

Facebook – written evidence (DAD0081)

Thank you for the opportunity to respond to this inquiry. Please find below our responses to the questions outlined in the call for evidence.

How can politicians and political institutions use technology to engage the public with national and local decision making and enhance democracy?

- Social media makes it easier for people to have a voice in government and politics.
- Facebook is a great resource for candidates and for enabling those who are elected to engage with constituents, share information about their work and hear feedback.
 - Our own research shows that 85% of UK MPs have a Facebook profile and around 390 MPs use the platform regularly to post updates on their Westminster or constituency work. And this works both ways: 2.6 million people in the UK like or follow an MP's page.
 - According to ComRes' latest Insider Insights Report on Social Media, almost four in five MPs (78%) agree that social media is now an essential communication tool for political debate and seven in ten (70%) agree that it has changed the way campaigning is done in their constituency.
- We also believe that democracies are stronger when more people are engaged, which is why we remind eligible voters on the day of the election to go vote and connect them with helpful resources.
 - On 10 April we launched a voter registration reminder ahead of the 12 April deadline for registration in the local elections, and an Election Day Reminder on the day of the election,
 - On 23 May, the day of voting for the European election in the UK, we launched an Election Day Reminder. The reminder directed people to the official election information website and allowed people to share that they voted.

What role should education play in helping create a healthy, active and digitally literate democracy?

- We want to help people make informed decisions on the information they're reading and empower them to decide for themselves what to read, trust, and share. We do so by partnering with local organisations, sponsoring media literacy programs, and sharing [tips to help spot false news](#).
- We are particularly interested in improving people's basic level of digital skills and familiarity with the digital world. This will be important if people are to make full use of the range of digital tools now available, including digital payment infrastructure.
- In August 2018 we [launched](#) our [Digital Literacy Library](#), a collection of lessons to help young people think critically and share thoughtfully online. There are [830 million young people online](#) around the world, and this library

is a resource for educators looking to address digital literacy and help these young people build the skills they need to safely enjoy digital technology.

- To help equip young people in the UK to spot false news, and improve their digital literacy, we worked with the APPG on Literacy, the National Literacy Trust, First News, and The Day to launch the Commission on Fake News and the Teaching of Critical Literacy Skills in Schools. As part of this work, we are working with these partners to survey young people on their experiences of fake news and help evidence gathering in this area.
- We are also a contributor to the Media Smart media literacy programme for 7 to 16 year olds. It has free educational resources for teachers, parents and young people - including a dedicated body image programme for boys and girls. Media Smart helps young people think critically about the advertising they come across in their daily lives using real and current examples of advertising that we're all familiar with to help teach core media literacy skills.

Should there be more transparency in online spending and political campaigning by political parties and other groups? What is the effect of targeted online advertising?

- Advertising makes it possible for a message to reach many people, so it is especially important that advertisers are held accountable for what they promote and that fake accounts are not allowed to advertise.
- Some people have argued that getting rid of political content and ads from Facebook is the only sure-fire way of guarding against foreign interference.
- We believe that banning political ads on Facebook would tilt the scales in favour of incumbent politicians and candidates with deep pockets. Digital advertising is typically more affordable than TV or print ads, giving less well-funded candidates a relatively economical way to reach their future constituents. Similarly, it would make it harder for people running for local office — who can't afford larger media buys — to get their message out. And issue ads also help raise awareness of important challenges, mobilising people across communities to fight for a common cause.
- As a result of changes we made in 2018, Facebook now has a higher standard of ads transparency than has ever existed with billboard or newspaper ads.
- All advertisers wanting to run ads in the UK that reference political figures, political parties, elections, legislation before Parliament and past referenda that are the subject of national debate, now need to verify their identity and location and carry a "Paid for by" disclaimer. Starting late October, we will also require advertisers wanting to run ads about eight social issues - Immigration, Political Values and Governance, Civil and Social Rights, Security and Foreign Policy, Economy, Environmental Politics, Health, and Crime. We see this as an important part of ensuring electoral integrity and helping people understand who they are engaging with.
- We recognise that this is ~~going to be~~ a significant change for people who use our service to publish this type of ad. While the vast majority of ads on Facebook are run by legitimate organisations, we know that there are bad

actors that try to misuse our platform. By having people verify who they are, we believe it will help prevent abuse.

