

---

# Right to Be Delisted? Attitudes Toward Delisting Public Interest Information

Ryan Lloyd and Rebecca Umbach

---

**Abstract.** There is a growing trend of legislation, regulation, and court rulings mandating the delisting of content from intermediary platforms. Few, if any, studies have evaluated user reactions to edge cases involving the delisting of content of public interest. We administered a vignette-based online survey experiment to a representative sample of over 20,000 participants in five countries. We sought to understand user perceptions of delisting content from search engine results and the factors that influence them. While leaving information accessible in search engine results generally leads to warmer feelings toward those search engines, we find that contextual elements also impact this resulting warmth. In addition, we analyze respondents' knowledge and attitudes about the "Right to be Forgotten" (RTBF), perhaps the most well-known legislation on delisting. We find that respondents in countries with active RTBF legislation are more likely to support delisting, know more about RTBF, and support RTBF, and that RTBF knowledge/attitudes affect respondents' answers to questions in our experiment. These results indicate a complex tension around delisting public interest content from search engines' results. Experts sensitive to local context should perform reviews to ensure that delisting requests are handled in a way that meets users' expectations.

---

## 1 Introduction

When a user inputs a query into a search engine in 2025, what do they *not* want to see? In many ways, this question is just as salient nowadays as what users want to see. Legislation, regulations, and court rulings in multiple countries seek to allow individuals to manage what information about them is accessible online. Often this includes a mandate that URLs be removed from search engines' results in certain circumstances within certain jurisdictions. But what exactly do users know about these dynamics, and how do they feel about them?

Perhaps the most emblematic example of this type of legislation is the European Union's General Data Protection Regulation (GDPR) (European Parliament 2016) with its associated "Right to be Forgotten" (RTBF). The literature on GDPR to date has focused on the utilization and understanding of GDPR (e.g., Kwon et al. (2023) and Bertram et al. (2019)), with little scholarship focused on how users understand and perceive the delisting of content, particularly content of potential public interest. How, then, do users prefer that search engines balance the concerns of privacy with the concerns of public interest in practice? What happens in edge cases—unusual instances that do not fit neatly into established categories and procedures—when, for instance, the data subject's conduct could potentially be relevant or important for the public to know about, but the data subject invokes a right to privacy? How should search engines behave when the rights of an individual requester conflict with the rights of information consumers (users), who could have a legitimate right to access this information? And when judicial systems get involved, how do issues such as national sovereignty and the rule of law affect individuals' calculations?

These are especially salient concerns not only because of the spread of RTBF-like legislation throughout the world (Ingram 2023), but also because of individuals' dual positioning in the internet's ecosystem as subjects *and* consumers of information. After all, individuals are not only consumers of searchable URLs—users—but can also be the subjects of searchable URLs, possibly submitting requests for removal themselves. Both search engines and government agencies seek to act on behalf of their users and constituents, but may struggle to decipher their preferences. This is particularly the case when the topic of delisting content from search engine results is so nascent in the public eye that many people do not even know about the existence of this very process.

In this article, we, researchers at Google, begin by giving a brief overview of the process of delisting content from search engines' results. We then describe how legislation or regulation (e.g., the RTBF) can interact with search engines' processes for delisting content. Next, we outline extant research on users' perceptions of such legislation and regulations. We then outline our research questions and methods, describing our vignette-based online survey experiment. Lastly, we share our results and discuss our conclusions.

We find that, when faced with edge cases with public interest content, respondents generally preferred that search engines leave the content accessible rather than remove it from their results. This preference was also reflected in another outcome variable: sentiment warmth. However, the country of the respondent and the particular conditions of our vignette did affect responses, with especially pronounced decreases in warmth when URLs were delisted in response to individuals' requests (as opposed to court orders), particularly in certain countries. This suggests that search engines must thoroughly take into account the specific circumstances of a demand before making a decision on delisting.

Lastly, although a sizable minority of respondents were not familiar with RTBF, attitudes about the RTBF were correlated with respondents' opinions about delisting: those who approved of the RTBF were more likely to feel warmer about delisting than those who did not approve of it. Even so, a majority of both pro- and anti-RTBF respondents still supported leaving the URL in our vignette accessible.

### 1.1 Delisting Content from Search Engine Results

The main role of a search engine is to make information that is published elsewhere on the internet accessible, typically organized via proprietary algorithms. This generally means that the most relevant and helpful results are presented in descending order, because the business models of search engines are most commonly predicated on user traffic. While all search engines are equipped with the same pool of information, product and policy choices mean that for the same query, different search engines in the same country may surface either different results, or the same results in a different order.

Most search engines have public-facing policies about which types of content are permissible to be displayed in their results (Google 2023b; Microsoft 2024d; Google 2023c); some of this material might be proactively removed by the company in question, while individual requesters can also demand the delisting of material that they believe violates these product policies (Google 2025b; Microsoft 2025a). Requesters can also appeal to government agencies or obtain court orders to demand delisting. In some cases, content may violate both the law of a country and platform policies (for example, child sexual abuse material), and will duly be delisted upon notice.

Typically, if content does not violate a platform's policy, it will remain accessible in search engine results unless a requesting individual, government body, or court has specific recourse under their local laws to mandate its delisting. Platforms' approaches regarding removal requests on legal grounds depend in large part on the law of a given country. In countries where internet platforms can be held liable for their users' activities, platforms may be more likely to accede to requesters' demands. In some countries, especially those in which immunity laws preclude an obligation to remove, platforms may perform removals in response to court orders (Keller 2018).

By far the biggest portion of legal removal requests relates to copyright; Google has received requests to remove more than 10 billion URLs from Search for copyright violations (Google 2024a). In the first half of 2024, Bing received requests to remove more than 260 million URLs for copyright violation, and reported a 99.68% acceptance rate (Microsoft 2024a). Other removal requests—often from governments and regulators—can relate to similarly uncontroversial topics such as security (e.g., phishing sites and scams), regulated goods and services, or criminal activity (Google 2024b). In these cases, both the search engine and the requester are aligned on the necessity of this removal.

In some cases, however, certain parties may request delisting for information that is

arguably in the public interest. For example, a teacher may not want information about their sexual misconduct with a minor (perhaps documented on a news source or via court filings) to be indexed on a search engine, easily discoverable. Even if that information still exists at the source, other individuals' ability to make informed decisions is impaired by removing easy access to that information.

In another case, an authoritarian government might demand the removal of opposition political parties' websites or posts for violating their country's laws of national security. This might guarantee compliance with the legal regime of the country in question, but also suppress the political rights of its citizens.

## 1.2 Government Legislation on Delisting

In the United States, Section 230 of the U.S. Communications Decency Act of 1996 (or simply Section 230 or CDA 230) is a well-known piece of legislation that states that “[n]o provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider,” which has effectively guaranteed broad protection for private intermediaries displaying user-generated content (Brannon and Holmes 2024). Since courts have considered search engines to be a type of intermediary, and not publishers of content, they have been largely shielded from many types of civil and criminal liability (although not necessarily given a “free pass”—see Ardia (2010)). However, many countries have extant legislation, regulation, or court rulings that not only do not provide safe harbor for search engines, but can even mandate content delisting under certain circumstances.

Within this wider universe of legal frameworks, perhaps the most well-known example of this type of legislation is the General Data Protection Regulation (GDPR). The GDPR is a European Union (EU) regulation that governs how personal data must be collected, processed, and erased. GDPR Article 17 codifies the “Right to be Forgotten” recognized by the European Court of Justice under prior data protection law (see GDPR.eu (2018)). This entitles citizens of the EU (and indeed, the European Economic Area, or EEA<sup>1</sup>) to request the delisting of URLs from results that appear when their name is used as a query under certain circumstances, subject to a public interest test (Article 29 Data Protection Working Party 2024). For these name-based, query-based removals, a URL would be removed from the results displayed for a query including the name “Jane Doe,” but not from the search engine results for other queries.

