environments, such as free live streaming platforms:

- **Zubair et al. (2016)** [59] examined deceptive `<iframe>` overlays that trick users into clicking ads by invisibly covering a streaming player on free live streaming sites. This is an example of research analyzing the dynamics of more high-risk Internet environments, rather than mainstream, trusted parts of the Internet such as Google search results.

3.2.3 Problematic Advertising on Mainstream News, Misinformation, and High-Traffic Sites

Later studies targeted more popular sites to capture broad malvertising behaviors:

- **Zeng et al. (2020)** [56] analyzed deceptive ads on mainstream news and misinformation websites, identifying tactics (e.g., fake endorsements and misleading financial claims) that circumvent platform policies.

- **Zeng et al. (2021)** [57] focused on political advertising via news and media sites. They used a classifier to identify political ads, and manually labeled ads for a systematic analysis that revealed patterns of problematic political advertising.

- **Ali et al. (2023)** [25] deployed a browser extension to passively track ads served on Facebook to participants. Their survey data highlight the user experience and perceptions of intrusive or deceptive ads.
- **Yan et al. (2023)** [53] introduced *ADHERE*, a large-scale crawl of Alexa's top 1M sites, aiming to detect intrusive ad behaviors (e.g., forced redirects, deceptive UI elements).

These works confirmed that malicious or misleading ads are not limited to low-tier networks but can appear on high-traffic sites. Note that especially [56, 57, 25] focus more on problematic ads that are somewhat deceptive in their claims, contain clickbait etc., while usually malvertising refers to ads that either distribute malware or pages containing other social engineering attacks.

3.2.4 Search Engine-Based Threats

Three notable efforts highlight social engineering in search results and brand queries, though from different angles:

- **Invernizzi et al. (2016)** [33] studied blackhat web cloaking techniques that detect and evade security crawlers. Their work focused on Google Search and Google Ads, highlighting how malicious advertisers hide harmful content from automated systems by using IP blacklists and contextual fingerprinting. They developed an anti-cloaking system to identify split-view content returned to different browsing profiles, achieving 95.5% accuracy on a dataset of 94,946 URLs. They studied the prevalence of cloaking in an unlabeled dataset of 135,577 Google Search and Ads URLs.
- **Carpineto et al. (2020)** [27] performed a study on counterfeit detection by submitting brand-related queries to search engines. They trained a model to classify results into legitimate or fake E-commerce and achieved high accuracy. While focused on search results instead of advertisements, this shows that brand-related keywords can help reveal fraudsters exploiting well-known brand identities. This study has a very similar methodology to ours, although our scope is broader than counterfeit detection for E-commerce.
- **Koide et al. (2020)** [35] built *STRAYSHEEP*, a system to collect and detect multi-step SE attacks by searching for high-risk keywords (e.g., *crack*, *stream online*) on search engines and social media platforms. Their method focuses on extended redirection chains and multi-step user interactions, not brand impersonation ads.

Although these studies show how search engines can be abused to distribute harmful or deceptive content, [27, 35] focus on organic search results rather than paid ads. [33] does include Google Ads to study the prevalence of cloaking in those, but only focuses on the practice of cloaking. Our work focuses on Google Ads shown for brand-related queries, a domain where brand impersonation can have a high impact due to user trust in familiar brands and in the integrity of the ads shown by Google.

Study	Data Collection	Focus
Invernizzi et al. (2016) [33]	135,577 search and advertisement URLs on high-risk search terms	Studying prevalence of cloaking in Google Search/Ads and developing an anti-cloaking system
Carpineto et al. (2020) [27]	Brand keywords, 34K search results	Counterfeit e-commerce site classification
Koide et al. (2020) [35]	High-risk keywords (e.g. "crack")	Automated crawling to detect multi-step SE attacks

Table 3.2: Related work on search-engine-based social engineering.

3.3 Techniques for Detecting Malicious Ads and Scams

Across malvertising and phishing research, multiple detection strategies have emerged. We provide an overview five primary approaches that researchers have used to detect social engineering attacks, malicious advertisements, or fraudulent brand impersonation.

3.3.1 Clustering Approaches

Clustering exploits the fact that malicious campaigns typically share content or behavioral similarities.

- **Vadrevu et al. (2019)** [52] crawled publisher sites relying on low-tier ad networks. By clustering malicious ads, they identified large-scale social engineering (SE) campaigns.
- **Stivala et al. (2023)** [48] used clustering to analyze *clickbait PDFs*—documents luring users into visiting malicious webpages. They compiled 176K PDFs, and by analyzing visual similarity, volumetric and temporal properties they identified 44 clickbait clusters.

Both works highlight the effectiveness of clustering to identify patterns of malicious content in large datasets.

3.3.2 JavaScript Anomaly Detection

By monitoring *client-side* behaviors for anomalies or known malicious patterns, researchers can often detect deceptive ads or domains:

- **Yang et al. (2023)** [54] introduced *TRIDENT*, which flagged JavaScript anomalies (e.g., suspicious function calls, redirects) as potential social engineering ads (SE-ads). They crawled over 100K sites, detecting 1,479 SE events among 259K JavaScript navigation actions.

- **Zubair et al. (2016)** [59], mentioned earlier in 3.2.2, analyzed deceptive `<iframe>` overlays on free live streaming sites by using heuristics to detect malicious client-side patterns they identified in their analysis.

3.3.3 Visual Classifiers

Visual similarity approaches are also common, especially in papers on phishing detection, where attackers copy brand logos or site layouts:

- **Liu et al. (2022)** [38] combined visual similarity with user interaction analysis (e.g., credential input fields). This fusion mitigates false positives when pages visually mimic a known brand but do not contain other elements confirming a phishing intention.
- **Ozen et al. (2024)** [41] introduced a deep-learning model for social engineering (SE) detection, collecting 650K screenshots. While shown to be highly effective for visually deviant pages, brand impersonation ads in search engines are completely text-based and therefore visually minimalistic, limiting purely visual signals at the ad level.
- Other examples include **Abdelnabi et al. (2020)** [22] who learn profiles for websites to detect phishing pages based on visual similarity and **Lin et al. (2021)** [37], who leverage recognition of logos and variants to improve upon page-based visual similarity detection.

3.3.4 Fingerprinting Techniques

Infrastructure-level fingerprinting can detect emerging scams:

- **Bijmans et al. (2021)** [26] monitored TLS certificate issuance to find domains likely to host phishing kits. They used fingerprinting of phishing kits they collected on Telegram to measure the prevalence and lifecycle of such kits on the Internet.

While a powerful technique to detect known phishing kits instantly, these methods detect landing signals on the landing page level rather than detecting brand impersonation or other signals on the ad level.

3.4 Detection Evasion Countermeasures

More sophisticated attackers use **cloaking** (serving one version of a web page to security crawlers and a different one to real users) or **redirection chains** to hide the final malicious page:

- **Invernizzi et al. (2016)** [33] studied the prevalence of cloaking in Google Search/Ads, see 3.2.4 for more details.

- **Oest et al. (2019)** [40] examined the effectiveness of cloaking techniques used by live phishing sites against browser phishing blacklists and **Zhang et al. (2021)** [58] did a new study on the prevalence of cloaking systems in phishing sites as well as their effectiveness to avoid browser blacklists.

- **Acharya et al. (2021)** [23] did a systematic study evaluating security crawlers against various cloaking techniques, finding multiple novel cloaking weaknesses in the crawling ecosystem.

In the context of Google Ads, Google prohibits cloaking and has rules against unexpected redirects, forcing advertisers to use the same final URL domain as the display URL. In our research, we discovered that malicious actors are able to abuse Google's own click trackers to circumvent these restrictions.

3.5 Brand Impersonation Studies

Two related works focus on brand impersonation:

- **Carpineto et al. (2020)** [27], showed the potential of using brand-related search queries to expose counterfeit e-commerce schemes in search engines. See 3.2.4 for more details.

- **Acharya et al. (2024)** [24] investigated brand impersonation *on social media*, analyzing username squatting on platforms such as X, Instagram, Telegram, and YouTube. They uncovered 349K squatted accounts for 2,625 major brands by mixing brand names with high-risk keywords (e.g., "recovery," "hacked"). This is a similar approach to ours, except that it studies social media accounts instead of search engine advertisements.

In both cases, attackers exploit user trust by closely mimicking brand names or appearances. However, brand impersonation ads, remain understudied.

3.6 Closest Prior Work

A recent study by Musteata et al. (2024) [39] investigates brand impersonation malvertising in Google Search Ads. Their work is the most methodologically similar to ours, using a Playwright-based crawler to search for brands and collect ad metadata and landing pages. However, their detection approach relies on a manually curated list of "known good" domains per brand (e.g., ads for Facebook must land on `facebook.com`). However, their detection approach is brand-specific: for each brand, the crawler has an allowlist of legitimate domains (e.g., ads for Facebook must land on `facebook.com`) and flags any deviations. Flagged ads are then labelled using a custom system that combines pre-defined domain categories with manual review and third-party reputation checks (e.g., VirusTotal, scammer.info). This approach does not account for brands with multiple legitimate domains, and

scaling it to larger, more diverse brand sets would require extensive manual tuning. In contrast, our detection relies on brand-agnostic heuristics which enable generalisable detection of different types of abuse across hundreds of brands without brand-specific rules.

Their results show 513 malicious ads across 10K brand-targeted search ads (5.1%), with 22 unique brandjacking domains identified. They also highlight a notable focus on antivirus brands (e.g., McAfee), which overlap with our own observations. However, their scope is narrower in both brand coverage (50 brands vs. our 605) and dataset size (10K ads vs. our 52K). Our work extends beyond their allowlist-based approach by developing six brand-agnostic heuristics and investigating advertiser identities and enforcement gaps, revealing systemic weaknesses in Google's verification and policy enforcement mechanisms.

Study	**Data Collection**	**Detection Method**	**Scope**
Musteata et al. (2024) [39]	Playwright crawler; 10K ads; 50 brands	Brand-specific allowlist for landing domains; manual review with VirusTotal/scammer.info	Focus on impersonation; 513 flagged ads; 22 unique domains
Our Work	Puppeteer crawler; 52K ads; 605 brands	Six brand-agnostic heuristics; scalable detection; advertiser-level OSINT	Broader scope incl. phishing, scams, affiliate brand bidding; 4,160 flagged ads

Table 3.3: Comparison of our work with Musteata et al. (2024), the closest study on brand impersonation in Google Search Ads.

3.7 Summary of Gaps and Relevance

The works reviewed above demonstrate that malvertising and social engineering are pervasive across the Internet, from low-tier networks and streaming sites to major news outlets. Furthermore, brand impersonation is a prevalent tactic affecting brands on social media platforms and in search engines. Still, several areas remain understudied:

- **Search Ads**: Existing malvertising research focuses either on display ads embedded on web pages, or on only one aspect of ads served through Google Ads (cloaking), while reports from security vendors indicate that specifically malvertising driven by brand impersonation in search engines is a growing threat [44, 30, 45, 6]. Furthermore, scam detection research tends to focus on "high-risk" websites or search keywords (e.g., "crack"), rather than brand-focused queries in a mainstream environment such as Google Search Ads. A notable exception is Musteata et al. (2024) [39], who target search-based brand impersonation using a Playwright crawler and domain allowlisting. However, their method requires brand-specific domain rules and manual labelling steps, limiting scalability and generalisability to broader threat patterns.

- **Broader scope**: While most prior work focuses on classifying or measuring the prevalence of a single scam vector (e.g., clickbait PDFs, counterfeit e-commerce, phishing, cloaking), our research represents a broader survey of the types of scams and tactics found in Google Ads targeting brands. We examine what platform-specific weaknesses exist that allow malicious actors to engage in various forms of abuse, going beyond the classification of a single scam category.
- **Advertiser Identity Abuse**: Prior research has paid limited attention to the role of advertiser identities in enabling or evading abuse. A gap remains in understanding how identity-level abuse facilitates malvertising. Our findings show that access to privileged advertiser accounts is being openly sold, bypassing Google's identity verification systems.

Therefore, while some progress has been made, especially by Musteata et al., a gap remains in systematically investigating brand impersonation malvertising in mainstream search engines. Our work addresses this by collecting ads from real queries on reputable brands and analyzing them for impersonation signals and potential scam indicators.

Chapter 4

Methodology

4.1 Methodology Overview

Our overall process has five main steps, a high-level overview of the pipeline is shown in Figure 4.1. At the core of our methodology, we continuously query Google for a list of brand + keyword combinations over a multi-week period, systematically scraping the ads that appear.

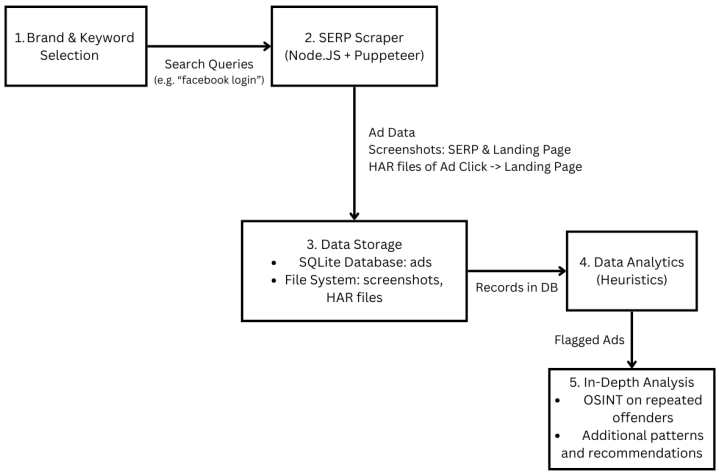

Figure 4.1: Methodology Overview

First, we pick which brands and keywords to search for on Google (See 4.2). Next, we use an automated scraper to collect ads from the Search Engine Results Pages (SERPs) returned for those search terms (See 4.3), saving the results in both a database and local files (See 4.4). Then, after manual inspection, we apply a set of heuristics to identify suspicious ads (See 4.5), and finally, we take a closer look at the advertisers we've flagged and patterns we've found (See 4.6).

4.2 Brand and Keyword Selection

4.2.1 Brand Selection

Our brand list consists of hundreds of brands drawn from a mix of random well-known companies and more targeted selections, chosen to increase our chances of encountering abuse. These include companies from a variety of sectors, software, marketing, fintech, e-commerce, and more. The brands include both established Fortune 500 firms and smaller startups, although the coverage of Fortune 500 companies is low (8% of our brand list). An additional scraping experiment focusing on the complete Fortune 500, was cut short due to scraping restrictions.

We used the following criteria to include brands:

- **Well Known Tech Brands.** We included a selection of well-known tech companies, as they often involve user logins (common bait for phishing or tech support scams) or offer downloadable programs that could be misused to trick users into installing malware.
- **Previously Reported Targets.** Some IP-scanner services and other brands were reported in news articles for being targeted by tech support scams [49]. We included those brands, plus closely related tools or brands, to see whether these attacks extended further.
- **Affiliate Program Presence.** In preliminary scraping runs, we observed ads where affiliates bid on a brand's own name, redirecting users to their affiliate links for that brand, which forces the brand to pay commissions unnecessarily. To better understand this phenomenon, we expanded our list to include brands that have affiliate programs, expecting these brands might also be targeted by this kind of abuse. We found these brands through public affiliate program directories.
- **Shared Advertiser Links.** Once we began finding suspicious ads in our initial runs, we noticed that some advertisers were running ads for multiple unrelated brands. We included some of these additional brands in our monitoring to try and capture those ads live, since the Google Ads Transparency Center does not reveal which search terms were targeted or where the ads actually lead.

Although this approach does not allow us to quantify the overall prevalence, it enabled us to capture actual malvertising scenarios, develop our heuristics and discover insights into malicious advertiser tactics. The full list of brands is provided in Appendix A.

4.2.2 Keyword Selection

To simulate realistic search behavior, we combined each brand in our list with a selection of 40 relevant keywords. These keywords were chosen to reflect common user intents such as logging in, signing up, or searching for support.

In each scraping cycle, we queried Google Search for the brand name alone, as well as the brand name combined with each keyword from our predefined list.

Keywords were selected based on the category of the brand to match typical use cases. For example, a logistics brand would be paired with "tracking" and "delivery", while a tech brand would be paired with "login", "password", and "download". The full mapping of brand categories to keyword sets is provided in Appendix A.1.

This keyword design not only helped us reflect realistic user queries, but also aligned with how malicious advertisers can configure their targeting. Google Ads allows advertisers to target exact keyword matches, and attackers can exploit this by targeting only specific combinations like "Facebook support" rather than just "Facebook". In practice, we found this to be the case for many scams. For example, all of the URL injection scams we captured (across multiple brands) only appeared when the search term included "support". Without using such keyword combinations, these abusive ads would have gone undetected.

Overall, this approach increased the likelihood of uncovering real-world attack scenarios, especially those designed to intercept users with specific intent.

4.3 SERP Scraper

To systematically collect Google Search Ads, we developed an automated scraping system that queries the search engine for each brand-keyword combination and records all sponsored results. This scraper was run continuously over a multi-week period in October and November 2024, collecting a dataset of ads for further analysis.

4.3.1 Scraper Design and Implementation

The scraper was implemented in Node.js using the Puppeteer browser automation library. Its purpose was to replicate realistic user interactions while capturing all relevant data associated with ads and their destinations.

