ISD | Institute for Strategic Dialogue

2019 EU Elections

# Cracking the Code: An Evaluation of the EU Code of Practice on Disinformation

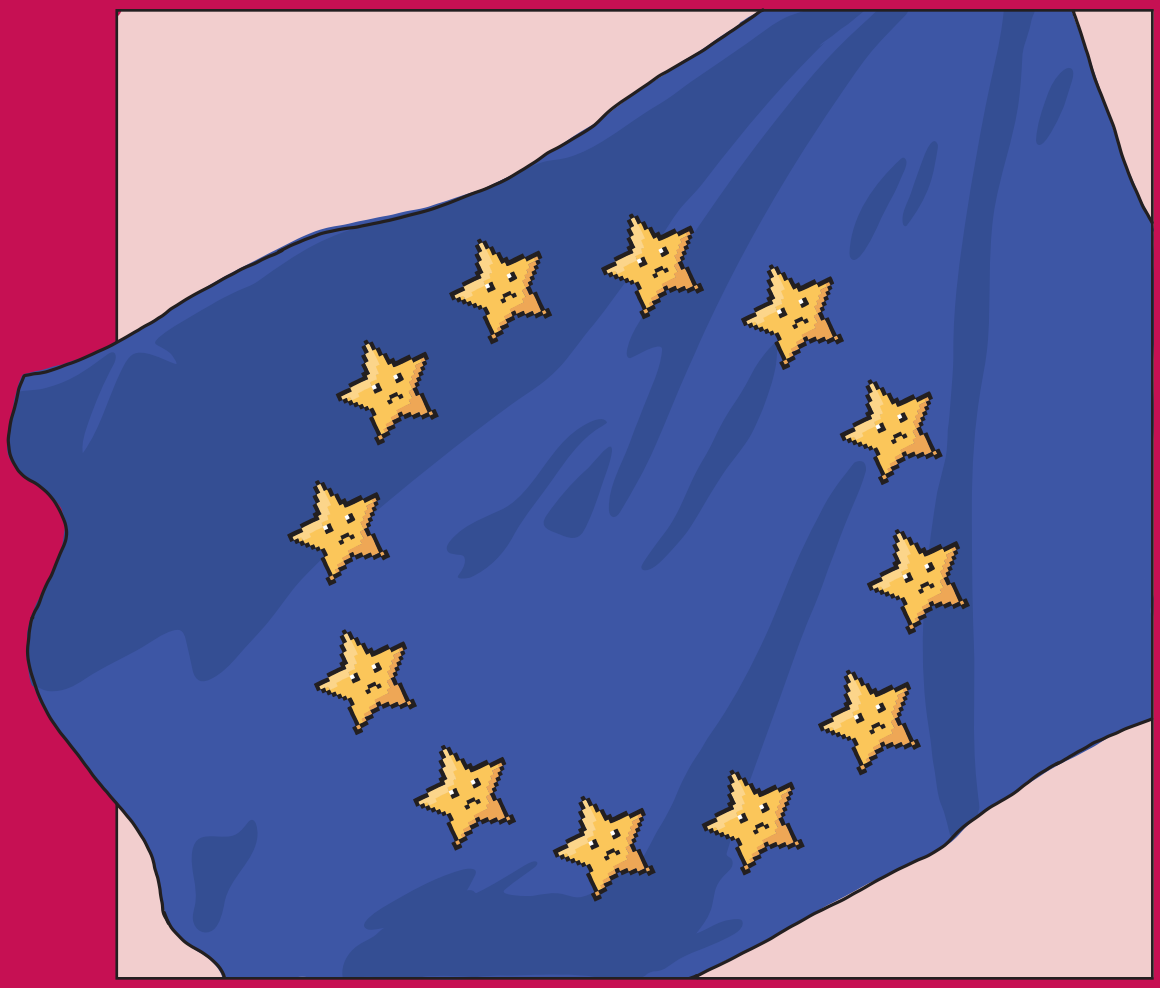

Chloe Colliver

## About this Report

In October 2018, Facebook, Google, Twitter, Mozilla and a selection of advertising industry companies signed up to the newly drafted EU Code of Practice on Disinformation (CoPD). This assessment attempts to evaluate the enforcement of the CoPD during the EU parliamentary elections. It finds that the CoPD prompted progress from tech companies in dealing with specific areas of disinformation risk, most notably transparency for political advertising. However, the effectiveness of the CoPD in achieving substantive changes was fundamentally challenged by its self-regulatory set-up and the lack of enforcement in place for non-compliance.

This report provides a set of recommendations, which seek to inform the continuing efforts to counter disinformation and online harms at the EU level through the upcoming Digital Services Act and European Democracy Action Plan in 2020 and beyond. The Institute for Strategic Dialogue (ISD) calls for policymakers in the EU to design and enforce systemic transparency for advertising, content moderation, appeals and redress systems, and algorithmic design and output in order to address the risks posed by disinformation in the European context.

This assessment is based on research conducted by ISD and additional insights from partner research organisations who evaluated the presence and scale of disinformation targeting the European parliamentary elections of May 2019 on the social media platforms that signed up to the CoPD. The full methods and findings of that research are available in the accompanying report, *Click Here for Outrage: Disinformation in the European Parliamentary Elections 2019*.[1]

## About ISD's Digital Analysis Unit

ISD's Digital Analysis Unit combines social listening and natural language processing tools with leading ethnographic researchers to better understand how technology is used by extremist and hateful groups. We use commercial tools that aggregate social media data to analyse broad trends in discussion and how they may be influenced by hateful groups and disinformation. Using tools co-developed by ISD, we are able to analyse specific types of hateful speech online and trace where this speech comes from. We use these insights to help policymakers and companies craft informed policy responses to hate and disinformation, and to help communities mount responses at the local level.

The research for this report was supported by a grant from Open Society Foundations. The report was produced with support from Luminate.

# Introduction

**In October 2018, Facebook, Google, Twitter, Mozilla and a selection of advertising industry companies signed up to the newly drafted EU Code of Practice on Disinformation (CoPD).[2] The document was drawn up to provide voluntary standards and commitments intended to reduce the spread of disinformation.**

It covers a broad swathe of the products and tactics known to be used to spread disinformation, from commitments that require signatories to improve the transparency of political advertising, to calls for more robust detection and removal of automated accounts (bots) that contravene the platforms' policies.

