Five Years of the Right to be Forgotten

Theo Bertram Elie Bursztein Stephanie Caro Hubert Chao Rutledge Chin Feman
 Peter Fleischer Albin Gustafsson Jess Hemerly Chris Hibbert Luca Invernizzi
 Lanah Kammourieh Donnelly Jason Ketover Jay Laefer Paul Nicholas Yuan Niu
 Harjinder Obhi David Price Andrew Strait Kurt Thomas Al Verney

Google, Inc.

ABSTRACT

The “Right to be Forgotten” is a privacy ruling that enables Europeans to delist certain URLs appearing in search results related to their name. In order to illuminate the effect this ruling has on information access, we conducted a retrospective measurement study of 3.2 million URLs that were requested for delisting from Google Search over five years. Our analysis reveals the countries and anonymized parties generating the largest volume of requests (just 1,000 requesters generated 16% of requests); the news, government, social media, and directory sites most frequently targeted for delisting (17% of removals relate to a requester’s legal history including crimes and wrongdoing); and the prevalence of extraterritorial requests. Our results dramatically increase transparency around the Right to be Forgotten and reveal the complexity of weighing personal privacy against public interest when resolving multi-party privacy conflicts that occur across the Internet. The results of our investigation have since been added to Google’s transparency report.

CCS CONCEPTS

• **Security and privacy** → **Human and societal aspects of security and privacy**;

KEYWORDS

Right to be forgotten; search delisting; data privacy

ACM Reference Format:

Theo Bertram, Elie Bursztein, Stephanie Caro, Hubert Chao, Rutledge Chin Feman, Peter Fleischer, Albin Gustafsson, Jess Hemerly, Chris Hibbert, Luca Invernizzi, Lanah Kammourieh Donnelly, Jason Ketover, Jay Laefer, Paul Nicholas, Yuan Niu, Harjinder Obhi, David Price, Andrew Strait, Kurt Thomas, Al Verney. 2019. Five Years of the Right to be Forgotten. In *2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*, November 11–15, 2019, London, United Kingdom. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3319535.3354208>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CCS '19, November 11–15, 2019, London, United Kingdom

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6747-9/19/11.

<https://doi.org/10.1145/3319535.3354208>

1 INTRODUCTION

The “Right to be Forgotten” (RTBF) is a landmark European ruling governing the delisting of information from search results [1]. It establishes a right to privacy, whereby individuals can request that search engines such as Google, Bing, and Yahoo delist URLs from across the Internet that contain “inaccurate, inadequate, irrelevant, or excessive” information surfaced by queries containing the name of the requester [7]. Critically, the ruling requires that search engine operators make the determination for whether an individual’s right to privacy outweighs the public’s right to access lawful information when delisting URLs.

After the RTBF came into effect in May 2014, the broad scope of the ruling raised significant questions about how search engine operators weighed and ultimately resolved delisting requests [23]. While Google and Bing both publicly report the volume of RTBF requests [13, 21], there has been a demand for more public data on the types of information that search engines delist as well as on the entities seeking RTBF delistings, as exemplified by a letter published from 80 academics [19].

In this work, we address the need for greater transparency by shedding light on how Europeans use the RTBF in practice. Our measurements cover over five years of RTBF delisting requests to Google Search, totaling nearly 3.2 million URLs. From this dataset, we provide a detailed analysis of the countries and anonymized individuals generating the largest volume of requests; the news, government, social media, and directory sites most frequently targeted for delisting; and the prevalence of extraterritorial requests that cross regional and international boundaries. These metrics strike a balance between transparency and respecting the privacy of the individuals involved. As such, we cannot provide more detailed information on the decision process and the exact content of removals. Ultimately, our study reveals the current trade offs between personal privacy, public interest in information, and the manual review required to support the Right to be Forgotten.

We frame the key findings of our analysis as follows:

There are two dominant intents for RTBF delisting requests:

29% of requested URLs related to social media and directory services that contained *personal information*, while 21% of URLs related to news outlets and government websites that in a majority of cases covered a requester’s *legal history*. The remaining 50% of requested URLs covered a broad diversity of content on the Internet.

Variations in regional attitudes to privacy, local laws, and media norms influence the URLs requested for delisting: individuals from France and Germany frequently requested to delist social media and directory pages, while requesters from Italy and the United Kingdom were 3 times more likely to target news sites.

Requests carry a local intent: over 75% of requests to delist URLs rooted in a country code top-level domain (e.g., peoplecheck.de) came from requesters in the same country; all of the top 30 European news outlets targeted for delisting received a minimum of 82% of requested URLs from requesters in the same country.

Only 44.5% of URLs meet the criteria for delisting: delisting rates ranged from 25–53% by country, and 3–100% based on the category of personal information referenced by a URL (e.g., political platforms versus medical status).

Requests skew toward a small number of countries and individuals: France, Germany, and the United Kingdom generated 50% of URL delisting requests. Similarly, just 1,000 requesters (0.2% of individuals filing RTBF requests) requested 16% of all URLs. Many of these frequent requesters were law firms and reputation management services.

Requests predominantly come from private individuals: 84% of requested URLs came from private individuals. Minors made up 6% of requesters. Non-government public figures such as celebrities requested the delisting of 74,602 URLs and politicians and government officials another 65,933 URLs.

As a result of our research, our new findings are now integrated into Google's official transparency report [13]. We hope our work serves as a more detailed guide to understanding the RTBF and the current scale and effect of its application.

2 THE RIGHT TO BE FORGOTTEN

Before diving into our analysis, we provide a short background on the RTBF and how it compares to other privacy controls. We also discuss the process for how Google arrives at delisting verdicts, how Google enforces delistings, and recent rulings that recognize a RTBF in other countries.

2.1 Origin

In May 2014, the Court of Justice of the European Union established a RTBF [1]. It allows Europeans to request that search engines delist links present in search results containing an individual's name, if the individual's right to privacy outweighs public interest in those results. The delisted information must be "inaccurate, inadequate, irrelevant or no longer relevant, or excessive in relation to those purposes and in the light of the time that has elapsed" [1]. The ruling requires that search engine operators conduct this balancing test and arrive at a verdict.

In the wake of the ruling, Google formed an advisory council drawn from academic scholars, media producers, data protection authorities, civil society, and technologists to establish decision criteria for particularly challenging delisting requests [8]. Google also added a transparency report to reveal the volume of requests

and domains targeted for delisting [13]. Transparency around this process is challenging, as the requester's identity cannot be disclosed. This has led to the development of a set of proposed best practices from Data Protection Authorities for handling delisting requests [4]. In parallel, researchers have examined de-anonymization risks surrounding the RTBF [29].

Compared to other privacy techniques such as access control policies in social networks or anonymization services, the RTBF is unique in how it addresses multi-party privacy conflicts. Under the RTBF, access to lawfully published, truthful information about a person may be restricted via search engines (e.g., a criminal conviction, or a photo). That older information is more likely to be restricted than newer information due to delisting is perhaps an important issue for both civil society and future generations.

2.2 Decision process

Individuals located in European Union countries, as well as Iceland, Liechtenstein, Norway, and Switzerland are eligible to submit a RTBF request, and can do so via Google's online delisting form [12]. As part of this form, requesters must both verify their identity by submitting a document (a government-issued ID is not required) and provide a list of URLs they would like to delist, along with the search queries leading to these URLs and a short comment about how the URLs relate to the requester. Requesters also choose a country associated with the request, which is often their country of residence.

