ISD | Institute for Strategic Dialogue

2019 EU Elections

Click Here For Outrage: Disinformation in the European Parliamentary Elections 2019

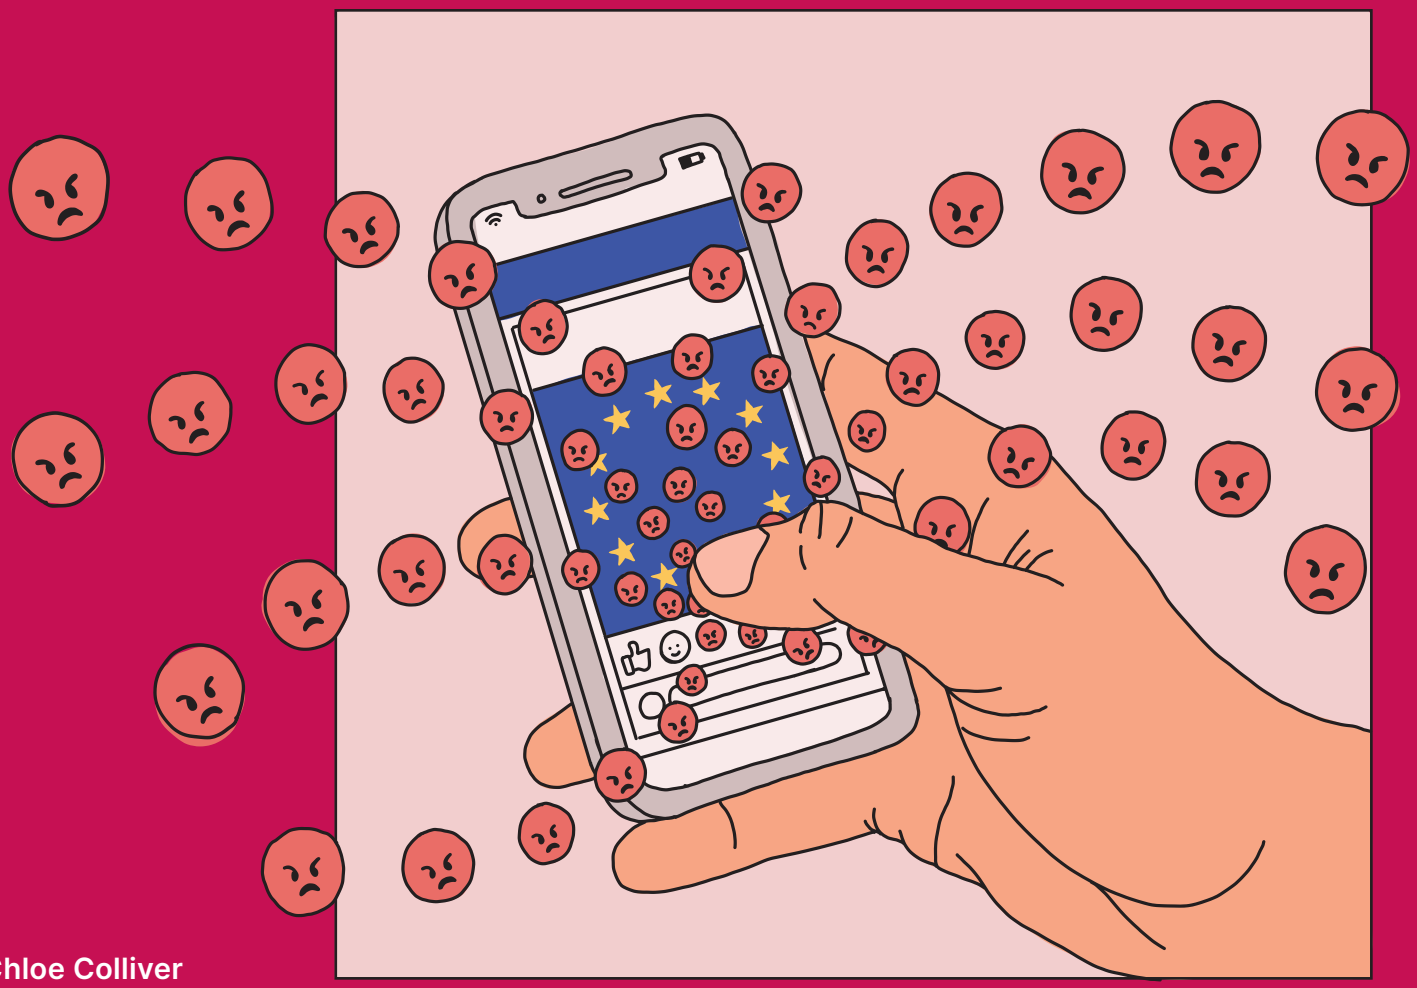

Chloe Colliver

Contributing Authors

Jakob Guhl
Cecile Guerin
Jacob Davey
Julia Ebner
Henry Tuck

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

The internet allows you to tell a new kind of lie. Intentionally trying to replace one belief with another is only the tip of the iceberg: tricks in the world of online deception take all shapes and forms.

Sometimes, it is one person pretending to be another. Other times, a whole community created out of nothing but software. It might look like popularity, but in reality, this is virtual popularity won through programmed clicks. The bar for entry is low, and getting lower.

In the aftermath of elections around the globe, academics, politicians and technologists have time and again unravelled tales of foreign states and special interest groups using disinformation to manipulate electorates. At senate hearings, committee evidence sessions and academic conferences, experts retrospectively debate and lament the potential impact that false or distorted information has had on democratic processes. These conversations have added nuance and much needed evidence to our understanding of the tactics and tools used to disrupt and deceive citizens and communities. But they remain retrospective: far less attention has been devoted to mounting real-time and systemic efforts to detect and respond to disinformation before votes are cast.

Pieces of the puzzle certainly exist. Individual research projects and dedicated investigative journalists have broken important stories to warn the public of ongoing efforts to deceive them online. Governments have set up teams to monitor the open web for signals of interference from foreign states before elections. Academic observatories closely monitor new tech products for exploitation or interference. Tech platforms have quietly built security teams dedicated to identifying anomalies that might signal foreign state operations. And the Institute for Strategic Dialogue (ISD) has built a three-year evidence base of how different groups – governments, extremists, paid agencies – are using disinformation tactics to sow hate and division and undermine democratic processes across Europe. However, real-time action to prevent harm to communities and the manipulation of voters by covert means has remained elusive.

Setting the Scene: European Parliamentary Elections 2019

In the context of these challenges, the European parliamentary elections in 2019 presented a unique opportunity to test new models for detecting, analysing and responding to disinformation at scale and speed. The complexity of the election – conducted across four days in 28 countries and almost as many languages – presented countless entry-points for disinformation actors and increased the already difficult job for researchers trying to identify them. Add to this the turmoil of the European political system – shaken by Brexit, a resurgent populism and political fracturing across the spectrum – and the field was set for a precarious confrontation between disinformation perpetrators and those committed to protecting electoral integrity. Spotting this moment, Steve Bannon announced the launch of his European operation, The Movement, designed to provide data and strategic support to Europe's burgeoning populist and far-right political parties.[1]

At the same time, the European context provided real opportunities. The EU and its member states have led the charge on global policy responses to disinformation and online harms, explicitly calling out perpetrators, and platforms they have manipulated. While efforts to crack down on the wild west of internet platforms hosting harmful content remain lacklustre, disproportionate or counter-productive almost everywhere else, the EU Commission and a number of individual member states have pursued determined policy agendas for digital regulation that provide real opportunities to build systemic responses to disinformation threats.

The EU's Code of Practice on Disinformation – a voluntary set of commitments signed up to by a number of major technology companies before the elections – set out a robust framework for preventative and reactive measures from tech companies to mitigate the harms of disinformation, albeit with no means of enforcement.[2] The Code of Practice provided a number of research organisations, including ISD, with a basis for evaluating the responses of tech companies to the threats unearthed across the election period, as laid out in this report.

It was in this context that ISD mounted an effort to detect, analyse and respond to disinformation activities targeting six European countries in as close to real time as possible. Engaging with a loose coalition of like-minded partners, each already working on disinformation and digital threats, ISD set out to test what a proportional operation to counter disinformation in an election might look like – including the processes, skills, resources and networks that might be necessary to push back against deceptive and distortive online efforts to influence electorates. In order to capture the broad spectrum of activities that disinformation now encapsulates, ISD's research sought to identify and respond to malign information activities writ large – defined by us as activities that use online products, media systems or platforms with the outcome of deceiving audiences, distorting the available flow of information or conducting illegal activities. Deceptive tactics include activities such as creating or promoting disinformation or using sock-puppet accounts. Distortive tactics include the use of bots or bot networks to disproportionately amplify content in online networks. Illegal activities differ across national legal contexts, but can include hate speech, harassment, defamation or the provision of foreign in-kind support to domestic political parties in elections.

This report details the findings of ISD's research between February and May 2019. It lays out the tactics and actors involved in covert disinformation campaigns, the targets of their activities, and what that might mean for the future of disinformation around elections and beyond. It also evaluates the responses from tech companies and governments to these challenges during the election campaign and in the immediate aftermath, culminating in a set of concrete proposals for filling the gaps that this assessment clearly signposts. And finally, it seeks to provide an honest review of the successes

and challenges of this kind of model for understanding and mitigating the impact of disinformation, highlighting what remained elusive as well as what was possible, and proposes a series of lessons to be taken forward into upcoming elections and ongoing digital regulation debates across the globe.

Methods: Collection - Analysis - Response

How have disinformation operations been discovered over the past few years? Sometimes, a small clue from a closed Facebook group or a Telegram channel provides an opening into a much bigger investigation unearthing co-ordinated activity. Other times, anomalies in huge volumes of data signal the outlines of suspicious networks. In either case, it starts with analysis of some form of dataset that is being collected or monitored, be it by a tech company, a research institute or a journalist. On top of this, quantitative or qualitative analysis unearths something of the scale, nature and source of the intended activity. Then the challenge remains what to do with such data and evidence? What type of action is most likely to have the desired effect of disrupting and nullifying the impact of a malign information operation: an attempt to get it removed from a tech platform; a report to law enforcement; exposure in the media to warn the public; or the debunking of a false claim by a fact-checking organisation? And how best to deliver on that strategic decision-making?

These considerations informed the four-stage experimental model we put in place to confront disinformation during the campaign: data collection, analysis and detection, strategy and responses (Figure 1). There is no single pathway through these four stages. Disinformation tactics are evolving as fast as the political context into which they are planted, and each example requires a specific and strategic assessment of the best tools for discovery, analytics and counter-measures. A broad network of inauthentic accounts amplifying legitimate news content requires different technology, analytical models and reactions to the intentional anonymous trolling of a public figure with defamation and threats. Exposing one kind of disinformation to the public might build resilience; exposing another might only help to spread a lie further.

