

Search + Seizure: The Effectiveness of Interventions on SEO Campaigns

David Y. Wang Matthew Der Mohammad Karami[†]
Lawrence Saul Damon McCoy[†] Stefan Savage Geoffrey M. Voelker

University of California, San Diego George Mason University[†]

Abstract

Black hat search engine optimization (SEO), the practice of abusively manipulating search results, is an enticing method to acquire targeted user traffic. In turn, a range of interventions—from modifying search results to seizing domains—are used to combat this activity. In this paper, we examine the effectiveness of these interventions in the context of an understudied market niche, *counterfeit luxury goods*. Using eight months of empirical crawled data, we identify 52 distinct SEO campaigns, document how well they are able to place search results for sixteen luxury brands, how this capability impacts the dynamics of their order volumes and how well existing interventions undermine this business when employed.

1. INTRODUCTION

Every new communications medium inevitably engenders a new form of abuse — telephones led to unsolicited “robocalls”, email beat spam, and so on. In turn, new mechanisms and policies are invariably brought to bear to restrict such activities (e.g., spam filters or, in the U.S., the national do-not-call registry). Today, one of the most dynamic such conflicts is playing out in the medium of online search.

In particular, as online marketing has become the leading mechanism by which sellers of goods and services engage potential consumers online, search engines, such as Google and Bing, have become the primary platform of this engagement. Because search engine results are presented directly in response to user queries, they offer the opportunity to precisely target consumers at the moment of their interest. As a testament to this, search engines received over \$16B in revenue in 2012 (46% of the total online advertising expenditures) for clicks on sponsored advertisements appearing in their search engine result pages (SERPs) [29].

However, while criminal use of sponsored advertisements occurs, the more fertile ground for abuse is the so-called “organic” search results, which are unpaid. These results are generated and ranked automatically based on the content and structure of the visible Web (e.g., based on the PageRank algorithm, the presence of user-generated content, etc.) and produce far more click traffic than sponsored ads. Unsurprisingly, techniques for improving the ranking of particular Web sites in these organic search results —

termed search engine optimization (SEO) — are extremely popular. While some SEO techniques are completely benign (e.g., keyword friendly URLs), quite a few are actively abusive (e.g., the use of compromised Web sites as “doorway” pages, “cloaking”, farms of “back links”, etc.). As a result, such “black hat” SEO campaigns are frequently able to poison search results so that one or more highly-ranked results for key search terms will direct traffic to their sites. This traffic can then be monetized by infecting the user with malware [11, 14, 30], defrauding the user via phishing [38], or through the marketing of counterfeit or illegal goods (e.g., pharmaceuticals [25]).

In this paper, we focus on a range of such SEO campaigns that are the principal means of marketing for organizations selling counterfeit luxury and lifestyle fashion goods. To wit, at the time of this writing, typing “cheap louis vuitton” into Google produces a list of ten results. Fully seven of these are fraudulent and ultimately direct user clicks to storefronts selling counterfeit knockoffs of Luis Vuitton products. This is no exception and similar search result poisoning is evident for a range of luxury brand names. Indeed, the combination of both high demand and high margins (a counterfeit of a handbag that might retail for \$2400 will sell for \$250, but will typically cost as little as \$20 to produce) make this a vibrant and profitable scam; we have evidence that a *single* fulfillment organization delivered over 250,000 such items over a nine-month period. However, such actors are not unopposed and there are a range of interventions they must contend with including labeling and deranking of their sites by search engine operators, and site or domain takedowns driven by brand holders. It is understanding the interplay of SEO campaigns and these interventions that motivates our research.

Concretely, our paper makes three contributions. First, we provide the first large-scale empirical characterization of SEO abuse for luxury brands. In particular, we explain how such scams work (identifying how they differ from existing markets technically and operationally), analyze their search placement success over time and, using the prior “purchase pair” technique [16], gather indicators of order flow volumes. Second, we develop a methodology for using this data to evaluate the impact of interventions on the effectiveness of these SEO campaigns. Finally, we apply our methodology to a range of existing anti-counterfeiting actions, identify why these prior efforts have had limited impact and make suggestions for improving their utility in the future.

The remainder of the paper is structured as follows. We begin by describing the background of search engine optimization and prior research in Section 2. In Section 3, we describe the technical details of how SEO campaigns are structured in the counterfeit luxury market as well as how current interventions by search engines and brand holders operate and the rationale behind them. Section 4 describes our data set and data collection methodology. Finally,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IMC'14, November 5–7, 2014, Vancouver, BC, Canada.
Copyright 2014 ACM 978-1-4503-3213-2/14/11 ...\$15.00.
<http://dx.doi.org/10.1145/2663716.2663738>.

in Section 5 we present our findings and their implications, while summarizing the most significant in the conclusion.

2. BACKGROUND

The term search engine optimization (SEO) covers a broad array of techniques, all designed to improve the ranking of organic search results in popular search engines. Given that the goal of search engines is to provide high-quality results, only the subset of these techniques that explicitly aid in search quality are viewed as benign by search companies. For example, in their “Search Engine Optimization Starter Guide”, Google suggests things such as using accurate page titles and the use of the “description” meta tag, which Google’s ranking algorithms view positively [8]. However, these benign techniques do not change the underlying Web link structure, and are insufficient to produce the large-scale changes in rankings required to capture significant traffic for popular queries.

Thus, “black hat” SEO campaigns will typically orchestrate thousands of Web sites operating in unison to achieve their goals. Each site will present the targeted keywords when visited by a search engine crawler.¹ However, when a visitor arrives at such a site via a targeted search, entirely different content will be revealed (this is one case of a general technique called “cloaking”). This content may be native to the site, may be proxied from a third site or, most commonly, will arise from a redirection to the true site being advertised. The popularity of this third approach is why such sites are commonly called “doorways” in the SEO vernacular. Doorways in turn obtain high-ranking either by mimicking the structure of high reputation sites (typically by creating backlinks to each other) or by compromising existing sites and exploiting the positive reputation that they have accrued with the search engine.

Poisoned search results (PSRs), or search results promoted by an attacker using black hat SEO with the intent of surreptitiously amassing user traffic, have been studied for a decade, with one of the best-known early empirical analyses due to Wang et al. [37]. More modern analyses have covered advances in detecting cloaking and poisoning techniques [19, 22, 27, 35] as well as deeper studies of the particular campaigns and their operational behavior [15, 36]. These efforts, which identify a range of technical behaviors implicated in abusive SEO, serve as the foundation for our own measurement work. However, our goals differ considerably from previous work as we are focused on understanding the overall business enterprise and through this lens evaluate the efficacy of existing interventions. In this, our work is similar to prior efforts to understand particular underground economies [16, 17, 18, 25, 34] and the economic issues surrounding various defenses and interventions [3, 12, 20, 21, 24, 26].

However, the ecosystem around luxury SEO abuse is quite distinct from these others and, as we will show, there are large differences in the underlying techniques, business structure, stakeholders and the kinds of interventions being practiced. Thus, we believe that our findings both serve to advance our understanding of how to best address search abuse, but also to expand our broader understanding about the interplay between technical countermeasures and the structure of online criminal enterprises.

3. LUXURY SEO AND INTERVENTIONS

Abusive SEO campaigns, by definition, can victimize two groups, users and search engine providers. The former because they may be convinced to purchase goods or services that are of low quality or

¹This is commonly done using the User-Agent string which self-identifies popular crawlers, but some SEO kits also include IP address ranges they have associated with the major search engines.

illegal, the latter because their ability to provide high quality search results is imperiled. However, within the niche of counterfeit luxury goods another potential victim is the luxury brands themselves (both in terms of lost potential sales and brand damage). Consequently, in addition to interventions from search engines (driven by general concerns about search quality), brands also drive interventions to protect their economic interests. In this section we discuss what makes this market distinct, both in terms of how counterfeit luxury SEO campaign are structured and the kinds of interventions used in response.

3.1 SEO Campaigns

The SEO campaigns funded by the counterfeit luxury goods market operate similarly to other SEO campaigns (see [36] for one such example), with a couple of noteworthy differences. First, they introduce distinct cloaking and evasion techniques designed to undermine existing defenses. Second, the businesses that ultimately fund these campaigns appear to be organized differently than the open affiliate marketing programs that have been endemic in prior studies of underground economies (e.g., counterfeit pharmaceuticals [25], software [24] or FakeAV [34]). We discuss each of these in turn.

3.1.1 Cloaking

At its essence, cloaking refers to any mechanism for delivering different content to different user segments. For the purposes of SEO, cloaking’s primary objective is to deceive search engines by providing different content to the search engine crawler than to users clicking on search results. For example, the most widely-used cloaking technique, called *redirect cloaking*, arranges that search engine crawlers (e.g., Googlebot) receive content crafted to rank well for targeted query terms, while normal users who access the site are instead redirected to another site hosting a particular scam (e.g., a storefront selling counterfeit goods). In some cases, particularly when the doorway is on a compromised site, a visitor will only be redirected after arriving via a search results page. Otherwise, the original legitimate site content is returned, enabling compromised sites to remain compromised longer by appearing unchanged to normal visitors.

However, cloaking is a violation of most search engine’s content guidelines and, when such activity is discovered, the cloaked sites are typically deranked automatically in search results. As with any adversarial process, though, attackers adapt to new defenses. In contrast to cloaking techniques we have previously observed [35, 36], we have identified a new method of cloaking, which we call *iframe cloaking*, which bypasses traditional means of detection. In particular, iframe cloaking does not redirect the user and frequently returns the same content to both search engines and users.² Instead of redirecting a user to a landing store site, the store is simply loaded within an iframe element on top of the existing doorway page content. Typically the iframe visually occupies the entire height and width of the browser to provide the illusion that the user is browsing the store (Figure 1 shows a simple example of iframe cloaking using JavaScript). The JavaScript implementation is frequently obfuscated to further complicate analysis and in some cases the iframe itself is dynamically generated. Taken together, these countermeasures require any detection mechanism to run a com-

²A complementary feature of iframe cloaking is that it reduces the requirements for cloaking on compromised sites. Traditional cloaking uses network features (e.g., IP address or user agent) to identify crawlers, requiring specialized server side code. In contrast iframe cloaking runs entirely on the client, relying on the assumption that crawlers do not fully render pages at scale.

Cheap Louis Vuitton Outlet Online, 2014 New arrival Free ...