Google assigns every request to at least one or more reviewers for manual review. There is no automation in the decision making process. In broad terms, the reviewers consider four criteria that weigh public interest versus the requester's personal privacy:

- (1) The validity of the request, both in terms of actionability (e.g., the request specifies exact URLs for delisting) and the requester's connection to an EU/EEA country.
- (2) The identity of the requester, both to prevent spoofing or other abusive requests, and to assess whether the requester is a minor, politician, professional, or public figure, all of which have implications for balancing public interest. For example, if the requester is a public figure, there may be heightened public interest surrounding the content compared to a private individual.
- (3) The content referenced by the URL. For example, information related to a requester's business may be of public interest for potential customers. Similarly, content related to a violent crime may be of interest to the general public. Other dimensions of this consideration include the sensitive, private nature of the content and the degree to which the requester consented to the information being made public.
- (4) The source of the information, be it a government site or news site, or a blog or forum. In the case of government pages, access to a URL may reflect a decision by the government to inform society on a particular matter of public interest.

In assessing each of these criteria, reviewers may seek additional details from the requester. Ultimately, a decision is made to delist the URL from Google Search or reject the request.

2.3 Delisting visibility

Delistings occur on result pages for queries containing a requester's name on (1) Google's European country search services¹; and (2) on all country search services, including *google.com*, for queries performed from geolocations that match the requestor's country [10]. However, in 2015 the French data privacy regulator CNIL notified Google to extend the scope of delisted URLs globally, not just within Europe [9]. Google appealed this decision, and the matter is now under consideration by the Court of Justice of the European Union [2, 24].

2.4 Similar rulings in other countries

While this study focuses solely on the European RTBF, for completeness we note that other countries have adopted similar rulings. In July 2015, Russia passed a law that allows citizens to delist links from Russian search engines that "violates Russian laws or if the information is false or has become obsolete" [26]. Turkey established its own version of the RTBF in October, 2016 [22]. As our subsequent findings identify large variations in the privacy attitudes of different countries, it is unlikely our results generalize to these other RTBF laws.

3 DATASET

Our dataset includes all European RTBF requests filed with Google Search from May 30, 2014 (when Google first implemented a delisting mechanism) to May 31, 2019. We provide a high level summary of this data in Table 1. Each record consists of two parts: (1) basic information related to URLs requested for delisting; and (2) additional annotations added by reviewers in the course of arriving at a verdict, described here.

3.1 Basic request data

Every request consists of the requester's self-reported name, email address, country of residence, the date of the request, and the URLs requested for delisting, which we refer to from here on out as the *requested URLs*. Over the five year period covered by our dataset, Google processed a total of 3,231,694 requested URLs submitted by 502,648 requesters². Additionally, each request includes a verdict per URL (i.e., delisted, rejected) and the date reviewers reached that verdict.

3.2 Annotations

Starting on January 22, 2016, Google reviewers began manually annotating each requested URL with additional categorical data for improving transparency around RTBF requests.

Category of site: The first annotation consists of a general purpose category that captures common strains of delisting requests related to social media, directory services, news, and government pages as shown in Table 2. These categories are not meant to be

¹At the time of the study, European country search services were determined by ccTLD, e.g., *google.fr*, *google.no*, etc.

²We omit any requested URLs that are still pending review at the end of our collection window. Over time, this may introduce discrepancies between metrics from Google's Transparency Report and our report.

Dataset	Summary
Requested URLs	3,231,694
Unique requesters	502,648
Unique hostnames present in requested URLs	509,051
Time frame	May 30, 2014–May 31, 2019

Table 1: Summary RTBF requests targeting Google Search included in our analysis.

exhaustive. Rather, they enable tracking trends within categories and divergent interests between countries.

In total, 49.9% of requested URLs from 2016 onward fell into one of these categories. Reviewers labeled the other 50.1% as Miscellaneous. The large volume of miscellaneous URLs stems directly from the inherent diversity of content on the Internet. Examples of sites in this category include archival sites like *archive.is*, forums and blogs like *indymedia.org*, and articles or discussions on *wikipedia.org*.

Content on page: In the process of reviewing a URL to determine whether it meets the criteria for delisting, reviewers will categorize the information found on a page into a high-level description that captures the intent of the delisting as shown in Table 3. Most of these categories relate to varying degrees of potentially private information. Two categories—self-authored and name not found—relate to how information was authored or whether content is anonymized or does not reference the requester by name. In total, reviewers determined a content category for 83.7% of requested URLs. The remaining 16.3% did not fall into one of these categories and were denoted Miscellaneous.

Requesting entity: The final annotation identifies special categories of individuals. These categories capture the degree to which information may be judged to have more of or less of a public interest, or the eligibility of the requester. We provide a category breakdown in Table 4. In total, reviewers labeled 100% of requesters as belonging to one of these categories.

3.3 Ethics & Reproducibility

Our analysis broadly expands on the information previously available in Google's Transparency Report [14]. Following the transparency report's same standard, in the course of our analysis we never discuss details that might de-anonymize requesters or draw attention to specific URL content that was delisted. This emphasis on privacy creates a fundamental tension with standards for reproducible science. We cannot directly reveal a sample mapping between URLs and our annotations, nor can we reveal the exact URLs requested and the justification for delisting verdicts. We recognize these limitations upfront—but nevertheless argue it is critical to increase transparency surrounding how Europeans apply the RTBF and how it influences access to information on the Internet.

Label	Description
Directory	The URL relates to a directory or aggregator of information such as postal addresses or phone numbers for businesses or individuals. Examples include <i>118712.fr</i> which acts as a French business yellow page, or <i>192.com</i> which aggregates information about UK persons and businesses.
News	The URL belongs to a non-government media outlet or tabloid. Examples include <i>repubblica.it</i> an Italian newspaper, or <i>tv3.lt</i> a television station in Lithuania.
Social media	The URL references an account profile, photo, comment, or other content hosted on an online social network or digital forum. Examples include <i>facebook.com</i> and <i>vk.com</i> .
Government	The URL references a regional government's legal or business records, or an official media outlet. Examples include <i>boe.es</i> , a bulletin for the Spanish government; and <i>thegazette.co.uk</i> , an official news outlet for the British government.

Table 2: Manually applied disjoint labels for different categories of sites. We denote URLs outside of these categories as Miscellaneous.

Label	Description
Personal information	The requester's personal address, residence, and contact information or images and videos.
Sensitive personal information	The requester's medical status, sexual orientation, creed, ethnicity, or political affiliation.
Professional information	A requester's work address, contact information, or neutral stories about their business activities.
Professional wrongdoing	References to the requester's convictions of a crime, acquittals, or exonerations in a professional role.
Crime	References to the requester's convictions of a crime, acquittals, or exonerations.
Political	Criticism of a requester's political or government activities, or information about their platform.
Self authored	Requester authored the content.
Name not found	No reference to the requester's name found in the content of the URL, though their name may appear in the URL parameters.

Table 3: Manually applied disjoint labels for the potentially private information appearing on a page requested for delisting. We denote content not falling into one of these categories as Miscellaneous.

Label	Description
Private individual	Default label for requesters not falling into a special category.
Minor	Requester is under the age of 18.
Government official, politician	Requester is a current or former government official or politician.
Corporate entity	Requester is filing on the behalf of a business or corporation.
Deceased person	Requester is filing on the behalf of a deceased person.
Non-governmental public figure	Requester known at an international level (e.g., a famous actor or actress), or has a significant role in public life within a specific region or area (e.g., a famous academic well known in their field).

