EC ACTION PLAN ON DISINFORMATION

GOOGLE JANUARY REPORT

This document comprises Google's first monthly report in response to the European Commission's Action Plan on disinformation, mapping out progress towards the implementation of the Code of Practice. It covers updates and actions for the month of January 2019 and builds on our initial 2018 baseline report on the CoP commitments, which we submitted to the European Commission and published in January 2019 (see Annex).

Since the start of the year, we have continued to work towards implementing and enforcing the commitments we made in the Code of Practice on Disinformation.

ADVERTISING:

Continuing to protect our users from the behaviors of malicious actors:
Over the course of January, we continued to enforce our policies against behaviors commonly used by malicious actors, who aim to deceive or scam our users, including (but not limited to) propagators of disinformation.

With respect to misrepresentative content, we detected 48,642 EU-based Google Ads accounts in violation of our misrepresentation policies in January 2019. We did not find any EU-based AdSense publishers in violation of our misrepresentative content policies in the monitored period. As per our usual process, content associated with those accounts was demonetized, meaning that we blocked ads from serving to users. In some egregious cases, advertiser accounts were disabled. Note that not all misrepresentation violations are associated with disinformation campaigns, but are meant to prevent users from being misled by the content they see from advertisers and publishers. Additionally, abuse patterns differ from product to product, so you should refer to the respective policies to see more information about the content we enforce on.

The Google Ads Misrepresentation policy, which applies globally to advertisers, prohibits ads that intend to deceive users by excluding relevant information or giving misleading information about products, services, or businesses. This includes making false statements about your identity or qualifications, or making false claims that entice a user with an improbable result. In addition to specific policies against misrepresentation, we have a number of other policies that protect our users from different types of abusive content -- including content that could be considered disinformation. Sometimes, these policies lead us to take action against content that is described as "fake news" in public debate. These

policies include prohibitions on monetization of valuable inventory policy and dangerous or derogatory content, both of which apply globally.

A country-by-country breakdown of the violations is available below:

EU member Countries	Google Ads accounts with Misrepresentation violations (by billing country)
AT	104
BE	84
BG	120
CY	81
CZ	221
DE	849
DK	88
EE	12295
ES	615
FI	55
FR	666
GB	16679
GR	57
HR	100
HU	155
IE	88
IT	6226
LT	51
LU	15
LV	52
MT	27
NL	275
PL	622
PT	244
RO	8652
SE	104
SI	51
SK	66
TOTAL	**48642**

We also continue to monitor for and take action on content that is insufficiently original and provides no value to users. In Google Ads, we detected and took action on 3,258 EU-based accounts in violation of our policy on insufficiently original content. And in AdSense, we detected and took site- or account-level action on 205 EU-based AdSense publishers who were trying to serve ads on non-valuable inventory on their site. Non-valuable inventory refers to websites that have more ads than content; use mirroring, framing, scraping or rewriting of content from other sources without adding value; automatically generated content without manual review or curation; hosted ad pages or pages without content; or pages that don't follow our Webmaster Quality Guidelines. As per our usual process, content associated with those accounts was demonetized, meaning that Google booked ads were stopped from serving to users and sites were prevented from showing ads on them through our system. In some egregious cases, advertiser or publisher accounts were disabled.

We also take action on individual pages that violate our AdSense policy, though those actions are not reflected in this month's report. As is the case for misrepresentation, not all of these policy violations are associated with disinformation campaigns, but are made to make sure that we are only enabling monetization for content that is valuable to users. And again, abuse patterns differ from product to product, so you should refer to the respective policies to see more information about the content we enforce on. A country-by-country breakdown is available below:

EU member countries	Google Ads accounts with Insufficient Original Content violations (by billing country)	AdSense accounts with Valuable Inventory violations (by billing country)
AT	43	6
BE	40	0
BG	26	8
CY	53	2
CZ	100	6
DE	515	58
DK	43	6
EE	9	0
ES	246	40
FI	17	1
FR	279	4
GB	1074	0
GR	29	0
HR	15	1

HU	14	8
IE	64	0
IT	214	2
LT	11	2
LU	6	0
LV	18	0
MT	11	0
NL	147	18
PL	145	10
PT	38	8
RO	38	10
SE	39	1
SI	6	2
SK	18	12
TOTAL	**3258**	**205**

POLITICAL ADVERTISING:

New policies mandating identity verification for advertisers who would like to run Election Ads:

On 29 January, we published our new ads policy for election ads that applies in all EU member states. This policy, which is available in 25 EU languages, requires that advertisers must be verified in order to run ads that feature a political party, or a current elected officeholder or candidate for the EU Parliament. These ads will feature in-ad disclosure of who paid for the ad and be included in our political ads transparency report. The announcement was made via a blog post and featured in our Change Log in our Advertising Policies Page.

To be verified, advertisers wanting to run EU parliamentary election ads on Google's platforms will be required to provide documentation to show they're an EU-based entity or citizen of an EU member country. More information about verification will be published in February 2019.

We expect to start publishing our transparency report as advertisers begin to run ads as we approach the elections.

INTEGRITY OF SERVICES:

As outlined in our 2018 baseline report, Google takes very seriously any threats to the integrity of our services. We operate policies that apply evenly across all types of actors and accounts, regardless of whether they are automated or human. They prohibit impersonation, misrepresentation, abusive account creation, engagement abuse and influence operations. You can find more detail about the operation of those policies in the annex below. We will continue to report on the implementation of these policies and will examine including relevant quantitative measures related to this area in a forthcoming report

Dedicated elections teams clamping down on abuse: Our work to prevent election-related abuse ahead of and during elections means teams and subject matter experts across Google are working closely together. These teams, many of whom are based in Europe, are trained to identify and stop a wide range of potential abuse from State Sponsored phishing attacks to attempts to alter Maps so people can't find their polling station. We're also constantly working to get people to authoritative content and improving our systems to combat the effects of disinformation. We're staffed so we can get ahead of abuse, clamp down on malicious activity, and react rapidly to breaking threats. Google's Trust & Safety teams have long worked in partnership with Google Security's Threat Analysis Group (TAG) and our partners at Jigsaw to identify malicious actors, disable their accounts, warn our users about them, and share intelligence with other companies and law enforcement officials.

