



Google UK
1 St Giles High Street
London
WC2H 8AG

31 May 2022

Dear Sir Roger, Ms Rees,

Thank you for inviting Google to give evidence to the Online Safety Bill Public Bill Committee this week, and for your subsequent invitation to provide written evidence. Google strongly supports the objectives of this legislation: to improve UK citizens' safety online and to protect their fundamental rights of freedom of expression and privacy.

As a company, we have not waited for new regulation before acting to keep our users safe. We have invested extensively in developing policies and products that protect users, while also developing programmes to equip young people with valuable media literacy skills. The measures we have taken include:

- Developing and updating new [community guidelines](#) on YouTube, that we strictly enforce, on a range of harmful content, including suicide and self harm, scams, elections misinformation, harrassment and hate speech;
- Developing a separate set of [policies](#) on Google Search for our Search features, including auto-complete predictions, to exclude hateful content, harassing content, self-harm, or sexually explicit content from appearing in auto-complete predictions;
- Developing [tools](#) to identify child sexual abuse material (CSAM) that we have shared with the wider digital industry;
- Implementing additional [protections](#) for children and teens on our platforms, including turning SafeSearch on for existing signed-in users under the age of 18 and users we believe to be under 18.
- Establishing longstanding partnerships aimed at strengthening media literacy: this includes our partnership with Parent Zone on [Be Internet Legends](#), the only PHSE-accredited online safety programme for 7-11 year olds in the UK.

We want to see this Bill work effectively and keep our users safe. While we support the Bill's systemic approach and risk-based framework, this remains one of the most complex pieces of digital regulation anywhere in the world. With additional clarity, we believe the Bill can deliver on its intended objectives in a way that is practical to implement:

- 1) Ensuring that the Bill will not lead to widespread automated monitoring of content, which would result in the over-removal of legal content that users should have access to;



- 2) Ensuring that Ofcom's assessment of proportionality protects UK users' privacy and rights to share/access information;
- 3) Considering the workability of some of the newer additions to the Bill, including complaints processes and user empowerment duties.

General monitoring and the over-removal of legal content

The language in the Bill to 'prevent individuals from encountering' illegal content and the practical challenges of distinguishing between illegal and legal content at scale mean that the duties in the framework appear to incentivise automated general monitoring – and over-removal – of content on a service. In our experience, algorithmic tools are effective at identifying CSAM and certain types of violent extremist content, but will struggle to accurately identify other, more context-dependent forms of illegal and harmful content. The current wording of the Bill, and the provisions that allow Ofcom to require companies to proactively monitor their services for illegal content, could force services to rely excessively on automated tools to identify illegal content, and significant amounts of legitimate content will be removed as a result. This would leave UK users with a poorer quality Internet relative to elsewhere.

Similar issues have been raised by a number of stakeholders:

- David Kaye, former UN Special Rapporteur for Human Rights, previously [warned](#) the House of Lords Select Committee on Communications and Digital that *"We do not want to see a move towards greater and greater use of algorithmic tools and automation because, when it comes to speech, context means so much"*.
- Harriet Harman MP, Chair of the Joint Committee on Human Rights, also [recently wrote to the Government](#) that *"there is a risk that even with careful guidance providers may err on the side of caution, or introduce systems that result in overly cautious responses, when it comes to content that is not easy to identify as illegal or not."*

Article 7 of the European Union's Digital Services Act avoids this consequence, by ruling out any general monitoring obligations. Google would welcome a similar provision to specify that the Online Safety Bill will not impose a general obligation to monitor content and therefore prevents the risk of over removal of content to the detriment of UK users.

Taking into account impacts on access to information and privacy

We strongly agree that protecting children should be a focus of the new regulatory framework, and have long invested in the safety of our products. We believe that protection of children should be considered holistically, and their right to access information and to privacy must be respected in much the same way that we would expect for adults. Indeed, this balance is a central part of the Age Appropriate Design Code.



In its current form, the duty to “prevent” children of any age encountering primary priority content raises further concerns about general monitoring and risks impacting the rights of children to access information. This is especially the case if the content in question is not clearly and narrowly defined.

In addition, the obligation to mitigate the risks of harm to children “in different age groups” would be complex to implement without collecting more data from children. Further, a system that mandates specific experiences for children according to their age does not take into account the differences in development and maturity that can take place during teenage years and has the potential to restrict children from accessing age-appropriate information and experiences.

We would welcome small additions to the criteria used by Ofcom to assess what is proportionate to comply with the child safety duties. The Bill does not currently require Ofcom to consider the impact of compliance measures on children’s rights to access information, and on both children and adults’ privacy. This is a significant omission, given the impact that any form of age-verification also has on adult users (who would also need to go through some form of age assurance to access services) . We believe these considerations should be taken into account in the legislation to directly align with the provisions already in place in the Age Appropriate Design Code.

