

The Role of Search Engines in the Amplification and Suppression of LGBTIQ+ Polarization

| | |
|------------------|---|
| Authors | Rönnback,Ronja; Emmery,Chris; Šafář,Marie; Milde,Filip et al |
| Published in | ACM Transactions on the Web |
| Publication Date | 2025-10-29 |
| Link | https://research.tilburguniversity.edu/en/publications/2dc5ef33-b431-4be1-954f-7d233db8f06d |
| Citation | Rönnback, R, Emmery, C, Šafář, M, Milde, F, Charvát, J & Brighton, H 2025, 'The Role of Search Engines in the Amplification and Suppression of LGBTIQ+ Polarization', ACM Transactions on the Web. |
| Download Date | 2026-04-22 02:41:07 |
| Rights | <p>General rights</p> <p>Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.</p> <ul style="list-style-type: none"> - Users may download and print one copy of any publication from the public portal for the purpose of private study or research. - You may not further distribute the material or use it for any profit-making activity or commercial gain - You may freely distribute the URL identifying the publication in the public portal" <p>Take down policy</p> <p>If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.</p> |

The Role of Search Engines in the Amplification and Suppression of LGBTIQ+ Polarization

Ronja Rönnback¹, Chris Emmery¹, Marie Šafář Postma¹, Filip Milde², Jan Charvát³, and Henry Brighton^{1,*}

¹Tilburg School of Humanities and Digital Sciences, Tilburg University, Tilburg, The Netherlands

²We Are Fair, Prague, Czech Republic

³Department of Political Science, Charles University, Prague, Czech Republic

*Corresponding author: h.j.brighton@tilburguniversity.edu

Abstract

Search engines are used and trusted by hundreds of millions of people every day. However, the algorithms used by search engines to index, filter, and rank web content are inherently biased, and will necessarily prefer some views and opinions at the expense of others. In this article, we examine how these algorithmic biases amplify and suppress polarizing content. Polarization refers to a shift toward and the acceptance of ideological extremes. In Europe, polarizing content in relation to LGBTIQ+ issues has been a feature of various ideological and political conflicts. Although past research has focused on the role of social media in polarization, the role of search engines in this process is little understood. Here, we report on a large-scale study of 1.5 million search results responding to neutral and negative queries relating to LGBTIQ+ issues. Focusing on the UK, Germany, and France, our analysis shows that the choice of search engine is the key determinant of exposure to polarizing content, followed by the polarity of the query. Location and language, on the other hand, have a comparatively minor effect. Consequently, our findings provide quantitative insights into how differences between search engine technologies, rather than the opinions, language, and location of web users, have the greatest impact on the exposure of web users to polarizing Web content.

1 Introduction

Search engines have the capacity to influence people's preferences, opinions, and decisions on an immense scale. In 2023, for example, Bing reported more than 100 million active daily users despite only achieving a single-digit market share [3]. This implies that Google, which has a market share of over 90% in many parts of the world, may attract over 1 billion daily users [59]. The scale of search engine use, along with findings that show that people often trust search engines more than social and traditional media [14, 15], has fueled the concern that search engine users are susceptible to influence and manipulation [17, 2]. For example, subtle changes to the ranking of politicized search results can shift the voting preferences of undecided voters from one candidate to another [18]. Given this capacity to influence web

users at scale, it is crucial to understand the ways in which search engines are biased toward certain views and opinions at the expense of others, and whether these biases undermine the integrity of the Web [30, 70].

A key barrier to studying search engines is that, like many online platforms, their algorithms are hidden from public view [53]. Moreover, countermeasures designed to block automated interactions make it technically challenging to conduct large-scale scientific analyses. While a number of large-scale monitoring studies have sidestepped this problem, they have tended to focus on the political bias of English language search results, considered only users located in the US, and restricted their attention to Google and occasionally Bing [56, 55, 45, 54]. Building on this work, we set out to examine the role played by four independent search engines in the propagation and suppression of polarizing content in Europe. Polarization, which refers to a shift toward and the acceptance of ideological extremes [10], is seen as a threat to the online information space by key European institutions [20, 19]. Polarizing content typically lacks credibility, tends to be misinformative, is frequently politicized, and often targets minority groups such as LGBTIQ+ people, which includes people who are lesbian, gay, bisexual, transgender, intersex, and queer [64, 27, 34, 16, 26].

In this article, we examine to what extent search engines return low credibility, politicized results likely to fuel polarizing, extremist views toward LGBTIQ+ people. Taking a data-driven approach, we monitored four independent search engines (Google, Bing, Yandex, and Mojeek) over a 3-month period and collected over 1.5 million search results. By controlling the location from which queries were issued, the language they were issued in, and whether the queries conveyed neutral or negative perspectives on LGBTIQ+ issues, we show that the choice of search engine has the greatest impact on exposure to low credibility, polarizing content. Yandex and Mojeek, for example, return over eight times as many low credibility results as Google and Bing. Our findings provide quantitative insights into how differences between search engine technologies, rather than the views, opinions, language, and location of web users, have the greatest impact on the exposure of web users to polarizing Web content.

2 Polarization, LGBTIQ+ rights, and search engine monitoring

Recent years have seen an increase in the polarization of traditional and social media [37, 4, 6]. This shift toward ideological extremes has negative consequences, such as threatening the functioning of democratic institutions [13], and undermining trust and adherence to the rule of law [51]. Although a range of social and political factors drive the process of polarization, the role of technologies such as social media platforms are increasingly under scrutiny [7, 71]. The questions of how and to what extent online platforms play a role in polarization have been posed in relation to several countries [5], but beyond the observation that different countries appear to be affected to different degrees, much remains uncertain [37, 20, 4, 66]. In this article, we address the role of search engines in the amplification and suppression of polarizing content relating to LGBTIQ+ people. LGBTIQ+ rights are of particular interest due to a growing body of research identifying a connection between anti-gender ideology (which seeks to oppose women’s rights as well as LGBTIQ+ rights [64, 27, 49]) and democratic backsliding, where democratic processes are undermined in favor of autocratic

rule [21, 41, 26].

Because the character of online anti-LGBTIQ+ discourse varies depending on the language and the country where it occurs, we developed a multilingual dataset of search engine queries that we use to monitor search engines in different countries. For example, in Germany, polarizing discourse on Twitter/X often invokes the idea of vaccine-induced homosexuality [40]. In contrast, in France, it is more common to use LGBTIQ+ slurs to insult political figures and athletes, and invoke disinformation narratives about homosexuality and disease [40]. Focusing on these cultural differences, our investigation of the amplification and suppression of polarizing content addresses a perceived threat to the online environment recently highlighted by key European institutions [20, 19, 64]. In contrast to our data-driven focus on search engines, previous work has tended to explore the issue of LGBTIQ+ polarization from a social, legal, and security perspective. This work often relies on case studies and aims to provide a detailed examination of specific regions [49, 23, 65, 1, 16], general narratives and strategies [64, 31, 16, 8, 9], and potential geopolitical motivations for the polarization of LGBTIQ+ issues [49, 50, 32, 31, 16, 9]. Our approach complements this body of work, and illustrates how its limitations can be addressed by large-scale algorithmic monitoring [28].

2.1 Search engines and search engine monitoring

Google has a search engine market share of over 90% in many parts of the world [63]. The main exceptions are Baidu which has a 46% share in China [60], and Yandex which has a 75% share in Russia [62]. In Europe, the main competitors to Google are Bing and Yandex, which both have a market share of roughly 4% [61]. Few independent alternatives exist although a number of search engine companies provide privacy-preserving interfaces to Google and Bing [12]. A notable exception is Mojeek, which is an independent search engine based in the UK that prioritizes web content from less well-known and less authoritative sources [46]. Because one of our goals is to investigate variation among search engine technologies, we will focus on four independent search engines that serve Europe: Google, Bing, Yandex, and Mojeek. Our approach to search engine monitoring has four key features:

1. *Active Monitoring.* Several large-scale studies have used technologies such as browser plug-ins to intercept search results delivered to human users [56, 55, 54]. The advantage of these approaches is that the search results being collected are responses to the queries of real users engaging with search engines. The disadvantage is that by passively monitoring search engines through the lens of human activity, there is no way of controlling which queries are issued, when, and from which location. In this study, we perform active monitoring by automating search engine interactions directly, and without any human activity driving the monitoring process. This allows us to design our monitoring study to address specific, experimentally controlled questions.
2. *Multilingual Query Dataset.* Although several studies have monitored search engines by issuing a controlled series of queries [45, 48], they have tended to be small-scale studies or focus on a narrow range of queries [68, 44, 39, 38, 43, 67]. In this study, we used a large, multilingual, and structured dataset of 3766 queries covering a range of LGBTIQ+ issues. This dataset was compiled in collaboration with experts in the field of LGBTIQ+ rights (see Section 3.1), and for each issue, we include queries conveying

neutral and negative perspectives on LGBTIQ+ people. The entire batch of queries was issued independently from 32 regional monitoring hubs spanning eight countries. For each hub, all queries were issued using ethically sourced residential proxies with a registered location within the specified region (for example, to monitor from London we used residential proxies registered in London).