As part of a larger project designed to identify and expose malign information campaigns targeting the European parliamentary elections, ISD, alongside a number of other independent organisations studying the spread of disinformation, set out in spring 2019 to assess as best it could the enforcement of these commitments around the European parliamentary elections. In addition to the work of teams from Avaaz, Counter Action, Who Targets Me and the Mozilla Foundation, among others, ISD used its detection and analysis of malign information operations during the election campaign to map the successes or failures of the CoPD. By synthesising our combined research findings about disinformation, we have evaluated the responses of signatory companies across three areas of the CoPD:

1. Political advertising and issue-based advertising
2. Integrity of services
3. Empowering the research community

The CoPD includes commitments to scrutinise of ad placements and empower consumers. The scope of ISD's research efforts during the European parliamentary elections did not include assessing companies' advertising placement businesses, or the information curation and prioritisation systems that are addressed in these two sections of the Code. ISD has therefore not assessed these commitments as part of this evaluation. Relevant research on the placement of ads by some of the CoPD signatories has been conducted by groups such as the Global Disinformation Index.[3] Restricted data access has limited the research community's ability to assess changes to information curation and prioritisation systems at scale. ISD's recommendations for transparency in this paper seek to address this gap in order to enable more comprehensive assessment of information sorting algorithms in the future.

We summarise the results of ISD's attempts to evaluate the enforcement of the CoPD and provides a set of recommendations, which seek to inform the continuing process in place to design counter-disinformation co-regulation at the EU level through the upcoming Digital Services Act and European Democracy Action Plan in 2020 and beyond. We look at the activities of Facebook, Twitter and Google, as these platforms were the main focus of ISD's research into malign information operations during the European parliamentary elections, the findings of which are laid out in the accompanying report, *Click Here for Outrage: Disinformation in the European Parliamentary Elections 2019*.

Overall, our findings underscore the limits of self-reporting by the tech companies, as well as the insufficient progress made in a number of key commitments laid out by the CoPD. The recommendations propose more effective and sustainable regulatory approaches to protect not only elections but democracies writ large from covert information attacks, both foreign and domestic. They point to the opportunity presented by the forthcoming Digital Services Act for the EU to adopt a systemic approach to regulation, one that enables the meaningful oversight of technology platforms while ensuring human rights are upheld, and lays responsibility for the online harms that are enabled and amplified by tech platforms and their products at the feet of the platforms themselves.

# How Effective Were Technology Companies at Enacting their Commitments?

## Political and Issue-based Advertising[4]

One of the easiest ways for companies to improve their practices around potential disinformation activities is to improve the transparency of their political advertising online. In Europe, political adverts are already stringently regulated in a variety of ways, with rules in some countries limiting political advertising altogether in the months preceding an election.[5] By the companies' own admission, revenue from political advertisements is negligible compared to other forms of advertising. In 2016, Republican nominee Donald Trump and Democratic nominee Hillary Clinton spent a combined $81 million on Facebook ads throughout their campaigns. That accounted for about 0.3% of Facebook's revenue that year.[6] The European share is even smaller.

Signatories to the CoPD committed to:

- clearly distinguish ads from editorial content
- enable public disclosure of political advertising (defined as advertisements advocating for or against the election of a candidate or passage of referenda in national and European elections)
- use 'reasonable efforts' to at least start the public disclosure of 'issue-based advertising'.

This was no small undertaking across 28 countries and almost as many languages. But Facebook, Twitter and Google had each already taken significant steps towards setting up the infrastructure required for these types of changes in 2018, through the advertising archives established in the US, as well as subsequent additions in the UK and Brazil on Facebook.[7] However, similar transparency features were not available in Europe (with the exception of the UK) before spring 2019, with the spend, purchasers and targeting criteria of political ads remaining opaque to users.

To evaluate the effectiveness of companies' efforts to ensure political advertising transparency, ISD supplemented the findings from our own research with insights from partners at Mozilla Foundation and Who Targets Me, as well as public reports from the Office of the French Ambassador for Digital Affairs, all of which conducted some level of evaluation of the advertising archives available for the election.[8]

In advance of the elections, it was reported by a number of groups that Facebook had blocked access to existing tools that allowed users and journalists to see how they were being targeted by political actors while using Facebook.[9] Such tools had been developed by organisations like ProPublica and Who Targets Me. It was in this context that companies were asked to improve their transparency for political advertising in the EU.

The resulting efforts were varied. While progress was certainly made across the EU, despite steps already taken towards the disclosure of political ads by the major platforms in 2018 and 2019, efforts to implement useful and accurate transparency at scale were found to be lacking in the lead up to the European parliamentary elections, from both a technical and conceptual standpoint:

- Facebook provided the timeliest as well as the broadest coverage of political and issue-based advertising across the EU, but with poor enforcement and often difficult access and analysis processes for researchers and the public.
- Google's advertising application programming interface (API) was well designed for researchers, but the database contained very limited types of advertising and was released only shortly before the elections were conducted.
- Twitter's approach was also limited in scope, though the platform took drastic steps in the months since the election to ban all political advertising.

Those recent policy decisions are not part of the scope of this study, but broadly continue to face the challenges of definition and enforcement that were encountered by platforms during the European parliamentary elections.

Mozilla's comparative assessments of the Facebook and Google political advertising APIs assessed the tools from a functional standpoint, demonstrating the significant limitations of Facebook's approach for researchers and the public, as well as the delays and drawbacks of Google's efforts.[10] ISD's research team uncovered examples of both false positives and false negatives in the Facebook political ads library, and provided insights into the difficulties faced by researchers in using the libraries and APIs for live disinformation research during the elections.

The Office of the French Ambassador for Digital Affairs conducted its own evaluation of the Facebook political advertising API in France. In the French context, where political advertising is illegal in the run-up to an election, transparency from digital services is critical to understanding potential illegal activity during election campaigns. The report points to the potential power but initial technical limitations of Facebook's efforts around the European parliamentary elections.[11] The experiences and insights of these teams are compiled below, building on initial briefing notes provided to the EU Commission in a joint submission document from ISD and Avaaz.[12]

## Commitment: Make Ads Clearly Distinguishable

All the platforms studied have basic provisions in place that allow a user to distinguish between a paid ad and unpaid content. The formats of these provisions differ across platforms and products, but are included for video-based ads or image and text ads that appear in search results or on social media platforms. ISD also analysed if and how the disclosure of political ads was different from other ads on the platforms, which highlighted some specific issues, which are discussed below.

**Facebook:** Despite the explicit commitment made to ensure ads were 'clearly distinguishable from editorial content' online, Facebook political ads that were shared by a user were no longer marked with the 'paid for' line that is present on the original version of the ad. This rendered those shared copies of political ads less transparent by obscuring the funding behind the ad, and caused

**CODE OF PRACTICE COMMITMENTS: POLITICAL ADVERTISING**

'Signatories commit to keep complying with the requirement set by EU and national laws, and outlined in self-regulatory Codes, that all advertisements should be clearly distinguishable from editorial content, including news, whatever their form and whatever the medium used. When an advertisement appears in a medium containing news or editorial matter, it should be presented in such a way as to be readily recognisable as a paid-for communication or labelled as such.'

'Relevant Signatories commit to enable public disclosure of political advertising (defined as advertisements advocating for or against the election of a candidate or passage of referenda in national and European elections), which could include actual sponsor identity and amounts spent.'