Query Execution. For each brand-keyword pair, we constructed a direct Google Search URL. The scraper automatically denied the cookie consent prompt and then loaded the results as a typical user might see them. After completing a brand's set of queries, we closed the browser context and launched a fresh Puppeteer instance, mitigating the risk of Google detecting continuous automated queries from a single browser. All queries were executed through a Mullvad VPN exit node in Seattle, United States, to maintain consistency and simulate a US-based user.

Ad Detection and Metadata Extraction. Using stable CSS selectors, the scraper identified each sponsored listing on the search results page. It extracted the ad's title and display

URL, then clicked the "About this ad" button to retrieve advertiser identity metadata from the resulting pop-up. Screenshots were saved for both the SERP and the advertiser disclosure window, enabling manual verification of ad content and identity claims.

Landing Page Capture. The scraper followed each ad to its final landing page, recording its HTML, redirection chain, and HAR logs. A screenshot of the landing page was also captured, helping us identify cloaking, affiliate tracking links, and other malicious or deceptive elements.

Scraping Duration. We continuously cycled through each brand-keyword pair over a multi-week period, which was chosen to:

- **Observe recurring ads and temporal variations.** Spanning weekdays, weekends, and different times of day allowed us to potentially discover temporal patterns.
- **Gather sufficient ad variety.** After collecting a large set of ads, we had encountered enough diverse malvertising scenarios to fulfill our main goal of understanding tactics. Beyond that point, additional ads offered diminishing returns for our exploratory focus.
- **Minimize advertiser burden.** Nearly every ad click incurs a cost for the advertiser, so continuously extending the scrape was not justifiable once further ads no longer contributed new insights.

4.4 Data Storage

4.4.1 Database Schema and Storage

All ad-related data is stored in an SQLite database (`ads.db`). The following schema is used:

- **Ads Table:** Each row represents one unique ad (unique combination of `title`, `final_url`, `location`, `advertiser`). The columns include:
  - `title`, `url`, `final_url`, `full_final_url` (retains query parameters).
  - `advertiser`, `location`, `verified`, `advertiser_page` (URL to this advertiser's Ad Transparency Center page).
  - `first_seen`, `last_seen`: timestamps.
- **SearchTerms Table:** Stores each unique search query (e.g., *[brand] login*). Ensures uniqueness with a `UNIQUE` constraint on `term`.
- **AdOccurrences Table:** Defines a many-to-many relationship between `Ads` and `SearchTerms`. Each row records {ad_id, search_term_id, timestamp} keeping track of when a particular ad was found for a particular search term.

This prevents duplication and allows us to measure how often a specific ad occurs. If an ad is encountered multiple times, `last_seen` is updated to the latest timestamp, and a new entry is appended to `AdOccurrences`.

Field	Description	Example
`Ad ID`	Unique identifier for each ad instance in the dataset.	1234
`Title`	Title text displayed in the ad (SERP view).	"Facebook Support 24/7 Help"
`Display URL`	The Display URL shown in the SERP.	`www.facebook.com`
`Final URL`	Actual URL that the user is taken to after clicking the ad.	`https://malvertising.com/landing?src=ad1`
`Location`	Location retrieved from the "About this advertiser" pop-up	"United States"
`Advertiser`	Name retrieved from the "About this advertiser" pop-up.	"Malvertising Inc."
`Verified`	Advertiser verification status retrieved from the "About this advertiser" pop-up.	"Advertiser identity verified by Google"
`Advertiser Page`	Link to the advertiser's Ads Transparency Center page.	`https://adstransparency.google.com/advertiser/...`
`Data RW`	A URL parameter that replays the ad click flow, allowing the same redirection sequence as if a user had just clicked the ad in the SERP.	`https://www.googleadservices.com/pagead/aclk?...`
`First Seen`	Timestamp of when this ad was first observed.	2024-10-10 12:03:21
`Last Seen`	Timestamp of when this ad was most recently observed.	2024-11-02 01:45:03
`Occurrence ID`	Unique identifier for each recorded ad occurrence.	33805
`Search Term`	The brand + keyword query used to scrape this occurrence.	"Facebook support"
`Timestamp`	Timestamp specifically for this occurrence of the ad.	2024-10-20 09:12:57

Table 4.1: Description of the combined dataset fields stored for each ad occurrence.

4.4.2 Local Assets: Screenshots and HAR Files

Screenshots of each SERP, "About this advertiser" pop-up, and final landing page are saved as image files, and HTTP Archive (HAR) logs of each clicked ad are stored in a local directory structure. The screenshots help us keep track of the exact way malvertising appeared,

the HAR logs are useful for manual verification of cloaking.

4.5 Designing Heuristics

We iteratively developed a set of heuristics by combining manual inspection of a subset of collected ads with insights from Google policies and existing security literature. Each heuristic focuses on brand-agnostic, structural indicators of deception or abuse, rather than relying on any single brand's known URL patterns to ensure they are generalizable outside our specific dataset as well. Although not exhaustive, these rules were effective at surfacing a variety of malvertising in our dataset. We performed manual checks on a sample of flagged ads to measure false-positive rates. Quantitative results (e.g. false positives, brand-level stats) are provided in Chapter 5.

4.5.1 Heuristic Discovery and Rationale

Our heuristic design followed an iterative process: we began with domain-focused checks for obviously deceptive or malicious URLs and then, by manually looking at advertisers that were flagged by these rules, and by exploring our dataset manually, we realized that attackers could also exploit legitimate brand domains. Below, we first cover the malicious domain heuristics, then explain how we discovered brand-domain abuses.

Malicious Domains.

- **Domain Mismatch (Display vs. Final).** Google policy requires an ad's Display URL to accurately represent its landing page domain. If the second-level domains do not match, it often indicates cloaking or other deceptive redirects. We flagged such ads for manual review.
- **VirusTotal Domain Check.** We also queried VirusTotal, a widely used antivirus scanning service that aggregates reputation data from multiple security vendors, to check each Final URL's domain for malicious flags. If the domain had more than 5 VirusTotal flags, we flagged that ad. Lower thresholds produced too many false positives. This lookup was done approximately one month after the ads were scraped, results therefore reflect the reputation each domain had acquired after the campaign was live, not real-time detection.
- **TLD in the Second-Level Domain** In early manual checks, we uncovered ads with lookalike domains such as `brand.com.maliciousdomain.com`, which can deceive users by embedding a familiar TLD inside the second-level domain. These ads do not trigger our domain mismatch heuristic if the Display URL is also set to `brand.com.maliciousdomain.com`, and in our dataset, none of these domains were flagged by VirusTotal either. Therefore we introduced a dedicated rule flagging final URLs whose second-level domain contains a known TLD substring.

Through these three checks, we were able to surface malvertising where advertisers are directing users to malicious domains when they click their ads. However, manual inspection of flagged advertisers revealed additional tactics that bypass the need for a malicious domain. Instead, attackers sometimes used the *actual brand's own domain* in deceptive ways. This led us to design the following three heuristics.

Abuse on Legitimate Brand Domains.

- **Phone Number Injection:** Through manual inspection, we noticed ads claiming to offer "support" for major brands. They displayed the brand's official domain in the ad, but came from suspicious advertiser names. Closer investigation revealed that malicious actors embedded a phone number in the Final URL path on the brand's site, exploiting a lack of user input sanitization. This tactic tricks users into calling a fraudulent help line. Our rule flags ads where (1) a phone number appears in the Final URL, and (2) the ad title references "support" or "help". This form of abuse is explained in more detail, with examples in Chapter 5.

- **Affiliate Brand Bidding:** We also discovered that some affiliates were buying ads on brand search terms, using an affiliate link as the Final URL. As discussed in Chapter 2, this tactic intercepts users already intending to visit the official brand site, forcing the brand to pay unearned commissions. We detect these ads by checking (1) whether the Final URL contains known affiliate parameters (see B.1 for the full list), and (2) whether the brand name from the search term also appears in that URL, to confirm that it is an affiliate ad for the specific brand of the search term, and not a benign affiliate ad.

- **Discount Site Brand Bidding:** While examining advertisers already flagged for affiliate brand bidding, we noticed a a more deceptive variant of brand bidding, that was not yet flagged by any of our heuristics, in which affiliates advertise what appears to be a "discount" or "coupon" site for branded queries like *"Nike discount"* or *"Spotify coupon"*. When users click a "get coupon" button on the site, they are redirected to the brand's actual website via an affiliate link. This generates commissions for the affiliate without offering real value to the brand, and there usually is no discount for the user. Our rule flags ads that include discount-related keywords in both the title and URL, unless the advertiser name includes the brand, an indicator that the ad may be legitimate rather than deceptive. The exact keywords and matching logic are listed in Appendix B.2.

Coverage and Stopping Criterion. After defining these six recurring tactics, we chose not to introduce further heuristics. Our manual reviews of flagged advertisers and brands frequently targeted by malvertising in our dataset, no longer revealed any new, consistently reproducible abuse patterns. While we did find a few unflagged cases of malvertising (for instance, a single tech support scam on a previously unknown domain), these ads did not have consistent, generalizable attributes. Making new rules for them would risk overfitting to rare scenarios and make the heuristics less generalizable for a different set of ads.

4.6 Advertiser Analysis

After applying our heuristics to flag abusive ads, we next examined *who* was running these ads. Our advertiser analysis relied on:

- **Advertiser Name**: As listed in the "About this ad" pop-up.
- **Verification Status**: Whether an advertiser was marked "verified" by Google.
- **Location**: The advertiser's registered business address (e.g., Hong Kong).
- **Ads Transparency Center (ATC) Data**: The total number of ad campaigns (beyond our dataset) and types of ads ran by each advertiser.

This enabled us to spot advertisers repeatedly engaging in abuse and see whether certain location or verification patterns existed.

4.6.1 Step 1: Advertiser Grouping

First we sorted the advertisers in our dataset by the number of flagged ads each advertiser had. For each advertiser, we recorded:

- **Total Ads in Our Dataset**: The number of ads by this advertiser.
- **Flagged Ads**: How many of those triggered any of our six heuristics.
- **Location & Verification Status**: As disclosed in the "About this ad" pop-up.
- **Ads Transparency Center Lookup**: For the 10 most frequently flagged advertisers, we visited their profile in the Ads Transparency Center to note the total ad campaigns associated with them (beyond our own dataset) and the types of ads they ran. Note that while the Ads Transparency Center does not allow visiting the landing pages these ads lead to, introducing some uncertainty in categorizing ads, we can still broadly distinguish between different ad types (e.g., search ads for SaaS products vs. video ads for mobile game downloads) to assess variety and scale.

Many smaller-scale advertisers only ran a handful of ads (or at most a few hundred ad campaigns) according to the Ads Transparency Center and typically appeared to focus on one tactic (e.g. brand bidding for a few different brands). By contrast, we found several larger-scale accounts with thousands or even tens of thousands of campaigns with ads for unrelated industries, which motivated us to look into these accounts further.

4.6.2 Step 2: Location- and Verification-Based Patterns

Next, we analyzed the distribution of flagged ads by region and advertiser verification:

- **Regional Analysis**: For each listed business location, we computed how many total ads vs. flagged ads appeared in our dataset. Certain locations such as Vietnam or Gibraltar, had a large fraction of abusive ads (see Section 5.3).
- **Verification Check**: We also measured how many flagged advertisers were "verified" vs. "unverified." As shown in Table 5.6, verified accounts dominated most abuse categories, indicating that identity checks alone are not deterring malvertising.

This regional data, combined with the advertiser-level data from 4.6.1, helped us prioritize certain accounts (those from regions with a high percentage of flagged ads, and a large number of ad campaigns according to the Ads Transparency Center) for further investigation.

4.6.3 Step 3: OSINT Investigation of Frequently Flagged Advertisers

From the previous steps, a clear pattern emerged: several advertisers responsible for many of the flagged ads in our dataset were operating at an unusually large scale. These accounts were often registered in regions like Hong Kong, Vietnam, or Gibraltar. These locations had been independently surfaced as outliers with high proportions of abusive ads. The combination of high ad volume and wide variety in ad types led us to hypothesize that these accounts may be shared or rented by multiple individuals or businesses. To explore this possibility, we did an open-source intelligence (OSINT) investigation of these advertisers, using the following steps:

- **Searching for Official Footprints**: We looked for legitimate websites, LinkedIn profiles, or other public-facing evidence of a corporate presence. In most cases, the only trace we found was a business registration (e.g., in Hong Kong's corporate registry), without any accompanying website or digital presence. This absence is unusual for companies operating in online advertising and raised suspicions about the legitimacy of the entities behind these accounts.
- **Forum and Social Media Queries**: We searched on scam-reporting sites, social networks (e.g., Facebook groups, X/Twitter) and underground marketing forum BlackHatWorld for complaints or references to these advertisers. In some cases, we found explicit posts advertising "verified Google Ads accounts for rent," matching the names we had flagged.
- **Combining Ad Transparency Center and OSINT Findings**: For advertisers showing tens of thousands of ad campaigns across multiple verticals, we compared their ad activity with external user reports and screenshots alleging e-commerce scams, phishing, or other abuses.

From these OSINT efforts, we uncovered multiple indications of an underground market for access to large-scale, verified Google Ads accounts. Evidence ranges from individual forum posts (e.g., on BlackHatWorld) to entire Facebook groups where sellers openly advertise access to the same accounts we flagged as large-scale abusers (see Section 5.3).

Chapter 5

Results

In this chapter, we present the empirical findings that address our overarching research question:

> **RQ:** *How does malvertising driven by brand impersonation manifest in Google Search Ads, and what can be learned from its observable types, prevalence, and advertiser behavior in scraped ad data?*

This question is broken down into three sub-questions:

(Q1) Can we systematically observe brand impersonation malvertising by collecting Google Search ads targeted at a broad selection of brands?

(Q2) What types of malvertising occur in Google Search ads?

(Q3) What patterns characterize malicious advertisers in Google Search?

We organize this results chapter accordingly:

- §5.1 (**Q1**) demonstrates how our scraping approach successfully captured real-world malvertising cases, showing an overview of flagged ads and confirming the feasibility of our data-collection methodology.
- §5.2 (**Q2**) provides the different types of malvertising we discovered, including phishing pages, tech support scams, affiliate brand bidding, and more, with both qualitative examples and quantitative prevalence data on a per-brand level.
- §5.3 (**Q3**) explores the advertiser-centric perspective: how scammers leverage verified vs. non-verified accounts and how large-scale rented/hijacked accounts operate.

We conclude in §5.4 with a summary of the findings.

5.1 Q1: Systematically Observing Malvertising

Our first sub-question asks whether we can *systematically observe brand impersonation malvertising* by scraping Google Search ads targeted at a broad selection of brands. Building on the methodology described in Chapter 4, we continuously cycled through 605 distinct brands over a 24-day period, combining each brand name with a keyword (e.g., "login," "support") to approximate real user queries, collecting a dataset of 52,946 ads in total. Here, we briefly highlight key outcomes confirming that this method does reveal malvertising in the wild.

5.1.1 Heuristic-Based Detection Overview

After storing each ad's metadata (display URL, advertiser identity, final URL, etc.) and landing-page content, we applied six detection heuristics (see §4.5 in the Methodology). Three of these (*Domain Mismatch*, *VirusTotal Malicious Domain*, *Lookalike Domain*) are more general 'red flags" and capture a variety of suspicious or deceptive behaviors (phishing pages, cloaking, tech support scams, etc.). The other three (*Phone Number Injection*, *Affiliate Brand Bidding*, *Discount-Site Brand Bidding*) are flags that detect specific scams we discovered by manually investigating ads.

Heuristic	Flagged Ads	% of 52K	Approx. FPR	No. Advertisers
Affiliate brand bidding	3,781	7.1%	3%	321
Discount-site brand bidding	153	0.29%	8%	45
VirusTotal malicious domain (≥ 5)	84	0.16%	0%	10
Domain mismatch (Display vs. Final)	77	0.15%	6.5%	12
Lookalike domain (TLD in 2nd-level)	47	0.09%	4.3%	4
Phone number in Final URL	18	0.03%	5.6%	5

Table 5.1: Summary of ads flagged by each heuristic, along with approximate false-positive rates (FPR).

Why False Positives Occur. We designed our heuristics to be brand-agnostic and broadly applicable, instead of tuned to every edge case of our dataset. This design choice inevitably leads to a small number of false positives. For example, we observed legitimate ads from a well-known company where the Display and Final URLs were different, something Google seems to allow in special cases, even though it conflicts with its official policy Adding specialized logic for known edge cases would result in an artificially lower FPR here, but would fail to reflect the heuristic's true performance in real-world scenarios with unknown brands. We also checked whether combining multiple heuristics (e.g., requiring that an ad be flagged by both Domain Mismatch and VirusTotal) might reduce false positives. In practice, only 4 ads in the entire dataset were flagged by more than one heuristic, so overlaps were too rare to improve our overall accuracy in a meaningful way.

Determining False-Positive Rates (FPR). For heuristics that flagged a relatively small number of ads (e.g., *Domain Mismatch*, *VirusTotal*, *Lookalike domain* and *Phone Number Injection*), we manually verified all flagged ads. This gave us an exact count of true vs. false positives and thus a precise FPR. It was not feasible to check all flagged ads manually for heuristics that flagged more than 100 ads, (*Affiliate Brand Bidding* and *Discount-site brand bidding*), so for those heuristics we took a random sample of 100 flagged ads and manually checked each, then extrapolated the FPR from how many turned out to be false positives in that sample.