A. Data Collection and Monitoring

We collected and monitored data through the following sources:

The social listening tool Crimson Hexagon	aggregates publicly available data from Twitter, YouTube comments, Reddit, fora and blogs
CrowdTangle	aggregates page and group-level activity from public Facebook pages and groups

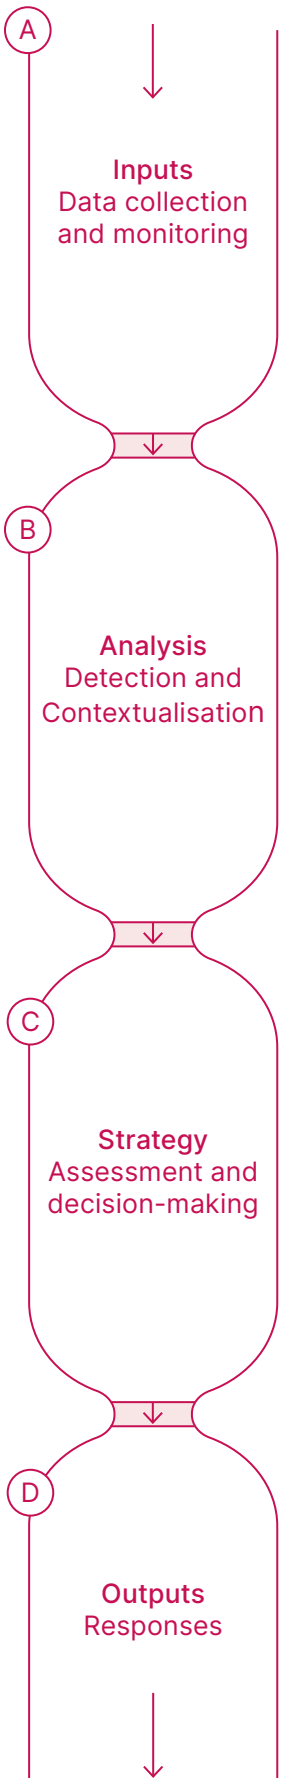

FIGURE 1
The input streams, analysis streams, strategy activities and output streams of the research project between March 2019 and May 2019

A political advertising library	API access to Facebook's political advertising database and web access to the online portal for Facebook advertising
Human intelligence on additional platforms	API access to Telegram groups and channels accessed by ISD researchers; manually archived content from researchers on platforms including Discord, VK and closed Facebook groups, as well as publicly available data from fora such as 4chan and 8chan
TV and online media monitoring	media monitoring and coding by partner organisation Memo 98 of online broadcast and print media coverage of the European parliamentary elections in European markets by state-affiliated and state-funded media channels
Polling data	consistent polling data from partner organisations assessing issues relevant to European publics running up to the elections.

B. Analysis and Detection

A team of researchers at ISD aggregated and analysed data drawn from all of the sources listed above. For each country, at least one data analyst and one political analyst worked together to assess signals of potentially malign behaviour or content and to prioritise leads for analysis along the following lines:

Legal framework	mapping of relevant laws on defamation, disinformation, extremism, hate speech, election campaigning, election finances, harassment, doxing, cybersecurity
Terms of service framework	mapping of platform policies, with a particular focus on disinformation, hateful content, spam and inauthentic behaviour
Network mapping	clustering accounts active on certain issues or in certain geographies to understand influencers and information vectors, and spot suspicious co-ordinated network behaviour
Models for assessing inauthentic patterns	volumetric analysis of content and account activity; sharing patterns for URLs, hashtags and keywords
Investigation	open source intelligence (OSINT) tools and techniques; offline investigative research to attempt to attribute disinformation activities.

C. Strategy

ISD's strategy team based their assessment of data trends and investigative leads on the relevance of each to the electoral context and potential threat to rights, safety or democratic processes. The volume of disinformation content and behaviours identified through the daily analysis required constant prioritisation and contextualisation to ensure the most pressing concerns were being addressed and communicated effectively. This process included:

Prioritisation	Assessing and prioritising research leads and findings, including relevant political and local expertise
Channels	deciding what was the most effective channel and mode of response
Messaging	developing a consistent policy and public messaging lines about disinformation threats and the narratives, actors and targets involved.

D. Responses

Depending on the urgency and potential harm identified in each case, ISD considered a number of options to try to mitigate the threat of malign

information activities, including:

Reporting to law enforcement	reports of suspected illegal activity to relevant authorities
Updates to policymakers (EU and national level)	regular trend insights on the threat and evaluations of technology company responses to EU level and national policymakers
Civil society early warnings and trend insights	alerts to potential target groups or individuals of upcoming disinformation attacks and ongoing communication around the nature of the threat to relevant issue areas and target communities
Direct reporting of content and account violations to social media platform security teams	direct reports of suspected co-ordinated disinformation campaigns to the security teams at Facebook, YouTube and Twitter, and regular reporting of individual instances of content or accounts violating the terms of service to request investigation and, where relevant, removal or downranking
Media exposure for public awareness	strategic media placement to build resilience to potentially widespread campaigns, or to explain already widespread campaigns
Media exposure for policy pressure	coverage of any perceived failures of company or government responses to incidents during the campaign to build public pressure for improved and timely responses to disinformation

Research Limitations

There are challenges at every stage of this process: poor data provision from online platforms limits collection and monitoring from the outset; rapidly evolving tactics of bad actors so analytical models for detection are constantly behind the curve; decision-making on responses being only partially informed as there is little information on the impact of disinformation activities; and the often opaque and ineffective responses of respondents from technology platforms to evidence of disinformation activities render some types of response futile.

Perhaps the biggest challenge is in understanding what the real impact of any one disinformation campaign might be. This project set out to weave together the normally disjointed pieces of evidence and data that might be required to at least start answering that question: online social media investigations, offline polling data and attitudinal surveys and broadcast media monitoring. The time and resource limits of the project did not allow for these streams of data to be integrated in the way they ideally could be to map the correlations between social media disinformation campaigns, media coverage and public attitudes around elections. But the sector is and needs to move in the direction of impact analysis if we are to develop truly useful responses to the threat of disinformation writ large.

These obstacles shouldn't prevent work being done to try to identify and mitigate harm before voters go to the polls. The fact that significant transgressions, co-ordinated malign activities and illegal behaviour can be identified despite these limitations hints that there is a lot more going on under the surface that researchers cannot see. The findings of this research laid out below help to build the case for greater data accessibility for researchers, the resourcing of ongoing independent research on disinformation, and concrete regulatory and legislative responses to step in where tech company responses continue to fall short in protecting voters online. We set out recommendations of how to move forward effectively and quickly – for researchers, governments and tech companies –at the end of the report, aiming to provide guidance in the run-up to upcoming elections across the globe.

The findings of this research laid out below help to build the case for greater data accessibility for researchers, the resourcing of ongoing independent research on disinformation, and concrete regulatory and legislative responses to step in where tech company responses continue to fall short in protecting voters online.

Research Findings

ISD and its partners uncovered 18 case studies of malign information operations targeting the European parliamentary elections during the campaign period.

Summary of key findings:

Actors: Types of digital deception most notoriously practiced by the Kremlin are now in the hands of a host of additional actors, including extremist groups and other states.

Tactics: A grab-bag of deceptive digital tactics was used to target the European Parliamentary Elections, going far beyond false information to include false identities, false communities, and false popularity.

Narratives and targets: The European Elections provided a short-term window onto much longer-running disinformation efforts aimed at undermining the equality and human rights of women, minorities and the lesbian, gay, bisexual and transgender (LGBT) communities, as well as efforts to address some of the major international challenges of our time through progressive, multilateral means, namely migration and climate.

The Tip of the Iceberg

ISD and its partners uncovered 18 case studies of malign information operations targeting the European parliamentary elections during the campaign period. These included sock-puppet networks in Poland, co-ordinated trolling attacks against vulnerable communities and activists in Germany and Italy, and automated and managed networks of accounts in Spain. There is little doubt that this is just the tip of the iceberg: if civic organisations and researchers are able to identify covert information operations without the requisite access to data or resources to do this work comprehensively, it suggests there is a much larger volume of activity still awaiting exposure under the surface, targeting not only elections but the wider information ecosystem online.

Many of the examples of co-ordinated disinformation that have been exposed over the past few years speak to the breakdown of binary categorisations of domestic or foreign, state or non-state, legal or illegal. The actors, tactics, targets and narratives involved in malign information campaigns continue to expand as the research sector attempting to detect them grows in scale and capability. While the perpetrators of malign information campaigns are hard to attribute definitively this project alone detected a number of new players, including organised hate communities and political activists from countries outside Europe. These shifts have been confirmed by more recent admissions of co-ordinated inauthentic behaviour disclosed by the platforms themselves, including the recent removal of significant domestic networks promoting both QAnon and VDARE[3] conspiracy and hate networks in the US on Facebook, alongside Iranian and Russian state-linked operations that have received greater attention over the past four years.[4] At the same time, the use of covert information tactics to target political campaigns and candidates is now well documented, but through this study we saw similar efforts waged against activists, non-governmental organisations (NGOs), and communities including Jews, Muslims and migrants across Europe.

A clear lesson learned from the multiple research efforts underway to detect and respond to disinformation during the European parliamentary elections is the residual difficulty of understanding the impact of covert or overt disinformation efforts online. Anecdotal evidence can help us understand the damaging impact of co-ordinated malign activity on individuals or specific communities targeted by such activity, including from misogynistic smear campaigns, conspiracy theory promotion or targeted harassment using anonymous accounts, doxing or hate speech. But any broader impact on the integrity of democratic processes or public attitudes remains little understood. There is a clear need to build methods that at least begin to chart the relationship between these deceptive efforts and public attitudes or behaviours around elections. Some ideas for the types of infrastructure and methods required for that type of research effort are laid out in the final section of this report.

The Digital State of Play in Elections

The field of play for information operations has been radically transformed by the onset of social media over the past decade. While our task was to find any covert activity online, the social media data we collected underscored important trends in political campaigning writ large in the digital era. We have witnessed a rapid uptake of digital electioneering by candidates and parties across the political spectrum, some legitimate and some questionably so. Despite the revolution we have experienced in our public information space and the impact this is having on political campaigning in the modern day, there has been no real public conversation about what rules should govern digital tactics in electoral campaigns. There remains little formal discussion of the acceptability of new tools and techniques of election campaigning online, including issues concerning the granularity of micro-targeting, the extent and nature of transparency for political advertising, and the use of automation to

help boost or amplify political messaging. This grey zone of acceptable political activity online provides crucial context for understanding the strategies, tactics and actors involved in disinformation efforts during 2019.

Social media data collected throughout the campaign period paints a stark picture of the extent to which parties considered to be on the political fringes outpaced more established parties in the online sphere. Social media platforms have played no small part in providing such parties, which often lack a long-term legacy of real-world grass roots mobilisation, with the ability to reach broader audiences. Data showing the frequency of party communications on Facebook at the outset of the campaign, and user engagement with these posts, demonstrates the complete imbalance of the field of play: in the run-up to the elections, right-wing populist parties and far-right parties dominated the online conversation about the elections. Parties such as Alternative für Deutschland (AfD) in Germany and Vox in Spain proved far more prolific and successful in engaging online audiences in direct materials about the European elections than their opponents (Figure 2). Through hyperactive levels of posting activity and similarly high engagement from users with their Facebook content about the elections, these parties radically outstripped their opponents on platforms like Facebook at the outset of the election campaign.[5] Whether such discrepancies in engagement were the result of highly resourced social media strategies, grassroots engagement from online

