

Libel Inc: An Analysis of the Libel Site Ecosystem

Rasika Bhalerao
New York University
United States
rasikabh@gmail.com

Damon McCoy
New York University
United States
mccoy@nyu.edu

Abstract—Libel sites publish anonymously submitted unproven libelous claims about individuals that often include personal information about the subject. The stated goal of the sites is to “warn” the public about an individual but the impact is harassment and ruining the subject’s reputation. These individual libelous posts are surfaced when searching for a person’s name using an online search engine and can cause a range of harms from emotional to economic. For example, the libelous posts might surface if a potential employer performs a Google search as part of a “background check.” There have been prior news reports of this troubling phenomena but no systematic analysis of the ecosystem.

In this paper, we conduct a rigorous analysis of these libel sites, supporting services, and intervention by Google. We discovered and analyzed 9 libel sites, 7 websites for reputation management services, and 12 related websites. We found that all of the libel websites included at least one method of generating revenue. The most common revenue generation method was including advertisements for “reputation management services” which require payment for the removal of a post. We found that all of these removal services were dubious in nature and that the removal policies were akin to extortion. Our analysis of Google’s intervention to reduce the visibility of these websites indicated that it appeared to only reduce the visibility of the specific libel post URL but that other URLs containing links to the post or the headline text of the post were still highly ranked. Based on our findings, we make recommendations to many of the stakeholders about potential approaches for mitigating this abusive ecosystem.

Index Terms—libel, harassment, hate speech, extortion

1. Introduction

“Cheaters hate this site!” “Report anonymously.” “Did you think she’d never find out?” For people who feel deeply wronged by another person, websites functioning as online billboards for anonymously posting libel are an easy way to air their frustrations, “warn” others about the target, publicize the target’s personal information, and ruin the target’s reputation. These widely circulated libel sites contain anonymously written libel about thousands of people with no easy way to remove the posts unless the post subjects can afford to pay for reputation management services. This libel appears in online search engine results, social media, and background search services, attaching a “cyber scarlet letter” [1] to its subjects. Subjects of online

libel report that these (unverified, often exaggerated, and sometimes completely fabricated) accusations disproportionately negatively affect multiple aspects of their lives indefinitely [2]. Libel sites and the ecosystem they foster become increasingly relevant amid the growing concern and research surrounding harassment, online privacy, and the publication of personal information [3]–[5]. As these libel sites tout their legal protection (in the United States) under Section 230 of the Communications Decency Act [6], our main defenses from this form of harassment are through search engine policies, payment platform policies, and the reputation management services which are advertised on libel sites as subjects’ only pathway to post removal. In this paper we investigate the efforts of 1) reputation management services and 2) internet search engines to protect libel victims, and discuss the ecosystem comprised of libel sites and these protection efforts.

We find that reputation management services are prohibitively expensive, with untrustworthy-appearing websites and advertisements. Our measurement of libel post removal over time finds that very few posts are actually removed in the way advertised by these services, indicating that the services are either ineffective or unappealing to libel subjects. As our main line of defense, Google, the most popular internet search engine, has recently included “exploitative removal practices” [7] as a legitimate reason to report a site to be de-indexed. We find that while these policies are motivated well, they remain insufficient. For example, the de-indexing only applies to the single post about the individual making the report, and moreover, the de-indexing is often insufficient because libel sites are continuously adapting to circumvent “pressures from law enforcement, federal legislation, and search engine algorithm changes simply to stay afloat” [8]. Not to mention, other internet search engines have no such policy.

To analyze this ecosystem, we collected and studied 9 libel sites, 7 websites for reputation management services, and 12 other websites related in the online libel ecosystem to analyze libel contents and motivations, reputation management services’ libel post removal and payment details, and online search engines’ interaction with libel posts. We find that reputation management services are insufficient for removing libel and may additionally be an integral part of the extortion inherent in the online libel ecosystem. We discuss online search engines’ policies to battle this form of extortion and find that Google’s limited intervention is the only effort designed to address this issue. Finally, we make recommendations for search engines, payment platforms, online libel’s intended audience and writers,

libel sites and reputation management companies, unwilling post subjects, and researchers and journalists. While it is unfortunate that our legislation and policies make the libel ecosystem inevitable, we make our recommendations to reduce the harm from this form of harassment.

2. Background and Motivation

Libel sites function by taking user reports, usually submitted via a short, straightforward online form, and posting these reports publicly and freely on their websites. The writer submitting the form remains anonymous¹, and there is no verification of the claims submitted. When published and circulated by the site, this libel becomes our modern version of public shaming, “result[ing] in an over-determined punishment with indeterminate social meaning” [9]. This assessment is in agreement with reports by large reputable news companies [10] on the financial and psychological harm and irreversible damage to the reputations of subjects of this “cyber scarlet letter” [1]. Further to that point, subjects of this libel are most often women or minorities [11] described with “misogynistic language and insulting implications” [8]. Subjects report fear about people targeting their families and stalking them at their workplaces because of their personal information being publicly available in an often socially irrelevant but strongly worded libel post [2].

Each post includes an image of the subject, a few paragraphs of text detailing the accusation, a descriptive and often explicit title, and other elements that attempt to legitimize the website, such as social media links, text boxes to reply to posts, search bars, advertisements, and links to related posts. These posts almost always include subjects’ personal identifying information, usually including images; even the title almost always includes the subject’s full name and includes accusations and derogatory descriptions. These posts are listed in newsfeeds on the libel sites’ homepages (Figure 1), which also include the advertisements, website tabs, and other legitimizing features. Each individual libel post has its own dedicated webpage that can be accessed by clicking on the post in the libel site’s newsfeed, directly with the URL, or by using the libel site’s search functionality. The post’s URL is usually similar to the post title, containing the subject’s full name and accusation. There are often multiple posts about the same person on a site, written at different times with slightly different text. Libel sites offer search functionality to find individual posts, and posts are also categorized and listed in sites’ menus and sitemaps. There is no functionality for requesting post removal; most libel sites instead display advertisements for reputation management sites and recommend that unwilling post subjects should seek “legal assistance” to have their posts removed.