- To further increase transparency, we launched a public [Ad Library](#) that contains a comprehensive, searchable collection of all ads that are active and running on Facebook and Instagram. This includes non-political ads. Political ads, however, are archived and remain in the Ad Library for seven years. The Ad Library shows range of impressions, range of spend, range of age, gender and geographical reach for each ad. Anyone can explore the Library, with or without a Facebook account.
- We've also introduced the [Ad Library Report](#), which provides aggregated insights about political ads in the Ad Library. The report shows the total number of ads and total spend on ads in the Ad Library by country, as well as total spend by advertiser, advertiser spend by day and top searched keywords from the past week. Users can explore, filter and download the data into a CSV file. The Ad Library Report is useful to people who are interested in understanding high-level activity in the Ad Archive since launch, while the Ad Library could be used to deep dive into specific ads.
- In June 2018, we launched Page Transparency, a feature which provides information about historical and admin information about Pages, including the date the Page was created; whether the Page name has changed; whether the Page has merged with other pages; and the primary country locations of people who manage the Page.
- We continue to update and improve [Why am I seeing this ad?](#) and Ad Preferences - tools we introduced over four years ago to provide people greater transparency and control over the ads they see. In the recent updates, people can see more detailed targeting, including the interests or categories that matched them with a specific ad. It will also be clearer where that information came from (e.g. the website they may have visited or Page they may have liked), and we'll highlight controls people can use to easily adjust their experience. We have also updated the [Ad Preferences](#) tool to show people more about businesses that upload lists with their information, such as an email address or phone number.

What steps can be taken to reduce the impact of misinformation online?

- Fighting false news and misinformation is a top priority for Facebook, and because the challenge is multifaceted, we take a variety of approaches through product and policy interventions. Our approach to misinformation is guided by the principle that we should not act as the arbiter of truth, and we should not have policies requiring everything our users post should be truthful. However, we also know that our users want to see high quality content on our platform which is why our strategy to combat misinformation has three parts: remove, reduce and inform.
- We think that reducing the distribution of misinformation—rather than removing it outright—strikes the right balance between giving people a place to express themselves and creating a community that's safe and authentic. There are some extreme forms of misinformation that violate our community standards that we remove from our platform.

- Fake accounts are not permitted under our community standards, and are one of the primary vehicles for spreading misinformation -- especially politically-motivated misinformation and propaganda. We try to stop fake accounts abusing our platforms in three distinct ways. Of the accounts we remove, both at sign-up and those already on the platform, 99.8% of these are proactively detected by us before people report them to us.
 - Blocking accounts from being created: The best way to fight fake accounts is to stop them from getting onto Facebook in the first place. We've built detection technology that can detect and block accounts even before they are created. Our systems look for a number of different signals that indicate if accounts are created in mass from one location. A simple example is blocking certain IP addresses altogether so that they can't access our systems and thus can't create accounts.
 - Removing accounts when they sign-up: Our advanced detection systems look for potential fake accounts as soon as they sign-up, by spotting signs of malicious behavior. These systems use a combination of signals such as patterns of using suspicious email addresses, suspicious actions, or other signals previously associated with other fake accounts we've removed. Most of the accounts we currently remove, are blocked within minutes of their creation before they can do any harm.
 - Removing accounts already on Facebook: Some accounts may get past the above two defenses and still make it onto the platform. Often, this is because they don't readily show signals of being fake or malicious at first, so we give them the benefit of the doubt until they exhibit signs of malicious activity. We find these accounts when our detection systems identify such behavior or if people using Facebook report them to us. We use a number of signals about how the account was created and is being used to determine whether it has a high probability of being fake and disable those that are.
- We have learned that in some cases, people post content that, on its face, doesn't violate. However, coupled with an understanding of ethnic and religious tensions on the ground or a more complete view of the publisher and the target, seemingly benign online content has the potential to incite or cause real world harm. To address this reality, we updated our Community Standards such that we remove misinformation that may contribute to [imminent violence or physical harm](#). To operationalize the policy, we have on-boarded local organizations and international institutions, who are often the first to become aware of inaccurate or misleading information that may contribute to real world harm. We have secure channels for these organizations to reach out to us and escalate content quickly.
- We also remove misinformation that contributes to voter suppression (actions that are designed to deter or prevent people from voting). This includes misrepresentation about how to vote, such as claims that

you can vote using an online app, and statements about whether a vote will be counted.