3. *Polarizing Content and Disinformation.* Recent search engine monitoring studies have tended to focus on the political bias of search results, particularly during critical periods of the US election cycle [56, 55, 45, 54, 43, 39]. One barrier to addressing the issue of polarization is that there is no formal definition of the term. On the other hand, prior work has considered how politically partisan views are a key feature of polarizing content [43, 22], as well as the reliability and credibility of the source of the content [71]. We will integrate these two features of polarizing content by combining measures of political bias and content credibility. Consequently, we focus not on user polarization, but on potentially polarizing content that web users are exposed to [5]. Underlying our approach is the idea that different groups of web users can be exposed to different kinds of content, and this may exacerbate ideological disagreement and hostility [52]. To determine whether search results reference potentially polarizing web content, we use the NewsGuard dataset. NewsGuard is a commercial dataset that details the political orientation and journalistic credibility of approximately 12,000 web domains, as well as providing flags that indicate cases of disinformation ([47], see Section 3.3).
4. *Analysis of Unpersonalized Organic Search Results.* When logged into a user account, Google, Bing, and Yandex deliver personalized search results (see, for example, [25]). Mojeek, in contrast, does not collect user information and only delivers unpersonalized search results. In order to compare different search engines on an equal footing, we will sidestep the complexities of comparing different personalization algorithms by focusing on organic, unpersonalized search results. Similarly, in order to make direct comparisons, we will not consider page components such as news, videos, and other potentially search engine specific aspects of the search results. As such, our approach examines the basic ability of search engines to supply ranked search results. Although personalization is a key feature of the major search engines, several studies report that the differences between generic results and personalized search results tend to quite be small [29, 33, 56].

In summary, our approach to search engine monitoring is specifically tailored to the European context, and it allows us to study the relative influence of location, language, query composition, and search engine technology on the amplification and suppression of polarizing LGBTIQ+ content.

3 Data collection through active monitoring

Between the 19th of April and the 11th of June 2024 we monitored Google, Bing, Yandex, and Mojeek in The United Kingdom, France, Germany, the Netherlands, Poland, Slovakia, Czech Republic, and Hungary. Within each country, the entire batch of queries for that country was issued independently, and from each of four separate regional hubs. This process

was repeated three times, resulting in the collection of more than 1.5 million search results. We then analyzed these search results using the NewsGuard dataset, described below. Of the 8 countries we monitored, however, NewsGuard only provides ratings for English, German, and French language web domains. For this reason, in the following analysis we focus on results collected in four locations in the UK (London, Cardiff, Edinburgh, and Belfast), Germany (Berlin, Dresden, Frankfurt, and Munich), and France (Paris, Marseilles, Bordeaux, and Strasbourg). This subset of our results includes 561,317 search results (120,906 collected in France, 220,084 in Germany, and 220,327 in the UK). To further clarify this process, we first detail the query dataset, how search engine interactions were automated, and which aspects of the NewsGuard dataset we used in our analysis.

3.1 A multilingual query dataset

To systematically study the extent to which search engines expose users to LGBTIQ+ related web content, we created a structured, multilingual dataset of queries relating to LGBTIQ+ issues. For each language, we employed a local specialist on LGBTIQ+ rights to guide the construction of queries relating to 7 categories: general LGBTIQ+ issues, LGBTIQ+ disinformation narratives, gender issues, LGBTIQ+ relationships, LGBTIQ+-related personalities, and known anti-LGBTIQ+ extremists (for examples, see Table 1; Table 5 in the Appendix provides a full description of each category and additional examples in English). All queries convey either a neutral perspective (e.g., “gay rights”), or a negative perspective (“gay agenda”), which we refer to as the polarity of the query. An approximately equal number of neutral and negative queries are included in each category. A key feature of the query dataset is that it reflects significant cultural differences in the issues, terminology, and personalities that feature in discussions and conflicts surrounding LGBTIQ+ issues [40, 31, 16]. As a result, it is rarely the case that a query in one language has a direct translation and has the same meaning and relevance in another language. For example, queries about LGBTIQ+-related personalities differ significantly from one language to another because the regional experts responsible for compiling the queries worked independently from each other, and often selected individuals who were only nationally relevant. In contrast, the category of anti-LGBTIQ+ extremists contains only minimal differences between languages because, for each language, the queries reference the same globally relevant set of individuals identified by the Southern Poverty Law Center [58]. In total, our analysis centers on 1,297 queries (288 in French, 528 in German, and 502 in English).

When presenting our results, we will draw attention to a subtlety in the interpretation of query polarity. For the 5 query categories that cover LGBTIQ+ topics, the query polarity refers to whether the query conveys a neutral or negative perspective on LGBTIQ+ people. In contrast, for the 2 categories that cover queries about individuals (LGBTIQ+-related personalities and anti-LGBTIQ+ extremists), a negative polarity refers to a negative perspective on the individual in question, rather than a negative perspective on LGBTIQ+ people (e.g., “Chaya Raichik is homophobic”). Similarly, a query conveying a neutral perspective on an individual means that the query is neutral with respect to that person, even though their views may be associated with anti-LGBTIQ+ sentiment (e.g. “Chaya Raichik is right”). During our analysis, we will separate the analysis of queries about people and queries about topics to avoid any confusion.

Table 1: Eight example queries, their associated query category, their polarity, and language. All queries were issued in lower case.

| Query | Query Category | Query Polarity | Language |
|--|------------------------|----------------|----------|
| homo diktatur | General LGBTIQ+ issues | Negative | German |
| lgbt bewegung | General LGBTIQ+ issues | Neutral | German |
| gays spread corona virus | LGBTIQ+ disinformation | Negative | English |
| gay people and stds | LGBTIQ+ disinformation | Neutral | English |
| gender nazis | Gender issues | Negative | English |
| gender dysphoria | Gender issues | Neutral | English |
| menace sur la famille par les ho-mosexuels | LGBTIQ+ relationships | Negative | French |
| gpa et lgbt | LGBTIQ+ relationships | Neutral | French |

3.2 Automated search engine interactions

We define a search engine interaction as the process of navigating to the homepage of the search engine, issuing a query, and retrieving the first ten ranked search results. In our study, all search engine interactions were automated, and all available localization options were selected for each search engine during these interactions (see Appendix A). These localization options range from selecting different services with different URLs in the case of Google (e.g., `Google.co.uk` vs `Google.de`), to different settings that restrict results to certain regions and languages in the case of Bing and Mojeek. Queries were always issued in the local language, and all internet traffic between our monitoring systems and the search engines was routed through ethically-sourced residential proxies located in the respective countries and regions. To allow direct comparisons between search engines, we will focus on ranked search results appearing in the main results panel. The structure of information boxes appearing in side panels, as well as separate panels focusing on news, images, and videos, often differs from one search engine to another, so we chose not to consider these additional information sources in our analysis. All search engine interactions were automated using the Chromium browser installed on the Linux operating system, and each interaction was automated using a clean browser session containing no stored cookies and no information relating to previous browser use. To minimize the possibility of personalization, no user accounts were used or associated with past use of the browser.

3.3 Domain-level political orientation and credibility ratings

To quantify the extent to which search results reference low-credibility and potentially polarizing web content, a reliable way of characterizing web content is required. For large-scale studies referencing millions of URLs, reliably determining these characteristics at the level of the webpage is infeasible [57]. Instead, previous studies have inferred document-level characteristics from properties of the web domains on which the documents were hosted [45, 54, 56]. For example, content appearing on `cnn.com` is assumed to be left-leaning, whereas content on `foxnews.com` is assumed to be right-leaning. While such inferences may not be valid for platforms that host varied content (such as YouTube, Facebook, and Twitter/X), for many web domains, inferences about specific pages can accurately


Table 2: The four features of the NewsGuard reliability dataset used in our analysis. For each web domain, we consider the credibility score, political orientation, topics covered, and disinformation flags. In Appendix C.1 we detail the degree to which the search results we have collected are documented by the NewsGuard dataset.

| NewsGuard feature | Description |
|--------------------------|---|
| Credibility score | A score ranging from 0 to 100 that indicates how credible the web domain is. A score of 100 indicates that the website complies with all of NewsGuard’s criteria of journalistic integrity, and is highly credible. A domain with a score below 60 is considered as untrustworthy by NewsGuard. Not all web domains are receive a score. For example, large platforms (e.g., Facebook) and domains associated with satire are not assigned a score. |
| Political orientation | The political orientation of the web domain, which is categorized as being either left-leaning, right-leaning, or having no political bias. |
| Topics covered | A list of the general topics covered by a web domain. For example, topics include political news or commentary, personal finance, health, religion, or conspiracy theories. |
| Disinformation flags | For each web domain, a list of flags that detail examples of misinformation or disinformation. Each flag refers to a category such as health information, the Ukraine conflict, or election integrity. |


be made from knowledge of the journalistic practices associated with the domain. In the following analysis, we use domain characteristics provided by a commercially available dataset developed by NewsGuard Technologies [47]. NewsGuard uses a systematic and transparent methodology to characterize the journalistic practices, credibility, and political orientation of over 10,000 web domains. Specifically, our analysis considers four features that are used to characterize web domains: credibility score, political orientation, associated topics, and disinformation flags (see Table 2). The credibility score ranges from 0 to 100, and is a weighted sum of nine component scores that cover various aspects of credibility and transparency, such as whether the web domain frequently publishes deceptive headlines. The political orientation of a domain is labeled as either left-leaning, right-leaning, or non-political. Associated topics are provided as a list of categories that describes the content of domain, such as personal finance, religion, or, most usefully for our purposes, conspiracy theories. Finally, the disinformation flags associated with a domain is a list of topics that this domain has published disinformation on. For reasons discussed above, all domains associated with platforms are excluded from our analysis. Many, but not all, domains referenced in our search results are covered by NewsGuard. We detail these coverage statistics per search engine and country in Appendix C.1.