2.1. Similar Sites

Many people search websites are closely linked with libel sites and reputation management sites. For example, L2 states that it is an affiliate of a popular people search

website that advertises itself as a means of “background searching” people by providing personal information such as birth date, physical address, and relatives. The key differences between people search websites and libel sites are that people search websites collect personal information and provide it only if a user searches for it specifically, they advertise free methods for unwilling subjects to have their information removed from the database, and they do not publish unvetted libel.

Of course, social media contains many communities that function as libel sites, without the need for revenue. For example, large popular social media platforms, contain groups of users focused on exposing infidelity in their local geographic region. These social media communities appear in search engine results, social media platform newsfeeds, and even other related social media platforms (i.e., a screenshot from one social media platform appearing on another). Interestingly, reputation management companies and libel sites often have social media presences, but none of their libel content appears in social media.

Other sites that reputation management sites or social media often refer to as libel sites are actually discussion forums. The key difference is that these forums do not include identifying information about subjects, and are designed to foster discussion more than a newsfeed. There are rules for engagement, and people aim to help each other work through personal struggles, some of which include infidelity. Several of these forums are targeted for specific demographics (i.e., Black women). The rules of these forums often allow forms of speech that would not be allowed on mainstream forums, which is likely a reason that they are included in lists of libel sites.

Finally, there are websites for reporting scams (mostly about products and services) that have sections for libel about individuals. For example, some websites for reviews of diverse companies and individuals aim to protect people from falling victim to fraud. These types of websites are too large and unfocused to include in our study.

3. Dataset Collection

3.1. Website Collection

For our analysis we collected 9 active libel sites, 7 reputation management sites, 2 aggregator sites, and 10 other websites in the libel ecosystem (e.g., blogs and social media platforms with libel site-like communities) through a snowballing method. We began by using internet search engines with search terms including “cheaters exposed,” “homewreckers,” and similar variations of these terms. Reputation management sites were advertised on libel sites, and also contained lists of libel sites from which they guaranteed removal. Similarly, social media platforms often contained posts and communities related to libel sites, and reputation management and libel sites often had presences on social media. Posts from libel sites often appeared in aggregator sites, which aim to archive content found on the internet. We reason that if a site was not surfaced in our snowballing method, then it is likely not popular enough to be a large part of the “public shaming” and harassment of the online libel

1. Sometimes the post text accidentally gives away the writer’s identity through their relationship with the subject.

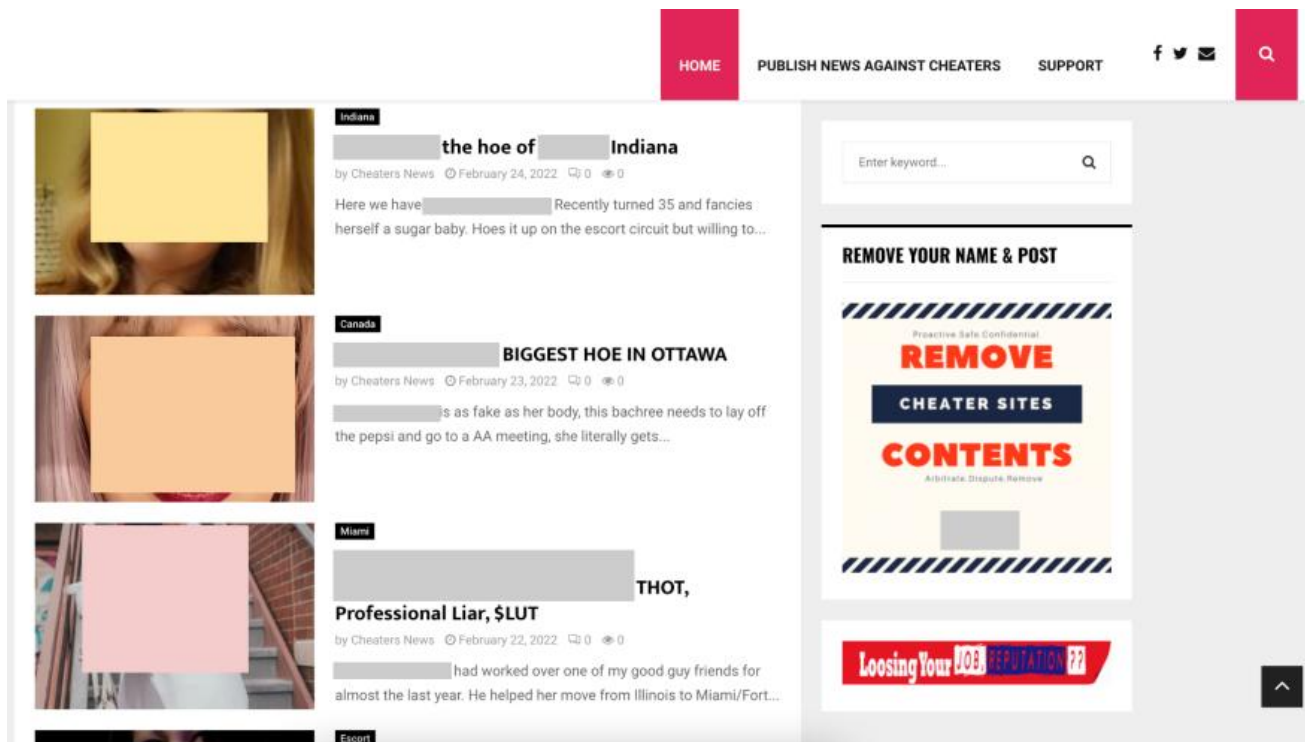


Figure 1: Screenshot of the newsfeed of a libel site (L8). Identifying details have been obscured. Three posts are shown, each with an image, a title, and the beginning of the text of the post. Each post also has a timestamp, the number of comments posted in reply to each post, and the number of views (all three posts are too recent for any views or posts). Each post has a category (three different locations), and the category of a fourth post is also seen (“Escort”). On the right are two prominent advertisements for reputation management services for post removal. At the top are tabs for the homepage (which is the current page), the form to submit posts for publication, and a support page. There is also an option to search within the site.