- As mentioned, we think that in most cases reducing the distribution of misinformation - rather than removing it outright - strikes the right balance between encouraging free expression and promoting a safe and authentic community, and we believe that down-ranking inauthentic content strikes that balance. In other words, we allow people to post it as a form of expression, but we're not going to show it at the top of News Feed.
 - When we demote stories rated false by fact checkers, we are able to reduce its future views. News Feed is ranked based on hundreds of thousands of signals, and some signals can identify clickbait, false news or engagement bait, which can factor into that post's ranking in a negative way if identified. Other signals we use to identify false content includes factors like the content's source, who is sharing it, and how people engage with the content.
 - One way we reduce the distribution of spammy content or false news in News Feed is through our third-party fact-checking program. Our fact-checking partners have all been certified through a non-partisan International Fact-Checking Network. Today we have 55 partners in 43 languages, including two in the UK (Full Fact and Fact Check NI), and we're working hard to continue expanding our third-party fact-checking program.
 - We identify content to be reviewed by our fact-checkers through a combination of technology and human review. Our machine learning models identify links to articles which might be false and we can use the model predictions to prioritize the links we show third-party fact-checkers. Since there are hundreds of millions of links per week being shared on Facebook, we prioritize third-party fact-checkers' time with algorithms to surface suspicious links.
 - Fact-checkers can also proactively identify the content they would like to review themselves. We ask our partners to focus their fact-checking for Facebook on identifying and addressing the worst of the worst - clear misinformation and fake news intended to harm and mislead people. After consulting the public guidelines of our fact-checking partners, we identified four consistently emphasized criteria for fact-checkers to consider when prioritizing work on Facebook: verifiability, importance, relevance and virality. These fact-checkers then review the stories, check their facts, and rate their accuracy. False stories are demoted in News Feed so fewer people see them. These confirmed fact-checker ratings then help further train our machine learning model, creating a continuous cycle of improvement.
- In addition to removing misinformation that violates our community standards, and reducing the distribution of false rated news in our News Feed, we are also investing in new ways to provide more context to users about the information they are seeing on the platform. For Page owners, we

notify them if they have shared something that was given a “false” or “mixture” rating by a fact-checker. We’ll also notify people who are about to share or have shared that post in the past.

- We also recently developed a feature to give people more information about the publishers and articles they see. For example, when a fact-checker writes an article that explains why a claim is false, misleading, or true, we’ll show that explainer article alongside the content in question in News Feed.
- Last year we launched a new feature called the [Context Button](#) to make it easier for people to view more information about websites and publishers they see on Facebook. This spring, we expanded the Context Button to images that have been fact-checked on Facebook.
- And finally when it comes to fighting misinformation, one of the most effective approaches is removing economic incentives for traffickers of misinformation. We’ve found that a lot of fake news is financially motivated. These spammers make money by masquerading as legitimate news publishers and posting hoaxes that get people to visit their sites, which are often mostly ads.
 - For Pages and domains that repeatedly share false news, their distribution will be reduced and their ability to monetize and advertise will be removed. These repeat offenders will no longer be allowed to advertise on Facebook. This policy includes Pages that are sharing false news but didn’t create it themselves because we believe that they, like publishers, have a responsibility for the content they share with their audiences. Once pages and domains are no longer repeat offenders, we no longer reduce their distribution and once again allow them to advertise and monetize.
- However, we’ve found that misinformation is spread in three main ways:
 - By fake accounts, including for political motivation;
 - By spammers, for economic motivation; and
 - By regular people, who often may not know they’re spreading misinformation.
- Misinformation that can contribute to real world violence has been one of the hardest issues we’ve faced. In places where viral misinformation may contribute to violence we now take it down. In other cases, we focus on reducing the distribution of viral misinformation rather than removing it outright.
- Economically motivated misinformation is another challenge that may affect elections. In these cases, spammers make up sensationalist claims and push them onto the internet in the hope that people will click on them and they’ll make money off the ads next to the content. This is the same business model spammers have used for a long time, and the techniques we’ve developed over the years to fight spammers apply here as well.
- The key is to disrupt their economic incentives. If we make it harder for them to make money, then they’ll typically just go and do something else instead. This is why we block anyone who has repeatedly spread misinformation from using our ads to make money. We also significantly reduce the distribution of any page that has repeatedly spread misinformation and spam. These

measures make it harder for them to stay profitable spamming our community.

- The third major category of misinformation is shared by regular people in their normal use of our services. This is particularly challenging to deal with because we cannot stop it upstream -- like we can by removing fake accounts or preventing spammers from using our ads. Instead, when a post is flagged as potentially false or is going viral, we pass it to independent fact-checkers to review. All of the fact-checkers we use are certified by the non-partisan [International Fact-Checking Network](#). Posts that are rated as false are demoted.
- Taking these strategies together, our goal around misinformation for elections is to make sure that few, if any, of the top links shared on Facebook will be viral hoaxes.
- We have launched a joint research project with the National Literacy Trust in the UK which we are funding. It works with youth publishers 'The Day' and 'First News' and will survey young people on their experiences of fake news.