Table 4: Manually applied disjoint labels for identifying categories of requesting entities.

4 DELISTING REQUESTS

To start, we examine how the public's use of the RTBF as a privacy tool has evolved over time. We explore the top countries generating requests, the impact of high-volume requesters on overall requests and delistings, and the number of corporations and government officials filing requests.

4.1 Timeline of requests

In order to understand whether RTBF requests are increasing in frequency, we calculate the number of URLs requested every month in Figure 1. During the first two months after RTBF enforcement began in May, 2014, requesters sought to delist roughly 282,000 URLs. From January 2015 onward, RTBF requests fell to an average of approximately 47,000 URLs requested each month. Broken down by year, 29% of URLs were requested in the first year of enforcement, while 16–18% of URLs were requested each year thereafter. We note this metric includes all requests, even those that reviewers ultimately declined to delist. In aggregate, only 44.5% of requested URLs met the criteria for delisting.

As an alternate metric of usage, we examine the arrival rate of new requesters as shown in Figure 2. After the initial spike of interest, the number of new requesters since January 2015 fell to an average of approximately 6,800 each month. Broken down by year, 35% of previously unseen requesters applied during the first year of enforcement, 19% in the second year, and 14% in the last year. This indicates that while new Europeans continue to seek delistings under the RTBF, the number of previously unseen requesters has trended down year after year.

4.2 Requests and delistings by country

We find that a handful of countries heavily skew the volume of requests, as shown in Table 5. Combined, requesters from France, Germany, and the United Kingdom generated 50.1% of requested URLs. While these are some of the most populous nations covered under the RTBF, only France is one of the highest volume in terms of requests per capita as shown in the Figure 3. For every thousand requesters in France, drawn from United Nations population estimates for 2015 [28], Google received approximately 10 requested URLs. If we restrict population to estimated Internet users, drawn from the International Telecommunication Union Internet penetration estimates for 2015 [18], then every thousand Internet users in France requested 12 URLs. In comparison, requesters from Germany and the United Kingdom generated roughly half that volume, at 7 requested URLs per thousand Internet-using residents. These findings suggest that usage of the RTBF differs across countries. As we explore later in Section 5, multiple factors may help explain this including regional news and government practices on reporting private information as well as social media penetration.

4.3 High-volume requesters

As individuals have varying privacy expectations and digital footprints, we examine the volume of requested URLs per requester. We find a small number of requesters made heavy use of the RTBF to delist URLs as shown in Figure 4. The top thousand requesters (0.2% of all requesters) generated 16.3% of requests and 23.4% of

Figure 1: URLs requested for delisting under the RTBF, including rejected requests. After an initial burst of requests as enforcement went into effect, there were an average of approximately 47,000 requested URLs per month since January 2015.

Figure 2: Timeline of previously unseen requesters based on the date of their first request. Year after year, the number of new requesters declined.

Country	Code	Requested URLs	Breakdown
France	FR	648,635	20.1%
Germany	DE	537,685	16.6%
United Kingdom	GB	433,520	13.4%
Italy	IT	280,923	8.7%
Spain	ES	257,805	8.0%
Netherlands	NL	177,587	5.5%
Poland	PL	106,634	3.3%
Sweden	SE	96,900	3.0%
Belgium	BE	85,353	2.6%
Switzerland	CH	62,817	1.9%

Table 5: Top 10 countries by volume of URLs requested under the RTBF.

delistings. These mostly included law firms and reputation management agencies, as well as some requesters with a sizable online presence. Of the entities in this group, 17.6% resided in Germany, 14.9% in France, and 14.7% in the United Kingdom. The most prolific requester sought to delist 8,222 URLs. In the remaining long tail, 35.0% of requesters sought to delist only a single URL, and 74.7%

Figure 3: URLs requested for delisting per capita, sorted by the Internet-using population of each country. For every thousand residents connected to the Internet, there were between 3–20 URLs requested for delisting per country.

Figure 4: CDF of all requested and delisted URLs, ranked by the highest volume requester. The top thousand requesters generated 16.3% of URL requests and 23.4% of delistings.

five or fewer URLs. These results illustrate that while a majority of Europeans rely on the RTBF to delist a handful of URLs, some law firms and reputation management services use the RTBF to delist hundreds of URLs to alter the content available about their clients in search results.

4.4 Requesting entities

As the RTBF explicitly outlines public interest as one of the balancing criteria when judging delistings, it is also important to understand the categories of entities requesting to delist URLs. Starting in January 2016, Google began categorizing requesters into coarse groups (discussed in Section 3). We provide a breakdown of all requested URLs since January 2016 based on these categories in Table 6. A majority of requested URLs—84.3%—were from private individuals. Minors constituted 5.6% of all requested URLs, a special group with a delisting rate nearly twice as high as private individuals. Government officials and politicians generated 3.5% of requested URLs and had a lower delisting rate than private individuals. The low delisting rates for government officials highlights one aspect of the public interest balance that Google strikes when applying the RTBF. We note that corporate entities never have content delisted under the RTBF.

4.5 Processing time

With tens of thousands of RTBF requests filed each month, we examine the time it takes Google to review each request and reach a verdict, where a request may include multiple URLs. In the first two months after Google began delisting URLs under the RTBF, it took a median of 85 days to reach a decision from the time a requester first filed a request. This delay resulted from reviewers establishing cases for what constituted inaccurate, inadequate, irrelevant, or excessive information. After January 2017, a URL took a median of 6 days to process.

Reviewing requests is critical for catching fraudulent and erroneous submissions. In one case, an individual, who was convicted for two separate instances of domestic violence within the previous five years, sent Google a delisting request focusing on the fact that their first conviction was “spent” under local law. The requester did not disclose that their second conviction was not similarly spent, and falsely attributed all the pages sought for delisting to the first conviction. Reviewers discovered this as part of the review process and the request was ultimately rejected. In another case, reviewers first delisted 150 URLs submitted by a businessman who was convicted for benefit fraud, after they provided documentation confirming their acquittal. When the same person later requested the delisting of URLs related to a conviction for manufacturing court documents about their acquittal, reviewers evaluated the acquittal documentation sent to Google, found it to be a forgery, and reinstated all previously delisted URLs.

More typically though, requests were rejected due to errors on the part of the requester. Common errors included requests for URLs not in Google’s search index and invalid or incomplete URLs (e.g., requesting facebook.com rather than a specific profile page). For requests filed after 2016 when reviewers started annotating requests, 25.8% were rejected due to errors on the part of the requester.

5 CONTENT TARGETED

In total, requesters have sought to delist URLs related to 509,051 different hostnames on the Internet. These requests capture a spectrum of potentially personal information: some requesters seek to control their digital footprint exposed through social networks and directory services, while others delist URLs related to news

Requesting entity	Requested URLs	Breakdown	Delisting rate
Private individual	1,582,249	84.3%	46.5%
Minor	104,976	5.6%	77.2%
Non-governmental public figure	74,602	4.0%	37.7%
Government official or politician	65,933	3.5%	13.2%
Corporate entity	40,575	2.2%	0.0%
Deceased person	8,008	0.4%	24.3%

Table 6: Breakdown of all requested URLs after January 2016 by the categories of requesting entities. Private individuals make up the bulk of requests.

sources and government reports. We explore these divergent applications in terms of the hostnames targeted, the categories of content present on requested pages, and ultimately whether there are better mechanisms for the public to resort to for controlling personal information on popularly requested sites.