ecosystem. Table 1 contains the (anonymized) relevant active sites found through this method. In the context of our website collection, “infidelity” refers specifically to posts complaining about a cheating partner. If the post instead complains about the person with whom a partner cheated, it is classified as “other women”; this is its own category because posts written by women complaining specifically about the women with whom their husbands cheated (not the husbands themselves) were popular, and multiple libel sites were dedicated to just this topic. Also in our taxonomy, the difference between libel and gossip is motivated by the way websites describe themselves: sites with libel advertise themselves as ways to “warn” about the subject or ruin their reputation, while sites with gossip advertise themselves as such, without the same goal. The contents of posts with libel and gossip are largely the same. Forums, sites for reporting scams, and social media platforms are different from libel sites, as explained in Section 2.1. Finally, blogs are different from libel sites because their website owners claim authorship of the content.

3.2. Post Collection

We also collected posts from each libel site. The total number of posts collected from each site can be found in the “Total” column of Table 3, with a total of 2,312 active posts overall. Table 3 also describes the methods for collecting URLs from each site. Because of the different

scraping methods, the majority of our dataset contains posts from L1 (705 posts), L9 (691 posts), and L5 (530 posts). L7 had the fewest, 47 posts. Sitemaps were downloaded when available in the website, and the most recent sitemap was used when archived versions were available. For sites without sitemaps available, we scraped the main newsfeeds; when pagination was available we scraped the five most recent pages, and when it was not, we scraped the most recent page. L7 did not have pagination, but it had the ability to filter posts using a calendar, and it returned all selected posts on a single page which we scraped. Websites with no pagination or calendar filtering used infinitely scrolling webpages that loaded new posts only when scrolled through, and they had no organization structure or calendar to facilitate scraping; these sites were likely attempting to evade automated crawling. For these sites, we scraped the top newsfeed page after searching based on each of the top 16 most popular cities (which were listed on the site).

3.3. Ethics

Our study was exempted by our institutional ethics review board (IRB). We do not publish any personally identifiable information in this paper, and we attempted to limit the amount of such information collected and stored. When collecting posts, we stored only post URLs without downloading the contents of each post; since most of these URLs contained subjects’ full names, this dataset

ID	Site Subject
L1	Infidelity and gossip
L2	"Other women"
L3	"Other women"
L4	Infidelity and similar libel
L5	"Other women"
L6	Infidelity and gossip
L7	Infidelity and similar libel
L8	Infidelity and similar libel
L9	"Other women"
R1	Reputation management
R2	Reputation management
R3	Reputation management
R4	Reputation management
R5	Reputation management
R6	Reputation management
R7	Post removal from L2
A1	Aggregator
A2	Aggregator
O1	Anonymous forum
O2	Social media platform
O3	Scam reports
O4	Infidelity and libel blog
O5	Libel copied from social media
O6	Anonymous forum
O7	Social media platform
O8	Social media platform
O9	Social media platform
O10	Social media platform

TABLE 1: Summary of websites studied including dedicated libel sites, reputation management company sites, and other related sites. Aggregator sites archive posts from sites including libel sites. "Other women" indicates posts describing women with whom married men had sexual relations, usually written by the married men's spouses.

is private and will not be shared. We also do not include the names of libel sites, reputation management services, or related websites to avoid increasing their viewership. While using personal information contained in the libel posts would yield a more realistic evaluation of libel sites' popularity and impact, we did not attempt to investigate targets or contact them as this may increase the psychological harm of this harassment on them.

4. Findings

4.1. Libel Sites

We used grounded theory to determine the main motivations or justifications of libel sites and purposes of posts. The first author generated the codes and regularly met with other subject matter experts to understand motivations in the libel site ecosystem.

4.1.1. Libel Site Motivations. Libel sites cite two main motivations or justifications for their existence. One is for people who felt deeply wronged to retaliate, warn others, and harm the reputation of the post subject. Statements of this motivation in site descriptions and guidelines contain real-sounding emotion in order to appear truthful about this intention of the site. With names like "Cheaters Exposed" and "Karma for Homewreckers"², these sites solicit reports to "expose fakes" or people deemed "crazy"

and display testimonials by readers who claim to have benefited from a warning about a subject on the site. L1 even allows free post removal if the post subject "makes amends" with the post writer and asks them to "post a new, positive review of you" or submits "proof of repentance"; after receiving this proof, L1 removes the post at their leisure, with a much longer wait time for people who play "games and people who cause trouble."

The other motivation is to generate revenue. Every libel site that we studied had a method to generate revenue, most often permanent advertisements for reputation management companies. Rather than having straightforward processes for requesting removal of personal information, libel sites hosted advertisements with headlines such as "Remove your name now!!" for post removal from reputation management services. Even if libel sites were not actually maintained by reputation management companies themselves, the libel sites still hosted prominent post removal service advertisements for which reputation management companies had paid. The same reputation management advertisements appeared on multiple libel sites. In addition to advertisement revenue, we found some libel sites also generate revenue through dox-for-hire services and people search services offered. For example, L1, which has the free post removal process, generates money through their offer: "For 0.01BTC (about \$50) we find people for you." L2's links to "find out more dirt" go to their affiliated people search website which offers "additional and supplemental data about a person for \$19.99." Each website had a financial incentive in addition to the emotional justification to remain online. Though in the website text, the goal to recover from people who hurt them was as prominent (if not more so) as the goal to generate money, it is difficult to quantify the true motivations to make this comparison.

