

**A revealing picture of how personal health information searches become the property of private corporations.**

BY TIMOTHY LIBERT

# Privacy Implications of Health Information Seeking on the Web

PRIVACY ONLINE IS AN increasingly popular field of study, yet it remains poorly defined. “Privacy” itself is a word that changes according to location, context, and culture. Additionally, the Web is a vast landscape of specialized sites and activities that may only apply to a minority of users—making defining widely shared privacy concerns difficult. Likewise, as technologies and services proliferate, the line between on- and offline is increasingly blurred. Researchers attempting to make sense of this rapidly changing environment are frequently stymied by such factors.

Therefore, the ideal object of study is one that is inherently sensitive in nature, applies to the majority of users, and readily lends itself to analysis. The study of health privacy on the Web meets all of these criteria.

Health information has been regarded as sensitive since the time of the ancient Greeks. In the 5<sup>th</sup> century B.C., physicians taking the Hippocratic Oath were required to swear that: Whatever I see or hear in the lives of my patients...I will keep secret, as considering all such things to be private.<sup>21</sup> This oath is still in use today, and the importance of health privacy remains universally recognized. However, as health-information seeking has moved online, the privacy of a doctor’s office has been traded in for the silent intrusion of behavioral tracking. This tracking provides a valuable vantage point from which to observe how established cultural norms and technological innovations are at odds.

Online health privacy is an issue that affects the majority of Internet users. According to the Pew Research Center, 72% of adult Internet users in the U.S. go online to learn about medical conditions.<sup>9</sup> Yet only 13% of these begin their search at health-specific sites. In fact, health information may be found on a wide spectrum of sites ranging from newspapers, discussion forums, to research institutions. In order to discover the full range of sites users may visit when seeking health information, I used a search engine to

## » key insights

- **Over 90% of the 80,142 health-related Web pages initiate HTTP requests to third-parties, oftentimes outside the view of the user.**
- **Some 70% of third-party requests transmit information on specific symptoms, treatments, and diseases in the URI string.**
- **Page visitors are at risk of their health interests being publicly identified as well as being blindly discriminated against by marketers.**
- **Extant policy and legal protections are few in number and weak in effect, demonstrating a need for interventions.**

...not - Google Search www...  
...flu - Google Search www.google...  
...std testing nyc - Google Search...  
...std information hotline - Google S...  
...https://www.google.com/webhp?sc...  
...std information - Google Search...  
...stds testing - Google Search www...  
...stds list - Google Search www.go...  
...stds - Google Search www.google...  
...std - Google Search www.google.c...  
...hiv medicine cost - Google Search...  
...hiv medicine - Google Search...  
...hiv medicine research - Google ww...  
...hiv research - Google Search...  
...hiv - Google Search www.google...  
...hiv treatment - Google Search...  
...hiv treatment cost - Google Se...  
...hiv symptoms - Google Search...  
...hiv information - Google Search...  
...hiv testing - Google Search...  
...hiv testing nyc - Google Search...  
...chemotherapy side effects - Go...  
...chemotherapy drugs - Google...  
...chemotherapy hair loss - Go...  
...chemotherapy informati...  
...lung cancer prognosis - Google...  
...lung cancer stage 4 - G...  
...lung cancer symptoms...  
...lung cancer stages - Go...  
...bronchitis curable - Go...  
...bronchitis cure - Goog...  
...bronchitis symptoms in...  
...bronchitis symptoms -...  
...bronchitis contagious -...  
...bronchitis home reme...  
...flu - Google Search...  
...flu symptoms - Go

...ung -...  
...list - Go...  
...Google S...  
...Google Sea...  
...medicine Sea...  
...medicine cost...  
...medicine - Go...  
...research - Go...  
...Google Search...  
...treatment - Go...  
...treatment cost - Goog...  
...symptoms - Google Sea...  
...information - Google Search...  
...ing - Google Search...  
...nyc - Goole

...enitis ne...  
...flu - Google Search...  
...flu symptoms - Google...  
...flu symptoms in kids - Go...  
...flu symptoms in children - G...  
...https://www.google.com/cancer...  
...flu shots - Google Search www...  
...flu shots nyc - Google Search...  
...flu shot cost - Google Search...  
...flu shot side effects - Google Sear...  
...flu shot safety - Google Search...  
...http://www.cancer.org/cancer/pa...  
...flu shot side effects - Google Sea...  
...flu shot - Google Search www...  
...std testing nyc - Google Search...  
...std information hotline - Google...  
...https://www.google.com/webhp...  
...testing - Google Search...  
...list - Google Search www...  
...Google Search...  
...medic

identify 80,142 unique health-related Web pages by compiling responses to queries for 1,986 common diseases. This selection of pages represents what users are actually visiting, rather than a handful of specific health portals.

Having identified a population of health-related Web pages, I created a custom software platform to monitor the HTTP requests initiated to third parties. I discovered that 91% of pages make requests to additional parties, potentially putting user privacy at risk. Given that HTTP requests often include the URI of the page currently being viewed (known as the “Referer” [sic]), information about specific symptoms, treatments, and diseases may be transmitted. My analysis shows 70% of URIs contains such sensitive information.

This proliferation of third-party requests makes it possible for corporations to assemble dossiers on the health conditions of unwitting users. In order to identify which corporations are the recipients of this data I have also analyzed the ownership of the most requested third-party domains. This has produced a revealing picture of how personal health information becomes the property of private corporations.