The concept of the Right to be Forgotten was launched by a landmark legal case in Spain. In 2009, Mario Costeja González, a Spanish lawyer, requested that Google remove links to an old government notice of the sale of his house due to social security debts. Google refused, and Costeja González sued Google. In May 2014, the European Court of Justice decided in favor of Costeja González, and went on to rule that search engines had to remove links from search results if requested by an individual when the links in question

---

1. i.e., the EU plus Iceland, Liechtenstein, and Norway.

went to web pages containing information about that individual that is “inadequate, irrelevant or no longer relevant, or excessive in relation to those purposes and in the light of the time that has elapsed” (Xue et al. 2016). Search engines were tasked with evaluating such requests and instructed to strike a fair balance between fundamental rights and interests, privileging the rights of the data subject “as a general rule, over the economic interests of the search engine and that of internet users to have access to the personal information through the search engine,” but also seeking to limit the impact of any delisting on access to information (Article 29 Data Protection Working Party 2024). There are public interest exceptions carved into the GDPR, including when “the data is being used to perform a task that is being carried out in the public interest or when exercising an organization’s official authority” and when the data “is being used to exercise the right of freedom of expression and information” (Wolford 2018). The GDPR therefore recognizes that freedom of expression and information often conflict with the Right to be Forgotten, and that the two need to be weighed appropriately.

The Costeja González case required search engines to develop the necessary processes to field and respond to such requests, usually but not exclusively from individuals as data subjects.<sup>2</sup> Google (Google 2023a) and Bing (Microsoft 2025b), for example, both have specific Right to be Forgotten personal data removal request forms, wherein they specify that requests will be balanced with the public interest in having access to the information. Moreover, both search engines publish a transparency report detailing the relevant data surrounding governments’ and courts’ requests to remove content (Google 2024b; Microsoft 2024c). Subsequent cases have resulted in the current system of geofenced implementation, wherein URLs are blocked from appearing in search engine results—or delisted—specifically for users who are searching from within the EU and inputting certain queries (Nunziato 2017). Requests for delisting are operationally costly, as more than 1.6 million requests regarding 6.6 million URLs have been submitted to Google since May 2014 (Google 2024b), and their complexity requires careful review to ensure appropriate compliance while detecting fraudulent or erroneous submissions (Google 2024c).<sup>3</sup>

In terms of legal frameworks around content delisting, RTBF has garnered the most

---

2. One study (Bertram et al. 2019) found that requesters from France, Germany, and the UK accounted for approximately half of the overall RTBF requests. Additionally, the top thousand requesters (commonly law firms and reputation management companies) generated 16.3% of requests and 23.4% of all delistings, generally as representatives of their clients.

3. Google’s rate of delisting under European privacy law as of November 2024 was 51% (Google 2024c), whereas Bing’s was 43% over the most recently reported period of January–June 2024 (Microsoft 2024b). Delisting may not happen in response to a request if the content has already been taken down by the publisher or if the search engine decides the content is in the public interest; this could be related either to the type of content itself or to details surrounding the individual in question (e.g., professional misconduct). Common reasons for requests include private or sensitive information, content that relates to criminal activity when the requester was a minor, and content that relates to criminal activity more broadly, particularly when the accused was exonerated, acquitted, or had a conviction that has now been spent (Google 2025a). Google also reports the percentage of requests removed by content category since January 2016, and one can see that removal rates range from 99% for requests involving personal information to 31% for requests involving political content as of June 2024 (leaving aside the considerable volume of requests not acted upon for insufficient information, which makes up 25% of all requests between January 2016 and November 2024).

attention in the scholarly literature.<sup>4</sup> However, the GDPR's Right to be Forgotten is merely one example of legislation that mandates the delisting of certain content, and its effects apply specifically and exclusively to EEA countries. Concepts related to delisting content are increasingly becoming instantiated in other jurisdictions as well (Erdos 2021; Nunziato 2017). In an analysis of the G20 countries, Erdos (2021) found that 15 of the 19 G20 states (inclusive of four countries covered at the time under GDPR) had data protection laws that could plausibly allow individuals to request delisting. Nunziato (2017) examined the expansion of RTBF, both beyond Europe and to media publishers themselves, raising concerns about a "global threat to freedom of expression." Some countries have cited the concepts in court rulings and/or legislation (e.g., Japan, India) (Nunziato 2017; Times of India 2023; Arakali 2025), while others have instantiated the concept into law (e.g., Türkiye, Russia) (Bertram et al. 2019; Nurullaev 2015). In short, the RTBF has taken on many guises in different contexts, both inside and outside the EEA (Van Calster, Gonzalez Arreaza, and Apers 2018; Nunziato 2017; Erdos 2021).

Indeed, while the RTBF may have inspired other legislation, circumstances can vary widely by country regarding the specifics and the scopes of such laws. It would be a stretch to say the RTBF, as interpreted within the EEA, is the norm. The RTBF, for example, has a delimited scope of delisting; whether a URL is removed from a list of results depends on the query (specifically, whether the respondent's name is included in it). This is not the case for other regulations around the world on the matter, in which delisting might be demanded for given URLs regardless of the query that the user inputs. It also relies largely on individuals (as data subjects) to initiate requests for content removal; government data protection authorities then represent these individuals and adjudicate the subsequent delisting process. This is not the only way to demand delisting, however, and government agencies and courts can demand delistings as well, both in one-off cases and as part of a sustained policy, legislation, or regulation. Search engines are generally protected from intermediary liability in most jurisdictions, though, as long as they respond promptly to valid legal notice. Lastly, the RTBF has a robust framework around the protection of the public interest in order to avoid scenarios of over-removal (see Article 29 Data Protection Working Party (2024)). In other situations, however, such nuances and subtleties might not always be incorporated into the legislation. In practice, there is considerable variation in how different government and judicial bodies demand delisting.

---

4. Several studies have assessed the effects of RTBF, specifically looking at the characteristics of requests to better understand who is leveraging this law and how. Of note for this study, news sites accounted for 18.5% of requested URLs, the largest named category (miscellaneous accounts for 50% of URLs), but had a relatively low delisting rate of 34.5% (as opposed to social media, which had a 52.9% delisting rate) (Xue et al. 2016). One study (Xue et al. 2016) analyzed articles in the UK media that were subject to RTBF delisting and found that the top six article themes were all criminal, including assault, murder, and sexual assault. Another study found that 33% of all requested URLs under RTBF were related to personal information on social media and directory services, while 20% were related to news outlets and government websites reporting on an individual's legal history (Bertram et al. 2019). Given that news sites tend to have a higher bar for publication than social media and could be more likely to publish information of public interest, this could indicate that the public interest carve-outs in RTBF are indeed leading to fewer removals of public interest content.

### 1.3 User Perceptions of Delisting Legislation

As legislation mandating delisting from search engines has diffused throughout the world, scholars and practitioners have become interested in individuals' perceptions and opinions about this legislation. However, this literature is still in its nascency, and there is little data available on this topic. As of now, this body of literature seems to be limited to Europe and the United States.

Scholarship indicates that awareness of the GDPR within the European Union has substantial variation at the individual and group levels, even when accounting for respondents' digital experience. While digital experience has a significant positive correlation with GDPR awareness on aggregate, there is a large cluster of respondents with extensive digital experience and little awareness of the GDPR or RTBF (Rughinis et al. 2021).

After the initial Costeja González ruling, various U.S. news outlets began to raise the possibility of applicability to the United States (Luckerson 2014). One 2017 vignette-based study (Bode and Jones 2017) of U.S. participants found mean support for a right to be forgotten law to be 3.7/5—slightly favorable toward the possibility. Across conditions, participants were more likely than not to endorse the law, with a slight decrease in agreement seen in participants who were told that a government agency (as opposed to a search engine or website) would be the one in charge of executing takedown requests. However, results should be interpreted in light of the recruitment strategy—respondents were recruited from Mechanical Turk and the authors conceded that this sample is neither representative nor random.

A Pew study in 2019 (Auxier 2020) found that while the majority of respondents endorsed the right to remove potentially embarrassing photos or videos, as well as potentially sensitive personal information (e.g., financial data, medical data), only 39% agreed that personal data collected by law enforcement (e.g., mugshots, criminal records) should be delisted from search engine results. Responses were affected by demographic characteristics, including political affiliation, race and ethnicity, age, and gender.

Across these two studies, it is clear that the type of content will affect U.S. respondents' willingness for it to be eliminated, with more latitude given toward arguably irrelevant media versus content that may be related to public safety or public interest. Additionally, there is less support when a government agency is the requester, which may speak to respondent concerns around bias or abusive government agencies. Given the United States' emphasis on freedom of information and speech, expansion and application of RTBF to the United States does not appear likely, at least in the immediate future.