Choosing a VirusTotal Threshold. The *VirusTotal Malicious Domain* heuristic relies on how many security vendors flag a domain as malicious. We encountered some domains that were flagged by multiple VirusTotal vendors, while our inspection revealed no actual malvertising in the specific ads leading to these domains. We decided to consider these instances false positives within our study's focus on malvertising. To find a good trade-off between sensitivity and precision, we tested each number of flags that occurred for the domains in our dataset as a potential threshold (from 1 to 18), measuring false positives at each. Whenever a threshold resulted in too many flagged ads to feasibly check all domains, we took a sample of 50 ads and manually reviewed those domains to estimate the FPR. Table 5.2 shows a high FPR when the threshold is too low, making it unreliable for detection. We ultimately used a threshold of 5, as it had zero false positives in our dataset, and decreasing the threshold by 1 would only introduce false positives without surfacing any additional malvertising cases. Although our chosen threshold is dataset-specific, the heuristic itself generalises well and can be easily re-tuned for other data sources, since the underlying logic (multiple independent detections by security vendors) is simple.

Threshold	Flagged Ads	False Positives	FPR
1	2435	∼48 out of 50 (sample)	96% (est.)
2	697	∼39 out of 50 (sample)	78% (est.)
3	136	52	38.2%
4	88	4	4.5%
5	84	0	0%
9	35	0	0%
10	5	0	0%
16	2	0	0%
18	1	0	0%

Table 5.2: VirusTotal threshold tests: flagged ads and approximate false positives at various thresholds.

5.1.2 Confirming Real-World Malvertising Cases

Overall, we identified **4,000+ flagged ads** (about 7.5% of our dataset). A large share of these were cases of *affiliate brand bidding*. Alongside these, we also found a smaller but significant subset of more traditional malvertising ads (e.g., phishing pages, tech support

scams), which can directly harm users. These instances of abuse are rarer in the overall dataset but can dominate results for certain brands, as we will show later.

So, **Q1** is answered affirmatively: our targeted brand list and automated scraping do capture real examples of malvertising on Google Search Ads. The next section (§5.2) explores the categories of these deceptive ads in more detail.

5.2 Q2: Types of Malvertising

Our second sub-question (**Q2**) addresses the forms of malvertising we found in Google Search Ads, focusing on how they operate at both the ad and landing-page levels. We categorize these types of malvertising into three broad groups:

1. **Phishing Pages**
2. **Tech Support Scams** (including a novel variant where malicious phone numbers are injected on legitimate brand domains)
3. **Affiliate-Based Abuses**: (i) affiliate brand bidding and (ii) discount-site brand bidding

In this section, we describe each category, highlight which heuristics often flagged them, and present screenshots captured by our scraper. We then provide a brand-level analysis showing that even low global flag rates can represent a large fraction of ads for certain brands (§5.2.5). Finally, we share results of a small test demonstrating that scam-injection vulnerabilities extend beyond our dataset, affecting a significant subset of Fortune 500 companies as well (§5.2.5).

5.2.1 Phishing Pages

Phishing pages impersonate a brand's official login or payment page to steal credentials. In our dataset, such pages were surfaced by *Domain Mismatch*, *VirusTotal Malicious Domain*, or *Lookalike Domain* heuristics.

Figures 5.1–5.3 show an example of a phishing ad captured by our scraper that appeared to be made by "DMarket," leading users to a highly convincing but malicious copy of the official DMarket site. This phishing campaign tries to trick users into providing their Steam user credentials, as can be seen in 5.3.

Note in Figure 5.1 that the display URL is the official domain (`dmarket.com`), yet once clicked, our scraper was redirected to a malicious site (`s-games.net`), as shown in Figures 5.2–5.3, causing our *Domain Mismatch* heuristic to flag this ad. By looking at the HAR files for this ad (and others in the same campaign), we confirmed that the advertiser used cloaking: sometimes directing our scraper to the genuine DMarket site, and other times to a malicious clone. This tactic helps avoid detection by Google, as the advertiser attempts to redirect automated visits to the genuine site (matching the display URL), and only real users to the malicious site.

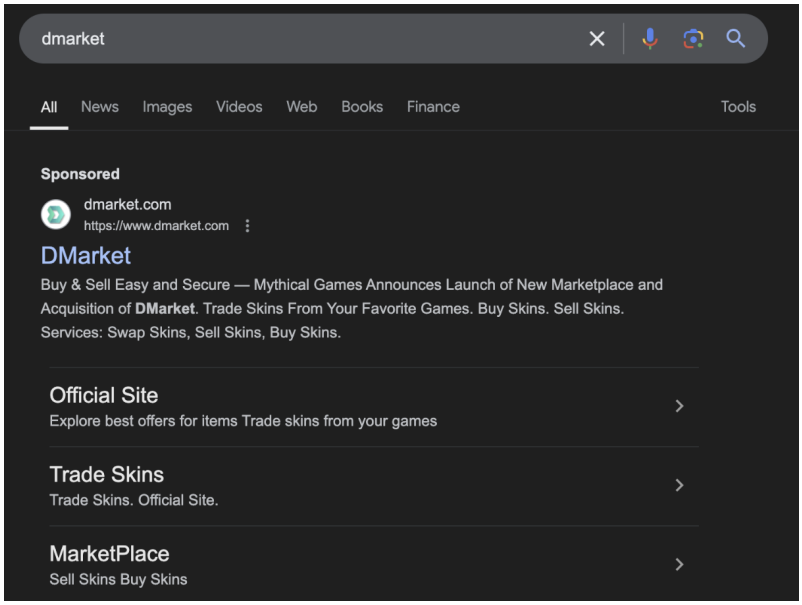

Figure 5.1: Phishing scam ad claiming to be "DMarket," placed above genuine search results.

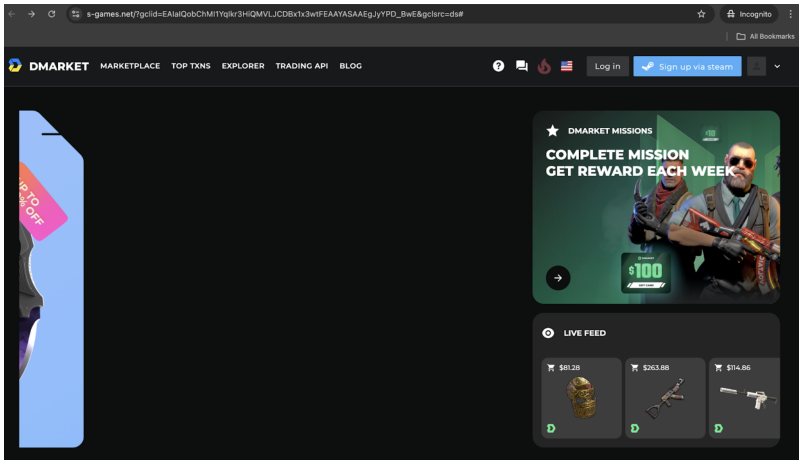

Figure 5.2: Fake landing page closely replicating DMarket's user interface.

5.2.2 Tech Support Scams

Tech support scams pose as official help desks for brands like Microsoft or Meta, typically prompting users to call a phone number. Victims are then pressured to download malware, share credentials, or make payments based on misleading claims. Our general heuristics (Domain Mismatch, VirusTotal, Lookalike Domain), surface typical examples of these scams, as shown in Figures 5.4–5.5.

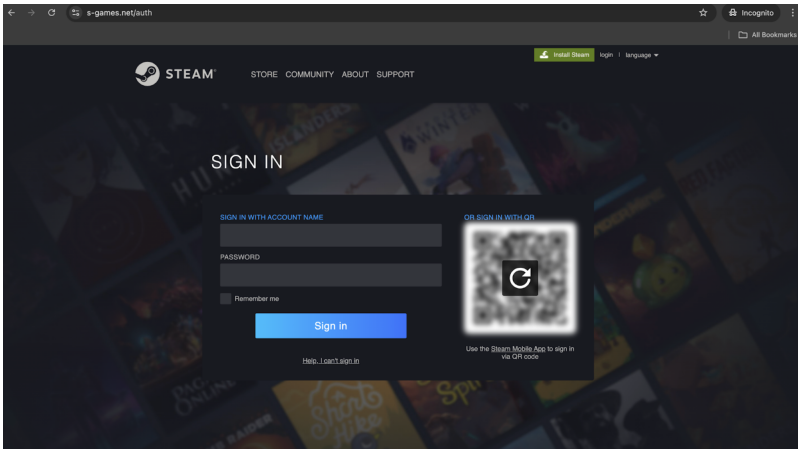

Figure 5.3: A fake Steam login prompt collecting user credentials.

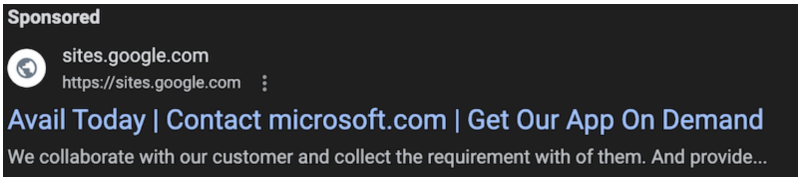

Figure 5.4: A fake "Microsoft Official Support" ad leading to a phishing helpline.

Phone Number Injections

Tech support scams are typically hosted on infrastructure owned by malicious actors. However, a novel trick we noticed, displays a fraudulent phone number on the legitimate brand's own website, deceiving users into believing it is officially provided. This relies on abusing some form of user input. We identified two main variants:

(1) Site-Search or URL Injection. Malicious advertisers abuse a brand's site-search or query parameters to append a phone number in the path (e.g., `facebook.com/search/Call-1-800-FAKE-NUM`). Since the second-level domain is the brand's actual domain, it bypasses Google's Display vs. Final URL check (and hence scammers can freely display e.g. `facebook.com` in the ad). Figures 5.6–5.7 show an example targeting Facebook, where a user clicking the ad is shown a search results page for a query containing the fraudulent phone number. Because the ad convincingly impersonates Facebook and the landing page is indeed on Facebook's domain, users may overlook that it is simply a search results page capable of displaying arbitrary user-submitted content, and fall for the scam.

(2) User Profile Injection. Attackers can exploit any user-generated section of a brand's website, such as profiles, groups, or playlists, by inserting a phone number into the content's title. Similar to the URL injection, the fact that the content resides on the brand's legitimate

Figure 5.5: The landing page urges the user to call a fraudulent phone line.

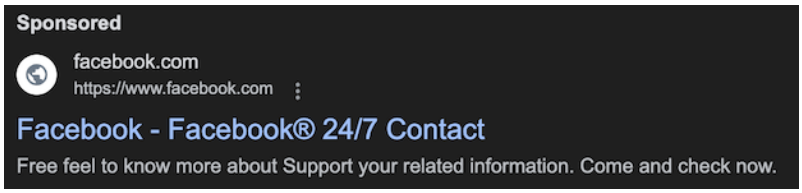

Figure 5.6: Tech support ad exploiting Facebook's search path to display a fraudulent phone number.

domain, combined with the ad impersonating the brand, can make the scam appear authentic. Figures 5.8–5.9 show an example where "Microsoft Official Phone Number", combined with a fraudulent phone number, is used as the profile name, deceiving visitors into believing it is an official support page.

Detection via Phone-Injection Heuristic. These tactics prompted us to create a dedicated *Phone Injection* heuristic (see §4.5), which checks for a phone number (`XXX-XXX-XXXX`) in the Final URL *and* keywords like "support" or "help" in the ad text. We flagged 18 ads (0.03 of our dataset) with a 5.6 false-positive rate. While small in absolute count, for several

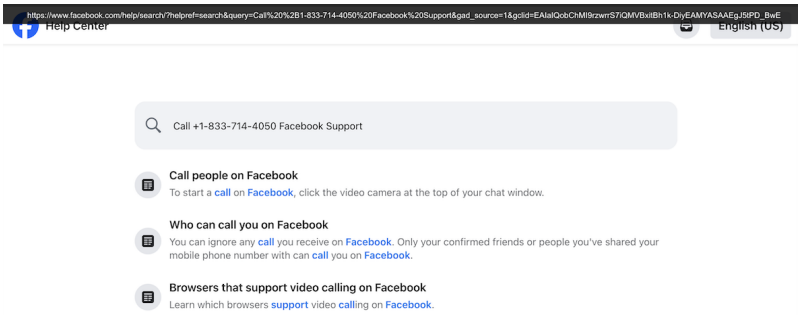

Figure 5.7: Landing page where a fake phone number is "injected" on a search results page on Facebook's official website.

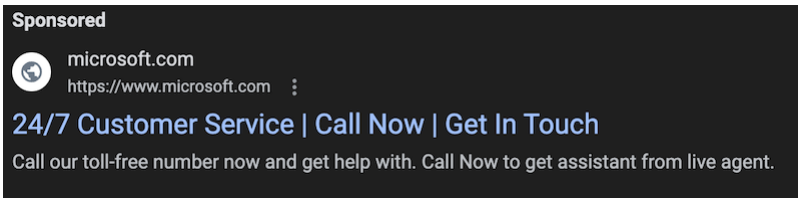

Figure 5.8: Tech support ad exploiting user generated content on Microsoft's website to display a fraudulent phone number.

brands a significant percentage of the ads targeted at their brand name were phone-injection scams, indicating the impact can be serious. Furthermore, in a follow-up test, most Fortune 500 websites we checked are also vulnerable to URL-injection tactics, due to insufficient sanitization of their internal site-search functionality (see §5.2.5).

5.2.3 Affiliate-Based Abuses

In contrast to phishing or tech support scams, the next forms of abuse do not directly harm users by stealing credentials or prompting malicious downloads. Instead, they exploit branded search traffic to claim affiliate commissions from that brand under false pretenses, imposing unfair costs on the brand, and usually violating the affiliate program's terms and conditions by doing so.

Affiliate Brand Bidding

Affiliate brand bidding occurs when affiliates buy ads on a brand's search terms, causing users who already intend to visit the brand's website to go there through the advertiser's affiliate link. This forces the brand to pay commissions for traffic that would otherwise be free. Our dedicated "Affiliate Brand Bidding" heuristic surfaced 3,781 flagged ads (7.1% of the dataset), of which we estimate 2–3% to be false positives based on sampling. These ads affected 189 out of 605 brands (31 %).

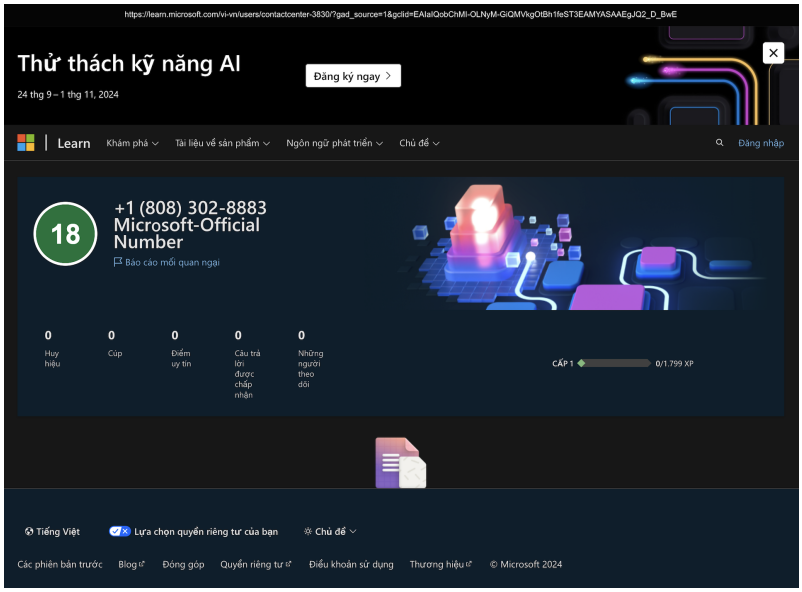

Figure 5.9: "Microsoft Official Phone Number" user profile tricking visitors into calling a fraudulent phone number.

Discount-Code Brand Bidding

Discount-code brand bidding is a similar scam where the ad copy promises coupons or "X% off" deals, and leads the user to a discount site. Usually, when the user clicks on 'Copy Code' or similar buttons, they are simply redirected through an affiliate link to the normal brand site, receiving no actual discount. The heuristic detecting these flagged 153 ads (0.29% overall) with a higher estimated false-positive rate (8%), as some genuine discount sites were wrongly flagged.

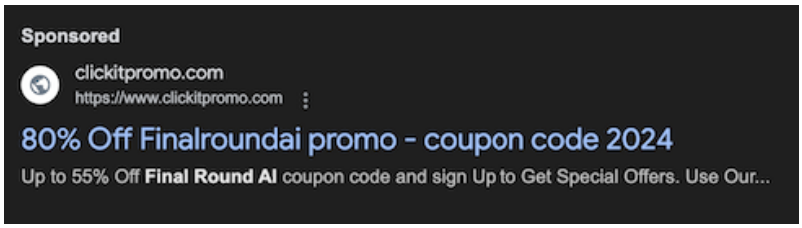

Figure 5.10: A discount-site brand-bidding ad placed above the real brand's result.