This article begins with a short primer on how third-party HTTP requests

work, reviews previous research in this area, details methodology and findings, and concludes with suggestions for protecting health privacy online.

**Background: Third-Party HTTP Requests**

A real-world example is the best way to understand how the information is leaked to third parties on a typical Web page. When a user searches online for “HIV” one of the top results is for the U.S. Centers for Disease Control and Prevention (CDC) page with the address <http://www.cdc.gov/hiv/>.<sup>a</sup> Clicking on this result initiates what is known as a “first-party” Hypertext Transfer Protocol (HTTP) request to the CDC Web server (Figure 1.1). A portion of such a request is as follows:

```
GET /hiv/
Host: www.cdc.gov
User-Agent: Mozilla/5.0 (Macintosh...
```

This request is sent to the CDC Web server (“Host: www.cdc.gov”) and is an instruction to return (“GET”) the page with the address “/hiv/.” This request also includes “User-Agent” information that tells the server what kind

<sup>a</sup> As of April, 2014

of browser and computer the user is on. In this case, the user employs the Mozilla Firefox browser on a Macintosh computer. Such information is helpful when loading specially optimized pages for smartphones or tablets.

Once this request has been made, the CDC Web server sends the user an HTML file. This file contains the text of the page as well as a set of instructions that tells the Web browser how to download and style additional elements such as images (Figure 1.2). In order to get the CDC logo, the following HTTP request is made:

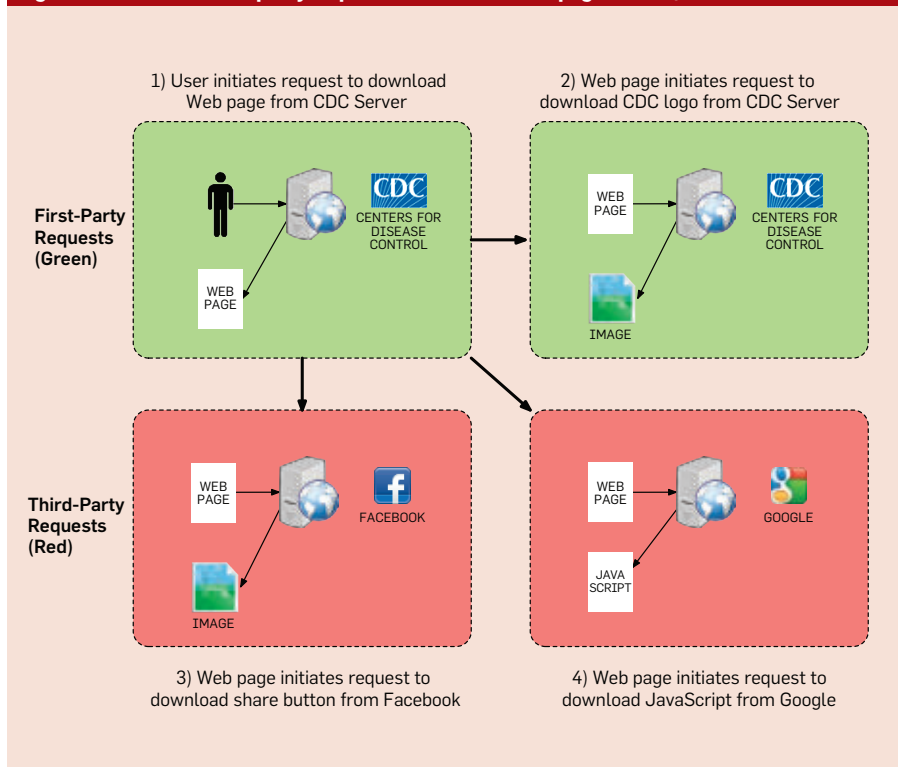
```
GET /TemplatePackage/images/cdcHeaderLogo.gif
Host: www.cdc.gov
User-Agent: Mozilla/5.0 (Macintosh...
Referer: http://www.cdc.gov/hiv/
```

This request introduces a new piece of information called the Referer, which contains the address of the page the user is currently viewing. The CDC Web server may keep records of all HTTP requests in order to determine what pages and content are being requested most often.

Because the “Host” for both requests is identical (www.cdc.gov), the user is only interacting with a single party and such requests are called “first-party requests.” The only two parties who know the user is looking up information about HIV are the user and the CDC. However, the HTML file also contains code that makes requests to outside parties. These types of third-party requests typically download third-party elements such as images and JavaScript. Due to the fact that users are often unaware of such requests, they form the basis of the so-called “Invisible Web.”

On the CDC’s HIV page, third-party requests are made to the servers of Facebook, Pinterest, Twitter, and Google. In the case of the first three companies, the requested elements are all social media buttons, which allow for the sharing of content via the “Recommend,” “Tweet,” or “Pin It” icons (Figure 1.3). It is unlikely that many users would understand the presence of these buttons indicates that their data is sent to these companies. In contrast, the Google elements on the page are entirely invisible and there is no Google logo present. One of

**Figure 1. First- and third-party requests on the CDC Web page for HIV/AIDS.**



these requests is sent to Google's Analytics service (Figure 1.4) to download a file containing JavaScript code:

```
GET /ga.js
Host: www.google-analytics.com
User-Agent: Mozilla/5.0 (Macintosh...
Referer: http://www.cdc.gov/hiv/
```


Again, the Referer field reveals the user is visiting a page about HIV. By pairing information about the User-Agent, Referer, and user's IP address, it is possible for companies like Google and Facebook to identify people who are concerned with HIV.<sup>34</sup> Those visiting this page likely are unaware of this fact, and would not be happy to find out.