### 1.4 This Study

Previous surveys are relatively far and few between, and have typically surveyed individuals from single countries. This study surveys participants in five democratic countries

(Economist Intelligence Unit 2022). We consider a number of factors in our vignette-based study. First, we vary whether the demand in question was issued by an individual requester (without the involvement of a data processing authority or a court), or through a court order. While these are not the only types of demands for delisting that search engines receive across the world (e.g., Data Protection Authorities and government agencies), they are the two most common. Second, we vary whether the URL in question was published by a blog (presumably less legitimate or well-researched) or by a well-known news outlet. Third, we varied the misconduct described to include two different types of misconduct often seen in these types of delisting requests: sexual misconduct and corruption. Fourth, we varied whether the search engine complied with the demand or refused to perform the delisting. These vignette scenarios are amalgamations of the types of demands with plausible public interest considerations often faced by search engines.

We used this research design to attempt to answer four questions:

- **RQ1:** *How do demographic characteristics interact with search engine decisions to influence user feelings toward a given search engine?*
- **RQ2:** *How do search engine decisions on content removal influence user feelings toward that search engine?*
- **RQ3:** *How do the characteristics of each demand (type of demand, type of media outlet, type of misconduct, search engine decision) and country of respondent affect user attitudes toward the search engine's decision?*
- **RQ4:** *How do respondents feel about the concept of RTBF and how does it affect their responses to our vignettes?*

Our first research question acknowledges that identity characteristics, including age and gender, may play an important role in influencing sentiment. We also expected that search engine decision may interact with these identity characteristics. For example, specific types of users may privilege freedom of information more than others, and look particularly askance at a search engine delisting content.

For our second research question, given the number of factors at play, we were *a priori* unsure about the effects of the varied conditions on respondent warmth. We did, however, expect a considerable degree of variation between countries.

For our third research question, we expected that delisting the URL would be associated with warmer feelings among respondents when the demand in question was a court order, was related to a blog posting (not a well-known news outlet), and was related to credible corruption allegations (rather than sexual assault).

Lastly, we expected to see that users with positive views of the RTBF would be more likely to feel warmly toward search engine companies when they delisted the content in

question.

## 2 Methods

We deployed a brand-blind (i.e., respondents were unaware of Google's involvement) online survey to 21,678 participants across five countries via a well-known market research company with global coverage. Participant demographics are available in Table 4. While we aimed for a minimum of 4,800 respondents in each country, the vendor company was only able to complete half that number in Poland.<sup>5</sup>

Our five countries were selected to ensure as broad a geographic coverage of countries as was feasible, with a variety of regimes for mandating content delisting (see Table 1). We also sought to specifically include respondents from countries with high levels of political polarization and differing levels of trust in political institutions, in order to examine possible effects of the politicization of content moderation.

Table 1: Breakdown of countries in our sample by whether a Right to be Forgotten is currently active in that country

Country	Active RTBF?
Canada	Only in Québec
France	Yes
India	No (but right to privacy)
Poland	Yes
USA	No

Participants first answered several demographic questions, then were asked to identify their preferred search engine (see Table 2) and to indicate their level of warmth toward that search engine (out of a possible 100 points). The majority of respondents (almost 90%) preferred Google Search, with between 1 and 3.8% of respondents choosing the remaining options.

The name of their preferred search engine was then piped into the vignette as the search engine in question, as part of the brand-blind research design.<sup>6</sup>

We fielded a fully crossed two (type of outlet: news outlet or blog) by two (financial corruption or sexual assault by religious figure) by two (request by person of interest or court order) by two (search engine attending to the demand and delisting, or refusing to attend and leaving the link accessible in the search engine's results) experiment centering

5. We also obtained a specific subsample in Québec ( $n=1,205$ ) given the existence of a Right to be Forgotten in Quebec, but not the rest of Canada ( $n=3,600$ )

6. It is worth noting that different search engines use different algorithms and methods to assess the relevance of a URL and rank it in results. This difference in products could also affect which search engine a consumer decides to use and why. However, because the vast majority of users selected Google as their preferred search engine (and the subsamples that preferred the remaining search engines were so small), we did not subdivide our sample by search engine for any analyses.

Table 2: Breakdown of percentage of respondents choosing each search engine as “preferred.” Almost 90% choose Google; only 3.2% chose the next-most popular search engine.

Search Engine	Percentage of Respondents
Bing	3.0%
DuckDuckGo	3.2%
Ecosia	1.0%
Google	89.0%
Other	3.8%

on a request made to a search engine to delist a URL. Each respondent received a single vignette at random. They were then asked their feelings toward their preferred search engine and to indicate what they thought the search engine **should have done** given the vignette scenario (answer options were “Attend to the takedown demand and make the link/content unsearchable,” “Refuse to attend to the takedown demand and leave the link/content accessible through [PREFERRED SEARCH ENGINE],” and “Other”). By experimentally manipulating the conditions for a public interest case, we can understand the greatest reputational risks associated with public interest delistings. At the end, participants answered questions about their knowledge and feelings about the Right to be Forgotten, an EEA-specific law under the General Data Protection Regulation. Two

Table 3: Vignette examples. These are two examples of the 16 vignettes used in our experiment.

Vignette
<p>Suppose that a local blog publishes a post claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf for the Minister of Health. This politician subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]’s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] attends to the court order, and removes the URL in question from its search results.</p>
<p>Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]’s search results. [PREFERRED SEARCH ENGINE] refuses to attend to the takedown demand and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].</p>

examples of the vignettes that were used are shown in Table 3. All vignette text can be seen in the Appendix in Table 8. All data analyses are conducted by the authors using R (R Core Team 2024). Packages used include the emmeans package (Lenth et al. 2022) and the ggplot (Wickham 2016) and xtable (Dahl 2016) packages for visualization.

## 2.1 Limitations

There are several limitations to our study. Perhaps the biggest limitation comes from the methodological choice to use a vignette-based survey experiment. While the scenarios used in this study are realistic amalgamations of edge-case requests received by content-moderation practitioners at search engines, they are not real requests due to legal and privacy considerations. The selection of our 16 experimental vignette conditions was made based on our experience of contentious edge cases, and aimed to inductively identify relevant contextual elements for users. As with any survey, however, presenting users with other scenarios could plausibly lead to different user opinions.

Another limitation is the menu of choices we provided to our respondents. We simplified the possible responses to our scenarios, limiting them to “attend to the takedown demand,” “refuse to attend to the takedown demand,” and “other.” Some respondents selected “other,” with many detailing sophisticated and nuanced positions that showed an admirable grasp of the complexities of these situations. As these respondents noted, actual responses from search engines might very well deviate from a dichotomy of either immediately removing or not removing content from search engine results. Our choice of responses purposefully simplified the situation, perhaps channeling some respondents into one of those two responses when they were convinced of neither one. That said, while we acknowledge this limitation, it is important to also recognize that this same dilemma can face practitioners at search engine companies, and that creative solutions or half-measures are often not viable options when faced with truly complex dilemmas.

A third limitation is that we only included two types of demands for removal from search engine results: requests from individuals and court orders. In reality, many countries—particularly in the EEA, where the RTBF is active—have Data Protection Authorities who can represent individual requesters and make removal demands on their behalf. In addition, government agencies can also make demands, often on individuals’ behalf. Given the presence of non-EEA countries in our sample, and the variation to which Data Protection Authorities and government authorities make these demands, we opted to simplify our scenario and restrict it to examples that would translate well to every country. However, the government often does play a role in removal demands in some form, and portraying that in vignette experiments would be a worthy exercise for future studies.

## 2.2 Measures

The key outcome variables were the post-vignette level of warmth for the search engine in question, and user opinion on what the search engine should do given the condition at hand. We were specifically interested in how the manipulated elements of the public interest case might affect the level of warmth, and whether user opinion on the search engine actions would vary by country. The former was operationalized

by asking respondents, “If this scenario happened in real life, what would your rating be?” with answers provided on a sliding scale from 0 to 100. The latter was assessed by asking “What action do you think [your preferred search engine] should take in response to this scenario?” Possible answers included “Attend to the takedown demand,” “Refuse to attend to the takedown demand,” and “Other.” If participants selected “Other,” they were prompted to elaborate further.