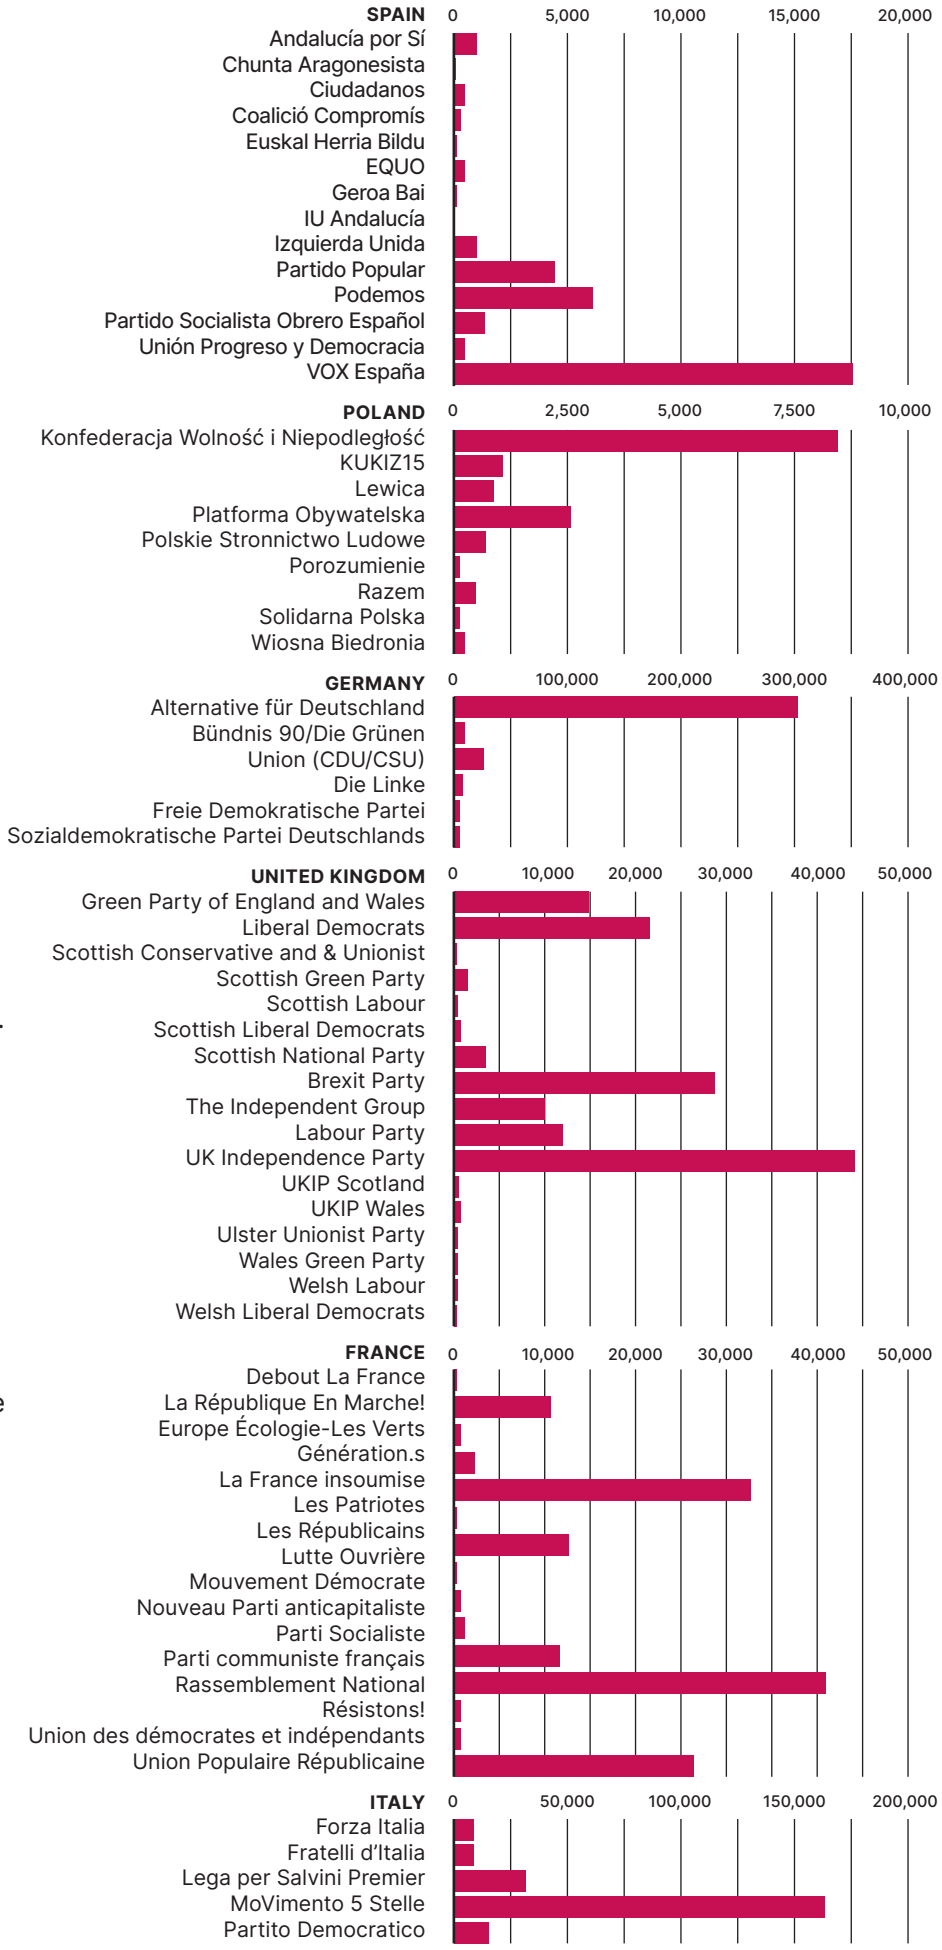

FIGURE 2
Total cumulative likes, comments and shares on official party public Facebook pages during European parliamentary elections 2019 in Spain, Poland, Germany, the UK, France and Italy.

supporters, hidden amplification tactics, or all three is impossible to tell from the top-line data alone, and remains opaque even to investigative researchers and journalists in part because of the limits on data access on these platforms.

These broad questions about legitimacy and norms in online political campaigning go beyond the scope of this report. Yet the scene-setting exercise is a necessary prerequisite for interpreting the aggressive far-right digital agenda that was subsequently launched across the election campaign.[6]

Findings: Actors

Ever since the revelations about the Kremlin's significant operation targeting the US 2016 election made their way into the public consciousness,[7] the focus of researchers, policymakers and the media has largely been trained on Russia when considering disinformation adversaries. The Internet Research Agency has been one of the most consistent highly resourced disinformation actors targeting Europe and the US in recent years.[8] Yet the disinformation playbook has rapidly been taken up by a whole host of other groups and organisations, as well as other states. The EU elections proved this more than most other examples so far studied: non-state extremist groups, a range of populist political parties and an array of nation states have made their way onto the disinformation field of play.

Problems Outside and Problems Inside

ISD's work with Media Monitoring NGO Memo 98 during the campaign found consistently poor journalistic standards from public broadcasters in some of the EU's own member states, providing a media context in these states whereby misrepresentation and media bias became the norm. Through quantitative and qualitative coding from 15 April 2019 to 26 May 2019, Memo 98 monitored the output of news shows on public broadcast channels to better understand state-backed media ecosystems active in the countries under study during the election campaign. This included Russia Today (RT) output in English, German, French and Spanish, Sputnik output in English, German, French, Spanish, Italian and Polish, the news output of TV channel M1 and Radio Kossuth in Hungary, TVP Wiadomości on TVP (Telewizja Polska) and TVP Polonia in Poland, and Rai Due and Rai Tre in Italy.

While much focus has been paid in recent years to the relevance of RT and Sputnik to influence campaigns and instances of disinformation targeting Europe, through traditional broadcast and online vehicles, two examples of intra-EU news media were identified through this research as providing an imbalanced and often misleading media environment during the campaign. M1 news programming in Hungary and TVP Wiadomości in Poland were found in the monitoring research consistently to lack balance in the presentation of news content and the sources and speakers provided with a platform through the programmes. Images and videos were often used out of context to promote narratives in favour of the ruling party in both contexts (Figures 3 and 4).

Turning to the social media analysis at the heart of the project, non-state activists and extremist groups came to the fore as perpetrators of organised malign information campaigns. There is increasing evidence that extremist and political non-state groups are taking up the full tactical playbook used by states such as Russia, China and Iran, using new media ecosystems to weaponise key issues and to launch deceptive or disinforming campaigns. The case study of the pro-Vox Twitter network (below) demonstrates the grey zones of foreign and domestic disinformation activities, with evidence of Venezuelan non-state activists managing accounts messaging in Spanish, directed at Spanish audiences, and promoted by Spanish party supporters and candidates.

Internationally networked online activists have mobilised across a range of recent European elections, using alt-tech platforms like 4chan, 8chan and Gab

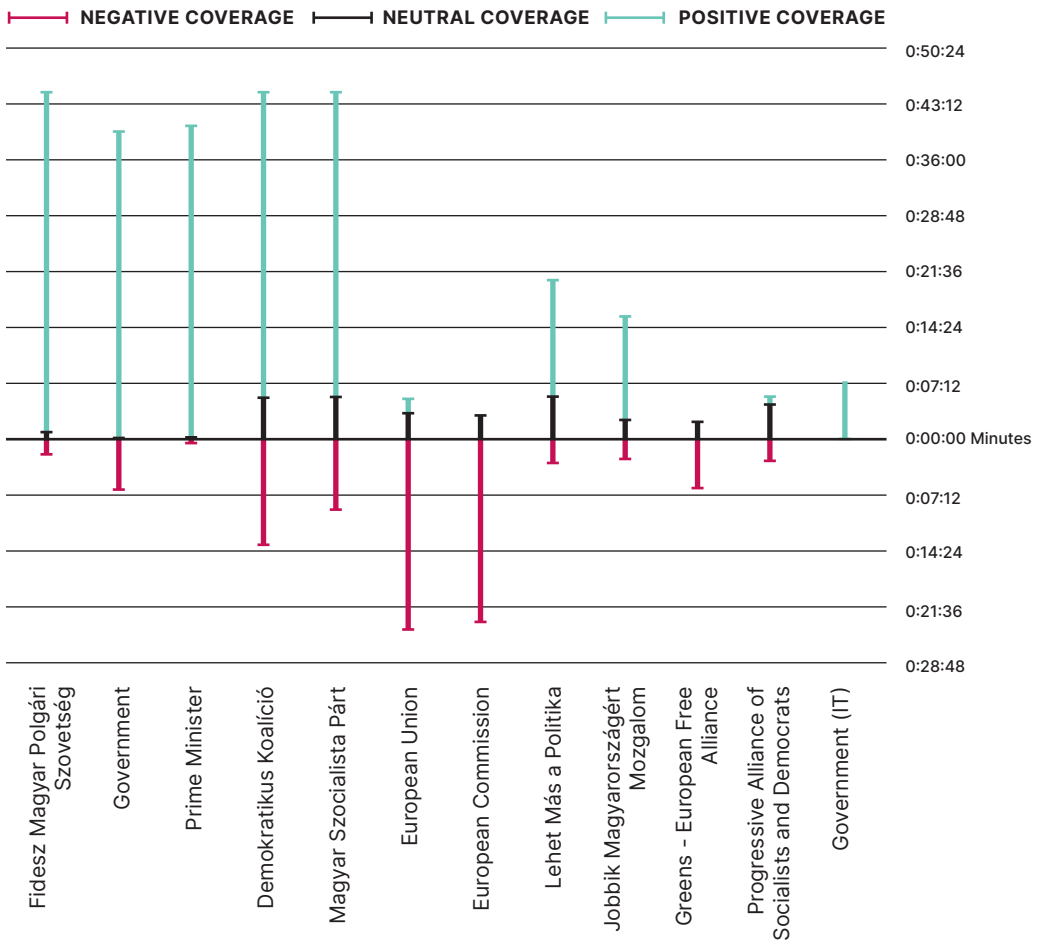

FIGURE 3
Number of minutes provided to the 12 most-covered topics on the main 18:00 news show on M1 in Hungary from 15 April to 26 May 2019, coded for positive (green), neutral (grey) and negative (red) coverage of the topics in question (Memo 98)