5.1 Frequently targeted sites

We present a breakdown of the top 20 hostnames requested for delisting from May 2014–May 2019 in Table 7.³ We rely on hostnames rather than second-level domains to differentiate multiple services hosted on the same domain (e.g., wordpress.com). Roughly half of the sites listed are social networks and forums, the most popular including Facebook, YouTube, Google+, Google Groups, and Twitter. In practice requests to these sites may be greater as many operate ccTLDs that are not reflected in the top 20 (e.g., facebook.fr, facebook.de). The other half includes directory sites that aggregate contact details and personal content from other sites, like 118712 and Profile Engine. Combined, the top 20 sites represent 10.6% of requested URLs. Delisting rates varied drastically based on site, ranging 13.3%–91.4%.

5.2 Categories of sites

In order to gain a broad perspective of the categories of sites targeted by requests, we rely on four years of labels from January 2016 onward assigned to each requested URL (discussed in Section 3). These labels cover four mutually exclusive categories: social media content, news sources, directory and information aggregators, and government pages. We provide a breakdown of requested URLs by category in Table 8. Directories that display names, addresses, and other personal information made up 16.0% of requested URLs. Social media made up 13.1% of requested URLs. Delisting rates for these two categories in aggregate surpassed 52%. News media, entirely absent from the top 20 requested hostnames, made up 18.5% of requested URLs; a reflection of the volume of independent media outlets on the Internet. Far less frequent, government pages made up only 2.4% of requested URLs. Given the public interest nature of news and government records, delisting rates were lower

³We exclude www. and m. prefixes from hostnames. We also omit 62,045 requested URLs during this period from this and all subsequent content analysis as they were not well-formed URLs.

Hostname	Description	Requested URLs	Delisting rate
facebook.com	Social	59,852	56.6%
plus.google.com	Social	35,892	40.6%
youtube.com	Social	35,798	42.2%
twitter.com	Social	33,015	49.1%
annuaire.118712.fr	Directory	22,379	80.5%
groups.google.com	Social	21,074	50.0%
flashback.org	Forum	15,333	13.3%
scontent.cdninstagram.com	Social	13,982	90.5%
profileengine.com	Directory	13,974	91.4%
badoo.com	Social	11,964	62.7%
societe.com	Directory	9,861	22.1%
myspace.com	Social	9,369	49.5%
pbs.twimg.com	Social	9,075	63.0%
copainsdavant.linternaute.com	Directory	8,148	48.8%
verif.com	Directory	7,972	22.7%
192.com	Directory	7,881	77.7%
pinterest.com	Social	7,340	61.3%
linkedin.com	Social	7,275	37.8%
wherevent.com	Directory	7,212	90.7%
infobel.com	Directory	6,853	74.3%

Table 7: Top hostnames requested for delisting during May 2014–May 2019 along with a description for content typically hosted on the domain. Social media (and its CDN equivalents), forums, and directory aggregation services dominate the list.

Category of site	Requested URLs	Breakdown	Delisting rate
News	341,821	18.5%	34.5%
Directory	295,349	16.0%	52.8%
Social Media	242,503	13.1%	52.9%
Government	44,380	2.4%	19.1%
Miscellaneous	925,944	50.1%	47.2%

Table 8: Breakdown of requested URLs during the four years from January 2016–May 2019 based on the category of site requested.

than average—34.5% and 19.1% respectively. The remaining 50.1% of URLs did not fall within a well-defined category.

To provide a perspective of the long tail of remaining hostnames, we measure the cumulative fraction of requests covered by successively adding the next most popular hostname in Figure 5. We find 19.7% of requested URLs target only 100 hostnames and 39.4% of requests target the top 1,000 hostnames. In contrast, 59.4% of hostnames received only a single request and 89.1% of hostnames at most 5 requests. Our results indicate that the privacy concerns of Europeans concentrate on a small fraction of the hundreds of millions of hostnames on the Internet.

We provide a monthly breakdown of requested URLs by the category of site in Figure 6. We find requests targeting news-related content increased from 14.7% on January 2016 to 18.2%, on May 2019 while social media requests generally declined from 15.5% to 11.7%. These trends may reflect requesters adapting their privacy

Figure 5: CDF of all requested and delisted URLs, ranked by the most frequently targeted hostnames. The top thousand hostnames received 39.4% of requests and 43.7% of delistings.

concerns over time. For social media, entities may be finding other mechanisms to proactively reign in their social media footprint (or are less concerned). In contrast, news services remain outside a requester’s control. The steep decline in directory-related sites after May 2018 is likely the result of the General Data Protection Regulation (GDPR) coming into effect. After this period, 80% of previously requested directory sites had no subsequent delisting requests. We checked the liveness of the top 500 directory domains by request volume prior to GDPR that received no requests after GDPR and found only 54.6% remain online⁴.

Cut another way, we find broad variations in the categories of sites requested by individuals of different countries as shown in Table 9. Focusing only on the top five countries by volume of requests, we find requesters from Italy and the United Kingdom were far more likely to target news media in their requests (33.0% and 26.2% of requested URLs respectively). This correlates with diverging journalistic practices: news sources in Italy and the United Kingdom are more prone to reveal the identity of individuals in relation to articles covering crimes. In contrast, news sources in Germany and France tend to anonymize the parties in articles covering crimes. Requesters in France and Germany were most concerned with information exposed on social media and via directory services compared to other countries. Finally, in Spain, 10.4% of requested URLs targeted government records. This may stem from Spanish law which requires the government to publish ‘edictos’ and ‘indultos’. The former are public notifications to inform missing individuals about a government decision that directly affects them; the latter are government decisions to absolve an individual from a criminal sentence or to commute to a lesser one. These variations expose how RTBF usage varies by country, illustrating the challenge of one-size-fits-all privacy policies.

5.2.1 News Media. Examining each category in more depth, we present a breakdown of the top 10 news outlets and tabloids targeted by delisting requests during January 2016–May 2019 in Table 10. We also include a count of all historical requests to the same domains

⁴Sites that were no longer online returned a 40X related error or timed out.

Figure 6: Monthly breakdown of the categories of sites requested for delisting. Requests to delist social media URLs have trended downward, while news media has increased.

Country	Directory	Social	News	Gov’t	Misc
France	26.1%	16.3%	12.1%	1.8%	43.7%
Germany	17.5%	15.2%	10.5%	0.9%	55.9%
Italy	8.0%	9.9%	33.0%	2.1%	47.0%
Spain	16.2%	11.8%	20.8%	10.4%	40.7%
United Kingdom	10.5%	10.6%	26.2%	2.2%	50.4%

Table 9: Categories of sites requested by the top five requesting countries. Requesters from Italy and the United Kingdom are the most likely to target news media, compared to France and Germany where the major concern is personal information exposed via social networks and directory aggregators.

for the entire period of our dataset. The institutions represented provide news coverage to some of the most populous nations in the European Union, including the United Kingdom, Italy, and France. Delisting rates for the January 2016–May 2019 period ranged from 22.3–40.5%.