4.1.2. Libel Post Taxonomy. Post analysis included manually looking at at least 50 posts from each libel site listed in Table 1 and generating codes for this taxonomy. We also checked reputation management sites' lists of services to ensure completeness of our taxonomy. We found that libel posts are usually written for at least one of the following purposes:

- 1) To complain about a partner's infidelity in a relationship. Most of these posts use strong language to describe women, with phrases ranging from the stereotyping "thirsty and desperate" to the imaginative "spineless worm" to the unoriginal "whore."
- 2) To complain about a third person with whom the writer's relationship partner cheated. These posts completely place all blame on the third person, with statements like "three children are growing up without a father because she couldn't keep her legs closed." Most posts of this category and the previous category solicit harassment towards the subjects with statements like "Feel free to let her know what you think," referring to the option to reply via comments on the post.
- 3) To expose someone as a sex worker. These posts expose people who use separate online personas to preserve their privacy with statements like

2. Both of these names have been slightly modified for this paper.

“[Subject’s legal name], aka [subject’s name used in sex work] is a professional escort who fancies herself a sugar baby.”³ Exposing someone both as a sex worker and as a member of the LGBTQ+ community at the same time (usually exposing an identity of gay or transgender) was also common, with post titles such as “[Name] Moonlights As A Gay Male Escort” and “[Name]... Former Gay PornStar.”

- 4) To use explicit language or slurs to describe someone’s appearance or personality, with no substantial content. For example, people are described as “rightfully insecure about being born so ugly” or “Gay a55 mother fuking d1ck sucking loser.” These descriptions are often full of racist, anti-LGBTQ+, or sexist slurs.
- 5) To warn the public about a scammer or a person who says racist or otherwise abusive things. Examples include “[Name] is behind the racist twitter account @Twitter handle)” and “[Name] is a fraud claiming to live in Monaco [...] His company is fake.”
- 6) To warn the public about serial [child] predators or other allegedly dangerous people who, for example, “set up spycameras in women’s restrooms of restaurants” or “calls and sounds like he’s doing PNP, hangs up the phone after getting off... Avoid him!” Quotations of graphic depictions involving child predators are omitted here.

Multiple motivations are often present in the same post, most often complaining about infidelity of both a partner and the third person involved. All accusations are unverified by the libel sites; several sites even include warnings that some posts claiming infidelity in relationships have been reported as untrue.

4.2. Reputation Management Sites

Libel sites had no functionality for removal by a post writer or subject. Instead, reputation management companies, most often R1 and R2, were advertised on libel sites as expert arbitrators. All libel sites except L1⁴ recommended that unwilling post subjects contact lawyers through reputation management companies to help them remove posts.

4.2.1. Reputation Management Company Types. The clearest division among reputation management companies was the appearance of legitimacy of the services offered: R1, R2, R3, and R6 advertised themselves as consulting services, with lawyers available to help in multiple aspects of online reputation including hiding negative online reviews for companies, promoting brands using search engine optimization, and removing libel posts for individuals. These companies listed employee

3. While it is likely that some of these post subjects are not actually sex workers, the post writers use strong language such as “serial mistress” and “always seeking for a sugar daddy” while intending to simply convey the meaning that the subject is promiscuous.

4. L1 had their own methods of post removal, claimed that “lawyers make us laugh,” and listed reasons why they believed themselves untouchable by American or European lawyers.

information, testimonials, and answers to frequently asked questions on their websites.⁵

R4, R5, and R7 were much less official in appearance. R7 was designed only for post removal from L2 and did not advertise itself as a company. R4 and R5 claimed to be companies, but their websites were poorly designed, with no language about lawyers; instead, they claimed to be “capable of having content permanently deleted from most cheater websites, and all of those mentioned [in the list of libel sites] above.” They also discussed hiding criminal records and removing posts from social media. R5 copied the logo from an unrelated company for their website and advertised services to “spread fake news about a political rival.”

4.2.2. Prices and payment methods. Reputation management companies generated revenue from clients paying for their services, mostly⁶ post removal. For example, R2 stated that clients’ money was used for “arbitration between [R2] and website owners for removals, guarantees, and monitoring.”

The websites for reputation management services were opaque about payment-related details, and instead requested that prospective clients contact them asking for quotes. We contacted R1 through R6 with the message detailed in Figure 2, asking what the price would be to remove one link from L4 and asking which payment methods would be accepted. All reputation management sites except R5 had forms embedded in them asking for a prospective client’s name, email address, and message detailing the request, and it was implied that the company would respond by email. We sent the message to R2 using their online form. The forms on R1, R3, R4, and R6 required more information including a link to the post that the client wanted removed, which we did not want to provide, so we instead contacted them through their advertised email address (along with emailing R5 as specified). We chose Sal Kavar as our fake name and salkavar85@protonmail.com as our email address because the name Sal Kavar has been used for several Nigerian Prince scams [12]; searching for the name online returns copious reports of email scams, none of which are likely using Sal’s real name. We created a new email address for the username salkavar85 on the privacy-centered email platform ProtonMail [13]. All inquiry emails were sent on the same afternoon for the appearance of price comparison.

We requested quotes for prices and payment methods from reputation management sites for the removal of one post from L4. We chose L4 because it is large yet difficult to scrape, decreasing the chances that reputation management sites could search our provided fake name and find that there are no matching posts. However, all responses from reputation management companies guaranteed relevant removal not only from L4 but from any other libel sites and search engines. They also specified that payment would only be required if the removal was

5. A quick Google search reveals that the people listed as employees do not really work there, and video testimonials are created by paid actors.