Scrutiny of new additions to the Bill

Google would welcome careful consideration from the Committee of the recent additions to the Online Safety Bill, to make sure they work in practice. This is important because these measures were not included in the draft legislation scrutinised by Parliamentarians, nor has there been the opportunity for the perspectives of other stakeholders, including industry and civil society, to be taken into account. The [latest impact assessment](#) shows that nearly 25,100 businesses will be impacted, at a total cost of £250 million. However, this is before the new additions. We would welcome an updated impact assessment.

Advertising

For Google, trust in online advertising - including from brands, consumers, and from society more broadly - is vital to the success of our business. We set a high policy bar for advertisers, publishers and content creators who wish to monetise or advertise, often going above and beyond the current legal framework to ensure a safe experience for users. We apply our advertising policies equally across Google Ads (which provides ads on Search and YouTube).

We welcome the intent of the new duties on fraudulent advertising to curb online fraud and have taken action ahead of legislation. The Online Safety Bill could ensure our enhanced



measures which restrict financial service ads to only those companies regulated by the FCA (or which qualify for a very limited number of exemptions) become the industry norm.

The new version of the Bill includes a requirement for platforms to proactively mitigate against fraudulent advertising. 'Fraudulent' is currently broadly defined in the Bill, including not just fraud by false representation, but also misleading statements or misleading impressions. In practice, identifying an advert that meets this definition may sometimes require knowledge that platforms simply do not have access to. Google will not have the ability to assess whether a specific product has been misleadingly described.

Under the current wording, it is possible that Google would need to establish intensive additional checks on its advertisers to minimise legal risk, which could prohibit or create new obstacles for legitimate UK businesses seeking to advertise online. This would add significantly to cost and stop some businesses from advertising altogether. The impact would disproportionately fall on small UK companies - firms such as Bloom Bakers in Leeds, who have used Google Ads to grow their online sales by 400% over the last 18 months - and the 270,000 small businesses who started advertising online during the pandemic.

Google believes targeted changes in wording would help ensure a more proportionate wider economic impact. This would include specifying that Ofcom's proportionality assessment must have regard to the risk of erroneous removal of legal paid-for advertising and the impact on the commercial rights of service providers and businesses.

Complaints procedures

The latest version of the Bill includes a new requirement to provide a complaints procedure when content has been given lower priority in other users' feeds or in search results or has been otherwise restricted as a result of a service's efforts to comply with the Bill's duties. To give a sense of volume, the Google Search index contains hundreds of billions of web pages and is well over 100,000,000 gigabytes in size, and we show 600 million videos per day in YouTube recommendations.

On both Search and YouTube, this involves constant automated updating of how content is presented to users - in practice, this means that billions of decisions of up-ranking and down-ranking of content are taking place everyday. If platforms are required to offer a new redress mechanism for cases of deprioritisation, this could see legitimate complaints drowned out by spurious ones. This will prevent us from providing fast and meaningful redress to users who have valid complaints.

Further, because decisions about what content to show are temporal in nature and change from one user and one moment to another, there is no way to materially reverse a decision to



not recommend or rank a piece of content above another, raising a major question about what remedy an upheld appeal would practically lead to.

Demoting content is also a key tool used by YouTube to reduce the spread of content that brushes up against our policy guidelines, but does not cross them. As an online platform, we face the persistent challenge of tackling nefarious actors looking to abuse our systems, from individual creators trying to test the boundaries of our content policies to organisations acting at the behest of authoritarian regimes seeking to drive large-scale misinformation narratives on our platforms. For example, we have restricted more than 19,000 videos related to the ongoing Ukraine crisis on YouTube. Individuals and organisations could use the information gained from appeals processes to game our systems. This would significantly undermine our efforts to stop the spread of misinformation.

Meanwhile, there are particular issues around the workability of complaints procedures for “demotions” on Search. Search ranking is different from a newsfeed (as on a user-to-user platform). Something that is ranked in the #1 spot for one query will not necessarily rank in that same spot for another query. This is not a “demotion” per se, but the nature of Search ranking operating in response to queries, rather than to populate a “feed”.

Google believes that the Bill should limit user complaints to removal of content (or account suspensions or terminations). This will ensure the Bill is effective in delivering redress to users who need it, without inadvertently allowing bad actors to game platforms’ efforts to reduce the spread of harmful content.