In order to illuminate aspects of the balancing tests that Google conducts for articles on news media outlets, we provide three examples of requests and their outcomes. In one case, a requester who held a significant position at a major company sought to delist an article about receiving a long prison sentence for attempted fraud. Google rejected this request due to the seriousness of the crime and the professional relevance of the content. In another request, an individual sought to delist an interview they conducted after surviving a terrorist attack. Despite the article’s self-authored nature given the requester was interviewed, Google delisted the URL as the requester was a minor and because of the sensitive nature of the content. Lastly, a requester sought to delist a news article about their acquittal for domestic violence on the grounds that no medical report was presented to the judge confirming the victim’s

Hostname	Requested URLs	Delisting rate	Requested URLs (all time)
dailymail.co.uk	3,831	35.1%	5,936
ouest-france.fr	1,964	22.3%	2,452
telegraph.co.uk	1,912	40.5%	3,140
entreprises.lefigaro.fr	1,778	23.0%	1,805
ricerca.gelocal.it	1,772	27.8%	3,678
leparisien.fr	1,666	34.9%	2,850
ricerca.repubblica.it	1,593	23.5%	2,896
bbc.co.uk	1,524	27.7%	2,358
mirror.co.uk	1,512	35.5%	2,202
247.libero.it	1,427	37.2%	2,470

Table 10: Top 10 news media sites requested from January 2016–May 2019.

Hostname	Requested URLs	Delisting rate	Requested URLs (all time)
facebook.com	29,541	49.4%	59,852
twitter.com	21,230	51.6%	33,015
youtube.com	19,463	40.1%	35,798
plus.google.com	15,467	31.9%	35,892
scontent.cdninstagram.com	12,541	91.1%	13,982
pbs.twimg.com	7,615	62.4%	9,075
groups.google.com	5,950	34.6%	21,074
i.ytimg.com	5,028	66.9%	6,081
myspace.com	4,473	49.5%	9,369
linkedin.com	4,146	32.0%	7,275

Table 11: Top 10 social media sites requested from January 2016–May 2019.

injuries. Given that the requester was acquitted, Google delisted the article.

We note that the RTBF is just one technical mechanism currently available to delist news stories. Websites can also add a `Disallow` directive in their `robots.txt` to instruct crawlers to ignore any listed URLs. For one popular Italian news site, roughly 20% of URLs requested under the RTBF for delisting appeared in their `disallow` directive, including URLs that Google ultimately rejected to delist on the grounds of public interest. We cannot definitively state what legal mechanism or internal decision lead the news site to add a `disallow` directive for these URLs, but the end result is that the news articles are now removed from *all* search results regardless of the search query or origin of the search request. Whether such site-level controls should play a role in enforcing the RTBF remains an area of discussion [11].

5.2.2 Social Media & Communities. We provide a breakdown of the top ten most popular social media sites targeted for delisting in the last four years in Table 11. Facebook, Twitter, YouTube, Google+, and Instagram were the most frequent target of delisting requests. Delisting rates for the top 10 sites from January 2016–May 2019 ranged from 31.9–91.1%.

All of these social networks provide privacy controls to restrict access to self-authored content, as well as mechanisms to take down abusive content (e.g., revenge porn) that violates the site’s terms

Hostname	Requested URLs	Delisting rate	Requested URLs (all time)
annuaire.118712.fr	16,608	78.7%	22,379
societe.com	6,829	20.7%	9,861
verif.com	4,621	21.6%	7,972
infobel.com	3,671	69.3%	6,853
profileengine.com	3,620	88.1%	13,974
copainsdavant.linternaute.com	3,527	46.3%	8,148
192.com	3,328	68.1%	7,881
gepatroj.com	3,101	59.2%	3,537
manageo.fr	2,492	17.5%	4,855
wherevent.com	2,394	88.5%	7,212

Table 12: Top 10 directory sites requested from January 2016–May 2019.

of service. These tools may be better suited to individuals seeking solely to control their social media footprint, with the exception of information leaked due to asymmetric privacy controls (e.g., photo tagging, mentions). Furthermore, these privacy controls also extend to search queries performed on the respective sites for an individual, whereas the RTBF only covers search engine providers.

Whenever such privacy controls are relevant, Google steers requesters towards their use. For example, in one request, an individual sought to delist their Facebook profile. Google noted that Facebook has procedures to limit the visibility of the content in question for all search engines, not just Google, and recommended that the requester utilize these controls. In another case, an individual requested to delist a URL to a page that had taken a self-published image and reposted it. With no social media privacy controls to intervene, Google delisted this URL.

5.2.3 Directories. Directory services index and aggregate information available publicly through social media channels and business listings (e.g., Yellow Pages) such as names, addresses, phone numbers, and more. We provide a breakdown of the most popularly requested sites in Table 12. The most popular site, 118712.fr, is a discovery service in France for looking up professionals and individuals. Delisting rates for all these sites during the January 2016–May 2019 period ranged from 17.5–88.5%. Delisting rates were higher for directories aggregation personal information compared to professional information.

Examples of requests include a requester that sought to delist several URLs leading to a directory page displaying their address and phone number. Google delisted all of the URLs. In another case, a requester sought to delist a URL that listed the requester’s directorships in various companies. Google denied the request as the URL contained information of professional relevance which appeared neither inaccurate nor outdated.

We note that most directory sites provide their own search functionality as well. While successful RTBF requests delist directory pages for individuals on Google Search, the public can still perform a direct search via any of the popular directory services if no additional RTBF action is taken on those sites directly. Discrepancies between search indexes can lead to possible privacy risks, such as identifying requesters [29].

Hostname	Requested URLs	Delisting rate	Requested URLs (all time)
beta.companieshouse.gov.uk	2,871	6.5%	2,965
boe.es	2,553	24.6%	4,956
infogreffe.fr	1,697	11.2%	2,677
bocm.es	944	58.4%	1,496
thegazette.co.uk	535	2.8%	850
juntadeandalucia.es	525	38.0%	798
boc.cantabria.es	385	58.9%	641
sede.asturias.es	382	48.1%	539
bodacc.fr	348	4.5%	538
legifrance.gouv.fr	343	10.7%	537

Table 13: Top 10 government sites requested from January 2016–May 2019.

Hostname	Requested URLs	Delisting rate	Requested URLs (all time)
flashback.org	7,830	11.4%	15,333
goo.gl	6,236	0.6%	6,696
wikileaks.org	3,906	31.0%	4,082
archive.is	3,441	30.0%	3,855
pressreader.com	3,107	45.2%	4,631
selecto-tiernahrung.de	3,000	0.0%	3,000
camgirl.gallery	2,874	13.8%	3,149
encrypted-tbn0.gstatic.com	2,766	2.4%	3,380
linksunten.indymedia.org	2,645	60.6%	5,661
issuu.com	2,280	39.7%	4,295

Table 14: Top 10 miscellaneous sites requested from January 2016–May 2019. We note that *googleusercontent.com* and *gstatic.com* are CDNs for multiple types of content and that *goo.gl* is a URL shortener.

5.2.4 Government Records. We provide a breakdown of the top 10 most popular government and government-affiliated websites requested for delisting over the last four years in Table 13. Popular hosts include *companieshouse.gov.uk* which serves as a government-operated directory of businesses in the UK and *boe.es* which posts government bulletins. Among the most popular sites, delisting rates for January 2016–May 2019 ranged from 2.8–58.9%.

As mentioned at the start of Section 5.2, many of these requests relate to ‘edictos’ and ‘indultos’ posted by the Spanish government. In one such case, a requester sought to delist 3 URLs for a Spanish government page containing a notice summoning the requester to city hall for a matter related to tax debt. Google delisted all of the URLs. In another request, an individual sought to delist a government page from 2015 that described how the requester had received a prison sentence for managing several companies whilst being subject to bankruptcy order and restrictions. Google denied the delisting request as the information was published by a government entity and there was significant professional relevance regarding the content.