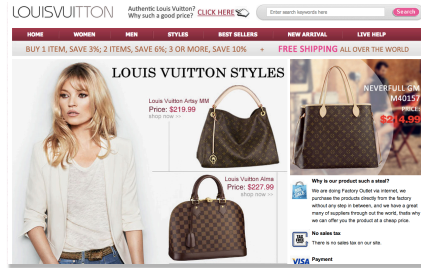
Order Louis Vuitton Outlet No taxes, louis vuitton have become focus in world, Louis Vuitton Outlet, 2013 New arrival Free Shipping, standing on the station early ...

Louis Vuitton Outlet Store: Louis Vuitton Handbags On S...


Welcome to Louis Vuitton Handbags Outlet online store, Choose the cheap Louis Vuitton Handbags save 70%, all the Louis Vuitton bags 100% free shipping.

LVMH: world leader in high-quality products, prestigious ...

www.lvmh.com/ LVMH Moët Hennessy Louis Vuitton. The Group - LVMH Group - Group mission and values - Key figures - LVMH companies and brands - Wines & Spirits - Fashion & Leather Goods ...



http://anonymized



```
..... http://anonymized .....
<title>Louis Vuitton Outlet Store: Louis Vuitton Handbags On Sale ,Real Louis
Price 80% Off. Sweet Clara's</title>
<meta name="description" content="Welcome to Louis Vuitton Handbags Outlet onl
Vuitton Handbags save 70% , all the Louis Vuitton bags 100% free shipping. Loui
best seller of Louis vuitton outlet."/>
<meta name="keywords" content="Louis Vuitton Outlet outlet,Louis Vuitton Outl
</head>
<body>
<iframe id="iframe" frameborder="0" scrolling="no" height="1760" width="100%"
<style type="text/css">
#iframe(position: absolute;top: 0px;left: 0px;z-index: 1000;ba
</style>
<script type="text/javascript" src="index_files/disclaim-element.js"></script>
<script type="text/javascript" src="index_files/graph-calc.js"></script>
<script type="text/javascript" src="index_files/jquery.js"></script>
```

Figure 1: An example of iframe cloaking where the same URL returns different content for different visitor types. Above, a user clicks through a search result and loads a counterfeit Louis Vuitton store. While, below, a search engine crawler visits the same URL directly, receiving a keyword-stuffed page because the crawler does not render the page. Our crawlers mimic both types of visits.

plete browser that evaluates JavaScript and fully renders a page (a set of requirements that greatly increase the overhead of detection at scale).³ We found the use of iframe cloaking to be pervasive within the domain of counterfeit luxury, but a more comprehensive study of the use of iframe cloaking for other domains remains an open question.

3.1.2 Business structure

Traditionally, a broad range of online scams have been organized around an affiliate marketing model in which an affiliate program is responsible for creating site content, payment processing and fulfillment, while individual affiliates are responsible for delivering the user to storefronts (e.g., via email spam, SEO, etc.). Core to this business model is the notion that affiliates are independent contractors agents paid on a commission basis, and thus affiliate programs work to attract a diverse set of affiliates. This model is commonly used today in a broad range of scams with a nexus in Eastern Europe and Russia including pharmaceuticals, pirated software, books, music and movies, herbal supplements, e-cigarettes, term paper writing, fake anti-virus and so on [32].

However, there are many indications (albeit anecdotal) that the structure of organizations in the counterfeit luxury market are distinct.⁴ First, the marketing portion of these scams can span both an array of brands and types of merchandise. For example, from infiltrating their command and control (C&C) infrastructure using the same approach as described in previous work [36], we find a single SEO campaign may shill for over ninety distinct storefronts selling thirty distinct brands ranging from apparel (Abercrombie), luxury handbags (Louis Vuitton), and electronics (Beats By Dre). Moreover, the same campaign will commonly host lo-

³Even after rendering a page, the ubiquity of iframes in online advertising make distinguishing benign from malicious content a challenge.

⁴There is a range of evidence suggesting that the big counterfeit luxury organizations have a nexus in Asia, unlike the Eastern European origin of many other scams. Our evidence includes the use of Asian language comments in SEO kit source code, the choice of Asian payment processors, fulfillment and order tracking from Asia and direct experience interviewing an Asian programmer working for one of these organizations. We surmise that a distinct cyber-crime ecosystem has evolved separately in East Asia with its own standard practices and behaviors.

calized sites catering to international markets (e.g., United Kingdom, Germany, Japan, and so forth). Unlike other kinds of counterfeit sales, which centralize payment processing within the affiliate program [16, 25, 34], we find each counterfeit luxury storefront allocates order numbers independently and engages directly with payment processors (merchant identifiers exposed directly in the HTML source on storefront pages allowed us to confirm). Finally, in the traditional affiliate program model, fulfillment is managed internally by the program, but in our investigations we have found at least one fulfillment site for luxury goods that appears to be designed to support outside sales on an *à la carte* basis (i.e., the site is designed to support wholesale ordering and allows each member to track the order status of their customer’s shipments). Overall, we suspect the counterfeit luxury ecosystem does not use an affiliate program and instead the ecosystem is composed of several independent advertisers (SEO campaigns) contracting with third parties for fulfillment and payment processing.

3.2 Interventions

As we have observed, the two groups with natural incentives to disrupt SEO campaigns targeting counterfeit luxury goods are search engine providers and luxury brand holders. Search engines maintain the value of their page views (and hence the pricing they can charge for advertising) by providing consistently high quality search results for their users. Thus, all major search engines have active anti-abuse teams that try to reduce the amount of search spam appearing in their results. When an SEO campaign is detected, search engines attempt to disrupt the campaign by either demoting their doorway pages in search results or even removing those pages from the search index entirely. Brand holders have a far less privileged technical position and they are neither able to analyze the Web at scale nor directly influence search results. However, as brand and trademark holders they have unique legal powers that allow them to target particular pieces of infrastructure from SEO campaigns. These two techniques, search and seizure, represent the de facto standard methods of intervention against SEO campaigns, with pressure applied at different strata in the business model.

3.2.1 Search Engine

In addition to allowing the search ranking algorithm to demote doorways performing black hat SEO, search engines commonly

have special handling for certain classes of malicious content. For example, starting in 2008, Google’s Safe Browsing service (GSB) has detected and blacklisted sites leading to malware or phishing sites with the aim of preventing users from being defrauded through search. GSB labels search results leading to malware or phishing pages as *malicious*, appends the subtitle “This site may harm your computer” to the result, and prevents the user from visiting the site directly by loading an interstitial page rather than the page linked to by the result.

In 2010, Google instituted a similar effort to detect compromised Web sites and label them as *hacked* by similarly appending the subtitle “This site may be hacked” in the result [9]. The motivation is to curb the ability of compromised sites to reach unsuspecting users, while simultaneously creating an incentive for innocent site owners to discover their site has been compromised and clean it. In principle, this notification could undermine black hat SEO since users may be wary to click on links with a warning label.

However, there are important differences between these two seemingly similar efforts, which more likely reflect policy decisions rather than technical limitations. First, contrary to malicious search results, users can still click through hacked search results without an interstitial page. Second, typically only the root of a site is labeled as hacked; e.g., while `http://anonymized` may be labeled as a hacked site, `http://anonymized/customize.php` will not. Unfortunately, often only the non-root search results are compromised and redirect users, while the root search result is clean. In Section 5.2 we examine the implications and limitations of these policy decisions on search interventions against SEO campaigns.

3.2.2 Seizure

As the name suggests, seizures reflect the use of a legal process to obtain control of an infringing site (typically by seizing their domain name, but occasionally by seizing control of servers themselves) and either shut it down or, more commonly, replace it with a seizure notification page. In the context of counterfeit luxury, seizures prevent users from visiting seized domains, thereby hindering the store’s ability to monetize traffic. Although we have witnessed brand holders performing seizures directly, typically they contract with third party legal counsel or with companies who specialize in brand protection, such as MarkMonitor [23], OpSec Security [28] and Safenames [31], to police their brand.

However, there are significant asymmetries in this approach. For example, a new domain can be purchased for a few dollars, but the cost to serve a legal process to seize it can cost 50–100 times more. Similarly, while a new domain name can be allocated within a few minutes and effectively SEO’ed in 24 hours [36], a seizure first requires finding the site, filing a legal claim and then waiting (from days to weeks) for the docket to be picked up by the federal judge to whom the case has been assigned. Presumably to amortize these costs, a manual review of court documents shows that domain name seizures commonly occur in bulk (hundreds or thousands at a time) and are not performed on a reactive basis. Finally, it is worth noting that doorway sites based on compromised Web servers present their own challenges since seizing the domain of an innocent third party can carry liability. Thus, while brands sometimes seize doorway pages, it is more common for them to target the storefront advertised. Section 5.3 explores these asymmetries in greater depth.

4. DATA SETS

The basis of our study relies upon extensive crawls of Google search results to discover poisoned search results that lead to counterfeit storefront sites. We then use a combination of manual label-

Vertical	# PSRs	# Doorways	# Stores	# Campaigns
Abercrombie	117,319	2,059	786	35
Adidas	102,694	1,275	462	22
Beats by Dre	342,674	2,425	506	16
Clarisonic	10,726	243	148	6
Ed Hardy*	99,167	1,828	648	31
Golf	11,257	679	318	20
Isabel Marant	153,927	2,356	1,150	35
Louis Vuitton*	523,368	5,462	1,246	34
Moncler	454,671	3,566	912	38
Nike	180,953	3,521	1,141	32
Ralph Lauren	74,893	1,276	648	27
Sunglasses	93,928	3,585	1,269	34
Tiffany	37,054	1,015	432	22
Uggs*	405,518	4,966	1,015	39
Watches	109,016	3,615	1,470	35
Woolrich	55,879	1,924	888	38
Total	2,773,044	27,008	7,484	52

Table 1: A breakdown of the verticals monitored highlighting the number of poisoned search results, doorways, stores, and campaigns identified throughout the course of the study. Note that the KEY campaign targeted all verticals except those with an ‘*’.

ing and supervised learning to map storefront sites into the different SEO campaigns that promote them. On a subset of storefront sites, we also use a combination of test orders and actual purchases to reveal information about customer order volume and payment processing. Finally, we crawl the site of a supplier to provide insight into the scale of order fulfillment and high-level customer demographics. This section describes each of these efforts in detail.