Does the increasing use of encrypted messaging and private groups on social media platforms present a challenge to democracy? What are the positive and negative effects of anonymity on online political debate?

- End-to-end encryption is a powerful tool for security and safety. It lets patients talk to their doctors in complete confidence. It helps journalists communicate with sources without governments listening in. It gives citizens in repressive regimes a lifeline to human rights advocates. And by end-to-end encrypting sensitive information, a cyber attack aimed at revealing private conversations would be far less likely to succeed.
- Encryption helps protect our information and safeguards its misuse by criminals and terrorists who could exploit it. Encryption plays an important role in securing our banking and e-commerce infrastructure and protects billions of communications and transactions every day, including those on Facebook.

How has the design of algorithms used by social media platforms shaped democratic debate? Should there be more accountability for the design of algorithms?

- Organizations of all shapes and sizes use algorithms -- a set of rules for solving a problem in a finite number of steps -- every day in countless ways. They help businesses, schools, hospitals and other institutions operate more efficiently and predictably.
- At Facebook, we use algorithms to assist our internal teams, offer customized user experiences, and help us achieve our mission of building a global and informed community. This includes algorithms used to generate and display search results; prioritize the content people follow with their personalized News Feed; and match people with the ads most relevant to them.

- Whether people are new to Facebook or have been using it for years, they should be able to easily understand and adjust how their information influences the ads and posts they see. We want to empower people to decide for themselves what to read, trust, and share, and we do this by providing greater transparency and control.
- We are committed to being transparent with our users, including as it relates to policies and our use of algorithms.
 - We publish a series of blog posts called [News Feed FYI](#) that highlights major updates to News Feed and explains the thinking behind them.
 - We are also promoting a series of AI educational initiatives and campaigns to help people learn about the technology that underlies our various products and features, which includes AI and machine learning
- We've also introduced tools and continue to add more information and features so that people can have more context and control over what they see in their News Feed.
 - In March, we introduced [Why am I seeing this post?](#) Similar to [Why am I seeing this ad?](#), it provides people with greater transparency and control over what they see from friends, Pages and Groups in their News Feed. This is the first time that we've built information on how our News Feed ranking works directly into the app. The new tool explains how people's past interactions impact the ranking of posts in their News Feed. Shortcuts to controls - such as See First, Unfollow, News Feed Preferences and Privacy Shortcuts - have also been added to help people more easily personalize their News Feed.
 - In August, to help shed more light on these practices that are common yet not always well understood, we have introduced a new way for people to view and control their off-Facebook activity. [Off-Facebook Activity](#) lets people see and control the data that other apps and websites share with Facebook.

Are people or organisations deliberately using social media to undermine trust in democracy? How can this be combatted?

- Protecting the integrity of elections is one of Facebook's highest priorities. It's why we've worked to develop a comprehensive strategy to close previous vulnerabilities, while addressing new and emerging threats.
- We've developed smarter tools, greater transparency, and stronger partnerships to help us do just that. We've blocked millions of fake accounts so they can't spread misinformation. We're working with independent fact-checkers to reduce the spread of fake news. And we've set a new standard for transparency in Pages and political ads so people can see who is behind them.
- We know that security is never finished and we can't do this alone - that's why we continue to work with policymakers and experts to make sure we are constantly improving.

- Security is an arms race, and as we continue to improve our defenses, our adversaries evolve their tactics. We will never be perfect, and we are up against determined opponents. But we are committed to doing everything we can to prevent people from using our platforms to interfere in elections.
- We are also improving our rapid response efforts. We have more than 30,000 people working on safety and security, with 40 teams contributing to our work on elections globally.

Smarter Tools

- Our teams are working to build innovative new tools, combining stronger artificial intelligence with expert investigators, to find and prevent abuse, including:
 - stopping millions of [fake accounts](#) from being created every day;
 - finding and removing thousands of Pages, Groups, and accounts involved in [coordinated inauthentic behavior](#); and
 - reducing the spread of [false news and misinformation](#), as independent [studies](#) have confirmed
- We've introduced:
 - sophisticated teams of specialized investigators to locate, analyze, and disrupt malicious actors;
 - better technology to proactively find and block voter suppression and other content that violates our [Community Standards](#);
 - improved capabilities to remove violating content in bulk and prevent it from being posted again;
 - and new abilities to detect tampered images, which we've already deployed for elections in the EU and India.