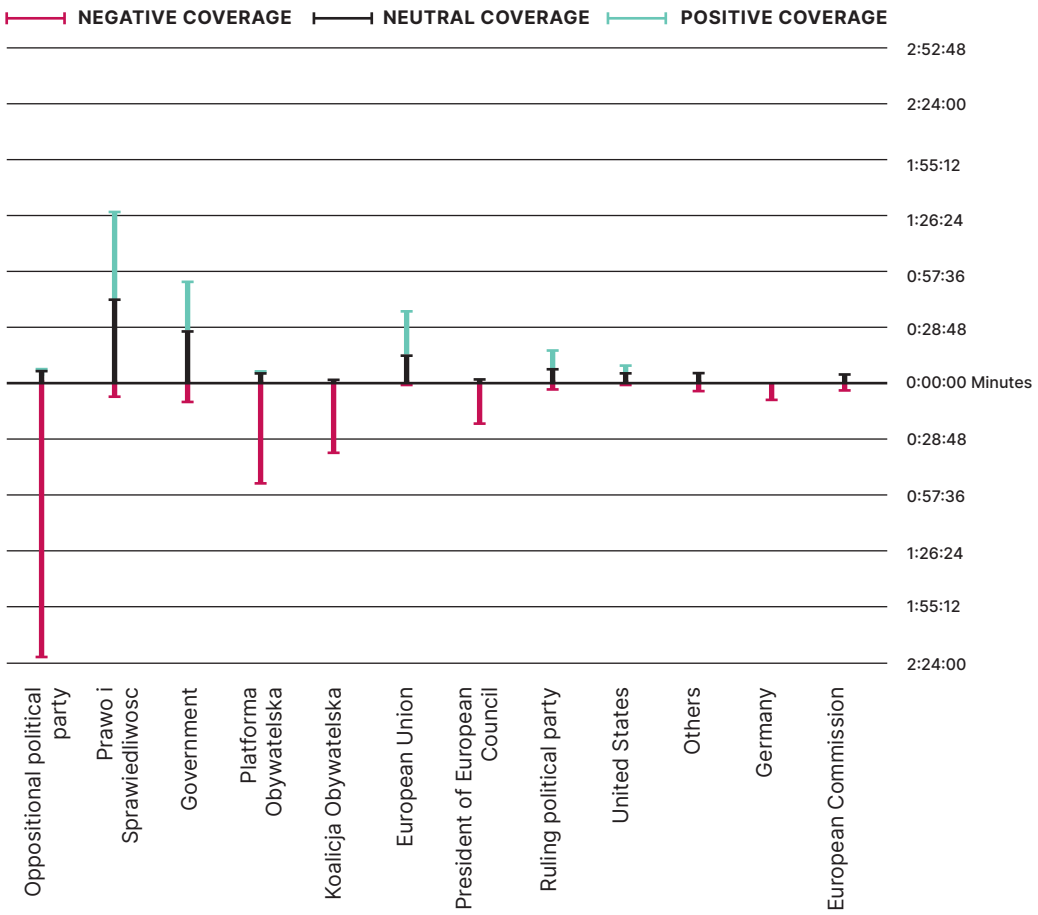

FIGURE 4
Number of minutes provided to the 12 most-covered topics on the TVP Wiadomości main 19:30 news show on TVP in Poland from 15 April to 26 May 2019, coded for positive (green), neutral (grey) and negative (red) coverage of the topics in question (Memo 98)

to co-ordinate disinformation campaigns. In recent related work, ISD supported the Polish NGO Reporters' Foundation to uncover the disinformation and media manipulation efforts of a private public relations firm.[9] Similar private companies have been exposed supporting covert information campaigns in the Philippines.[10] ISD has consistently evidenced the use of disinformation by far-right groups across Europe and the US to spread conspiracy theories including the 'great replacement' and 'white genocide' conspiracies.[11]

Our research in this election cycle exposed the use of a flourishing ecosystem of blogs, vlogs, news sites and channels by international far-right activists across Europe. A parallel ISD research study, focusing on the variety of online platforms used by the far-right in Germany specifically, confirmed the breadth of this activity online.[12] Campaigns by non-state organised groups were used to discourage Europeans from voting for centrist and left-leaning parties, to drive the political discussions towards far-right populist themes, and to put pressure on politicians in power, moderate frontrunners and journalists reporting for traditional news outlets through harassment and hate speech.

ISD researchers observed mutual amplification between pan-European extreme-right networks such as the New Right and Identitarian Movement, alternative news sites and far-right blogs, and far-right and populist parties and frontrunners for the European Parliament elections online. These three groups were found coalescing around common rallying points such as terrorist attacks in Christchurch and Utrecht, but also around events such as the Notre Dame fire. Local and national issues often became international rallying points for these actors through online campaigns. Examples of the tactics and narratives launched or promoted by these groups are described in the case studies that follow.

Venezuelan Twitter account network active in Spain

Case Study

In Spain, ISD uncovered a network of 2,882 accounts, including suspected bots and managed sock-puppet accounts, which promoted anti-Islam, anti-immigration, anti-Soros and pro-Vox content in the run-up to the European parliamentary elections, and was managed by activists in Venezuela. In 2019 alone, the combined network published an average of 152,907 tweets per day, spreading pieces of disinformation linked to migrants and Muslims. In one example, a video of a riot in Algeria was shared, after its caption was changed to imply that the footage was filmed in a predominantly Muslim neighbourhood in France. The content was promoted by both the network and an official Vox account.[13] Between 23 April 2018 and 23 April 2019, the network tweeted 4.4 million posts mentioning Vox, with coding of samples of this content suggesting that these posts supported the party. Similarly, the network mentioned Vox leader Santiago Abascal around 460,000 times in tweets during the same year year. The network consistently posted outbound URLs to unofficial Vox supporter Telegram channels.

Analysis of the network suggested automated behaviour. In the two years before identifying the network, 33% of all its outbound URLs linked to the Twitter management tool tuitutil.net, which allows users to semi-automate key actions such as following other accounts quickly. In addition, the ten most active accounts had been created relatively recently, with very similar usernames, identical Spain and Venezuelan flag-emojis, with each having tweeted hundreds of thousands of times. These signals suggest at least semi-automated activity.

The network published more than 400,000 anti-Islam posts, with hashtags such as #StopIslam and #NoAllslam. Over 10,000 posts using the term #NoAllslam were published between February and May 2019. A joint investigation by ISD and El País found that the network was originally used to oppose the Venezuelan government but was reactivated in 2017 after a period of silence.[14] The accounts were reported to Twitter, whose staff removed 39 of

TACTICS SPOTTED
Bots, sock-puppet accounts, account management tools, co-ordinated hashtag hijacking, media taken out of context.

RELEVANT PLATFORMS
Twitter, Telegram

ACTORS
Political party supporter groups, foreign non-state activists

the accounts. Many of the accounts have not been active since the elections in April and May 2019, despite previously high levels of activity daily.

Findings: Tactics

Outright false information was rarely evident in disinformation targeting European audiences during the European Parliament campaign, despite its clear resurgence around other key moments such as the recent coronavirus crisis.[15] It certainly still exists, and in certain contexts remains a threat to public safety and human rights. The recent examples of lynching and violent attacks against individuals in India are just one example of such harm done through falsified content at massive scale.[16] Disinformation about vaccines and Ebola pose serious risks to public health.[17] The UK general election of December 2019 clearly showed that explicit and intentional political lies are far from a thing of the past, and are now supercharged by micro-targeted ads and social media promotion.[18]

But there is a much wider and often subtler grab-bag of tactics available to all of the actors mentioned above. These tactics present a different set of challenges in detection and response. Examples of actors misrepresenting communities, people and popularity were all identified in the European elections context, each falling within the frame of 'deceptive, distortive or illegal' activity that was the exploratory subject of the research.

The tactics identified across case studied unearthed in the research included:

- ↗ false content and conspiracy theories
- ↗ bots
- ↗ sock-puppet accounts
- ↗ co-ordinated account networks
- ↗ co-ordinated hashtag hijacking
- ↗ media taken out of context
- ↗ harassment

The case study of disinformation activity surrounding the Notre Dame fire neatly encapsulated the breadth of tactics deployed in an opportunistic disinformation campaign. Actors mixed false information with decontextualised photos, shared alongside harassment and conspiracy theories about Muslims. It also provided an example of the co-ordination that is often so invisible to researchers: activists on boards on 4chan used by extremist groups and hosting extremist content purposefully shared images of Muslims smiling near Notre Dame from completely irrelevant contexts in order to claim Muslim involvement or glorification of the fire.

Disinformation about Notre Dame fire

Case Study

The instrumentalisation of news events to spread anti-Muslim hate was exemplified in the disinformation campaigns that followed the fire of Notre Dame in April 2019. Far-right activists in Spain, France, Germany and Italy disseminated disinformation about the fire, including claims that Muslims were rejoicing at the news and that the fire was orchestrated by Islamist extremists or by the authorities themselves.

On 4Chan, users encouraged each other to share pictures of smiling Muslims near Notre Dame in order to 'expose them'.[19] Pictures of Muslims near Notre Dame, taken out of context, were posted to the thread as material for the campaign. French Identitarian Telegram channels were used to publish extensive anti-Muslim content in reference to the fire. For instance, the English-speaking channel Europe Lives published a picture of Notre Dame in flames with the captions: 'They will take from you everything: Your culture. Your history. Your existence.'

In Spain, Vox influencer accounts spread a range of false details and pieces of disinformation, including that the fire was announced before the event, that a purported terrorist was arrested with gas canisters and that Muslims celebrated the fire. Vox leader Santiago Abascal wrote on Twitter: 'Islamists

TACTICS SPOTTED
False information, media taken out of context, conspiracy theories, hashtag hijacking

RELEVANT PLATFORMS
4Chan, Telegram, Twitter, Facebook

ACTORS
Political candidates, political party supporter groups, extremist movements

who want to destroy Europe and Western civilization by celebrating the #Notre Dame fire. Let's take note before it's too late.'[20]

German far-right activists, politicians and media outlets proactively engaged in speculation and disinformation campaigns about Notre Dame. Our analysis of German Twitter traffic in connection with Notre Dame-related hashtags and keywords shows that articles by far-right outlets such as Philosophia Perennis[21] and Tichy's Einblick[22] were among the top-shared URLs about the fire. Similarly misleading narratives were amplified by official regional AfD accounts on Facebook. The AfD Herne page claimed, for example, that 'it is another fact that the fire of Notre Dame de Paris was literally celebrated by Arabs on the street and in social networks!!'.[23] Meanwhile, the AfD page for North Rhine Westphalia claimed that Islamic associations were silent about the fire and asked: 'What would be happening had one of the world's great Mosques been on fire?'[24]

A consistent trend in ISD's work researching disinformation around elections is the use of false information, weaponised anonymity and co-ordinated hate speech to harass female political candidates and journalists. ISD identified several concerted campaigns against female politicians, journalists and activists in France, Germany, Italy and Spain. The tactics deployed ranged from meme warfare to smear campaigns and offline intimidation, all exploiting the ability of social media platforms and content hosting sites to amplify, target or anonymously harass public figures. It should be noted that our analysis can only scratch the surface of the actual volume and quality of abuse directed towards these figures. Victims of such attacks have often said in interviews that most harassment material is directed through private messaging channels.

FIGURE 5
The number of posts on German political parties' group admin Facebook pages mentioning Chebli, 4 April 2018 to 4 April 2019, by party affiliation

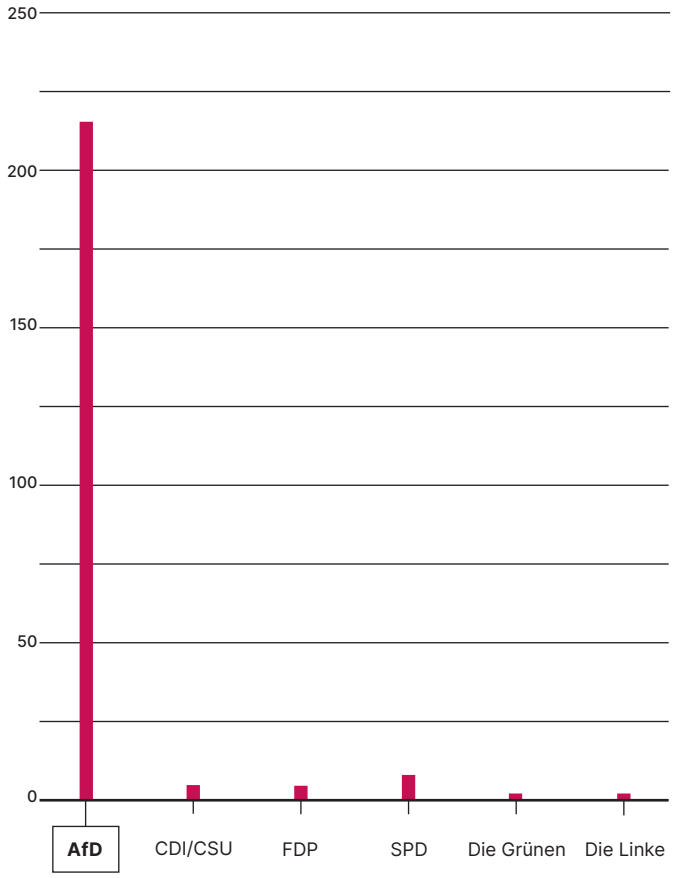

Ongoing attacks against female politicians

Case Study

Across Europe, ISD identified several co-ordinated misogynistic campaigns targeting female politicians, journalists and activists. In Germany, the far-right has been attacking Sawsan Chebli, a Sozialdemokratische Partei Deutschlands (SPD) politician with a migration background. The far-right populist party AfD has been at the forefront of the attacks. The volume of attacks dramatically rose after it emerged that Chebli wore a Rolex in October 2018. These attacks included threats as well as insults. During the peak of the outrage between 20 and 24 October 2018, more than 26,000 tweets were sent directly mentioning the 'Rolex-Chebli' issue. Sustained harassment via social media over this period forced Chebli to delete her Facebook account.[25] AfD was particularly active in its criticism of Chebli. It published disproportionately about her activities between April 2018 and April 2019 (Figure 5).