4.1 Google Search Results

In general, it can be challenging to unambiguously determine if the site pointed to by a given search result is a counterfeit storefront or represents a legitimate reseller. To address this issue we focus our analysis on “cloaked” sites (i.e., that present different results to different visitors) using a crawler we developed previously for detecting such sites [35]. This approach largely removes the problem of false positives (i.e., legitimate sites advertising brands do not cloak) and in our experience most counterfeit landing pages do make use of cloaking. Finally, while this definition will cause us to miss those counterfeit sites that do not use cloaking, our primary goal is to identify how various interventions impact sites over time, and thus any bias here is only important if non-cloaking counterfeit sites are both large in number and respond substantively differently to interventions. We do not believe either to be true. Given these assumptions, each day we issue queries to Google using search terms targeted by counterfeit sites, crawl the sites listed in the search results, and identify sites using cloaking as depicted in Figure 1. We repeat this process for five months from November 13, 2013 through July 15, 2014.

In the rest of this section we define the notion of counterfeit luxury *verticals* for organizing search queries, and describe our methodology for selecting the search terms that comprise the verticals, the implementation of our crawlers and the information they collect, and our heuristics for detecting counterfeit stores in poisoned search results.

Note that we search exclusively using Google for a couple of reasons. In prior work we found that Google is the most heavily targeted search engine by attackers performing search poisoning and black hat SEO [35]. Furthermore, Google is the leading search engine for the United States and many European countries, the pre-eminent markets receiving counterfeit products (based on shipping data from a large supplier as discussed in Section 4.5).

4.1.1 Search Terms

Any work measuring search results is biased towards the search terms selected because the selected terms represent just a subset of the entire search index. In our study, we monitor search results for counterfeit luxury *verticals*, a set of search terms centered around a single brand (e.g., Ralph Lauren) or a category composite of several brands (e.g., Sunglasses is a composite of Oakley, Ray-Ban, Christian Dior, etc.). For our study, each vertical consists of a static set of 100 representative terms that we determined were targeted by SEO campaigns.

Due to the early prominence of the KEY campaign, a large SEO botnet responsible for most of the PSRs manually observed in September 2013, we initially compiled terms for each of the 13 verticals it targeted as listed in Table 1. Similarly, we followed the KEY campaign’s approach in determining whether to center a vertical’s terms around a single brand or a composite. To select these terms we extracted keywords from the URLs of the doorway pages of the KEY campaign. For a given vertical, we manually queried Google to find ten KEY doorways redirecting to the same store selling counterfeit merchandise (related to the vertical). Then we issued site queries (e.g., “site:doorway.com”) for each doorway to collect all search results originating from the doorway. And for each search result we extracted search terms from the URL path (e.g., “cheap beats by dre” from `http://doorway.com/?key=cheap+beats+by+dre`) to assemble a large collection of terms. We then randomly selected 100 unique terms as a representative set for each vertical.

To extend the scope of our study to other campaigns, we included three additional verticals that we saw counterfeiters targeting: Ed Hardy, Louis Vuitton, and Uggs. Since the KEY campaign does not target these brands, we adopted a different approach in selecting search terms by using Google Suggest, a keyword autocomplete service. We first fetched suggestions for a targeted brand (e.g., “Louis Vuitton wallet”). Then we recursively fetched suggestions for the suggestions. In addition, we fetched suggestions for the concatenation of a commonly used adjective (e.g., cheap, new, online, outlet, sale or store) and the brand name to form search strings (e.g., “cheap Louis Vuitton”). From the combined set of these various search strings, we randomly selected 100 unique strings as our search set for each vertical.

To evaluate any bias introduced from these two different approaches, we take the ten original KEY verticals that are not composites, generate alternate search terms using the Google Suggest approach, and run the crawlers using those alternate terms for one day on April 25, 2014. Among the ten verticals, we find four out of a thousand total terms overlap. Additionally, when comparing the percentage of PSRs detected after crawling, for both classified and unknown, and the distribution of PSRs associated to specific campaigns, we find no significant difference between results from the original and alternate terms over the same time range. Despite using two different approaches for selecting search terms, in the end we find the same campaigns poisoning search results. This overlap highlights both the pervasiveness of these campaigns and the representativeness of terms selected in spite of the KEY campaign’s early influence on our methodology.

4.1.2 Crawling Search Results

For each search term, we query Google daily for the top 100 search results. For each search result, we crawl each page link using an updated version of the Dagger cloaking detection system from previous work [35]. Dagger uses heuristics to detect cloaking by examining semantic differences between versions of the same page fetched first as a user and then as a search engine crawler (distinguished by the User-Agent field in the HTTP request).

A previous limitation of Dagger was that it did not render the page and, as a consequence, did not follow JavaScript (JS) redirects. Thus, we extended Dagger by rendering each cloaked search result detected using HtmlUnit [13], essentially a headless browser complete with a JavaScript interpreter. (Since rendering a page is an expensive operation, we only render pages we detect as cloaked.)

To detect iframe cloaking (Section 3.1.1), we implemented a second crawler, VanGogh. VanGogh also uses HtmlUnit to render pages. To detect iframe cloaking, it identifies any iframes attempting to occupy the entire page visually (hiding the original content). Specifically, we classify pages as using iframe cloaking if they load iframes where the height and width attributes are both either set to 100% or larger than 800 pixels.

And again, due to the high overhead of rendering pages, we only crawl a subset of search results using VanGogh. In particular, for each measurement we crawl at most three randomly selected pages from the same doorway domain to reduce the crawling workload. We further trim the workload by not crawling domains previously seen and not detected as poisoned by either VanGogh or Dagger. This approach has proved reasonable due to the low daily churn in search results for each vertical (on average 1.84% newly seen domains are found in search results each day).

4.1.3 Store Detection

Ultimately we want to identify counterfeit luxury storefronts advertised through PSRs. We detect stores by applying two heuristics to the set of PSRs discovered from crawling. First, we inspect cookies from each landing site (the page eventually loaded in a user’s browser after redirection through the doorway page) to look for cookies commonly used by counterfeit luxury storefronts such as those related to payment processing (e.g., Realex, Mallpayment), e-commerce (e.g., Zen Cart, Magento), and Web analytics (e.g., Ajstat, CNZZ). Second, we search for either of the substrings “cart” or “checkout” on the landing pages. If either of the heuristics succeed, we treat the landing site as a counterfeit luxury store advertised through search poisoning. Note that this approach identifies stores from the search results within a vertical irrespective of brand. For example, we may identify a counterfeit Christian Louboutin store within Louis Vuitton search results.

We validate our detection methodology by manually inspecting sampled search results from three popular verticals, Beats By Dre, Isabel, and Louis Vuitton. For each vertical, we randomly chose three search terms, and compared the search results for those terms from two measurements taken at least two months apart (e.g., one from November 23, 2013 and one from February 24, 2014). In total we examined 1.8K search results and detected 532 storefronts advertised using cloaked search results. Among these we found no false positives (instances where a benign page is mistakenly labeled as a doorway to a storefront) and 21 (1.2%) false negatives (instances where a doorway to a storefront is not labeled). These results are reassuring because errors are likely skewed towards underrepresenting the number of storefronts.

4.2 Campaign Identification

Our targeted crawls of Google search results produce a large collection of doorway pages and counterfeit storefronts. We know that behind these thousands of doorways and storefronts lurk a much smaller number of distinct SEO campaigns, and the goal of our work is to understand the full ecosystem of campaigns operating in this counterfeit luxury market rather than focus on a singular campaign, e.g., the KEY campaign.

A brute-force approach to this understanding would require a domain expert to examine each Web page in our collection and

use domain-specific heuristics to infer the SEO campaign behind it. The manual labeling of Web pages, however, is a time-consuming and laborious endeavor that does not scale well to the many thousands of examples in our collection. Instead we take a statistical approach, and the rest of this section describes an automated, data-driven method to identify the SEO campaigns behind individual doorway and storefront Web pages.

To build a statistical model, we need a data set of *labeled examples*. Though manual labeling is tedious, we created such a data set by identifying the SEO campaigns behind a small subset of 491 Web pages in our much larger collection of crawled results. From this small data set, we learned a classifier that mapped the remaining thousands of doorway and storefront Web pages to the 52 SEO campaigns for which we had manually labeled examples. The results of this analysis (discussed in later sections) provide a comprehensive understanding of the ecosystem of SEO campaigns in the counterfeit luxury market.

Our classifier makes its predictions by extracting textual features of HTML content and analyzing the statistics of these features that distinguish Web pages from different campaigns. The following subsections describe the classifier in more detail, focusing in particular on the individual stages of feature extraction, model estimation, and model validation.

4.2.1 Feature Extraction

The premise of our statistical approach is that doorway and storefront Web pages contain predictive signatures of the SEO campaigns behind them. Motivated by previous work [2], we looked for these signatures in their HTML source. We expect HTML-based features to be predictive in this domain for two reasons: first, because SEO campaigns use highly specialized strategies to manipulate the search rankings of doorways [36], and second, because campaigns often develop in-house templates for the large-scale deployment of online storefronts (e.g., customized templates for Zen Cart or Magento providing a certain look and feel).

To extract HTML features, we follow a conventional “bag-of-words” approach. In particular, we construct a dictionary of all terms that appear in the HTML source code, and for each Web page, we count the number of times that each term appears. In this way, each Web page is represented as a sparse, high-dimensional vector of feature counts. We implemented a custom bag-of-words feature extractor based on tag-attribute-value triplets [5] for the Web pages in our data set.

One might also expect to find predictive signatures of SEO campaigns in network-based features (e.g., IPv4 address blocks, ASes). However, we found that such features were ill-suited to differentiate SEO campaigns due to the growing popularity of shared hosting and reverse proxying infrastructure (e.g., CloudFlare). Therefore, after a brief period of experimentation, we did not pursue the use of such features.

4.2.2 Model Estimation

We learned linear models of classification from our data set of labeled examples. Specifically, we used the LIBLINEAR package [7] to learn L1-regularized models of logistic regression. The L1-regularization encourages sparse linear models in which the predictions of SEO campaigns are derived from only a handful of HTML features. Thus the resulting models are highly interpretable: for each campaign, the regularization serves to identify the most strongly characteristic HTML features from the tens of thousands of extracted ones.

We evaluated the predictive accuracy of the classifier by performing 10-fold cross-validation on the data set of labeled exam-

ples. The average accuracy on held-out data was 86.8% for multi-way classification of Web pages into 52 different SEO campaigns. (Note that uniformly random predictions would have an accuracy of $1/52 = 1.9\%$.) Our model’s high accuracy on held-on examples gave us confidence to classify the remaining (unlabeled) Web pages that we collected from poisoned search results.

