

# **Google – written evidence (DAD0086)**

## **Introduction**

Google welcomes the opportunity to provide written evidence to the House of Lords Committee on Democracy and Digital Technologies. As it is in many areas of life, the internet is playing an increasing role in our democratic debate and it is right that this is properly considered by Parliament. The Committee has asked for evidence in a number of areas. This response sets out:

- Google’s view on the impact of digital technology on how democracy works in the UK.
- The specific steps that Google has taken to help support the democratic process in the UK, including how we support elections and Parliament through Search, Google News and YouTube.
- Our efforts to tackle disinformation, including how we support local news organisations to continue to thrive and play their important role in the democratic process.
- The processes that are in place to prevent abusive and other inappropriate content on Google’s platforms, particularly in relation to the democratic process.
- How Google supports educational programmes around the country, which help to equip internet users with the skills they need to stay safe online, including by identifying potential disinformation.

We take our commitment to tackling abuse on our platforms, including disinformation, incredibly seriously. We are also committed to our role in supporting a healthy democratic process in the UK and elsewhere, and believe our technology can contribute positively to the democratic process. Google hopes this evidence is helpful to the Committee as it conducts its inquiry.

## **The Impact of Digital Technology on the Democratic Process**

Google’s mission is to organize the world’s information and make it universally accessible and useful. Efforts to undermine the integrity of democratic elections, whether by attacking election websites, promoting false information about voting stations and hours of operation, compromising candidate or campaign accounts, or otherwise, are antithetical to that mission. We remain committed to working with government, industry, and civil society to address this challenge around the world.

This is also why Google and YouTube have always had a deep commitment to freedom of expression. This not only includes creative content, but also academia and journalism. YouTube and other Google products provide an opportunity for many people to have a voice and to engage in democratic discourse by democratizing whose stories get told and how. For example, YouTube offers an opportunity to groups that are often more marginalised and whose voices may not have been represented by traditional media, such as young people in the LGBTQ+ community or the victims of domestic abuse, to share their experiences and opinions with a wider audience. Digital platforms have also allowed different groups to form communities, share their experiences and seek support in ways that may not have been available to them before. All of these are important ways that digital technology has broadened access to the democratic process and expanded the ability of people to take part in debate.

To protect election integrity globally, we are building products and programs designed to detect and prevent efforts to subvert elections. Our response focuses on four areas where we are making progress to help ensure the integrity of elections: (i) empowering people with information they can trust when going to the polls; (ii) securing elections; (iii) combating misinformation; and (iv) improving transparency of election advertisements.

## **Technology and Democratic Engagement**

The Committee asked about positive examples of where technology has been used to enhance the democratic process, and how technology can be used to support and enhance democratic engagement.

Google aims to make civic information more easily accessible and useful to people globally as they engage in the political process. We have been building products for over a decade that provide timely and authoritative information about elections around the world and help voters make decisions that affect their countries.

For example, during the 2019 European Parliamentary elections we introduced features across Search in European Union countries to help users find the information they need. When users searched for instructions for how to vote, they could see those details right on the results page. We sourced this data directly from the European Parliament to ensure users got trusted information.

On Google Search, we make algorithmic updates every day to ensure we surface authoritative content that's useful to users. On YouTube, too, we've invested in new product features to make authoritative sources more prominent, including launching Top News and Breaking News in the UK. These features make it easier for

users to find news from verified sources by highlighting videos in the Top News shelf, or showing the Breaking News shelf on the YouTube homepage for key news events.

YouTube also offers additional ways to engage internet users in the democratic process. UK Parliament has a YouTube channel, for example, which hosts and shares videos of parliamentary events and select committee hearings. It also allows the Commons to live stream major events like PMQs, expanding the audience that can access the workings of our democracy on a much wider scale.

## **Tackling Disinformation**

As the Committee's call for evidence notes, disinformation is an area of concern in relation to elections, and one we take seriously at Google. We are a signatory of the European Code of Practice on Disinformation and work with governments and regulators to find ways of removing this type of content from online platforms.

