Protecting Election Integrity

Protecting the integrity of elections while making sure people can have a voice is a priority for Facebook. Learning from every election over the last two years, we have increased our capabilities to take down fake accounts, reduce false news, increase ads transparency, disrupt bad actors and support an informed and engaged electorate. We have made investments in both people and technology, and our teams continue to develop smarter tools, greater transparency, and stronger partnerships.

Today, we have more than 30,000 people dedicated specifically to safety and security, with 40 teams contributing to our work on elections. With each election we get better, and we are committed to doing everything we can to prevent bad actors from interfering in the democratic process.

Transparency is a big priority for us to help prevent interference in elections. We believe people should have more information about the ads they see. We're also working to bring more authenticity to advertisers that want to place issue, electoral or political ads. Our goal isn't to penalize organizations but to help people know more about who is trying to influence them.

While these efforts represent an improvement over the past few years, we know that our work is never finished. There have always been people trying to undermine democracy. We are up against determined adversaries who try to attack on many fronts, and we recognize our role and responsibility. We will never stop all the bad actors, but we're making real progress and we are committed to improving. [For more on our work to protect elections around the world in 2019, see here.]

Defining ads about Social Issues, Elections and Politics

Making changes to the way ads with political content work on Facebook and putting in place transparency measures to better hold advertisers accountable is one of the key actions we have taken to better protect elections and integrity on our platform.

Finding the right approach has been a highly complex task given the variety of political systems, electoral regulations and the number of local, regional, and national elections - each with its own complexity. Aware of these complexities, we have undertaken a long and careful process, involving legal experts, researchers, electoral authorities, as well as advertisers around the globe, to help us develop our policies and definitions for political advertisement. We are learning from each election and continue to make improvements to our ads transparency measures.

Today, we have made our ad transparency tools available in 190 + countries and territories and strongly encourage advertisers to get authorised and add the proper disclaimers on their ads - especially in a rapidly evolving regulatory landscape. In over 30 countries, advertisers who run ads about Social Issues, Elections or Politics are required to complete the ads authorization process before they can set up an ad campaign. This includes the US, Canada, India, Argentina, the UK, and countries in the EU. We will continue rolling out enforcement to more countries around the world and remain committed to introducing more enforcement mechanism.

How we define, detect and review political ads vary by country due to different understandings, criteria and definitions of what a political ad is. In general, our definition of political ads include:

- Ads made by, on behalf of or about a current or former candidate for public office, a political party, a political action committee or advocates for the outcome of an election to public office
- Ads about any election, referendum or ballot initiative, including "get out the vote" or election information campaigns
- Ads about social issues, which vary by country where there is enforcement
- Ads regulated by law as political advertising

But in all cases, the advertiser is the one that must comply with any applicable electoral and advertising laws and regulations in the countries they want to run ads in.

Legislation is indispensable for protecting elections but it not yet adapted to the modern threats and online environments. Facebook tries to contribute by introducing significant changes in how the platform deals with political ads: advertisers in many countries must verify their identities before purchasing political ads; we built a searchable archive that shows who pays for ads, what other ads they ran and what audiences saw the ads and users may use a reporting mechanism. In any event, deciding whether an ad is political and that this definition work in different jurisdictions isn't straightforward. Our systems would be more effective if regulation creates common standards for verifying political actors.

Transparency

We believe that transparency leads to increased accountability and responsibility over time – not just for Facebook but advertisers as well. It is why we continue to introduce tools that allow elected officials, those seeking office, and organizations aiming to influence public opinion to show more information about the ads they run and who is behind them. Transparency is a big priority for Facebook to help prevent interference in elections and we believe people should have more information about the ads they see.

We have rolled out our transparency tools globally for advertisers wanting to place ads about social issues, elections or politics. The new ads transparency tools are a critical part of our election integrity efforts and include:

- **Identity Confirmation and Authorization.** To help prevent abuse and foreign interference, the admins of Pages must provide an identification document that can be verified/confirmed. Admins may only be authorised in one country and must provide an ID document for the country which they would like to target for advertising.
- **Labeling Political and Issue Ads.** To increase advertiser transparency, political ads on Facebook and Instagram must be clearly labeled with a "Paid for by" disclaimer. The disclaimer has several components that is verified before it is approved - page or organisation name; website; email address; phone number; and address. This information is included in each ad and publicly available in the Ad Library.
- **Ad Library.** The Ad Library provides advertising transparency by offering a comprehensive, searchable collection of all ads that are active and running on Facebook and Instagram. This includes non-political ads. Political ads, however, are archived and remain in the Ad Libary for seven years. The Ad Library shows ranges of range of impressions, range of spend, range of age and gender reached for each ad. People have the ability to report ads that they believe should have a "Paid for by" disclaimer and don't, or for violating other Facebook advertising policies, directly from the Library. Anyone can explore the Library, with or without a Facebook account. The Ad Library today houses over 4 million ads and continues to be a critical tool to many, especially regulators, journalists, watchdogs and researchers.
- **Ad Library API.** The Ad Library application programming interface (API) allows regulators, journalists, watchdog groups and other people to analyse political ads and help hold advertisers and Facebook accountable. Users can perform customized keyword searches of ads stored in the Ad Library. The results from the API include the creatives and ad performance data that is archived in the Ad Library. Since the EU launch in March, we have made improvements to our API so people can easily access ads from a given country and analyze specific advertisers. We are also working on making it easier to programmatically access ad images, videos and recently served ads.
- **Ad Library Report.** The Ad Library Report is a report of aggregated insights about ads about social issues, elections or politics in the Ad Library. It provides cumulative statistics about ads in the Ad Library. For example, the report shows the total number of ads and total spend on ads in the Ad Library by country, as well as total spend by advertiser, advertiser spend by day and top searched keywords from the past week. Users can explore, filter and download the data into a CSV file. The Ad Library Report is useful to people who are interested in understanding high-level activity in the Ad Archive since launch, while the Ad Library could be used to deep dive into specific ads.

Remarks on report content

PI: How did the companies define "issue-based advertising"?

How we define, detect and review political ads vary by country due to different understandings, criteria and definitions of what a political ad is. In the EU, for example, there aren't laws or agencies that list specific issues that are subject to regulation. But to have a policy that our reviewers can enforce, they need a list explaining what's covered and what's not. So, like in the US, we looked to the non-partisan Comparative Agendas Project (CAP). For decades, CAP has collected information on the policy processes of governments around the world and used that information to develop a list of common terms related to politics and issues. In addition to the CAP list, we leveraged the Eurobarometer survey that the European Parliament published last year in May and engaged with in-country policy stakeholders and trusted third-party advisors throughout the region. We used all of these inputs to develop our policy for issue ads in the EU that take a position with the goal of influencing public debate on topics such as immigration.

PI: ***"In a limited number of countries Facebook requires those running political ads to become authorised. These countries are Brazil, Canada, all EU member states, including the United Kingdom, as well as India, Israel, Ukraine, and the USA. Facebook does not require political advertisers in other countries to become authorised."***

In the markets where we have already introduced our ads transparency tools, authorization and disclaimers are required and enforced upon when setting up the ad. For all other countries, we are providing these tools and strongly encourage advertisers to authorize and provide disclaimers --especially in a rapidly evolving regulatory landscape.

We will continue rolling out enforcement to more countries around the world and remain committed to introducing more enforcement mechanism.

In all cases, the advertiser is the one that must comply with any applicable electoral and advertising laws and regulations in the countries they want to run ads in. In any event, if we are made aware of an ad that is in violation of a law, we will act quickly to remove it.

Advertisers must comply with applicable electoral and advertising laws and regulations. Even if advertisers believe there aren't clear legal requirements that apply to them, we strongly encourage advertisers to become authorized, especially in a rapidly evolving regulatory landscape. Regardless, candidates and those holding public office should use these tools to provide their constituents with more information about who's influencing their vote — and we suggest voters and local regulators hold elected officials and other influential groups accountable, as well.

One useful step to protecting elections is ensuring that advertising is transparent. Right now, all active ads on Facebook are available in the Ad Library. With elections happening all over the world, we are committed to bringing authorization and disclaimer enforcement to more places around the world, starting in countries with imminent elections or regulations.Now that advertisers across the globe can give people more information about the issue, electoral or political ads they see, policymakers and regulators are better positioned to consider how to protect elections with sensible regulations. We can and should not do this task alone. We are open to continue partnering with governments, civil organizations and electoral authorities to contribute to the protection of the integrity of the democratic process.

PI: ***"Facebook provides heightened transparency for political ads in 34 countries (roughly 17% of the countries in the world). This means that for roughly 83% of the countries in the world, the company does not require political advertisers to become authorised, for political ads to carry disclosures, or for ads to be archived."***

In the markets where we have already introduced our ads transparency tools, authorization and disclaimers are required and enforced upon when setting up the ad. For all other countries, we are providing these tools and strongly encourage advertisers to authorize and provide disclaimers --especially in a rapidly evolving regulatory landscape.

We will continue rolling out enforcement to more countries around the world and remain committed to introducing more enforcement mechanism.