### Prior Research


Prior research has demonstrated that while users are uncomfortable with this type of tracking, it is performed in a number of highly sophisticated ways, and it is increasingly widespread.

**Attitudes.** There has long been anxiety about how personal data will be used on the Web. A 1999 study determined that "only 13% of respondents reported they were 'not very' or 'not at all' concerned" about their privacy online.<sup>2</sup> Such anxiety remained in 2003 when 70% of survey respondents reported they were nervous that websites had information about them.<sup>29</sup> A 2009 follow-up study revealed that 67% of respondents agreed with the statement they had "lost all control over how personal information is collected and used by companies."<sup>30</sup> These surveys demonstrate the activities of many businesses run directly counter to public preferences.

As with general concerns with online privacy, there is excellent research exploring attitudes toward health information. In 2012, Hoofnagle et al. determined that only 36% of survey respondents knew that advertisers are allowed to track their visits to health-related websites.<sup>12</sup> An extensive study from the year 2000 found that 85% of Internet users in poor health were concerned that websites would share their data, and only 3% were comfortable with websites sharing their data with other sites, companies, and advertisers.<sup>10</sup> Despite these fears, 44% of respondents felt their information was safe



**The proliferation of third-party requests makes it possible for corporations to assemble dossiers on the health conditions of unwitting users.**



with institutions such as the National Institutes of Health (NIH).<sup>10</sup> The CDC example detailed earlier indicates this trust is potentially misplaced.

**Mechanisms.** Once a third-party request is made, a user may be tracked using a number of ever-evolving technical mechanisms. Researchers have been tracing the development of such mechanisms for years, often analyzing the code and behaviors that take place within the Web browser. These are often called "client-side" techniques for they take place on the user's computer. Traditional client-side techniques typically involve storing data on the user's computer in small text files known as cookies—this functions as a sort of digital name tag.<sup>3</sup> Users are getting more adept at evading such practices, therefore newer techniques often employ so-called "browser fingerprinting" to identify users based on characteristics of their computers. This area of research has proven very popular of late with numerous studies investigating fingerprinting techniques.<sup>1,7,13,14,23</sup> In addition, Miller et al. have recently demonstrated sophisticated attacks on HTTPS that are able to reveal "personal details including medical conditions."<sup>20</sup>

Turning attention to the server-side, Yen et al. have recently demonstrated a tracking technique that utilizes a combination of IP address, User-Agent string, and time intervals when HTTP requests were made. This team was able to identify users 80% of the time, which is on par with what is typically accomplished with client-side cookies.<sup>34</sup> Furthermore, identification rates remained essentially static even when removing the final octet of the IP address, which is a common technique by which major advertisers claim to anonymize data. Yen et al.'s findings indicate that while novel techniques may be needed on the client-side, the lowly HTTP request is sufficient for advanced server-side techniques.

**Measurement.** The final area of related research is measurement. Measurement of Web tracking generally entails two steps: selecting a population of pages, and performing automated analysis of how user data is transmitted to third parties. Many studies have relied on popular site lists provided by the Alexa

company,<sup>4,17,18,25</sup> but often utilize their own methodologies for analysis. Krishnamurthy and Wills have conducted many of the most important studies in this area<sup>18</sup> and developed the idea of a privacy footprint<sup>17</sup> based upon the number of nodes a given user is exposed to as they surf the Web. This team has consistently found there are high levels of tracking on the Web, including on sites dealing with sensitive personal information such as health.<sup>17</sup> Other teams have performed comparative analyses between countries<sup>4</sup> as well as explored general trends in tracking mechanisms.<sup>19,25</sup> A common theme among all measurement research is the amount of tracking on the Web is increasing, and shows no signs of abating. The data presented in this article updates and advances extant findings with a focus on how users are tracked when they seek health information online.

### Methodology

In order to quickly and accurately reveal third-party HTTP requests on health-related Web pages, my methodology has four main components: page selection, third-party request detection, request analysis, and corporate ownership analysis.

**Page selection.** A variety of websites such as newspapers, government agencies, and academic institutions provide health information online. Thus, limiting analysis to popular health-centric sites fails to reach many of the sites users actually visit.<sup>16</sup> To wit, the Pew Internet and American Life Project found “77% of online health seekers say they began at a search engine such as Google, Bing, or Yahoo”<sup>9</sup> as opposed to a health portal like WebMD.com. In order to best model the pages a user would visit after receiving a medical diagnosis, I first compiled a list of 1,986 diseases and conditions based on data from the Centers for Disease Control, the Mayo Clinic, and Wikipedia. Next, I used the Bing search API in order to find the top 50 search results for each term.<sup>b</sup> Once duplicates and binary files (pdf, doc, xls) were filtered out, a set of 80,142 unique Web pages remained. A major contribution of this study to prior work is the fact that my analysis is focused on the pages that users seeking medical information

<sup>b</sup> Search results were localized to U.S./English.



**Prior research has demonstrated that while users are uncomfortable with this type of tracking, it is performed in a number of highly sophisticated ways, and it is increasingly widespread.**



are most likely to visit, irrespective of if the site is health-centric.

**Third-party request detection.** To detect third-party HTTP requests, my methodology employs a “headless” Web browser named PhantomJS.<sup>24</sup> PhantomJS requires no GUI, has very low resource utilization, and is therefore well suited for large-scale analyses. Due to the fact it is built on WebKit, PhantomJS’s underlying rendering engine is capable of executing JavaScript, setting and storing cookies, and producing screen captures. Most important for this project, PhantomJS allows for the direct monitoring of HTTP requests without the need to resort to browser hacks or network proxies.