Greater Transparency

- We've made some [big changes](#) around political and issue ads so people can see who's trying to influence their vote:
 - Advertisers that want to run political and issue ads now need to [get verified](#), proving who they are and where they are located;
 - We label political and issue ads so people know who has paid for them;
 - We archive political ads in an [Ad Library](#) so people can find more details like what types of people saw an ad and how much was spent.
 - The Ad Library now houses over 4 million ads and continues to be a critical tool for many - especially regulators, journalists, and researchers - to scrutinize and hold advertisers to account

Stronger Partnerships

- Today, we have over 55 independent fact-checking partners around the world covering 43 languages.

- We continue to improve our coordination and cooperation with government authorities, to allow for better information sharing and threat detection. We are also working with academics, civil society groups, and researchers, including the Atlantic Council's Digital Forensic Research Laboratory, to get the best thinking on these issues.
- We know we can't do this alone, and these partnerships are an important piece of our comprehensive efforts to fight election interference.

Supporting an Informed Electorate

- People are already using Facebook to talk about politics and issues that matter to them and to communicate with their elected officials. We want to support them and make it easier to vote and connect to reliable and accurate information.
- On election day, we will launch a reminder that it's time to vote and help people find their [polling information](#). This reminder will be displayed on people's News Feed.

Prohibiting Coordinated Inauthentic Behaviour

- Accounts, Pages and Groups removed for violations of our [Coordinated Inauthentic Behavior](#) (CIB) policy is a small subset of the overall accounts and content we remove from our platforms.
- We're committed to continually improve to stay ahead by building better technology, hiring more people and working more closely with law enforcement, security experts and other companies. We have built an expert team of analysts with experience in law enforcement, cybersecurity, and intelligence to find, investigate, and disrupt these operations. We are also growing our technical teams to build scaled solutions to finding and preventing these behaviours.
- We're constantly working to detect and stop CIB activity because we don't want our services to be used to manipulate people. When we take down Pages, Groups, accounts, or Instagram accounts we remove them for misleading behaviours, not for the content they post. These behaviours include the use of fake accounts, often in coordination with real accounts, designed to mislead people about the identity of people behind an operation or the source or origin of the content.
- We prioritise building partnerships with civil society organisations and sharing information with government partners about the types of operations we anticipate and key opportunities for joint action to disrupt these threats.
- Security is a constant mission that is never complete. But we're making significant progress and are committed to staying ahead in protecting our platforms against sophisticated information operations.
 - In 2017, we only had 1 CIB operation taken down; in 2018, we had a couple dozen; in H1 2019 alone, we had more takedowns than 2017 and 2018 combined.
- In 2019, the following CIB takedowns were announced in our Newsroom:

- [Banning Twinmark Media Enterprises in the Philippines from Facebook](#) (January 10)
- [Removing Coordinated Inauthentic Behavior from Russia](#) (January 17)
- [Removing Coordinated Inauthentic Behavior From Iran](#) (January 31)
- [Taking Down Coordinated Inauthentic Behavior in Indonesia](#) (January 31)
- [Removing Coordinated Inauthentic Behavior From Moldova](#) (February 13)
- [Removing Coordinated Inauthentic Behavior From the UK and Romania](#) (March 7)
- [Removing Coordinated Inauthentic Behavior from Iran, Russia, Macedonia and Kosovo](#) (March 26)
- [Removing Coordinated Inauthentic Behavior From the Philippines](#) (March 28)
- [Removing Coordinated Inauthentic Behavior and Spam From India and Pakistan](#) (April 1)
- [Removing More Coordinated Inauthentic Behavior From Russia](#) (May 6)
- [Removing Coordinated Inauthentic Behavior From Israel](#) (May 16)
- [Removing More Coordinated Inauthentic Behavior From Iran](#) (May 28)
- [Removing Coordinated Inauthentic Behavior in Thailand, Russia, Ukraine and Honduras](#) (July 25)
- [Removing Coordinated Inauthentic Behavior in UAE, Egypt and Saudi Arabia](#) (August 1)
- [Removing Coordinated Inauthentic Behavior From China](#) (August 19)
- [Taking Down More Coordinated Inauthentic Behavior in Myanmar](#) (August 21)
- [Removing Coordinated Inauthentic Behavior From Iraq and Ukraine](#) (September 16)
- [Removing Coordinated Inauthentic Behavior in Spain](#) (September 20)
- [Removing Coordinated Inauthentic Behavior in UAE, Nigeria, Indonesia and Egypt](#) (October 3)