### 3 Findings

Participant demographics can be seen in Table 4,<sup>7</sup> and descriptive statistics (including average mean post-vignette warmth by identity characteristics and country) can be seen in Table 5.

Table 4: Participant demographic characteristics, divided by age, sex, and country

Variable	Group	%(#) of Participants
Age	18–25	14.8% (n = 3215)
	26–35	19.5% (n = 4232)
	36–45	18.6% (n = 4034)
	46–55	16.4% (n = 3566)
	56–65	15.3% (n = 3315)
	66+	15.3% (n=3316)
Sex	Female	49.1% (n = 10637)
	Male	50.4% (n = 10928)
	Non-Binary	0.3% (n = 72)
	Prefer not to say	0.2% (n = 41)
Country	Canada	22.2% (n = 4805)
	France	22.2% (n = 4818)
	India	22.2% (n = 4818)
	Poland	11.2% (n = 2420)
	USA	22.2% (n = 4817)

#### 3.1 Demographic Characteristics and Warmth

The age group of the respondent played a crucial role in the mean search engine warmth, as seen in Table 5 ( $F(5, 21672) = 6.837, p < 0.001$ ). Tukey HSD tests (not shown) revealed that the youngest age group (18-25-year-olds) was entirely responsible for this finding; as compared to every other age group, this youngest group felt the least warmth toward search engines. A follow-up analysis of variance (ANOVA) examining the interaction between search engine action and age revealed a significant interaction term

7. Within our Canada sample, we had a subsample of 1,205 in Québec, with the remaining 3,600 coming from the rest of Canada.

Table 5: Average warmth (with standard errors) toward preferred search engine by country, age, and sex. The lowest level of warmth is from French respondents, at 69.38, and the highest level is from Indian respondents, at 77.13.

Variable	Group	Mean Warmth	<i>n</i>	<i>SE</i>
Country	Canada	69.45	4805	0.37
	France	69.38	4818	0.34
	India	77.13	4818	0.33
	Poland	73.48	2420	0.48
	USA	70.28	4817	0.36
Age	18–25	69.62	3215	0.41
	26–35	72.31	4232	0.37
	36–45	72.45	4034	0.39
	46–55	72.37	3566	0.41
	56–65	72.16	3315	0.43
	66+	71.34	3316	0.43
Sex	Female	72.53	10637	0.23
	Male	71.11	10928	0.24

( $F(5, 21666) = 5.54, p < 0.001$ ), showing that this youngest age group has equivalent warmth with the older age groups when the search engine action is to delist, but lower warmth when the search engine leaves the content accessible.

Examining sex (and subsetting the data to include just female and male due to limited sample sizes), females felt significantly warmer toward search engines as compared to males ( $F(1, 21563) = 18.25, p < 0.001$ ). An additional ANOVA with a search engine action\*sex interaction term ( $F(1, 21561) = 14.27, p < 0.001$ ) reveals that the differences are only significant in conditions where the search engine chooses to delist, in which case females' warmth is higher than males' warmth ( $M_{female} = 68.6, M_{male} = 65.9, p < 0.001$ ).

Results across these demographic factors suggest that identity characteristics do play a role in perceptions toward search engines broadly, and are also likely to influence reactions to search engine responses to demands for content delisting.

### 3.2 Treatment Effects on Warmth by Vignette

To determine whether we needed to run an analysis of covariance (ANCOVA), we first confirmed with a simple ANOVA that the vignette groups did not significantly differ in pre-test warmth ( $F(1, 21676) = 1.35, p = 0.25$ ). Between-groups ANOVA was then used to identify significant main and interaction effects of outlet type, action by search engine, country, crime reported, and type of request on warmth, followed by post hoc Tukey HSD tests when ANOVAs revealed significant main effects. Mean warmth across all 80 possible vignette combinations are depicted in Figure 1 on the following page ( $2_{outlet}^*$

$$2_{crime} * 2_{search\ engine\ response} * 2_{request\ type} * 5_{country}$$




Figure 1: Warmth post-vignette with 95% confidence intervals. The most important driver of differences in warmth is the decision to delist or not by the search engine, rather than the type of crime, publication outlet, or nature of the demand.

Across conditions, participants generally felt more warmly toward the search engine when it chose to leave the content accessible; this was particularly true when the type of demand was a request from an individual (as opposed to a court order). Perhaps unsurprisingly, the inverse is also true—participants felt least warmly toward the search engine when it decided to delist a URL from its results by request only (Table 6).

Table 6: Average warmth by country and search engine action toward preferred search engine. Across all countries, warmth is higher when the search engines leave the content available as compared to when they delist, although India shows the least difference.

Country	Action	Mean	SE	df	95% confidence interval
Canada	Delist	63.12	0.48	21598	62.18–64.06
France	Delist	64.32	0.48	21598	63.38–65.26
India	Delist	74.29	0.48	21598	73.35–75.23
Poland	Delist	70.07	0.67	21598	68.75–71.40
USA	Delist	65.55	0.48	21598	64.61–66.49
Canada	Leave up	75.84	0.48	21598	74.90–76.78
France	Leave up	74.40	0.48	21598	73.46–75.34
India	Leave up	79.95	0.48	21598	79.02–80.89
Poland	Leave up	76.84	0.68	21598	75.51–78.16
USA	Leave up	75.02	0.48	21598	74.08–75.96

Results are averaged over the levels of: Type of Request, Crime, Outlet

As seen in Table 7, a five-way ANOVA showed a significant main effect for the type of

demand, type of misconduct, and the action the search engine took. The type of outlet that originally published the URL in question did not have a significant main effect. Of these main effect results, the most important one to note is that the search engine action had an effect; the rest are not relevant to our experimental manipulation.

The main effects also should be interpreted in light of significant interaction effects for the type of demand and type of misconduct, for search engine action and outlet, as well as for the type of demand and outlet. When the action taken is part of the interaction term, these results show how the particular conditions of the vignettes affected respondents' receptions of the search engines' actions.

The partial  $\eta_p^2$  reveal that the largest effects are seen in how the search engine reacted (small–medium effect), in the country of the respondent (small effect), and in the interaction between search engine action and requester (small effect).

Table 7: ANOVA results, with partial  $\eta_p^2$ . Excludes those interaction coefficients where  $p > 0.05$ . Significant main effects were seen for every vignette condition, with the exception of outlet. The strongest main effect was seen in the search engine response. There is also a significant interaction effect between type of request and search engine action.

	<i>df</i>	F value	Pr(>F)	partial $\eta_p^2$
Type of request	1.00	33.08	0.00	0.00
Crime	1.00	15.54	0.00	0.00
Action	1.00	834.30	0.00	0.04
Outlet	1.00	0.59	0.44	0.00
Country	4.00	94.76	0.00	0.02
Type of request:Action	1.00	341.76	0.00	0.02
Type of request:Outlet	1.00	5.26	0.02	0.00
Action:Outlet	1.00	34.15	0.00	0.00
Action:Country	4.00	15.90	0.00	0.00
Type of request:Crime:Country	4.00	3.03	0.02	0.00
Crime:Action:Country	4.00	3.10	0.01	0.00
Action:Outlet:Country	4.00	2.76	0.03	0.00
Residuals	550.98			

Tukey's HSD post hoc tests were conducted to probe these results, and in the interest of readability, we focus on those with partial  $\eta_p^2 > 0.01$ . Respondents felt more harshly toward search engines if they attended to the demand ( $M_{\text{delist}} = 67.5$ ) as opposed to refusing to do so ( $M_{\text{leave up}} = 76.4$ , estimated difference = 8.94,  $t = 26.98$ ,  $p < 0.001$ ). There was a main effect of country, such that the countries had different levels of warmth toward the search engines, aggregated across the different vignettes. This is most easily seen in Tables 5 and 6. Tukey HSD tests confirm that Canada, the United States, and France are all equivalent, and that Poland and India differ from every other country (India having the most warmth and Poland falling in the middle), all  $p < 0.001$ .