5.2.4 How Many Ads Each Category Includes

We put together Table 5.3 to show how many ads belong to each malvertising category we described (phishing, tech support, and affiliate-based scams), and which of our heuristics

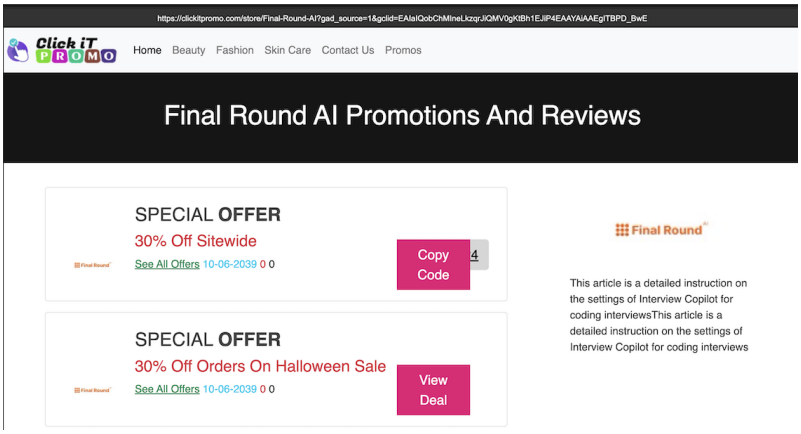

Figure 5.11: The landing page claims to offer coupons but simply redirects to the brand's website through an affiliate link.

flagged them. For some ads, it was not possible to determine the exact scam type, due to cloaking or unavailability of the landing page.

Category	# Ads (%)	Heuristics (Flag Counts)
Affiliate Brand Bidding	3781 (7.1%)	Affiliate Brand Bidding: 3781
Discount-Site Brand Bidding	153 (0.29%)	Discount-Site Brand Bidding: 153
Phishing Pages	135 (0.25%)	Domain Mismatch: 47 VirusTotal: 84 Lookalike Domain: 4
Phone Injection	18 (0.03%)	Phone Number Injection: 18
Tech Support Scams	14 (0.03%)	Domain Mismatch: 2 Lookalike Domain: 12

Table 5.3: Number of ads in each malvertising category (with % of the total 52,946 ads), plus how many times each relevant heuristic flagged those ads.

5.2.5 Per-Brand Prevalence and Impact

Although some categories (e.g., phone injection) represent a small fraction of all ads, they can be a significant percentage of the ads shown to users searching for specific brands. Table 5.4 shows examples where a high percentage of ads targeted at a brand's name were flagged. In Appendix C we include an extended version of this table including brands with a lower percentage of flagged ads.

These brand-level insights highlight why a seemingly small fraction of total ads can still pose a serious threat to a given company.

Heuristic	Brand	Ads Flagged	Total Ads	% Flagged
Domain Mismatch (Display vs. Final URL)				
Domain Mismatch	Photovibrance	16	57	28.1%
Domain Mismatch	DMarket	49	546	9.0%
Phone Number Injection				
Phone Injection	Facebook	4	39	10.3%
Phone Injection	Meta	9	132	6.8%
Phone Injection	Process Explorer	1	20	5.0%
VirusTotal Malicious Domain (Threshold ≥ 5)				
VirusTotal	McAfee	79	774	10.2%
Lookalike Domain (TLD in 2^{nd}-level)				
Lookalike Domain	Advanced Port Scanner	23	109	21.1%
Lookalike Domain	IP Scanner	17	206	8.3%
Affiliate Brand Bidding				
Affiliate	BodyGraph Chart	83	88	94.3%
Affiliate	AppSumo	77	84	91.7%
Affiliate	EverBee	92	105	87.6%
Affiliate	MaxAI.me	83	130	63.8%
Affiliate	Fireflies.ai	76	133	57.1%
Discount-Site Brand Bidding				
Discount	Stealthwriter	29	149	19.5%
Discount	Kikkerland	16	87	18.4%
Discount	Madaz	9	74	12.2%
Discount	Final Round AI	12	172	7.0%
Discount	Cloudways	10	153	6.5%

Table 5.4: Brands whose flagged percentage exceeded 5% under each heuristic, showing the number of flagged ads, total ads collected for this brand, and the resulting percentage.

Fortune 500 Phone Number Injection Test

To see how widely the phone number injection attack we observed could affect major brands, we tested the top 50 US Fortune 500 websites, checking their internal site-search features for vulnerability to phone-number injection. We found that 86% (43 out of 50) of these websites allowed a phone number to be injected through a URL-based search query, leaving them open to the same malvertising tactic we captured in our dataset. Although our main dataset only captured active attacks targeting Meta, Microsoft, and their sub-brands, these findings reveal that many large companies share the same underlying flaw.

5.3 Q3: Patterns in Malicious Advertisers

The third sub-question (**Q3**) asks what *advertiser behavior* underlies these threats. In other words, *who* is running these ad campaigns and what techniques let them remain active?

5.3.1 Step 1: Grouping by Advertiser

To identify which advertisers were most active in running malicious ads, we sorted our dataset by the number of flagged ads per advertiser. Many of these advertisers ran malicious ads targeted at dozens of different brands. For the top 10 advertisers, we also retrieved their total number of ad campaigns according to Google's Ads Transparency Center (ATC), and noted any observed diversity in the types of ads they ran (e.g., SaaS search ads vs. mobile app promotions). In practice this meant that the Hong Kong advertisers with over 20K ads, shown in Table 5.5, all ran a wide variety of ad types, while the others had a narrower range.

Advertiser	Location	Total Ads	Flagged Ads	ATC Campaigns
Xun Meng International Limited	Hong Kong	479	471	∼20K
Bringme Consultant Services Ltd.	Hong Kong	263	263	∼30K
Traffic Heroes Ltd.	Gibraltar	185	173	∼76
HERA SEO SRL	Romania	107	107	∼200
Saba Fatima	Pakistan	98	98	∼300
Simply LLC	United States	85	84	∼500
BlueVision Interactive Limited	Hong Kong	107	82	∼600K
AFFI	France	79	79	∼55
IMOBSESSION TECH LLP	India	78	78	∼41
Sky Cosmo Limited	Hong Kong	90	74	∼21

Table 5.5: Top 10 advertisers with the highest number of flagged ads in our dataset, including their known campaign volume on Google Ads Transparency Center (ATC).

Note that in some cases, the number of ad instances we observed in our dataset exceeds the number of ad campaigns shown in the Ads Transparency Center (ATC). This can be attributed to two main factors: (1) multiple ads with slight variations may be grouped under a single campaign in the ATC, and (2) we observed several cases where even currently active ads were not listed at all, suggesting that the ATC's reporting may be incomplete or unreliable. As a result, these campaign counts should be treated as approximate.

5.3.2 Step 2: Regional and Verification Patterns

Verified vs. Non-Verified Accounts. As Chapter 2 noted, Google Ads requires certain advertisers to verify their identity. Table 5.6 provides an overview of how many *flagged* advertisers were verified vs. unverified, by heuristic:

Although non-verified accounts appear, verified advertisers dominate most abuse categories, showing that identity verification alone is not enough to prevent malvertising.

Regional Patterns. We also observed significant differences in the prevalence of abusive ads by the advertiser's registered location. Table 5.7 shows the top five such locations, based on the number of flagged ads originating from advertisers who listed these regions as their business address. Percentage wise, mainly Hong Kong, Vietnam and Gibraltar stand out.

Heuristic	Flagged Advertisers	Verified	% Verified
Affiliate brand bidding	324	231	71.3%
Discount brand bidding	45	42	93.3%
Domain mismatch	12	11	91.7%
VirusTotal malicious	10	10	100%
Phone number injection	7	6	85.7%
Lookalike domain	4	4	100%

Table 5.6: Flagged advertisers (by each heuristic) and their verification status.

Location	Flagged Ads	Total Ads	% Flagged	% Unverified
Hong Kong	1161	1982	58.6%	3.4%
United States	794	34113	2.3%	1.5%
Vietnam	575	701	82.0%	33.4%
India	310	2271	13.70%	1.4%
Gibraltar	168	224	75.0%	0%

Table 5.7: Top 5 locations with the highest number of flagged ads for any heuristic.

Hong Kong. Hong Kong advertisers accounted for 1,161 flagged ads out of 1,982 total (58.6%). Looking at the advertisers listed in Table 5.5 revealed that the most frequently flagged Hong Kong advertisers were all verified and had collectively run approximately 650K ads on Google. In our dataset, they were responsible for hundreds of affiliate brand bidding ads, with some also flagged for other forms of abuse (e.g., tech support scams). The scale of these verified accounts and the variety of ads they ran prompted a more extensive OSINT investigation (§5.3.3).

Vietnam. With 575 flagged ads out of 701 total (82.0%), Vietnam showed a particularly high ratio of abusive ads. We noted a large number of unverified advertisers here, many engaging in affiliate brand bidding.

Gibraltar. In Gibraltar, 75% (168 of 224) of ads were flagged. Notably, 96% of all ads from Gibraltar originated from a single account, *Traffic Heroes Ltd.*, which had 80% of its ads flagged for affiliate brand bidding. This one account skewed the region's overall ratio of abusive ads, and led us to investigate it as part of our OSINT investigation.

5.3.3 Step 3: OSINT Investigation of Frequently Flagged Advertisers

As noted in the methodology (§4.6.3), several advertisers in our dataset had unusual characteristics: high flag counts, registration in regions associated with a large proportion of abuse (e.g., Hong Kong, Gibraltar), and thousands of ad campaigns listed in the Ads Transparency Center, with a wide variety of ad types. These patterns raised the possibility that the accounts were being shared or rented by multiple actors, and were the reason for a more extensive OSINT investigation into these advertisers.

Lack of Official Web Presence. For the most frequently flagged advertisers, we attempted to find any official web presence, such as a corporate website or verified LinkedIn profile. In most cases, however, the only available trace was a company registration in a regional business registry (e.g., Hong Kong's Companies Register). This is unusual for companies supposedly operating at scale in online advertising, and raised strong suspicions about the legitimacy of the underlying entities.

Abuse Reports and Public Complaints. Another signal we tracked was the presence of abuse reports on scam-reporting platforms and user complaint forums. We found such reports for multiple advertisers in our top-flagged group, notably *Xun Meng International Limited*, *Bringme Consultant Services Ltd.*, and *Traffic Heroes Ltd.*, describing a wide variety of scams advertised by these accounts. Detailed examples, screenshots, and links to these findings are provided in Appendix D.

Evidence of Account Rental. We also examined external sources for signs of account sharing or abuse. On the underground marketing forum BlackHatWorld, we found posts by users describing the sale or rental of Google Ads accounts, including references to *Traffic Heroes Ltd.* as a buyer or intermediary. In addition, we discovered Vietnamese-language Facebook groups with thousands of members where access to advertiser accounts was openly advertised, some of which matched the names of advertisers we flagged for large-scale abuse, including *Xun Meng International Limited*, *Bringme Consultant Services Limited*, and *Sky Cosmo Limited*. Other accounts that were advertised in these groups were also associated with ads in our dataset that were flagged, often for affiliate brand bidding. These account offers typically involved payment in cryptocurrency, after which the seller would fund the account and charge a service fee.

Note that both *Traffic Heroes Ltd.* and *BlueVision Interactive Ltd.* maintain a legitimate web presence offering online advertising services. Malicious actors may be exploiting their infrastructure for abusive ad campaigns without their knowledge. Note that also for the advertisers without a public web presence, they may not knowingly be complicit to the rental of these accounts, but could be hijacked instead.

The availability of verified accounts for rent or resale makes it trivially easy for malicious actors to bypass Google's identity verification requirements, undermining the intended safeguards of the platform.

Advertiser	Location	Web Presence	Variety	Abuse	Rental
Xun Meng Intl. Ltd.	HK	False	True	True	True
Bringme Consultant Services Ltd.	HK	False	True	True	True
Traffic Heroes Ltd.	Gibraltar	True	False	True	True
BlueVision Interactive Ltd.	HK	True	True	True	False
Sky Cosmo Limited	HK	False	True	False	True

Table 5.8: Advertisers included in OSINT investigation, and whether specific signals were present: legitimate web presence, diverse ad types, abuse complaints, and traces of rental activity.

5.4 Summary of Findings

Bringing together the above sections:

- **(Q1) Systematic Observation:** Our targeted scraping of 605 brands over 24 days collected a total of 52,946 Google Search ads. Using six heuristics, we flagged 4160 of these ads (7.9%), including hundreds of phishing and scam ads, verified through manual inspection. We confirm that this approach does surface real-world malvertising on Google Search Ads.
- **(Q2) Malvertising Characteristics:** We discovered multiple scam categories: *phishing pages*, *tech support scams* (including a novel delivery mechanism of content injection on the actual brand's domain), *affiliate brand bidding*, and *discount site brand bidding*. These methods often exploit trust in the impersonated brand and in the legitimacy of Google search results to mislead users. Some rely on domain cloaking or injection of content on the brand's own domain to avoid detection.
  - **3,781 ads** (7.1%) were flagged as *affiliate brand bidding*, where affiliates hijack branded search traffic to claim illegitimate commissions.
  - **153 ads** (0.29%) were flagged as *discount-site brand bidding*, promising coupons or deals that ultimately redirect users through affiliate links.
  - **135 ads** (0.25%) were identified as *phishing pages*, mimicking login flows to steal user credentials.
  - **14 ads** (0.03%) were *traditional tech support scams*, directing users to fake helplines.
  - **18 ads** (0.03%) involved a *novel phone-number injection technique*, where malicious actors display fake phone numbers on the brand's legitimate website through URL or content injection, bypassing Google's Display vs. Final URL check.

While these categories vary in overall prevalence, some dominated the ad space for specific brands. Additionally, a manual test showed that out of a sample of 50 Fortune 500 brands, 86% were vulnerable to the same phone-number injection tactics we observed in practice.

- **(Q3) Advertiser Patterns:** We identified a small set of verified advertisers responsible for large volumes of flagged ads, often with tens of thousands of ad campaigns in the Ads Transparency Center and diverse targeting across verticals. These high-flag accounts were frequently registered in regions that stood out for abuse, such as Hong Kong or Gibraltar. OSINT revealed limited legitimate web presence for several of these advertisers, alongside public abuse complaints and references in forums. We discovered Vietnamese Facebook groups offering access to the same (Hong Kong-based) accounts for rent. In contrast, a separate cluster of smaller, often unverified advertisers, particularly from Vietnam, primarily engaged in affiliate brand bidding. These findings suggest that malicious actors are not only exploiting Google's ad systems directly, but also circumventing identity checks through shared or rented accounts.

In summary, our findings confirm that a wide variety of brand-impersonation malvertising is actively being exploited by malicious advertisers, targeting unsuspecting users who explicitly search for a trusted brand or service. Some of these types of abuse occur on a large scale, and advertisers can easily circumvent identity verification measures from Google. The next chapter discusses these results in a broader context, proposes recommendations for both Google and impacted brands, and outlines opportunities for future work.

Chapter 6

Conclusion and Discussion

The ad ecosystem is a key component of the modern Internet, used to fund free content and services for users worldwide. This growing industry, due to its size, has naturally also become a target for malicious actors. Cybercriminals use malvertising, the spread of malware or other malicious content through ads as a way to deliver various types of schemes. While academic research has measured the threat of malvertising on publisher sites in the past, and developed detection techniques for those threats, the growing threat of malvertising in search engines is left largely unexplored.

Several weaknesses in Google's enforcement techniques and policies, leave their search engine vulnerable for deceptive search engine ads impersonating brands that often harm users.

This thesis addresses this gap by answering the research question:

> **How does malvertising driven by brand impersonation manifest in Google Search Ads, and what can be learned from its observable types, prevalence, and advertiser behavior in scraped ad data?**

This chapter provides the main findings of our research, answering the research questions. We also reflect on these results, placing them in a broader context, provide recommendations, and highlighting opportunities for future work. We also discuss some of the limitations of our research.

6.1 Main Findings

We collected over 52K Google Search ads, shown for search terms consisting of brand + keyword combinations like "Facebook support." Using this dataset, we developed six heuristics that help us flag malicious ads.

These heuristics surfaced a wide variety of malvertising, notably: phishing, tech support scams, affiliate brand bidding, and discount site brand bidding.

We discovered a novel way of delivering tech support scams, where attackers exploit insufficient sanitization of user input on internal site-search pages, and can inject fraudulent

tech support phone numbers on the actual brand's domain. We found that 86% of Fortune 500 brands we checked were vulnerable to this exact exploit.

In total, our heuristics flagged over 4,000 ads, or roughly 7.5% of the 52K collected. The most common category was affiliate brand bidding (3,781 ads), followed by discount-site brand bidding (153), phishing pages (135), phone-number injection cases (18), and tech support scams (14).

Most malvertising comes from advertisers that have verified their identity with Google. However, in comparison with other forms of abuse, affiliate brand bidding ads are relatively more likely to come from unverified advertisers, many of them based in Vietnam.

We discovered a notable subset of verified advertisers from Hong Kong, that are often flagged for abuse in our dataset, and run tens of thousands of ad campaigns according to the Ads Transparency Center. The large scale of these accounts, combined with the wide variety of ads they ran (according to the ATC) prompted an OSINT investigation, leading us to a Vietnamese Facebook community that functioned as an underground market for Google Ads accounts, where access to these exact advertiser accounts, and many others, was being advertised.

6.2 Conclusion

Based on these findings, we can answer the sub-questions of our overarching research question:

Q1: Can we systematically observe brand impersonation malvertising by collecting Google Search ads targeted at a broad selection of brands?