ELECTION ACTIVITY:

New trainings and initiatives to protect the integrity of European elections:
- *Trainings for journalists including to identify and tackle disinformation:* Over the course of January 2019, we have provided 1,034 journalists across Europe with in-person training on a range of digital tools to help strengthen their research and reporting skills and equip them with best practice techniques in the fields of data journalism, visualisation and verification. We've so far hosted free workshops for local and national newsrooms, journalism schools and for freelance members of journalism associations.

- *In-person trainings from Google experts for those involved in elections:* Because it can be hard to know how to be safe online, we're running in-person and online security trainings for those most at risk around the upcoming elections. Like how to use our Advanced Protection Program, which gives Google's strongest security for those that need it most. So far we've trained close to 1,000 campaign and election officials, journalists and people from election-related NGOs in Europe in-person, so they can learn which security tools they need and how to use them. Our goal is to support these groups in keeping their information secure and enable them to publish freely so that people can access the stories, debates, policies and results when it matters most. In January 2019, we organized 15 in-person training sessions for

academics, practitioners, policy-makers and industry representatives, focusing on the EU 2019 elections. These trainings cover using Google and YouTube tools to communicate to voters as well as election integrity and security issues. Overall, we have had 258 participants in these trainings in January 2019. As the education on security risks is rather low among campaigners and those involved in elections, we welcome the Commission's help in scaling education around the key security precautions.

- *Media literacy programs:* Google continues supporting media literacy programs across the EU.
    - For instance, in January, The Student View, a Google.org grantee from the UK, scaled their pop–up newsrooms to schools across London and Manchester, training 60 schoolchildren from underserved communities as local journalists. During these workshops, pupils were trained to spot misinformation and research local issues to write about in pairs.
    - Until now, we have trained more than 1 million people in person through our Be Internet Citizens and Be Internet Awesome in the EU. Be Internet Citizens has been designed to teach teenagers about media literacy, critical thinking and digital citizenship, with the aim of encouraging young people to have a positive voice online. Be Internet Awesome is a program from Google that aims to help kids be safe, confident explorers of the online world. In the subsequent monthly reports we will be posting actual numbers related to these activities.

Additional Election Work:
- *Project Shield protection against DDOS attacks***:** Journalists, campaigns and political parties, NGOs and election monitoring groups ensure people can stay informed during election periods. It's never been more necessary to defend these groups from digital attacks that can exploit many thousands of computers to overwhelm a website's servers and take it offline—preventing voters from getting official information when they need it most. Project Shield uses Google's infrastructure to protect independent news sites from distributed denial of service attacks (DDoS) and Jigsaw is now offering strong, free DDoS-protections to the organizations across Europe that are vital to free and fair elections. You can find out more about Jigsaw and apply for Shield protection here.

Encouraging and advancing research on synthetic media ('deep fakes')
Over the last few years, research on using neural networks to simulate a human voice has significantly advanced, including through the publication of open models, leading to growing capabilities for creators and companies around the world to synthesize increasingly realistic, human-like speech. While these developments have promising potentials in fields like translation, entertainment, and empowering those affected by speech or reading disabilities, they could also lead to impersonation or spoofing – a type of content sometimes referred to as audio 'deep fakes'.

To help advance research on detecting and countering these types of malicious behaviors, Google AI and Google News Initiative have partnered to create a body of synthetic speech containing thousands of phrases spoken by our deep learning 'Text-To-Speech' (TTS) models. On January 31st, we have made this dataset available to all participants in the third party & independent 2019 ASVspoof challenge. This open challenge invites researchers all over the globe to submit countermeasures against fake (or "spoofed") speech, with the goal of making automatic speaker verification (ASV) systems more secure. By training models on both real and computer-generated speech, ASVspoof participants can develop systems that learn to distinguish between the two. The results will be announced in September at the 2019 Interspeech conference in Graz, Austria.

ANNEX - EU CODE OF PRACTICE ON DISINFORMATION

Google baseline report

On April 26th 2018, the European Commission published a Communication outlining that "*The exposure of citizens to large scale disinformation, including misleading or outright false information, is a major challenge for Europe*". In response to this threat, the Communication laid out an Action Plan against disinformation, part of which involved the development of a self-regulatory Code of Practice on disinformation for online platforms and the advertising industry, in order to increase transparency and to better protect European citizens.

In May 2018, the Commission convened a "*Multistakeholder Forum on Disinformation"* to draft this self-regulatory Code of Practice. Within the Forum, a "Working Group" comprising Google, other online services, and representatives of the advertising industry and major advertisers, was tasked with drafting the Code.

The Code was published on September 26th, along with the opinion of a "Sounding Board" composed of representatives of the media, academia and civil society. On October 16th, signatories (including Google) provided their formal subscriptions to the Code, including a list of the commitments to which they adhere, a table listing company best practices, and future milestones for the overall implementation of the Code in the European Union. In addition, a number of the signatories (including Google) have outlined a number of specific activities they will carry out with a view to, and in advance of, the European Parliament elections of 2019, focusing on the security and integrity of elections.

On December 5th, the Commission [adopted a Report](#) assessing the progress made in the implementation of the actions set out in the April Communication on online disinformation. In the Report, the Commission considers that "*the Code provides an appropriate framework for pursuing the objectives for the Code set out in the Communication*", and that "*the Code is consistent with the Commission's Principles for Better Self-and Co-Regulation*". It however emphasises that "*Signatories must now work to ensure the robust implementation (...) of the policies and actions they have identified (...) on a full, effective, and timely basis*".

As a first step, Code signatories are to share "*up-to-date information about measures taken towards the implementation of the Code by the end of 2018*", to be published in January 2019. Moreover, "*between January and May 2019, the Commission will carry out a targeted monitoring of the Code on a monthly basis to verify that effective policies with particular pertinence to the integrity of the electoral processes are in place before the European elections*".

Throughout 2018, Google acknowledged the Commission's concerns with regards to disinformation and was a committed participant to the consultation processes and

multi-stakeholder forums it set out. We fully engaged with the drafting of the Code of Practice on disinformation (see blogpost), and are pleased that the Commission agrees that it represents an appropriate framework to pursue the objectives it set out in April. We intend to continue to work together with the European Commission in the run up to the 2019 European Parliamentary elections and beyond (see blogpost).

This report represents Google's first response to the request outlined by the Commission in its December 5th report. It covers measures taken by Google and YouTube towards the implementation of the Code of Practice.