4.2.3 Model Validation and Refinement

We used the above models, trained on a small subset of labeled Web pages, to infer the SEO campaigns behind the remaining unlabeled Web pages. To do so, we extracted HTML features from the unlabeled Web pages and used the classifiers to predict the most likely campaign behind each example. To validate these predictions, we manually inspected additional subsets of unlabeled examples. This step can be done efficiently by first validating the top-ranked predictions for each SEO campaign (as reflected by the probabilities that the logistic regressions attach to each prediction).

We briefly describe how we validated the classifier’s predictions on unlabeled Web pages. Primarily we assume that distinct SEO campaigns are unlikely to share certain infrastructure such as SEO doorway pages and C&Cs, payment processing, and customer support. We also consider less robust indicators such as unique templates, WHOIS registrant, image hosting, and Web traffic analytics (e.g., 51.la, cnzz.com, statcounter.com, etc.).

A final stage is to refine the model, using the manually verified predictions to expand the set of labeled Web pages, retraining the classifier on this expanded set, and repeating this process in rounds. With each iteration of this process we obtain a more accurate classifier and also one with greater coverage of distinct SEO campaigns. Though some manual labeling is unavoidable, this overall approach (of repeated human-machine interaction) is far more efficient than a brute-force expert analysis.

4.3 Purchases

From previous work studying email spam advertising illicit pharmaceutical and software storefronts [16, 20], we found that making orders on sites can shed light on normally opaque facets of underground businesses: order volume, payment processing, and order fulfillment. This information serves two important roles. First, it reveals the interplay between the various actors in the counterfeit luxury ecosystem (SEO campaigns, payment processors, and suppliers). Second, the estimated order volume serves as a vital metric in measuring the effectiveness of interventions (e.g., does labeling doorway search results as “hacked” lead to lower campaign order volume?). Similar to this prior work, we created test orders on counterfeit stores to estimate their order volume over time, and made actual purchases to reveal the payment processors used by these storefronts and the quality of the merchandise they sell.

4.3.1 Order Volume

We use the “purchase pair” technique [16] to estimate the order volume of individual stores over time. This technique exploits the fact that stores use monotonically increasing order numbers, where the difference between order numbers represents the total number of orders created over the time delta between orders.

Note that stores give users an order number before processing their credit cards, and users still have an opportunity to back out. Therefore this metric represents an upper bound on orders placed and overestimates the absolute number of orders at a given store. However, it is still useful to quantify the rate at which orders are created, as well as the changes in order rate over time when correlated with interventions.

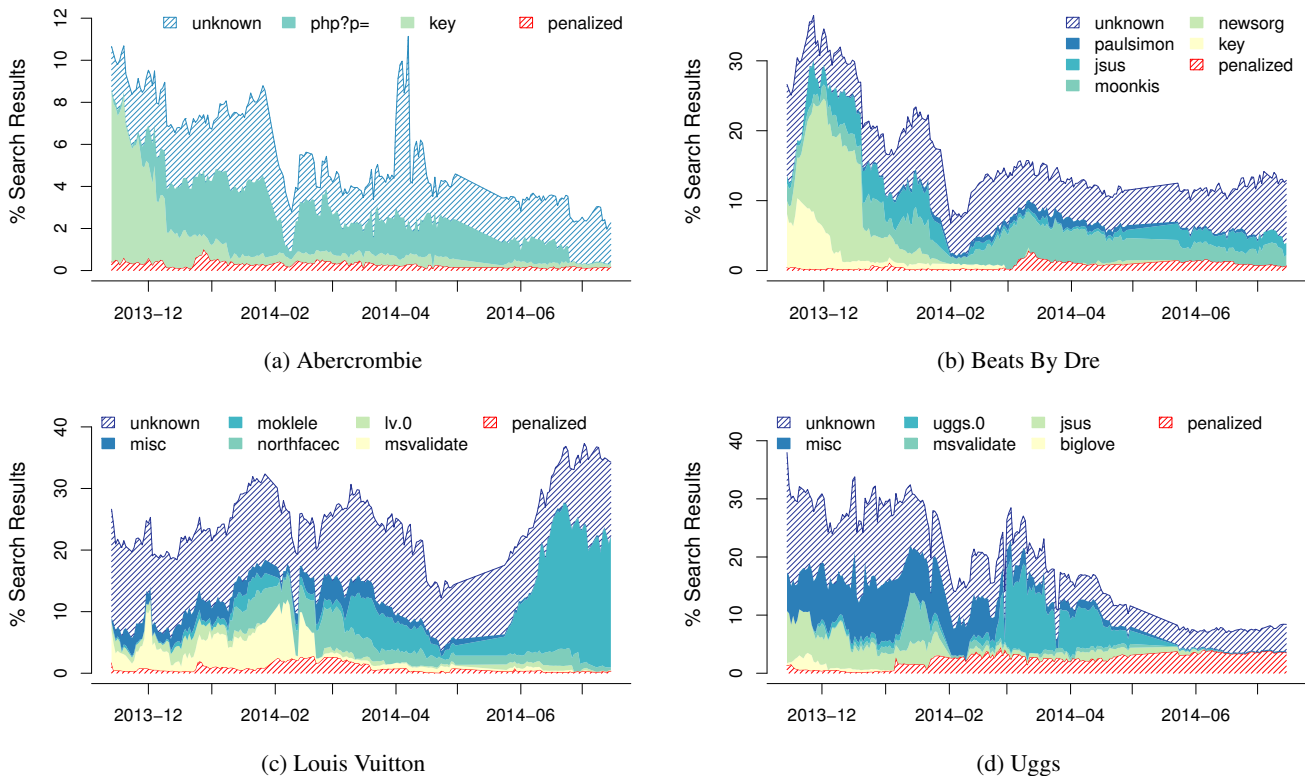


Figure 2: Stacked area plots attributing PSRs to specific SEO campaigns within the labeled vertical. The red area represents the percentage penalized, either through search or seizure. The remainder of the areas represents active PSRs, where the filled areas are attributed to specific campaigns and the unfilled area is the remainder unclassified.

Using this technique, we created 1,408 orders from 290 stores touching 24 distinct campaigns and 13 verticals, between November 29, 2013 and July 15, 2014. We created 343 orders by hand and 1,065 orders using scripts. Operationally, for both manual and automated orders, we visit each store using via TOR and create orders at weekly intervals, and we limit orders to three per day per campaign to reduce the chance of being detected by the store or payment processor. We take the orders all the way to the payment processing page, which requires credit card details, before finally leaving the site. The order and customer information we provide are semantically consistent with real customers, but fictional and automatically generated [6].

4.3.2 Transactions

To shed light on payment processing and order fulfillment in the counterfeit luxury ecosystem, we successfully placed product orders from 16 unique stores covering 12 different campaigns. In total, we received 12 knock offs of low to medium quality, all shipped from China. From the bank identification numbers (BINs) in our transactions, we found that our purchases were processed through three banks (two in China, one in Korea). This concentration suggests payment processing is another viable area for interventions as in [24], but investigating such an intervention remains future work.

4.4 User Traffic

As surveyed in Section 2 and described in depth in our previous study [36], SEO campaigns poison search results to acquire user traffic that can then be monetized through scams — in this case, counterfeit luxury stores. The order volume data shows that counterfeit luxury stores do successfully convert user traffic into

sales, and indirectly measures an SEO campaign’s effectiveness in attracting traffic via PSRs.

For a small number of stores, we were also able to collect user traffic data that directly measures an SEO campaign’s effectiveness in attracting customers to their stores. Specifically, we were able to periodically collect AWStats data for 647 storefronts in 12 campaigns. AWStats is a Web analytics tool [1] that uses a Web site’s server logs to report aggregated visitor information (e.g., the number of visitors, visitor durations, visitor geolocations, referrers of visitors, etc.). From our crawled data, we discovered that these stores left their AWStats pages publicly accessible, and we were able to fetch visitor data for each store by visiting the publicly accessible default AWStats URL (e.g., <http://<site>/awstats/awstats.pl?config=<site>>).

4.5 Supply Side Shipments

To better understand the suppliers, customers, and operational relationship among storefronts and suppliers, we collected longitudinal shipment data from a supplier partnering with MSVALIDATE, one of largest SEO campaigns peddling counterfeit Louis Vuitton.

We discovered the supplier site from the packing slip of two of our purchases. Upon visiting the site, we noticed it contains a scrolling list of fulfilled orders and a mechanism to lookup shipping records for valid order numbers in bulk (20 orders at a time). Each record contains a timestamp and information regarding current location and delivery status.

Using this mechanism, we collected over 279K shipping records for nine months of orders placed through the supplier between July 5, 2013 and March 28, 2014. In summary, 256K orders successfully reached their destination, 4K were seized at the source (China), 15K were seized at the destination, and of the delivered, 1,319 were

returned by the customer. From country data listed in the records, the three largest destinations are the United States, Japan, and Australia, with 90k, 57K, and 39K orders, respectively. If we combine these with the countries from Western Europe (41K), these regions account for over 81% of orders.

5. RESULTS

In this section we use our crawler data to characterize the activities of SEO campaigns that use search to promote stores selling counterfeit luxury goods, and we further use our order data to study the effects of both search engine and domain seizure interventions on these SEO activities. In short, we find instances where both can have the desired effect of disrupting counterfeit sites, but they need to be far more reactive in time and comprehensive in coverage to undermine the entire ecosystem of SEO campaigns exploiting search engines for customers.

5.1 Ecosystem

We start with classifying poisoned search results into campaigns, how those campaigns target verticals, and what the PSRs reveal about the operations of the campaigns.

In terms of raw data, we crawled search results for eight months from November 13, 2013 through July 15, 2014 and detected 2.7M PSRs, across all verticals, using 27K doorways from unique domains and sending users to 7,404 different stores selling counterfeit luxury merchandise. Applying the classifier described in Section 4.2 to this data, we identified 52 distinct SEO campaigns that account for 828 stores, 11K doorway domains, and 1.6M PSRs. Table 2 lists the campaigns using a name we derived from a pattern in their URLs, the domain names used for their C&C, or some other telltale aspect of their operation. For each campaign, the table lists the number of doorway domains, storefronts, and brands targeted. Note that, although we ascribed more than half (58%) of all PSRs to their respective campaigns, these PSRs only account for 11% of all stores. This disparity suggests that the ecosystem has a skewed distribution where a handful of large campaigns account for the majority share of PSRs that redirect users to a concentrated set of storefronts.