Yes. By scraping over 52K Google Search ads over a 24-day period using brand + keyword combinations (e.g., "Facebook support"), we were able to systematically surface hundreds of verifiably malicious ads impersonating brands. These included phishing attempts, fake tech support pages, affiliate brand bidding, and deceptive discount sites. The results confirm that brand-targeted scraping is an effective method for surfacing real-world malvertising in Google Search Ads.

Q2: What types of malvertising occur in Google Search ads?

We observed several distinct scam categories, including phishing, tech support scams, affiliate brand bidding, and discount site brand bidding. Some attacks used novel techniques such as injecting fake phone numbers into the actual brand's domain through unsanitised internal search functionality. Others relied on cloaking, redirects, or domain similarity. What they share is the exploitation of user trust in both the impersonated brand and the perceived legitimacy of Google's advertising platform. These threats operate in ways that make them difficult for users to detect and for automated systems to block reliably.

Q3: What patterns characterize malicious advertisers in Google Search?

Most advertisers flagged for malicious ads had verified their identity with Google, and several regions had a relatively large percentage of their ads in our dataset flagged (Hong Kong, Vietnam, Gibraltar). Of these, Vietnamese advertisers had a comparatively large share of unverified advertisers, who were mostly running affiliate brand bidding scams.

A small set of verified advertisers stood out due to the large number of flagged ads linked to them. These advertisers often operated tens of thousands of ad campaigns across diverse industries and were frequently registered in regions such as Hong Kong or Gibraltar—regions that were disproportionately responsible for flagged ads in our dataset.

OSINT investigations found little legitimate web presence for several of these advertisers, but did find public complaints and references to them on underground forums. We also discovered Vietnamese-language Facebook groups advertising access to these same (Hong Kong-based) accounts.

This indicates two parallel abuse strategies: one where access to large-scale, verified accounts is being shared to hide behind the scale of these advertisers, and another relying on unverified throwaway accounts. Both highlight weaknesses in Google's identity verification and enforcement mechanisms.

6.3 Implications of Our Findings

6.3.1 Implications for User Trust and Safety

Malicious ads observed in this study directly target users with scams such as phishing and fake tech support. These threats are often designed to appear indistinguishable from legitimate ads, exploiting both the reputation of trusted brands and the perceived reliability of Google Search. This deception is made even more effective by the visual similarity between search engine ads and organic search results, which makes it difficult for users to tell them apart.

As a consequence, users can be directly harmed, financially, through phishing scams or fraudulent payments, or through the exposure of personal and account information. These incidents may not only result in monetary loss, but also lead to compromised online accounts or other privacy risks.

As a result, users may develop a reduced level of trust in search engine ads overall, especially when these experiences lead to financial harm or data exposure.

While malicious ads that directly harm users are a small percentage of the total ads collected, considering the scale of Google Search these risks are still significant.

6.3.2 Implications for Brands

Implications of User Harming Attacks for Brands

The forms of abuse that directly harm users also have a negative impact on brands, as users may not always realise they have been phished or scammed, and may associate the negative experience with the impersonated brand. One brand that appeared frequently in our dataset,

DMarket, an in-game item marketplace. As shown in 5.2.5, DMarket was the target of a relatively high number of flagged ads, often leading to phishing pages. DMarket has been the subject of user accusations of malicious behavior in the past [46, 47]. While it is unclear whether these reports stem from phishing specifically via malvertising, the overlap between this brand being targeted by brand impersonation ads that lead to phishing pages, and the existence of these scam complaints, is noteworthy.

Brands whose websites are exploited in tech support scams leveraging the Phone Number Injections (see 5.2.2) also face reputational risks. Although the scam content is injected without their intent, users encountering fake tech support information on a legitimate domain may hold the brand responsible.

Another way brands are at risk is when Google Ads accounts bearing their name, or belonging to them directly, are used in malvertising ad campaigns, whether through hijacking or intentional misuse. In one case we observed, a verified Google Ads account associated with a legitimate travel agency was used to run both URL injection-based tech support scams (targeting Facebook) and a phishing ad (targeting DMarket). While we cannot conclusively determine whether the account was compromised or misused with intent, the pattern is consistent with previously reported abuse of hijacked advertiser infrastructure [11, 45]. The affected brand received negative Google Business reviews, illustrating the reputational harm that can result when malvertising ad campaigns are ran through trusted advertiser accounts.

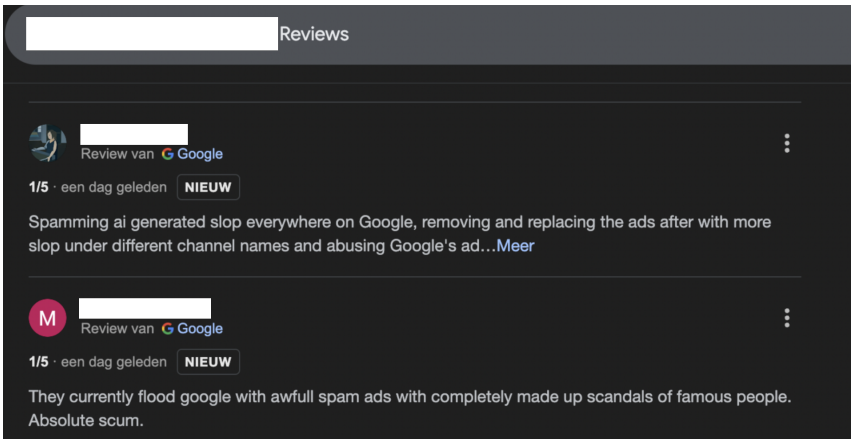

Figure 6.1: Google Reviews for the hijacked travel agency (names redacted)

A similar case involves BlueVision, a verified advertiser that appears to be an established Chinese advertising company. In our dataset, this account was observed running many flagged ads that led to phishing pages, fake discount sites, and affiliate links (brand bidding). While it seems unlikely that the company itself is directly responsible for placing these ads, it is more plausible that malicious actors are leveraging its infrastructure. Because this activity occurs at large scale under the guise of BlueVision's advertiser identity, the brand has been named in public scam accusations online (see D.4). Such associations contribute to reputational harm.

Implications of Affiliate Brand Bidding for Brands

Affiliate brand bidding causes financial harm to brands by forcing them to pay commissions for traffic they would otherwise receive organically. It also raises the cost of branded keywords through competition, increasing advertising expenses. While users are not directly harmed, the practice undermines trust and confuses attribution, especially when affiliate ads mimic the brand's own messaging.

Legitimate affiliates may lose their fairly earned commissions if the affiliate program uses a last-click attribution model, and a user visits a brand-bidding affiliate's link after initially clicking on theirs.

To assess whether this abuse persists over time, we conducted a manual check on April 10th, 2025, months after our scraping experiment. For 7 out of the 10 most targeted brands in our dataset, we immediately observed a brand bidding ad when searching for the brand name on Google. This shows that affiliate brand bidding is not only widespread but also persistent.

Separately, we examined whether these brands prohibit affiliate brand bidding in their terms and conditions. 7 out of the same 10 brands explicitly forbid the practice. The fact that these brands are among the most heavily targeted despite clearly prohibiting this in their terms indicates that such terms alone are not an effective defense mechanism without active monitoring and enforcement.

The results of this small experiment are shown in 6.1.

Brand	Ongoing Brand Bidding	Forbids Brand Bidding in T&C
EverBee	True	True
BodyGraphChart	True	False
MaxAI	False	True
AppSumo	True	True
Fireflies	False	False
SaneBox	False	True
Landerbolt	True	True
ThriveCart	True	True
VideoTap	True	False
GetMunch	True	True

Table 6.1: Top 10 most targeted brands: manual check for ongoing brand bidding ads and whether affiliate terms forbid brand bidding

6.3.3 Implications for Platform Integrity (Google)

Ineffective Identity Verification. Our findings show that Google's enforcement mechanisms are being systematically circumvented. The identity verification process is undermined by an underground market of access to privileged accounts. Advertisers with long-standing, verified accounts appear able to run a wide variety of ad campaigns with little oversight, even when abuse has been repeatedly documented.

The ability to operate at scale through rented or shared advertiser accounts introduces long-tail risks that go beyond brand impersonation. These accounts could be used to launder money (e.g., purchase ad credit with illicit crypto, run a product ad, and collect clean cash), promote political messaging without attribution (as observed in one political ad attributed to "Xun Meng Intl Ltd", see D.1), or carry out other ad campaigns with no accountability. Without reliable traceability of who controls these accounts and how funds flow through them, such abuse represents a serious challenge to Google's ability to govern its advertising platform effectively and maintain accountability at scale.

We also saw some ads using cloaking tricks that shouldn't work on new Google Ads accounts. These techniques rely on redirects that are normally blocked. This suggests that some of the abuse may be happening through what are called "grandfathered" accounts, older ad accounts that were created before Google introduced stricter rules. We found discussions on forums like BlackHatWorld where people talk about such accounts [20].

Limitations of the Ads Transparency Center. Additionally, the Ads Transparency Center (ATC), which could serve as a valuable tool for investigation and public accountability, falls short in practice, especially for large-scale advertisers. While the ATC aims to increase transparency, in its current state it is not usable for systematic investigation:

- For large advertisers, the interface becomes practically unusable. Ads cannot be searched by keyword or filtered by brand, and the only way to browse is by endless scrolling.
- During our investigation, we observed that scrolling through a large advertiser's history often triggers glitches, causing the interface to reset to the top, making it impossible to perform a full review of their ad activity.
- Advertisers that have been removed by Google (due to policy violations) also disappear from the ATC entirely. This harms transparency. Even if removal is justified, the fact that their advertising history is removed as well undermines accountability.

Failures of the Domain Comparison Check. Another implication of our findings for Google is the limitation of the existing check that compares the domain of the Display URL to the domain of the Final URL. While this mechanism is intended to detect misleading ads, our research shows that it fails to prevent several forms of abuse. First, malicious advertisers can circumvent the check through cloaking, particularly when using "grandfathered" accounts. Second, in cases of affiliate brand bidding, the Display and Final URLs are identical, pointing to the brand's legitimate domain, making the abuse indistinguishable under this rule. Third, the check is ineffective against URL injection attacks, or all similar attacks where malicious content is injected into pages hosted on the legitimate domain. Such abuse entirely bypasses this safeguard. These findings suggest that the Display vs Final URL domain comparison offers limited protection and should not be relied on as a primary enforcement mechanism.

Lack of Enforcement for Affiliate Brand Bidding. Furthermore, affiliate brand bidding abuse appears to be largely unpoliced by Google. While the majority of detected cases involved verified advertisers, this category included a significantly higher proportion of unverified accounts than other forms of malvertising, suggesting weaker scrutiny or enforcement in practice.

Google's advertising policies do not explicitly forbid affiliate brand bidding. As a result, even when affiliates advertise on branded search terms and redirect users to the brand's own website, this behavior does not trigger automatic enforcement mechanisms, despite violating most affiliate program rules. This further supports our finding that affiliate brand bidding is not actively policed by the platform.

Finally, the tools Google provides to report trademark abuse, such as the ad reporting forms, are poorly suited to affiliate brand bidding. Reports must be submitted manually, are limited to individual ads, and automatic enforcement mechanisms only work when repeat offenses are linked to the same advertiser and domain. This approach does not work in affiliate brand bidding, where the destination domain is legitimate (e.g., nike.com), and many different advertisers participate in the abuse.

6.3.4 Implications for the Ad Ecosystem and Regulation

The misuse of verified advertiser accounts weakens the ad ecosystem by undermining the assumption that a verified ad implies safety or legitimacy. When verified advertisers are repeatedly involved in abuse, the integrity of the whole verification process is called into question.

While our dataset focused on Google ads shown to U.S. users and targeting primarily U.S.-based brands, a significant portion of the abuse originates from international actors. We observed that many abusive brand bidding ad campaigns were run by unverified advertisers based in Vietnam. Furthermore, in a Vietnamese-speaking Facebook community with thousands of members access to Hong Kong based, verified Google Ads accounts is openly being advertised. Some of these accounts have a long track record of abuse, raising questions about why Google continues to allow them to operate. Instructional video content tied to the usage of these accounts has accumulated thousands of views on YouTube [19], suggesting an organized ecosystem that supports these activities.

This points to a deeper problem: as abuse becomes more international and coordinated, enforcement systems built around individual accounts or local accountability start to fail. For users or brands that fall victim to these attackers, attribution becomes difficult, compensation or recourse is nearly impossible, and reporting fraud across borders is highly impractical. These patterns are consistent with a broader trend of digital fraud operations emerging from Southeast Asia [51].

6.4 Recommendations

6.4.1 Recommendations for Users

Given the current limitations in detecting and blocking deceptive ads, especially those that closely mimic legitimate brand ads, users should consider using an ad blocker as a protective measure.

In many cases, malicious ads appear indistinguishable from legitimate ones up to the display URL. In some instances, such as URL injection attacks, even the final landing page is hosted on the brand's real domain. This makes it extremely difficult for users to determine whether an ad is trustworthy based on visual cues alone.

Until platforms improve detection and enforcement mechanisms, ad blockers remain one of the most effective defenses against malicious ads in search results from a user standpoint.

6.4.2 Recommendations for Brands

Brands affected by impersonation, affiliate abuse, or content injection vulnerabilities can take several steps to reduce exposure and improve resilience:

- **Monitor brand-targeted ads.** Brands should monitor the ads shown on search results for their name and related keywords, either manually or by using a dedicated brand protection service. Malicious ads can be reported to Google through the trademark complaint system [13], although this remains a manual and ad-specific process.
- **Encourage internal use of ad blockers.** Employees, especially those with privileged access, should use ad blockers to reduce the risk of encountering phishing ads that impersonate internal services or support channels.
- **Secure publicly exposed features that reflect user input.** Brands should review and secure publicly accessible components of their websites, such as internal site-search tools, user-generated content, or profile pages, to prevent abuse. Features that reflect user input back into the interface must sanitise this input to avoid abuse. For example, attackers can exploit unsanitised search queries or profile fields to insert scam content (e.g., fake support numbers) directly into a legitimate looking brand page. While our research observed this abuse via Google Ads, the same vulnerability can be leveraged through other channels that bring users to such a page, like phishing emails or social media links. As shown in 6.2, Facebook mitigated this risk by stripping phone numbers from search result pages.
- **Protect against affiliate brand bidding abuse.** Brands using affiliate platforms should enforce policies that prohibit brand bidding and educate affiliates accordingly. Tools like Rewardful suggest clearing the affiliate tracking cookie when traffic is identified as originating from Google Ads by checking for specific URL parameters [9].

However, this approach is limited: Google allows advertisers to disable this parameter [1], and sophisticated affiliates can bypass such detection by redirecting clicks through intermediate domains.

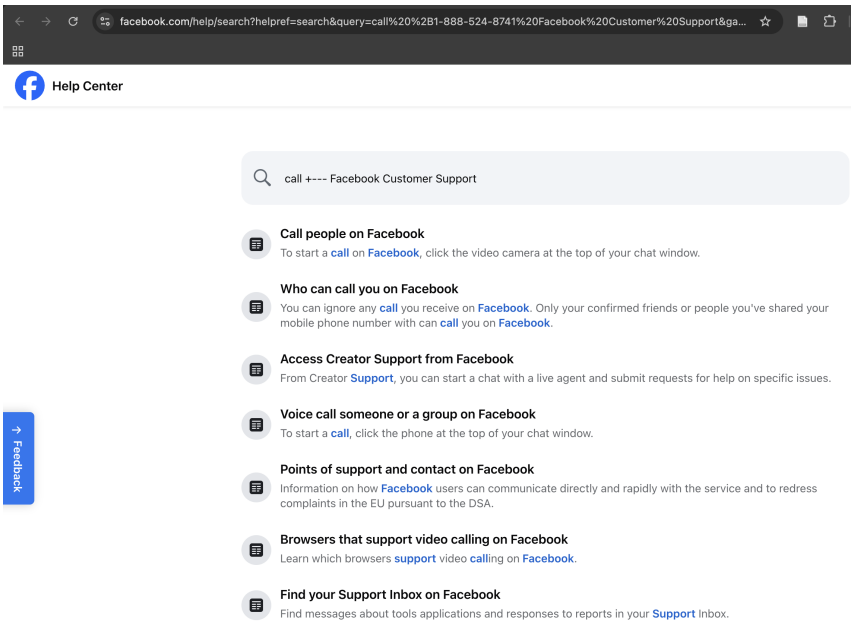

Figure 6.2: Facebook internal site-search with phone number visibly stripped from search query

6.4.3 Recommendations for Google

Our findings highlight several areas where Google's current enforcement and transparency mechanisms fall short. Addressing these issues is critical to prevent the abuse we documented.

- **Improve identity verification.** The current identity verification process is undermined by account sharing and underground markets that offer access to privileged accounts. Google should strengthen enforcement by verifying not only the entity but also access patterns, such as location consistency or sudden shifts in ad campaign behavior, and flag suspicious activity, also on large-scale accounts.
- **Fix the Ads Transparency Center (ATC).** The ATC is currently not usable for large-scale advertisers. It lacks essential features like search and filtering, and the interface becomes unreliable when browsing thousands of ad campaigns. Additionally, when advertisers are removed for policy violations, their ad history is wiped from public view, undermining transparency and accountability. Google should keep historical ad

data for flagged accounts, potentially anonymised in cases of proven account compromise, instead of removing all records.