5.2.5 Miscellaneous. We provide a breakdown of the top hostnames that fell into the miscellaneous category in Table 14. Delisting rates for all of the most popular miscellaneous pages ranged from 0.0–60.6% for the period of January 2016–May 2019. We note URLs shortened via *goo.gl* are not in Google’s search index, leading to the low delisting rate shown. Instead, Google asks requesters to provide the full URL. Similarly, the low delisting rate for adult websites stems from a bias where a small number of requesters generated the vast majority of requests, where the content was still relevant to their ongoing career as an adult performer.

Examples of requests for this category include a requester that sought to delist multiple URLs to a website displaying nude images from the individual’s previous profession at an adult cam website. Google delisted all of the URLs given the sensitive nature of the content and the lack of professional relevance due to the individual’s change to a career outside the adult industry. In another case, a prominent former politician requested to delist a URL leading to a Wikipedia page that detailed a conviction the requester claimed was spent. Google denied the request due to the public role of the requester and the significant public interest regarding the conviction referenced at the URL.

5.3 Information present on page

As a final measure of how Europeans utilize the RTBF in practice, we examine the relation between each requester and the information present at URLs requested between January 2016–May 2019. For the purposes of this analysis, we exclude all misfiled URLs (discussed in Section 4.5). We provide a general summary of the content referenced by URLs in Table 15 using the categories we previously outlined (see Section 3). The most commonly requested content related to professional information, which rarely met the criteria for delisting (20.7%). Many of these requests pertain to information which is directly relevant or connected to the requester’s current profession and is therefore in the public interest to have indexed in Google Search. Personal information (e.g., a home address, contact information) made up only 7.7% of requests, while more sensitive personal information (e.g., medical information, sexual orientation) made up only 2.5% of requests. The delisting rate for both of these categories exceeds 92%, indicating that any potential public interest is usually outweighed by the sensitive nature of the content⁵. Likewise, instances of when a URL appeared in the search results for a requester’s name, but nevertheless had no reference to the requester in the content, had a delisting rate of 100% due to no competing public interest. In contrast, only 3.4% of requests targeting a requester’s political platform met the criteria for delisting.

We find that where potentially private information appears is heavily dependent on the category of site involved in a RTBF delisting request, as shown in Table 16. In the case of news, requests most frequently targeted articles covering crimes or professional wrongdoing. Conversely, individuals seeking to delist social media content most often targeted self-authored posts. These results again highlight two classes of individuals using the RTBF: those

⁵We note that sensitive personal information appearing to have a lower delisting rate than personal information is the result of a bias where that information appears: personal information most commonly appears on directory sites with limited public interest, whereas sensitive personal information appears more often on news articles where there are public interests.

Content on page	Requested URLs	Breakdown	Delisting rate
Professional information	331,693	23.9%	20.7%
Self authored	127,130	9.2%	34.4%
Crime	119,318	8.6%	48.2%
Professional wrongdoing	115,438	8.3%	19.5%
Personal information	106,996	7.7%	96.8%
Political	48,784	3.5%	3.4%
Sensitive personal information	34,530	2.5%	92.6%
Name not found	278,087	20.0%	100.0%
Miscellaneous	226,681	16.3%	70.4%

Table 15: Potentially private information appearing on URLs requested for delisting from January 2016–May 2019. Most delistings relate to professional information and self-authored content (e.g., a social media post).

Content on page	Directory	Social	News	Gov't	Misc
Crime	1.8%	2.5%	22.3%	7.6%	6.9%
Personal information	22.4%	5.3%	3.2%	3.7%	5.1%
Political	0.5%	1.2%	7.3%	4.9%	3.6%
Professional information	35.2%	8.2%	17.6%	49.9%	25.3%
Professional wrongdoing	1.3%	2.0%	22.1%	4.6%	6.9%
Self authored	3.0%	32.5%	4.7%	1.4%	7.3%
Sensitive personal information	0.7%	1.0%	1.9%	0.8%	3.9%
Name not found	18.4%	29.4%	10.3%	7.8%	22.9%
Miscellaneous	16.8%	17.9%	10.6%	19.3%	18.0%

Table 16: Breakdown of information found per category of site, summed vertically. The majority of social media sites requested for delisting were self-authored posts, whereas news articles predominantly dealt with criminal or professional wrongdoing.

seeking to delist personal information appearing on social media and directory sites, and those targeting news and government sites that report crime-related or professionally relevant information respectively.

6 EXTRATERRITORIAL REQUESTS

As a final measurement, we examine how information requested for delisting under the RTBF relates to the audience of that information. In particular, we explore whether requests span beyond the national boundaries or even the continental boundaries of the requester. However, mapping geographic boundaries to Internet sites or audiences in a generalized way is extremely difficult. We consider two proxy metrics: requests to news outlets and other pages headquartered outside the requester’s country, and requests to delist content from sites whose country code top-level domain (ccTLD) differs from the requester’s country.

6.1 Regional & International News

For the top 30 news hostnames that received delisting requests from January 2016–May 2019, we calculate the fraction of requests that originated from entities residing in the same country as the media outlet’s headquarters. As shown in Table 17, we find all but 2 of the top 30 news hostnames received over 87% of requests from local individuals. Of the two exceptions, The Irish Independent, headquartered in Ireland, received 14.7% of URL requests from UK individuals (possibly in Northern Ireland) while the same was true for 16.3% URLs requested for delisting from The Irish Times. This concentration indicates that most requesters have local intent for their delistings.

While requests from Europeans to non-European media outlets occurred, as shown in Table 18, they happened an order of magnitude less frequently compared to European media outlets (previously shown in Table 10). The most popular targets included Bloomberg, the Financial Times, and the New York Post. This mirrors our previous finding that the majority of requests to news outlets originate from local requests.

6.2 Directories, Social, and Government

We repeat the same local analysis, this time for the top 10 hostnames categorized as directory, social media, and government pages from January 2016–May 2019. We present our results in Table 19. For directory services and government pages, there is a clear banding of local requests. More than 91.6% of requests to 7 of the top 10 directory services originated from local requesters. The same was true of 81.7% of requests to all of the top 10 government pages. Only social media and three directory pages with a global user base received delisting requests from a variety of European countries. These results illustrate again the local nature of a majority of requests.

6.3 Country code domains

As an alternate measure of extraterritorial requests, we examine the relationship between the locale of a requester (e.g., France) and the ccTLDs of the URLs requested for delisting (e.g., *.fr* in the case of *number.fr*). We compare this to requests for non-EU ccTLDs and non-ccTLDs altogether such as *.com*. We note these ccTLDs relate only to the publisher’s site and not to any country specific search portal.

We present our results in Table 20 for the top five countries by volume of requests. We find 32.6–51.1% of requests targeted a host registered to a ccTLD, of which 75.0–86.6% were within the requesters locale. As with our analysis of news media, our results highlight the local nature of a majority of RTBF requests (excluding URLs with no clear locale).

7 RELATED WORK

Previous studies on the RTBF have largely focused on examining the competing interests between the right to privacy and freedom of expression, as well as the technical challenges of enforcement [3, 5, 15, 25]. Apart from these legal and policy analyses, information on how the RTBF has been used in practice is limited to the transparency reports from Google and Bing [13, 21].