4 Results


Recall that our goal is to understand the extent to which search engines return results that reference politically partisan, low-credibility content that is likely to fuel polarization.




(a) Political orientation for queries about topics.



(b) Political orientation for queries about individuals.



(c) Credibility scores for queries about topics.




(d) Credibility scores for queries about individuals.

Figure 1: Search results per country returned by all four search engines combined. Panel (a) and (b) show how the political orientation of these search results are distributed per country, for queries about topics and individuals, respectively. A greater proportion of politicized content is seen in the UK in comparison to France and Germany, and there tends to be more left-leaning content for queries about individuals. Panel (c) and (d) show how the same search results are distributed according to credibility score, for queries about topics and individuals, respectively. These search results are in response to both neutral and negative queries. Political orientations and credibility scores are provided by NewsGuard (see Section 3.3).


Specifically, we will examine the relative contribution of location, query category, query polarity, and the choice of search engine.

4.1 Differences between countries


Focusing first on the role of location, Figures 1a and 1b show how search results are distributed according to political orientation (characterized as either left-leaning, no bias, or right-leaning) in Germany, France, and the UK. These distributions, which are estimated from search results for all four search engines and for both neutral and negative queries, indicate that users in Germany, France, and the UK are exposed to slightly different proportions of politicized content. Those differences vary depending on whether the queries concern topics (Figure 1a) or individuals (Figure 1b). For queries about topics, we note that in France we found more right- than left-leaning content. In all three countries, however, the majority of the domains referenced by the search results lacked a political orientation. For queries about individuals, there is comparatively more left-leaning content. But again, the majority of search results were politically neutral. In the UK, though, this politically neutral majority amounts to only slightly more than half of the search results. Turning to




(a) Political orientation for topics queries.



(b) Political orientation for queries about individuals.



(c) Credibility scores for topics queries.



(d) Credibility scores for queries about individuals.

Figure 2: Focusing on search results for all three countries combined, panels (a) and (b) show how the political orientation of these search results are distributed by search engine, for queries about topics and queries about individuals, respectively. Panels (c) and (d) show how the same search results are distributed by credibility score. Mojeek and Yandex show more content from low credibility sources. These search results include responses to both neutral and negative queries.

the issue of exposure to low-credibility content, Figures 1c and 1d show how search results are distributed according to credibility scores (ranging from 0 to 100). Whether the query concerned a topic or an individual has little effect on the credibility distributions. Here, we find that the vast majority of search results refer to web domains with a credibility score greater than 60, which is considered the credibility threshold by NewsGuard. What this course-grained view on our results indicates is that the majority of search results delivered to users in the UK, France, and Germany are credible and, more often than not, hosted by domains that are not associated with politically partisan content.

4.2 Differences between search engines

In Figure 1, we analyzed the search results of all four search engines combined. In Figure 2 we pay close attention to differences between the search engines, focusing again on the political orientation and credibility of the search results.

4.2.1 Political orientation

Figures 2a and 2b show that the search results returned by Bing had the highest proportion of sources with no political orientation (between 80% and 90%), followed by Google (between

77% and 83%). Both of these search engines return more left-leaning than right-leaning sources, and queries about individuals appear to amplify this tendency. In contrast, Mojeek and Yandex returned more results from politically slanted sources, surpassing Bing and Google in terms of the proportion of left-leaning content. Crucially, though, both Mojeek and Yandex show considerably more right-leaning content than left-leaning content (with the exception being Yandex, when queried about individuals). For both Mojeek and Yandex, nearly twice as many references to right-leaning outlets than left-leaning outlets are returned for queries about topics. What these results tell us is that the choice of search engine has a major impact on the degree of exposure to politicized content. Specifically, Google and Bing are behaviorally similar and tend to be biased to the left, whereas Mojeek and Yandex tend to return more politicized content than both Google and Bing, and tend to be biased to the right.

4.2.2 Credibility

Again, for all three countries combined, Figures 2c and 2d show how the same search results are distributed by credibility scores ranging from 0 to 100. Here we see only minor differences between queries about topics and queries about individuals, with the majority search results referencing web domains with credibility scores between 80-100. Notably, Bing and Google very rarely reference content with a credibility score below NewsGuard’s trustworthiness threshold of 60. On the first page of search results, Bing returns a mean of 0.1 results from sources under this threshold, and Google a mean of 0.3 search results. In contrast, on the first page of search results, Mojeek references a mean of 2.9 search results and Yandex references a mean of 1.9 search results from sources below the credibility threshold. To better understand the prevalence of low-credibility and other suspicious content, we used NewsGuard’s metadata to extract the proportions of search results that reference web content that either: (1) had a credibility score less than 60; (2) was hosted on domains known to promote conspiracy theories; or (3), was hosted on domains that have been flagged for publishing content relating to one or more forms of disinformation (see Table 2).

For each of these three markers of low-credibility content, Figure 3 shows the proportion of search results returned by each search engine. Because these results are very similar for both queries about topics and queries about individuals, we will focus only on results relating to queries about topics. First, we see that less than 0.1% of the results returned by Bing and Google reference web domains associated with conspiracy theories. For Google, only 1.6% of results reference domains associated with disinformation, and for Bing this figure drops to less than 0.1% of results. In contrast, approximately 3.8% of the results returned by Mojeek and 3.6% of the results returned by Yandex’s reference domains associated with conspiracy theories. Focusing on domains associated with disinformation, 9% of the results returned by Yandex have been flagged for disinformation, and for Mojeek this proportion is slightly lower at 7%.

4.2.3 Credibility scores and political orientation

So far, we have viewed the credibility of a web domain and its political orientation as being unrelated. Is it the case, though, that low-credibility content is distributed equally across left leaning, right leaning, and apolitical web domains? For the search results returned by each search engine, Figure 4 shows the distribution of credibility scores for each of these




Figure 3: For each search engine, the proportions of search results associated with one of 3 markers of problematic content: a credibility score below 60, a record of hosting conspiracy-related content, and a record of disinformation-related content. All markers are based on NewsGuard metadata. These results combine results from Germany, France, and the UK, and include search results relating to both neutral and negative topic queries.




Figure 4: For each search engine, the distribution of credibility scores for each political orientation label. These distributions have been normalized to account for differences in the absolute numbers for each orientation label and search engine. Appendix C.3 details these absolute numbers.

respective political orientations. Because the results for queries about individuals and results for queries about topics are very similar, we will focus only on the results for queries about topics. First, notice how the credibility scores associated with right-leaning search results are widely distributed, and feature outlets with both very low and very high credibility scores. In contrast, left-leaning and politically neutral search results are concentrated above the credibility threshold. For Bing and Google, low-credibility content is overwhelmingly from right-leaning domains. In general, though, the majority of right-leaning content is hosted by domains that are above the threshold of credibility (roughly 16% of right-leaning results returned by Google and Bing are from domains with low credibility scores, compared to under 1% of domains for left-leaning and politically neutral search results). Strikingly, Mojeek and Yandex have a strong tendency to display low-credibility right-leaning content, with 72% of right-leaning content referenced by Mojeek being labeled as low-credibility, and 63% of right-leaning content referenced by Yandex being low-credibility. In summary, on the first page of search results, Bing and Google respectively display a mean of 0.1 and 0.3 results that are both low credibility and have a political bias. For Mojeek and Yandex, these frequencies increase to a mean of 2.6 and 1.5 results per page, respectively.

4.3 The impact of query polarity

Recall that the query dataset was structured such that, for each query category, there is an approximately equal number of neutral and negative queries. The issue we consider now is to what extent this property of the queries, which we refer to as their polarity, influences exposure to potentially polarizing web content.

4.3.1 Query polarity and political orientation

For each search engine in each of Germany, France, and the UK, Figure 5 examines the mean political bias of the search results returned for both negative and neutral queries

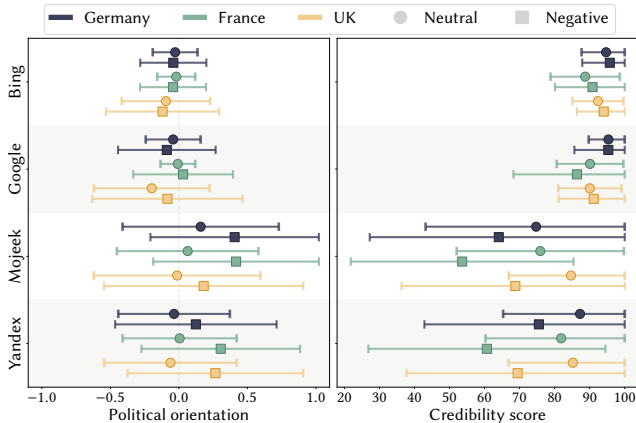


Figure 5: For each country, and for neutral and negative queries, the mean political orientation and credibility scores for queries about topics. The political orientation score assumes that left-leaning domains have a score of -1, right-leaning domains have a score of 1, and domains without a political orientation have a score of 0. Error bars represent standard deviations.