TACTICS SPOTTED
Harassment, false information

RELEVANT PLATFORMS
Twitter, Facebook, 4Chan, YouTube, Telegram

ACTORS
Extremist groups, political candidates

Chebli has also been a target of prominent right-wingers and extremist influencers in the country. In one example, Akif Pirinçci, a Turkish-born far-right influencer, attacked Chebli's religious beliefs: 'You're not a secretary of state at all. You're an annoying Muslim woman to fulfil a quota drinking taxpayers' money for doing nothing. And if you don't drink and eat for 12 hours, it's not fasting, it's bullshit. Actually, it doesn't matter what you do.'[26] The anti-Chebli campaign was also picked up by the international far-right. Before the May 2019 elections, users on 4Chan created memes of Chebli in the lead-up to the Bavarian 2018 state election. One of the first videos that appeared on YouTube when searching for Chebli's name during the election period is a video by a YouTuber notorious for publishing nationalist, xenophobic and anti-refugee videos, of which more than a dozen are about Chebli herself. Some of these videos have received upwards of 300,000 views.[27]

In France, the far-right launched a co-ordinated harassment campaign following the nomination of Sibeth Ndiaye, a black woman with dual Senegalese–French citizenship, as government spokesperson on 31 March 2019. Far-right media and influencers spread disinformation within minutes of Ndiaye's nomination. One report claimed that Ndiaye had responded to the death of Holocaust survivor and veteran politician Simone Veil in 2017 with a text saying, 'Yes the chick is dead.' This claim was spread without context to suggest that the message was a celebration of her death. Fact-checkers showed that the text message was misquoted, with Ndiaye in fact responding to a question about the date of a funeral for the late Simone Veil with the words, 'Aucune idée, la meuf est morte il y a moins de vingt-quatre heures' ['No idea, the girl died less than twenty-four hours ago'].[28] ISD's analysis of social media data in the 24 hours following Ndiaye's nomination showed that comments from Ndiaye on controversial issues were often misquoted, and went hand in hand with explicitly racist and sexist content. On Twitter, online conversations were driven by far-right influencers who disseminated disparaging comments about Ndiaye's nationality, appearance and skin colour. Ndiaye's promotion also became a talking point for the international far-right. ISD found multiple examples of messages and tweets in English-speaking Telegram channels, which have not been included here owing to their graphic nature.[29]

Findings: Narratives and Targets

Elections only provide a short-term window onto much longer-running disinformation efforts that we have seen attempt to shift public discourse and opinion so as gradually to undermine faith and trust in liberal, international(ist), democratic norms. The central perceived 'acquis' of liberal democratic societies – the equality and human rights of women, minorities and the lesbian, gay, bisexual and transgender (LGBT) community – as well as efforts to address some of the major international challenges of our time through progressive, multilateral means (migration and climate) have become ongoing targets of a constellation of disinformation actors. Attacking 'scape goat' communities and skilfully exploiting existing fears and grievances within society, disinformation actors sow division, confusion and mistrust, with a view to building an ideational social fabric more supportive of authoritarian and nationalist political behaviours. This is where state and non-state actors are aligned in their objectives.

During these elections, we saw culture war dynamics, more frequently recognised as a characteristic of US public discourse than a European one, flourishing online in countries such as Italy, Spain and Poland, boosted by disinformation actors. Malign activities were directed towards promoting anti-LGBT sentiment and discussions of 'family values' and women's rights, especially in Italy, Spain and Poland. Elsewhere, existing hate groups used disinformation tactics to target issues around integration, immigration and race, in order to weaponise their existing anti-Muslim, anti-Semitic, anti-migrant and anti-Roma agendas across Europe. These tactics were used to give age-old

conspiracy theories a modern-day digital revamp, and to target Muslim and Jewish communities. Finally, climate change entered the field as a prime target for disinformation in Europe, with a concerted attempt by some to enshrine it as the new wedge issue of choice. We expect to see continuing jockeying around the issue of climate and the environment as younger generations in Europe take it up as a central political concern of their age.

The Weaponisation of Hate

Hate groups and extremist movements have long used disinformation to promote their causes. Deception lies at the heart of much of the propaganda used to mobilise support and recruit members, both before and since the dawn of social media. Those efforts have been hypercharged by social media tools – widely available, usually for little to no cost – which can falsely amplify and intensely micro-target content.

In monitoring social media and alt-tech platforms for co-ordinated disinformation activity during the European elections 2019, the most prominent trend in the incidents that ISD identified as illegal, deceptive or distortive was their relationship to hate and extremism.

In monitoring social media and alt-tech platforms for co-ordinated disinformation activity during the European elections 2019, the most prominent trend in the incidents that ISD identified as illegal, deceptive or distortive was their relationship to hate and extremism. Across all six countries studied, inauthentic account networks, false content or illegal harassment tactics were identified as part of wider efforts to promote anti-Semitic, anti-Muslim, anti-immigrant, sexist or anti-LGBT ideas. Disinformation activities are intricately interlinked with long-term efforts to instil division and prejudice in societies and to promote political parties that sympathise with ideologies that diminish the rights of minorities or vulnerable communities. Thus the intention and methods of disinformation present potential harm to European citizens, inside and outside elections.

The battleground of disinformation in Europe is now less riven by obvious bot networks and completely false news stories than in the past, even though we still found some evidence of these. In line with the shifts in tactics and strategies of disinformation evidenced elsewhere,[30] ISD witnessed in 2019 the resurgence of a culture war dynamic in Europe, sitting firmly on the back of disinformation tactics that help to promote confusion and polarisation in online communities. While completely false stories and clearly false accounts are still one part of the disinformation playbook, the major weight of efforts promoted polarisation and confusion around key social, cultural and political wedge issues. In addition to the issues of hate explored above, the climate change debate became a clear target of malign information campaigns in Europe in early 2019, linking to offline efforts to support climate change denial across the continent.

Co-ordinated Disinformation Targeting Muslims and Jews

Anti-Semitic disinformation campaigns in Poland

Case Study

ISD identified a seemingly co-ordinated network of accounts in Poland spreading anti-Semitic content in Polish on Twitter. The network comprised 803 accounts, whose posts referred to Jews over 92,900 times in the year running to May 2019. The network used hashtag pairing and hashtag hijacking tactics to disseminate messages promoting disinformation about Jews in Poland, including disinformation about the history of the Jews in Poland, for example through constant use of the hashtag #JewishTruth from the start of 2018 until May 2019, identified in 33,000 tweets from the network.

This network of accounts spread disinformation about historic events on social media, for example around the hashtag #ustawa447. The hashtag referred to the Justice for Uncompensated Survivors Today (JUST) Act, a piece of legislation which requires the State Department to report to the US Congress on the effort made by European countries to compensate survivors of the Holocaust and their descendants for property seized by Nazi Germany

TACTICS SPOTTED
Hashtag hijacking, co-ordinated account network, false information

RELEVANT PLATFORMS
Twitter

ACTORS
Unknown

and post-war communist states. The JUST Act is controversial in Poland as some argue it portrays Poland as a nation of collaborators instead of as a victim of fascism. The far-right party Konfederacja claimed that it could cost Poland between $65 billion and $300 billion.31 And the Independence March Association created a petition to appeal to Donald Trump to repeal the JUST Act, with the main slogan of the petition being: 'Do you know, that they [Jews] want to rob you?'.[32]

See 'Case Study: Venezuelan Twitter account network active in Spain', above, for details on the case of sock-puppet accounts and bots used to promote anti-Muslim hashtags in advance of the Spanish general election in April 2019.

The New Divide: Disinformation and Climate Change

In the lead-up to the European elections, climate change became one of the most hotly debated political issues, driven in part by highly visible environmentalist movements such as Fridays for Future and Extinction Rebellion staging major protests across the continent. In an EU-wide survey conducted by the European Commission, 93% of EU citizens see climate change as a serious problem and 79% see it as a very serious problem. Among the countries we were monitoring at ISD in the lead-up to the European elections, respondents in Poland (70%) and the UK (75%) were less likely to view climate change as a very serious problem while those in Spain (89%), Italy (84%), France (82%) and Germany (81%) were more likely to view it as a very serious problem than the average across all EU countries.[33] In Germany, polls before the election even showed that climate change led the list of topics that survey respondents said would shape their voting decisions (48%).[34] Climate change was one of the main topics discussed in the video 'The destruction of the CDU [Christian Democratic Union]' by the German YouTuber Rezo, which was viewed 12.7 million times within a week before the May election, and became a key discussion point in German political conversation preceding the vote.[35]

In response, the German and French far-right actively engaged in ongoing campaigns targeting climate change activists and disputing climate science. The AfD's campaign concerning climate change included vicious personal attacks on young climate activist Greta Thunberg, and campaigns that dispute climate science and support for green energy. These recent campaigns fit into a broader anti-environmentalist slant which the far and extreme-right have taken in recent years.[36]

In Germany, Karsten Hilse, the AfD's environmental spokesperson, invited a group of speakers to a symposium to challenge and question climate science hosted in the German parliament. The event was publicised by a German-based political

FIGURE 6
Official AfD public Facebook page admin posts mentioning 'Klimawandel' (climate change), 19 April 2016 – 18 April 2017, 19 April 2017 – 18 April 2018, and 19 April 2018 – 18 April 2019.

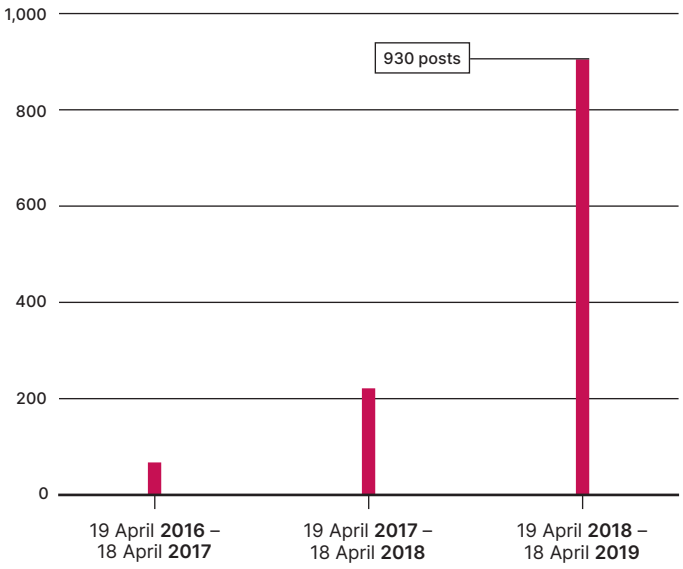

organisation called the European Institute for Climate and Energy (EIKE).[37] EIKE's annual climate conference is co-sponsored by the Heartland Institute,[38] a fossil fuel industry-funded U.S. think tank that has a history of funding projects aimed at weakening public confidence in climate science.[39]

The AfD's social media attack on climate change campaigns and campaigners[40]

Case Study

ISD researchers used CrowdTangle to gather all AfD posts about climate change (Klimawandel) in the past three years on official party Facebook pages (Figure 6). Since 2016, AfD pages on Facebook have posted content denying human-made climate change. While the AfD has not shifted its position, the party has used social media to communicate more frequently about this over time, especially in the run-up to the election in spring 2019.