From the perspective of brands, we attributed 16% to 69% of PSRs in each vertical to known campaigns. Figure 2 visualizes our classification results for four verticals: Abercrombie (64.2% of PSRs classified to campaigns), Beats By Dre (62.2%), Louis Vuitton (66%), and Uggs (58%). We chose these verticals for their diversity in merchandise, campaigns, and search term selection methodology. For each vertical, the filled areas in the stacked area plots show the fraction of search results poisoned by the major campaigns targeting the vertical; note that the “misc” label collapses multiple campaigns into a single category to reduce clutter.

Each graph presents the PSRs detected, classified, and penalized over time at the granularity of a day. For example, in Figure 2b on December 1, 2013, 34.6% of search results for the Beats By Dre vertical were poisoned. Of these PSRs, 85.3% redirect users to counterfeit stores operated by five campaigns: KEY (16.8%), NEWSORG (53.8%), MOONKIS (5.8%), JSUS (8.0%), and PAULSIMON (0.3%). The remaining 14.7% PSRs redirect users to counterfeit stores we have yet to classify. The bottom shaded area shows that just 0.6% are penalized either through Google labeling the search result as “hacked” (Section 5.2.2) or a brand has seized the storefront domain name (Section 5.3).

5.1.1 Verticals

Figure 3 shows the percentage of search results that were poisoned for each vertical as pairs of sparklines. Each sparkline is a

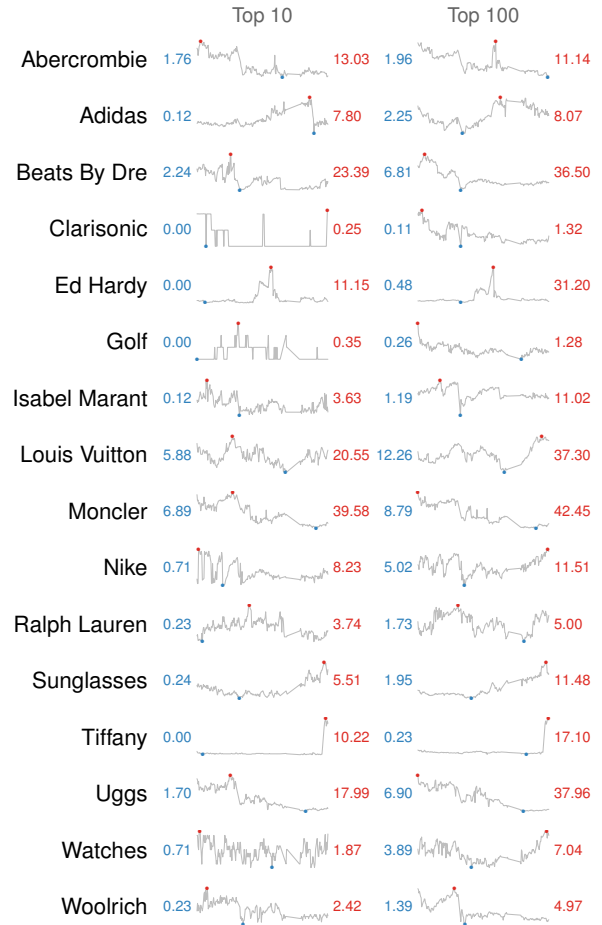


Figure 3: Percentage of search results poisoned for each brand vertical, shown as sparklines. Each sparkline is a daily time series showing relative values over five months. The left number is the minimum value across time and the right is the maximum (also shown as dots on the line).

time series showing relative values over the five-month time span of the study at the granularity of days. The left number is the minimum value across time and the right is the maximum (also shown as dots on the line). Each column of lines shows the percentage of PSRs among top 10 search results (left) or top 100 (right). For example, in the Abercrombie vertical in the top left, at most 13% of the top 10 search results in the vertical were poisoned, while at least 2% were poisoned. The sparkline shows that the first three months were closer to the 13%, while the latter five months were much lower.

Overall, heavily targeted verticals are particularly vulnerable to poisoned search results. In 13 out of 16 verticals, about 5% of search results are poisoned at some point in time. But for the five verticals most vulnerable to search poisoning, at different points in time 31–42% of the top 100 search results in those verticals were poisoned. Also, as expected, it is easier in general to poison search results from outside the top 10; Beats By Dre, for instance, had at most 23% of its top 10 results poisoned while at one point 37% of its top 100 results were poisoned.

Brands face multiple “adversaries”. Whether targeted by many campaigns (14 and 17 for Louis Vuitton and Uggs, respectively) or just a few (three and six for Abercrombie and Beats By Dre), all

Campaign	# Doorways	# Stores	# Brands	Peak
171760	30	14	7	44
ADFLYID	100	18	4	66
BIGLOVE	767	92	30	92
BITLY	190	40	15	89
CAMPAIGN.02	26	4	3	61
CAMPAIGN.10	94	18	5	99
CAMPAIGN.12	118	5	1	59
CAMPAIGN.14	39	8	2	67
CAMPAIGN.15	364	10	10	8
CAMPAIGN.17	61	8	3	44
CHANEL.1	50	10	4	24
G2GMART	916	28	3	53
HACKEDLIVEZILLA	43	49	9	56
IFRAMEINJS	200	2	1	39
JAROKRAFKA	266	55	3	87
JSUS	439	59	27	68
KEY	1,980	97	28	65
LIVEZILLA	420	33	16	70
LV.0	42	3	1	62
LV.1	270	12	9	90
M10	581	35	8	30
MOKLELE	982	15	4	36
MOONKIS	95	7	4	99
MSVALIDATE	530	98	6	52
NEWSORG	926	7	5	24
NORTHFACEC	432	2	1	60
NYNY	29	14	5	40
PAGERAND	122	7	4	43
PARTNER	62	9	5	33
PAULSIMON	328	33	12	128
PHP?P=	255	55	24	96
ROBERTPENNER	56	7	12	50
SCHEMA.ORG	46	17	7	54
SNOWFLASH	271	14	1	48
STYLESHEET	222	9	6	63
TIFFANY.0	26	1	1	4
UGGS.0	428	6	5	30
VERA	155	38	12	156

Table 2: Classified campaigns along with # doorways seen redirecting on behalf of a specific campaign, # stores monetizing traffic from the campaign, # brands whose trademarks are abused by the campaign, and # days of peak poisoning duration, for campaigns with 25+ doorways.

verticals are targeted by multiple campaigns all competing to SEO their doorway pages into search results to lure customers for their counterfeit goods.

5.1.2 SEO Campaigns

SEO campaigns employ considerable infrastructure to maintain their businesses. As shown in Table 2, SEO campaigns use hundreds to thousands of doorway sites to redirect users to dozens of storefronts (similar in scale to other abusive SEO botnets [36]). Interestingly, we do not find a strong correlation between the number of doorways and the campaign’s efficacy in poisoning search results. For example, as shown in Figure 2b, MOONKIS poisoned search results for Beats By Dre from the start of 2014 onwards with 95 doorways, while two larger campaigns, JSUS and NEWSORG, used 439 and 926 doorways, respectively, in the same time period.

The operators of the campaigns successfully SEO their doorways in concentrated time periods. Although we observe campaigns poisoning search results for multiple months, their SEO effectiveness varies over time as exemplified by the campaigns targeting Beats By Dre in Figure 2b. To capture this notion of bursty SEO behavior, we compute a “peak range” for each campaign defined as the shortest contiguous time span that includes 60% or more of all

PSRs from the campaign. For example, NEWSORG’s peak range lasts 24 days from November 23 to December 17, 2013, with a daily average of 1676 PSRs during this span. Table 2 summarizes the peak duration in terms of number of days for each campaign. Using this metric, we find campaigns run at their peak for 51.3 days on average.

The campaign operators also run a diversified business that gives them flexibility in the face of disruption. A single campaign, for example, will use its doorways to poison search results from multiple verticals simultaneously. For instance, the MSVALIDATE and BIGLOVE campaigns both successfully poison search results for Louis Vuitton (Figure 2c) and Uggs (Figure 2d). As a result, campaigns possess multiple revenue streams, giving them flexibility in the event a setback disrupts one revenue stream (e.g., domain seizures from one brand, problems with a supplier for Beats By Dre headphones, etc.). The campaign can adjust and continue monetizing traffic by simply reallocating resources towards stores selling counterfeit merchandise from other verticals.

Moreover, campaigns often operate multiple storefronts targeting the same vertical and selling the same merchandise. Sometimes the goal is to localize for a market, such as a Japanese Uggs storefront catering to Japanese consumers. More often, though, these redundant stores can serve as backups in the event of interventions, which we explore further in Section 5.3.

5.2 Search Engine Interventions

Since poisoned search results manipulate and degrade user experience, search engines have an incentive to identify and penalize PSRs used by the SEO campaigns that lead users to counterfeit sites. Two options available to search engines for reacting to PSRs are to demote them in search rank, and to add warning labels to search results to alert users before clicking through.

5.2.1 Search Result Demotion

Figure 4 shows the prevalence of poisoned search results for four SEO campaigns over time, and the corresponding order activity at storefront sites gleaned from creating test orders as described in Section 4.3.1. The bottom two rows of graphs show the number of PSRs per day for each SEO campaign: the lowest row focuses on PSRs in just the top 10 search results, and the row above focuses on the full top 100 search results. The dark portion of the bars at the bottoms of the graphs corresponds to PSRs labeled by Google as “hacked”.

The top two rows show the results of sampling order numbers from a handful of representative stores promoted by each campaign; the stores that are both visible in PSRs and have high order activity relative to other stores from within the same campaign. The top “Volume” row shows the actual samples over time for the handful of representative stores and reflects the combined cumulative volume of order numbers created (recall that these numbers are an upper bound of actual orders made by customers). As another way of looking at the same data, the lower “Rate” row shows the order data as a histogram: we bin the number of estimated orders per week, interpolating in regions where we lack samples. The number at the top of each graph is the maximum value across the time series.