- **Reconsider URL-based safeguards.** The current check comparing the Display URL and Final URL is easily circumvented and fails to prevent key forms of abuse, including cloaked landing pages, affiliate brand bidding, and URL injection. Google should not rely on this check as a primary enforcement mechanism. One possible way to reduce abuse is to restrict the use of official brand domains in ads. For example, only verified partners or the brand itself could be allowed to use domains like `facebook.com` as a Display URL. While whitelisting has downsides, such as limiting flexibility for legitimate ad campaigns, it would help protect users from misleading ads and reduce the effectiveness of impersonation-based abuse.

6.5 Limitations and Future Work

6.5.1 Limitations

While this study provides new insight into the characteristics of brand impersonation driven malvertising in Google Search Ads, there are some limitations that we will discuss here.

- **Brand-focused scope.** This research intentionally focused on brand-targeted abuse by querying Google Search with brand + intent keywords (e.g., "Facebook support"). As a result, it does not capture broader categories of ad-based abuse that are not related to brands, such as e-commerce scams, deepfake video content, or politically motivated ads. This focus limits the range of abuse types observable in the dataset.

- **Keyword selection bias.** The search terms used for scraping were manually curated and focused on high-risk or high-intent queries (e.g., login, support, signup). This means certain attack vectors, such as misleading ads shown for informational queries, may have been missed.

- **Heuristic-based detection.** Our detection methods rely on a set of heuristics, designed to be general and dataset-agnostic rather than perfectly tuned to the collected data. This design choice was intentional: our goal was to capture a wide range of abuse patterns using interpretable and transferable rules, rather than optimise for precision on this particular dataset. As a result, some heuristics may produce false positives or false negatives. For example, we noticed a few legitimate ads from a large company that were flagged by the Display URL vs Final URL rule. While this company was seemingly allowed to run ads that break Google's policy, we chose not to add a special exception. Doing so would have introduced a false sense of accuracy since we cannot account for similar exceptions we are not aware of. Instead, we accepted a small number of false positives in favour of maintaining rules that could be easily understood and applied broadly across datasets.

  Only the VirusTotal-based heuristic was tested with different thresholds, which makes it more specific to our dataset. However the threshold can be easily adjusted to suit

other datasets if needed. Because VirusTotal labels were retrieved a month after scraping, some domains may have been flagged after our collection window, slightly inflating the malicious count compared with same-day look-ups.

- **Focused brand selection.** The dataset was collected from a curated list of 605 brands including brands with a higher likelihood of attracting malicious ads. This targeted approach enabled the capture of real-world abuse but does not reflect a random sample of the broader ad ecosystem. An attempted follow-up scrape using a wider set of Fortune 500 brands was blocked by Google's anti-bot measures early in the process, limiting our ability to assess prevalence at scale.

- **Google-specific focus.** This study exclusively targeted Google Search Ads. The extent to which similar abuse occurs on other search engines remains unexplored here.

- **Scraping viability over time.** During the early phase of our data collection, Google's search results could be scraped without triggering bot protection. However, later attempts were blocked by Google's detection systems. This limits the ability to reproduce or extend the data collection in its current form.

6.5.2 Future Work

Several opportunities exist to extend or complement this research:

- **Broaden scope and diversify queries.** To capture a wider range of malvertising beyond brand impersonation, future work could expand beyond brand-focused queries and use more varied search terms. This would allow researchers to uncover other categories of abuse such as e-commerce scams, misleading health or financial products, deepfake content, or politically motivated ads, of which we found anecdotal reports during our OSINT investigation.

- **Improve detection methods.** Future work could focus on combining multiple heuristic signals to strengthen detection, or using the flagged ads from this study as training data for supervised machine learning models. This could help reduce false positives and uncover abuse patterns not captured by simple rule-based approaches. In a preliminary experiment, we trained a simple classifier using these labels, which showed promising results, although this was not extensively evaluated and is not presented as a contribution.

- **Scale to a broader brand set.** If scraping constraints can be overcome, repeating this methodology on a broader or more representative sample, such as all Fortune 500 companies, would help assess the general prevalence of abuse across brand tiers and verticals.

- **Compare across platforms.** Applying the same analysis to other ad platforms (e.g., Bing, Facebook, or YouTube) would help determine whether the observed patterns are

unique to Google or more widespread. Notably many of the large-scale advertisers we flagged were also running large numbers of app-install ads, raising the question of whether similar forms of abuse may be occurring in other ad formats or ecosystems.

6.6 Ethical Considerations

During this research, we tried to balance transparency with ethical responsibility. Below, we outline how we approached decisions related to the publication of advertiser identities and disclosure of vulnerabilities.

6.6.1 Naming Abusive Advertisers

We have chosen to include the names of certain advertisers involved in the observed abuse. This decision is justified on the basis that most of these advertisers have been publicly flagged multiple times for abusive or deceptive practices, some dating back several years. Even in cases where the advertiser had a legitimate-looking web presence, the scale and consistency of the abuse, along with public complaints and available OSINT (open-source intelligence), made the risk to users and brands clear.

In many cases, the advertiser accounts are being openly traded or resold in underground markets, posing significant risks regardless of whether a legitimate business entity stands behind them. We were unable to find public-facing contact information to verify the intent or legitimacy of these advertisers, despite reasonable effort.

In our analysis of a hijacked advertiser account belonging to a legitimate travel agency, we observed that this account was used to run both phishing and tech support scams, though on a smaller scale. To avoid reputational harm to this entity that may itself have been a victim, we have redacted both the agency's name and the names of the users writing negative reviews about them in the published material.

6.6.2 Disclosure of URL Injection Vulnerability

We also decided to include a technical discussion of a content injection vulnerability affecting internal site-search pages of many large brands. However, we deliberately omitted the list of specific affected brands.

This vulnerability is already being actively exploited in live ad campaigns, and we believe that disclosure in this case does more good than harm. The attackers already possess the asymmetrical advantage: they can trivially discover other brands that are vulnerable to it, while defenders are unaware, though a fix is relatively easy, as seen in the Facebook example (see 6.2). Making the issue publicly known helps affected companies take action and raises awareness.

Bibliography

[1] About auto-tagging - Google Ads Help, . URL https://support.google.com/google-ads/answer/3095550?hl=en.

[2] Ads transparency - Advertising Policies Help, . URL https://support.google.com/adspolicy/answer/13733850?hl=en.

[3] Advertiser verification - Advertising Policies Help, . URL https://support.google.com/adspolicy/answer/9703665?hl=en.

[4] BlueVision, . URL https://www.bluevision.com/.

[5] Bringme Consultant Services Limited, . URL https://www.hongkongdir.org/companies/2175746/.

[6] Cloud and Threat Report - 2025, . URL https://www.netskope.com/netskope-threat-labs/cloud-threat-report/cloud-and-threat-report-2025.

[7] Digital Advertising - Worldwide | Statista Market Forecast, . URL https://www.statista.com/outlook/dmo/digital-advertising/worldwide.

[8] Final URL: Definition - Google Ads Help, . URL https://support.google.com/google-ads/answer/6080568?hl=en.

[9] How to block referrals from Google Ads | Rewardful Help Center, . URL https://help.rewardful.com/en/articles/8237717-how-to-block-referrals-from-google-ads.

[10] Maximum CPC bid: Definition - Google Ads Help, . URL https://support.google.com/google-ads/answer/6326?hl=en.

[11] My account got hacked. and the hacker used my google ads - Google Ads Community, . URL https://support.google.com/google-ads/thread/191560725/my-account-got-hacked-and-the-hacker-used-my-google-ads?hl=en.

[12] Report a cloaked ad - Google Ad Manager Help, . URL https://support.google.com/admanager/answer/11089715?hl=en.

[13] Trademarks - Advertising Policies Help, . URL https://support.google.com/adspolicy/answer/6118?hl=en.

[14] Traffic Heroes, . URL https://www.trafficheroes.agency/.

[15] Hong Kong Company Directory, . URL https://www.hkcompanydirectory.com/en/xun-meng-international-limited-.

[16] Flyazuh is rated "Poor" with 2.3 / 5 on Trustpilot, December 2023. URL https://www.trustpilot.com/review/flyazuh.com.

[17] Tessavo is rated "Bad" with 1.5 / 5 on Trustpilot, November 2023. URL https://www.trustpilot.com/review/tessavo.com.

[18] Any known Traffic Heroes LTD Media buyer? | BlackHatWorld, January 2024. URL https://www.blackhatworld.com/seo/any-known-traffic-heroes-ltd-media-buyer.1558625/.

[19] Hng dn cách mua và chi tài khon Google Ads invoice Hong Kong nhn mã 385$ - YouTube, 2024. URL https://www.youtube.com/watch?v=uw1lk63walE&list=LL.

[20] Most difficult cloaking google ads issue google ads expert help needed who knows cloaking | BlackHatWorld, January 2025. URL https://www.blackhatworld.com/seo/most-difficult-cloaking-google-ads-issue-google-ads-expert-help-needed-who-knows-cloaking.1600265/page-2.

[21] SKY COSMO LIMITED | Hong Kong Companies Directory, April 2025. URL https://www.ltddir.com/companies/sky-cosmo-limited-2/.

[22] Sahar Abdelnabi, Katharina Krombholz, and Mario Fritz. VisualPhishNet: Zero-Day Phishing Website Detection by Visual Similarity. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, CCS '20, pages 1681–1698, New York, NY, USA, November 2020. Association for Computing Machinery. ISBN 978-1-4503-7089-9. doi: 10.1145/3372297.3417233. URL https://dl.acm.org/doi/10.1145/3372297.3417233.

[23] Bhupendra Acharya and Phani Vadrevu. {PhishPrint}: Evading Phishing Detection Crawlers by Prior Profiling. pages 3775–3792, 2021. ISBN 978-1-939133-24-3. URL https://www.usenix.org/conference/usenixsecurity21/presentation/acharya.

[24] Bhupendra Acharya, Dario Lazzaro, Efrén López-Morales, Adam Oest, Muhammad Saad, Antonio Emanuele Cinà, Lea Schönherr, and Thorsten Holz. The Imitation Game: Exploring Brand Impersonation Attacks on Social Media Platforms. pages

4427–4444, 2024. ISBN 9781939133441. URL https://www.usenix.org/conference/usenixsecurity24/presentation/acharya.

[25] Muhammad Ali, Angelica Goetzen, Alan Mislove, Elissa M. Redmiles, and Piotr Sapiezynski. Problematic Advertising and its Disparate Exposure on Facebook. pages 5665–5682, 2023. ISBN 978-1-939133-37-3. URL https://www.usenix.org/conference/usenixsecurity23/presentation/ali.

[26] Hugo Bijmans, Tim Booij, Anneke Schwedersky, Aria Nedgabat, and Rolf van Wegberg. Catching Phishers By Their Bait: Investigating the Dutch Phishing Landscape through Phishing Kit Detection. pages 3757–3774, 2021. ISBN 978-1-939133-24-3. URL https://www.usenix.org/conference/usenixsecurity21/presentation/bijmans.

[27] Claudio Carpineto and Giovanni Romano. An Experimental Study of Automatic Detection and Measurement of Counterfeit in Brand Search Results. *ACM Trans. Web*, 14(2):6:1–6:35, February 2020. ISSN 1559-1131. doi: 10.1145/3378443. URL https://dl.acm.org/doi/10.1145/3378443.

[28] Fernando Cassia. Fernando Cassia on X: "@Takahiro_nagaki this ..., March 2024. URL https://x.com/fcassia/status/1770309716570173467.

[29] Fernando Cassia. Fernando Cassia on X: XUN MENG ..., March 2024. URL https://x.com/fcassia/status/1770308830695395495.

[30] Jerome Dangu. Malvertiser Makes the Big Bucks on Black Friday, January 2023. URL https://blog.confiant.com/malvertiser-makes-the-big-bucks-on-black-friday-637922cd5865.

[31] Rand Fishkin. New Research: We analyzed 332 million queries over 21 months to uncover never-before-published data on how people use Google, December 2024. URL https://sparktoro.com/blog/new-research-we-analyzed-332-million-queries-over-21-months-to-uncover-never-before-published-data-on-how-people-use-google/.

[32] Anthony Higman. Anthony Higman on X: "FYI: FAKE AD ALERT! Some company..., September 2024. URL https://x.com/AnthonyHigman/status/1866215866838241577.

[33] Luca Invernizzi, Kurt Thomas, Alexandros Kapravelos, Oxana Comanescu, Jean-Michel Picod, and Elie Bursztein. Cloak of Visibility: Detecting When Machines Browse a Different Web. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 743–758, May 2016. doi: 10.1109/SP.2016.50. URL https://ieeexplore.ieee.org/abstract/document/7546533?casa_token=vqlTH0v_aqIAAAAA:HDVbZckgffI8ffVvKhwKYaw_DbK-9CAjzudc55y05AgqFQAQxhbm4GZA7aNnN3pS77Ultvaa.

[34] Jason Koebler ·. Deepfaked Celebrity Ads Promoting Medicare Scams Run Rampant on YouTube, January 2024. URL https://www.404media.co/joe-rogan-taylor-swift-andrew-tate-ai-deepfake-youtube-medicare-ads/.

[35] Takashi Koide, Daiki Chiba, and Mitsuaki Akiyama. To Get Lost is to Learn the Way: Automatically Collecting Multi-step Social Engineering Attacks on the Web. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, ASIA CCS '20, pages 394–408, New York, NY, USA, October 2020. Association for Computing Machinery. ISBN 978-1-4503-6750-9. doi: 10.1145/3320269.3384714. URL https://dl.acm.org/doi/10.1145/3320269.3384714.

[36] Zhou Li, Kehuan Zhang, Yinglian Xie, Fang Yu, and XiaoFeng Wang. Knowing your enemy: understanding and detecting malicious web advertising. In *Proceedings of the 2012 ACM conference on Computer and communications security*, CCS '12, pages 674–686, New York, NY, USA, October 2012. Association for Computing Machinery. ISBN 978-1-4503-1651-4. doi: 10.1145/2382196.2382267. URL https://doi.org/10.1145/2382196.2382267.

[37] Yun Lin, Ruofan Liu, Dinil Mon Divakaran, Jun Yang Ng, Qing Zhou Chan, Yiwen Lu, Yuxuan Si, Fan Zhang, and Jin Song Dong. Phishpedia: A Hybrid Deep Learning Based Approach to Visually Identify Phishing Webpages. pages 3793–3810, 2021. ISBN 978-1-939133-24-3. URL https://www.usenix.org/conference/usenixsecurity21/presentation/lin.

[38] Ruofan Liu, Yun Lin, Xianglin Yang, Siang Hwee Ng, Dinil Mon Divakaran, and Jin Song Dong. Inferring Phishing Intention via Webpage Appearance and Dynamics: A Deep Vision Based Approach. pages 1633–1650, 2022. ISBN 978-1-939133-31-1. URL https://www.usenix.org/conference/usenixsecurity22/presentation/liu-ruofan.

[39] Chris Musteata. Detecting brandjacking-based malvertising, 2024. URL https://www.cs.ru.nl/bachelors-theses/2024/Chris_Musteata__1035303___Detecting_Brandjacking-Based_Malvertising.pdf.

[40] Adam Oest, Yeganeh Safaei, Adam Doupé, Gail-Joon Ahn, Brad Wardman, and Kevin Tyers. PhishFarm: A Scalable Framework for Measuring the Effectiveness of Evasion Techniques against Browser Phishing Blacklists. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1344–1361, May 2019. doi: 10.1109/SP.2019.00049. URL https://ieeexplore.ieee.org/document/8835369.

[41] Irfan Ozen, Karthika Subramani, Phani Vadrevu, and Roberto Perdisci. SENet: Visual Detection of Online Social Engineering Attack Campaigns, January 2024. URL http://arxiv.org/abs/2401.05569. arXiv:2401.05569 [cs].

[42] William Salusky. Malvertising, June 2007. URL https://isc.sans.edu/diary/3727.

[43] Sarah Miller. A surge of malvertising across Google Ads is distributing dangerous malware, February 2023. URL https://www.spamhaus.com/resource-center/a-surge-of-malvertising-across-google-ads-is-distributing-dangerous-malware/.

[44] Jérôme Segura. Malvertising via brand impersonation is back again, May 2023. URL https://www.malwarebytes.com/blog/threat-intelligence/2023/05/malvertising-its-a-jungle-out-there.

[45] Jérôme Segura. The great Google Ads heist: criminals ransack advertiser accounts via fake Google ads, January 2025. URL https://www.malwarebytes.com/blog/cybercrime/2025/01/the-great-google-ads-heist-criminals-ransack-advertiser-accounts-via-fake-google-ads.

[46] SlapSlay3r. DMARKET running scams from inside, October 2024. URL www.reddit.com/r/D_Market/comments/1gb1j91/dmarket_running_scams_from_inside/.

[47] SomeRevolution9849. I just got a trade scam from Dmarket, October 2023. URL www.reddit.com/r/cs2/comments/17gsxjx/i_just_got_a_trade_scam_from_dmarket/.

[48] Giada Stivala, Sahar Abdelnabi, Andrea Mengascini, Mariano Graziano, Mario Fritz, and Giancarlo Pellegrino. From Attachments to SEO: Click Here to Learn More about Clickbait PDFs! In *Proceedings of the 39th Annual Computer Security Applications Conference*, ACSAC '23, pages 14–28, New York, NY, USA, December 2023. Association for Computing Machinery. ISBN 979-8-4007-0886-2. doi: 10.1145/3627106.3627172. URL https://doi.org/10.1145/3627106.3627172.