With the rise in publicity and media on the issue surrounding Greta Thunberg's campaigns, the AfD started to post more about climate change in general and Thunberg in particular in early 2019. In the year between April 2018 and April 2019, the AfD mentioned 'Greta' 791 times and 'Greta Thunberg' 452 times.[41]

TACTICS SPOTTED
False information, conspiracy theories, harassment

RELEVANT PLATFORMS
Facebook, offline think tanks and advocacy organisations

ACTORS
Political parties, political party supporter groups

Since the emergence of Thunberg as a public figure, the AfD has increasingly presented belief in climate change as irrational, hysterical or cult-like, or depicted it as a replacement religion. ISD's team gathered AfD posts using the following words suggestive of such narratives: CO2Kult, Ersatzreligion, Klimasekte, Klimapanik, Klima-Panik, Klimawandelpanik, Klimawandel-Panik, Klimahysterie, Klima-Hysterie, Klimawandel-Hysterie, Ersatzreligion, Klimawandelhysterie, Klimareligiöse, Klimareligiösen, Klimagehirnwäsche, Klimareligion, Klimaschwindel, Klimatismus, Ökoterrorismus.[42] Figures 7 and 8 show how references to climate change as irrational have risen dramatically since the AfD started posting about Greta Thunberg and her climate campaigns. A joint investigation between ISD and Unearthed at Greenpeace sheds light on the offline links between the AfD and groups seeking to undermine faith in climate science, including the European Institute of Climate and Energy, which is in turn supported by the US organisation the Heartland Institute.[43]

Online attacks on Greta Thunberg

Case Study

The harassment of the Swedish climate activist Greta Thunberg from both right-wing populist parties and far-right groups cuts across a number of trends in relation to the online attacks on public figures. Thunberg was attacked as a very prominent female public figure and mocked for her disability, but also became a convenient target in the discussions around climate change. Some of the attacks against her originated from alt-tech platforms such as Telegram, while others gained traction from the frequent attacks from right-wing populist party pages on major platforms such as Facebook.

TACTICS SPOTTED
Harassment, false information, conspiracy theories

RELEVANT PLATFORMS
Telegram, Facebook, YouTube

ACTORS
Political parties, political candidates, extremist movements

In Germany the AfD mounted a highly personal campaign against Greta Thunberg claiming she is being manipulated, denigrating her for being autistic and comparing her to a Nazi. On 8 February 2019, the AfD's local Facebook page for Salzgitter promoted the conspiracy that Greta Thunberg is being manipulated by 'eco-fascists'. On 8 March 2019, AfD MP Frank Pasemann tweeted in support of this conspiracy theory, also denigrating Thunberg for being autistic.[44] On 23 March 2019, Martin Schiller, an AfD candidate for the European Parliament, posted a meme of Thunberg wearing the uniform of the female wing of the Hitler youth, which read 'Youth serves the climate', a reference to Nazi slogans.[45]

Pegida activists have similarly engaged in a harassment campaign that uses grotesque caricatures promoting the narrative that Thunberg is disabled. On 31 March 2019, Pegida-founder Lutz Bachmann posted a distorted picture of Greta on his Telegram channel, asking: 'Has the mentally retarded climate-Greta [wordplay with frog] already received an award today or is that still taking some time?'[46]

FIGURE 7
Number of admin posts from official AfD public Facebook page mentioning “Greta Thunberg” or “Greta” from Jan 2018 to April 2019.

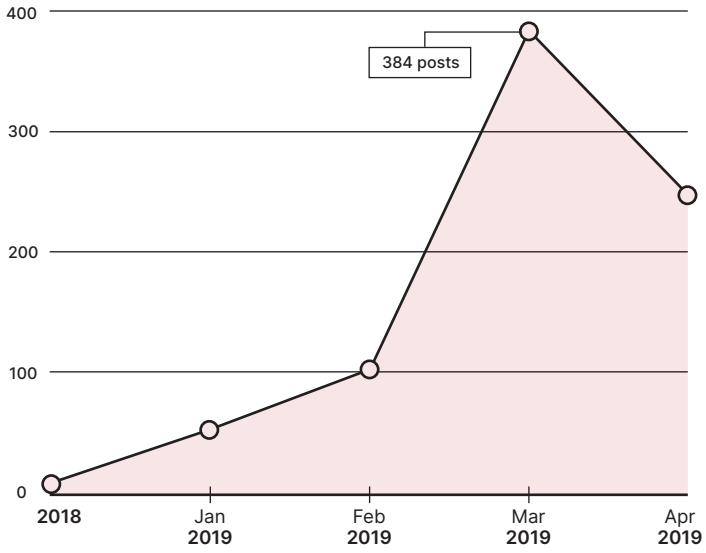

FIGURE 8
Number of admin posts from official AfD public Facebook page presenting climate change as irrational, cult-like or a replacement religion,48 19 April 2016 – 18 April 2017, 19 April 2017 – 18 April 2018, 19 April 2018 – 18 April 2019.

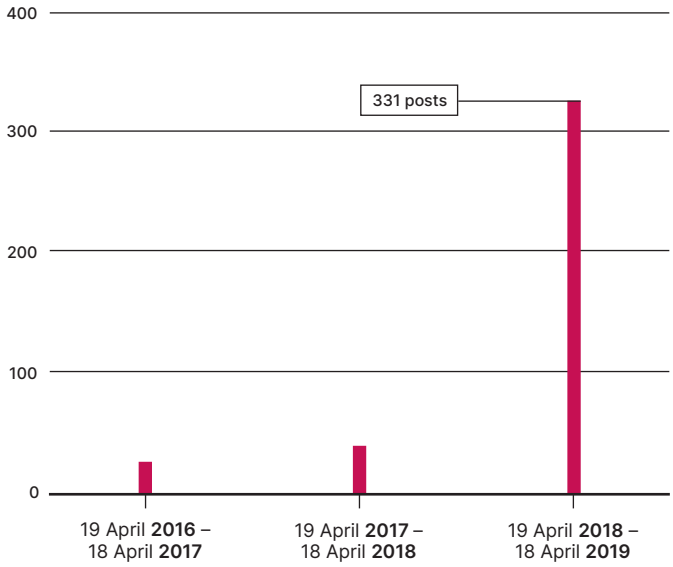

In a similar vein, French far-right influencers and politicians mobilised against Greta Thunberg, spreading disinformation about, and in some cases abuse against, the young activist. YouTuber Bruno Le Salé (52,000 followers) posted a video entitled ‘The scam of the climate march: GRETA THUNBERG’ on 4 April 2019. The video, which has been viewed more than 71,000 times, shows Bruno ‘strangling’ a figure with Thunberg’s face. The YouTuber goes on to describe Thunberg as a ‘pure marketing product’ manipulated by entrepreneurs who are using her action to make money through greenwashing.[47]

Implications: What Does It Mean and What Can We Do?

Between 23 and 26 May 2019, over 50% of Europeans eligible to vote went to the polls to cast their ballot for the next European Parliament.[49] A reverse from four decades of declining voter participation, turnout levels suggest a still healthy democratic participation across Europe. The results did not tell a coherent story across the continent. While both dominant blocs from the previous Parliament lost seats – the centre-right European People's Party and the centre-left Socialists and Democrats – they remained the two largest blocs in the new set-up. Seats were not taken from them by a unified grouping but by a broad range of parties across the centre, the left and the populist and nationalist right, varying markedly from country to country.[50]

While in France, Italy and the UK, populist right-wing and nationalist right-wing parties took the lead, their parallels across the rest of Europe failed to make a significant dent in the overall presence of nationalist parties on the European stage. There was a surge in support for Liberal and Green parties, notably so in Germany. Fragmentation may be the only clear way to characterise the overarching pattern of voting behaviour in the May elections, then. The weakening of centre-right and centre-left groupings, the rise of the Greens, the burgeoning presence of the Alliance of Liberals and Democrats for Europe (ALDE) and Salvini's substantial gains tell a wealth of stories. If disinformation and malign online campaigns had a direct impact on voter behaviour, it would certainly be a difficult relationship to draw out from such a complex new arrangement of political power.

What Does This Mean for Research? Methods and Tools

The European parliamentary election of 2019 was an opportunity to try new ways of detecting and responding to disinformation in real time, across six countries and on even more platforms. Along with its successes, the attempt to mount a collaborative project to detect and deal with disinformation in near real-time also shed light on the requirements for any similar efforts in the future. In the first place, lessons learned from the experience highlight the need to build the capacity and methodological range of the disinformation sector in the context of ever-evolving tactics and actors deploying them. This goes beyond current reliance on commercial tools and platform-specific analytics.

Laid out below are six principles for the kind of infrastructure that civic society requires to do this work comprehensively and accurately. This weaves together research strategy, technological architecture and team skills – all of which are integral considerations in approaching this field of work in the long run. A research infrastructure based on these principles is now being designed to help the civic sector identify and respond to disinformation threats in the US presidential elections in 2020, based on these insights from the European experience.

1 Research should be plugged into civic society in a number of ways: its priorities and direction should be informed by what many different groups and people see; it should work in ways that civic society understands; and it should produce outputs that allow for a civic societal response.

2 Research must leverage the full opportunities available for civic society to acquire data from all the platforms and online spaces relevant to illicit online manipulation, beyond the scope of internet platforms most often researched.

3 Research should have a detection function to identify and filter social media data according to whether it conforms to one of a series of behaviours that relate to illicit online influence. This should remain sensitive to platform, but also operate across platforms.

4 Research systems should never be in a settled state, but should have a reactive technology development capacity to add technological capability on the basis of requests driven by analytical teams who are face-to-face with the trends that are being exposed.

5 Research in this area must have a cyclical discovery function: the system must help analysts find online behaviours which are not known, but which are similar to behaviours that are known. As far as it is possible, the system must be designed cyclically, so its outputs can be used as further inputs. Then over time the system will be able to evolve as the phenomena that it tracks themselves change in nature and scope.

6 Lastly, the empirical outputs of the system contribute to, and draw from, conceptual and definitional work. This area continues to suffer from overlapping and poorly delineated definitions of the problem phenomenon itself, and the tactics, techniques and strategies related to that phenomenon. Any detection system must be integrated into a systematic and continuous effort to develop the concepts that define what it is intended to detect.

Understanding Impact

While researchers have been able to find hundreds of individual examples of disinformation activities across a range of elections, connecting these to a tangible outcome has been nigh-on impossible. The difficulty of getting close to an understanding of the real-world impact of digital disinformation on the public fundamentally undermines our ability to respond strategically and proportionately to cases that arise.

But there are methodological possibilities that could help to push disinformation research towards something of an understanding of potential impact. Although imperfect, a comparative analysis of attitudinal polling, broadcast and online news media, and social media would start to map potential links (or the lack of them) between disinformation efforts and the media, public or political ecosystem. These research opportunities are far from new in political communication studies or social science, but have barely been tested on the issue of disinformation because of the sheer volume of data required, the limits on data access, and the speed at which such analysis must be made to mitigate possible harm. Finding ways to examine the relationship between disinformation activities, media discussion and public attitudes or concerns can begin to put the audience back at the heart of conversations around disinformation and its potentially dangerous role in democratic processes, social cohesion and public safety.[51]

Other options also present opportunities: turning the research methodology for detecting illicit influence on its head, efforts could start instead by finding evidence of where disinformation has had an 'impact', and then deploying data analytical and OSINT investigative resource from that point to trace back to the identification of precursor online experiences. This approach would go through four steps:

1	Does it matter?	Find people who have been 'impacted'.
2	What is it?	Identify how the social media experiences of people who have been so impacted are different from a control group.
3	Where is it?	Understand the prevalence of this different content or messaging over time and/or during important events.
4	Who is it?	Investigate the different content for inauthenticity or illicit influence and attempt attribution.

These kinds of experiments would require a new approach to data collection, instead of relying on the limited accessibility provided by companies themselves, for example through online panels or targeted communities. Online panels enable researchers to gather new kinds of data about real online experiences. Panels, as used time and time again by polling companies, ensure explicit consent from volunteers for data collection and can provide a rich picture of online experience alongside human attitudes and behaviours. Another option would be for researchers to work closely with specific targeted communities at risk from disinformation efforts. These frontline groups could play a similar role to online panels in providing visibility on content and activities online that would otherwise be invisible to researchers, and providing insights into the penetration of disinformation campaigns into specific communities. ISD is trialling this approach in 2020 in the context of the US presidential election.