Figure 6: For each search engine, and for both query polarities relating to topics, the mean proportion of search results that reference conspiracy-related and disinformation-related web domains. Error bars represent standard deviations.

about topics (as opposed to individuals). We see that Google and Bing return results which are largely politically balanced, although with slight left-leaning bias for neutral queries, particularly in the UK. For both of these search engines, the polarity of the query had a relatively minor effect, although Google shows a slight shift towards right-leaning results for negative queries. Turning to Yandex and Mojeek, the variation in the political orientation of the search results is notably greater. Both search engines have a slight right-leaning bias that is particularly pronounced for negative queries. In short, the query polarity appears to have a consistently observed impact on the political orientation of the search results for Mojeek and Yandex, but less so for Bing and Google. Turning to queries about individuals, Figure 7 shows, as we would expect, that for queries referencing individuals, query polarity has the opposite effect. All search engines shift towards left-leaning content when the queries are negative, because the negative polarity refers to a negative perspective on the individual, rather than LGBTIQ+ people.

4.3.2 Query polarity and credibility

Turning to the relationship between query polarity and credibility, Figures 5 and 7 show that for queries about both topics and individuals, the query polarity has a relatively minor impact on the behavior of Bing and Google. However, topic queries result in a clear reduction in the mean credibility of the search results returned by Mojeek and Yandex. For queries about individuals, though, the query polarity either has little impact or slightly improves the mean credibility of the results. These differences, however, are minor in comparison to the differences that arise due to the choice of search engine, where there is a clear behavioral distinction between Google and Bing on the one hand, and Mojeek and Yandex on the other. Google and Bing consistently return search results referencing high-credibility domains, while Mojeek and Yandex return results that span high- and low-credibility domains.

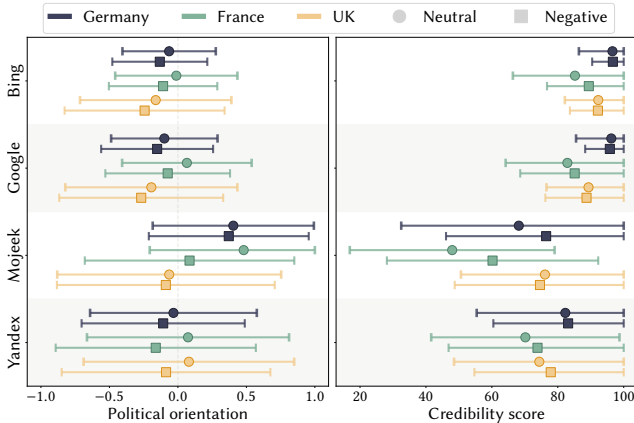


Figure 7: For each country, and for neutral and negative queries, the mean political orientation and credibility scores for queries about individuals. The political orientation score assumes that left-leaning domains have a score of -1, right-leaning domains have a score of 1, and domains without a political orientation have a score of 0. Error bars represent standard deviations.


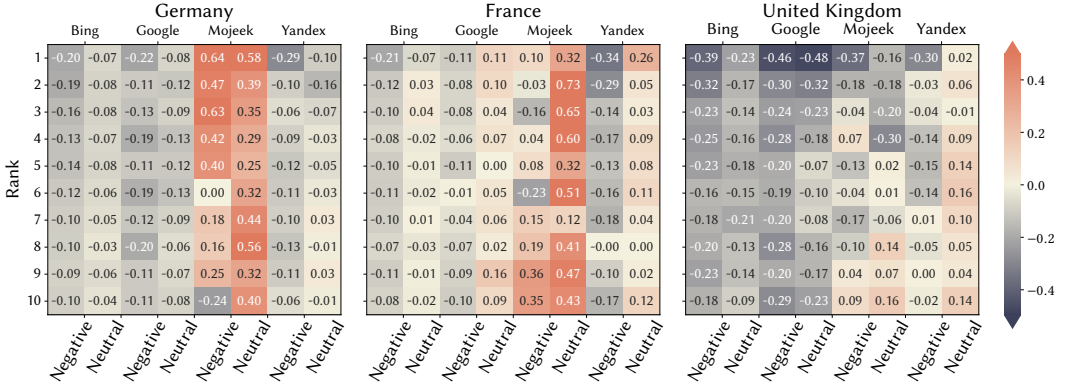


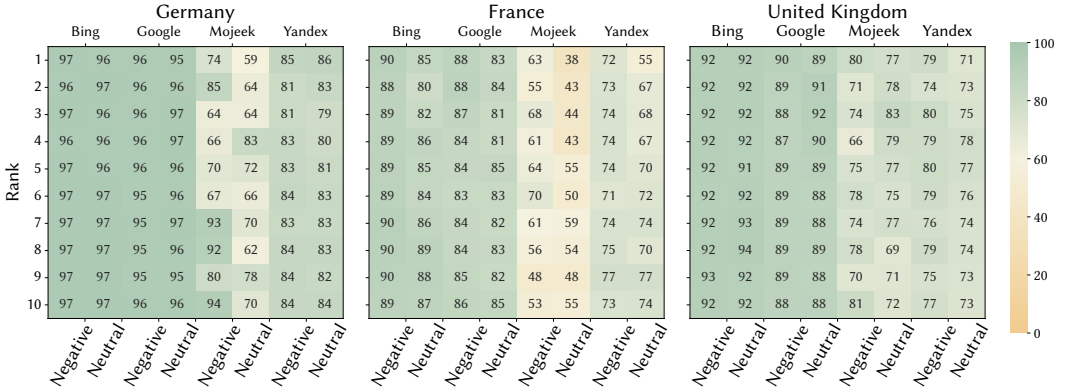
Figure 8: For each search engine, and for both query polarities relating to individuals, the mean proportion of search results that reference conspiracy-related and disinformation-related web domains. Error bars represent standard deviations.

4.3.3 Query polarity and exposure to domains associated with conspiracy theories and disinformation

Figure 5 suggests that Mojeek and Yandex are more reactive to query polarity than Google and Bing, because they return results that differ in both political bias and credibility depending on the query polarity. Much like the analysis of Section 4.2.2, in Figures 6 and 8 we examine the characteristics of low-credibility content arising from negative versus neutral queries in more detail. Specifically, we examine the proportions of search results from domains associated with content flagged by NewsGuard as containing disinformation or relating to conspiracy theories. For Mojeek and Yandex, Figure 6 shows that search results returned for negative queries about topics contain more references to conspiracy-related web domains and domains associated with disinformation (particularly for Yandex). In numbers, negative queries issued to Mojeek increased the proportion of search results referencing domains associated with disinformation from 3% to 9%, while for Yandex the increase was from 2% to 14%. The number of conspiracy-related domains increase from 2% to 5% for Mojeek, and from 1% to 6% for Yandex. A very different picture emerges for Bing and Google. Bing shows less than 0.5% of such content regardless of query polarity. Google shows slightly more disinformation for negative queries (1.4% compared to 0.1% for neutral queries). Once again, we see an opposite effect for queries about individuals in Figure 8. Neutral queries have higher amounts of conspiracy and disinformation content for all search engines (around as much as negative queries about LGBTIQ+ topics), and negative queries result in a drop of such problematic content, although the differences are minor.



(a) Political bias scores



(b) Credibility scores

Figure 9: Examining the role of search result rankings. In (a), for each search engine and for each of Germany, France, and the UK, we consider the mean political orientation per search result rank for queries about individuals. In (b), we conduct the same comparison for credibility scores. As before, the political orientation score assumes that left-leaning domains have a score of -1, right-leaning domains have a score of 1, and domains without a political orientation have a score of 0.

4.3.4 Query polarity and the ranking of search results

The ranking of search results is a key consideration that we have so far glossed over. For instance, if politically biased or low-credibility search results tend to appear toward the bottom of the first page of results, then this is less concerning than if they tend to appear at the top. Indeed, previous research has identified cases where the ranking of search results is consistently biased toward one political orientation at the expense of another [55]. To examine this possibility, we considered the first ten positions on the first page of results for each search engine. We then compared the mean political orientation and the mean credibility score of the results appearing at each of these positions. For queries about LGBTIQ+ topics, we found no systematic and consistent bias in the ranking of search results, for neither political bias scores nor credibility scores, and this absence of a consistent

bias was found for both negative and neutral queries.¹

However, when examining the political orientation of results for queries about individuals, we observe what appears to be some consistent patterns. For each search engine, Figure 9a details the mean political bias score at each of the first 10 search result positions. Figure 9(b) shows the same information for credibility scores. What we observe are slight rank-sensitive biases for: (1) negative queries issued to Bing in all three countries; (2) search results returned by Google in the UK, for both neutral and negative queries; (3) search results returned by Mojeek in Germany; and (4), negative queries returned by Yandex in France. In summary, the impact of rank is generally inconsistent, but can affect the political leaning of results to some degree.

4.4 The impact of query category

The final factor to consider is the impact of query category on exposure to polarizing content. Until now, our analysis of negative and neutral queries includes queries from all seven query categories detailed in Table 5, which we have treated separately depending on whether the queries relate to LGBTIQ+ topics or to anti-LGBTIQ+ individuals. Focusing instead on all seven query categories individually, Figure 10 examines the mean political leaning of search results for each search engine, within each country. As before, we consider the impact of negative and neutral queries separately. For the first four categories (disinformation narratives, gender issues, LGBTIQ+ topics, and LGBTIQ+ relationships), we see a familiar pattern where Google and Bing are largely unreactive to the polarity of the query, while Mojeek and Yandex typically return more right-leaning and more low credibility content for negative queries. Figure 11 conducts the same comparison for the final two query categories about individuals (anti-LGBTIQ+ personalities, and anti-LGBTIQ+ extremists). Here, we find that while both query categories show a left-leaning shift when queries are negative, the effect is larger and more consistent for anti-LGBTIQ+ extremists. On the other hand, the query polarity has little effect on the credibility of the results when considering queries about individuals.