6. The exception is the services associated with L1, which has its own removal policy which is free. L1 makes its revenue from paid requests to dox people.

From: salkavar85@protonmail.com
Subject: Inquiry for personal reputation management
Message: Hello,

How much would it cost to remove a post from [L4]? Which payment methods do you accept?

Regards,
Sal Kavar.

Figure 2: Message sent to R1 through R6 to inquire about prices and payment methods.

successful (or, for those that required payment in advance, it would be refunded if removal was unsuccessful).

ID	Price	Payment methods	Time
R1	"2000 to 2500 USD"	"credit cards and bank wires"	
R2	"\$2,500"	"wire transfer, cashier's check and Zelle"	
R3	"\$2700 USD"	"Wire transfer"	
R5	"\$ 5800 USD"	"international bank wire transfer"	"90 days"
R6	"\$3500 CAD + tax"	"credit card or email transfer"	"2-3 weeks"

TABLE 2: Quotes for prices, payment methods, and post removal times provided by reputation management sites when requested. R7 is omitted because it specifically targets removal from L2, requires sending personally identifiable information, and does not mention money on its website. R4 did not respond to a request for a quote.

As seen in Table 2, prices provided by reputation management companies ranged from \$2000 to \$5800, and those that guaranteed post removal by a certain time had estimates ranging from 2 weeks to 90 days. Every reputation management company preferred payment by wire transfer. Much like cashier's checks and Zelle, wire transfers provide no fraud protection; once funds are transferred, they can only be returned if the reputation management company wishes to initiate another transfer. Two companies also offered payment by credit card.

The main limitations of this analysis are that 1) the quotes provided were presumably specific to our request and 2) we did not verify the claims that they would accept these payment methods or the other guarantees stated. While we attempted to make our request as general as possible, using a gender-neutral name of ambiguous race with a popular libel site and limiting the extra information available, we cannot know how prices, timelines, and payment method guarantees generalize. We also know that while two services offered payment by credit card, it is possible that when it is time to pay, they will fake technical issues or use some other excuse in order to require payment by an irreversible method with no fraud protection. Similarly, we received no legally enforceable way to ensure post removal within the time frame. Be-

cause of the large prices for unverified guarantees and use of payment platforms that may attract fraud, it is understandable why many may choose not to use these reputation management services.

ID	URL collection method	Total	2w	4w
L1	Entire sitemap	705	"	"
L2	Newest newsfeed page	60	"	"
L3	5 newest newsfeed pages	50	"	"
L4	16 most popular cities' newest page	90	89	"
L5	Newest sitemap	530	"	"
L6	16 most popular cities' newest page	65	"	"
L7	All posted in December 2021	47	"	"
L8	5 newest newsfeed pages	74	"	"
L9	Most recent sitemap	691	690	"

TABLE 3: Method for collecting URLs of individual posts from each site, number of live individual post URLs collected, and the number out of those remaining live after two and four weeks. " " indicates that the number is the same as the cell to its left.

4.2.3. Libel site post removal. To quantify post removal, we first collected active URLs of individual posts from nine libel sites. We then checked the status codes of the collected URLs in the following weeks and found that two libel sites removed one post each in the first two weeks, and none were removed during the following two weeks. Table 3 shows the methods for collecting URLs for each site along with the numbers of URLs that were active at the time of collection and the subsets of those that remained active after two and four weeks. This low rate of removal is supported by historical data: in addition to the post removed during our observation period, the sitemap for L9 also included 22 additional URLs for posts that were no longer active (they returned status code 403), and the date of the oldest post in L9 is 40 months before our scraping. We conclude that either the vast majority of people are not willing to pay for a post removal service, or the post removal is not being carried out as promised.

The main limitation of this analysis is that the status code may not be the most accurate indicator of a "removed" post, depending on the removal method. For example, some reputation management services advertise the option to remove a post only from search engines for a lower price. Libel sites could have also removed the content from a libel post page without changing the status code (though this would be inconsistent on the part of libel sites since there were prior posts which had been "removed" by removing access, changing the status code). Other options for "removal" that we would not detect include only taking the post off the libel site's sitemap or newsfeed (which would also be inconsistent on the part of libel sites) without affecting the individual post page, or simply adding a warning to the post itself saying that it has been found to be false or misleading. We note that though these other forms of "removal" would be undetected by our methods, they also would not be considered sufficient removal, since libel post URLs often contain enough information to deduce the contents of the "removed" post.

4.3. Internet Search Engines

As noted by news articles, libel sites are efficient because minimal searching on the internet surfaces existing posts about a person [14]. By default, internet search engines index each post from each libel site. While some search engines ban libel sites for “exploitative removal practices” [7] removal requires the subject to report the post and prove that it fits the criteria for removal; furthermore, de-indexing the post is not always sufficient. To illustrate, we found multiple examples where a libel post about a person was de-indexed from search engines but not removed from the libel site, so it still appeared in search results as text on a different page on the same libel site. For those who cannot afford the steep price for complete post removal, R6 offers “to have the post only removed from Google” within “1 weeks” for “\$1000 CAD + tax,” implying that it is easier to simply de-index the post, but still worth \$1000 to have professionals help remove a post that is already known to violate the terms of service. To further investigate the interactions between libel sites and search engines, we studied five popular search engines: Google, Bing, Yahoo!, YANDEX, and DuckDuckGo.