As we've all experienced over the past few years, the words "misinformation", "disinformation", and "fake news" mean different things to different people and can become politically charged when they are used to characterize the propagators of a specific ideology or to undermine political adversaries. However, there is something objectively problematic and harmful to our users when malicious actors attempt to deceive them. It is one thing to be wrong about an issue. It is another to purposefully disseminate information one knows to be inaccurate with the hope that others believe it is true, or to create discord in society. We refer to these deliberate efforts to deceive and mislead using the speed, scale, and technologies of the open web as "disinformation".

The lines between both types of misleading content can be blurry. For instance, content that is created as part of an organized disinformation campaign can be unwittingly propagated by users who believe it to be real. For that reason, the solutions Google adopts for both issues tend to be similar.

Users rely on Google to provide the most relevant, reliable and highest quality information as quickly as possible. This is especially important during elections when voters come to Google for trustworthy information on prospective candidates, party positions, and voting procedures. However, there are those who seek to game the system or flood search results with material that is poor quality, misleading, and deceptive.

We aim to tackle disinformation in two ways. Firstly, we aim to provide users with the most accurate and reliable information, instead of unreliable or misleading content. Secondly, we work to remove disinformation content and to prevent attempts to manipulate our platforms. We have outlined our approach to both in turn below.

### ***Providing access to high quality content***

Google Search and Google News take a pragmatic approach that focuses on giving prominence to high quality media sources and providing users with choice in the news sources that are available to them. We also support a number of initiatives to help support publishers to tackle misinformation. These efforts, taken together, demonstrate Google's commitment to support a fair and open democratic process.

### *Making Quality Count*

We use ranking algorithms to elevate authoritative, high-quality information in our products, and we also take additional steps to improve the quality of our results for contexts and topics that our users expect us to handle with particular care. Google uses algorithms, not humans, to determine the ranking of the content they show to users. Our algorithms are geared toward ensuring the usefulness of our services, as measured by user testing, not fostering the ideological viewpoints of the individuals who build or audit them. As a result, these systems do not make subjective determinations about the truthfulness of webpages, but rather focus on measurable signals that correlate with how users and other websites value the expertise, trustworthiness, or authoritativeness of a webpage on the topics it covers.

We take additional steps to improve the trustworthiness of our results for contexts and topics that our users expect us to handle with particular care. For instance, when our systems detect that a user's search may be news-related, we give more weight in our ranking systems to factors like our understanding of the authoritativeness, expertise, or trustworthiness of the pages we present in response.

We've also been working to make sure that when people come to YouTube looking for information, they get it from authoritative sources. When topics trend in the news, for example, YouTube automatically triggers "Breaking" and "Top News" results. This means that our algorithms kick in and automatically start elevating videos from reliable news sources in dedicated shelves on YouTube's home page or in our Search results.

### *Giving Users More Context*

When users search for news on Google, they are always presented with multiple links but in many cases, they are also presented with additional elements that help them get more context about their search. For instance, Knowledge or Information panels might appear in Google or YouTube Search results to provide context and basic information about people, places, or things. Fact-check tags or snippets might show below links in Google Search and Google News, outlining that a specific piece of content purports to fact-check a claim made by a third party. Or we might call out related searches or questions that users tend to ask about the topic of a search query.

In Google News, additional cues may help users pick up on points of context that are particularly relevant to News stories, such as “Opinion” or “User-Generated Content” tags under articles that news publishers want to signal as such; or algorithmically generated story timelines that let users explore at-a-glance the milestones of a news story over the weeks or months that led to the day’s events. In addition, to help users access context and diverse perspectives about the news stories they read, the “Full Coverage” feature in Google News lets users explore articles and videos from a variety of publishers related to an article or news story of their choice. The “Full Coverage” feature is not personalized and is accessible in one click or tap from most articles in Google News’ “For You” and “Headlines” tabs.