4.5 Queries that expose users to the greatest degree of polarizing content

In Section 2.1 we characterized polarizing content as content that tends to stress ideological extremes, and often does so by deviating from mainstream journalist norms. Indeed, our analysis of polarizing content has considered both whether the content is politically partisan and whether it is of low credibility, as defined by NewsGuard. The final step in our analysis is to identify which queries expose users to the greatest degree of polarizing content by defining a polarization score that measures, for a particular search engine and a particular query, the degree to which the returned search results are of both low credibility and politically slanted. Specifically, for a given search engine, and for a given query q , we first calculate the mean credibility score of the returned search results, denoted c_q (which ranges between 0 and 100). Because the query is likely to have been issued multiple times (e.g., from different locations), this mean is estimated from multiple pages of search results. From the same set

¹These findings are in fact consistent with previous work that found political bias in the ranking of search results, because these biases were only evident when looking beyond the first ten search results [55]. Here, we focus on the first ten search results and also find no rank-sensitive biases.

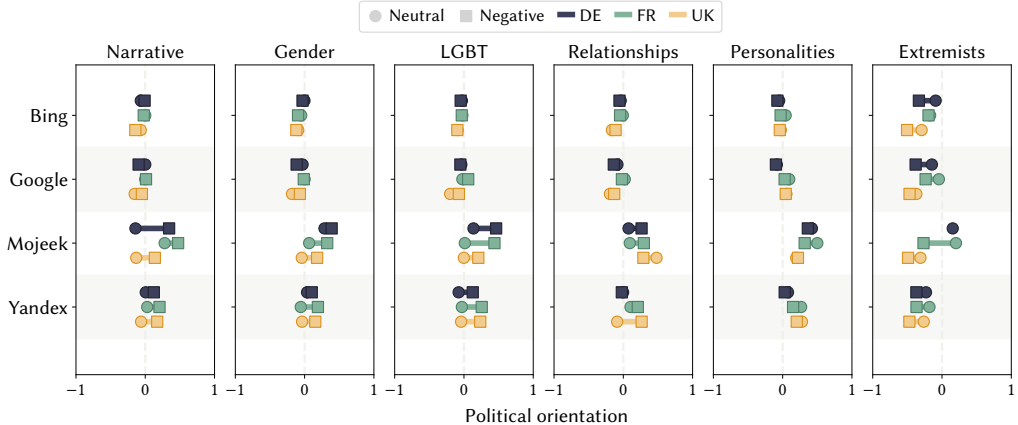


Figure 10: An analysis of all seven query categories. For each of the seven query categories, and for each search engine, we compare the mean political orientation of the search results in Germany, France, and the UK. The political orientation score assumes that left-leaning domains have a score of -1, right-leaning domains have a score of 1, and domains without a political orientation have a score of 0. See Appendix B for descriptions of the categories.

of search results, we then calculate the mean political leaning, denoted p_q (which ranges between -1 and 1, where -1 represents a bias to the political left and 1 represents a bias to the political right). The polarization score $P(q)$ of a query q is then simply the product of the two quantities, after scaling the mean credibility score to range between 0 and 1:

$$P(q) = p_q \cdot \left(1 - \frac{c_q}{100}\right) \quad (1)$$

The polarization score ranges from -1 (polarizing to the political left) to 1 (polarizing to the political right), and approaches zero when the search results tend to be of high credibility and are politically neutral.

For each search engine, Table 3 lists the three queries that yielded polarization scores with the greatest magnitudes, both to the political left (negative scores) and to the political right (positive values). First, notice that the queries with greatest absolute polarization scores (above 0.4) are polarizing to the political right, were issued to Mojeek and Yandex, and relate to topics rather than individuals. Google, on the other hand, yields polarization scores that are less than or equal to zero for all queries. Bing exhibits more varied behavior and yields low polarization scores in both political directions. To summarize, there are cases where all four search engines polarize to the left with low magnitude, but the greatest polarization occurs to the political right for Bing, Yandex, and Mojeek.

5 Discussion

To what extent do search engines expose users to polarizing content? To examine this question, we monitored four independent search engines in eight European countries over a three-month period, and gathered over 1.5 million search results. Using the NewsGuard

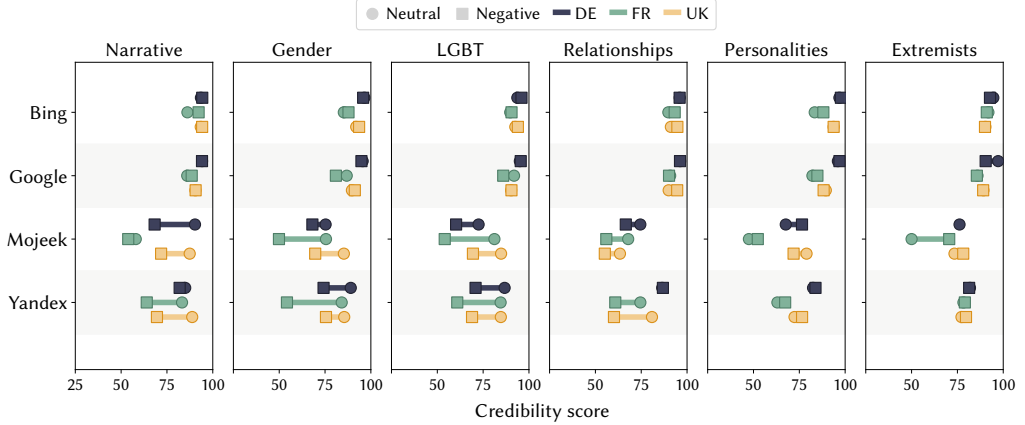


Figure 11: An analysis of all seven query categories. For each of the seven query categories, and for each search engine, we compare the mean credibility scores of the search results in Germany, France, and the UK. See Appendix B for descriptions of the categories.

dataset, our analysis has focused on search engine behavior in the UK, France, and Germany. We have characterized polarizing content as content that, according to NewsGuard, has both low-credibility and is politically slanted. For queries relating to LGBTIQ+ people, we have examined how location, query category, query polarity, and choice of search engine contribute to exposure to potentially polarizing content. Our results show that, although all search engines have the capacity to return results that reference polarizing content, the frequency with which users are likely to encounter this content is determined in large part by their choice of search engine. The location of the user and the composition of the query play a relatively minor role. Before considering the implications and limitations of our findings, we will first outline what they mean, in simple quantitative terms, for search engine users.

5.1 Factors influencing people’s exposure to polarizing content

In numbers, our results (detailed in Section 4.2) indicate that a Google user who switches to Mojeek can expect to encounter over 8 times as many search results that reference content that is both low-credibility and politically partisan. A Bing user switching to Mojeek can expect to see approximately the same increase. In absolute terms, Google returns a mean of less than 1 potentially polarizing search result on the first page of results (3 in every 100 search results), while Mojeek returns a mean of over 2 potentially polarizing results on the first page of results (26 in every 100 search results). When switching to Yandex, Google and Bing users can expect to encounter at least 5 times as many results that reference potentially polarizing content. In short, we have observed a clear behavioral distinction between Bing and Google on the one hand, and Mojeek and Yandex on the other. This quantitative summary relates only to searches about LGBTIQ+ people and topics, and is based on results that have been aggregated over languages, locations, and query polarities. Could it be that these aggregated findings mask additional factors that search engine users should be aware of? While other factors certainly play a role, they are of comparatively minor importance. For example, if a user of Google or Bing switches from issuing only

Table 3: For each search engine, the queries that resulted in the highest polarization scores, as defined by Equation 1. Negative scores indicate polarization to the political left, and positive scores indicate polarization to the political right.

| | Query | Polarization score |
|--------|---|--------------------|
| Bing | matt walsh is an extremist | -0.09 |
| | bryan fischer is homophobic | -0.09 |
| | chaya raichik is homophobic | -0.06 |
| | ... | ... |
| | ludovine de la rochère a raison | 0.04 |
| | ludovine de la rochère protège les valeurs | 0.08 |
| | ludovine de la rochère n'a pas peur d'appeler les choses par leur nom | 0.13 |
| Google | bryan fischer protects values | -0.07 |
| | bryan fischer is an extremist | -0.05 |
| | transgender | -0.05 |
| | ... | ... |
| | cyril hannouna est homophobe | 0.0 |
| | aymeric chauprade est extrémiste | 0.0 |
| | aymeric chauprade a raison | 0.0 |
| Mojeek | chaya raichik is homophobic | -0.09 |
| | trans people were not persecuted by the nazi regime | -0.03 |
| | julia hartley brewer is an extremist | -0.02 |
| | ... | ... |
| | transgender | 0.13 |
| | aymeric chauprade a raison | 0.44 |
| | surgically mutilate trans confused kids | 0.65 |
| Yandex | roy moore is an extremist | -0.06 |
| | roy moore ist transphobisch | -0.06 |
| | andrzej duda ist lgbt feindlich | -0.01 |
| | ... | ... |
| | trans prison rape | 0.12 |
| | andrew tate ist transphobisch | 0.14 |
| | lgbt lobby | 0.42 |

neutral queries to issuing only negative queries in relation to LGBTQI+ people, this shift would result in no discernible difference in their exposure to polarizing content. For Mojeek and Yandex users, though, this same behavioral shift would lead to a consistent but minor difference. Specifically, they should expect see an increase from approximately 13 polarizing results in every 100 first page search results to approximately 16 in 100 (see Section 4.3). The impact of language and location, and differences resulting from the choice of query topic result in similarly minor differences. As we have shown, some of these differences appear to lack a systematic pattern, which suggests that in order to make firmer claims, a larger-scale study focusing in greater detail on factors such as query topic would be needed. Nevertheless, it is clear that the search engine a user chooses to use can influence their

exposure to polarizing content to a far greater extent than what they search for, the language they use, or where they are located.