'Relevant Signatories commit to use reasonable efforts towards devising approaches to publicly disclose "issue-based advertising". Such efforts will include the development of a working definition of "issue-based advertising" which does not limit reporting on political discussion and the publishing of political opinion and excludes commercial advertising.'

them to appear more like organic content than the original political ad. An investigation by Mother Jones revealed the loophole in May 2019,[13] and noted that the feature also removed a user's ability to click through to the advert in the advertising library, removing any other transparency information that should otherwise be available for a political ad.

**Google:** In general, all paid ads on Google Search and YouTube are marked with an 'ad' label to distinguish them from unpaid content. Google reported to the Commission that in 'some ad formats', the company 'adds a built-in disclaimer based on the advertiser's verification information. For other ad formats, the advertiser is responsible for incorporating the disclaimer into the ad on their own.'[14] The reporting to the Commission does not clarify which kinds of ad format include automatic disclaimers and which do not, rendering it difficult to determine how many political ads are signposted as such on Google's platforms. It is therefore still unclear on which products and platforms Google enables users to distinguish political ads from other types of ad through the disclaimer.

### Commitment: Ensure 'Public Disclosure of Political Advertising'

This section focuses on the specific request for companies to provide transparency on ads that advocated 'for or against the election of a candidate or passage of referenda in national and European elections', as per the Code's definition. All companies implemented some level of transparency for political ads, narrowly defined, before the vote in May 2019. This included disclaimers on political ads and ad archives, libraries or APIs that provided access to collections of political ads run on the platforms.

Some of the challenges that researchers at ISD and partner organisations faced in accessing and using political ads data from the companies are listed below:

Data access:

- **Lack of targeting and engagement data:** Across platforms, users are not provided with information to help them understand how they are being reached through advertising and why. Researchers could not tell whether an ad is particularly popular and potentially getting additional free promotion through shares on any of the platforms reviewed.
- **Placement information missing:** Google did not specify whether adverts were placed on YouTube or displayed next to search results on Google Search. Twitter's Ads Transparency Centre seemingly only included results from ads shown through 'promoted tweets', with no ads hosted through 'in-stream video' available in the library, or not labelled as such.[15]
- **Irregularities in the returned search results:**Researchers were hard-pressed to use the Facebook ad archive's aggregated data to assess real trends, as search results changed when queried at different times, and items appeared and disappeared from the library seemingly without explanation. Mozilla's evaluation of the ad archive provides more details on the issues faced by researchers in this regard,[16] as does the French Ambassador for Digital Affairs' analysis of the Facebook system's 'poor data integrity'.[17]
- **Limits to bulk-access data requests:** A user of the Facebook API is allocated on average 171 API requests per country per day. However, there were 751 seats contested during the elections, bringing the total number of seats and candidates alone far above the query limit for the API.[18] The API returned only 25 ads per page by default, and each request for an additional page counted against a user's limit of API requests, making more comprehensive, pan-European analysis impossible. On Twitter, researchers in Europe could not download data directly from the Ad Transparency Centre.
- **Keyword interface on Facebook and Twitter:** The public-facing portal for the Facebook ad archive limits complete searches per country or per advertiser, instead relying on a keyword search to define results. As

it is impossible to predict ahead of time the complete and precise set of words and variations used by all advertisers in all countries in all languages, researchers had to develop ad-hoc workarounds to access complete datasets of political ads. For members of the public, this obstacle was even more limiting. On Twitter, data cannot be downloaded from the Ads Transparency Centre, with analysis therefore tied to the user interface on the site.

- **Deleted content:** Ads removed from the Facebook and Twitter ad archives were deleted, either by the platforms or by the advertisers, and disappeared from view for researchers. This disabled researchers' ability to perform accurate retrospective analysis and made it easier for those spreading disinformation through political ads to hide their tracks. In the French context, this was particularly problematic in preventing oversight of illegal campaign ads in the week before the election: the report by the French Ambassador for Digital Affairs' office states that Facebook 'removed 31% of ads in the French library over the week of the European parliamentary elections', which included 'at least 12 ads that were illegal under French law'.[19] Advertisers on Twitter could delete their ads and the record of the ad would automatically be removed from the Ads Transparency Centre.[20]

Enforcement:

- **False negatives:** ISD's research found many political ads were not disclosed as political and failed to be captured by Facebook's transparency systems, across numerous countries. Examples identified by ISD were reported to Facebook and provided to journalists, including examples from Germany and Italy (see Figure 1).[21]
- **Late timing for roll-out from Google:** Little research could be conducted on the use of Google's advertising platforms for malicious political advertising or disinformation through political advertising because of the short timeframe before the elections where researchers or the public had access to transparency information. The ad transparency report and associated searchable library were launched on 2 May 2019.[22]
- **Mandatory registration process failures:** The processes put in place by Facebook to ensure that foreign sources could not buy political ads in another EU country were shown to be insufficient to stop even basic attempts to subvert them. This was demonstrated by a Bits of Freedom investigation in the Netherlands, where journalists posed as foreign nationals and yet were still able to purchase political ads targeting Dutch users. On Twitter, certification processes remained unclear outside the US.[23]


**FIGURE 1**
Two examples of pro-AfD political ads on Facebook running in April 2019 that were not disclosed as 'political' ads

## Commitment: Build Transparency Products for 'Issue-Based Advertising'

Adverts that mention a specific candidate, party or electoral process constitute only a small part of the paid political communications leading up to any election. The Code included the development of transparency for issue-based advertising relating to politics around the elections, in addition to the strict definition of political ads referencing candidates or parties.

Given the findings of research from ISD, Avaaz and others around disinformation in the European elections 2019, which demonstrated the importance of wedge issues to those wishing to use social media for malicious and deceptive purposes, the inclusion of transparency on issue-based ads is an important, if ambitious, undertaking in the context of an election. The difficulties in defining 'issue-based advertising' point to the need for leadership from governments or election regulation bodies in establishing clear parameters for companies to work from in initiating transparency for issue-based ads.

An evaluation of the companies' efforts on this front found the following:

**Google and Twitter's efforts were limited in scope:** Google did not include issue-based ads in its library. Twitter did not define issue-based ads in the EU for the elections, focusing only on direct political ads mentioning parties or candidates.

**Definitional issues:** Facebook included issue-based ads in its library, though the choice of issues for inclusion was broad and Facebook used the same criteria for all of Europe. The issues chosen were immigration, political values, civil and social rights, security and foreign policy, economy, and environmental politics. There was no tailored country approach to defining 'issues', so certain subjects pertinent in specific European countries in the lead up to the election were missed. However, Facebook was ambitious in undertaking to attempt to address issue-based advertising across all 28 countries of the EU before the election, and its measures went a step beyond the transparency efforts of Google and Twitter.