It should be noted that the most recent versions of PhantomJS (1.5+) do not support the Adobe Flash browser plug-in. To address this potential limitation, I conducted testing with an older version of PhantomJS (1.4) and Flash. The inclusion of Flash led to much higher resource utilization, instability, and introduced a large performance penalty. While this method successfully analyzed Flash requests, I determined that Flash elements were comparatively rare and had negligible effect on the top-level trends presented below. Therefore, I made the decision to forgo analysis of Flash requests in favor of greater software reliability by using the most recent version of PhantomJS (1.9).

In order to fully leverage the power of PhantomJS, I created a custom software platform named WebXray that drives PhantomJS, collects and analyzes the output in Python, and stores results in MySQL. The workflow begins with a predefined list of Web page addresses that are ingested by a Python script. PhantomJS then loads the given Web address, waits 30 seconds to allow for all redirects and content loading to complete, and sends back JSON-formatted output to Python for analysis. This technique represents an improvement over methods such as searching for known advertising elements detected by popular programs such as Ghostery or AdBlock.<sup>4</sup> As of March 2014, Ghostery reports the WebMD Web page for “HIV/AIDS” contains four trackers. In contrast, WebXray detects the same page initiating requests to thirteen distinct third-party domains.

This is due to the fact that Ghostery and Adblock rely on curated blacklists of known trackers, rather than reporting all requests.

**Request analysis.** The primary goal of WebXray is to identify third-party requests by comparing the domain of the Web page being visited to the domains of requests being made. For example, the address “http://example.com” and the request “http://images.example.com/logo.png” both share the domain “example.com,” thus constituting a first-party request.

Alternately, a request from the same page to “http://www.googleanalytics.com/ga.js,” which has the domain “google-analytics.com,” is recognized as a third-party request. The same technique for HTTP requests is also applied toward evaluating the presence of third-party cookies. The method is not flawless, as a given site may actually use many domains, or a subdomain may point to an outside party. However, when evaluating these types of requests in aggregate, such problems constitute the statistical noise that is present in any large dataset.

Finally, in order to evaluate larger trends in tracking mechanisms, third-party requests are dissected to extract arguments (for example, “?SITEID=123”) and file extensions such as .js (JavaScript), .jpg (image), and .css (cascading style sheet).

Removing arguments also allows for a more robust analysis of which elements are the most prevalent, as argument strings often have specific site identifiers, making them appear unique when they are not.

**Corporate ownership.** A specific focus of this investigation is to determine which corporate bodies are receiving information from health-related Web pages. While it is possible to programmatically detect requests to third-party domains, it is not always clear who belongs to the requested domains. By examining domain registration records, I have been able to pair seemingly obscure domain names (for example, “2mdn.net,” “fbcdn.net”) with their corporate owners (for example, Google, Facebook). This process has allowed me to follow the data trail back to the corporations that are the recipients of user data. To date, the literature has given much more attention to technical

mechanisms, and much less to the underlying corporate dynamics. This fresh analytical focus highlights the power of a handful of corporate giants.

**Limitations.** While this methodology is resource efficient and performs well at large scale, it comes with several potential limitations, many of which would produce an under-count of the number of third-party requests. First, given the rapid rate by which pages are accessed, it is possible that rate-limiting mechanisms on servers may be triggered (that is, the requests generated by my IP would be identifiable as automated), and my IP address could be blacklisted, resulting in an under-count. Second, due to the fact I use PhantomJS without browser plugins such as Flash, Java, and Silverlight, some tracking mechanisms may not load or execute properly, resulting in an under-count. Third, many tracking mechanisms are designed to be difficult to detect by a user, and an under-count could result from a failure to detect particularly clever tracking mechanisms. Therefore, the findings presented here constitute a lower bound of the amount of requests being made.

## Findings


In April 2014, I scanned 80,142 Web pages that were collected from search results for 1,986 common diseases with the intent of detecting the extent and the ways in which the sensitive health data of users was being leaked.

**General trends.** I have broken up my top-level findings into five general categories based on information gleaned from the TLDs used. They are: all pages,

commercial pages (.com), non-profit pages (.org), government pages (.gov), and education-related pages (.edu). This information is illustrated in Figure 2. Of all pages examined, 91% initiate some form of third-party HTTP request, 86% download and execute third-party JavaScript, and 71% utilize cookies. Unsurprisingly, commercial pages were above the global mean and had the most third-party requests (93%), JavaScript (91%), and cookies (82%). Education pages had the least third-party HTTP requests (76%) and JavaScript (73%), with a full quarter of the pages free of third-party requests. Government pages stood out for relatively low prevalence of third-party cookies, with only 21% of pages storing user data in this way. Figure 2 details these findings.

**Mechanisms.** Given that 91% of pages make third-party HTTP requests, it is helpful to know what exactly is being requested. Many third-party requests lack extensions, and when viewed in a browser display only blank pages that generate HTTP requests and may also manipulate browser caches. Such requests accounted for 47% of the top 100 requests and may point toward emerging trends in the ongoing contest between user preferences and tracking techniques. The second most popular type of requested elements were JavaScript files (33%). These files are able to execute arbitrary code in a user’s browser and may be used to perform fingerprinting techniques, manipulate caches and HTML5 storage, as well as initiate additional requests. The third most popular type of content is the tried-and-true image file, which ac-

Figure 2. Prevalence of third-party requests, JavaScript, and cookies by TLD.



counts for 8% of the top requested elements. Table 1 presents additional detail into the file extensions found.