5.2 Study limitations

Rather than choose to passively capture real-world search engine interactions that involve human users with a history of online behavior, we have opted to automate and actively control all search engine interactions. Because these automated interactions are not linked to any user accounts, our approach cannot shed light on personalized search results. This might be seen as a disadvantage, because for search engines such as Google, real-world users could have been exposed to different search results to those that we have collected and analyzed. The advantage of our approach, though, is that we could systematically control the queries, locations, and most importantly, the choice of search engine. If we took a passive monitoring approach, where queries are spontaneously created by users during course of their everyday online activities, a controlled analysis of query polarity would be infeasible. Similarly, monitoring search engines with a minimal market share, like Mojeek, is unrealistic within a passive monitoring framework. Indeed, our focus on understanding the basic ability of search engines to deliver ranked results means that we can compare search engines on an equal footing: we have factored-out those design features that differ between search engines, such as whether they provide personalized results, and how they adapt the page layout to include, for example, separate sections detailing related news stories and social media posts.

Our analysis rests on some simplifying assumptions. First, we have assumed that the (approximately) 12,000 domains documented by NewsGuard provide a representative and sufficiently large sample of web domains from which to draw our conclusions. Second, we used this dataset to infer document-level features (those relating to a document referenced by a search result) from domain-level features (those relating to a web domain that hosts potentially many documents). These assumptions are to a certain extent unavoidable due to a lack of reliable, web-scale datasets detailing the journalist practices and political biases associated with many hundreds of thousands of web resources [57]. This is why previous search engine monitoring studies have made the same, or very similar, assumptions (e.g., [56, 55, 54]). One implication of these assumptions is that our findings are likely to represent a best-case scenario. Specifically, NewsGuard has a tendency to document domains with significant audience exposure, and includes an approximately balanced distribution of politically left- and right-leaning web domains. This means that while the larger and well-known news outlets are well represented, the long tail of lesser known and more obscure web domains is under-represented. If we then assume that polarizing content is more likely to originate from domains occupying the long tail, then our analysis is likely to have under-estimated the prevalence of polarizing content in search results. Another factor that limits the scope of our conclusions is that we have only studied how search engines respond to queries about LGBTIQ+ topics and people. Extrapolating our conclusions to other ideologically sensitive social issues is an inference that would require additional, and targeted, monitoring studies. Similarly, we have analyzed the issue of polarization, which is a social phenomena that is far more complex and nuanced than our analysis of only two aspects, namely source credibility and political slant [10, 71].

5.3 Explaining search engine behavior

In light of our findings, it is natural to ask why Google and Bing return very few low-credibility, politicized results in comparison to both Yandex and Mojeek. What causes these systematic differences? Because the algorithms used by search engine technologies are hidden from public view, any attempt to answer these questions will necessarily be speculative. Such speculation can nevertheless make use of information in the public domain. Mojeek, for example, is a relatively small UK operation with the expressed goal of differentiating itself from more established search engines by prioritizing web content from less well-known and less authoritative sources [46]. Yandex, on the other hand, is a much larger and well-established operation based in Russia, and has repeatedly been accused of Kremlin influence [35], disregard for privacy concerns [24], censorship [36], and the promotion of misinformation [35]. We have found that these two search engines, Mojeek and Yandex, despite operating under very different circumstances, expose their users to similar degrees of potentially polarizing content. Given our findings, perhaps the least speculative explanation is that the differences between search engines we observe may reflect a limited range of opportunities for offering a substantive alternative to Google and Bing. Google and Bing excel at delivering relevant results from high-credibility sources. To compete with Google and Bing, and offer something that these services lack, the path of least resistance may be to deliver relevant results, but from a greater number of low-credibility sources, and thereby cater to a growing audience that trusts and seeks information from sources outside of the mainstream media [14, 15].

6 Acknowledgements

This research is part of the project “The dynamics of LGBTQ+ misinformation across Europe” which was funded by a grant from the European Media and Information Fund (<https://gulbenkian.pt/emifund/>). All inquiries should be directed to the principal investigator, Henry Brighton (h.j.brighton@tilburguniversity.edu).

References

- [1] AHONEN, A., BILAL, A., GJØRV, G. H., JURĀNS, J., KRAGH, M., KRŪMIŅŠ, K. K., LANGE-IONATAMISHVILI, E., LIUBINAVIČIUS, B., MIKULSKI, K., MÖLDER, H., ÓLAFSSON, J. G., ÓMARSÓTTIR, S. B., SAZONOV, V., AND SERRITZLEV, J. Russian influence operations in nordic baltic, 2024.
- [2] ASLETT, K., SANDERSON, Z., GODEL, W., PERSILY, N., NAGLER, J., AND TUCKER, J. A. Online searches to evaluate misinformation can increase its perceived veracity. *Nature* 625, 7995 (2014), 548–556. Publisher: Nature Publishing Group.
- [3] BING SEARCH BLOG. The new bing and edge – progress from our first month. https://blogs.bing.com/search/march_2023/The-New-Bing-and-Edge----Momentum-from-Our-First-Month, March 2023.
- [4] BOXELL, L., GENTZKOW, M., AND SHAPIRO, J. M. Cross-country trends in affective polarization. *The Review of Economics and Statistics* 106, 2 (2024), 557–565.

- [5] BUDAK, C., NYHAN, B., ROTHSCHILD, D. M., THORSON, E., AND WATTS, D. J. Misunderstanding the harms of online misinformation. *Nature* 630, 8015 (2024), 45–53.
- [6] CHINN, S., HART, S., AND SOROKA, S. Politicization and polarization in climate change news content, 1985–2017. *Science Communication* 42, 1 (2020), 112–129. Publisher: SAGE Publications Inc.
- [7] CHUECA DEL CERRO, C. The power of social networks and social media’s filter bubble in shaping polarisation: an agent-based model. *Applied Network Science* 9, 1 (2024), 69.
- [8] CORRALES, J., AND KIRYK, J. Homophobic populism. *Oxford Research Encyclopedia of Politics* (August 2022), 48. Retrieved June 4, 2025.
- [9] DATTA, N. Tip of the iceberg: Religious extremist funders against human RIghts for sexuality and reproductive health in europe, 2021.
- [10] DIMAGGIO, P., EVANS, J., AND BRYSON, B. Have american’s social attitudes become more polarized? *American Journal of Sociology* 102, 3 (1996), 690–755.
- [11] DOMCOP. What is Open PageRank? <https://www.domcop.com/openpagerank/what-is-openpagerank>, 2025. Retrieved June 4, 2025.
- [12] DUCKDUCKGO. Where do DuckDuckGo search results come from?, 2025. Retrieved November 8, 2024.
- [13] ECKER, U., ROOZENBEEK, J., VAN DER LINDEN, S., TAY, L. Q., COOK, J., ORESKES, N., AND LEWANDOWSKY, S. Misinformation poses a bigger threat to democracy than you might think. *Nature* 630, 8015 (2024), 29–32.
- [14] EDELMAN TRUST INSTITUTE. 2024 edelman trust barometer, 2024.
- [15] EDELMAN TRUST INSTITUTE. 2025 edelman trust barometer, 2025.
- [16] EDENBORG, E. ‘traditional values’ and the narrative of gay rights as modernity: Sexual politics beyond polarization. *Sexualities* 26, 1 (2023), 37–53. Publisher: SAGE Publications Ltd.
- [17] EPSTEIN, R. Manipulating minds: The power of search engines to influence votes and opinions. In *Digital dominance: Implications and risks*, M. Moore and D. Tambini, Eds. Oxford University Press, Oxford, 2018, pp. 293–318.
- [18] EPSTEIN, R., AND ROBERTSON, R. The search engine manipulation effect (seme) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences* 112, 33 (2015), E4512–E4521.
- [19] EUROPEAN UNION AGENCY FOR FUNDAMENTAL RIGHTS. *LGBTQI equality at a crossroads: progress and challenges : EU LGBTQI survey III*. Publications Office of the European Union, Luxembourg, Luxembourg, 2024.
- [20] FLETCHER, R., AND JENKINS, J. Polarisation and the news media in europe. *Publications Office of the EU* (2019), 1–49.