Looking at the interaction effects, there was a significant interaction between search

engine action and type of request. Namely, respondents felt less warm toward search engines when they delisted content ( $M_{\text{Delist: Court order}} = 71.4$ ,  $M_{\text{Delist: Request}} = 63.5$ ) than when they left it up ( $M_{\text{Leave up: Court order}} = 74.4$ ,  $M_{\text{Leave up: Request}} = 78.4$ ), regardless of requester; however, the drop in warmth was most notable when a search engine chose to delist (rather than leave the content accessible) in response to a request from an individual (estimated difference=14.89,  $t(21598)=31.826$ ,  $p<0.001$ ).

### 3.3 Preferred Actions

Respondents were also asked about what they would *prefer* that the search engine does in response to the scenario. The answer options were dichotomous, but with a third “other, fill in the blank” option. Responses are depicted graphically in Figure 2, broken out by country.

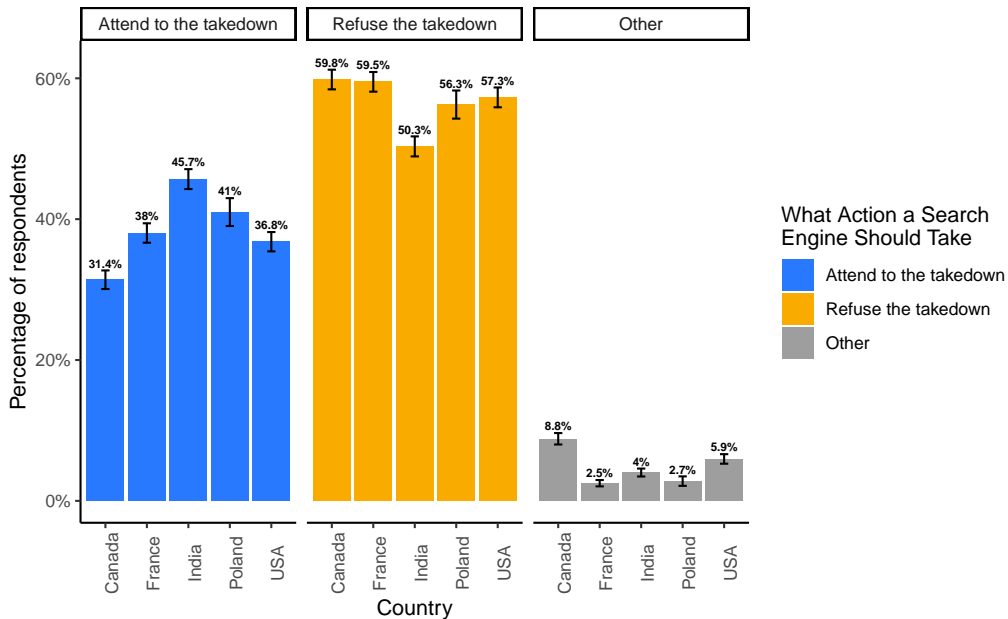


Figure 2: What respondents think search engines should do, by country. 95% confidence intervals in black. A majority in all five countries wanted the search engine to refuse the takedown.

Aggregated across vignette conditions, respondents were more likely in every country (including those covered by the GDPR) to say that the search engine should leave the information accessible in their results (56.7%, 95% CI=56.0–57.3, not pictured). India was the outlier here—only 50.3% of our respondents in India preferred that the content remain accessible, while this rate stayed between 56% and 60% in the other four countries. A noteworthy minority (about 5% of respondents overall) chose the third, “other” option, choosing to write in their own response; in fact, 8.8% of Canadian respondents chose this option.

Fill-in-the-blank responses could be summarized as falling into one of four buck-

ets:

1. The search engine should verify the veracity of the statement and act accordingly.

*“What I’d like the search engine to do is remove the URL if the claims are actually false, and keep the URL if the claims are actually true, with zero margin for error. Which is difficult! But when you’re powerful, you’re called upon to do difficult things.”*

2. The search engine should appeal (in the case of a court order).

*“There are means to get a court order overridden, though it will cost \$\$\$.”*

3. The search engine may abide but notify the user that there *used to be* available information, though it’s been removed.

*“This is a matter of public information. The search engine has to abide by laws, but this take down notice should be mentioned if I search for that event.”*

4. The outlet itself should be the one held accountable.

*“Legal action should be taken against the blogger...if a newspaper is slandering someone, you do not take action against the person delivering the mail.”*

These responses reject the premise of a simple dichotomous response, and acknowledge both additional relevant parties (the outlet), and additional options that they believe the search engine could take.

### **3.4 Right to Be Forgotten Knowledge, Opinions, and Effects**

After the vignettes were presented, participants were asked “how familiar are you with the concept of ‘Right to be Forgotten’” (RTBF). Given the recent adoption of RTBF by the EU and the subsequent discussions about it in countries throughout the world, we were interested in understanding more about users’ familiarity and opinions on it. In particular, we were interested in whether respondents’ opinions on this matter were associated at an individual or country level with downstream opinions on our dependent attitudinal variables. This was an especially interesting question to consider in terms of the extent to which content moderation had entered the public consciousness and begun to serve as a fault line for polarization.

Results are presented in Figure 3. Unsurprisingly, France and Poland both had relatively high awareness of a law that directly impacts them, with around 50% of participants saying they were quite familiar or knew a little bit about it; India also reported relatively high awareness, which may be related to past discussions related to RTBF in the country. The majority of North American participants were completely unfamiliar with the term (about 54% in both Canada and the United States selected “have never heard of this until now”).

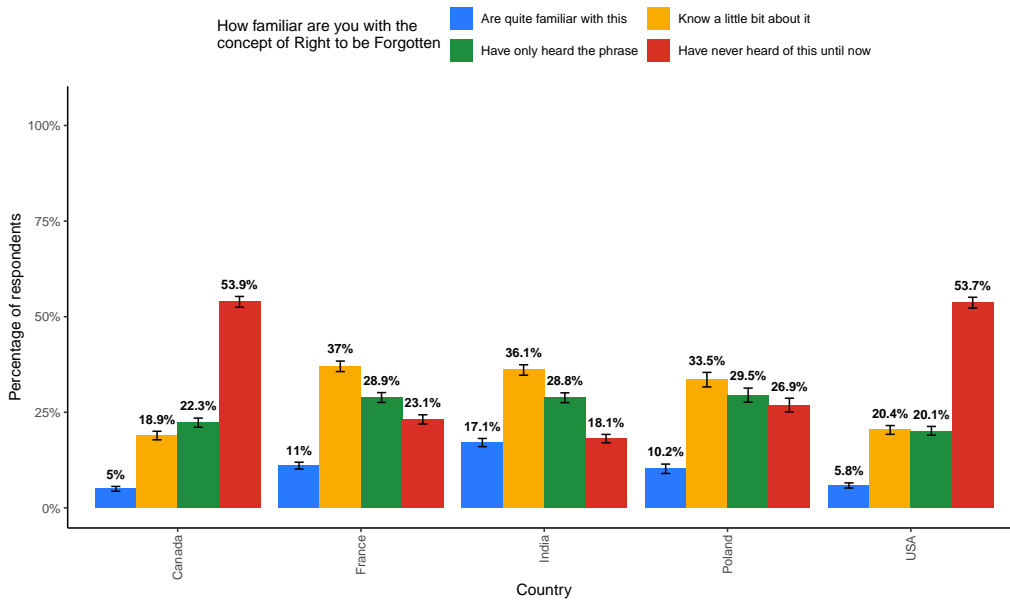


Figure 3: Familiarity with the Right to be Forgotten by country. 95% confidence intervals in black. Respondents selected one of the four options seen in the legend. Majorities in non-RTBF countries had never heard of RTBF, while large minorities in India and RTBF countries were also unfamiliar with the term.




Figure 4: Opinions on the Right to be Forgotten by country. 95% confidence intervals in black. Countries with relevant legal frameworks have higher rates of agreement. A large minority of respondents (between 14.9% and 25.9%) indicated they weren't sure or couldn't say.

We also asked participants for their opinion on the concepts surrounding the Right to be Forgotten. Findings can be seen in Figure 4.