For the purposes of this initial report, we organize our response in five sections, one for each commitment within the Code. Each section comprises a reminder of the commitments we made as part of the Code, and outlines actions we undertook to further those goals.

This initial report will be followed by others between January and May, where we will outline more implementation details and metrics as elections-related product and policy changes are rolled out.

CONTENT:

1. SCRUTINY OF ADS PLACEMENTS

Signatories to the Code of Practice agreed to the following Commitment:
1. *Relevant Signatories commit to deploy policies and processes to disrupt advertising and monetization incentives for relevant behaviours, such as misrepresenting material information about oneself or the purpose of one's properties. These policies and processes can include, for example, the restriction of advertising services or limiting paid placements, and could potentially take place in partnership with fact-checking organizations. Such policies and processes may, as appropriate:*
- *(a) Promote and/or include the use of brand safety and verification tools;*
- *(b) Enable engagement with third party verification companies;*
- *(c) Assist and/or allow advertisers to assess media buying strategies and online reputational risks;*
- *(d) Provide advertisers with necessary access to client-specific accounts to help enable them to monitor the placement of ads and make choices regarding where ads are placed.*

How we are complying with this Commitment by end of 2018:
In order to protect our users and advertisers who trust our services with the distribution of their ads, we have rolled out a number of policies and processes geared towards disrupting the monetization incentives of malicious and misrepresentative actors.

Publishers and advertisers who wish to use our ad networks are obligated to comply with our policies against misrepresentation that apply globally. The Google Ads Misrepresentation policy, which applies globally to advertisers, prohibits ads that intend to deceive users by excluding relevant information or giving misleading information about products, services, or businesses. This includes making false statements about your identity or qualifications, or making false claims that entice a user with an improbable result. Similarly, publishers who are members of our AdSense network are prohibited via our Misrepresentative Content policy from misrepresenting, misstating, or concealing information about themselves, their content or their primary purpose of their web property.

Google Ads Misrepresentation policy (available here):
We don't want users to feel misled by ads, so we strive to ensure ads are clear and honest, and provide the information that users need to make informed decisions. We don't allow ads or destinations that intend to deceive users by excluding relevant information or giving misleading information about products, services, or businesses.
This policy applies globally to all Ads served by Google.

AdSense Misrepresentative content policy (available here):
Users don't want to be misled by the content they engage with online. For this reason, Google ads may not be placed on pages that misrepresent, misstate, or conceal information about you, your content or the primary purpose of your web property.
This policy applies globally for all webmasters using Google Ads to monetize their websites.

In addition to specific policies against misrepresentation, we have a number of other policies that protect our users from different types of abusive content -- including content that could be considered disinformation. Sometimes, these policies lead us to take action against content that is described as "fake news" in public debate. These policies include prohibitions on monetization of valuable inventory policy and dangerous or derogatory content, both of which apply globally.

Valuable inventory policy (available here):
a
Examples of unacceptable pages include but are not limited to:
- Mirroring, framing, scraping or rewriting of content from other sources without adding value;
- Pages with more advertising than publisher-provided content;
- Automatically generated content without manual review or curation;
- Hosted ad pages or pages without content;
- Pages that don't follow our Webmaster Quality Guidelines.
This policy applies globally to all Ads served by Google.

Google Ads Inappropriate content policy (available here):
We value diversity and respect for others, and we strive to avoid offending users, so we don't allow ads or destinations that display shocking content or promote hatred, intolerance, discrimination or violence. This includes dangerous or derogatory content, such as:

- Content that incites hatred against, promotes discrimination of, or disparages an individual or group on the basis of their race or ethnic origin, religion, disability, age, nationality, veteran status, sexual orientation, gender, gender identity or other characteristic that is associated with systemic discrimination or marginalisation
- Content that harasses, intimidates or bullies an individual or group of individuals
- Content that threatens or advocates for harm on oneself or others
- Content that seeks to exploit others

This policy applies globally to all ads served by Google.

To make sure that advertisers and publishers on our networks comply with our policies, we continuously monitor our advertising networks for compliance with these policies around the world, using a combination of algorithmic and human reviews. You will find an overview of our efforts in this table:

2017 Takedowns

Overall Policy Violations	3.2 billion ads disapproved 2M pages / month blocked from the publisher network 320,000 publishers terminated from our network 90,000 websites blacklisted from our network 700,000 mobile apps blacklisted from our network
Unoriginal content, including "scraping," duplicating and copying content from other sites	12,000 websites blocked from the publisher network (up from 10,000 in 2016)
Misrepresentative content policy	Of the 11,000 websites we reviewed for potentially violating the misrepresentative content policy, we blocked over 650 of those sites and terminated 90 publishers from our network.

We also take steps to make sure that our customers are aware of all the content policies that apply to them and how they are enforced. Our [Google Ads](#) and [AdSense](#) content policies are publicly available in those products' respective help centers, including translations into many of the languages spoken in EU member states. Those same help centers also provide information about how we enforce our policies and what customers should do if they want to fix a policy violation. Additionally, we maintain a [change log](#) for Google Ads that keeps track of any updates to our policies.

On YouTube, creators who meet our eligibility criteria can apply to join the [YouTube Partner Program](#), which makes their videos eligible to run advertising and earn money through other sources on the platform. Not only does a channel have to apply, but just because creators are eligible to monetize does not mean that all of their content is necessarily eligible to serve ads and earn money. We have a set of [Advertiser-friendly content guidelines](#) and disable advertising from running on videos that violate these policies. Our advertisers are understandably sensitive to what content their ads are placed against. These policies ensure that we strike a balance where advertisers and creators are both able to thrive.

Over the last few years, we've taken steps to strengthen our requirements for monetization so spammers, impersonators, and other bad actors can't hurt our ecosystem or take advantage of good creators. To apply for membership in the YouTube Partner Program, we increased the thresholds needed for channels to be deemed eligible: channels must have generated 4,000 watch hours in the previous 12 months and have more than 1,000 subscribers. After they apply, our review team will ensure their channel has not run afoul of our monetization, content and copyright policies prior to admitting them to the program. As a part of this process, we believe that only creators with sufficient history and demonstrated advertiser safety will receive access to ads and our other monetization products. In changing these thresholds, we have significantly improved the protections we had in place against impersonating creators.