What Does This Mean for Tech Companies?

The research was not only set up to identify the perpetrators of disinformation, but also to evaluate the response of the tech companies on whose platforms such activities take place. 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.[52] 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 transparency over political advertising, to calls for more robust detection and removal of inauthentic behaviour.

ISD and partners from Avaaz, Digital Action, Counter Action, Who Targets Me and Mozilla combined forces to assess the enforcement of these commitments around the EU parliamentary elections through their research. In a separate policy evaluation and recommendations report, the successes, failures and challenges of these commitments are laid out in detail. There were some wins, but a significant selection of failures or disappointments in the responses from tech companies to the challenges posed during this election period.

The most concrete improvements from companies **came in political advertising transparency**, with Facebook attempting the most fully fledged ads library of the major companies signed up to the Code. Issues with API data access, false negatives, false positives and vague definitions of political content plagued the otherwise helpful improvements. Google and Twitter's attempts at transparency fell well short of the EU Commission and researchers' hopes, both in the timelines of their release and the scope of the data contained within those transparency efforts. **Transparency over 'co-ordinated inauthentic activity'** and meaningful **co-operation with researchers** proved limited, for the most part. Without genuine risks of sanction for non-compliance, the Code of Practice largely failed to engender a significant improvement in tech company responses to disinformation during the election.

The findings of ISD's accompanying policy report point to the failure of self-regulation around election integrity and disinformation and the pressing need for liberal democratic governments in Europe to upgrade their election laws and develop more systemic business regulation for tech companies for the digital age. Without clear definitions, laws and guidelines from democratically elected governments, we are relying on tech companies to create the norms for democracies in Europe from offices in Silicon Valley. Currently they are not meeting that challenge with satisfactory answers.

What Does This Mean for Law and Regulation?

ISD's attempt to identify and respond to illegal, deceptive and distortive activities around the election also brought under the spotlight another set of standards supposedly in place to protect citizens from harm and fraud inside and outside elections. ISD was interested to map how relevant or irrelevant existing law proved in preventing or countering the kinds of threat levelled at modern election processes and voters from the online domain.

ISD conducted an audit of relevant existing laws in the six countries under study for the European elections research, as well as relevant EU law. This included a scoping of:

- electoral law and electoral campaign law (funding, transparency, foreign and in-kind support, etc.)
- disinformation and defamation law
- hate speech and hate crime law
- harassment, stalking or doxing law
- privacy and data protection law.

Old Rules – New Tricks

Matching online findings from the project to this legal mapping, it became clear that certain challenges sit squarely within the remit of existing democratic law. The process underscored the availability of existing legal frameworks to challenge some types of malicious activity in the online world in theory, with the obstacle proving to be the effectiveness and accessibility of enforcement mechanisms for such laws in practice. Hate speech amplified by distortive and deceptive tactics is one example of activity discovered by ISD during the elections that is seemingly fairly well covered by existing laws. The application of these laws in digital circumstances is, however, challenged by the ease of perpetrator anonymity, the deletion of evidence by companies in some circumstances, and the difficulty in proving trauma and harm as a result of online incidents.

However, these challenges in enforcement should not inhibit attempts to enhance existing legal channels in confronting malicious uses of the internet. In many contexts, existing democratic law could prove a transparent and powerful tool to deal with certain kinds of malign activity online, if capacity, skills and processes in law enforcement and the judiciary can be upgraded for the digital age. In many instances, there is no need to reinvent the wheel entirely, focusing instead on new kinds of application of existing law. The UK Crown Prosecution Service made an initial symbolic gesture towards this approach in summer 2017 by adjusting guidelines on the treatment of hate crime in the UK courts to approach online hate crime incidents 'with the same robust and proactive approach used with offline offending'.[53] Given the clear intersection of disinformation with targeted hate crime and hate speech, as demonstrated in ISD's research, such approaches should not be forgone in the excitement around opportunities for new regulation.

Initial work has been done by scholars and legal experts such as Heidi Tworek to scope out potential models such as e-courts that might help to increase the speed, efficiency and availability of justice in the digital era.[54] ISD recommends that a thorough audit of the application and enforcement of existing law by European justice systems be initiated and practical suggestions be developed looking at how current legal and practical obstacles to its enforcement for online illegal activity might be addressed and remedied.

It is increasingly clear that our legal frameworks and processes haven't caught up with all things digital. Advances in digital technology bring both opportunities and challenges for the rule of law, including legislation relevant to elections. Digitalisation offers new solutions for increasing accountability by democratising access to information and enabling (some types of) transparency. Yet the new availability of opaque micro-targeting, anonymous political advertising, digital currencies vulnerable to foreign campaign funding, and a social media infrastructure tilted towards misinformation and sensationalism has rendered electoral law weak in protecting the integrity of the democratic process against manipulation. The lack of public rules for private companies now intricately connected with our information systems is a serious obstacle to free and fair democracy.

New Rules?

There are nonetheless emerging gaps in European and nation-state law or independent regulation to protect elections from new kinds of threat enabled by the digital infrastructure that has sprung up in recent years. Existing laws on foreign in-kind or direct contributions to election campaigns are insufficient to cover new digital modes of donating, from PayPal to Bitcoin. The role of electoral commissions, for one, requires a serious reboot in order to address new kinds of digital campaigning and funding. Election commissions are poorly equipped to analyse spending on digital advertising from unregistered campaign groups or the use of amplification technologies to achieve free exposure deceptively, through inauthentic account promotion. The UK Electoral Commission's review into the systems in place within the Brexit Party to vet funding contributions through PayPal exemplifies the limited powers of enforcement or review currently granted to such entities to identify and challenge potential malicious uses of technology to interfere with transparent and fair electoral processes.[55]

Public attitudes towards deceptive digital campaign tactics, whether used by parties, external groups or activists to support or attack parties or politicians, are still poorly understood. Electoral commissions can and must play a role in understanding public sensibilities towards amplification, micro-targeting and voter profiling tactics online, as well as mapping their compliance with new data protection standards under the General Data Protection Regulation in the EU. Recent public opinion research by Open Rights Group in the UK has started to unearth the level of public discomfort with the kinds of digital electioneering

that is possible in the modern day, along with relevant public opinion research by groups like Doteveryone and Open Knowledge Foundation.[56] All show a public appetite for more concerted government action to deal with the treat of harms directed at users online.

However, as the research above shows, threats to democracy do not take place solely in the domain of elections or electoral campaigning as traditionally defined: the distortion of information systems online is a threat to individual safety, public safety and democratic integrity during and after election periods. Private companies that are profiting from such threats to public life and free and fair information provision must be responsibly regulated, as so many other elements of the private sector are when engaging with citizens' rights and safety.

The recommendations in the accompanying policy report address each of these three channels for improving the public information space and the safety of online users in liberal democracies: improvements to existing legal process in the digital age; new powers for electoral commissions; and new regulation to ensure responsibility for safety and risk prevention in private technology companies. This is a domain in flux, contending with seismic shifts in how people consume information, form relationships and communicate across communities and borders. Those seeking to do harm have quickly adjusted to this new normal. For the most part, democratic governments, legislators and justice systems have yet to find their footing in the digital environment and all the potential harms and threats that it brings with it. Political parties have been slow to agree on what makes acceptable and unacceptable political campaigning online, with little understanding of public attitudes towards the new tools available to those in positions of influence and power online. Threats of digital distortion are not slowing, and practical responses from those with a duty to respect and protect rights must catch up quickly.

1 L. Cerulus, 'Steve Bannon Plans Right-wing Group in Europe', *Politico*, 22 July 2018, https://www.politico.eu/article/steve-bannon-the-movement-plans-right-wing-group-in-brussels/.

2 The Code of Practice was signed by the online platforms Facebook, Google, Twitter, Mozilla, and advertisers and others in the advertising industry in October 2018, and signatories presented their roadmaps to implement the Code. In May 2019, Microsoft subscribed to the Code of Practice and presented its roadmap.

3 QAnon is a conspiracy theory that surfaced in 2017, which is popular among some far-right extremist movements. See Anti-Defamation League, https://www.adl.org/resources/backgrounders/qanon. VDARE is designated as a white supremacist hate group by the Southern Poverty Law Center in the U.S., see Southern Poverty Law Center, 'VDARE', https://www.splcenter.org/fighting-hate/extremist-files/group/vdare.

4 See Facebook, *April 2020 Coordinated Inauthentic Behaviour Report*, May 2020, https://about.fb.com/news/2020/05/april-cib-report/.

5 Facebook remains an important platform for communication in Europe: Facebook reported 286 million daily active users in Europe in the first quarter of this 2019. See E. Schulze, 'Facebook's User Growth in Europe is Bouncing Back, Defying Stricter Privacy Laws', CNBC, 25 April 2019, https://www.cnbc.com/2019/04/25/facebook-q1-2019-user-growth-in-europe-is-bouncing-back-despite-gdpr.html.

6 Analysis of the official political party pages on Facebook across all six countries under study (France, Germany, the UK, Poland, Italy, Spain) shows a consistent trend of disproportionate right-wing populist and far-right party discussion and engagement about the elections online, with Italy the one exception to this rule. For this, we gathered all posts by the main national party pages on Facebook, and included in our analysis just those posts from these pages that mentioned the European elections through CrowdTangle. We analysed the numbers of posts, compared the total of likes, comments and shares within these posts, and the average of likes, comments and shares gained per post. These graphs show the total cumulative number of likes, shares and comments on all posts for each party.

7 Special Counsel Robert S. Mueller, III, 'Report On The Investigation Into Russian Interference In The 2016 Presidential Election Volume I of II', US Department of Justice, March 2019, https://www.justice.gov/storage/report.pdf.

8 Ibid., pp. 22-23.

9 K. Pruszkiewicz, J. Dauksza, K. Szczygieł and W. Cieśla, '"Po Trynkiewiczu też byś płakał?": Tak farma trolli broniła TVP po śmierci Pawła Adamowicza', *Newsweek Poland*, 28 October 2019; C. Davies, 'Undercover Reporter Reveals Life in a Polish Troll Farm', *Guardian*, 1 November 2019, https://www.theguardian.com/world/2019/nov/01/undercover-reporter-reveals-life-in-a-polish-troll-farm.

10 S. Mahtani and R. Cabato, 'Why Crafty Internet Trolls in the Philippines may be Coming to a Website Near You', *New York Times*, 26 July 2019, https://www.washingtonpost.com/world/asia_pacific/why-crafty-internet-trolls-in-the-philippines-may-be-coming-to-a-website-near-you/2019/07/25/c5d42ee2-5c53-11e9-98d4-844088d135f2_story.html.

11 J. Davey and J. Ebner, *The Great Replacement: The Violent Consequences of Mainstreamed Extremism*, ISD, July 2019, https://www.isdglobal.org/wp-content/uploads/2019/07/The-Great-Replacement-The-Violent-Consequences-of-Mainstreamed-Extremism-by-ISD.pdf.

12 J. Guhl, J. Ebner and J. Rau, *The Online Ecosystem of the German Far Right*, ISD, February 2020, https://www.isdglobal.org/isd-publications/the-online-ecosystem-of-the-german-far-right/.

13 The original tweet sharing the falsely labelled video was sent by @Timido762 on 22 April 2019. This tweet was then retweeted by the official Vox Moncloa-Aravaca account @VoxMoncloa. The tweet is no longer available but an archived copy of the tweet is available to researchers on request.

14 F Peinado, 'Una Red de Cuentas Falsas de Twitter Promueve a Vox en Campaña', El País, 26 April 2019, https://elpais.com/politica/2019/04/25/actualidad/1556203502_359349.html.

15 ISD, 'Covid-19 Disinformation Briefing No. 1', Institute for Strategic Development, 27 March 2020, https://www.isdglobal.org/isd-publications/covid-19-disinformation-briefing-no-1/.

16 S. Banaji and R. Bhat, 'WhatsApp Vigilantes: An Exploration of Citizen Reception and Circulation of WhatsApp Misinformation Linked to Mob Violence in India', blog, 11 November 2019, https://blogs.lse.ac.uk/medialse/2019/11/11/whatsapp-vigilantes-an-exploration-of-citizen-reception-and-circulation-of-whatsapp-misinformation-linked-to-mob-violence-in-india/.