These efforts are intended to help provide users with the diversity of content, and the context, they need to make informed judgements and have a wide-ranging view on the news of the day.

### *Working with industry experts to tackle misinformation*

We know high-quality journalism is an essential part of a healthy society and democratic process. It is also essential to tackling sources of misinformation. We have taken a number of steps to support industry experts, news publishers and fact checking projects.

As a founding partner, Google continues to support the work of First Draft, an independent UK nonprofit that provides ways for news publishers to collaborate in their efforts to identify and debunk misinformation. In addition, through the Google News Lab, we have trained 20,000 journalists in-person and online, providing tools training in news organisations of all sizes.

We’re committed to working with local publishers to ensure communities are served with high quality local news. Local news is an important way of people engaging with the democratic process and political issues in their areas. That is why

we have a local news tab on Google News, to make it easier for people to find stories that are relevant to their part of the UK. But for local news outlets to be able to thrive and to play this leading role in the democratic process, they must be able to adapt for the modern era. That is why we are directly funding digital innovation in news to the tune of more than £100m, enabling news organisations from the Bureau Local to Cardiff University's Centre for Community Journalism to test and scale new ways to bring in traffic.

We recently announced that we are partnering with Archant through our Local Experiments Project to help local journalism find ways to thrive in the digital age.

Across the world, we are investing \$300m in a series of programmes and tools to help publishers as part of the Google News Initiative, all of which will help ensure that these publishers can continue to play their important role in the democratic process and help in the fight against misinformation.

### ***Counteracting disinformation from malicious actors and promoting a safe democratic process***

We also look for and take action against attempts to deceive our ranking systems or circumvent our policies. This includes disinformation content on either Search and YouTube, as well as other issues such as phishing or other forms of cybercrime.

#### *Removing disinformation on our platforms*

There are synergies between our work to tackle disinformation and our efforts to tackle the wide range of actions that scammers have attempted to use over the past decades to trick our systems in order to promote their own content (via a set of practices we refer to as "spam"). While not all spammers engage in disinformation, many of the malicious actors who try to distribute disinformation (at all levels of sophistication or funding) engage in some form of spam. The tactics they use are similar to those of other spammers. Therefore, our work against spam goes hand-in-hand with our work against disinformation. Our algorithms can detect the majority of spam, including disinformation content, and demote or remove it automatically. The remaining spam is tackled manually by our spam removal team, which reviews pages (often based on user feedback) and flags them if they violate our guidelines.

We also remove content from YouTube that breaches our Community Guidelines. One example of a type of content that we deem to be abusive and that would breach our guidelines would be content that deliberately seeks to spread disinformation that could suppress voting or otherwise interfere with democratic or

civic processes. For example, demonstrably false content that claims one demographic votes on one day while another votes on a separate day would be in violation of our policies. Other examples of content and channels we remove include content containing hate speech or that has been used to harass other people, or channels that are deliberately set up to mislead by impersonating other users. Similarly, we take action in Google News and our Advertising services against publishers that misrepresent material details about themselves such as their ownership or country of origin.

### *Security and working together to tackle abuse*

We also know that elections pose particular challenges that require all of our teams across Google and YouTube to work together. We have established dedicated teams across Europe who specialize in preventing abuse of our systems during elections, ranging from phishing schemes to attempts to hack Google products. These teams are trained to get ahead of abuse, clamp down on malicious activity, and react rapidly to breaking threats. Our teams work in partnership to identify malicious actors, disable their accounts, warn our users about them, and share intelligence with other companies and law enforcement officials.

Google has written to all MPs to inform them of the security tools we make available to them. We included information on how to leverage our strongest security features for Google accounts through our Advanced Protection Program, and Project Shield, a free service that uses Google technology to help protect news, election and campaign sites from DDoS attacks.