- [21] FLORES, A. R., FUENTES, M., AND SHAW, A. Democratic backsliding and lgbt acceptance, 2023. Retrieved June 4, 2025.
- [22] GARIMELLA, K., SMITH, T., WEISS, R., AND WEST, R. Political polarization in online news consumption. *Proceedings of the International AAAI Conference on Web and Social Media 15* (2021), 152–162.
- [23] GEGUCHADZE, G., AND URUSHADZE, M. The lgbtq issue as a supplementary tool for polarization in post-soviet countries—the case of georgia. In *STRATEGIES XXI: The Complex and Dynamic Nature of the Security Environment* (Bucharest, Romania, 2021), CAROL 1 National Defence University, pp. 84–93.
- [24] GERBET, T., KUMAR, A., AND LAURADOUX, C. A privacy analysis of google and yandex safe browsing. In *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)* (2016), pp. 347–358.
- [25] GOOGLE. Why your Google Search results might differ from other people. <https://support.google.com/websearch/answer/12412910?hl=en>, 2025. Retrieved June 4, 2025.
- [26] GRABOWSKA-MOROZ, B., AND WÓJCIK, A. Reframing LGBT rights advocacy in the context of the rule of law backsliding: The case of poland. *Intersections. East European Journal of Society and Politics* 7, 4 (2021), 85–103. Number: 4.
- [27] GRAFF, A., KAPUR, R., AND WALTERS, S. D. Introduction: Gender and the rise of the global right. *Signs: Journal of Women in Culture and Society* 44, 3 (2019), 541–560. Publisher: The University of Chicago Press.
- [28] HAIM, M. Agent-based testing: An automated approach toward artificial reactions to human behavior. *Journalism Studies* 21, 7 (2020), 895–911.
- [29] HANNAK, A., SAPIEZYNSKI, P., MOLAVI KAKHKI, A., KRISHNAMURTHY, B., LAZER, D., MISLOVE, A., AND WILSON, C. Measuring personalization of web search. In *Proceedings of the 22nd international conference on World Wide Web* (2013), ACM, pp. 527–538.
- [30] INTRONA, L. D., AND NISSENBAUM, H. Shaping the web: Why the politics of search engines matters. *The Information Society* 16, 3 (2000), 169–185.
- [31] JALONEN, O., AND HOOGENSEN GJØRV, G. *Identity as a tool for discrimination: Exploiting social divisions in modern societies*. Hybrid CoE, Helsinki, Finland, 2023.
- [32] KARLSEN, G. H. Divide and rule: ten lessons about russian political influence activities in europe. *Palgrave Communications* 5, 1 (2019), 19.
- [33] KLIMAN-SILVER, C., HANNAK, A., LAZER, D., WILSON, C., AND MISLOVE, A. Location, location, location: The impact of geolocation on web search personalization. In *Proceedings of the 2015 Internet Measurement Conference* (New York, NY, USA, 2015), IMC '15, Association for Computing Machinery, p. 121–127.

- [34] KOVÁTS, E., AND PÕIM, M. Gender as symbolic glue: the position and role of conservative and far right parties in the anti-gender mobilizations in europe. <https://www.genderportal.eu/resources/gender-symbolic-glue-position-and-role-conservative-and-far-right-parties-anti-gender>, 2015.
- [35] KRAVETS, D., BESELER, A., TOEPFL, F., AND RYZHOVA, A. The kremlin-controlled search engine Yandex as a tool of foreign propaganda. *Russian Analytical Digest* 313 (2024), 11–15.
- [36] KRAVETS, D., AND TOEPFL, F. Gauging reference and source bias over time: how russia’s partially state-controlled search engine yandex mediated an anti-regime protest event. *Information, Communication & Society* 25, 15 (2022), 2207–2223.
- [37] KUBIN, E., AND VON SIKORSKI, C. The role of (social) media in political polarization: A systematic review. *Annals of the International Communication Association* 45, 3 (2021), 188–206.
- [38] KUZNETSOVA, E., MAKHORYTKH, M., SYDOROVA, M., URMAN, A., VITULANO, I., AND STOLZE, M. Algorithmically curated lies: How search engines handle misinformation about US biolabs in ukraine, 2024.
- [39] LEÓN, E. D., MAKHORYTKH, M., URMAN, A., AND ULLOA, R. Googling the big lie: Search engines, news media, and the US 2020 election conspiracy, 2024.
- [40] LOCATELLI, D., DAMO, G., AND NOZZA, D. A cross-lingual study of homotransphobia on Twitter. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)* (Dubrovnik, Croatia, May 2023), S. Dev, V. Prabhakaran, D. I. Adelani, D. Hovy, and L. Benotti, Eds., Association for Computational Linguistics, pp. 16–24.
- [41] LOMBARDO, E., AND KANTOLA, J. European integration and disintegration: Feminist perspectives on inequalities and social justice. *JCMS: Journal of Common Market Studies* 57, S1 (2019), 62–76.
- [42] LONG, J. About Ranking on Mojeek. <https://blog.mojeek.com/2024/08/about-ranking-on-mojeek.html>, 2024. [Accessed 12-11-2024].
- [43] MAKHORYTKH, M., RORHBACH, T., SYDOROVA, M., AND KUZNETSOVA, E. Search engines in polarized media environment: Auditing political information curation on google and bing prior to 2024 US elections, 2025.
- [44] MAKHORYTKH, M., URMAN, A., AND WIJERMARS, M. A story of (non)compliance, bias, and conspiracies: How google and yandex represented smart voting during the 2021 parliamentary elections in russia. *Harvard Kennedy School Misinformation Review* 3, 2 (2022), 1–16.
- [45] METAXA, D., PARK, J. S., LANDAY, J. A., AND HANCOCK, J. Search media and elections: A longitudinal investigation of political search results. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (November 2019), 1–17.
- [46] MOJEEK. Independent and Unbiased Search Results. <https://blog.mojeek.com/2018/08/independent-and-unbiased-search-results.html>, 2018. [Accessed 07-11-2024].

- [47] NEWSGUARD. News Reliability Ratings. <https://www.newsguardtech.com/solutions/news-reliability-ratings/>, 2025. Retrieved June 4, 2025.
- [48] NOROCEL, O. C., AND LEWANDOWSKI, D. Google, data voids, and the dynamics of the politics of exclusion. *Big Data & Society* 10, 1 (2023), 20539517221149099.
- [49] NOROCEL, O. C., AND PATERNOTTE, D. The dis/articulation of anti-gender politics in eastern europe: Introduction. *Problems of Post-Communism* 70, 2 (2023), 123–129.
- [50] PATERNOTTE, D., AND KUCHAR, R. Disentangling and locating the “global right”: Anti-gender campaigns in europe. *Politics and Governance* 6, 3 (2018), 6–19.
- [51] PURI, A., AND KEYMOLEN, E. The doors of janus: A critical analysis of the socio-technical forces eroding trust in the rule of law. *Cardozo Arts and Entertainment Law Journal* (2024), 1 – 55.
- [52] REKKER, R. Political polarization over factual beliefs. In *Knowledge resistance in high-choice information environments*. Routledge, London, UK, 2022, pp. 222–236.
- [53] RIEDER, B., AND HOFMANN, J. Towards platform observability. *Internet Policy Review*, 9, 4 (2020), 28.
- [54] ROBERTSON, R., GREEN, J., RUCK, D. J., OGNANOVA, K., WILSON, C., AND LAZER, D. Users choose to engage with more partisan news than they are exposed to on google search. *Nature*, 618 (2023), 342–348.
- [55] ROBERTSON, R., JIANG, S., JOSEPH, K., FRIEDLAND, L., LAZER, D., AND WILSON, C. Auditing partisan audience bias within google search. *Proceedings of the ACM on Human-Computer Interaction* 2 (2018), 1–22.
- [56] ROBERTSON, R. E., LAZER, D., AND WILSON, C. Auditing the personalization and composition of politically-related search engine results pages. In *Proceedings of the 2018 World Wide Web Conference* (Republic and Canton of Geneva, CHE, 2018), WWW ’18, International World Wide Web Conferences Steering Committee, p. 955–965.
- [57] RÖNNBACK, R., EMMERY, C., AND BRIGHTON, H. Automatic large-scale political bias detection of news outlets. *PLoS One* 20, 5 (2025), e0321418.
- [58] SOUTHERN POVERTY LAW CENTER. Extremists, groups, & ideologies, December 2025. Accessed 7 July 2025.
- [59] STATCOUNTER.COM. Search engine market share united states of america: Oct 2023 - nov 2024. <https://gs.statcounter.com/search-engine-market-share/all/united-states-of-america>, 2024.
- [60] STATCOUNTER.COM. Search engine market share china. <https://gs.statcounter.com/search-engine-market-share/all/china>, 2025.
- [61] STATCOUNTER.COM. Search engine market share europe. <https://gs.statcounter.com/search-engine-market-share/all/europe>, 2025.