17 D. Fidler, 'Disinformation and Disease: Social Media and the Ebola Epidemic in the Democratic Republic of the Congo', *Center on Foreign Relations*, 20 August 2019, https://www.cfr.org/blog/disinformation-and-disease-social-media-and-ebola-epidemic-democratic-republic-congo.

18 Computational Propaganda Project et al., 'UK General Election 2019: Digital Disruption by the Political Parties, and the Need for New Rules', ISD, December 2019, https://www.isdglobal.org/isd-publications/uk-general-election-2019-digital-disruption-by-the-political-parties-and-the-need-for-new-rules-joint-paper/.

19 ISD has not included the direct link to the content under discussion here due to privacy concerns and the graphic nature of content on the site in question. ISD is able to provide archived copies of this thread on 4Chan to researchers on request.

20 @Santi_ABASCAL, Twitter, 15 April 2019, https://twitter.com/Santi_ABASCAL/status/1117890168340586497.

21 Philosophia Perennis speculated about the fire being an arson attack that is covered up, Muslims celebrating the fire and supposed plans for the cathedral to be replaced by a minaret: Philosophia Perennis, 'Notre Dame – Suchergebnisse', April and May 2019, https://philosophia-perennis.com/?s=notre+dame.

22 Tichys Einblick accused social media users with 'Arabic-Muslim names' of celebrating the fire. See Tichys Einblick, 'Notre Dame: Europas Seele brennt', 15 April 2019, https://www.tichyseinblick.de/kolumnen/aus-aller-welt/notre-dame-europas-seele-brennt/.

23 See https://www.facebook.com/AfD.KVHerne/posts/2452059264845962.

24 See https://www.facebook.com/AfDfuerNRW/posts/2637694759635822.

25 Der Tagesspiegel, 'Chebli deaktiviert Facebook-Account wegen Hass-Nachrichten', 22 October 2018, https://www.tagesspiegel.de/berlin/berliner-staatssekretaerin-chebli-deaktiviert-facebook-account-wegen-hass-nachrichten/23216792.html.

26 @ArifPirincci, Twitter, replying to @SawsanChebli tweet of 10 May 2019. A copy of the archived tweet is available to researchers on request, as the the @ArifPirincci account has since been removed from Twitter.

27 The YouTube channel of Tim Kellner has more than 30 videos about Swansan Chebli, with at least two having more than 300,000 views and two more having over 260,000 views. Links to these videos are archived and are available to researchers on request, but have not been included here due to the potentially offensive content of the material. Tim Kellner was acquitted in court in February 2020 of charges of insult against Swansan Chebli. Von Wiebke Ramm, '"Islamische Sprechpuppe" - YouTuber nach Hasskommentaren freigesprochen', *Der Spiegel*, 27 February 2020, https://www.spiegel.de/politik/deutschland/sawsan-chebli-youtuber-timm-k-nach-hasskommentaren-vor-gericht-freigesprochen-a-90c56d9c-56e5-4e9b-bdad-3d2981bd2898.

28 P. Moullot, 'No, Sibeth Ndiaye did not say "Yes, the Girl is Dead" to Confirm the Death of Simone Veil', Liberation, 1 April 2019, https://www.liberation.fr/checknews/2019/04/01/non-sibeth-ndiaye-n-a-pas-dit-yes-la-meuf-est-dead-pour-confirmer-la-mort-de-simone-veil_1718694.

29 ISD has not included any images of graphic hate speech in this report but archived copies of these examples are available to researchers on request.

30 See Oxford Internet Institute's research on junk news in the UK general election 2019: N. Marchal et al., 'Junk News & Information Sharing During the 2019 UK General Election', Oxford Internet Institute, 2019, https://comprop.oii.ox.ac.uk/wp-content/uploads/sites/93/2019/12/UK-Data-Memo_FINAL.pdf; and Demos' report on the nature of Internet Research Agency campaigns in the UK: 'New Demos Analysis finds Russian Influence Operations on Twitter Targeted at UK were most Visible when Discussing Islam', press release, Demos, 1 November 2018, https://demos.co.uk/press-release/new-demos-analysis-finds-russian-influence-operations-on-twitter-targeted-at-uk-were-most-visible-when-discussing-islam/.

31 J. Shotter, 'Polish nationalists protest at law on restitution of Jewish property', *Financial Times*, 12 May 2019, https://www.ft.com/content/c47fcb02-749d-11e9-bbad-7c18c0ea0201.

32 Marsz Niepodległości , 'NIE dla żydowskich roszczeń – petycja do polskiego rządu', 15 February 2019, https://medianarodowe.com/nie-dla-zydowskich-roszczen-petycja-do-polskiego-rzadu/

33 European Commission, 'Europeans' Attitudes on EU Energy Policy', Eurobarometer Special 492, 11 September 2019, https://ec.europa.eu/info/news/eurobarometer-survey-confirms-public-support-energy-policy-objectives-2019-sep-11_en.

34 Tagesschau, '86 Prozent Sagen, der Mensch sei Schuld', 5 May 2019, https://www.tagesschau.de/inland/deutschlandtrend-1645.html.

35 D. Scally, 'Germany's CDU Leader Finds Election Scapegoat: A 27-year-old YouTuber', *The Irish Times*, 28 May 2019, https://www.irishtimes.com/news/world/europe/germany-s-cdu-leader-finds-election-scapegoat-a-27-year-old-youtuber-1.3906948.

36 S. Schaller and A. Carius, *Convenient Truths:*

Mapping Climate Agendas of Right-wing Populist Parties in Europe, Adelphi, 2019, https://www.adelphi.de/en/publication/convenient-truths.

37 '14. Mai 2019 – Bundestagsabgeordnete setzen Fakten gegen CO2-Hysterie und Klima-Aktionismus', EIKE, 4 May 2019, https://www.eike-klima-energie.eu/2019/05/04/14-mai-2019-bundestagsabgeordnete-setzen-fakten-gegen-co2-hysterie-und-klima-aktionismus/.

38 The Heartland Institute co-sponsored the EIKE climate conference at the NH Munich East Conference Center in 2018, [https://www.heartland.org/news-opinion/news/skeptic-climate-scientists-gather-for-conference-in-germany-before-un-meeting]. The two organisations co-hosted the "Climate Reality Forum" in Madrid in 2019, [https://climaterealityforum.com/about/].

39 The Heartland Institute hosts an annual "International Conference on Climate Change", aimed at bringing together those who challenge the scientific consensus on the causes, consequences, and implications of climate change. According to the conference's website at the time of writing, speakers for the event are "brave enough to publicly tell the truth about the CLIMATE DELUSION", and "so-called academic elites—with help from the mainstream media—have made a living peddling the delusion that human activity is causing an impending global catastrophe", [https://climateconference.heartland.org/about/]. According to Greenpeace project 'Exxonsecrets', the Heartland Institute has received $676,500 from the oil and gas company ExxonMobil since 1998, [https://exxonsecrets.org/html/orgfactsheet.php?id=41#src26]. The Heartland Institute does not dispute these claims, but clarifies on its website that "the gifts [from ExxonMobil] never exceeded 5 percent of Heartland's annual budgets", [https://www.heartland.org/about-us/reply-to-critics/index.html].

40 Some of the following data was published as part of a joint investigation with Unearthed: K. Connolly, 'Germany's AfD Turns on Greta Thunberg as it Embraces Climate Denial', *Guardian*, 14 May 2019,

41 As only four AfD posts from April 2017 to April 2018 were found to have used the keyword 'Greta', we are confident that most of these 791 posts mentioning 'Greta' from April 2018 to April 2019 are in fact about Greta Thunberg.

42 We excluded terms like 'Greta-Mania' and 'Greta-disciples' as they might have biased the results, even though they are used in a similar way.

43 Connolly, 'Germany's AfD Turns on Greta Thunberg as it Embraces Climate Denial'; D. Kayha, 'German Far Right Targets Greta Thunberg in Anti-climate Push', *Unearthed*, 14 May 2019, https://unearthed.greenpeace.org/2019/05/14/germany-climate-denial-populist-eike-afd/.

44 @Frank_Pasemann, Twitter, 8 March 2019. An Archived copy of the tweet is available to researchers on request.

45 Image found on Facebook, *Martin Schiller*,23 March 2019. An archived copy of the image on Facebook is available to researchers on request, but the image has not been included here as it is potentially offensive.

46 Image found on Telegram, *Lutz Bachmann offiziell*, 31 March 2019. An archived copy of the image is available to researchers on request, but the image has not been included here as it is misleading and potentially offensive.

47 Bruno Le Salé 'L'ARNAQUE DE LA MARCHE POUR LE CLIMAT: GRETA THUNBERG', 6 April 2019. An archived copy of the video posted on YouTube is available at http://archive.fo/rKWXo.

48 The search query used to conduct this analysis identified posts that contained any of the following keywords/phrases: "CO2Kult, Ersatzreligion, Klimasekte, Klimapanik, Klima-Panik, Klimawandelpanik, Klimawandel-Panik, Klimahysterie, Klima-Hysterie, Klimawandel-Hysterie, Ersatzreligion, Klimawandelhysterie, Klimareligiöse, Klimareligiösen, Klimagehirnwäsche, Klimareligion, Klimaschwindel, Klimatismus, Ökoterrorismus".

49 European Parliament, 'European Parliament 2019–2024', 2 July 2019, https://www.europarl.europa.eu/election-results-2019/en.

50 BBC News, 'European Election 2019: Results in Maps and Charts', 27 May 2019, https://www.bbc.co.uk/news/world-europe-48417191.

51 Other research has called similarly for more of a focus on cross-platform dynamics of campaigns, for example T. Wilson and K. Starbird, 'Cross-Platform Disinformation Campaigns: Lessons Learned and Next Steps', Harvard Kennedy School, *Misinformation Review*, 14 January 2020, https://doi.org/10.37016/mr-2020-002. Initial attempts to bring tabloid or print media together with social media to assess patterns of disinformation sharing include A. Chadwick, C. Vaccari and B. O'Loughlin, 'Do Tabloids Poison the Well of Social Media? Explaining Democratically Dysfunctional News Sharing', *New Media & Society*, 20 April 2018, https://hdl.handle.net/2134/33261.

52 Microsoft also signed up in May 2019: European Commission, 'Codes of Practice on Disinformation', September 2018, https://ec.europa.eu/digital-single-market/en/news/code-practice-disinformation.

53 A. Cowburn, 'Lockdown on "Corrosive" Online Hate Crimes Launched by Crown Prosecution Service', *Independent*, 21 August 2017, https://www.independent.co.uk/news/uk/politics/hate-crimes-social-media-crown-prosecution-service-home-office-prejudice-a7903166.html.

54 H. Tworek, R. Ó Fathaigh, L. Bruggeman and C. Tenove, *Dispute Resolution and Content Moderation: Fair, Accountable, Independent, Transparent, and Effective. Transatlantic High Level Working Group on Content Moderation Online and Freedom of Expression*, Transatlantic Working Group, 14 January 2020, https://www.ivir.nl/publicaties/download/Dispute_Resolution_Content_Moderation_Final.pdf.

55 Electoral Commission, 'Recommendations for The Brexit Party – Financial Procedures for Incoming Funds', June 2019, https://www.parliament.uk/documents/commons-committees/culture-media-and-sport/Background-Electoral-Commission-FOI.pdf.

56 Open Rights Group, 'Poll Finds Strong Support for Policies Combating Nefarious Online Campaigning Activity', press release, 22 November 2019, https://www.openrightsgroup.org/press/releases/2019/poll-finds-strong-support-for-policies-combating-nefarious-online-campaigning-activity; Doteveryone, 'People, Power and Technology: The 2020 Digital Attitudes Report', 11 May 2020, https://www.doteveryone.org.uk/project/people-powertech/; Open Knowledge Foundation, 'Brits Demand Openness from Government in Tackling Coronavirus', blog, 5 May 2020, https://blog.okfn.org/2020/05/05/brits-demand-openness-from-government-in-tackling-coronavirus/.

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