Journalistic content

Google is absolutely committed to supporting established news publishers through our platforms, recognising the vital role they play in providing authoritative news to consumers. We are one of the world’s biggest financial supporters of journalism and our products and services create significant value for publishers via traffic, advertising funding and funding. We continue to invest in products and programs such as Google News Showcase to provide even more support for journalism.

We acknowledge that news publisher content is excluded from the scope of the regulatory framework. We share the Government’s aim to ensure appropriate protection for journalistic content on user-to-user services and already build these protections into a number of our review processes. On YouTube, content that would otherwise violate the Community Guidelines but which we assess as having [Educational, Documentary, Scientific or Artistic \(EDSA\) value](#) will remain on the platform. To help determine whether a video might qualify for an EDSA exception, we look at multiple factors, including the video title, descriptions and the context provided in the video’s audio or imagery.



The Online Safety Bill currently includes protections for a very broad class of ‘journalistic content’, defined as any UK-linked content generated “for the purposes of journalism” on a user to user service. This content benefits from an expedited appeals process. Under the current wording of the Bill, we see two main risks.

- Firstly, the broad definition could leave the system open to abuse. We have already seen how a number of high-profile individuals posting harmful material in violation of our policies have presented themselves as citizen journalists. We are concerned that the Bill in its current form will further encourage users with no established journalistic credentials to claim that they can publish content in violation of our policies, because they will be able to claim that their content has journalistic value. [As HOPE Not Hate have pointed out](#), a broad definition of journalistic content will provide cover for bad actors and risks creating a loophole in an otherwise rigorous safety regime.
- Secondly, we believe this will lead to users seeking to exploit the expedited appeals process for journalistic content. Without a tighter definition of journalistic content, the expedited appeals mechanism could be overwhelmed with spurious complaints from these individuals, meaning that we would be unable to provide quick redress to genuine news publishers.

We would recommend narrowing the breadth of content and creators covered by these protections to ensure they work effectively and do not inadvertently expose more UK users to harmful content. The current definition of journalistic content should be amended to apply only to recognised news publishers.

In addition, the Government has also committed to introducing an amendment to ensure that journalistic content remains up on a service until an appeal has been resolved. We believe this will expose UK users to harm. For example, we would have been unable to act as swiftly as we did to block RT on YouTube whilst it retained its Ofcom licence. We would also be forced to delay acting on removing harmful content that violates our policies from established outlets (such as claims that dead people voted in the 2020 US election).

User empowerment duties

We support the Government’s aim to better empower users to manage their online experience. However there is the risk that the new user empowerment duties become unworkable. We believe that technical amendments are necessary to ensure that:

- The provisions around reducing the likelihood of encountering priority content (content filtering provisions) do not lead users to lose access to legitimate content.



- The provisions around giving users the option to filter out unverified users (user identification provisions) can work in practice and do not disproportionately impact on UK users' privacy.

For the content filtering provisions, if broad categories of context-dependent content are selected as the priority content that platforms must reduce the likelihood of users seeing, users who choose not to be exposed to any content of one particular category could lose access to important and legitimate content. This is because automated tools used by platforms to filter out specific content cannot accurately assess context and will struggle to tell the difference between legitimate and harmful content. Those using the filtering function would lose access to legitimate content as a result, leaving them with a poorer internet experience than users in other countries, and making the UK an outlier. Further, there is a wider risk of driving users into their own echo chambers, as the implementation of these duties is likely to lead to the exclusion of diverse perspectives on controversial issues.

We would welcome focused changes to the wording of these provisions to help ensure that they are technologically feasible, do not inadvertently restrict users' access to legitimate and useful content, and enable service providers some discretion as to the most appropriate compliance measures to allow the service to deliver on the policy objectives.

We also have two concerns about the user identification provisions. Firstly, we think that it is important to avoid disproportionate impacts on users' privacy. UK Creators and users may have to choose whether to hand over more personal data to platforms or have their voice restricted, without clear evidence for how it will make the online community safer. Secondly, the provision may have been conceived with specific platforms in mind, but a one-size-fits-all approach may not be workable on a platform like YouTube. For example, it is unclear whether this would mean that UK users would lose access to all YouTube videos produced in other countries where users are not required to verify their identity.

We would welcome further refinement to these provisions, to ensure that they do not conflict with the data minimisation principles in the Data Protection Act or limit users' access to legitimate information and views, including voices outside of the UK.

Thank you, once again, for the opportunity to contribute to such a vital stage in the development of this Bill. We would be happy to follow up with amendment suggestions to reflect these points. If you have further questions, do not hesitate to contact me.

Yours sincerely,

Katie O'Donovan

Director of Government Affairs and Public Policy, Google UK