**False positives:** The broadly defined issue areas used to label ads on Facebook led to the inclusion of false positives in the political advertising library, including adverts relating to video games and shopping. Despite its limited definition of political ads, Twitter's Ad Transparency Centre included a number of false positives, stretching far beyond specific political content. The French government's evaluation of Twitter's ad library found examples of windshield dealers and pizzerias in the political ad centre.[24]

## Improve Integrity of Services[25]

**CODE OF PRACTICE COMMITMENTS: INTEGRITY OF SERVICES**

'Relevant Signatories commit to put in place clear policies regarding identity and the misuse of automated bots on their services and to enforce these policies within the EU.'

'Relevant Signatories commit to put in place policies on what constitutes impermissible use of automated systems and to make this policy publicly available on the platform and accessible to EU users.'

As the findings of ISD and Avaaz research about disinformation in the EU elections demonstrated, most visible examples of disinformation are not simply false content.[26] Instead, misrepresentation of sources, communities and popularity lie at the heart of most of the disinformation tactics detected. Commitments to dealing with inauthentic and misrepresentative behaviour are therefore at the heart of tackling the reality of disinformation at present.

The CoPD limited the commitments around this issue to identifying and removing misused 'bots'. While this is an important part of addressing threats of disinformation activity online, it fails to consider the range of misrepresentative tactics deployed by those seeking to deceive users; ISD and its partners also assessed responses to suspected sock-puppet accounts and pages or groups misrepresenting communities. Furthermore, the definitions of the CoPD did not consider the increasing sophistication of automated, semi-automated and co-ordinated manual networks of accounts deployed by deceptive actors. Civil society, research organisations and academic institutions use a range of thresholds to classify 'bot' activity; these behaviours can often be hard to distinguish from hyperactive or spam-like but human accounts.

Overall, since signing up to the EU's Code of Practice, platform operators claim significant progress in acting against inauthentic behaviour to protect the integrity of their services. According to the EU Commission's *Report on the Implementation of the Action Plan Against Disinformation*, Facebook disabled 2.19 billion fake accounts, including many targeting Europe, YouTube removed

over 3.39 million channels for violating its spam and misleading content policy, and Twitter challenged 77 million fake accounts, all between January and April 2019.[27] The relevance of these efforts to the European parliamentary elections is difficult to establish as companies bulk report statistics of account removals or channel removals. Moreover, despite these enormous numbers of accounts being removed, evidence from ISD's research, combined with findings from partner organisations, demonstrates poor and sporadic enforcement of policies to protect services against unlabelled bots, sock-puppet accounts and accounts or channels consistently spreading disinformation.

Even when suspicious pages and groups were reported, platforms generally acted slowly to respond, or were opaque in their response to reported violations. This runs not only against the commitment to identify and remove 'bots', but also against the commitments to ensure transparency for researchers (see below).

The lack of shared language, thresholds or methods for detecting covert automation on different platforms is a challenge for both companies responding to researcher or civil society claims and those conducting research. These issues point to the need for companies to develop more effective channels for communication between their enforcement teams and those attempting to identify covert automated activity on their platforms in civil society and academia, with assurances of responsiveness and trasparency over if and how they respond to investigations and claims of covert automation on their platforms.

### Identifying and Removing Inauthentic Accounts and Bots

**Spain – Twitter**
In Spain, ISD reported to Twitter a network of over 2,000 accounts seemingly co-ordinated from Venezuela, which were boosting anti-Islam hashtags and amplifying support for Vox, including a mixture of bots and inauthentic accounts.[28] Despite in-depth reporting of the network to Twitter, only 39 of the accounts have since been removed. Many have not been active since the elections in April and May 2019, despite identified hyperactivity during the election campaign period. Twitter staff were communicative about how they analysed and responded to the network over email, and engaged in a timely and direct conversation with ISD on the methods they used and the different classifications employed to identify bots or inauthentic accounts.

**Poland – Facebook**
In Poland, ISD identified a suspected co-ordinated network of pages, accounts and groups on Facebook used to promote nationalist party Konfederacja and to amplify anti-Semitic and pro-Kremlin content.[29] This network included 60 pages with a total of 194,675 followers and 5 groups with a total of 23,187 members. Facebook did not respond to requests to communicate what action, if any, was taken on this network.

**UK – Twitter**
In the UK, ISD demonstrated that support for most major parties on Twitter was boosted by bot-like accounts showing activity levels, activity patterns and profile features indicative of automated accounts. Nearly half (42%) of the most active accounts supporting official party Twitter handles showed signs of bot-like, hyperactive behaviour with largely anonymous profiles. In May 2019 it was reported that 8 of the top 10 most prolific accounts engaging with the Brexit Party on Twitter showed signals of bot-like activity.[30] Six months later, in October 2019, only two of these accounts had been suspended.

**Avaaz Research – Facebook**
Avaaz uncovered disinformation networks in France, UK, Germany, Spain, Italy and Poland, posting content that was viewed an estimated 763 million times between January and April 2019.[31] Avaaz reported to Facebook a total of nearly 700 suspect pages and groups, followed by more than 35 million people and

generating over 76 million 'interactions' (comments, likes, shares) during the three-month research period. Ahead of the European elections, Facebook took down 132 of the pages and groups reported, accounting for around 30% of all interactions across these networks, and 230 suspicious profiles. The content racked up an average of 6 million views per day in total while it was live. Some examples of the findings include: 60 pages and groups identified in France (one network spreading disinformation and others posting dehumanising, racist and white nationalist content targeting migrants); and 14 networks discovered in Italy, many supporting the League and the Five Star Movement. Again, these were not all involved in spreading false information, with some posting divisive anti-migrant or hate content.

Germany – Facebook
Professor Trevor Davis, George Washington University, studied a network of 200,000 pro-AfD accounts in Germany spreading electoral content, which were seemingly co-ordinated and misrepresentative.[32] It is unclear whether the accounts under study were reported to Facebook or if any action was taken following the investigation.

Where networks of inauthentic activity were identified and removed by platforms, data on their reach, audiences and activities was limited, and rarely contextualised to ascertain their relevance to European audiences during the election campaign. The potential impact of those activities on EU electorates was therefore impossible to ascertain. For example, in their March 2019 report on the implementation of the Code of Practice, Facebook reported on taking down eight co-ordinated inauthentic networks originating in North Macedonia, Kosovo and Russia, but provided very minimal evidence to help readers understand the nature of engagement with those networks or the audiences that they reached, and there was therefore little visibility on whether these networks were targeting the EU or the electoral process.[33]

Google did not report in detail to the EU on the status of their efforts to counter abusive account creation and abusive engagement between December 2018 and the election in May 2019. In its March monitoring report, the EU Commission noted that more up to date information is required from Google about how many inauthentic accounts or bots are functioning on their platforms and which countries they are targeting, as well as the types of content being shared and number of people being reached by these accounts.[34]

## Empower the Research Community[35]

The Code of Practice included a commitment for companies to take 'reasonable measures' to enable appropriate access to data for fact-checking and research activities by academics. The combined work of ISD, Avaaz and Mozilla across the campaign period points to clear and continuing gaps in companies' support for those seeking to detect, call out and respond to disinformation online.