Given that tracking occurs on the so-called Invisible Web, it initially appears odd that so many mechanisms are images. However, when investigating the images themselves, it is clear they provide little indication as to whom they belong to, and thus users are kept in the dark as to their purpose or presence. An examination of the top 100 requested images determined that only 24% contained information that would alert the user they had initiated contact with a third party. Many images were only a single pixel in size, and are often referred to as tracking pixels as their only purpose is to initiate HTTP requests. The most popular image, found on 45% of pages, was a single tracking pixel with the name `utm.gif`, which is part of the Google Analytics service. The second most popular image is the clearly identifiable Facebook “Like” button that was found on 16% of pages. It is unclear how many users elect to “Like” an illness, but Facebook is able to record page visits regardless if a user clicks the “Like” button, or if they

even have a Facebook account in the first place. Google and Facebook are not alone, however, there are a number of companies tracking users online.

**Corporate ownership.** While security and privacy research has often focused on how user privacy is violated, insufficient attention has been given to who is collecting user information. The simple answer is that a variety of advertising companies have developed a massive data collection infrastructure that is designed to avoid detection, as well as ignore, counteract, or evade user attempts at limiting collection. Despite the wide range of entities collecting user data online, a handful of privately held U.S. advertising firms dominate the landscape of the Invisible Web.

Some 78% of pages analyzed included elements that were owned by Google. Such elements represent a number of hosted services and use a variety of domain names: they range from traffic analytics (`google-analytics.com`), advertisements (`doubleclick.net`), hosted JavaScript (`googleapis.com`), to videos (`youtube.com`). Regardless of the type of services provided, in some way all of these HTTP requests funnel information back to Google. This means a single company has the ability to record the Web activity of a huge number of individuals seeking sensitive health-related information without their knowledge or consent.

While Google is the elephant in the room, they are far from alone. Table 2 details the top 10 firms found as part

of this analysis along with the rankings of two data brokers. In second place is comScore who are found on 38% of pages, followed by Facebook with 31%. It is striking that these two companies combined still have less reach than Google.

Additionally, companies were categorized according to their type of revenue model. Some 80% of the top 10 companies are advertisers. The only exceptions to this rule are Adobe and Amazon. Adobe offers a mix of software and services, including traffic analytics. Amazon is in the business of both consumer-retail sales as well as Web hosting with the Amazon Web Services (AWS) division. At present it is unclear if AWS data is integrated into Amazon product recommendations or deals, but the possibility exists.

While advertisers dominate online tracking, I was also able to detect two major data brokers: Experian (5% of pages), and Acxiom (3% of pages). The main business model of data brokers is to collect information about individuals and households in order to sell it to financial institutions, employers, marketers, and other entities with such interest. Credit scores provided by Experian help determine if a given individual qualifies for a loan, and if so, at what interest rate. Given that a 2007 study revealed that “62.1% of all bankruptcies ... were medical,”<sup>11</sup> it is possible that some data brokers not only know when a given person suffered a medical-related bankruptcy, but perhaps even when they first searched for information on the ailment that caused their financial troubles.

**Health information leakage.** The HTTP 1.1 protocol specification warns the source of a link [URI] might be private information or might reveal an otherwise private information source and advises that “[c]lients SHOULD NOT include a Referer header field in a (non-secure) HTTP request if the referring page was transferred with a secure protocol.”<sup>8</sup> In simpler terms, Web pages that include third-party elements, but do not use secure HTTP requests, risk leaking sensitive data via the Referer field. Of the pages analyzed, only 3.24% used secure HTTP, the rest used non-encrypted HTTP connections and thereby potentially transmitted sensitive information to third parties. Unsurprisingly, a significant amount of

**Table 1. Types of file extensions.**

Type	%
No Extension	47
JavaScript	33
Image	8
Dynamic Page	4
Other	8

**Table 2. Corporate ownership and risk assessment (N=80,142).**

Rank	% Pages	Company	Revenue	Identification	Blind Discrimination
1	78	Google	Advertising	X	X
2	38	comScore	Advertising	—	X
3	31	Facebook	Advertising	X	X
4	22	AppNexus	Advertising	—	X
5	18	Add This	Advertising	—	X
6	18	Twitter	Advertising	—	X
7	16	Quantcast	Advertising	—	X
8	16	Amazon	Retail and Hosting	—	X
9	11	Adobe	Software and Services	—	X
10	11	Yahoo!	Advertising	—	X
...	—	—	—	—	—
31	5	Experian	Data Broker	X	—
...	—	—	—	—	—
47	3	Acxiom	Data Broker	X	—

sensitive information was included in URI strings.

Based on a random sample of 500 URIs taken from the population of pages analyzed ( $N=80,142$ ), 70% contained information related to a specific symptom, treatment, or disease. An example of an URI containing specific symptom information is:

[http://www.nhs.uk/conditions/breast-lump/...](http://www.nhs.uk/conditions/breast-lump/)

a URI containing no such information is:


<http://www.ncbi.nlm.nih.gov/pubmed/21722252>

Given the former type of URI was by far the most prevalent, it may be seen that third parties are being sent a large volume of sensitive URI strings that may be analyzed for the presence of specific diseases, symptoms, and treatments. This type of leakage is a clear risk for those who wish to keep this information out of the hands of third parties who may use it for unknown ends.