Media outlet	BE	DE	ES	FR	GB	IE	IT	Other
hln.be	88.5%	0%	0%	1.0%	1.6%	0.1%	0.1%	8.7%
nieuwsblad.be	93.4%	0.1%	0.1%	0.3%	0.4%	0.1%	0.3%	5.4%
bild.de	0.1%	96.4%	0%	0.8%	1.9%	0%	0%	0.8%
elmundo.es	0%	0.6%	96.2%	1.1%	0.7%	0%	0.6%	0.8%
elpais.com	0.4%	0.1%	96.8%	0.9%	0.3%	0%	0.7%	0.8%
expansion.com	0.6%	0%	98.0%	0.2%	0.2%	0%	0.1%	0.7%
entreprises.lefigaro.fr	0.2%	0.1%	0.1%	98.8%	0.3%	0%	0.1%	0.6%
estrepubicain.fr	0.1%	0.3%	0%	99.6%	0%	0%	0%	0%
ladepeche.fr	0.9%	0%	0.1%	98.2%	0%	0%	0.1%	0.6%
lavoixdunord.fr	1.0%	0%	0.1%	98.6%	0%	0%	0%	0.3%
leparisien.fr	0.9%	0.1%	0%	97.2%	0.5%	0%	0.5%	0.8%
letelegramme.fr	0.1%	0.2%	0%	99.5%	0.2%	0%	0%	0.1%
ouest-france.fr	0.2%	0.1%	0%	99.5%	0%	0%	0.1%	0.2%
bbc.co.uk	0.1%	0.3%	0.2%	0.3%	96.7%	0.5%	0.3%	1.7%
dailymail.co.uk	0.4%	1.4%	0.4%	1.6%	91.5%	0.8%	0.9%	2.9%
independent.co.uk	0.7%	1.2%	0.4%	1.5%	90.7%	0.9%	1.5%	3.1%
mirror.co.uk	0.5%	0.6%	0%	0.9%	93.5%	0.9%	0.7%	3.0%
news.bbc.co.uk	0.1%	0.6%	0.8%	1.9%	93.1%	1.0%	0.7%	1.7%
standard.co.uk	0.3%	0.9%	0%	0.5%	94.9%	0.1%	1.1%	2.2%
telegraph.co.uk	0.3%	1.0%	0.4%	2.7%	92.1%	0.5%	0.8%	2.2%
theguardian.com	0.3%	0.9%	0.4%	2.8%	88.8%	1.5%	1.0%	4.3%
thesun.co.uk	0.4%	0.6%	0.4%	1.7%	91.8%	0.7%	1.9%	2.4%
thetimes.co.uk	0.3%	1.0%	0.2%	0.9%	87.0%	6.4%	1.5%	2.6%
independent.ie	0.2%	0.4%	0.2%	0.5%	14.7%	82.3%	0.1%	1.4%
irishtimes.com	0.5%	0.2%	0.6%	1.2%	16.3%	77.5%	0.7%	2.3%
247.libero.it	0%	0.1%	0.1%	0.1%	0.6%	0%	98.7%	0.4%
ilgazzettino.it	0%	0.4%	0%	0.5%	0.5%	0%	98.1%	0.4%
lastampa.it	0%	0.3%	0.6%	0.7%	1.0%	0%	95.9%	1.5%
ricerca.gelocal.it	0%	0.1%	0.1%	0.1%	0.1%	0%	99.4%	0.3%
ricerca.repubblica.it	0.1%	0%	0.1%	0.1%	0%	0%	98.7%	1.0%

Table 17: Top 30 hostnames categorized as news targeted for delisting, broken down by the country of the requester. We find requesters were overwhelmingly located in the same country as the outlet. We note some hostnames belong to the same outlet.

Hostname	Requested URLs	Delisting rate	Requested URLs (all time)
bloomberg.com	200	30.8%	407
ft.com	185	24.9%	331
nypost.com	160	42.2%	289
nydailynews.com	149	35.4%	251
reuters.com	126	37.6%	189
nytimes.com	125	33.6%	292
vice.com	100	22.8%	176
timesofindia.indiatimes.com	96	24.1%	172
huffingtonpost.com	90	40.0%	214
businessinsider.com	85	36.2%	178

Table 18: Top 10 non-EU news hostnames targeted for delisting from January 2016–May 2019. These requests occurred an order of magnitude less frequently than EU news outlets.

In the closest study to our own, Xue et al. analyzed 283 URLs publicly disclosed by media outlets as being delisted through the RTBF [29]. They found 62 of the articles related to miscellaneous

matters, while the remaining 221 articles referenced events surrounding sexual assault, murder, terrorism, and other criminal activity. Our own work expands on the analysis of the categories of content delisted under the RTBF. Our dataset is also unbiased to any particular segment of delistings.

A subset of RTBF requests relate to multi-party privacy conflicts. These arise due to diverging views of how content should be shared and how broadly [6, 16, 17, 20, 27]. The RTBF expands the surface area for such conflicts to occur, even beyond social networks, due to the inclusion of any websites appearing in search results. Equally challenging, the RTBF leaves the resolution of these privacy conflicts to both search engines and Data Protection Authorities rather than the parties in conflict.

8 CONCLUSION

In this paper, we presented an in-depth analysis of how the RTBF affects access to information on Google Search. We identified two dominant categories of RTBF requests: delisting personal information found on social media and directory sites; and delisting legal history and professional information reported by news outlets and

Hostname	BE	DE	ES	FR	GB	IE	IT	Other
infobel.com	10.0%	16.1%	52.9%	4.3%	0.4%	0.1%	3.6%	12.6%
annuaire.118712.fr	0.1%	0.0%	0.1%	99.3%	0.0%	0%	0.0%	0.4%
copainsdavant.linternaute.com	1.2%	0.5%	0.1%	97.0%	0.3%	0%	0.1%	0.9%
manageo.fr	0.0%	0%	0%	99.2%	0.2%	0%	0%	0.5%
gepatroj.com	0.2%	0.1%	0.0%	98.9%	0%	0%	0%	0.8%
societe.com	0.3%	0.1%	0.1%	98.5%	0.3%	0%	0.1%	0.6%
verif.com	0.3%	0.0%	0.2%	98.5%	0.2%	0%	0.1%	0.7%
192.com	0.4%	0.9%	0.5%	2.3%	91.6%	0.5%	0.3%	3.5%
wherevent.com	8.6%	29.3%	2.9%	24.7%	3.2%	0.1%	2.6%	28.4%
profileengine.com	3.6%	11.7%	6.6%	30.4%	13.6%	0.7%	4.0%	29.4%
facebook.com	3.6%	18.8%	7.1%	21.1%	10.9%	1.1%	6.3%	31.2%
groups.google.com	1.3%	37.7%	3.4%	13.9%	7.1%	0.1%	14.0%	22.4%
i.ytimg.com	2.7%	14.3%	2.4%	19.4%	7.3%	0.3%	5.0%	48.7%
linkedin.com	3.5%	9.8%	9.3%	31.3%	12.9%	0.9%	5.0%	27.3%
myspace.com	3.1%	21.1%	4.7%	33.3%	14.4%	0.4%	4.5%	18.6%
pbs.twimg.com	2.5%	16.1%	7.1%	32.6%	11.9%	0.6%	5.7%	23.5%
plus.google.com	3.8%	19.5%	8.4%	28.6%	7.9%	0.5%	4.8%	26.4%
scontent.cdninstagram.com	3.3%	27.7%	5.2%	17.6%	2.9%	0.1%	17.1%	26.1%
twitter.com	2.9%	11.6%	6.6%	31.4%	17.5%	1.0%	4.5%	24.5%
youtube.com	3.1%	13.8%	5.7%	23.2%	12.9%	0.8%	8.0%	32.4%
boc.cantabria.es	0%	0%	100.0%	0%	0%	0%	0%	0%
bocm.es	0%	0.2%	97.9%	1.2%	0.1%	0%	0%	0.6%
boe.es	0%	0.1%	99.2%	0.2%	0.1%	0%	0.1%	0.3%
juntadeandalucia.es	0.2%	0%	99.4%	0%	0%	0%	0%	0.4%
sede.asturias.es	0%	0%	99.7%	0%	0.3%	0%	0%	0%
bodacc.fr	0%	0.3%	0%	99.4%	0%	0%	0%	0.3%
legifrance.gouv.fr	0.3%	0%	0.3%	98.3%	0.3%	0%	0%	0.9%
infogrefe.fr	0.1%	0%	0%	99.3%	0.3%	0%	0.1%	0.3%
beta.companieshouse.gov.uk	0.4%	3.1%	1.7%	3.0%	81.7%	0.8%	1.5%	7.8%
thegazette.co.uk	0.2%	3.6%	0.2%	0.7%	93.1%	0.2%	0.2%	1.9%