Steps have been taken in the right direction, but only incrementally and far more slowly than the pace at which disinformation is evolving as a threat to democracies. Twitter has consistently provided users and researchers with data access to enable computational analysis of trends on the platform. This is a somewhat simpler task for a platform that is, for the most part, a public space where users have an expectation of low privacy. Facebook has begun to provide access to CrowdTangle, a tool enabling computational analysis of public page and public group admin activity on the platform. This is a very welcome step, but access remains a privilege that only some academic, research and media organisations are provided with and comes with some technical limitations. YouTube is still a relatively opaque and difficult platform to study for researchers, despite its critical role as a source of news and information for millions of Europeans.

**CODE OF PRACTICE COMMITMENTS: THE RESEARCH COMMUNITY**

'Relevant Signatories commit to support good faith independent efforts to track Disinformation and understand its impact, including the independent network of factcheckers facilitated by the European Commission upon its establishment. This will include sharing privacy protected datasets, undertaking joint research, or otherwise partnering with academics and civil society organizations if relevant and possible.'

'Relevant Signatories commit not to prohibit or discourage good faith research into Disinformation and political advertising on their platforms.'

'Relevant Signatories commit to encourage research into Disinformation and political advertising.'

'Relevant Signatories commit to convene an annual event to foster discussions within academia, the fact-checking community and members of the value chain.'

There are undoubtedly challenges in developing safe and transparent processes for anonymised data-sharing at scale. The increasing frustrations around the Social Science One project, designed to provide vetted academics with access to large anonymised datasets to study disinformation on Facebook, demonstrate that there is still limited co-operation between companies and independent researchers in this area.[36] However, companies sell detailed data about users to advertisers legally and widely in Europe and beyond. This is a central part of the business model of social media and advertising platforms. The argument that privacy-protective data-sharing with vetted research organisations is an impossibility does not hold up in that context.

During the European elections, there was limited access to appropriate public data for research, and to swift and timely communication and transparency from platforms to researchers about emerging disinformation threats. Companies worked sporadically and inconsistently with external research organisations to react to activity flagged as potentially malign on their platforms. When ISD reported suspicious networks to the platforms during the campaign, Twitter was responsive to direct email communications, while Facebook's teams were not. Avaaz staff found that Facebook teams only responded to their requests after significant and sustained human effort from the researchers over time.

A joint statement by then Commissioners for Justice, Consumers and Gender Equality Věra Jourová, former Commissioner for the Security Union Julian King, and former Commissioner for the Digital Economy and Society Mariya Gabriel on the one-year anniversary of the CoPD states:

> *While progress has been reported on the commitments monitored by the Commission from January to May ahead of the 2019 EP elections, less is reported on the implementation of the commitments to empower consumers and the research community. The provision of data and search tools is still episodic and arbitrary and does not respond to the needs of researchers for independent scrutiny.*[37]

ISD's research and response to detected instances of malign activity online during the elections confirmed this assessment.

## Sharing Privacy-Protected Data

The expectation of companies to share privacy-protected data was and is ambitious. The critical need to ensure effective privacy protections is no small task. However, apart from Twitter, major tech companies signed up to the Code of Practice have rolled back on providing even basic access to public data or metadata, let alone to data requiring stringent privacy protections. This not only hampers detection efforts, but also impedes evaluation of the impact of current company responses to disinformation. The Code of Practice as a whole is hard to evaluate when access to data is so limited.

Overall, while Twitter continues to provide relatively thorough access to data for researchers and the public through its API, Facebook and YouTube continue to fall short in their provision of accurate and timely data for researchers studying disinformation. Failures were identified in three areas:

- sharing data and/or information to support detection efforts
- sharing data and/or information on any company responses to disinformation reported to platforms
- retrospective data on the activities, reach and engagement with disinformation networks identified by platforms during the European elections.

# Recommendations

**The EU faces a significant task in 2020: it has the opportunity to forge an approach to digital regulation that can help protect users from a range of online harms while also bolstering human rights to free expression and information.**

There are few other institutions with the opportunity to lead this charge: governments with lesser respect for international human rights have already put a stake in the ground on their vision of the information sphere under authoritarian leadership; private sector companies, largely based in Silicon Valley, have no democratic mandate and little incentive to make themselves accountable or transparent to their users.

The EU has examples and lessons to draw on in considering its own design for digital regulation, with many of its member states (and recent ex-member state) recently rolling out their own national approaches on the matter, with varying levels of success. Germany has trialled sanctions for illegal content remaining on platforms once flagged, on the one hand, and approaches tying together competition law with data privacy obligations to enforce higher standards of responsibility on large technology companies on the other. The UK is working towards a draft bill for a statutory duty of care for online harms that attempts to grapple with the systems and processes that underpin the use of technology for harmful ends. France has discussed radical transparency as a central principle for digital regulation, but has simultaneously moved ahead with content-specific laws banning 'disinformation' around elections and imposing time limits for the removal of hate speech and terrorist content.

Regulatory approaches alone will not suffice in dealing with the breadth of harm directed at users online: investment in sustainable research on detecting and analysing these threats is critical; the empowerment of individuals to shape their own digital experiences should be a central tenet of the European plan of action; and education for those old and young on how they can use online expression to help prevent and counter harm will doubtless be necessary to create a more equitable and enjoyable online environment for all. While ISD has provided recommendations and tested programmes of work in each of these areas, the recommendations in this paper focus on the lessons learned from the CoPD and their relevance for potential future versions of the Code or parallel regulation designed through the Digital Services Act.

In its convening of research and civil society actors in this space, ISD has found that civil society groups concentrated on the protection of free expression and groups focused on identifying and addressing specific types of online harm have begun to find some interesting common ground in their requests of governments and platforms. In line with this emerging consensus from the research and civil society community in Europe, ISD recommends there should be a move away from siloed regulatory responses attempting to address each type of illegal activity or breach of rights separately. Instead, ISD calls for regulation requiring transparency and accountability for the processes and systems that order, curate, promote, target, amplify or, in many cases, profit from user-generated content. Transparency over the decisions made

by technology companies (either by humans or through automation) is a prerequisite to understanding how these products are involved in promoting illegal or harmful activity, and therefore how to potentially mitigate these negative externalities through regulation. Democratic governments must move from efforts largely centred on enforcing content removals in relation to specific types of harm towards a joined up approach dealing centrally with the means of distribution and decision-making of content-hosting platforms.

Recommendations for that approach are laid out below, concentrating on three key principles: the requirement for regulation; the centrality of transparency to any regulation agenda; and the need for investment in a multi-stakeholder system, including independent research, to provide evidence of the evolving threat and to hold policymakers and companies to account on their responses.