## Discussion

Defining privacy harms is a perennially difficult proposition. Health information, however, presents two main privacy risks that are interrelated. The first is personal identification, where an individual's name is publicly associated with their medical history. The second is blind discrimination, where an individual's name is not necessarily revealed, but they may be treated differently based on perceived medical conditions.

**Personal identification.** While most people would probably consider details of their health lives to be of little interest or value to others, such details form the basis of a lucrative industry. In 2013, the U.S. Senate Committee on Commerce, Science and Transportation released a highly critical review of the current state of the so-called data broker industry. Data brokers collect, package, and sell information about specific individuals and households with virtually no oversight. This data includes demographic information (ages, names, and addresses), financial records, social media activity, as well as information on those who may be suffering from "particular ailments, in-



**While security and privacy research has often focused on how user privacy is violated, insufficient attention has been given to who is collecting user information.**



cluding Attention Deficit Hyperactivity Disorder, anxiety, depression ... among others."<sup>26</sup> One company, Medbase200, was reported as using proprietary models to generate and sell lists with classifications such as rape victims, domestic abuse victims, and HIV/AIDS patients.<sup>6</sup>

It should also be noted that such models are not always accurate. For example, individuals looking for information on the condition of a loved one may be falsely tagged as having the condition themselves. This expands the scope of risk beyond the patient to include family and friends. In other cases, an individual may be searching for health information out of general interest and end up on a data broker's list of sufferers or patients. Common clerical and software errors may also tag individuals with conditions they do not have. The high potential for such errors also highlights the need for privacy protections.

Furthermore, criminals may abuse poorly protected health information. The retailer Target has used datamining techniques to analyze customers' purchase history in order to predict which women may be pregnant in order to offer them special discounts on infant-related products.<sup>5</sup> Even if shoppers and surfers are comfortable with companies collecting this data, that is no guarantee it is safe from thieves. In 2013, 40 million credit and debit card numbers were stolen from Target.<sup>15</sup> While a stolen credit card may be reissued, if Target's health-related data were leaked online, it could have a devastating impact on millions of people. Merely storing personally identifiable information on health conditions raises the potential for loss, theft, and abuse.

**Blind discrimination.** Advertisers regularly promise their methods are wholly anonymous and therefore benign, yet identification is not always required for discriminatory behavior to occur. In 2013, Latanya Sweeney investigated the placement of online advertisements that implied a given name was associated with a criminal record.<sup>27</sup> She found the presence of such ads were not the result of particular names being those of criminals, but appeared based on the racial associations of the name, with African-American names more often resulting in an implication of criminal record. In this way, extant societal injustices may be replicated



through advertising mechanisms online. Discrimination against the ill may also be replicated through the collection and use of browsing behavior.

Data-mining techniques often rely on an eclectic approach to data analysis. In the same way a stew is the result of many varied ingredients being mixed in the same pot, behavioral advertising is the result of many types of browsing behavior being mixed together in order to detect trends. As with ingredients in a stew, no single piece of data has an overly deterministic impact on the outcome, but each has some impact. Adding a visit to a weather site in the data stew will have an outcome on the offers a user receives, but not in a particularly nefarious way. However, once health information is added to the mix, it becomes inevitable it will have some impact on the outcome. As medical expenses leave many with less to spend on luxuries, these users may be segregated into data silos<sup>28</sup> of undesirables who are then excluded from favorable offers and prices. This forms a subtle, but real, form of discrimination against those perceived to be ill.

**Risk assessment.** Having collected data on how much tracking is taking place, how it occurs, and who is doing it, it is necessary to explicate how this constitutes a risk to users. As noted earlier, there are two main types of harm: identification and blind discrimination. Table 2 shows a breakdown of how data collection by 12 companies (top 10 and data brokers) impacts the two types of risk. The two data brokers most obviously entail a personal identification risk as their entire business model is devoted to selling personal information. It is unlikely they are selling raw Web tracking data directly, but it may be used as part of aggregate measures that are sold.

Despite the fact that Google does not sell user data, they do possess enough anonymous data to identify many users by name. Google offers a number of services that collect detailed personal information such as a user's personal email (Gmail), work email (Apps for Business), and physical location (Google Maps). For those who use Google's social media offering, Google+, a real name is forcefully encouraged. By combining the many types of information held by

Google services, it would be fairly trivial for the company to match real identities to anonymous Web browsing data. Likewise, Facebook requires the use of real names for users, and as noted before, collects data on 31% of pages; therefore, Facebook's collection of browsing data may also result in personal identification. In contrast, Twitter allows for pseudonyms as well as opting-out of tracking occurring off-site.

The potential for blind discrimination is most pronounced among advertisers. As noted here, online advertisers use complex data models that combine many pieces of unrelated information to draw conclusions about anonymous individuals. Any advertiser collecting and processing health-browsing data will use it in some way unless it is filtered and disposed of.

### Policy Implications

The privacy issues raised by this research are of a technical nature and invite technical solutions. These solutions often come in the form of add-on software users may install in their Web browsers. Such browser add-ons have proven effective at blocking certain types of behavioral tracking.<sup>19,25</sup> However, this type of solution places a burden on users and has not been broadly effective. As measurement research has shown, tracking has only increased over the past decade despite technical efforts to rein it in.

Purely technical solutions are problematic, as they require a relatively high level of knowledge and technical expertise on the part of the user. The user must first understand the complex nature of information flows online in order to seek out technical remedies. Next, the user must be proficient enough to install and configure the appropriate browser additions. This may seem trivial for the well educated, but many who use the Internet have little education or training in computing. Despite this, these users deserve to have their health privacy protected.