[49] Roy Tay and Sudeep Singh. Malvertising campaign targeting IT teams with MadMxShell, April 2024. URL https://www.zscaler.com/blogs/security-research/malvertising-campaign-targeting-it-teams-madmxshell.

[50] Ryan Tomcik, Adrian McCabe, Rufus Brown, and Geoff Ackerman. Opening a Can of Whoop Ads: Detecting and Disrupting a Malvertising Campaign Distributing Backdoors, December 2023. URL https://cloud.google.com/blog/topics/threat-intelligence/detecting-disrupting-malvertising-backdoors.

[51] United Nations Office on Drugs and Crime. Transnational organized crime and the convergence of cyber-enabled fraud, underground banking and technological innovation in southeast asia: A shifting threat landscape, 2024. URL https://www.unodc.org/roseap/uploads/documents/Publications/2024/TOC_Convergence_Report_2024.pdf. Accessed: 2025-04-10.

[52] Phani Vadrevu and Roberto Perdisci. What You See is NOT What You Get: Discovering and Tracking Social Engineering Attack Campaigns. In *Proceedings of the Internet Measurement Conference*, IMC '19, pages 308–321, New York, NY, USA,

October 2019. Association for Computing Machinery. ISBN 978-1-4503-6948-0. doi: 10.1145/3355369.3355600. URL https://doi.org/10.1145/3355369.3355600.

[53] Yutian Yan, Yunhui Zheng, Xinyue Liu, Nenad Medvidovic, and Weihang Wang. AdHere: Automated Detection and Repair of Intrusive Ads. In *Proceedings of the 45th International Conference on Software Engineering*, ICSE '23, pages 486–498, Melbourne, Victoria, Australia, July 2023. IEEE Press. ISBN 978-1-6654-5701-9. doi: 10.1109/ICSE48619.2023.00051. URL https://dl.acm.org/doi/10.1109/ICSE48619.2023.00051.

[54] Zheng Yang, Joey Allen, Matthew Landen, Roberto Perdisci, and Wenke Lee. TRIDENT: towards detecting and mitigating web-based social engineering attacks. In *Proceedings of the 32nd USENIX Conference on Security Symposium*, SEC '23, pages 6701–6718, USA, August 2023. USENIX Association. ISBN 978-1-939133-37-3.

[55] Apostolis Zarras, Alexandros Kapravelos, Gianluca Stringhini, Thorsten Holz, Christopher Kruegel, and Giovanni Vigna. The Dark Alleys of Madison Avenue: Understanding Malicious Advertisements. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, pages 373–380, New York, NY, USA, November 2014. Association for Computing Machinery. ISBN 978-1-4503-3213-2. doi: 10.1145/2663716.2663719. URL https://doi.org/10.1145/2663716.2663719.

[56] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites. 2020.

[57] Eric Zeng, Miranda Wei, Theo Gregersen, Tadayoshi Kohno, and Franziska Roesner. Polls, clickbait, and commemorative $2 bills: problematic political advertising on news and media websites around the 2020 U.S. elections. In *Proceedings of the 21st ACM Internet Measurement Conference*, IMC '21, pages 507–525, New York, NY, USA, November 2021. Association for Computing Machinery. ISBN 978-1-4503-9129-0. doi: 10.1145/3487552.3487850. URL https://dl.acm.org/doi/10.1145/3487552.3487850.

[58] Penghui Zhang, Adam Oest, Haehyun Cho, Zhibo Sun, RC Johnson, Brad Wardman, Shaown Sarker, Alexandros Kapravelos, Tiffany Bao, Ruoyu Wang, Yan Shoshitaishvili, Adam Doupé, and Gail-Joon Ahn. CrawlPhish: Large-scale Analysis of Client-side Cloaking Techniques in Phishing. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1109–1124, May 2021. doi: 10.1109/SP40001.2021.00021. URL https://ieeexplore.ieee.org/document/9519414.

[59] M. Zubair Rafique, Tom Van Goethem, Wouter Joosen, Christophe Huygens, and Nick Nikiforakis. It's Free for a Reason: Exploring the Ecosystem of Free Live Streaming Services. In *Proceedings 2016 Network and Distributed System Security Symposium*, San Diego, CA, 2016. Internet Society. ISBN 978-1-891562-41-9. doi: 10.14722/ndss.2016.23030.

URL https://www.ndss-symposium.org/wp-content/uploads/2017/09/free-reason-exploring-ecosystem-free-live-streaming-services.pdf.

Appendix A

Experimental Input Data

- Freshdesk
- Zendesk
- mcafee
- PRTG Network Monitor
- Rapid7 Nexpose
- Sendinblue
- Veeam Backup & Replication
- dmarket
- SolarWinds IP Address Manager
- MailerLite
- Acunetix
- Monday.com
- ManageEngine OpManager
- Zapier
- getresponse
- SugarCRM
- Elasticsearch
- Zoho CRM
- OFX
- Qualys
- Asana
- Mailchimp
- Later
- Hotjar
- mylowes
- Kubernetes
- GanttProject
- lodgify
- Automate.io
- Tenable Nessus
- SENDX
- Mouseflow
- Sucuri
- PingPlotter
- WSJ
- SocialBee
- synthflow ai
- HubSpot
- landerbolt
- SAP concur
- ClickUp
- AWS
- Power Automate
- Burp Suite
- Serpstat
- OpenVAS
- engagebay
- microsoft
- surveysparrow
- WiFi Analyzer
- Hootsuite
- fastbots ai
- sintra ai
- segmetrics

- IP Scanner
- FullStory
- vmware
- instantproxies
- google authenticator
- OWASP ZAP
- brilliantdirectories
- Checkmk
- activecampagin
- mailshake
- Surfer SEO
- Jira
- blaze ai
- Ramp
- fellahealth
- buff
- trendrocket
- timetastic
- podpage
- chatbase
- Splunk
- Harvest
- Cisco AnyConnect
- reipro
- smartleadai
- Nagios
- skedsocial
- Bambee
- finalround ai
- Jenkins
- databox
- synthflow
- Lyca Mobile
- voicespin
- omnisend
- Trello
- kahunas
- salespype
- collectchat
- myask ai
- TeamViewer
- Metasploit
- cloudways
- uplead
- apple
- Google Tag Manager
- danelfin
- grouptrack crm
- thrivecart
- leetify
- calendly
- buyproxies
- sendible
- MRTG
- passion io
- apollo
- alitu
- anytrack
- Rebrandly
- aragon ai
- greatlysocial
- sprintful
- Bloomberg
- zoof
- fireflies
- paperbell
- leadcapture
- Meta
- Crazy Egg
- Spyder
- aweber
- maxai
- Paypal
- Microsoft Teams
- colddms
- jobboardly
- uploadcare
- WorldRemit
- crosslist
- sillaby
- wonderchat
- shipito

- Spiceworks
- Buffer
- backblaze
- BlackRock
- Plausible Analytics
- kalodata
- AstraChat
- Graylog
- heropost
- Tradingview
- klap app
- clickexpose
- AOMEI Backupper
- Matomo
- expirationreminder
- Postman
- Rank Math
- collosyan
- videotap
- Remitly
- Advanced Port Scanner
- copygram
- speechify
- bookingkoala
- proxyline
- Chime
- hoppycopy
- openbiz
- everbee
- sanebox
- frontly
- searchapi
- AnyDesk
- Bitly
- userback
- anyword
- Qonto
- Netdata
- getmunch
- harpoonapp
- sourcemogul
- SyncBack
- uplisting
- walmart
- Screaming Frog SEO Spider
- cheq ai
- ownr
- stealthwriter
- Amazon
- stealthgpt
- videogen
- Kibana
- Miro
- MongoDB Compass
- soundraw
- heysummit
- madgicx
- webshare
- Mercury
- notion
- TickTick
- bookafy
- textcortex
- Zabbix
- Bitwarden
- akool
- quizgecko
- Ubersuggest
- bodygraphchart
- thereadystate
- Open Broadcasting Studio
- kikkerland
- rewardful
- epinium
- vendoo
- OpenStack
- heygen
- mixoio

- Proxmox VE
- Toggl
- OpenProj
- outranking
- appsumo
- enterprisedna
- pipefile
- Payoneer
- Censys
- Yoast SEO
- Ahrefs Webmaster Tools
- quicknode
- CoSchedule
- viralsweep
- CircleCI
- memberspace
- synthesia
- Wunderlist
- fusebase
- TeamCity
- gmapsextractor
- easysong
- LEARNINGSTUDIO AI
- buybotpro
- udimi
- Nikto
- syllaby
- 7taps
- jenni ai
- bigspy
- madaz money
- GitHub Desktop
- Basecamp
- crayo ai
- nudgify
- Slack
- vizzlo
- Transferwise
- Hyper-V
- joelister
- HWiNFO
- builtright
- Clockify
- RescueTime
- DBeaver
- KiCad
- salefreaks
- Wireshark
- Logstash
- Lucidchart
- anstrex
- chatpdf
- postplanner
- Blender 3D
- checkoutpage co
- decktopus
- kaspr
- mindstamp
- apitemplate
- Cloudflare
- beeyondai
- pgAdmin
- followr
- Macrium Reflect
- MySQL Workbench
- humata
- Grafana
- SmartGit
- bounceban
- Webex
- Nmap
- Docker
- Microsoft To Do
- Android Studio
- Travis CI
- chaindesk
- instawp
- nuelink
- Clonezilla
- timedoctor
- adcreative.ai

- Instagram
- SQLiteStudio
- ProjectLibre
- zopto
- usesuperflow
- StarUML
- FreeCAD
- photovibrance
- Jetpack
- Bitgo
- ChatGPT
- Belarc Advisor
- choppity
- stealthwriter ai
- Binance
- trueclicks
- BGInfo
- FreeFileSync
- phpMyAdmin
- trymaverick
- pabbly
- soonworks
- ddevi
- uniswap
- Uphold
- Prometheus
- Diagrams.net
- Jupyter Notebook
- tugan ai
- breeew
- anbernic
- anydo
- kits ai
- magai
- syncee
- taskade
- wondercraft
- adalo
- finalscout
- groovecm
- taja
- GitKraken
- VirtualBox
- Skype
- Bitbucket
- MozBar
- gerpc
- etsyhunt
- graphy
- boost
- Cyberduck
- screenprotech
- Zoom
- Cacti
- highnote
- yomu ai
- Snort
- Sysinternals Suite
- scoreapp
- memberstack
- foreplay
- quartr
- logoai.com
- signaturely
- magicslides
- CNN
- specterr
- Facebook
- fansmetric
- findniche
- opusvirtualoffics
- planetexpress
- Devolutions Remote Desktop Manager
- SEOquake
- Monzo
- OneNote
- Todoist
- diffy
- vidyo ai
- Adidas
- Copy ai
- Skrill

- Fiddler
- Icinga
- minea
- captivatefm
- freshlearn
- lovo ai
- interexo
- heidihealth
- System Ninja
- Vagrant
- Landbot
- airops
- manning
- vidnoz
- sleeknote
- Transmit
- saasrock
- speakflow
- Go To Meeting
- Open Hardware Monitor
- elevenlabs
- jupitrr
- quizbreaker
- publit io
- sms-activate
- Revolut
- IntelliJ IDEA
- Sourcetree
- calendafy
- FedEx
- Radient
- Octave
- servanamanaged
- seospace
- coinrule
- neuraltext
- playht
- USPS
- Uber
- Zenmap
- Robocopy
- Grammarly
- Visual Paradigm Community
- QCAD
- rundiffusion
- typefully
- ceartas
- Gimp
- CPU-Z
- ExamDiff
- MAMP
- theretrokit
- Bybit
- FileZilla Pro
- framer com
- pocket option
- quoteplicity
- droxy
- VLC media player
- Addsearch
- plexicam
- stubhub
- dataspark
- vetrec
- Process Explorer
- Visual Studio Code
- Figma
- Draw.io
- exportyourstore
- DHL
- LAN Speed Test
- elai
- saaspegasus
- filezilla
- WampServer
- proxy-n-vpn
- skeddly
- UPS
- IFTTT
- AnswerThePublic
- XAMPP

- HeidiSQL
- Coinbase
- RStudio
- ppspy
- reroom ai
- sellerlabs
- contentharmony
- pitchdb
- N26
- Lido
- TCPView
- Roam Research
- Let's Encrypt
- propel io
- proxy6
- wordgalaxy
- typecast
- Monese
- Bitvavo
- Hemingway Editor
- binom
- shophunter
- bulltrixtrading
- pingbell
- clipgen
- dropchat
- hypermedialab
- RRDTool
- rsync
- UltraCompare
- copyleaks
- Signify
- Bunq
- Speccy
- FastCopy
- PyCharm
- Git
- Element
- Obsidian
- LibreCAD
- JustPark
- 7zip
- Beyond Compare
- Notability
- keygensh
- Neo Financial
- Kucoin
- Poloniex
- Google Meet
- Tcpdump
- Kismet
- Shodan
- Eclipse
- PlantUML
- Spotify
- TinyURL
- pancakeswap
- Koho
- Netcat
- CrystalDiskMark
- Atom
- Amphenol
- bycasino
- WinSCP
- putty
- Duplicati
- QEMU
- Inkscape
- narrato
- makereels
- divjoy
- immigo
- Atlassian
- betpublic
- Nike
- GPU-Z
- askyourpdf
- koddosnet
- lifepixel
- webgazer
- vondy
- Starzbet

- bitget
- crypto.com
- Brave
- DefiLlama
- Evernote
- Pocket
- gizzmo
- otomatic
- aimlapi
- traderspost
- Philips
- Exness
- onwin
- xslot
- draftkings
- 010 Editor
- Bear
- wizer
- pairrd
- emed
- Loom
- fanduel
- Arc Browser
- Simplenote
- Raindrop.io
- Anaconda
- simplebackups
- moovd
- Mastercard Foundation
- Camper and Nicholsons
- betmatik
- Kraken
- HWMonitor
- WinMerge
- NetBeans
- illow io
- profitl
- corsix
- xnapper
- youthfully
- nexlev
- VVA LLC
- matadorbet
- fixbet
- supertotobet
- upbit
- PayByPhone App
- RingGo
- WinDirStat
- faceit
- CrystalDiskInfo
- SpeedFan
- TeraCopy
- Overleaf
- chapple
- cornerrr
- gremi
- applevel
- ringly
- Ledwork
- BancaProgetto
- betmgm
- Airbnb
- magic eden
- opensea

A.1 Keyword List

Category	Keywords
Banking	login, account, support, deposit
Financial Services	login, account, invoice, support
Technology Platforms	login, account, sign up, support, password, hacked
Internet Services	login, account, sign up, password, support, profile, hacked
Telecom	login, account, bill, verify, plan, recharge
E-Commerce	login, account, password, support, order, checkout, shipping
Government Services	login, account, official, support
Logistics	tracking, shipment, delivery, order, support
Entertainment	login, account, subscription, password, premium, free, streaming, support
Postal Services	tracking, shipment, delivery, parcel, package, support
Crypto	login, wallet, account, transaction, withdrawal, deposit, support, hacked
Gambling	login, account, bet, withdrawal, deposit, bonus, support
Downloadable	
Open Source Tools	download, install, setup, update, latest version, free, support
Affiliate	sign up, login, bonus, code, account, pricing, demo
Rest	login, account, password, support

Table A.1: Search keyword list by industry category.

Appendix B

Heuristic Details

B.1 Affiliate Parameter List

The following URL parameters and patterns were used to flag ads potentially containing affiliate links:

- `via`
- `?ref`
- `fpr`
- `afmc`
- `irgwc`
- `affiliate`
- `pscd`
- `ps_partner_key`
- `fp_ref`
- `sca_ref`
- `gr_pk`
- `_go=`
- `_r=`
- `_pac=`
- `acode=`
- `lmref`

- a_aid
- ?rfsn
- ?deal=
- #_r
- clickid (only when combined with irgwc)

B.2 Discount-Site Brand Bidding Keyword List

The following keywords were used to flag discount-site brand bidding ads. An ad was flagged if:

- Its title **and** display URL both contained at least one of the following words:
  - deals
  - discount
  - coupon
  - promo
  - bundle
  - bonus
- The brand name also appeared in the title

Ads were **not** flagged if the advertiser name included the brand name, to avoid penalising legitimate ads by the brand itself.