In all four campaigns, we see correlation between PSR prevalence and order activity, which suggests that search penalizations can be effective. This is most evident with the KEY campaign. The rate of orders (slope of the line in the “volume” graph) decreases in mid-December, soon after its PSR activity drops precipitously. We do not know the actual cause of the drop in PSRs: whether the KEY campaign stopped actively performing SEO on its

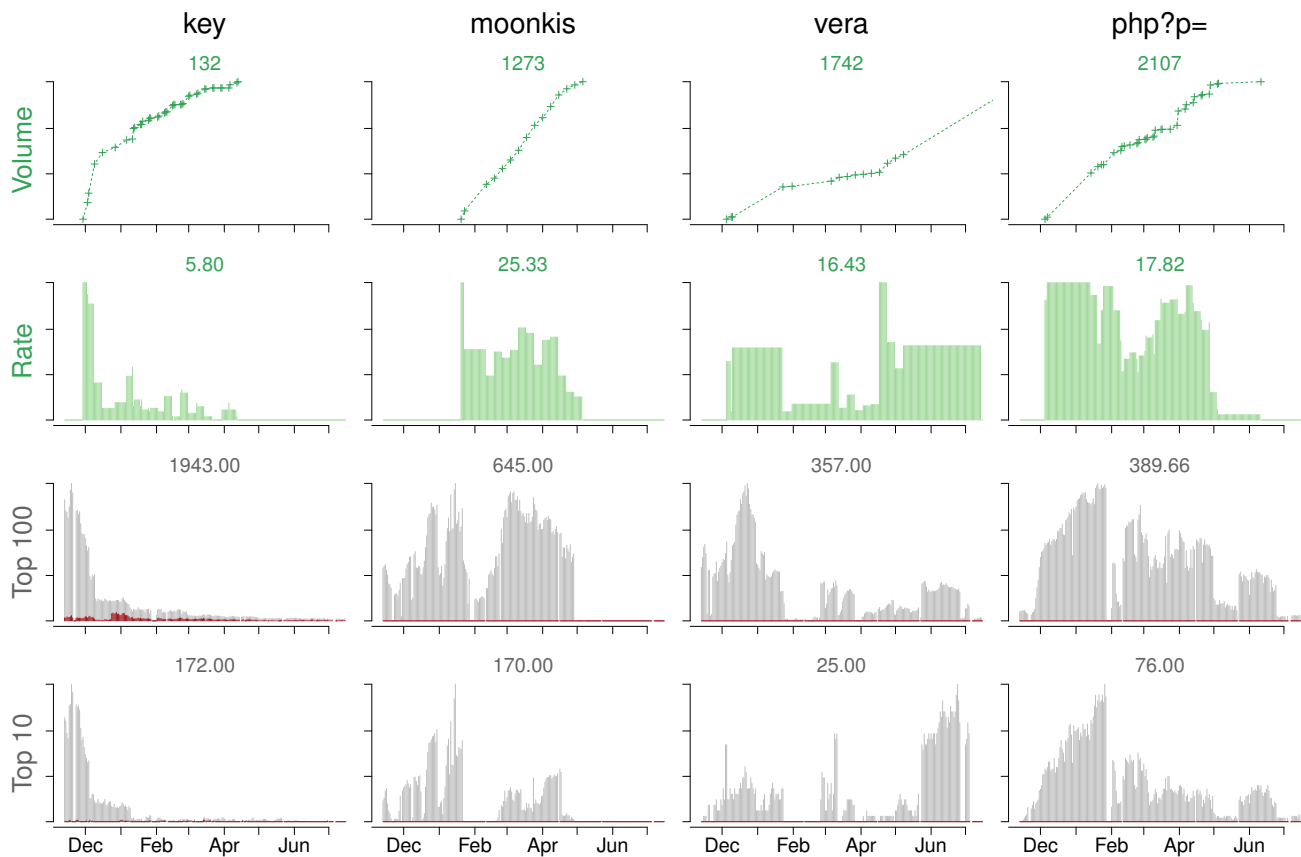


Figure 4: Correlation between a store’s visibility in PSRs and order activity for four SEO campaigns. Each column of graphs is associated with an SEO campaign. Bottom two rows of graphs depict the prevalence of PSRs among the top 100 and top 10 search results, respectively. Top two rows reveal cumulative changes in sampled order numbers, as well as histograms binning order number changes into extrapolated daily rates, respectively.

doorway sites, Google aggressively demoted its doorways in search result rank, or the “hacked” warning added to its PSRs dissuaded users from clicking on search results. However, it appears that the penalization pressure from Google—demoting most of the PSRs out of the top 100 and labeling half of the remaining as hacked—did have an effect. From attempts to actually purchase items, the stores promoted by the KEY campaign stopped processing orders after the drop in PSRs.

Penalizing PSRs by demoting them in search rank follows the conventional SEO wisdom that highly-ranked results are by far the most valuable. On this topic, the bottom two rows of Figure 4 show the prevalence of PSRs in the top 10 and top 100 results. For the most part, the shapes of both histograms are similar: campaigns are successful in SEOing poisoned search results throughout search page ranks, and it is difficult to conclude whether order volumes seen at stores are primarily due to the much smaller number in the top 10 or the much larger number across the top 100. One example, though, suggests that there is value in having PSRs across the full top 100. For the MOONKIS campaign, during most of March 2014 it had negligible PSRs in the top 10 but hundreds in the top 100. Nonetheless, order volumes seen at its stores remained high and steady. In this instance at least, search rank penalization would need to be even more aggressive to demote the PSRs from the top 100.

5.2.2 Warning Labels

Google uses the “hacked” label on search results to warn users about suspicious sites. This form of intervention faces two key challenges—coverage and reaction time—and, based upon our crawling results, overall appears to be ineffective for this type of abusive SEO activity.

Although most doorways are hacked sites, Google only penalizes 2.5% of the PSRs we crawled with a “hacked” label. From the perspective of brands, Figure 2 showed that penalized PSRs labeled with the “hacked” warning were a small fraction of all PSRs at any point in time for four large brand verticals. From the perspective of campaigns, Figure 4 shows a similar result: except for the KEY campaign, both the absolute number and fraction of penalized PSRs are quite small.

One issue that undermines coverage is that Google only labels the root of a Web site as “hacked”, and does not label search results that link to sub-pages within the same root domain. In the PSR data set, we found 68,193 “hacked” search results. When counting the number of PSRs that share the same root as a penalized site, Google could have labeled 102,104 search results (an additional 49%).

A second challenge is reaction time. A key metric of any reactive intervention is the time delay between when a campaign starts SEOing a doorway and when the search provider detects and penalizes the doorway with a label. This delta represents the window

Company	# Cases	# Brands	# Seized	# Stores	# Classified Stores	# Campaigns
Green, Burns, & Crain	69	17	31,819	214	40	17
SMGPA	47	11	8,056	76	20	12

Table 3: Summary of identified domain seizures initiated by brand holders found in our data set from Feb. 2012 – Jul. 2014. Note that these seizures are ones we identified from seizure notice pages in PSRs we crawled, and do not include other seizure activity not evident via PSRs (e.g., seized domains that did not use cloaked PSRs to redirect to them, seized domains that do not display notice pages, are shut down, etc.). From this PSR viewpoint, we aggregate the following per seizing entity: number of court cases initiating seizures (# Cases), number of brands protecting their trademarks through such cases (# Brands), and total number of store domains seized as reported in cases (# Seized). For overlap with the eight months of our crawled data set (Nov. 2013 – Jul. 2014), we also list the subset of store domains seized and directly observed in our crawled PSRs (# Stores), the number of those stores we classified into campaigns (# Classified Stores), and the number of SEO campaigns affected by seizures (# Campaigns).

of opportunity for an SEO campaign to monetize traffic obtained through PSRs without any warnings to users.

For doorways penalized with a label, campaigns have multiple weeks in which to monetize traffic through PSRs. Of the 1,282 “hacked” doorways in the PSRs data, 588 doorways were already labeled when we first saw them and we cannot determine when they were first labeled. The remaining 694 have lifetimes between 13–32 days on average until Google labeled them as “hacked”.

Note that the variance in the lifetime is due to the difficulty in determining exactly when Google penalizes a site. Using crawled search results, we know when we last saw a doorway prior to the penalty and when we first saw a doorway after the penalty. However we cannot always determine when the penalty occurred because it may be the case the doorway does not appear in our results for an extended period of time. As a result, we present two numbers, the smaller of which is the lifetime ending when we last saw the doorway actively redirecting, while the larger number is the lifetime when we first observed the labeling.

5.2.3 User Traffic

The correlation between search result visibility and order volume, observed in Figure 4, is an indirect measure of the ability of campaigns to attract and convert traffic via PSRs. Combining the AWStats data described in Section 4.4 with the crawled data and test purchases, we are able to examine this relationship in greater detail with a case study of a counterfeit Chanel store run by the BIGLOVE campaign that rotates across three storefront domains over time (*cocoviphandbags.com*, *cocovipbags.com*, and *cocoLovebags.com*).

As above, in Figure 5 we present both the prevalence of PSRs attributable to this store and the corresponding extrapolated order activity from June 10, 2014 to August 31, 2014. Using the AWStats data, in the bottom-most graph we also present the daily user traffic seen by the store in terms of the number of HTML pages fetched by users each day. We use color gradients in the PSRs and traffic graphs of Figure 5 to distinguish separate instances of *coco*.com*, where each instance represents a different domain name used for the storefront. As seen by the change in gradients, the BIGLOVE campaign rotated domains for this storefront twice, at the end of June and the middle of August, updating its doorways found in PSRs to redirect to the new instances. We see similar changes in traffic coinciding with each of the domain name changes.