Countries fell into one of three patterns. In most countries, more respondents agree with the principles of RTBF. In some countries, however, (e.g., Canada, India, the United States), it is not clear whether the differences between two main groups of interest are statistically significant. Accordingly, we run two-sided proportion tests (excluding the unsure group) for those three countries. In Canada, the proportion of respondents agreeing is equivalent to the proportion of respondents disagreeing ( $\chi^2 = 1.64$ ,  $df = 1$ ,  $p = 0.20$ ). In the United States, the proportion of respondents disagreeing is significantly greater than the proportion agreeing ( $\chi^2 = 5.61$ ,  $df = 1$ ,  $p < 0.05$ ), and in India, the opposite is the case ( $\chi^2 = 20.09$ ,  $df = 1$ ,  $p < 0.001$ ), meaning India falls more in line with EU countries. The United States was the only country where significantly more participants indicated disagreement rather than agreement with the principles of RTBF.

There does appear to be a relationship between respondents' opinion on the RTBF and the action they prefer search engines to take given the vignette scenario. As shown in Figure 5, more of those who agreed with the concept of a right to be forgotten wanted search engines to delist the URL in our vignette than those who did not support a RTBF (by 10 percentage points); those who selected "other" as their response were excluded from this graph. It is notable, however, that majorities of both the pro-RTBF and anti-RTBF respondents still wanted the content in the vignette to remain accessible on the search engine.




Figure 5: Support of the Right to be Forgotten by preferred action. More respondents who supported a Right to be Forgotten believe that the search engine should delist the material in our vignette than respondents who did not support a RTBF.

## 4 Discussion

This vignette-based survey experiment aimed to better understand how respondents felt about the actions a search engine may take in response to delisting requests, specifically when the content in question has plausible public interest relevancy. We obtained results in several different areas related to our research questions.

### 4.1 Identity and Demographic Characteristics

First, we find that the demographic characteristics of our respondents did play a role in determining how they reacted to our treatment. Younger respondents (ages 18–25) felt less warmth toward search engines compared to other age groups, specifically when the search engine chose to leave the URL accessible in our vignette. In addition, females were less punitive as compared to males when the search engine opted to delist the content in question.

To our knowledge, there is no extant research on the relationship between identity and users' preferences for search engine action. This could be a useful area for future research to verify this finding and examine the reasons behind it.

### 4.2 Effects of Search Engine Action on User Sentiment

Second, the action taken by the search engine in our vignette had significant effects on our outcome variables. In terms of our warmth measure, leaving content accessible led to lesser drops in warmth than delisting the content across most of our experimental conditions in most countries. The action taken had a significant main effect on warmth, as shown by a five-way ANOVA test. In terms of preferred action, a majority of respondents (56.7%) preferred that the search engine keep the content accessible. Furthermore, a majority of respondents in all five of our countries preferred that the content remain available through search engines.<sup>8</sup>

While the effect of search engine action varied with the circumstances of the demand (see Subsection 4.3 below), the consistency of this finding was noteworthy, and indeed surprising. In five vastly different countries, with two different outcome measures, across 16 different experimental conditions, respondents felt more warmly toward search engines when content remained accessible. Given the variety of conditions and the heterogeneity of our sample, this was not necessarily our *a priori* expectation.

### 4.3 Effects of Case Characteristics on User Sentiment

Third, we find that the particular characteristics of the delisting demand affected users' opinions. Our five-way ANOVA test showed that there was a significant interaction effect

---

8. Note, however, that the confidence interval for leaving the content accessible in India does cross the 50% threshold.

on respondent warmth between the action the search engine took in our vignette and both the type of demand (individual request or court order) and the outlet publishing the content (newspaper or blog). The country of the respondent also affected how the search engine's action was seen. Additionally, there were relevant interactive effects between the action taken, the offense committed, and the country of the respondent, as well as between the action taken, the outlet reporting, and the country of the respondent.

In practice, the action taken by the search engine had the largest effect on the respondents' levels of warmth (see Subsection 4.2 above). After that effect, the second-largest was the type of demand (court order vs. individual demand), depending on whether the content was left accessible or delisted. In other words, search engines suffered an even larger penalty in warmth if they delisted pursuant to an individual demand (as opposed to a court order). Certain countries also punished search engines more for delisting the sexual assault case or the newspaper report (as opposed to the corruption case and the blog post, respectively). This shows that respondents are sensitive to the specific aspects of a case, and that search engine delisting is viewed as more or less reasonable depending on the additional context surrounding the situation.

One can see these differences in practice in Figure 1. On the one hand, the difference in warmth between leaving the URL accessible and delisting it tends to be uniformly high, regardless of country, when there is an individual request; the difference tends to be more than 10 percentage points in most cases. That is not to say there is no variation within this category; the difference reaches more than 26 percentage points for corruption in Canada, whereas corruption in India is also an outlier, with a difference of 6.3 percentage points. On the other hand, the difference in warmth between leaving the URL available and delisting it is often insignificant when there is a court order.

In addition, country effects can be seen. India has the highest baseline of warmth for search engines (see Table 5), and tends to see lesser drops in support than other countries for delisting content, while Canada often penalizes search engines especially heavily for delisting content. This is especially true when the outlet in question is a reputable newspaper, for instance. As an example, the difference in warmth in India between delisting the URL and leaving it accessible is 0.2 percentage points (insignificant at 95% confidence) for a corruption case under a court order, while the corresponding difference for Canada is 9.9 percentage points.

These contextual characteristics are especially important for the practice of content moderation. Knowing the particulars of a demand (e.g., the trustworthiness of a news outlet, the type of request being made, or the seriousness of the misconduct in question) makes a difference when it comes to understanding what users' expectations will be in a given country. These characteristics can change from demand to demand and from country to country.

#### 4.4 RTBF Knowledge, Opinions, and Effects

Lastly, this study sought to understand respondents' familiarity with and opinions about RTBF, and how this affected their attitudes about public interest content. We found relatively limited knowledge of RTBF, even among respondents living in countries with active RTBF laws. While it was not unexpected that majorities of respondents in Canada and the United States—countries without the RTBF—would not have heard of this term, it was more noteworthy that more than one-fifth of respondents had never heard of the term in France and Poland, two countries that do have the RTBF. This may indicate a limited awareness of which types of information are accessible through search engines. More explicitly, respondents who depend on search engines as their main way of accessing information may be unaware that the information presented to them is influenced or limited by delisting demands.

We also found that respondents in the two countries with active RTBF laws (France and Poland) were more likely to approve of the RTBF. India, a country that has discussed the implementation of such a regime, serves as an intermediary point between the EU countries and the United States and Canada, neither of which have the RTBF at a federal level. Large minorities in every country of our sample were unable to form an opinion on the RTBF, with more than a quarter of Canadian respondents saying they were unsure.

Lastly, respondents' opinions on the RTBF seem to be correlated with their expectations of search engines in edge cases for delisting demands. While 35% of respondents who opposed the RTBF preferred that the search engine attend to the delisting demand in our vignette, a full 45% of respondents who supported the RTBF wanted this outcome. It is still notable, however, that majorities of both groups of respondents preferred that the URL in the vignette remain accessible and available through the search engine. This illustrates that the appropriate evaluation of public interest remains critical for search engines to serve user needs.

#### 4.5 Future Research

Our findings, while useful in and of themselves, also point to new opportunities for research in this nascent area. These avenues for research would provide useful insights into the generalizability and scope of our results.

First, while we surveyed a broad, heterogeneous sample of respondents, we selected five countries. Expanding this survey to other countries in other regions would be useful for stress-testing our finding that respondents prefer that public-interest content be left accessible in search engine results in most cases. Our results could, of course, be simply a function of the five countries we chose. In addition, replicating this study in other countries could prove useful in further examining the particularities of the countries included in the sample. It could be inherently useful for trust and safety

practitioners to understand the preferences of their users on when and where to perform delistings.

Second, our demographic findings were fascinating and relatively unexpected. While we imagined that different demographic groups might have different preferences regarding delisting, we were surprised to see that our youngest cohort was consistently skeptical of search engines, even when they left content accessible. While both sexes preferred that search engines left content accessible, in delisting cases, warmth suffered less for female respondents as compared to their male counterparts. We do not have any particularly convincing explanations for this finding, and would be intrigued to know more about demographic differences in attitudes toward content moderation in general.

Third, we simplified the possible sources of delisting demands to the two that are most country-agnostic: individual demands and court orders. In practice, there are other types of demands as well: government agencies can make demands, and Data Protection Authorities can mediate on behalf of individual data subjects. Repeating this experiment with other sources of delisting demands could help apply these findings to other circumstances.