We are aware that we have constant work to do to keep our users safe and we actively welcome the thoughts and feedback of researchers, policymakers, civil society, and journalists around the world who help us identify and solve problems with our algorithms and other efforts to tackle disinformation.

### *Algorithms and security*

In its call for evidence, the Committee also asked about the role of algorithms. Algorithms are a central part of Google's approach and we have indicated areas throughout our response where they play a central role in providing our users with high quality content and to keep them safe from harm.

Because the malicious actors who propagate disinformation have the incentive to keep doing so, they continue to probe for new ways to game our systems, and it is incumbent on us to stay ahead of this technological arms race. A compounding factor in that challenge is that our systems are constantly confronted with searches

they have never seen before. Every day, 15% of the queries that our users type in the Google Search bar are new. For these reasons, we regularly evolve our ranking algorithms, our content policies, and the partnerships we enter into as part of our efforts to curb disinformation.

We have an important responsibility to our users and to the societies in which we operate to curb the efforts of those who aim to propagate false information on our platforms. At the same time, we respect our users' fundamental human rights (such as free expression) and we try to be clear and predictable in our efforts, letting users and content creators decide for themselves whether we are operating fairly. Of course, this is a delicate balance, as sharing too much of the granular details of how our algorithms and processes work would make it easier for bad actors to exploit them. History has taught us that providing too much detail in public about our processes can provide opportunities for malicious actors to try and manipulate Google's search results. The earliest example of such attempts dates back to 1999, when Google's founders published a seminal paper on PageRank, a key innovation in Google's algorithm. The paper described how our algorithms use links between websites as an indicator of authority. Once that paper was published, spammers immediately tried to manipulate Google search results by paying each other for links.

### *Transparency in Political Advertising*

To help people better understand the election ads they see online and support the integrity of elections, earlier this year we implemented a new process to verify advertisers for the EU Parliamentary election. This policy covers:

- A political party, a current elected officeholder or candidate for the EU Parliament;
- A political party, a current officeholder or candidate for an elected national office within an EU member state. Examples include members of a national parliament and presidents that are directly elected; or
- A referendum question up for a vote, a referendum campaign group or a call to vote related to a national referendum or a state or provincial referendum on sovereignty.

We also require that these verified election ads incorporate a clear "paid for by" disclosure. We have expanded our portfolio of transparency reports to include an EU Political Advertising on Google Transparency Report to show voters who is purchasing election ads on Google in the EU and how much money is being spent. This report includes a searchable election Ad Library that provides valuable data such as which ad had the highest impressions, what the latest election ads running

on the Google platform are, and how the ads are targeted in terms of age, gender, and location.

### **Policy Enforcement: Keeping our users, including those in public life, safe from harm**

As the Committee rightly notes in its call for evidence, disinformation is not the only area where companies such as Google have to take action to keep our platforms and users safe. We are aware that abusive content aimed at people in public life, and particularly political figures and election candidates, is an area of particular concern and we take our responsibilities in this area incredibly seriously.

At YouTube, we have Community Guidelines that set the rules of what is not allowed on the platform. We remove content that violates these guidelines, whether in videos or comments. These policies prohibit, for example, hate speech, harassment, and cyberbullying. We have a dedicated policy development team that systematically reviews all of our policies to ensure that they are current, keep our community safe, and do not stifle YouTube's openness.

While we have always had a longstanding policy against hate speech; for example, this year we updated our policy to prohibit videos that allege that a group is superior in order to justify discrimination, segregation or exclusion based on qualities like gender, race, religion and sexual orientation - this includes white nationalist and supremacists content.

We have a rigorous approach when it comes to taking enforcement action against controversial content. Our strict policies are combined with investments in technology and people to detect and remove content that does not comply with our Community Guidelines. Alongside an easy system for any YouTube user to flag content, we use machine learning to detect violative content at scale, and employ specialist teams to review the content flagged by those machines to make decisions based on context. We sometimes use hashes (or "digital fingerprints") to catch copies of known violative content before they are ever made available to view. For some content, like child sexual abuse images (CSAI) and terrorist recruitment videos, we contribute to shared industry databases of hashes to increase the volume of content our machines can catch at upload.