## 1. Voluntary Self-regulation Is Not Enough

Across the past five years, governments have attempted to nudge tech companies to take on more responsibility and accountability for illegal or legal but potentially harmful activity on their platforms, including terrorist materials, hate speech and disinformation, all of which ISD has researched in detail. Largely as a consequence of this approach, actions from companies have been almost entirely reactive and siloed between issue areas, leaving the overarching structures and processes that promote, amplify or recommend harmful content or enable and encourage harmful activities opaque and untouched. As a result, in Europe, representatives from the major tech companies themselves have publicly asked national governments and the EU Commission for a clearer set of definitions and regulatory standards that they are able to follow.

The Code of Practice was a formalised and well-documented attempt at voluntary self-regulation. It laid out requests and standards in more detail than many previous efforts (though the level of detail varied from section to section, with advertising commitments clearer and detailed at greater length than the parallel commitments for 'integrity of services'). These standards were not accompanied by the threat of sanctions or financial penalties in the event of the signatories' failure to comply with the commitments. Instead, the CoPD relied entirely on the will of the companies to institute changes.

The Code of Practice was therefore a litmus test for whether a robust and formalised attempt at self-regulation was enough to compel tech companies to change fundamentally their policies and practices on disinformation. It was a test that needed to be undertaken, if only to confirm the growing sense among researchers and many European governments that self-regulation has proved inadequate as an approach to protecting citizens from illegal or harmful conduct online.

Overwhelmingly, the Code of Practice proved the limits of voluntary efforts to bring about systemic change from the signatory companies. On the one hand, companies were able to work with the Commission to craft and sign up to substantial commitments to issues where requests were defined and the scope relatively contained. However, the follow through and enforcement of these commitments was variable and hard to track for external researchers and governments alike owing to the lack of data for evaluation and verification. Reactive, largely opaque processes remained for the removal of inauthentic account networks or the detection of covert automated accounts. Outside the limited sphere of political advertising, little data was made available for researchers studying disinformation, aside from Twitter's efforts at transparency for public social media data, in place long before the Code of Practice was adopted.

The lessons learned from the Code of Practice point to the need for EU leadership on a regulatory agenda for Europe to hold technology companies

to account for the negative externalities that arise from their products, policies and practices. It is not just the research community calling for such action: recent public opinion research conducted by UK think tank Doteveryone found that most members of the public surveyed about their attitudes towards people, power and technology believe the technology sector to be under-regulated.[38] Recommendations for the focus of that regulation follow here.

## 2. Transparency Must Be the First Principle, But It Must Be Rigorous and Regulated

Transparency is commonly used as a comprehensive solution to improving visibility and understanding, and providing accountability for disinformation and a wealth of additional online harms. It is of central importance that governments, civil society and the public are able to understand better the ways in which the internet is impacting society and democracy in order to encourage its positive effects and curb negative externalities. But what does transparency really mean and how could it be used as a practical tool to enhance users' safety and understanding of online information?

The requirements and expectations associated with transparency are often poorly articulated. ISD has drawn out four central areas where transparency should be prioritised by any future regulators:

- content and content moderation
- advertising
- complaints and redress
- algorithms.[39]

### Content and Content Moderation

Platforms that have become public spaces must make that space as intelligible as possible. As web platforms and their users play an increasing role in shaping our culture, informing our political decision-making, and driving societal change, the activities taking place in these spaces should be observable.

### Advertising

Advertising – particularly political advertising – has been shown to be a key vector through which the public can be manipulated. It is in the public interest for internet users to understand how and why they are being targeted online, and for regulators to be able to understand and respond to malpractice. The Code of Practice made significant moves in this direction, but there is work to be done in ascertaining the scope, usability, accessibility and accuracy of transparency data on advertising from companies. Definitional clarity on what political advertising means and any subsequent transparency requirements made of companies are a job for democratic governments' independent regulatory bodies, including election commissions where relevant.

### Complaints and Redress

A significant gap exists in the public's understanding of platforms' ability to moderate and respond to abuses of their platforms. Visibility of complaints made to platforms is essential to accountability, to support the victims of online harms, to raise awareness of challenges facing users online and to provide evidence in redress.

### Algorithms

There is significant concern that platform architectures contribute to negative outcomes. Central to this disquiet is the fact that the algorithms dictating a user's experience and journey have led to unintended consequences and have been challenging to scrutinise or evaluate.

A recent collaborative briefing from ISD, Open Rights Group, Doteveryone, Demos, Global Partners Digital and Digital Action lays out a number of viable approaches for algorithmic inspection and the associated transparency requirements.[40] It is feasible to design regulatory oversight for the decision-making systems that have a structural impact on rights to information and of free expression. Standards for this type of oversight already exist. The briefing includes suggestions for:

- consensual and warrant-based models for algorithmic inspection by a regulator, referencing existing models used by bodies such as the Information Commissioner's Office and Investigatory Powers Commissioner's Office in the UK
- independent expert third party audit powers, whereby a regulator could instruct an expert to undertake an audit, referencing existing models used by bodies such as the Financial Conduct Authority in the UK
- the role of a regulator in ensuring that commitments made by companies to share data with accredited academics and researchers are upheld.

Each of these areas of transparency requires a different set of frameworks. Detailed technical recommendations for each of these can be found in ISD's 2019 consultation response to the UK's *Online Harms White Paper*.[41] However, drawing these areas together, there are two broad principles of digital transparency underlying each: transparency must be computational and must complement rights to data privacy, not erode them.

### Transparency Must Be Computational

For an online space to be transparent, it must be possible to observe it computationally. For instance, Twitter's API allows for a holistic view of what takes place on that platform. Were that API not to exist, an otherwise nominally 'public' platform would not be transparent as a direct result of its scale – it overwhelms human capacity.

### Transparency Must Complement Rights to Data Privacy, Not Erode Them

A good model for transparency will protect individuals' data privacy while enabling a macro understanding of the nature and scale of technology platforms' processes and any potential infringement of rights that stems from the use of the platform. Access to transparency data may be contentious. Although we believe that it is in the public interest to have oversight over all four areas, it is possible that there should be exceptions to the types of data and access available to the general public. A 'tiered' access structure (by which a regulator or institutions accredited by a regulator or other body have increased access to transparency data) may be advisable in light of data protection and privacy expectations. However, we believe the starting point should be public access. Existing proposals from researchers and expert institutions should be tested as additional safeguards for data acquired through transparency regulation, including the model of data trusts currently being piloted by the Office for Artificial Intelligence and the Open Data Institute in the UK.

Good models for transparency exist and should be used by the EU as best practice when developing the Digital Services Act as a framework for regulating online platforms.