PI: ***"Facebook's Ad Library provides an amount spent range for ads in the Library [for example £100-499], instead of exact impression data."***

With any new undertaking we're committed to taking feedback, and learning and improving our tools to make them more useful. For example, we made updates to our Ad Library in March and added more transparency information and search capabilities as well as a reporting feature for users to report ads they think violate our political or general advertising policies.

In may, we updated the Ad Library report [announced in blog] making it available daily, rather than weekly, and added weekly, monthly and quarterly reports that are downloadable for anyone. We also added spending by location across different date ranges and exchange rate conversion for advertisers that have spend across multiple countries that have different currencies to compute total spend or have spend in a currency different from the primary currency they advertised in.

PI: ***"Facebook provides a range of how much was spent on an ad in its Ad Library, as well as general information about the age, gender, and region an ad was shown in. The Ad Library***

also shows the total spent by a Facebook Page on political or political issues ads. Facebook does not break down the amount spent by targeting category, provide information about the intended v. actual audience, or provide granular targeting information via the Ad Library"

We show information and demo breakdown of people who actually saw the ad. We believe the actual impact of an ad with political or issue content offers more transparency than its intent, which is targeting. While we're committed to transparency, we won't do it at the expense of people's privacy. Providing exact targeting categories for ads could expose information about people. We're open to improving our transparency tools for researchers, journalists and people, but it's important we do so in a privacy-safe way.

PI: ***"At present it is not possible to type in a country name and see all the ads that have been or are being targeted at users in a certain country. The results show only a partial view of all the ads. As noted by the Broadcasting Authority of Ireland, this means that "without having access to the complete set of words used by all advertisers, a user cannot retrieve the full set of political adverts from the Facebook Ad Library API".*** **PI:** ***"Facebook's Ad Library provides an amount spent range for ads in the Library [for example £100-499], instead of exact impression data."***

With any new undertaking we're committed to taking feedback, and learning and improving our tools to make them more useful. For example, we heard feedback on the Ad Library Report and plan to add regional spend aggregation, rolling presets and more up-to-date data. [taken from Ads FAQ]

PI: ***"Facebook says that the company requires the name or 'entity' behind Facebook Pages running ads to be available to users. However the Guardian has reported that a powerful political lobbying company was able to run multiple Pages without telling users that they were behind it or who paid for the work. Therefore, at present, users cannot have faith that the person listed on a Facebook Page is the actual person behind the Page" When a political or political issue ad runs without a disclaimer, the ad does not appear to include a sponsor identity, even when the ad is removed and archived by Facebook."***

Our review to-date suggests the majority of the Pages are operated by real people and do not currently violate our coordinated inauthentic behaviour policy. However, we take seriously the information shared by The Guardian and are continuing to review the activity of the Pages mentioned. Ensuring the safety of people on our platform, and that they can trust what they see on Facebook, is paramount.

We recently introduced greater transparency to help people on Facebook understand more about Pages they follow. This includes showing all the ads a Page is running, when it was

created, a history of name changes, and, for Pages with a large audience, the locations of admins. We are exploring additional transparency that would show people more information about who is running the page.

Real people running Facebook Pages with a point of view is not against our terms, and is common across the internet through websites, blogs and other platforms. We are already working with Governments and regulators across a number of areas where we believe regulation is needed, including ensuring that the rules of online political campaigning are clearly defined. These examples could highlight the case for new rules since internet platforms are not currently required to find out who is paying page admins or managers.

PI: ***"In countries where authorisation is required, Facebook requires political ads to be bought in local currency and advertisers must have a payment source for ads with an address in country in which political ads are running. Advertisers can run political ads only in the countries in which they are authorised. If an advertiser wants to run political ads in multiple countries – they must complete the authorisation process in each respective country. However, two weeks prior to the 2019 EU Parliamentary elections, the Dutch digital rights organisation Bits of Freedom was able to run ads in Germany using a Dutch Facebook and bank account, apparently skirting the policy."***

Deciding whether an ad is political is not straightforward. This is why we are working with governments, watchdogs, regulators and others, and calling for sensible regulation. Our systems would be more effective if regulation created common standards for verifying political actors. Feedback from organisations like Bits of Freedom gives us the opportunity to learn and improve. And we encourage our user community to report ads that they find violate our political advertising policies - this is why we recently added a user reporting feature in the Ad Library.

No one and, we are not an exception, will be able to prevent abuse completely. Our societies are dealing with smart, creative and well-funded opponents who change their tactics as soon as abuse is discovered. But this will not change our determination to continue working on measures that we reasonably think will contribute to prevent interference with elections in the future.