In addition to these policies, we provide advertisers with additional controls and help them exclude types of content that, while in compliance with our policies, may not fit their brand or business. These controls let advertisers apply content filters or exclude certain types of content or terms from their video, display and search ad campaigns, advertisers can exclude content such as politics, news, sports, beauty, fashion and many other categories. These categories are listed in our Google Ads Help Center[1].

<u>What to look for in subsequent reports</u>
Over the course of 2019, we will provide new metrics (including metrics specific to EU member states) on our enforcement of our advertising policies (including on misrepresentative content), and continue to upgrade our enforcement mechanisms in anticipation of the European elections.

---

2. POLITICAL ADVERTISING AND ISSUES-BASED ADVERTISING

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

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

4. Relevant Signatories commit to use reasonable efforts towards devising approaches to publicly disclose "issue-based advertising". Such efforts will include the development of a working definition of "issue-based advertising" which does not limit reporting on political discussion and the publishing of political opinion and excludes commercial advertising. Given the implications related to freedom of expression, Signatories encourage engagement with expert stakeholders to explore approaches that both achieve transparency but also uphold fundamental rights. The work to develop this definition shall not interfere with the areas covered by advertising self-regulatory organisations.


How we are complying with these Commitments by end of 2018:

→ **Commitment #2 - ads clearly distinguishable from editorial content**: Across all of our products and surfaces and in all European Union Member States, advertisements are clearly labeled as such. This clarity is paramount to the trust of our users.
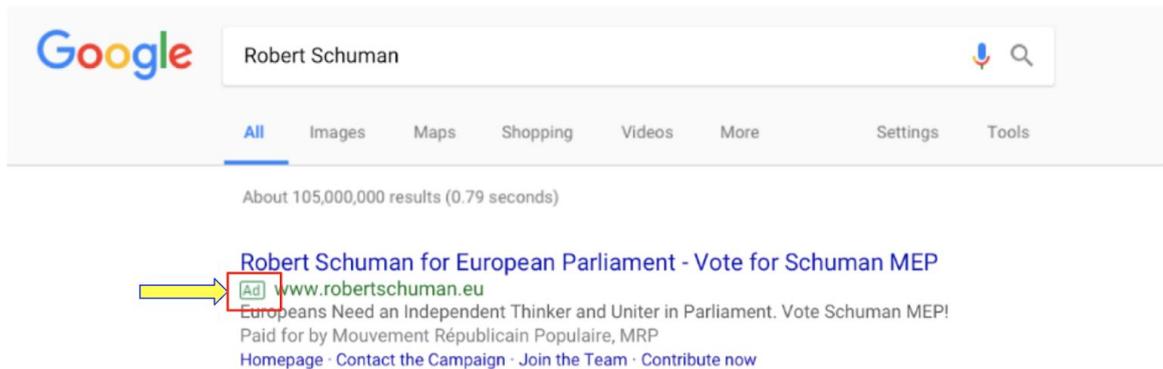
→ **Commitment #3 - public disclosure of political advertising:** As announced in November of 2018, we are hard at work rolling out products and programs to help people get important voting information ahead of the European Union Parliamentary elections. We will require that ads that mention a political party, candidate or current officeholder make it clear to voters who's paying for the advertising. We'll also introduce a new process to verify EU election advertisers to make sure they are who they say they are. And there's more to come: we'll introduce an EU-specific Election Ads Transparency Report and searchable ad library to provide more information about who is purchasing election ads, whom they're targeted to, and how much money is being spent, and make it available via download to the public. Our goal is to make this information as accessible and useful as possible to citizens, practitioners, and researchers.
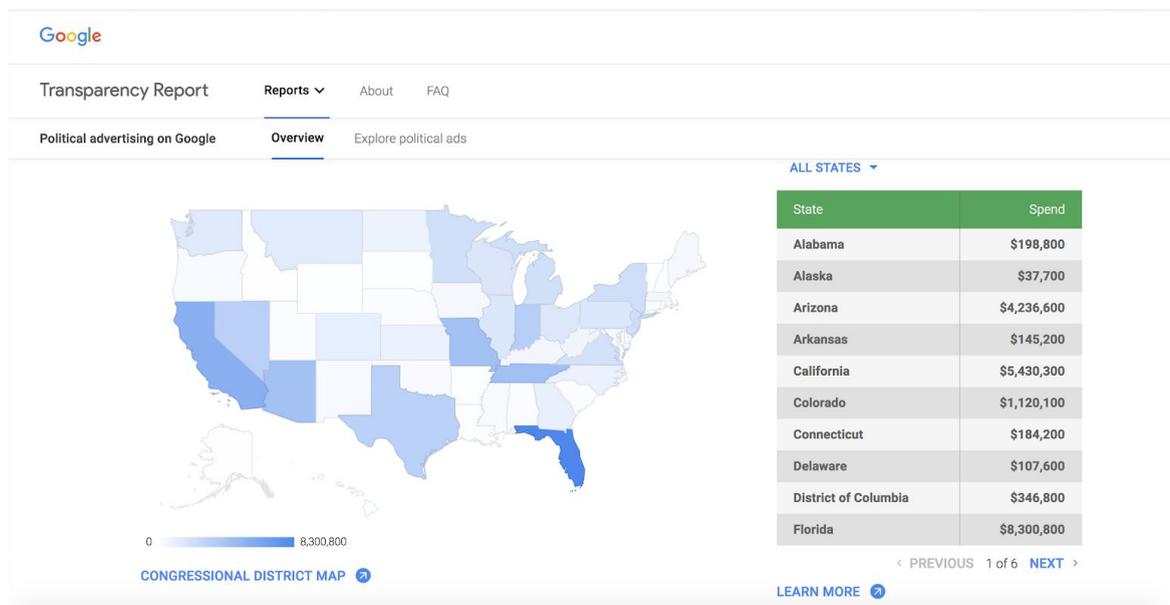
Example: Political advertising transparency report in the U.S (please note that the European Transparency might differ).

This work builds on our experience during the United States 2018 midterm elections, where Google processed for verification more than 2,200 advertisers whose ads featured a Federal candidate or current Federal officeholder. More than 20% of all the political ads for 2018 were submitted during the seven days leading to and including election day, and tens of thousands of advertisement creatives were submitted for Google's creative repository of political advertisements. Google's processes prevented many election ads from unverified advertisers from showing during the U.S. election. We will be bringing these processes to the EU for the election in 2019, along with many learnings from the United States election.