Furthermore, add-ons are often unavailable on the default browsers of smartphones and tablets, making it difficult for even the highly skilled to protect their privacy. A final reason that browser add-ons provide insufficient remedy is the fact that advertisers devote significant resources to

overcoming such barriers and will always find creative ways to bypass user intent. Thus, on one hand we have users who are poorly equipped to defend themselves with available technical measures, and on the other, highly motivated and well-funded corporations with cutting-edge technologies.

In order to effectively tackle the issue of tracking on health-related pages, attention toward the underlying social dynamics is needed. Government and corporate policies formalize these dynamics. By addressing policy issues directly, rather than combating obscure tracking techniques, we may produce durable solutions that outlast today's technology cycle. Unfortunately, extant policies are few in number and weak in effect.

### Extant policies and protections.

Health information is one of the few types of personal information that has been granted special protections. The Health Insurance Portability and Accountability Act (HIPAA)<sup>31</sup> is a U.S. law that stipulates how medical information may be handled, stored, and accessed. HIPAA is not meant to police business practices in general; rather it is tailored to those providing health-specific services such as doctors, hospitals, and insurance claims processors. Yet, even within this realm, HIPAA provides incomplete protections. Contrary to popular perceptions, HIPAA permits the disclosure of patient information between health providers and insurance claims processors without patient notification or consent. HIPAA generally does not allow patients to restrict the flow of their sensitive data; therefore, extending HIPAA in the online domain does not present an effective approach to privacy protection.

Nevertheless, the U.S. Federal Trade Commission (FTC) has established a Health Breach Notification Rule that requires entities holding personally identifiable health records to notify users if such records have been stolen.<sup>32</sup> However, merely providing health information (rather than storing doctor's notes or prescription records) does not place a business under the jurisdiction of HIPAA or associated rules. Many businesses that handle health information are subject to virtually no oversight and the main source of policy regarding the use of health informa-

tion online comes in the form of self-regulation by the parties that stand to benefit the most from capturing user data: online advertisers.

However, self-regulation has proven wholly insufficient. No lesser authority than the FTC determined that “industry efforts to address privacy through self-regulation have been too slow, and up to now have failed to provide adequate and meaningful protection.”<sup>33</sup> When self-regulations are present, there are no serious sanctions for violating the rules that advertisers draw up among themselves. Nevertheless, the Network Advertising Initiative (NAI) has produced a Code of Conduct that requires opt-in consent for advertisers to use precise information about health conditions such as cancer and mental health.<sup>22</sup> Yet the same policy also states that “member companies may seek to target users on the basis of such general health categories as headaches.”<sup>22</sup> Given the range of ailments between cancer and a headache is incredibly broad, this directive provides virtually no oversight. Likewise, the Digital Advertising Alliance (DAA) provides rules that also appear to protect health information, but legal scholars have determined that “an Internet user searching for information about or discussing a specific medical condition may still be tracked under the DAA’s principles.”<sup>12</sup>

**Potential interventions.** Although this problem is complex, it is not intractable and there are several ways health privacy risks may be mitigated. First, there is no reason for non-profits, educational institutions, or government-operated sites to be leaking sensitive user information to commercial parties. While advertising revenue keeps commercial sites running, non-profits gain support from donors and grants. Fixing this situation could be as simple as an internal policy directive on a per-institution basis, or as expansive as adopting language that would deny funding to institutions that leak user data.

As for commercial-oriented sites, it is true they rely on ad-tracking revenue. However, regulatory and legislative bodies have the authority to draft and implement policies that would require a mandatory limitation on how long information from health-related websites could be retained and how it could

be used. Such policy initiatives could have significant impact, and would reflect the preferences of the public.

Finally, talented engineers may devote a portion of the time they spend analyzing data to developing intelligent filters to keep sensitive data quarantined. The spark of change could be the result of a single engineer’s 20% time project. If the mad rush to ingest ever more data is tempered with a disciplined approach to filtering out potentially sensitive data, businesses and users may both benefit equally.

## Conclusion

Proving privacy harms is always a difficult task. However, this study has demonstrated that data on health information seeking is being collected by an array of entities that are not subject to regulation or oversight. Health information may be inadvertently misused by some companies, sold by others, or even stolen by criminals. By recognizing that health information deserves to be treated with special care, we may mitigate what harm may already be occurring and proactively avoid future problems.

**Acknowledgments.** The author thanks the anonymous reviewers for wise revisions. Thanks to A. Blanford, M. Delli Carpini, S. González-Bailón, J. Goodwin, B. Hoffman, B. Kroeger, D. Liebermann, N. Maruyama, T. Patel, V. Pickard, J. Poinsett, J. Rosen and J. Smith for their invaluable feedback. ☐