Although there is not sufficient evidence to discern the campaign’s intent, one possibility is that these domain name changes are a proactive countermeasure against domain name seizures. As discussed in Section 3.2.2, luxury brand holders frequently seize domain names to curtail counterfeit sales. However, as we will show in Section 5.3, SEO campaigns are well aware of the ongoing seizures and oftentimes react within days of the initial seizure by simply redirecting to another domain. And being proactive ensures

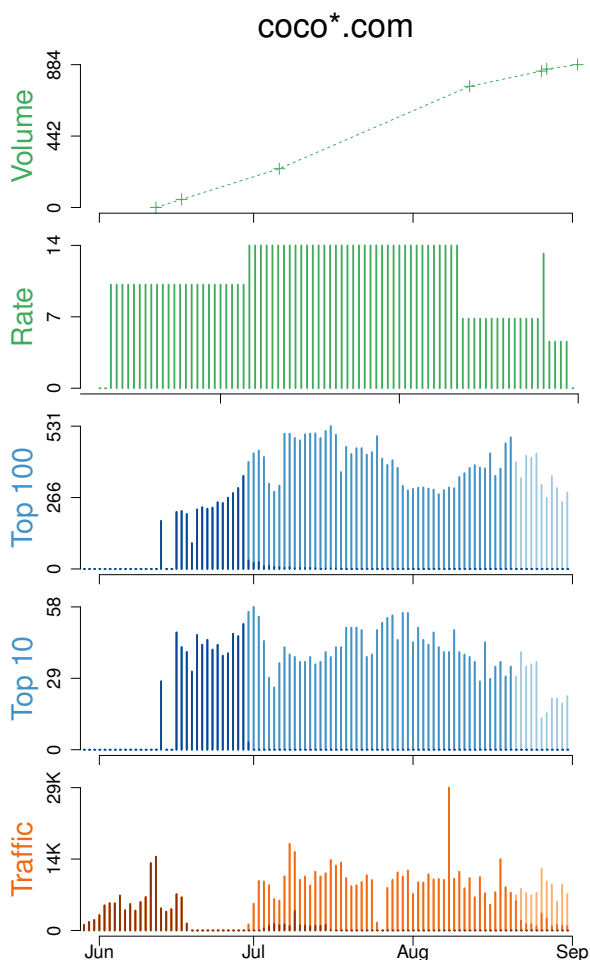


Figure 5: A detailed example of the correlation between a store’s prominence in search results (Top 100, Top 10), the resulting user traffic seen by the store (Traffic), and the monetization of user traffic through orders (Volume, Rate), for a counterfeit Chanel store run by the BIGLOVE campaign from June – September 2014. Each color gradient in the PSRs and traffic graphs is associated with separate instances of *coco*.com*, where each instance used a different domain name.

that there is no downtime: the first domain *cocoviphandbags.com* was seized on July 11, yet by that time the doorways were already redirecting users to the second domain *cocovipbags.com*.

Inspecting the detailed user traffic data from AWStats, we make rough estimates on conversion metrics from *coco*.com* that are

consistent with those reported by marketers [4]. During the months of July and August 2014, `coco*.com` combined received 93,509 visits, 60% of which properly set the HTTP referrer header.⁵ Extracting the referrers reveals the complete set of doorways supplying traffic for this store, and we find 83 out of 174 doorway domains (47.7%) were seen in our crawled PSRs data (recall that we limit the number of terms we search for a given vertical, and so it is not surprising that we do not capture all doorways). Examining user visits more closely, we find each visit generates 5.6 HTML page fetches on average. And when combining the traffic data with the order data from test purchases, we estimate this store had a 0.7% conversion rate, roughly a sale every 151 visits.

5.3 Domain Seizure Interventions

As discussed in Section 3.2.2, brands have the most incentive for undermining online counterfeiters, and a highly visible intervention they can use is to seize the domains of counterfeit storefronts. With this intervention, brands use legal means to seize domain names of stores violating brand holder trademarks, thereby preventing users from visiting sites selling counterfeit merchandise.⁶

We use two sources of data for studying domain seizure by brand holders. The first is the set of PSRs from our crawled search data. Mechanistically, if the site identified in a PSR redirects to a serving notice site provided by one of the third-party brand protection services (e.g., Greer, Burns & Crain [10], SMGPA [33]) or the brand holders themselves, this is a clear signal that the site has been seized. Such notices are very common, but their use is not guaranteed and we have found some seized sites that simply are shut down. The second source of data comes from the set of seized domains listed in court documents embedded in each of these serving notice pages; these documents typically list the other domains seized as a part of a given action and thus enable us to obtain a broader view of domain seizure activity by brand holders spanning up to two years. This list is not comprehensive, however, for two reasons. First, the subset of serving notice sites identified via our search terms may not capture all legal actions (this is particularly true for SMGPA; for GBC we believe we were able to enumerate these sites comprehensively). Second, if the list of domains is not complete or if additional domains are added to a legal action without updating these documents (or after we crawl them) then we will undercount the number of domains being seized. Indeed, we understand through personal communications with brand representatives that these represent a subset of the domains seized through legal process.

By extracting the brand holders and the timestamps from seizure notices, we can also infer how brands use brand protection services. For both GBC and SMGPA, we see evidence that brand holders initiate domain name seizures on a periodic basis, typically on the order of months between rounds of seizures. Although our data set suggests that a handful of brands seized domains more frequently—we identified Oakley issuing 6 court cases at monthly intervals, Ugg's issuing 19 court cases at bi-weekly intervals, and Chanel issuing 18 court cases also at bi-weekly intervals—they tend to be the exception in our data set.

⁵The HTTP referrer header is not properly set in many situations, including transitioning from HTTPS to HTTP, visiting through an email client, visiting through a proxy that strips the header, or simply typing the URL directly into the browser.

⁶Another option would be to seize doorway domains, but doing so presents two obstacles: there are two orders of magnitude more doorway domains than stores (Table 2), and the doorways are often compromised sites.

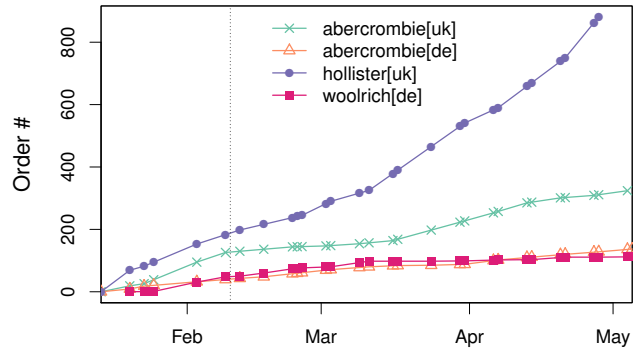


Figure 6: Order number samples over time in early 2014 for the PHP?P= campaign. Each curve corresponds to one of four international stores, where three sell Abercrombie (United Kingdom, Germany) while the remaining sells Woolrich (Italy).

To assess the completeness of observing domain seizures using PSRs, we compared the court cases seen in PSRs against ground truth we collected by enumerating all court orders from GBC, which are publicly accessible through their Web site. During the time frame of our study (November 2013 to July 2014), we observe 47 cases in PSRs out of the 50 total cases initiated by GBC (94%) during the same time frame. If these results from GBC are typical, this suggests that our crawled search data likely captures the bulk of such seizure activity.

5.3.1 Coverage

Brand holders have been aggressive in seizing storefront domains. Table 3 shows a breakdown of seized domains across brands, storefronts, and campaigns. From manually examining the court documents embedded in seizure notice pages, brands arranged to have almost 40,000 domains seized over two years. Specifically, our data set shows that while representing 17 brand holders GBC seized 31K domains using 69 court orders from December 2012 to June 2014. Similarly, while representing 11 brand holders, our data set identifies 8K domains seized by SMGPA using 47 court orders from February 2012 to July 2014.

Seizing domains can disrupt online counterfeit stores. For example, Figure 6 shows order numbers over time for four stores promoted by the PHP?P= campaign. The domain for its Abercrombie UK store was seized on February 9, 2014 (vertical dotted line), and its rate of new order numbers declined immediately (it did not stop completely due to reaction by the SEO campaign, as discussed below in Section 5.3.2).

Despite their aggressiveness, though, brand holders must be far more aggressive when seizing domains. From the PSRs crawled, we directly observed 290 seizures over our eight-month period: 214 seized by GBC and 76 by SMGPA. Compared with the total number of storefronts we observed (7,484), though, these domain seizures represent just a small percentage (3.9%) of stores used by SEO campaigns. As a result, unless brands comprehensively seize domains, stores promoted by SEO campaigns remain unaffected and they continue to attract traffic and customers via search. Referring back to the example in Figure 6, the orders remained steady at other stores whose domains were not seized and the SEO campaign remained effective overall.

5.3.2 Reaction Time

Even if brands eventually seize all storefront domains, though, they take such a long time to seize domains, and attackers respond so quickly to having domains seized, that the current environment still strongly favors the attackers.

As with labeling sites as “hacked”, the time from when a storefront goes live and when a brand seizes the store domain represents the window of opportunity for a counterfeit store to monetize traffic. As in Section 5.2.2, we can compute the lifetime of seized stores to measure their earning potential before seizure. We define the lifetime of seized stores as the delta between the first time that the storefront appears in PSRs to the time the domain was seized.

We find the average lifetime of seized stores lies between 58–68 days for GBC and 48–56 days for SMGPA. Again, due to the nature of our data collection, we can only observe seizures when crawling search results redirects our crawlers to seized stores. Therefore, we cannot determine exactly when a seizure takes place. Instead, we present two bounded numbers approximating the true lifetime. The smaller is the duration ending when we last saw the store actively poisoning search results, while the larger is the lifetime ending when we first definitively observed the seizure.⁷

However, even when brands seize storefront domains, the SEO campaigns possess backup domains in anticipation of such an intervention and can quickly react to continue to monetize traffic without significant interruption. Returning to the example in Figure 6, when the Abercrombie UK domain was seized, the PHP?P= campaign changed their doorways to forward to a new store domain within 24 hours. This domain agility represents a critical weakness of seizures: even though a store domain is seized, SEO campaigns can easily modify their doorways to redirect users to their backups rather than the seized domains.

Indeed, we found widespread evidence of attackers exploiting this weakness as a countermeasure to domain seizures. Specifically, of the 214 seized stores from GBC, 130 were redirected to new stores (59 of which were subsequently seized) and, among the 76 seized from SMGPA, 57 were redirected to new stores (22 of which were subsequently seized). These responses by the counterfeiters happened on average within 7 and 15 days of the initial seizure, respectively, for GBC and SMGPA. Such domain agility suggests the counterfeiters are well prepared for domain seizures, and as a result such interventions are not likely to undermine their business.

6. CONCLUSION

Online business in counterfeit luxury goods is brisk: from the site of just one supplier, we saw over 250,000 successfully delivered orders in nine months. Such businesses prosper by poisoning search results for popular luxury goods to attract customers to their online storefronts; for heavily-targeted brands, a third of the top 100 search results are so poisoned for months at a time.

In this paper we presented techniques for detecting poisoned search results that lead users to counterfeit stores, and a classification approach for mapping the Web sites of these stores into distinct SEO campaigns that promote these sites. From eight months of crawled search results for 16 brand verticals, we detected 2.7 million PSRs using 27 thousand doorway pages that redirect users to 7,484 storefronts, and classified over half of the PSRs into 52 distinct SEO campaigns. Simultaneously, we created test orders on stores to sample their order number sequence space to estimate their order volume over time.