## 3. Invest in a Multi-Stakeholder System, Including Independent Research, to Hold Policymakers and Companies to Account on their Responses

The scale of digital media is so large that no government, institution or technology company can possibly detect and understand all disinformation

threats alone. Neither can any one institution alone assess all the measures taken to respond to these threats and their possible impacts. Technology company self-assessments are only of limited use as a tool for charting progress in mitigating threats from disinformation, as there are few incentives for companies to admit gaps or mistakes. A future regulator could certainly take on a significant weight of such efforts to hold platforms accountable but may well lack the subject matter expertise required to identify evolving disinformation threats. Regulatory responses themselves require oversight and checks from independent experts, in order to ensure they are not themselves stifling human rights or proving ineffective in dealing with the threat.

There is therefore a critical role for independent researchers from academia and civil society to hold technology companies and governments accountable for responses to disinformation online. As the threat picture of disinformation changes, there will be a role for them to identify new kinds of threat and to use that insight to inform better policy responses. As responses to disinformation evolve, there will be a role for them to ensure the accountability of governments and that companies respond effectively and with respect for human rights.

Independent experts will require access to relevant data to make such assessments. Any future regulatory efforts from the EU must include provisions securing access to privacy-protected data for vetted independent research organisations. The transparency principles presented above lay out the parameters of providing data safely and responsibly and the four areas in which that is required.

The Social Science One experiment attempted to share large datasets relevant to the study of disinformation safely without compromising either user privacy or business competition. Lessons must be learned from that effort in order to avoid running into similar obstacles to progress:

- The legal and technical complexity of data-sharing efforts should not be underestimated.
- The effort has led to the development of a new framework for 'scholarly and ethical review of networked data research', which provides industry standards for social media data research that should be considered in future data-sharing models.[42]
- The new statistical method for differential privacy developed by the scholars involved in the project should be considered in order to design data-sharing systems that can 'preserve the privacy of end users while enabling scholars to draw valid statistical inferences on the questions they are investigating'.[43]

In the near term, there are creative options for increasing researchers' capability in detecting disinformation campaigns that do not require regulatory design or government oversight. There is an opportunity here for technology companies to take the initiative. Technology companies know more than anyone about the types of signals that help detect co-ordinated disinformation efforts on their platforms. Opportunities should be explored to set up new kinds of collaborations, where technology companies produce 'dummy data' to artificially simulate examples of platform manipulation. This could protect both data privacy and business competition by creating invented disinformation scenarios, but would still improve knowledge-sharing with the independent research sector about methods and tools for detecting disinformation activity on specific platforms. The EU Commission should strongly encourage such collaboration between technology company threat intelligence teams and the academic and civil society research sector in order to move the field forward. It should also support mechanisms for better exchange of expertise and practice among European governments over digital policy and regulation, and foster pan-European research and analysis efforts and networks.

1 ISD, Click Here for Outrage: Disinformation in the European Parliamentary Elections 2019, June 2020.

2 Microsoft signed up in May 20 19, two days before polling day and therefore were not included in ISD's research methods.

3 https://disinformationindex.org/.

4 Based on contributing research from ISD, Mozilla Foundation and Who Targets Me.

5 French law prohibits paid advertising directly for political campaigning for six months preceding an election, for example see here.

6 Wagner, K., Facebook's Political Ad Business Is Lots of Pain and Little Gain', Bloomberg, 16 October 2019, https://www.bloomberg.com/news/articles/2019-10-16/facebook-s-political-ad-business-is-lots-of-pain-and-little-gain.

7 Facebook announced the launch of the US political advertising library in May 2018; Twitter and Google announced the launch of their US political advertising libraries in June 2018.

8 Ambassadeur Pour le Numérique, 'Facebook Ads Library Assessment', 2019, https://disinfo.quaidorsay.fr/en/facebook-ads-library-assessment; Ambassadeur Pour le Numérique, 'Twitter Ads Transparency Center Assessment', 1 Introduction, 2019, https://disinfo.quaidorsay.fr/en/twitter-ads-transparency-center-assessment#introduction.

9 Merill B. J. and Tobin, A., 'Facebook Moves to Block Ad Transparency Tools – Including Ours', ProPublica, 28 January 2019, https://www.propublica.org/article/facebook-blocks-ad-transparency-tools.

10 Mozilla, 'Facebook's Ad Archive API is Inadequate', The Mozilla Blog, 29 April 2019, https://blog.mozilla.org/blog/2019/04/29/facebooks-ad-archive-api-is-inadequate/; Mozilla, 'Google's Ad API is Better Than Facebook's, But…', The Mozilla Blog, 10 May 2019, https://blog.mozilla.org/blog/2019/05/10/googles-ad-api-is-better-than-facebooks-but/.

11 Ambassadeur Pour le Numérique, 'Facebook Ads Library Assessment'.

12 ISD and Avaaz, 'Disrupted: Evidence of Widespread Digital Disruption of the 2019 European Parliament Elections', June 2019, https://www.isdglobal.org/wp-content/uploads/2019/06/Joint_Submission-070619-letter.pdf.

13 Levy, P., 'The Facebook Loophole That Makes Political Ads Look Like Regular Content', Mother Jones, 17 May 2020, https://www.motherjones.com/politics/2019/05/the-facebook-loophole-that-makes-political-ads-look-like-regular-content/.

14 European Commission, 'Third Monthly Intermediate Results of the EU Code of Practice Against Disinformation, ec.europa.eu, 2019, https://ec.europa.eu/digital-single-market/en/news/third-monthly-intermediate-results-eu-code-practice-against-disinformation.

15 Ambassadeur Pour le Numérique, 'Twitter Ads Transparency Center Assessment', 2.4 Only a subset of political ads is shown.

16 Mozilla, 'Facebook's Ad Archive API is Inadequate'.

17 Ambassadeur Pour le Numérique, 'Facebook Ads Library Assessment', 2.8 Poor data integrity, 2019, https://disinfo.quaidorsay.fr/en/facebook-ads-library-assessment#poor-data-integrity.

18 European Parliament FAQs, 'How Many MEPs?', ec.europa.eu, 2020, https://www.europarl.europa.eu/news/en/faq/12/how-many-meps.

19 Ambassadeur Pour le Numérique, 'Facebook Ads Library Assessment', 2.8 Poor data integrity.

20 Twitter, 'Ads Transparency Center FAQs', nd, https://business.twitter.com/en/help/ads-policies/ads-transparency-center-faqs.html.

21 Scott, M., 'Facebook's European Election War Room', Politico, 5 May 2019, https://www.politico.eu/article/facebook-european-election-war-room-dublin-political-advertising-misinformation-mark-zuckerberg/.

22 Google, 'Google Annual Report', [2019], p. 13, https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=62680.

23 Martins, F., 'Facebook Lies to Dutch Parliament about Election Manipulation', Bits of Freedom, 21 May 2019, https://www.bitsoffreedom.nl/2019/05/21/facebook-lies-to-dutch-parliament-about-election-manipulation/.