Table 19: Breakdown of requests for the top directory, social media, and government pages. As with news media, we find requesters are overwhelmingly located in the same country as the site operator, with the exception of social media.

	GB	FR	DE	ES	IT
Non-ccTLD	64.3%	67.4%	50.9%	64.1%	48.9%
ccTLD	35.7%	32.6%	49.1%	35.9%	51.1%
Local ccTLD	75.0%	76.8%	77.9%	82.0%	86.6%
Other EU ccTLD	11.3%	10.2%	12.5%	6.0%	6.1%
Non-EU ccTLD	13.7%	13.1%	9.6%	11.9%	7.4%

Table 20: Breakdown of extraterritorial requests based on the ccTLD of requested URLs. We note the ccTLD refers only to a publisher's site and is unrelated to country specific search portals.

government pages. Our results showed the intent behind these requests was nuanced and stemmed in part from variations in regional privacy concerns, local media norms, and government practices. While a majority of URLs were requested by private individuals—84%—we found that politicians and government officials requested to delist 65,933 URLs and non-governmental public figures another 74,602 URLs. Overall, the RTBF can lead to a reshaping of search results for certain individuals, where just 1,000 entities (0.2% of roughly 502,000 requesters) requested to delist over 526,000 URLs.

Our results highlight the challenges of implementing privacy regulations in practice, both in terms of the thousands of hours of human review required and the fundamentally challenging process of weighing individual privacy against access to information. In order to ensure better transparency to enable discussions on this topic, all our results are now reflected on Google's transparency report [13].

ACKNOWLEDGEMENTS

We thank Nina Taft, Beckett Madden-Woods, and Thibault Guiroy from Google and Luciano Floridi from the Oxford Internet Institute for their valuable feedback in drafting this report. We also thank all the other Googlers involved in the RTBF delisting process.

REFERENCES

- [1] Google Spain SL and Google Inc. v Agencia Española de Protección de Datos (AEPD) and Mario Costeja González. <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A62012CJ0131>.
- [2] Google Inc. v CNIL (C-507/17). <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62017CN0507>, 2017.
- [3] Meg Leta Ambrose and Jef Ausloos. The right to be forgotten across the pond. *Journal of Information Policy*, 2013.
- [4] Article 29 Data Protection Working Party. Guidelines on the implementation of the Court of Justice of the European Union judgment on "Google Spain and Inc v. Agencia Española de Protección de Datos (AEPD) and Mario Costeja González"

- c-131/12. http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp225_en.pdf, 2014.
- [5] Jef Ausloos. The 'right to be forgotten'—worth remembering? *Computer Law & Security Review*, 2012.
- [6] Andrew Besmer and Heather Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the Conference on Human Factors in Computing Systems*, 2010.
- [7] European Commission. Factsheet on the “right to be forgotten” ruling. http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf, 2016.
- [8] David Drummond. Searching for the right balance. <https://europe.googleblog.com/2014/07/searching-for-right-balance.html>.
- [9] Julia Fioretti. France fines google over 'right to be forgotten'. <http://www.reuters.com/article/us-google-france-privacy-idUSKCN0WQ1WX>, 2016.
- [10] Peter Fleischer. Adapting our approach to the european right to be forgotten. <https://www.blog.google/topics/google-europe/adapting-our-approach-to-european-right/>, 2016.
- [11] Luciano Floridi. "the right to be forgotten": a philosophical view. *Jahrbuch für Recht und Ethik-Annual Review of Law and Ethics*, 23(1):30–45, 2015.
- [12] Google. Google search index removal form. https://support.google.com/legal/contact/lr_eudpa?product=websearch.
- [13] Google. Search removals under european privacy law. <https://transparencyreport.google.com/eu-privacy/overview>.
- [14] Google. Transparency Report. <https://transparencyreport.google.com/eu-privacy/overview>, 2018.
- [15] David Hoffman, Paula Bruening, and Sophia Carter. The right to obscurity: How we can implement the google spain decision. *North Carolina Journal of Law & Technology*, 2015.
- [16] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. Multiparty access control for online social networks: model and mechanisms. *IEEE Transactions on Knowledge and Data Engineering*, 2013.
- [17] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the Conference on Computer and Communications Security*, 2015.
- [18] International Telecommunication Union. Percentage of individuals using the Internet. <https://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx>, 2017.
- [19] Jemima Kiss. Dear google: open letter from 80 academics on 'right to be forgotten'. <https://www.theguardian.com/technology/2015/may/14/dear-google-open-letter-from-80-academics-on-right-to-be-forgotten>, 2015.
- [20] Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. We're in it together: interpersonal management of disclosure in social network services. In *Proceedings of the conference on human factors in computing systems*, 2011.
- [21] Microsoft. Content removal requests report. <https://www.microsoft.com/en-us/about/corporate-responsibility/crrr>, 2016.
- [22] Begüm Yavuzdoğan Okumuş and Bensu Aydın. Turkey's court of constitution officially recognizes right to be forgotten. <http://gun.av.tr/turkeys-court-of-constitution-officially-recognizes-right-to-be-forgotten/>, 2016.
- [23] Julia Powles and Rebekah Larsen. Academic commentary: Google spain. <http://www.cambridge-code.org/googlespain.html>, 2016.
- [24] Mathieu Rosemain and Julia Fioretti. French court refers 'right to be forgotten' dispute to top eu court. <http://www.reuters.com/article/us-google-litigation-idUSKBN1A41AS>, 2017.
- [25] Jeffrey Rosen. The right to be forgotten. *Stanford Law Review Online*, 2011.
- [26] RT. Russia's 'right to be forgotten' bill comes into effect. <https://www.rt.com/politics/327681-russia-internet-delete-personal/>, 2016.
- [27] Kurt Thomas, Chris Grier, and David M. Nicol. unFriendly: Multi-Party Privacy Risks in Social Networks. In *Proceedings of the Privacy Enhancing Technologies Symposium*, 2010.
- [28] United Nations, Department of Economic and Social Affairs, Population Division. World population prospects (2017). <https://esa.un.org/unpd/wpp/>, 2017.
- [29] Minhui Xue, Gabriel Magno, Evandro Cunha, Virgilio Almeida, and Keith W Ross. The right to be forgotten in the media: A data-driven study. *Proceedings on Privacy Enhancing Technologies*, 2016.