Thanks to the transparency database, journalists in the United States were able to more fully cover candidates' advertising strategies and educate voters on candidates' policy positions, outreach efforts and voter engagement strategies (see for instance here). Voters were able to be more aware of candidates' efforts to persuade voters, and therefore have a more comprehensive view of a candidate when deciding which candidate to support with their vote.

Additional to advertiser verification, we will be requiring advertiser disclosures for political advertisers in the EU for the 2019 election. As we enforce all of our policies, we will be using a combination of automated classifiers and manual review to ensure that our standards are met.

Below is an example of a mock Google search ad for a candidate with the disclosure in the text of the search ad:
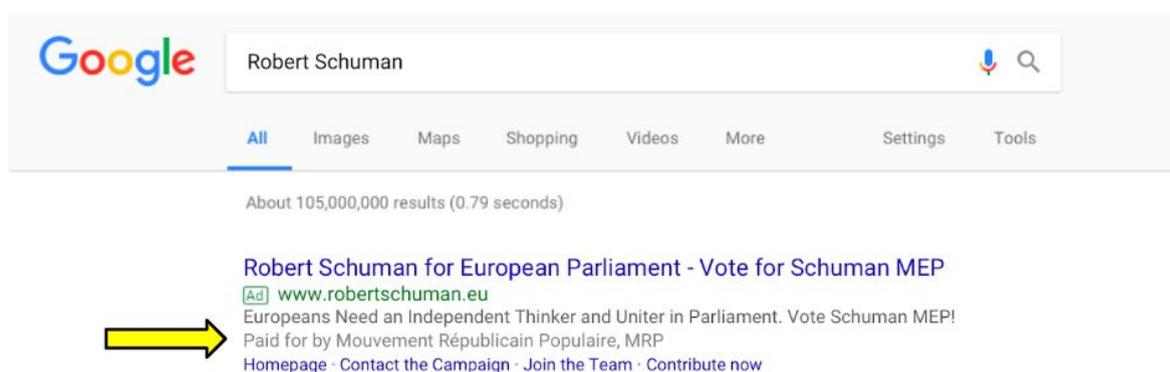
Below is an example of a mock YouTube ad for a candidate with the disclosure in the text of the video ad:
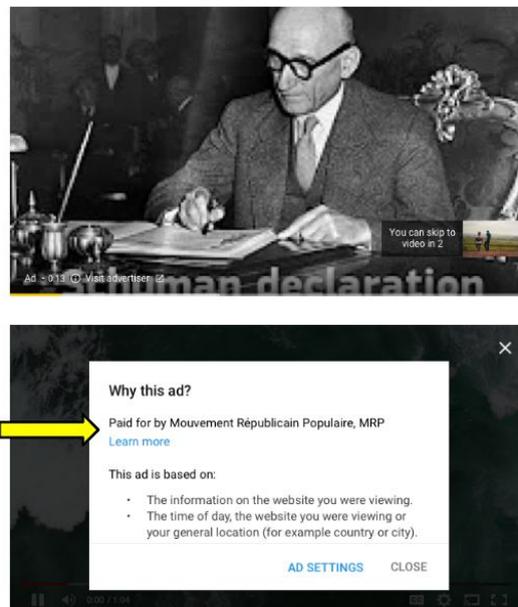
→ **Commitment #4 - public disclosure of issue-based advertising:** We are investigating solutions that would improve transparency of strictly issue-based political advertising (ie not specifically related to an EU election campaign). It remains a challenge amongst the signatories and others in the civil or public sector to find a workable definition of a "political issue," and implementation of such a definition in practice, that would address the challenges laid out by the Code.

What to look for in subsequent reports
We will soon roll out our election ads policy in all EU member states in advance of the May 2019 elections and expect to start publishing our transparency report as advertisers begin to run ads as we approach the elections.

---

3. INTEGRITY OF SERVICES

Signatories to the Code of Practice agreed to the following Commitments:
5. Relevant Signatories commit to put in place clear policies regarding identity and the misuse of automated bots on their services and to enforce these policies within the EU. Such measures could include some of the measures in the Annex 2 to this Code.

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

How we are complying with these Commitments by end of 2018:
→ **For both commitments #5 (policies on identity and misuse of automated bots) and #6 (impermissible uses of automated systems):** Our policies apply evenly across all types of actors

and accounts, regardless of whether they are automated or human. They prohibit the following behaviors:

- *Impersonation*: Across our monetized and consumer products (notably including YouTube), we do not allow accounts that attempt to mislead or deceive users by pretending to be someone they are not, or falsely implying an affiliation with a business or individual. We will take action when we find this kind of activity.

- *Misrepresentation in Google News and in our advertising products:* Google News requires that publications be accountable and transparent about their identity, to aid users in determining the credibility of articles and publishers. Publications which falsify, misrepresent, or omit basic information about who they are or where they are based may be removed from News products. Misrepresentation is similarly prohibited for advertisers and publishers using our advertising network, as outlined in section 1 (above).

- *Abusive account creation:* Our systems rely on a host of inputs about historical use and pattern recognition across various services in an effort to detect if an account creation or login is likely to be abusive and to detect and prevent other types of suspicious conduct. As examples, our systems operate to block "bad" account creation or to close groups of such accounts (our consumer Google accounts, upon creation, work across most Google services -- including YouTube). We prevent users from creating a large number of Google Accounts in a short time period. We also require verification, aimed at detecting if a bot is attempting to access or create an account, if we detect suspicious conduct. There is no panacea, however, for this challenge. We continue to work collaboratively with others to identify inauthentic accounts that are used for improper purposes, but our systems are not perfect, and the discovery and deactivation of inauthentic accounts is not always immediate.

- Engagement Abuse: We take any abuse of our systems, on any of our products and services, very seriously. This includes, for instance, attempts to artificially manipulate engagement such as video dislikes or view counts on YouTube, or attempts to game Search ranking algorithms by manufacturing inbound links to one's web property. We have strict policies, including our [Webmaster Guidelines for Search](#) and [Community Guidelines](#) for YouTube, that govern the rules of the road for what content and behavior is allowed on the platform. Artificial manipulation of engagement such as video dislikes and views violates our policies and is prohibited. Furthermore, for well over a decade, YouTube and Search have invested in, built, and deployed proprietary technology to address such efforts to abuse our systems. While no anti-spam system will ever be perfect, we have extensive safeguards in place to detect and neutralize the impact of these types of abuse on our systems.