4.3.1. Libel appearance in search results. We found that internet search engines have largely not blocked sites in the libel ecosystem. Search engines’ advertisements recently stopped including libel sites, but reputation management sites and people search sites (often with names that indicate they are linked to libel sites) still appear in ads. Regarding non-ad search results, all libel sites except L4 have their landing pages and newsfeeds highly ranked by the five popular search engines that we investigated. (L4 was de-indexed by Google but L6, which is a copy of L4 with a slightly modified URL, was not.) Individual posts were also surfaced in search results, and the search engine policies are such that individual posts are only de-indexed if the post subject submits the relevant form to the search engine; this holds even for L4, which had its homepage de-indexed but had individual posts appear in search results. One cannot submit a search engine de-indexing request for a libel post about a different person or for any other page on a libel site of which they are not the subject. This means that if a post subject is unaware of the post about them, does not know how to submit a de-indexing request or work through a reputation management service, or chooses not to address the post, the libel post may still continue to be surfaced in search engine results and they may still continue to be affected by the libel.

It is easy for a lay user to surface a target’s libel post. For all libel sites except L3, posts appear in reverse image search results if we search using a copy of the image in the post.⁷ Post images include pictures of individuals such as headshots, pictures of people with their family, and generally images that could easily exist elsewhere and be used for a reverse image search. L3 watermarked all images, so they did not appear in reverse image search results. Searching only somebody’s name usually yields too many results for the libel post to be ranked very

high; however, a post will sometimes be surfaced with the addition of a few carefully chosen keywords, and will very easily be surfaced if the libel site’s name or URL is included in the search terms. We note that since our site discovery method used search engines to generate the seeds for snowballing, these results regarding libel’s appearance in search engines are likely more focused on sites that have not been de-indexed and are still actively contributing to the libel ecosystem.

4.3.2. Insufficient removal. Posts that were likely removed from libel sites through reputation management companies still spread the libel. At the time of collection, we found that 691 out of 713 URLs collected from L9 were active. The 22 inactive URLs from L9 returned status code 403 (meaning access was forbidden by administrative rules). By contrast, URLs for posts that never existed in the same domain and in the domains of other libel sites returned status code 404 (meaning the page did not exist). This shows that when L9 removed posts (likely through the services offered by R1, which is prominently advertised on L9), they only removed access permissions – the post was still listed on their sitemap. Sure enough, posts that were no longer live on L9 in the following weeks also returned status code 403 rather than 404. Similarly, posts from other libel sites that were no longer live in subsequent weeks were still listed in the site’s main newsfeed, advertised as “relevant” links on libel posts about other subjects, and therefore also surfaced in search engine results. Generally, in L9 and L4, “removed” posts were still listed in sitemaps, menus, and newsfeeds as if they still existed.

If a libel site “removed” posts simply by changing the content of the post page and leaving the status code unchanged, we did not detect it as “removed”; however, we do not consider this a complete removal anyways since the libel is often fully available in the URL itself. Post titles include the subject’s full name and a few words summarizing the accusation, and each post’s URL is usually a copy of the title. This yields an incomplete removal because knowledge that the subject at one point had a post about them on a libel site can still harm their reputation without access to the text of the post. Hints in the title and URL provide enough information that one can use an internet search engine to surface posts about the subject on other libel sites or aggregator sites, which are websites that archive posts from libel and other sites; we found that all subjects of posts that were “removed” from L9 (with status code changes) still had fully active posts about them on other sites that were surfaced in search engine results. Further, posts that were “removed” but still listed on libel sites as if they existed also contained text that was returned in search engine results for other pages; e.g., “Related: Two faced whore [name]” may appear as relevant text on a page for a post about a different person surfaced by a search engine. None of the subjects of any “removed” posts that we studied had managed to completely erase their accusation from internet search engines.

5. Discussion: The Online Libel Ecosystem

Search engines and reputation management sites both describe policies they implemented to protect subjects

7. For privacy reasons, we did not test to see if different images of the same person also surface the post.

of libel. However, while search engines likely sympathize with libel subjects, the same cannot be said about reputation management companies. Libel sites offer no explanation for why reputation management companies have the unique ability to remove libel posts other than “expert arbitration.” While reputation management companies are advertised as specialized lawyers or marketing consultants, they pay for permanent advertisements on libel sites and “guarantee” libel post removal within a time frame; lawyers usually cannot make such guarantees with confidence. Even if reputation management services and libel sites are not run by the same people, there is likely an agreement between the two that allows both to stay financially afloat.

Furthermore, reputation management companies are definitely not who they say they are. The people listed as employees on their websites do not actually work there, and their testimonials are made by paid actors. They require payment through methods that keep them anonymous with no fraud protection. R5 even copied the logo from a different company, changed the background color, and posted it on their website as their own logo.

Even if reputation management services are simply sending very well-written letters as lawyers for arbitration (unlikely given their poor grasp of the English language evidenced through their websites and email correspondence), they have repeated the process enough that they have an established reliable system in place, an informal agreement. Given that reputation services are the only pathway to post removal, the agreement makes the online libel ecosystem one that enforces extortion. If libel sites and reputation sites are not the same companies, they are still accomplices.

6. Recommendations

We provide advice for all stakeholders that might be involved in the libel ecosystem for policies and practices to better protect subjects of libel.

6.1. For Internet Search Engines

We highly recommend that internet search engines allow users to report websites for “exploitative removal practices” if they do not already. We urgently suggest that reports should be submitted anonymously and without the requirement that the subject of the post must be the one to submit the report. We also recommend that if a website receives multiple reports for this reason or for displaying personal information of unwilling subjects, then that entire website should be investigated and likely de-indexed for the same reason. Similarly, we recommend using machine learning techniques to avoid sites that resurface with slight modifications after being de-indexed, and to avoid surfacing social media communities that are effectively informal local libel sites.