- [62] STATCOUNTER.COM. Search engine market share russian federation. <https://gs.statcounter.com/search-engine-market-share/all/russian-federation>, 2025.
- [63] STATCOUNTER.COM. Search engine market share worldwide. <https://gs.statcounter.com/search-engine-market-share>, 2025.
- [64] STRAND, C., SVENSSON, J., BLOMEYER, R., AND SANZ, M. *Disinformation campaigns about LGBTI+ people in the EU and foreign influence*. European Parliament, Policy Department for External Relations, Brussels, Belgium, 2021.
- [65] UNAL, D. Political homophobia as a tool of creating crisis narratives and ontological insecurities in illiberal populist contexts: lessons from the 2023 elections in turkey. *New Perspectives on Turkey* 71 (2024), 1–22.
- [66] URMAN, A. Context matters: political polarization on twitter from a comparative perspective. *Media, Culture & Society* 42, 6 (2020), 857–879. Publisher: SAGE Publications Ltd.
- [67] URMAN, A., MAKHORTYKH, M., AND ULLOA, R. The matter of chance: Auditing web search results related to the 2020 u.s. presidential primary elections across six search engines. *Social Science Computer Review* 40, 5 (2022), 1323–1339.
- [68] URMAN, A., MAKHORTYKH, M., ULLOA, R., AND KULSHRESTHA, J. Where the earth is flat and 9/11 is an inside job: A comparative algorithm audit of conspiratorial information in web search results. *Telematics and Informatics* 72 (2022), 101860.
- [69] VARANGOULI, E. Google PageRank in 2024: What Google Search Leak Reveals. <https://www.semrush.com/blog/pagerank/why-pagerank-still-matters>, 2024. [Accessed 12-11-2024].
- [70] WALLACE, J. Modelling contemporary gatekeeping: The rise of individuals, algorithms and platforms in digital news dissemination. *Digital Journalism* 6, 3 (2017), 274–293.
- [71] WEISMUELLER, J., GRUNER, R. L., HARRIGAN, P., COUSSEMENT, K., AND WANG, S. Information sharing and political polarisation on social media: The role of falsehood and partisanship. *Information Systems Journal* 34, 3 (2024), 854–893. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/isj.12453>.

A Appendix: Search Engine Localization Options

Table 4: Localization options used per search engine and country, as well as the four major cities where data was collected.

| | France (Bordeaux, Marseilles, Paris, Strasbourg) | Germany (Berlin, Cologne, Dresden, Munich) | United Kingdom (Belfast, Cardiff, Edinburgh, London) |
|--------|---|--|--|
| Bing | .com/?setlang=fr&cc=fr&cc=FR | .com/?setlang=de&cc=de&cc=DE | .com/?setlang=en&cc=GB |
| Google | .fr | .de | .co.uk |
| Mojeek | .com/?reg=fr&lb=fr&arc=fr | .com/?reg=de&arc=de&lb=de | .com/?reg=uk&lb=en&arc=gb |
| Yandex | .com | .com | .com |

B Appendix: Query Categories

Each query belongs to an overarching category. Table 5 provides a brief description of what each category entails. Additionally, there are varying numbers of unique queries per country, topic category, and query polarity (neutral or negative). The table 6 provides an overview.

Table 5: Overview of the categories of queries, with some representative examples.

| Category | Description | Example |
|-----------------|---|--|
| Narratives | Query prompts search engine about common disinformation narratives. These are predominantly related to LGBT-related disinformation, but also include coronavirus- and vaccine-disinformation. | gays will traffic children, doctors and pharma forcing children into transition |
| Gender | Query relates to traditional and other genders, homo- and transphobia. | gender dysphoria, 76 genders, transgender quotas |
| LGBT | Query concerns topics surrounding LGBTIQ+ people, such as important events, discrimination and general information. | violence against trans people, homo propaganda, lgbt rights, lgbt meaning |
| Relationships | Query relates to family relationships, particularly the relationship between parents and children. | same-sex parenthood, same sex couples can't bring up children correctly, surrogate motherhood and lgbt |
| Personalities | Query is about a specific individual with some anti-LGBTIQ+ stances, particularly whether they are considered as being correct, or problematic. | jk rowling speaks the truth, allison pearson is homophobic, lord moonie protects values |
| Extremist | Query relates to an individual that have been designated by the Southern Poverty Law Centre (SPLC) as extremists. | matt walsh speaks the truth, david lane is homophobic, fred Phelps is a fascist |

Table 6: Overview of number of queries per topic category, country, and query polarity.

| Category | France | | Germany | | United Kingdom | |
|---------------|---------|----------|---------|----------|----------------|----------|
| | Neutral | Negative | Neutral | Negative | Neutral | Negative |
| Narratives | 7 | 8 | 13 | 17 | 9 | 52 |
| Gender | 10 | 10 | 35 | 31 | 29 | 53 |
| LGBT | 43 | 34 | 65 | 71 | 46 | 71 |
| Relationships | 16 | 14 | 24 | 22 | 9 | 25 |
| Personalities | 36 | 36 | 93 | 93 | 66 | 66 |
| Extremist | 39 | 39 | 39 | 39 | 39 | 39 |

C Appendix: Supplementary Data Exploration

C.1 NewsGuard Coverage

Our data focuses on a relatively specific topic that may be under-represented in the NewsGuard data. To examine whether there is sufficient coverage by the NewsGuard metadata, we provide tables of the amounts of search results and unique web-domains with or without labels, both per country (see Table 7) and per search engine (see Table 8).

Table 7: Coverage of NewsGuard per country

| | Search results with metadata | Search results without metadata | Percentage of search results with metadata | Unique domains with metadata | Unique domains without metadata | Percentage of domains with metadata |
|----|------------------------------|---------------------------------|--|------------------------------|---------------------------------|-------------------------------------|
| DE | 82732 | 137352 | 37.6% | 692 | 6299 | 9.9% |
| FR | 43608 | 77298 | 36.1% | 539 | 4422 | 10.9% |
| UK | 88345 | 131982 | 40.1% | 949 | 7094 | 11.8% |

Table 8: Coverage of NewsGuard per search engine.

| | Search results with metadata | Search results without metadata | Percentage of search results with metadata | Unique domains with metadata | Unique domains without metadata | Percentage of domains with metadata |
|--------|------------------------------|---------------------------------|--|------------------------------|---------------------------------|-------------------------------------|
| Bing | 77290 | 75698 | 50.5% | 761 | 3901 | 16.3% |
| Google | 47333 | 88174 | 34.9% | 697 | 4714 | 12.9% |
| Mojeek | 29440 | 107922 | 21.4% | 691 | 7107 | 8.9% |
| Yandex | 60622 | 74838 | 44.8% | 1013 | 5153 | 16.4% |

As is evident in Tables 7 and 8, most search results and domains do not have NewsGuard metadata. This is not surprising, since our queries may result in obscure sites that have

been of lower priority for NewsGuard to evaluate. To validate this, we examined the Open PageRank² of the websites in our dataset with and without metadata. Websites without NewsGuard metadata generally have a lower Open PageRank [11], meaning they are ranked lower by search engines and therefore less likely to be shown (see Fig 12).

There is also a lower coverage of unique domains compared to search results. This indicates that there are a lot of "long-tail" web-domains in our data. Domains with NewsGuard metadata are more popular and therefore appear more frequently, boosting the coverage of search results. Mojeek’s search results (and unique domains) are covered the least by NewsGuard, which reflects the search engine’s commitment to providing more diverse search results, even if the sites are not necessarily deemed authoritative [42]. Surprisingly Google has lower NewsGuard coverage compared to Bing or Yandex. Overall, given the extraordinary number of websites in existence, NewsGuard’s coverage, though not complete, is more than adequate.




Figure 12: Open PageRank scores of search results with and without NewsGuard metadata. An Open PageRank of 10 represents the most popular and authoritative web-domains.




Figure 13: Credibility per political orientation label for full NewsGuard database.

C.2 Distribution of credibility scores per political orientation in NewsGuard

As a point of comparison, we provide an overview of the credibility score distribution for left-, right-, and no political orientations. As can be seen, the credibility scores range the entire spectrum for all orientations, but right-leaning content tends to also score lower in terms of credibility. While it is tempting to relegate the difference in credibility distributions per political leaning to labelling bias, NewsGuard is currently one of the most extensive, detailed, and high-quality databases covering online sources. As such, the database presents one of the better existing approximations of the online sphere.

²Open PageRank is based on Google’s PageRank algorithm, which is no longer publicly visible since 2016. It is apparently still in use as a minor part of their ranking algorithm [11, 69].

C.3 Absolute counts per political orientation label per credibility score range

To provide thorough insight into the dataset, we also provide the counts of the political orientation labels per credibility score range, for each search engine. These results combine all data across all query topic categories, and all countries.

Table 9: Counts per political orientation label and credibility score range for Bing.

| Bing | 0-20 | 20-40 | 40-60 | 60-80 | 80-100 |
|-------------|------|-------|-------|-------|--------|
| Left | 0 | 0 | 11 | 1201 | 8530 |
| No Bias | 17 | 13 | 278 | 2619 | 58619 |
| Right | 192 | 128 | 223 | 387 | 1263 |

Table 10: Counts per political orientation label and credibility score range for Google.

| Google | 0-20 | 20-40 | 40-60 | 60-80 | 80-100 |
|---------------|------|-------|-------|-------|--------|
| Left | 28 | 2 | 139 | 780 | 6223 |
| No Bias | 6 | 60 | 171 | 1729 | 31755 |
| Right | 213 | 145 | 357 | 768 | 1072 |

Table 11: Counts per political orientation label and credibility score range for Mojeek.

| Mojeek | 0-20 | 20-40 | 40-60 | 60-80 | 80-100 |
|---------------|------|-------|-------|-------|--------|
| Left | 91 | 94 | 20 | 1482 | 3562 |
| No Bias | 244 | 279 | 361 | 1273 | 10887 |
| Right | 3519 | 2484 | 1092 | 1326 | 1699 |

Table 12: Counts per political orientation label and credibility score range for Yandex.

| Yandex | 0-20 | 20-40 | 40-60 | 60-80 | 80-100 |
|---------------|------|-------|-------|-------|--------|
| Left | 13 | 46 | 450 | 2100 | 6677 |
| No Bias | 687 | 931 | 353 | 2389 | 22264 |
| Right | 3364 | 1787 | 1302 | 2727 | 1830 |