- Influence operations: We undertake a broad range of approaches combat political influence operations, which we have regularly reported on over the course of the past two years. Over the course of 2018 only, we provided an [update](#) in August about our ongoing efforts to combat political influence campaigns from the Internet Research Agency (IRA) in Russia, as well as our identification of an influence campaign linked to the Iranian government. Our efforts included removing 42 YouTube channels and 1 blog on Blogger associated with the IRA, and removing 39 YouTube accounts and 6 blogs linked to the Islamic Republic of Iran Broadcasting. We provided another update on our efforts to combat political influence operations in November 2018 (same [blogpost](#)), noting that we had found limited additional activity linked to the IRA and Iranian operations. Specifically, we removed 1 YouTube channel

linked to IRA, as well as 34 YouTube channels, 6 G+ pages, and 1 blog linked to Iranian operations.

<u>What to look for in subsequent reports</u>
We will continue to report on the deployment and implementation of policies tackling impersonation, misrepresentation, and abusive account creation.

---

4. EMPOWERING CONSUMERS

<u>Signatories to the Code of Practice agreed to the following Commitments:</u>
7. Relevant Signatories commit to invest in products, technologies and programs such as those referred to in Annex 2 to help people make informed decisions when they encounter online news that may be false, including by supporting efforts to develop and implement effective indicators of trustworthiness in collaboration with the news ecosystem.

8. Relevant Signatories commit to invest in technological means to prioritize relevant, authentic and authoritative information where appropriate in search, feeds, or other automatically ranked distribution channels.

9. Relevant Signatories commit to invest in features and tools that make it easier for people to find diverse perspectives about topics of public interest.

10. Signatories commit to partner with civil society, governments, educational institutions, and other stakeholders to support efforts aimed at improving critical thinking and digital media literacy.

11. Signatories commit to encourage market uptake of tools that help consumers understand why they are seeing particular advertisements.

<u>How we are complying with these Commitments by end of 2018</u>:

→ **Commitment #7 - products, technologies, and programs to help people make informed decisions when they encounter online news that may be false**: Part of our approach to this goal lies in the very nature of Google Search, which by design presents users with a wide array of webpages and information. This gives users access to a plurality of perspectives, helping them get the full picture on news stories of interest to them. In addition, we strive to elevate context and information that help users make judgments about the quality of information they encounter online. We are investing in many efforts to that end, including:
- *Surfacing Fact-Checks on Google*: in 2016, we [launched] the Fact Check label in Google News and Search to make it easier for people to find articles that fact check public information, ranging from claims to public statements to statistics. Initially available in the US only, this feature was expanded globally in April 2017 and is used today by fact-checkers in France, Germany, the UK, Italy, Spain, Poland, Austria, Sweden, Croatia, Netherlands, Denmark, and more. To discover Fact-Checks on Google, users merely have to browse Google Search or News and to look for links that show a "fact-check" tag or snippet. This feature is open to all publishers who signal their use of fact-checking by using a dedicated html mark-up and comply with our [policies on fact-checking]. Over time, we have made it even easier for fact-checkers to help Google find and distribute fact-checking content across Google, as Alphabet's Think/Do tank Jigsaw and the Duke University Reporters' Lab created

a "Share the Facts" widget, which provides an easier alternative for publishers of fact-checks than adding HTML mark-ups.

- *Working with the Trust Project to help surface indicators of credibility*: Google is participating in and providing financial support to the Trust Project, which explores how journalism can signal its trustworthiness in an increasingly chaotic media landscape. The Trust project has developed eight indicators of trust publishers can use to better convey why their content should be seen as credible, with promising results for the publishers who have adopted them on an early trial basis. For example, after the Trinity Mirror in the UK implemented the Trust Project, consumer trust in the newspaper increased by 8 percent. Trust Project participants in Europe include Die Zeit, The Economist, La Repubblica, La Stampa, El Mundo, El Pais, Corriere Della Sera, Deutsche Presse-Agentur (DPA), BBC, FourFourTwo, and more.
- *Working with the Credibility Coalition and the World Wide Web Consortium on markers of information credibility online*: With funding from Google News Initiative, the Credibility Coalition set out more than a year ago to build a project that would develop, test, research, and generate data to support the creation and use of third party credibility signals. CredCo has since built an interdisciplinary community that has worked to define and validate qualitative and quantitative hypotheses about the most meaningful signals of content and source credibility.

→ **Commitment #8 - technological means to prioritize relevant, authentic and authoritative information:** Google Search, News, and YouTube's businesses rely upon continuously meeting our users' expectations that we provide them with the best results for their queries, and as such, we continuously invest in the testing and improving of our ranking algorithms -- rolling out thousands of updates every year (e.g. more than 2,400 updates to Google Search algorithms in 2017 alone).

Every time a user searches for something on Google, our algorithms work to understand the meaning of their query and to rank the pages in our index according to signals including their *relevance* to the user's expectation (does the webpage content cover the subject matter the user is asking about? Are the user's keywords present in this webpage?) and their authoritativeness (Google's estimate of the expertise, authority, and trustworthiness of a web page, based on factors like whether users and other sites on the web seems to trust this web page)[2]. YouTube similarly requires that the Breaking News and Top News shelves only include videos from news sources that satisfy the Google News standard. For subject matters related to news and civic issues, we realize that our users have high expectations when it comes to the trustworthiness of the results we present them. We give more importance to factors related to authoritativeness in response to such queries. This, as well as all other goals and principles of Google Search, is outlined in our Search Quality Raters Guidelines.

This approach, ensures that we prioritize authoritative information. In addition, we have implemented a number of other measures towards the same goal:
- *Additional ranking safeguards during breaking news events*: during breaking news or crisis situations, stemming the tide of misinformation can be challenging. Speculation can outrun facts as legitimate news outlets on the ground are still investigating. At the same time, malicious actors are publishing content on forums and social media with the intent to mislead and capture people's attention as they rush to find early information online. To reduce the visibility of this type of content during crisis or breaking news events, we are improving our systems across Google and YouTube in Europe to put even more emphasis on authoritative results over factors like freshness or relevancy for these breaking news situations.