Our teams also run rigorous quality checking to be sure our machine systems are effective. We now have over 10,000 people addressing violative content and we have hired highly trained experts across specific categories, including hate speech.

This group includes Intelligence Analysts, Data Scientists, Engineers and Product Managers who help us detect emerging trends in how people and organisations are manipulating platforms across the internet to spread hate and terror. We have also created an Intelligence Desk for the proactive identification of issues and have updated our policies and our systems, making thirty policies changes in the last year alone. We also continue to invest in our network of over 180 academics, government partners, and NGOs who bring valuable expertise to our enforcement systems, including through our Trusted Flagger program, which gives eligible individuals, government agencies, and non-governmental organisations additional tools for flagging content.

These efforts have been a success - problematic content represents a fraction of one percent of the content on YouTube and we're constantly working to reduce this even further. We are committed to demonstrating further progress through our quarterly Community Guidelines Enforcement Reports. For example, in the second quarter of 2019, we removed 9.0 million videos for violating our Community Guidelines. 87 per cent of the content that was removed for violating our guidelines was first flagged by our machine learning technology and over 80% of these machine flagged videos were removed without receiving a single view.

We are continuously improving our systems and processes, allowing us to respond to situations faster and more effectively. Whenever an election campaign occurs, we continue to use our combination of people and technology to enforce our Community Guidelines.

We also develop tools to give creators options to manage their content, like enhanced moderation controls to manage comments on YouTube channels. YouTube account holders can delete inappropriate comments and block a user so they can't view videos or leave more comments. Comments can also be turned off for any video by the uploader or managed by requiring pre-approval before they are posted publicly.

In addition to removing content that violates our guidelines, we raise up authoritative voices when people are looking for breaking news and information, especially during breaking news moments. Our breaking and top news shelves are available in 40 countries and we're continuing to expand that number.

We also reduce the spread of content that brushes right up against our policy line. Already, in the U.S. where we made changes to recommendations earlier this year, we've seen a 50% drop in watch time to this type of content, meaning quality content has more of a chance to shine. We have begun experimenting with this change in the UK, Ireland, South Africa and other English-language markets.

We are committed to living up to our responsibility and have invested significantly in the teams and systems that protect YouTube using these principles.

## **Education**

The investment that Google makes in technology and staff, and the policies and processes we put in place, are an incredibly important aspect of our work. But as the Committee has recognised in its terms of reference, just as important are the steps taken by ourselves, the Government and other partners to educate internet users in order to create a healthy, active, digitally literate democracy. This includes education on recognising misinformation and disinformation and being a good internet citizen.

We provide a range of free skills training schemes for people of all ages, including Be Internet Citizens. This is a programme run by Google and the Institute for Strategic Dialogue which encourages young people aged 13-15 years old to have a positive voice online, and provides training on how to navigate social media, strengthen critical thinking, and escape social media bubbles.

Be Internet Citizens has been designed to teach 13 to 15-year-olds about media literacy, critical thinking and digital citizenship. The programme provides participants with a strong foundation, empowering them with the confidence to become producers of online content, where they can express their identities, share their stories, make a social impact, and bring communities together.

The programme is delivered as a mix of our PSHE Association accredited curriculum pack, regional train the trainer events for teachers and youth workers, and selected workshops in high schools. The programme provides participants with a strong foundation, empowering them with the confidence to become producers of online content, where they can express their identities, share their stories, make a social impact, and bring communities together.

Since launching the programme in 2017, we have trained over 35,000 young people through our resources and workshops. According to the Institute of Strategic Dialogue's 2018 report into the programme, 92% of teenagers who participated felt they had gained new knowledge and 86% felt they had gained new skills. We are particularly proud that 71% said they would behave differently online as a result of the programme.