We invest a lot in people and better technology to proactively identify abuse. We cannot and should not do this alone. We are working with governments, watchdogs, regulators and others; calling for sensible regulation; and bringing transparency and authenticity to ads so people can keep us and advertisers accountable.

PI: Researchers have identified 96,106 political ads without "Paid for by" disclosure -- Do we know more about this?

If we determine that an ad is related to social issues, politics or elections after it has started running without the disclaimer label, we would disable the ad. The ad would be archived in the Ad Library with a label stating that "This ad ran without a disclaimer" and would remain searchable in the Ad Library for 7 years.

As noted in the NYU report, dated 12 February 2019, Facebook is the only company that has a mechanism to retroactively mark an ad as political if discovered. As part of these efforts, we also updated reporting capabilities in the Ad Library for users to report ads they believe are violating our Ads Policies. This is one of the many ways we are helping to hold advertisers accountable.

PI: How users can report political advertising for the company to act upon?

Users can report ads that do not have a disclaimer, that they believe should, directly in the Ad Library. Users can click on the three dots on the top right corner of the ad and select Report Ad.

PI: What policies did the companies put in place regarding what constitutes impermissible use of automated systems and to make this policy publicly available on the platform and accessible to users?

All of our policies can be found here: https://www.facebook.com/policies

Authenticity is the cornerstone of our community. We believe that people are more accountable for their statements and actions when they use their authentic identities. That's why we require people to connect on Facebook using the name they go by in everyday life. Our authenticity policies are intended to create a safe environment where people can trust and hold one another accountable.

Preventing fake accounts is one way to stop abuse on our platforms. We focus our enforcement against abusive accounts to both prevent harm and avoid mistakenly taking action on good accounts.

In the six months between Q4 2018 and Q1 2019, we've seen a steep increase in the creation of abusive, fake accounts on Facebook. The number of fake accounts disabled spiked up from 1.2 billion accounts in Q4 2018 to 2.19 billion in Q1 2019, largely due to increased automated attacks by bad actors who attempt to create large volumes of accounts at one time. The majority of these accounts were caught within minutes of registration, before they became a part of our monthly active user population. For the accounts that are not initially detected and disabled, we find that many of them are used in spam campaigns and are financially

motivated.

We try to stop fake accounts abusing our platforms in three distinct ways. Of the accounts we remove, both at sign-up and those already on the platform, 99% of these are proactively detected by us before people report them to us.

- **Blocking accounts from being created:** The best way to fight fake accounts is to stop them from getting onto Facebook in the first place. We've built detection technology that can detect and block accounts even before they are created. Our systems look for a number of different signals that indicate if accounts are created in mass from one location. A simple example is blocking certain IP addresses altogether so that they can't access our systems and thus can't create accounts. The data we include about fake accounts does not include unsuccessful attempts to create fake accounts that we blocked at this stage. This is because we can't know the number of attempts to create an account we've blocked as, for example, we block whole IP ranges from even reaching our site. While these efforts aren't included in the report, we can estimate that every day we prevent millions of fake accounts from ever being created using these detection systems.

- **Removing accounts when they sign-up:** Our advanced detection systems look for potential fake accounts as soon as they sign-up, by spotting signs of malicious behavior. These systems use a combination of signals such as patterns of using suspicious email addresses, suspicious actions, or other signals previously associated with other fake accounts we've removed. Most of the accounts we currently remove, are blocked within minutes of their creation before they can do any harm. We include the accounts we disable at this stage in our accounts actioned metric for fake accounts. Changes in our accounts actioned numbers are often the result of unsophisticated attacks like we saw in the last two quarters. These are really easy to spot and can totally dominate our numbers, even though they pose little risk to users. For example, a spammer may try to create 1,000,000 accounts quickly from the same IP address. Our systems will spot this and remove these fake accounts quickly. The number will be added to our reported number of accounts taken down, but the accounts were removed so soon that they were never considered active and thus could not contribute to our estimated prevalence of fake accounts amongst monthly active users or our publicly stated monthly active user number or even any ad impressions.

- **Removing accounts already on Facebook:** Some accounts may get past the above two defenses and still make it onto the platform. Often, this is because they don't readily show signals of being fake or malicious at first, so we give them the benefit of the doubt until they exhibit signs of malicious activity. We find these accounts when our detection systems identify such behavior or if people using Facebook report them to us. We use a number of signals about how the account was created and is being used

to determine whether it has a high probability of being fake and disable those that are. The accounts we remove at this stage are also counted in our accounts actioned metric. If these accounts are active on the platform, we would also account for them in our prevalence metric. Prevalence of fake accounts measures how many active fake accounts exist amongst our monthly active users within a given time period.