We do advise caution in proactively using machine learning to detect and automatically de-index libel sites. Libel sites attempt to distance themselves from the extortion by offering no explicit removal procedure and simply including advertisements for relevant services. While in the libel ecosystem this is clearly not a coincidence, a

more general rule implemented by search engines de-indexing any webpage with ads for services related to extortion may cause public backlash without careful consideration. We must also be careful that fraudsters cannot misuse the de-indexing rule to remove negative product reviews and other legitimate feedback to companies.

Removal of reputation management sites would be more controversial and require consideration for each site. While allowing them to be surfaced perpetuates the circulation of libel sites, de-indexing them also makes post removal more difficult for those who can afford it. We recommend investigating each reputation management site and ensuring that they are a legitimate service and not a facade for another branch of the same service that hosts libel sites, which would make them part of the “exploitative removal practices.” Similarly, we recommend discouraging advertisements from people search sites, which drive traffic to the libel ecosystem and cause harm on their own [5].

6.2. For Social Media Platforms

Libel posts will likely be driven to social media platforms if libel sites are effectively banned from internet search engines. In fact, major social media platforms already contain many communities meant to function as informal local libel sites. Just as other forms of harassment or doxxing are banned on social media, we recommend implementing rules and content moderation policies to battle this growing threat. The revenue of libel sites and reputation management services would likely decrease if users could report libel communities and have them removed, combating this exploitative ecosystem.

6.3. For Banks and Payment Platforms

Banks and payment platforms are the third and final party with the ability to deter the online libel ecosystem, other than directly involved stakeholders. Statistical fraud detection algorithms can be trained to report suspicious collections of transactions similar to those described in Table 2. In many cases, large irreversible transactions between parties with no previous contact may be considered suspicious.

6.4. For Libel Post Writers

We strongly recommend rational thinking and empathy. In most cases, regretful post writers cannot rescind published libel. Writers should first check to see if there are any other ways to handle the situation other than publishing identifiable details. We also urge caution about truthfulness and legal liability: are you completely sure that the content to be posted is true, and have you checked that this libel will not lead to a civil lawsuit in your jurisdiction? In most cases involving strongly emotional content, we urge rethinking whether it is necessary to post libel.

6.5. For the Intended Audience of Libel Sites

Libel sites’ intended audience includes anybody who performs a Google search as a background check, e.g. for

potential job applicants, apartment lessees, loan recipients, romantic partners found online, or service providers. We urge caution in believing these results as they are unverified and often irrelevant. It is also likely that images and names are taken that do not actually belong to the post subject. Moreover, potential viewership is what fuels the libel ecosystem, so discouraging the audience of libel sites would substantially affect this form of harassment.

6.6. For Libel Sites and Reputation Management Companies

Libel sites, if they wish to continue pursuing their goals, would highly benefit from content moderation. The derogatory terms used to describe marginalized populations such as BIPOC community members, sex workers, and members of the LGBTQ+ community harm those populations by perpetuating a culture that vilifies these identities and harms post subjects when they are involuntarily exposed. Furthermore, believing only the post writer, irreversibly placing blame on the subject, and bullying subjects fosters a sense of immaturity and attracts negative attention from news companies and other media. Attempts to circumvent de-indexing from search engines only serve to prove that libel sites are an unwanted plague in their current form.

For reputation management companies we recommend transparency and fostering a sense of legitimacy through verifiable facts. Many of the testimonials of past removals are easily provable as false, and supposed employees are available online as paid models. Websites that claim to have developed platforms with dashboards for monitoring libel posts could provide screenshots of the platforms to provide a sense of transparency to prospective clients without revealing trade secrets or company practices.

6.7. For Unwilling Post Subjects

It is less than ideal to place the burden of fixing this ecosystem on post subjects. However, we provide some recommendations for those who find themselves unwilling subjects of libel posts. Though it is tempting, we strongly advise against posting a reply to the libel post, even to state that the accusations are false. Based on our large sample of posts, this will only attract hateful discussion without having the post removed, and the increased activity will push the libel up in search engine results. The first step is to report the webpage of your individual post to search engines through forms such as [7]; this will only be successful if you are the subject of the post on the page in question. Next, if you wish to, you may follow the recommendations advertised on the libel site for post removal. This step is costly and often the results are incomplete; however, because of the “exploitative removal practices” [7] it is the only way to reduce the spread of this libel outside of search engines. Finally, we recommend prioritizing personal safety, practicing self-care, and finding comfort in the knowledge that libel sites themselves are the ones with poor reputations.

6.8. For Researchers and Journalists

Firstly, we urge caution in publicizing libel sites to avoid increasing their popularity. We do recommend further research and media attention surrounding the libel ecosystem as a whole to apply pressure to parties with the power to deter this form of public harassment.

7. Related Work

While this specific form of harassment to our knowledge has not been studied in academic research, we believe libel sites fall between Toxic Content and Content Leakage on the taxonomy of online harassment compiled by [4]. The individual posts on libel sites fit the narratives described in [15], specifically from the perspective of somebody “surviving” a partner’s infidelity. [16] studied incitements and calls to harassment in communities focused on harassing others and compiled a taxonomy of coordinated harassment; they found that the majority of coordinated harassment occurs through reporting the target to appropriate authorities, and made recommendations. [5] also made recommendations and suggestions based on a study interviewing people who had attempted to remove their data from people search websites, with the alarming finding that most experiences indicated that “removal is not final.” [17] compiled resources regarding digital abuse from a legal perspective, noting that it is “on the rise” and discussing “whether and how to regulate digital abuse.” [9] studied content moderation policies, surveying targets of harassment regarding policies that they wish had happened to the harasser, instead of the actual policies implemented; among their findings was that public shaming is imprecise in its targets, disproportionately harmful to any crimes committed, and ineffective at best.

Outside of academic research, the libel ecosystem has garnered concern organically [2], [11]. Sources such as [18] and [19] compiled lists of tips and resources for people whose personal information is available online without their consent. Journalists [10] performed an in-depth study of a small set of libel sites and reputation management services, with the probable conclusion that both are managed by the same person who hires employees that are unaware of the depth of harassment inherent in this ecosystem.