## References

1. Acar, G. et al. Fpdetective: Dusting the Web for fingerprinters. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*. ACM, 1129–1140.
2. Ackerman, M.S., Cranor, L.F. and Reagle, J. Privacy in e-commerce: examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce*. ACM, 1999, 1–8.
3. Ayenson, M., Wambach, D., Soltani, A., Good, N. and Hoofnagle, C. Flash cookies and privacy ii: Now with HTML5 and etag respawning. Available at SSRN 1898390, 2011.
4. Castellucia, C., Grumbach, S., Olejnik, L. et al. Data harvesting 2.0: From the visible to the invisible Web. In *Proceedings of the 12th Workshop on the Economics of Information Security*, 2013.
5. Duhigg, C. How companies learn your secrets. *New York Times*, (2012), 2012.
6. Dwoskin, E.D.E. Data broker removes rape-victims list after journal inquiry. *Wall Street Journal*, 2013.
7. Eckersley, P. How unique is your Web browser? *Privacy Enhancing Technologies*. Springer, 2010, 1–18.
8. Fielding, R. et al. Hypertext transfer protocol (1999), <http://1.1>.
9. Fox, S. and Duggan, M. Health online 2013. Pew Internet and American Life Project.
10. Grimes-Gruczka, T., Gratz, C. and Dialogue, C. Ethics: Survey of Consumer Attitudes about Health Web Sites. California HealthCare Foundation, 2000.
11. Himmelstein, D.U., Thorne, D., Warren, E. and Woolhandler, S. Medical bankruptcy in the United

- States, 2007: Results of a national study. *Amer. J. Med.* 122, 8 (2009), 741–746.
12. Hoofnagle, C., Urban, J. and Li, S. Privacy and modern advertising: Most us Internet users want do not track to stop collection of data about their online activities. In *Proceedings of the Amsterdam Privacy Conference*, 2012.
13. Jackson, C., Bortz, A., Boneh, D. and Mitchell, J.C. Protecting browser state from Web privacy attacks. In *Proceedings of the 15th International Conference on World Wide Web*. ACM, 2006, 737–744.
14. Jang, D., Jhala, R., Lerner, S. and Shacham, H. An empirical study of privacy-violating information flows in JavaScript Web applications. In *Proceedings of the 17th ACM conference on Computer and Communications Security*. ACM, 2010, 270–283.
15. Krebs, B. Sources: Target investigating data breach (2013); <http://krebsonsecurity.com/2013/12/sources-target-investigating-data-breach/>.
16. Krishnamurthy, B., Naryshkin, K. and Wills, C. Privacy leakage vs. protection measures: The growing disconnect. In *Proceedings of the Web 2.0 Security and Privacy Workshop*, 2011.
17. Krishnamurthy, B. and Wills, C. Privacy diffusion on the Web: A longitudinal perspective. In *Proceedings of the 18th International Conference on World Wide Web*. ACM, 2009, 541–550.
18. Krishnamurthy, B. and Wills, C.E. Generating a privacy footprint on the Internet. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*. ACM, 2006, 65–70.
19. Mayer, J.R. and Mitchell, J.C. Third-party Web tracking: Policy and technology. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*. IEEE, 413–427.
20. Miller, B., Huang, L., Joseph, A. and Tygar, J. I know why you went to the clinic: Risks and realization of https traffic analysis. arXiv preprint arXiv:1403.0297, 2014.
21. National Institutes of Health, History of Medicine Division. Greek medicine (2002); [http://www.nlm.nih.gov/hmd/greek/greek\\_oath.html](http://www.nlm.nih.gov/hmd/greek/greek_oath.html)
22. Network Advertising Initiative. NAI code of conduct, 2011.
23. Nikiforakis, N. et al. Cookieless monster: Exploring the ecosystem of Web-based device fingerprinting. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2013.
24. PhantomJS. PhantomJS is a headless Webkit scriptable [browser] with a JavaScript API, 2013; <http://phantomjs.org/>.
25. Roesner, F., Kohno, T., and Wetherall, D. Detecting and defending against third-party tracking on the Web. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*. USENIX Association, 2012, 12.
26. Staff of Chairman Rockefeller. A review of the data broker industry: Collection, use, and sale of consumer data for marketing purposes. U.S. Senate, 2013.
27. Sweeney, L. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (May 2013), 44–54.
28. Turow, J. *The Daily You: How the New Advertising Industry is Defining Your Identity and Your Worth*. Yale University Press, 2012.
29. Turow, J. and Center, A.P.P. Americans & online privacy: The system is broken. Annenberg Public Policy Center, University of Pennsylvania, 2003.
30. Turow, J., King, J., Hoofnagle, C.J., Bleakley, A. and Hennessy, M. Americans reject tailored advertising and three activities that enable it. Available at SSRN 1478214, 2009.
31. United States. Health Insurance Portability and Accountability Act of 1996. Public Law, 1996, 104–191.
32. U.S. Federal Trade Commission. Complying with the FTC’s health breach notification rule, 2010; <http://www.business.ftc.gov/documents/bus56-complying-ftcs-health-breach-notification-rule/>
33. U.S. Federal Trade Commission. Protecting consumer privacy in an era of rapid change preliminary staff report, 2010; [http://www.ftc.gov/sites/default/\\_les/documents/reports/federal-trade-commission-bureau-consumer-protection-preliminary-ftc-sta-report-protecting-consumer/101201privacyreport.pdf](http://www.ftc.gov/sites/default/_les/documents/reports/federal-trade-commission-bureau-consumer-protection-preliminary-ftc-sta-report-protecting-consumer/101201privacyreport.pdf)
34. Yen, T.-F., Xie, Y., Yu, F., Yu, R.P. and Abadi, M. Host fingerprinting and tracking on the Web: Privacy and security implications. In *Proceedings of NDSS*, 2012.

**Timothy Libert** (tlibert@asc.upenn.edu) is a doctoral student in the Annenberg School for Communication at the University of Pennsylvania, Philadelphia, PA.

Copyright held by author.  
Publication rights licensed to ACM. \$15.00.

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.