Fourth, we selected a particular set of vignettes for our experiment based on the contentious edge cases that had been seen by practitioners at a major search engine company. However, reactions might vary for other scenarios. It could be interesting to gauge user reactions to new and different edge cases to see the extent to which our findings are specific to the vignettes we used.

#### **4.6 Conclusion**

In short, these results begin to quantify the reputational cost of over-removing in response to delisting requests, as users tend to prefer freedom of information when considering public interest edge cases, regardless of a court order mandating delisting. If users implicitly expect that search engines in their countries display accurate and comprehensive representations of all topics of public interest (a definite possibility, especially considering respondents' limited knowledge of RTBF), then they may feel disappointed, or even betrayed, when they learn that relevant public information was not made accessible to them. This study shows how delicate it can be to balance the competing concerns of individuals as both internet users and data subjects (or even requesters for delisting) in more complex, realistic situations. It also speaks to the value of verifying every delisting demand, despite the operational costs associated with doing so.

Many respondents, particularly those living within the jurisdiction of RTBF laws, agreed with the concept of RTBF in principle; however, once faced with scenarios wherein cases with a public interest bent were requested for removal under the law, the majority of individuals across all surveyed countries preferred that a search engine leave the content

accessible (even in the presence of a court order). This may also reflect a perception that delisting from search engines in effect surpasses the “right to be forgotten” and moves toward the “right to oblivion” (Ambrose and Ausloos 2013), given the public’s heavy reliance on search engines for discovery (Zhang 2008).

These results suggest a possible gap between the intention behind laws and regulations and their effects once implemented. Teams that deal with high-level delisting demands from governments and courts may be conducting manual reviews of thousands of requests a day (Google 2024c) to decide whether a specific request contains information that is strongly in the public interest (and therefore should remain accessible through search engines). Given the importance of country-specific context, as shown in this study, manual review of high-level edge cases is an essential part of the content moderation process.

Indeed, there is a concerning possibility that search engines may err on the side of delisting in an effort to be compliant with the law, particularly as more countries pass laws establishing deadlines to respond to complainants and harsher penalties for noncompliance. In practice, the Right to be Forgotten would supersede freedom to information considerations. As legislation and regulation mandating content delisting spread into additional jurisdictions, search engine companies will increasingly find themselves in uncomfortable positions when the content in question has some interest to the public.

## References

- Ambrose, Meg Leta, and Jef Ausloos. 2013. "The Right to Be Forgotten Across the Pond." *Journal of Information Policy* 3:1–23. <https://doi.org/10.5325/jinfopoli.3.2013.0001>.
- Arakali, Harichandan. 2025. "The Personal Data Protection Bill Could Be a Serious Threat to Indians' Privacy." *Forbes India*, February 9, 2025. <https://www.forbesindia.com/article/leaderboard/the-personal-data-protection-bill-could-be-a-serious-threat-to-indians-privacy/56623/1>.
- Ardia, David S. 2010. "Free Speech Savior or Shield for Scoundrels: An Empirical Study of Intermediary Immunity Under Section 230 of the Communications Decency Act." *Loyola of Los Angeles Law Review* 43 (2): 373–506. <https://digitalcommons.lmu.edu/llr/vol43/iss2/1>.
- Article 29 Data Protection Working Party. 2024. "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." European Commission, November 17, 2024. [https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp225\\_en.pdf](https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp225_en.pdf).
- Auxier, Brooke. 2020. "Most Americans Support Right to Have Some Personal Info Removed from Online Searches." Pew Research Center, January 27, 2020. <https://www.pewresearch.org/short-reads/2020/01/27/most-americans-support-right-to-have-some-personal-info-removed-from-online-searches/>.
- Bertram, Theo, Elie Bursztein, Stephanie Caro, Hubert Chao, Rutledge Chin Feman, Peter Fleischer, Albin Gustafsson, Jess Hemerly, Chris Hibbert, Luca Invernizzi, et al. 2019. "Five Years of the Right to Be Forgotten." In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 959–72. Association for Computing Machinery, November 9, 2019. <https://doi.org/10.1145/3319535.3354208>.
- Bode, Leticia, and Meg Leta Jones. 2017. "Ready to Forget: American Attitudes Toward the Right to Be Forgotten." *The Information Society* 33, no. 2 (March 8, 2017): 76–85. <https://doi.org/10.1080/01972243.2016.1271071>.
- Brannon, Valerie C., and Eric N. Holmes. 2024. "Section 230: An Overview." Congressional Research Service, January 4, 2024. <https://crsreports.congress.gov/product/pdf/R/R46751>.
- Dahl, David B. 2016. *xtable: Export Tables to LaTeX or HTML. R package version 1.8-2*. <http://xtable.r-forge-project.org/>.
- Economist Intelligence Unit. 2022. *Democracy Index 2022. Frontline Democracy and the Battle for Ukraine*. [https://www.eiu.com/n/wp-content/uploads/2023/02/Democracy-Index-2022\\_FV2.pdf](https://www.eiu.com/n/wp-content/uploads/2023/02/Democracy-Index-2022_FV2.pdf).

- Erdos, David. 2021. "The 'Right to Be Forgotten' Beyond the EU: An Analysis of Wider G20 Regulatory Action and Potential Next Steps." *Journal of Media Law* 13, no. 1 (February 20, 2021): 1–35. <https://doi.org/10.1080/17577632.2021.1884947>.
- European Parliament and Council of the European Union. 2016. "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)." OJ L 119, 4.5.2016, 1–88, May 4, 2016. <https://data.europa.eu/eli/reg/2016/679/oj>.
- GDPR.eu. 2018. "Art. 17 GDPR: Right to Erasure ('Right to Be Forgotten')." November 14, 2018. Accessed November 17, 2024. <https://gdpr-info.eu/art-17-gdpr/>.
- Google. 2023a. "Personal Data Removal Request Form," April 10, 2023. [https://reportcontent.google.com/forms/rtbf?visit\\_id=638167594089677144-941706239&rd=1](https://reportcontent.google.com/forms/rtbf?visit_id=638167594089677144-941706239&rd=1).
- . 2023b. "Remove Images of Minors from Google Search Results," April 10, 2023. <https://support.google.com/websearch/answer/10949130?hl=en>.
- . 2023c. "Remove Your Personal information from Google," April 10, 2023. <https://support.google.com/websearch/troubleshooter/3111061?hl=en>.
- . 2024a. "Content Delistings Due to Copyright," November 11, 2024. <https://transparencyreport.google.com/copyright/overview?hl=en>.
- . 2024b. "Government Requests to Remove Content," November 11, 2024. [https://transparencyreport.google.com/government-removals/government-requests?hl=en&lu=country\\_request\\_amount&country\\_item\\_amount=group\\_by:reasons&country\\_request\\_amount=group\\_by:reasons](https://transparencyreport.google.com/government-removals/government-requests?hl=en&lu=country_request_amount&country_item_amount=group_by:reasons&country_request_amount=group_by:reasons).
- . 2024c. "Requests to Delist Content Under European Privacy Law," November 11, 2024. [https://transparencyreport.google.com/eu-privacy/overview?hl=en&delisted\\_urls=start:1401321600000;end:1731455999999;country:&lu=delisted\\_urls](https://transparencyreport.google.com/eu-privacy/overview?hl=en&delisted_urls=start:1401321600000;end:1731455999999;country:&lu=delisted_urls).
- . 2025a. "European Privacy Requests Search Removals FAQs," February 3, 2025. <https://support.google.com/transparencyreport/answer/7347822/#zippy=%2Cwhat-are-some-common-scenarios-for-delisting-pages%2Cwhat-are-some-common-scenarios-where-you-do-not-delist-pages%2Ccan-you-provide-more-detailed-statistics-about-the-nature-of-these-requests-and-removals>.
- . 2025b. "Personal Content and Product Policies, and Removal Requirements," February 13, 2025. <https://support.google.com/websearch/troubleshooter/3111061>.
- Ingram, Mathew. 2023. "The EU's Right to Be Forgotten Is Migrating to Other Countries." *Columbia Journalism Review* (October). [https://www.cjr.org/the\\_media\\_today/the-eus-right-to-be-forgotten-is-migrating-to-other-countries.php](https://www.cjr.org/the_media_today/the-eus-right-to-be-forgotten-is-migrating-to-other-countries.php).