8. Conclusion

Widely available online libel harms people’s reputations and puts them in danger through publishing personal information along with strongly worded accusations, with no repercussions to the anonymous libel writers. Post removal is nearly impossible, with the only pathway through reputation management services that charge thousands of dollars or through incomplete removal via search engines. We collected and studied 9 libel sites, 7 websites for reputation management services, and 12 related websites in the online libel ecosystem. Our findings included how and why this ecosystem functions, how reputation management services’ policies for libel post removal contribute to this ecosystem, and why internet search engines’ policies are inefficient at addressing this extortion in their current

state. We discussed the exploitative nature of this system and made general suggestions for stakeholders and policy recommendations for parties with the power to mitigate this harassment.

Acknowledgment

This research was supported by NSF grants 1916126 and 2016061.

We gratefully acknowledge reviewer feedback and suggestions. As part of the open-report model followed by the Workshop on Attackers & CyberCrime Operations (WACCO), all the reviews for this paper are publicly available at <https://github.com/wacco-workshop/WACCO/tree/main/WACCO-2022>.

References

- [1] H. A. Popkin, "Website exposes 'homewreckers' — but doesn't break the law," 2013. [Online]. Available: <https://www.nbcnews.com/technology/website-exposes-homewreckers-doesnt-break-law-8c11554126>
- [2] ABC News, "Shesahomewrecker.com: Revenge website targets alleged mistresses," 2013. [Online]. Available: https://youtu.be/BNN_mt80Efo
- [3] H. Nissenbaum, "Privacy as contextual integrity," *Washington Law Review*, vol. 79, 2004. [Online]. Available: <https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10>
- [4] K. Thomas, D. Akhawe, M. Bailey, D. Boneh, E. Bursztein, S. Consolvo, N. Dell, Z. Durumeric, P. G. Kelley, D. Kumar, D. McCoy, S. Meiklejohn, T. Ristenpart, and G. Stringhini, "Sok: Hate, harassment, and the changing landscape of online abuse," in *2021 IEEE Symposium on Security and Privacy (SP)*. San Francisco, CA, USA: IEEE, 2021, pp. 247–267. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9519435>
- [5] K. Take, K. Gallagher, A. Forte, D. McCoy, and R. Greenstadt, "'it feels like whack-a-mole': User experiences of data removal from people search websites," in *Proceedings on Privacy Enhancing Technologies (PoPETs)*. Sydney, Australia: Privacy Enhancing Technologies Symposium (PETS), 2022.
- [6] Electronic Frontier Foundation, "Section 230 of the communications decency act," 1996. [Online]. Available: <https://www.eff.org/issues/cda230>
- [7] Google, "Remove content about me on sites with exploitative removal practices from google," 2022. [Online]. Available: <https://support.google.com/websearch/answer/9172218>
- [8] M. Earle, "What happened to thedirty.com? was their website shut down?" 2021. [Online]. Available: <https://www.reputation.ca/what-happened-to-thedirty/>
- [9] S. Yardi, O. Haimson, and L. Nakamura, "Drawing from justice theories to support targets of online harassment," *New Media and Society*, vol. 23, no. 5, pp. 1278–1300, 2020. [Online]. Available: <https://doi.org/10.1177/1461444820913122>
- [10] A. Krolik and K. Hill, "The slander industry," 2021. [Online]. Available: <https://www.nytimes.com/interactive/2021/04/24/technology/online-slander-websites.html>
- [11] Women24, "She's a home wrecker discussion with dorothy black," 2014. [Online]. Available: <https://www.youtube.com/watch?v=tONSXpMsXPY>
- [12] Y. Genc, H. Kour, H. T. Arslan, and L.-C. Chen, "Understanding nigerian e-mail scams: A computational content analysis approach," *Information Security Journal: A Global Perspective*, vol. 30, no. 2, pp. 88–99, 2021. [Online]. Available: <https://doi.org/10.1080/19393555.2020.1804647>
- [13] Proton Technologies AG, "Protonmail," 2022. [Online]. Available: <https://protonmail.com/>
- [14] K. Hill and D. Wakabayashi, "Google seeks to break vicious cycle of online slander," 2021. [Online]. Available: <https://www.nytimes.com/2021/06/10/technology/google-algorithm-known-victims.html>
- [15] R. Bellini, E. Tseng, N. McDonald, R. Greenstadt, D. McCoy, T. Ristenpart, and N. Dell, "'so-called privacy breeds evil': Narrative justifications for intimate partner surveillance in online forums," *Proceedings of the ACM on Human-Computer Interaction*, vol. 4, no. CSCW3, jan 2021. [Online]. Available: <https://doi.org/10.1145/3432909>
- [16] M. Aliapoulos, K. Take, P. Ramakrishna, D. Borkan, B. Goldberg, J. Sorensen, A. Turner, R. Greenstadt, T. Lauinger, and D. McCoy, "A large-scale characterization of online incitements to harassment across platforms," in *Proceedings of the 21st ACM Internet Measurement Conference*. New York, NY, USA: Association for Computing Machinery, 2021, p. 621–638. [Online]. Available: <https://dl.acm.org/doi/abs/10.1145/3487552.3487852>
- [17] T. Kadri, "Digital abuse," 2022. [Online]. Available: <https://www.thomaskadri.com/digital-abuse/syllabus>
- [18] M. Bazzell, "Inteltechniques," 2021. [Online]. Available: <https://inteltechniques.com/workbook.html>
- [19] yaelwrites, "Big ass data broker opt-out list," 2022. [Online]. Available: <https://github.com/yaelwrites/Big-Ass-Data-Broker-Opt-Out-List>