⁷Given how quickly campaigns react to domain seizures, it is also possible for campaigns to redirect doorways to new store domains faster than our crawler can detect that the initial domain was seized.

Finally, we used our crawler and order data to study the effects of both search engine and domain seizure interventions on these abusive SEO activities. Although we find instances where both can have the desired effect of disrupting counterfeit sales activity, overall neither are currently employed with the level of coverage or responsiveness necessary to be broadly effective against the actors in this market. Search engines and brand holders should take into account that these activities are organized as business campaigns, that effective interventions should target their infrastructure at the granularity of these campaigns, and that they are being targeted by dozens of campaigns. Otherwise campaigns have shown great agility in adapting to partial intervention, and in filling in gaps left by the disappearance of other campaigns. We believe that the measurement and classification techniques we describe in this paper for identifying campaigns and their infrastructure could provide the improved targeting required for more robust intervention.

7. ACKNOWLEDGMENTS

We are grateful to Aaron Hurtado and Kelsey Ma, supported by the National Science Foundation Research Experience for Undergraduates (REU) program, for scripting the process of gathering order volume data. We also thank Andreas Haerberlen and the anonymous reviewers for their valuable feedback. This work was supported in part by National Science Foundation grant NSF-1237264, by the Office of Naval Research MURI grant N00014-09-1-1081, and by generous support from Yahoo, Google, Microsoft, and the UCSD Center for Networked Systems (CNS).

8. REFERENCES

- [1] AWStats — Free log file analyzer for advanced statistics (GNU GPL). <http://awstats.sourceforge.net/>.
- [2] BANNUR, S. N., SAUL, L. K., AND SAVAGE, S. Judging a site by its content: learning the textual, structural, and visual features of malicious Web pages. In *Proceedings of the ACM Workshop on Artificial Intelligence and Security (AISEC)* (Chicago, IL, Oct. 2011).
- [3] CHACHRA, N., MCCOY, D., SAVAGE, S., AND VOELKER, G. M. Empirically Characterizing Domain Abuse and the Revenue Impact of Blacklisting. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)* (State College, PA., June 2014).
- [4] CHAFFEY, D. Ecommerce conversion rates. <http://www.smartinsights.com/ecommerce/ecommerce-analytics/ecommerce-conversion-rates/>, March 2014.
- [5] DER, M. F., SAUL, L. K., SAVAGE, S., AND VOELKER, G. M. Knock It Off: Profiling the Online Storefronts of Counterfeit Merchandise. In *Proceedings of the ACM SIGKDD Conference* (New York, NY, August 2014).
- [6] Generate a Random Name — Fake Name Generator. <http://www.fakenamegenerator.com/>.
- [7] FAN, R. E., CHANG, K. W., HSIEH, C. J., WANG, X. R., AND LIN, C. J. LIBLINEAR: A Library for Large Linear Classification. <http://www.csie.ntu.edu.tw/~cjlin/liblinear/>.
- [8] Google Search Engine Optimization Starter Guide. <http://www.google.com/webmasters/docs/search-engine-optimization-starter-guide.pdf>.
- [9] GOOGLE. Results labeled “This site may be hacked”. <http://support.google.com/websearch/answer/190597>.
- [10] GREER, BURNS, & CRAIN. Anti-counterfeiting legal strategies, enforcement and remedies. <http://gbclaw.net/practiceareas/anti-counterfeiting>.
- [11] GRIER, C., BALLARD, L., CABALLERO, J., CHACHRA, N., DIETRICH, C. J., LEVCHENKO, K., MAVROMMATIS, P., MCCOY, D., NAPPA, A., PITSILLIDIS, A., PROVOS, N., RAFIQUE, M. Z., RAJAB, M. A., ROSSOW, C.,

- THOMAS, K., PAXSON, V., SAVAGE, S., AND VOELKER, G. M. Manufacturing Compromise: The Emergence of Exploit-as-a-Service. In *Proceedings of the ACM Conference on Computer and Communications Security* (Raleigh, NC, October 2012).
- [12] HERLEY, C., AND FLORENCIO, D. Nobody Sells Gold for the Price of Silver: Dishonesty, Uncertainty and the Underground Economy. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)* (London, UK, June 2009).
- [13] HtmlUnit — Welcome to HtmlUnit. <http://htmlunit.sourceforge.net/>.
- [14] INVERNIZZI, L., BENVENUTI, S., COVA, M., MILANI-COMPARETTI, P., KRUEGEL, C., AND VIGNA, G. EvilSeed: A Guided Approach to Finding Malicious Web Pages. In *Proceedings of the IEEE Symposium and Security and Privacy* (San Francisco, CA, May 2012).
- [15] JOHN, J. P., YU, F., XIE, Y., KRISHNAMURTHY, A., AND ABADI, M. deSEO: Combating Search-Result Poisoning. In *Proceedings of the USENIX Security Symposium* (San Francisco, CA, August 2011).
- [16] KANICH, C., WEAVER, N., MCCOY, D., HALVORSON, T., KREIBICH, C., LEVCHENKO, K., PAXSON, V., VOELKER, G. M., AND SAVAGE, S. Show Me the Money: Characterizing Spam-advertised Revenue. In *Proceedings of the USENIX Security Symposium* (San Francisco, CA, August 2011).
- [17] KARAMI, M., GHAEMI, S., AND MCCOY, D. Folex: An Analysis of an Herbal and Counterfeit Luxury Goods Affiliate Program. In *Proceedings of the eCrime Researchers Summit (eCRS)* (San Francisco, CA., September 2013).
- [18] LEONTIADIS, N., MOORE, T., AND CHRISTIN, N. Pick Your Poison: Pricing and Inventories at Unlicensed Online Pharmacies. In *Proceedings of the ACM Conference on Electronic Commerce* (Philadelphia, PA., June 2013).
- [19] LEONTIADIS, N., MOORE, T., AND CHRISTIN, N. A Nearly Four-Year Longitudinal Study of Search-Engine Poisoning. In *Proceedings of the ACM Conference on Computer and Communications Security* (Scottsdale, AZ, Nov. 2014).
- [20] LEVCHENKO, K., PITSILLIDIS, A., CHACHRA, N., ENRIGHT, B., FÉLEGYHÁZI, M., GRIER, C., HALVORSON, T., KANICH, C., KREIBICH, C., LIU, H., MCCOY, D., WEAVER, N., PAXSON, V., VOELKER, G. M., AND SAVAGE, S. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *Proceedings of the IEEE Symposium and Security and Privacy* (Oakland, CA, May 2011).
- [21] LIU, H., LEVCHENKO, K., FELEGYHAZI, M., KREIBICH, C., MAIER, G., VOELKER, G. M., AND SAVAGE, S. On the Effects of Registrar-level Intervention. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)* (Boston, MA, March 2011).
- [22] LU, L., PERDISCI, R., AND LEE, W. SURF: Detecting and Measuring Search Poisoning. In *Proceedings of the ACM Conference on Computer and Communications Security* (Chicago, IL, October 2011).
- [23] MarkMonitor Brand Protection. <https://www.markmonitor.com/services/brand-protection.php>.
- [24] MCCOY, D., DHARMDASANI, H., KREIBICH, C., VOELKER, G. M., AND SAVAGE, S. Priceless: The Role of Payments in Abuse-advertised Goods. In *Proceedings of the ACM Conference on Computer and Communications Security* (Raleigh, NC, October 2012).
- [25] MCCOY, D., PITSILLIDIS, A., JORDAN, G., WEAVER, N., KREIBICH, C., KREBS, B., VOELKER, G. M., SAVAGE, S., AND LEVCHENKO, K. PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs. In *Proceedings of the USENIX Security Symposium* (Bellevue, WA, August 2012).
- [26] MOORE, T., AND CLAYTON, R. The consequence of non-cooperation in the fight against phishing. In *Proceedings of the eCrime Researchers Summit (eCRS)* (Atlanta, GA., October 2008).
- [27] MOORE, T., LEONTIADIS, N., AND CHRISTIN, N. Fashion Crimes: Trending-Term Exploitation on the Web. In *Proceedings of the ACM Conference on Computer and Communications Security* (Chicago, IL, October 2011).
- [28] OPSEC. Brand protection from manufacturing to retail. <http://www.opsecsecurity.com/brand-protection>.
- [29] PRICEWATERHOUSECOOPERS. IAB Internet Advertising Revenue Report, 2012 Full Year Results. <http://www.iab.net/media/file/IABInternetAdvertisingRevenueReportFY2012POSTED.pdf>.
- [30] RAJAB, M. A., BALLARD, L., MARVROMMATHIS, P., PROVOS, N., AND ZHAO, X. The Nocebo Effect on the Web: An Analysis of Fake Anti-Virus Distribution. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)* (San Jose, CA, April 2010).
- [31] SAFENAMES. Mark Protect. <http://www.safenames.net/BrandProtection/MarkProtect.aspx>.
- [32] SAMOSSEIKO, D. The Partnerka – What Is It, And Why Should You Care? In *Proceedings of the Virus Bulletin Conference* (September 2009).
- [33] SMGPA. <http://smgpa.net/>.
- [34] STONE-GROSS, B., ABMAN, R., KEMMERER, R., KRUEGEL, C., STEIGERWALD, D., AND VIGNA, G. The Underground Economy of Fake Antivirus Software. In *Proc. of the 10th Workshop on the Economics of Information Security (WEIS)* (Washington D.C., 2011).
- [35] WANG, D. Y., SAVAGE, S., AND VOELKER, G. M. Cloak and Dagger: Dynamics of Web Search Cloaking. In *Proceedings of the ACM Conference on Computer and Communications Security* (Chicago, IL, October 2011).
- [36] WANG, D. Y., SAVAGE, S., AND VOELKER, G. M. Juice: A Longitudinal Study of an SEO Campaign. In *Proceedings of the Network and Distributed System Security Symposium* (San Diego, CA, February 2013).
- [37] WANG, Y.-M., MA, M., NIU, Y., AND CHEN, H. Spam Double-Funnel: Connecting Web Spammers with Advertisers. In *Proceedings of the International World Wide Web Conference (WWW)* (Banff, Alberta, May 2007).
- [38] WHITTAKER, C., RYNER, B., AND NAZIF, M. Large-Scale Automatic Classification of Phishing Pages. In *Proceedings of the Network and Distributed System Security Symposium* (San Diego, CA, February 2010).