- Keller, Daphne. 2018. "The Right Tools: Europe's Intermediary Liability Laws and the EU 2016 General Data Protection Regulation." *Berkeley Technology Law Journal* 33 (October 14, 2018): 287–364. <https://doi.org/10.15779/Z38639K53J>.
- Kwon, Yujin, Ella Corren, Gonzalo Munilla Garrido, Chris Hoofnagle, and Dawn Song. 2023. *SoK: The Gap Between Data Rights Ideals and Reality*, December 3, 2023. <https://doi.org/10.48550/arXiv.2312.01511>. arXiv: 2312.01511 [cs.CY].
- Lenth, Russell, et al. 2022. *emmeans: Estimated Marginal Means, aka Least-Squares Means. R package version 1.7. 2*. <https://CRAN.R-project.org/package=emmeans>.
- Luckerson, Victor. 2014. "Americans Will Never Have the Right to Be Forgotten." *Time* (May 14, 2014). <https://time.com/98554/right-to-be-forgotten/>.
- Microsoft. 2024a. "Copyright Content Removal Requests," November 17, 2024. <https://www.microsoft.com/en-us/corporate-responsibility/copyright-removal-requests-report>.
- . 2024b. "Digital Safety Content Report," November 17, 2024. <https://www.microsoft.com/en-us/corporate-responsibility/digital-safety-content-report>.
- . 2024c. "Government Requests for Content Removal Report," November 17, 2024. <https://www.microsoft.com/en-us/corporate-responsibility/government-content-removal-requests-report>.
- . 2024d. "How Bing Delivers Search Results," November 17, 2024. <https://support.microsoft.com/en-us/topic/how-bing-delivers-search-results-d18fc815-ac37-4723-bc67-9229ce3eb6a3>.
- . 2025a. "Report a Concern," February 13, 2025. <https://www.microsoft.com/en-us/concern/bing>.
- . 2025b. "Request to Block Bing Search Results in Europe," February 3, 2025. <https://www.bing.com/webmaster/tools/eu-privacy-request>.
- Nunziato, Dawn Carla. 2017. "The Fourth Year of Forgetting: The Troubling Expansion of the Right to Be Forgotten." *University of Pennsylvania Journal of International Law* 39:1011. <https://doi.org/10.2139/ssrn.3191068>.
- Nurullaev, Ruslan. 2015. *The Right to Be Forgotten in the European Union and Russia: Comparison and Criticism*. Research paper WP BRP 54/LAW/2015. Higher School of Economics Research, October 7, 2015. <https://doi.org/10.2139/ssrn.2669344>.
- R Core Team. 2024. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Rughinis, Razvan, Cosima Rughinis, Simona Nicoleta Vulpe, and Daniel Rosner. 2021. "From Social Netizens to Data Citizens: Variations of GDPR Awareness in 28 European Countries." *Computer Law & Security Review* 42 (August 6, 2021). <https://doi.org/10.1016/j.clsr.2021.105585>.

- Times of India. 2023. "Digital Personal Protection Bill 2023 Passed in Rajya Sabha: Key Points," August 11, 2023. <https://timesofindia.indiatimes.com/gadgets-news/digital-personal-data-protection-bill-2023-passed-in-rajya-sabha-key-points/articleshow/102579315.cms>.
- Van Calster, Geert, Alejandro Gonzalez Arreaza, and Elsemie Apers. 2018. "Not Just One, but Many 'Rights to Be Forgotten'." *Internet Policy Review* 7 (2 2018). <https://doi.org/10.14763/2018.2.794>.
- Wickham, Hadley. 2016. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN: 978-3-319-24277-4. <https://ggplot2.tidyverse.org>.
- Wolford, Ben. 2018. "Everything You Need to Know About the 'Right to Be Forgotten'," November 5, 2018. <https://gdpr.eu/right-to-be-forgotten/>.
- Xue, Minhui, Gabriel Magno, Evandro Cunha, Virgilio Almeida, and Keith W Ross. 2016. "The Right to Be Forgotten in the Media: A Data-Driven Study." *Proceedings on Privacy Enhancing Technologies* 2016 (4): 389–402. <https://doi.org/10.1515/popets-2016-0046>.
- Zhang, Yan. 2008. "Undergraduate Students' Mental Models of the Web as an Information Retrieval System." *Journal of the American Society for Information Science and Technology* 59 (13): 2087–98. <https://doi.org/10.1002/asi.20915>.

## Authors

**Ryan Lloyd** and **Rebecca Umbach** are user experience researchers at Google. They can be reached at [ryanlloyd@google.com](mailto:ryanlloyd@google.com) and [rumbach@google.com](mailto:rumbach@google.com).

## Acknowledgements

We would like to thank the colleagues and anonymous reviewers who took the time to give us invaluable suggestions and feedback. We would especially like to thank Dave Price, Ali Lange, Olivia Sklarew, and Tom Slovak for the additional review of this manuscript.

## Data availability statement

Not applicable.

## Funding statement

This study was funded by Google, the employer of both authors.

## Ethical standards

Not applicable.

## Keywords

Right to be forgotten; delisting; public interest; search engines; GDPR; freedom of information; data privacy.

# Appendices

## Appendix A: Vignette Wording

Table 8: Vignette wording for all experimental conditions

Vignette	Outlet	Crime	Type	Action
1. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] attends to the court order and removes the URL in question from its search results.	newspaper	corruption	court order	delist
2. Suppose that a local blog publishes a post claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] attends to the court order and removes the URL in question from its search results.	blog	corruption	court order	delist
3. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] refuses to attend to the court order and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	newspaper	corruption	court order	leave up

Continued on next page

Table 8 – continued from previous page

Vignette	Outlet	Crime	Type	Action
4. Suppose that a local blog publishes a post claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] refuses to attend to the court order and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	blog	corruption	court order	leave up
5. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] attends to the court order and removes the URL in question from its search results.	newspaper	sexual assault	court order	delist
6. Suppose that a local blog publishes a post claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] attends to the court order and removes the URL in question from its search results.	blog	sexual assault	court order	delist
7. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] refuses to attend to the court order and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	newspaper	sexual assault	court order	leave up

Continued on next page

Table 8 – continued from previous page

Vignette	Outlet	Crime	Type	Action
8. Suppose that a local blog publishes a post claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently obtains a court order demanding that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results under the pain of a significant fine. [PREFERRED SEARCH ENGINE] refuses to attend to the court order and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	blog	sexual assault	court order	leave up
9. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] attends to the takedown demand and removes the URL in question from its search results.	newspaper	corruption	request	delist
10. Suppose that a local blog publishes a post claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] attends to the takedown demand and removes the URL in question from its search results.	blog	corruption	request	delist
11. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] refuses to attend to the takedown demand and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	newspaper	corruption	request	leave up

Continued on next page

Table 8 – continued from previous page

Vignette	Outlet	Crime	Type	Action
12. Suppose that a local blog publishes a post claiming that a national politician accepted kickbacks from a pharmaceutical company in order to lobby on their behalf to the Minister of Health. This politician subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] refuses to attend to the takedown demand, and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	blog	corruption	request	leave up
13. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] attends to the takedown demand and removes the URL in question from its search results.	newspaper	sexual assault	request	delist
14. Suppose that a local blog publishes a post claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] attends to the takedown demand and removes the URL in question from its search results.	blog	sexual assault	request	delist
15. Suppose that a well-known national newspaper publishes an in-depth, well-researched investigative report claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this investigative report from [PREFERRED SEARCH ENGINE]'s search results. [PREFERRED SEARCH ENGINE] refuses to attend to the takedown demand and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	newspaper	sexual assault	request	leave up

Continued on next page

Table 8 – continued from previous page

Vignette	Outlet	Crime	Type	Action
16. Suppose that a local blog publishes a post claiming that a religious leader in your city groomed and sexually assaulted a 15-year-old member of his congregation. This religious leader subsequently demands that [PREFERRED SEARCH ENGINE] remove the URL to this post from [PREFERRED SEARCH ENGINE]’s search results. [PREFERRED SEARCH ENGINE] refuses to attend to the takedown demand and leaves the link/content accessible through [PREFERRED SEARCH ENGINE].	blog	sexual assault	request	leave up