Appendix C

Full Per-Brand Breakdown of Flagged Ads

Brand	Ads flagged	Total ads	Percentage
Photovibrance	16	57	28.1%
DMarket	49	546	9.0%
SMS-Activate	1	28	3.6%
Zoom	1	43	2.3%
Calendly	2	147	1.4%
Amazon	1	93	1.1%
Meta	1	132	0.8%
Cloudways	1	153	0.7%

Table C.1: Ads flagged by Final URL domain $\neq$ Display URL domain heuristic, shown by brand (True Positives only)

Brand	Ads flagged	Total ads	Percentage
Meta	9	132	6.8%
Facebook	4	39	10.3%
Microsoft Teams	2	128	1.6%
OneNote	1	37	2.7%
Process Explorer	1	20	5%

Table C.2: Ads flagged by phone number heuristic, shown by brand (True Positives only)

Brand	Ads flagged	Total ads	Percentage
McAfee	79	774	10.2%
DMarket	3	546	0.55%

Table C.3: Ads flagged by VirusTotal heuristic, shown by brand (True Positives only)

Brand	Ads flagged	Total ads	Percentage
Advanced Port Scanner	23	109	21.1%
Cisco AnyConnect	1	183	0.5%
IP Scanner	17	206	8.3%
Trend Rocket	4	189	2.1%

Table C.4: Ads flagged by TLD in second-level domain heuristic, shown by brand (True Positives only)

Brand Name	Ads flagged	Total Ads	Percentage Flagged	Abusing Advertisers
EverBee	92	105	87.6%	10
BodyGraph Chart	83	88	94.3%	6
MaxAI.me	83	130	63.8%	1
AppSumo	77	84	91.7%	8
Fireflies.ai	76	133	57.1%	11

Table C.5: Overview of the 5 brands with the greatest number of affiliate brand bidding ads

Brand	Ads flagged	Total ads	Percentage
Stealthwriter	29	149	19.5%
Kikkerland	16	87	18.4%
Final Round AI	12	172	7.0%
Cloudways	10	153	6.5%
Madaz	9	74	12.2%

Table C.6: Overview of the 5 brands with the greatest number of Discount Site Brand Bidding ads

Appendix D

OSINT Examples on Malicious Advertisers

This appendix provides additional details on the open-source intelligence (OSINT) findings described in 5.3.3, presenting evidence gathered from public forums, social media groups, and consumer complaint sites. The findings show how certain advertisers in our dataset are tied to reported scams, account-rental offers, or other suspicious activity.

D.1 Xun Meng International Limited

Lack of Public Web Presence. Despite the large scale of Xun Meng's ads, we found no official corporate website or LinkedIn page. A public Hong Kong Company Directory [15] entry lists its incorporation but offers no details on actual business activities.

Complaints and Scam Reports. Multiple Trustpilot reviews describe alleged scams connected to Xun Meng's ads [16, 17]. Users report being directed to fake online shops or malicious pages after clicking on seemingly legitimate ads.

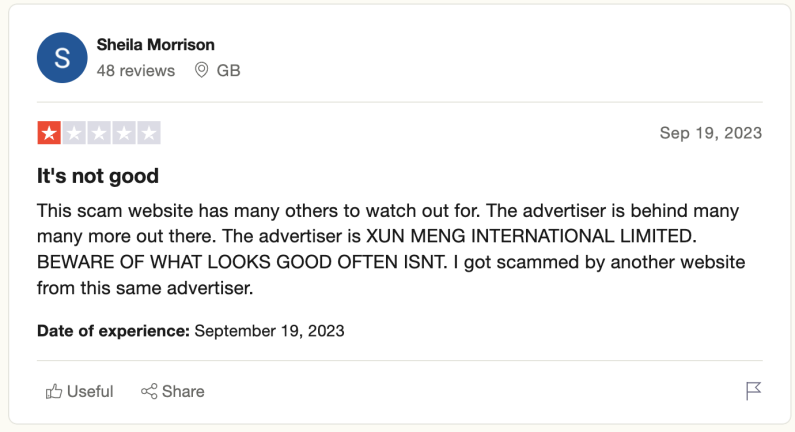

Figure D.1: An example of a review alleging a scam connected to Xun Meng Intl. Ltd.

Account Sharing Signals Access to Xun Meng Intl. Ltd. Google Ads accounts was openly being advertised in the Vietnamese Facebook communities we discovered and mentioned in 5.3.3 (See D.2, D.3 and D.4).

Figure D.2: Example 1 of access to Xun Meng Intl. Ltd. being advertised

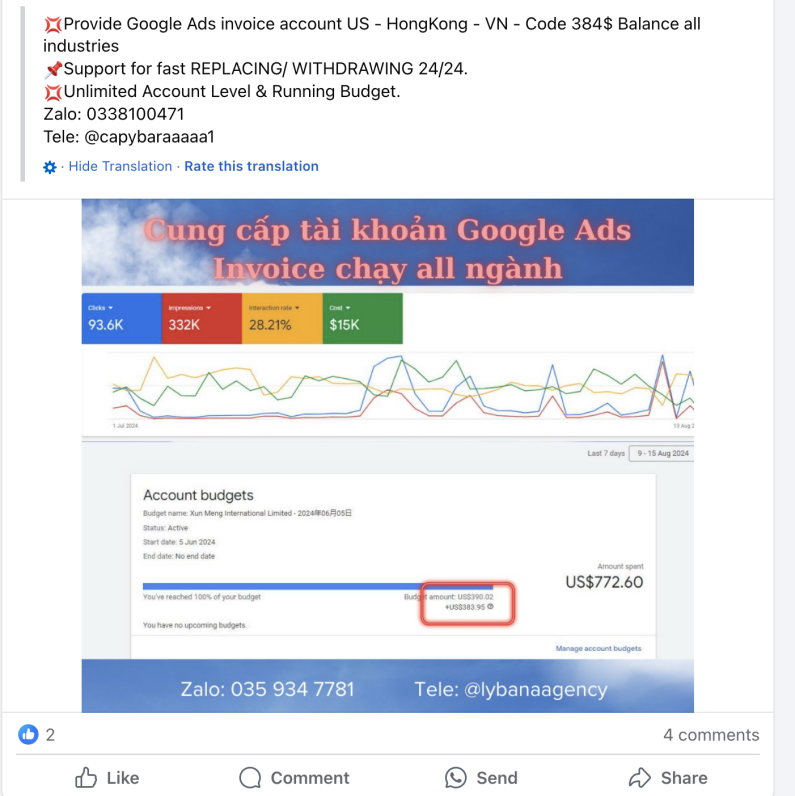

Figure D.3: Example 2 of access to Xun Meng Intl. Ltd. being advertised

Figure D.4: Example 3 of access to Xun Meng Intl. Ltd. being advertised

Political Ads Controversy. Multiple X/Twitter posts claim that Xun Meng's verified Google Ads account ran election ads in Argentina [28, 29]. Figure D.5 shows a user-posted screenshot indicating that one such ad last appeared on 13 October 2023, in the middle of the 2023 Argentinian election period. Under Google's policies, running political ads typically requires additional vetting and geographic restrictions. The fact that a Hong Kong–based advertiser account, advertised for rent in the Facebook groups mentioned earlier, apparently displayed these ads suggests that malicious actors could bypass the usual identity checks and regulations around political advertising through these types of accounts.

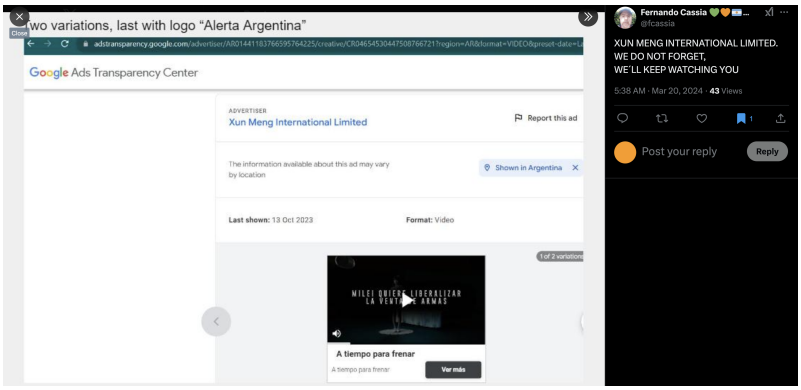

Figure D.5: A user on X.com sharing a screenshot of an Argentine election ad attributed to Xun Meng Intl. Ltd., shown on 13 Oct 2023 (within the 2023 election period).

D.2 Bringme Consultant Services Ltd.

Lack of Public Web Presence. This advertiser also ran tens of thousands of ads according to the ATC, despite this large footprint, we found no corporate website or public profile, besides a listing in a Hong Kong company directory [5].

Complaints and Scam Reports. Figure D.6 shows user-posted screenshots of a scam allegedly being advertised by Bringme Consultant Services Ltd.

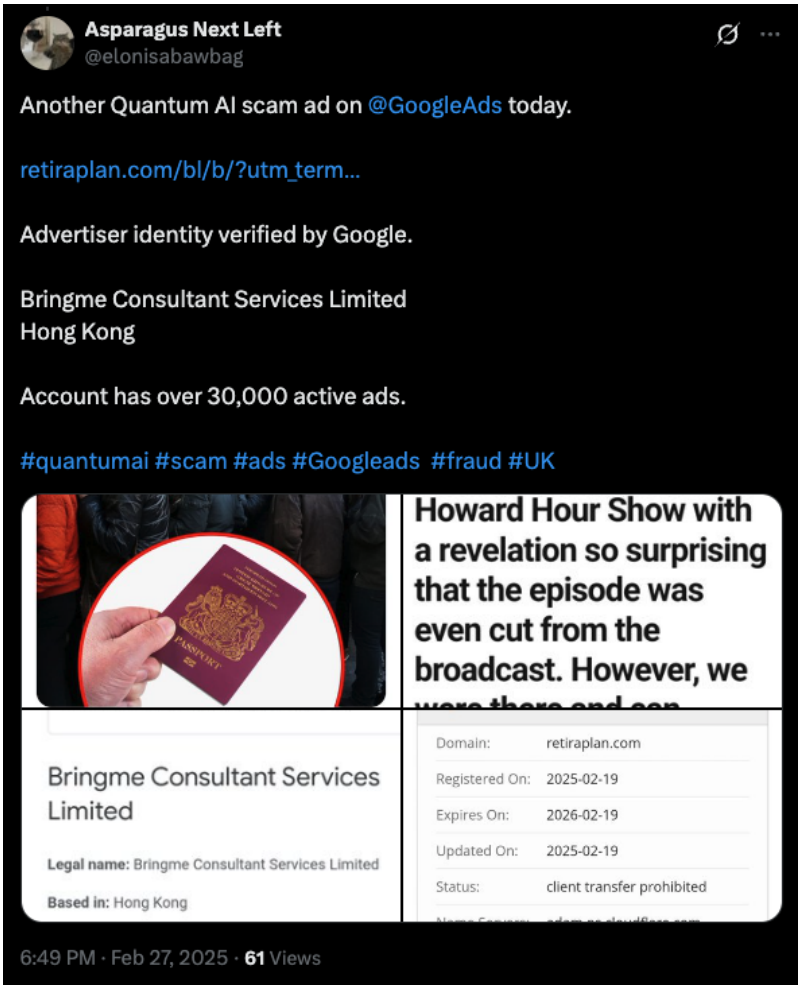

Figure D.6: An example of an alleged scam connected to Bringme Consultant Services Ltd.

Account Sharing Signals. Figure D.7 shows a screenshot of access to Bringme Consultant Services Ltd. being advertised.

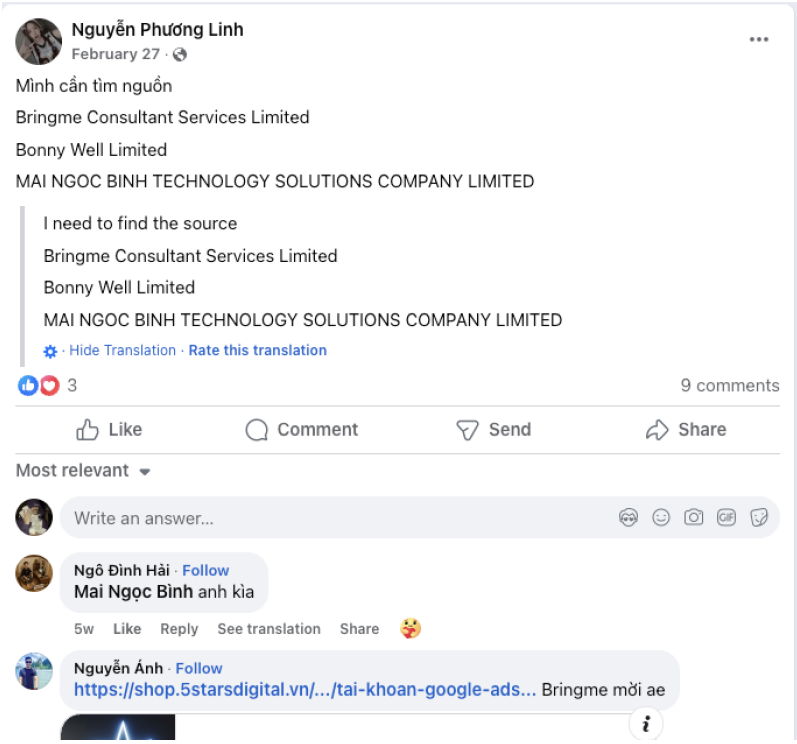

Figure D.7: Example of access to Bringme Consultant Services Ltd. being advertised

D.3 Traffic Heroes Ltd.

Public Web Presence. Traffic Heroes Ltd. is one of the advertisers that does have a public web presence, offering digital advertising services [14]. This indicates they may unknowingly be facilitating the abuse we've flagged and other users have reported online.

Complaints and Scam Reports. On social media, users have accused Traffic Heroes of promoting malicious links [32].

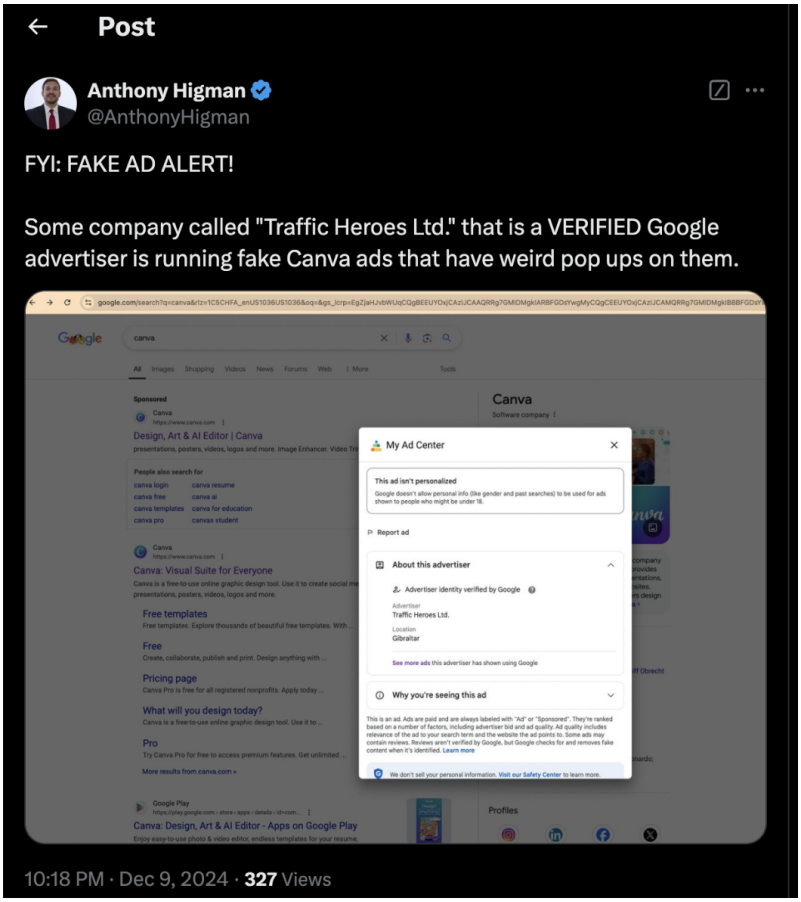

Figure D.8: A user on X.com sharing a screenshot of an ad ran by Traffic Heroes Ltd., allegedly leading to a suspicious landing page.

Account Sharing Signals. We found discussions on marketing forums indicating that individuals have bought or sold Google Ads accounts to Traffic Heroes Ltd [18].

D.4 BlueVision Interactive Ltd.

Public Web Presence. BlueVision also has a public web presence, offering digital advertising services [4]. Similarly to Traffic Heroes, their advertising infrastructure may unknowingly be abused by malicious advertisers.

Complaints and Scam Reports. BlueVision has been accused by this Reddit post to have run Deepfake Celebrity video ads that were reported by 404 Media [34] for leading to a Medicare Scam on a large scale.

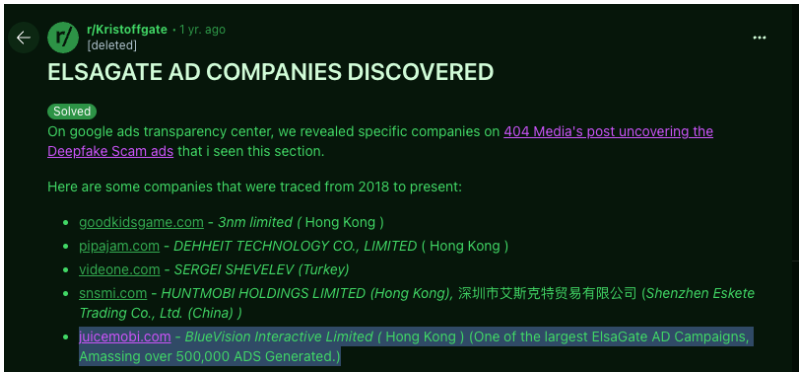

Figure D.9: A user on reddit.com accusing BlueVision, of participating in a Deepfake Celebrity video ad scam.

Account Sharing Signals. This account was not found openly being advertised in any of the underground communities we've discovered and searched.

D.5 Sky Cosmo Limited.

Public Web Presence. This advertiser's only publicly traceable information is a listing in a Hong Kong Company Directory [21].

Complaints and Scam Reports. For this advertiser, we found no public complaints or scam reports, although it was flagged by our heuristics for abuse.

Account Sharing Signals. We did find this account being advertised in the Facebook communities mentioned earlier, as shown in D.10

Figure D.10: Example of access to Sky Cosmo Limited being advertised