---

[2] For more information on Search ranking, see: www.google.com/search/howsearchworks

● *Breaking News and Top News Shelves on YouTube*: On YouTube, we have worked to launch several product features that prominently surface news content from only authoritative sources. Two such cornerstone products have been the Top News shelf in YouTube search results and the Breaking News shelf on the YouTube homepage. The Top News shelf triggers in response to certain search queries that have news-seeking intent. Our systems rely on a variety of signals that we derive from Google News and from our own systems when a user might be seeking information on news-related topics. Given the user's intent is to find out more information in this domain where the authority of content and source is particularly important, we display a clearly delineated shelf in the product where we only surface content from authoritative, verified news channels. The Breaking News shelf triggers when there is a significant news event happening in a specific country. This shelf triggers automatically and is shown prominently on the YouTube homepage to all users in the country. Similar to the Top News shelf, only content from authoritative sources is eligible to be surfaced in the Breaking News shelf. In the EU today, these products are launched in the UK, Germany, France, Italy, Poland, Sweden, the Netherlands, Ireland, and Spain, with more countries to be covered over the course of 2019.

Breaking News Shelf on YouTube Homepage

Top News Shelf on YouTube Search
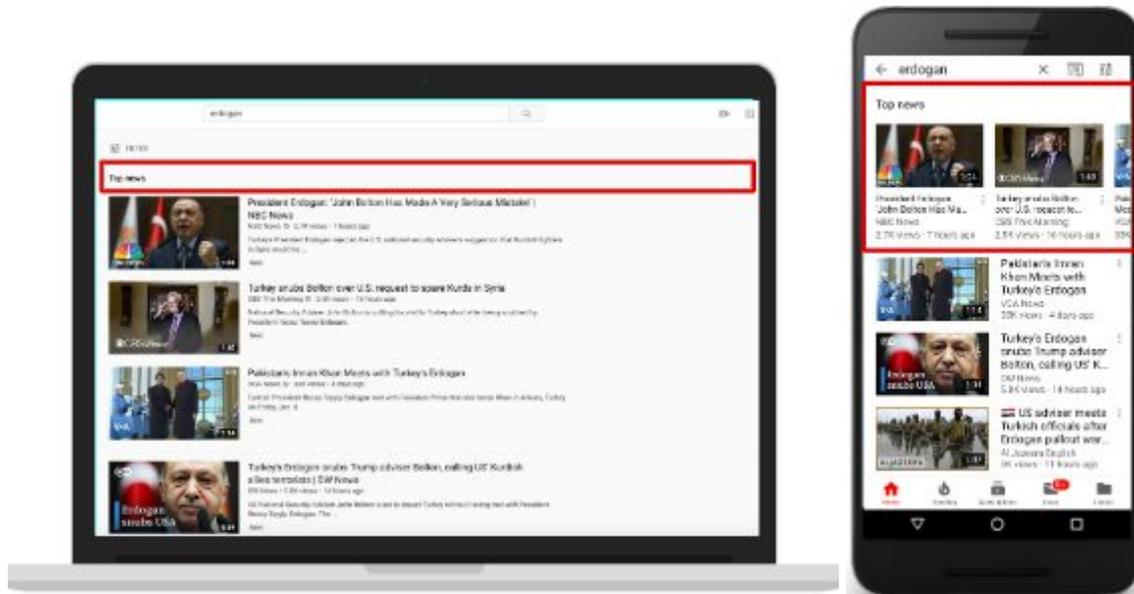
We are also working with candidates, political parties and public authorities to surface relevant and authoritative content during the election period. During 2018 we trained over 800 campaigners on how to manage their online presence using Google Search, including Google's Knowledge Panels, and YouTube. We also trained these campaigners on how to protect their campaigns through our offering of the [Advanced Protection Program](#) for email and [Project Shield](#) for websites. In November, we held trainings specifically for MEPs, candidates and their staffers directly in the European Parliament or in their political groups' annual congresses. We will continue and expand such outreach efforts to political parties, campaigners, NGOs, civil society and others ahead of the 2019 elections.

→ **Commitment #9 - features and tools that make it easier for people to find diverse perspectives about topics of public interest:** as mentioned above, the very nature of our products means that users searching for a topic see a plethora of results from various perspectives, whether they're looking on Google Search, News, or YouTube. That said, we acknowledge the importance of ensuring that users are not locked into a single set of perspectives, and have launched a number of additional efforts to that end:
- *Full Coverage in Google News*: to help users access context and diverse perspectives about the news stories they read, the *"Full Coverage"* feature in Google News lets users explore articles and videos from a variety of publishers related to an article or news story of their choice. The *"Full Coverage"* feature is not personalized, and accessible in one click or tap from most articles in Google News' "*For You*" and "*Headlines*" tabs. It is live in all EU28 Member States except for Spain and Denmark.
- *Breaking News and Top News Shelves on YouTube:* The news shelves outlined above always provide a set of news videos from a wide array of authoritative sources on a particular news topic. Providing multiple videos from a range of channels with varying perspectives provides users with easy access to engage with diverse news content on topics of public interest.
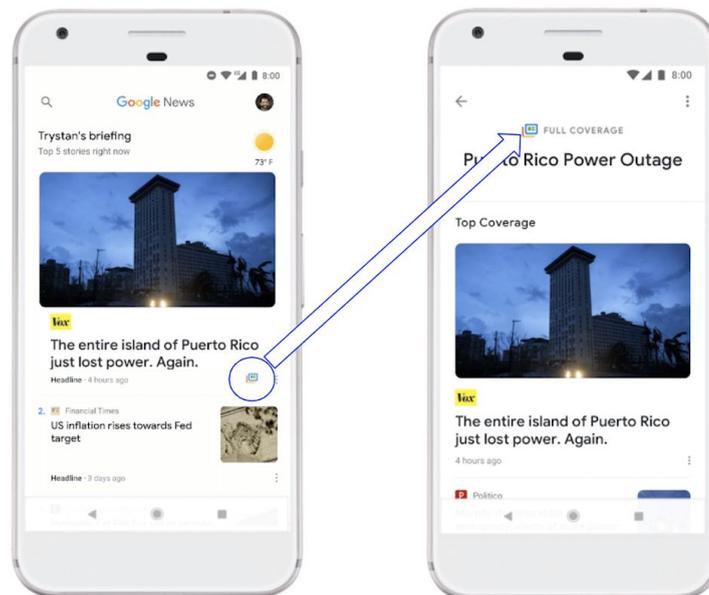
Full Coverage in Google News

→ Commitment #10 - Signatories commit to partner with civil society, governments, educational institutions, and other stakeholders to support efforts aimed at improving critical thinking and digital media literacy.