24 Ambassadeur Pour le Numérique, 'Twitter Ads Transparency Center Assessment', 2.3 No comprehensive access to political ads, 2019, https://disinfo.quaidorsay.fr/en/twitter-ads-transparency-center-assessment#introduction.

25 Contributions from ISD and Avaaz.

26 ISD and Avaaz, 'Disrupted: Evidence of Widespread Digital Disruption of the 2019 European Parliament Elections', June 2019, https://www.isdglobal.org/wp-content/uploads/2019/06/Joint_Submission-070619-letter.pdf.

27 European Commission, 'Report on the implementation of the Action Plan Against Disinformation', June 2019, p. 4, https://eeas.europa.eu/sites/eeas/files/joint_report_on_disinformation.pdf.

28 Peinado, F., 'Una red de cuentas falsas de Twitter promueve a Vox en campana', El País, 26 April 2019, https://elpais.com/politica/2019/04/25/actualidad/1556203502_359349.html.

29 Ćwiklak, D., Sieć antysemickich trolli na polskim Twitterze wspiera Konfederację, Newsweek, 23 May 2019, https://www.newsweek.pl/swiat/polityka/siec-antysemickich-trolli-na-polskim-twitterze-wspiera-konfederacje/65nw7g6.

30 Smith, M. and Boyd, M., '8 of the Top 10 Brexit Party Promoting Twitter Accounts Appear To Be Bots', Mirror Online, 21 May 2019, https://www.mirror.co.uk/news/politics/8-top-10-brexit-party-16179913.

31 Avaaz, *Far Right Networks of Deception*, May 2019, https://secure.avaaz.org/campaign/en/disinfo_network_report/.

32 Davis, T., Livingston S. and Hindman, M., 'Suspicious Election Campaign Activity on Facebook: How a Large Network of Suspicious Accounts Promotes Alternative für Deutschland in the 2019 EU Parliamentary Elections', July 2019, https://smpa.gwu.edu/sites/g/files/zaxdzs2046/f/2019-07-22%20-%20Suspicious%20Election%20Campaign%20Activity%20White%20Paper%20-%20Print%20Version%20-%20IDDP.pdf

33 Facebook, 'Facebook March 2019 Monthly Update on Implementation of the Code of Practice on Disinformation', p. 6, https://ec.europa.eu/digital-single-market/en/news/third-monthly-intermediate-results-eu-code-practice-against-disinformation.

34 European Commission, 'Third Monthly Intermediate Results of the EU Code of Practice Against Disinformation', ec.europa.eu, 2019, https://ec.europa.eu/digital-single-market/en/news/third-monthly-intermediate-results-eu-code-practice-against-disinformation.

35 Contributions from ISD, Avaaz and Digital Action.

36 Alba, D., 'Ahead of 2020, Facebook Falls Short on Plan to Share Data on Disinformation', *New York Times*, 29 September 2019, https://www.nytimes.com/2019/09/29/technology/facebook-disinformation.html.

37 European Commission, 'Code of Practice on Disinformation One Year On: Online Platforms Submit Self-assessment Reports', 29 October 2019, https://ec.europa.eu/commission/presscorner/detail/en/STATEMENT_19_6166.

38 Miller, C., 'People, Power and Technology: The 2020 Digital Attitudes Report', doteveryone, 11 May 2020, https://www.doteveryone.org.uk/2020/05/people-power-and-technology-the-2020-digital-attitudes-report/.

39 ISD, 'Extracts from ISD's Submitted Response to the UK Government Online Harms White Paper', July 2019, https://www.isdglobal.org/wp-content/uploads/2019/12/Online-Harms-White-Paper-ISD-Consultation-Response.pdf

40 Demos, doteveryone, Global Partners Digital, ISD and Open Rights Group, 'Joint Report: Algorithm Inspection and Regulatory Access', April 2020, https://www.isdglobal.org/wp-content/uploads/2020/04/Algo-inspection-briefing.pdf.

41 ISD, 'Extracts from ISD's Submitted Response to the UK Government Online Harms White Paper'.

42 SSRC, 'A New Standard of Scholarly and Ethical Review for Networked Data Research', April 2018, https://www.ssrc.org/fellowships/view/social-media-and-democracy-research-grants/a-new-standard-of-scholarly-and-ethical-review-for-networked-data-research/.

43 Social Science One, 'Analyzing Data From Facebook', blog, 1 February 2020, https://socialscience.one/blog/analyzing-data-facebook; Evans, G. and King, G., 'Statistically Valid Inferences from Differentially Private Data Releases, with Application to the Facebook URLs Dataset', 16 May 2020, https://gking.harvard.edu/dpd.

## About the Institute for Strategic Dialogue

**We are a global team of data analysts, researchers, innovators, policy-experts, practitioners and activists - powering solutions to extremism, hate and polarisation.**

The Institute for Strategic Dialogue (ISD) is an independent nonprofit organisation dedicated to safeguarding human rights and reversing the rising global tide of hate, extremism and polarisation. We combine sector-leading expertise in global extremist movements with advanced digital analysis of disinformation and weaponised hate to deliver innovative, tailor-made policy and operational responses to these threats.

Over the past decade, we have watched hate groups and extremist movements deploy increasingly sophisticated international propaganda, influence and recruitment operations, skillfully leveraging digital technology, and often boosted by hostile state actors. Alongside an exponential spike in violence (conflict, hate crime, terrorism), societies around the world are being polarised. At ballot boxes, populists have made significant gains and authoritarian nationalism is on the rise. If left unchecked, these trends will existentially threaten open, free and cohesive civic culture, undermine democratic institutions and put our communities at risk. Progress on the major global challenges of our time – climate change, migration, equality, public health – threatens to be derailed.

We can and must turn the tide. Help us build the infrastructure to safeguard democracy and human rights in the digital age. We believe it is the task of every generation to challenge fascistic and totalitarian ideologies and to invest in reinforcing open, democratic, civic culture.

ISD draws on fifteen years of anthropological research, leading expertise in global extremist movements, state-of-the-art digital analysis and a track record of trust and delivery in over 30 countries around the world to:

1. Support central and local governments in designing and delivering evidence-based policies and programmes in response to hate, extremism, terrorism, polarisation and disinformation

2. Empower youth, practitioners and community influencers through innovative education, technology and communications programmes.

3. Advise governments and tech companies on policies and strategies to mitigate the online harms we face today and achieve a ‘Good Web’ that reflects our liberal democratic values

Only in collaboration with all of these groups can we hope to outcompete the extremist mobilization of our time and build safe, free and resilient societies for generations to come. All of ISD’s programmes are delivered with the support of donations and grants. We have the data on what works. We now need your help to scale our efforts.

**If we succeed in empowering just a small minority of the silent majority with the insights, knowledge and tools they need, we have won.**

**PO Box 75769 | London | SW1P 9ER | UK**

**www.isdglobal.org**