Media literacy continues to be one of the most important issues of our digital age: a recent [study](#) from the Stanford History Education Group found that 93 percent of college students couldn't flag a lobbyist's website as biased, and 82 percent of middle schoolers couldn't distinguish sponsored content from real news. Google is committed to helping tackle this challenge.

Our 'Be internet awesome[3]' campaign aims at empowering children to be safe and confident digital citizens is critical – but isn't something any one group can solve alone. We worked with the Family Online Safety Institute, ConnectSafely and iKeepSafe to build a solution that aims to engage with parents, educators and kids.g This programme was developed in the US, it has been launched in the UK and will be rolled out to other countries in Europe.

We also hope to work with media platforms and cultural influencers to help change the narrative and elevate the voices that are most credible in speaking out against hate and xenophobia. YouTube's Creators for Change program highlights online stars taking a stand against xenophobia and extremism. The 2018 cohort had 61 creators from over 20 countries (including France, Germany, and the UK) representing a combined audience of over 40 million fans.

We have also recently announced a $10m global Google.org commitment to media literacy projects. We have already supported media literacy projects in the U.K. and in Finland and we are going to support more in the months to come.[4]

---

[3] [beinternetawesome.withgoogle.com](#)
[4] In Finland we have supported The Mannerheim League for Child Welfare and Save the Children; in the U.K. we have supported The Student View to promote media literacy.

→ **Commitment #11- Signatories commit to encourage market uptake of tools that help consumers understand why they are seeing particular advertisements:**
We provide tools to our users to help them better understand why they are seeing an ad in Google's network. On Google Services like Search or YouTube, users can click on Why This Ad. On non-Google websites and apps that partner with Google to show ads, users can click on the Ad Choices icon. Users can also edit their preferences to control the kinds of ads they see, block certain ads, and give us feedback on ads that they think are inappropriate. We are also giving control to our users on their ads settings and preferences via Google Account (1.6 billion unique visits in 2016 and just under 2 billion in 2017).

How to access tools that help users understand why they are seeing an ad:

On Google services, like Google Search, Youtube, or Gmail: Click ⓘ , then "Why This Ad".

On Non-Google websites and apps that partner with Google to show ads: Click AdChoices ▷ .

We also help users understand how advertising works on Google by providing details and examples in our "How We Make Money With Advertising" website, available at https://howwemakemoney.withgoogle.com/.

What to look for in subsequent reports
We will continue to inform the European Commission on the progress of our efforts on fact-checking, trust and credibility signals, media literacy, and elevating authoritative information. We are developing partnerships with European fact-checking organizations ahead of the elections that we will share details on in due course.

---

5. EMPOWERING THE RESEARCH COMMUNITY

Signatories to the Code of Practice agreed to the following Commitments:

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

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

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

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

How we are complying with these Commitments by end of 2018:

Commitment #12 - Supporting good faith independent efforts to track Disinformation and understand its impact. In addition to the work outlined as part of the reporting on commitment #7 to collaborate with fact-checkers, we run a number of ongoing projects to collaborate with fact-checkers and researchers who explore the propagation of Disinformation in the European Union and beyond:

- *Collaboration with the International Fact-Checking Network*: we have partnered with Poynter's International Fact-Checking Network (IFCN), a nonpartisan organization whose Code of Principles for fact check organizations gathers signatories from multiple EU Member States. Our partnership with IFCN focuses on three key areas with a global point of view: helping train more fact-checkers around the world, translating the IFCN's Code of Principles in ten languages, and providing tools and training for free to the fact-checking community. We look forward to exploring how we can collaborate with the European Network of Fact-Checkers.
- *Making it easier for researchers to explore the corpus of fact-checks around the world*: with the DataCommons project, we have begun to work on a "fact-check explorer" which will ultimately let researchers and fact-checkers explore the body of fact-check content that our systems have observed around the world. The project is currently in a test phase, and we hope to have more to announce in 2019.
- *Funding research*: We support the work of researchers who explore the issues of disinformation and trust in journalism, by supporting research at organizations like First Draft, the Oxford University's Reuters Institute, Michigan University's Quello Center, and more. Furthermore, our own research teams work closely with academics around the world on emerging issues related to disinformation, such as the potential threats raised by synthetic media (also known as "deep fakes")
- *Engaging with researchers who explore the propagation and impact of disinformation*: we receive frequent inquiries from academic researchers who investigate these issues as they seek to better understand our products and policies or have questions with regards to data they've collected as part of their independent research efforts. We respond to these requests as frequently as the availability of our teams permit, and aim to inform the work of these researchers to the best of our ability. Where possible, we make it easier for researchers to access the information that they'd otherwise ask of us, by publicizing important information like the functioning of our Search ranking algorithms (see www.google.com/search/howsearchworks) or the Search Quality Rater Guidelines which outline in great details the goals of Search as well as the means by which we ask evaluators to assess the quality of our search results.
- Google News Lab, a team within the Google News Initiative whose mission is to collaborate with journalists and entrepreneurs to drive innovation in news, is going to host verification training workshops in the EU member states - we have recently organized one in Denmark and there are more being scheduled.

Commitment #13. Relevant Signatories commit not to prohibit or discourage good faith research into Disinformation and political advertising on their platforms. Google and YouTube do not prohibit or discourage such good faith research.

Commitment #14. Relevant Signatories commit to encourage research into Disinformation and political advertising. As mentioned as part of our reporting on commitment #12, we currently fund research into disinformation and political advertising, and we will explore additional opportunities to expand our efforts in 2019.

Commitment #15. Relevant Signatories commit to convene an annual event to foster discussions within academia, the fact-checking community and members of the value chain.

Google and YouTube have organized events with all of these communities in 2018, notably including convenings led by the Google News Initiative to hear from these groups which products and policy changes they believe would be most helpful coming from our services. We will continue such efforts in 2019.

What to look for in subsequent reports
We will continue to keep the Commission informed on the progress of ongoing projects to support the research and fact-checking communities, such as the Fact-Check Explorer, as well as of major events and product or policy changes that would benefit the European fact-checking and research communities.