

Helping to Protect the US 2020 Election

Table of Contents

01	Introduction.....	03
02	Key Investments and Improvements	04
03	Fighting Influence Operations.....	07
	Coordinated Inauthentic Behavior	07
	Inauthentic Behavior Policy	08
	Protecting Accounts Through Facebook Protect.....	09
04	Increasing Transparency.....	10
	Making Ads More Transparent	10
	Bringing Greater Transparency to Political Ads Around the World.....	10
	Making It Easier to Understand Political Ads	11
	Expanded Transparency and More Control for Political Ads.....	12
	Our Political Ads Policy.....	13
	Making Pages More Transparent.....	14
	Labeling State-Controlled Media.....	14
	Increasing Transparency With CrowdTangle.....	15
05	Reducing Misinformation.....	17
	Remove: Enforcing Against Manipulated Media and Voter Suppression	17
	Voter Suppression and Intimidation	17
	Misleading Manipulated Videos	18
	Reduce: Preventing the Spread of Viral Misinformation	20
	Fact-Checking Program	21
	Inform: Helping People Better Understand What They See Online	22
06	Supporting Political Speech	24
	Newsworthiness Exemption.....	25
07	Global Partnerships and Coordination	26
08	Conclusion	28

01. Introduction

This overview provides a look at Facebook's comprehensive efforts over the past three years to help protect the democratic process ahead of the 2020 US elections. We know that elections have changed, and so has Facebook. We've worked to develop a comprehensive strategy to close previous vulnerabilities while addressing new and emerging threats. And we've developed smarter tools, greater transparency, and stronger partnerships to help us do just that. We have more than 35,000 people dedicated to safety and security issues, with about 40 teams contributing to our work on elections. We block millions of fake accounts each day so they can't spread misinformation. We continue to improve our coordination and cooperation with law enforcement, including DNI, DHS, FBI, as well as other federal officials, state election officials, and technology companies, to support better information sharing and threat detection in service of public safety. And we've set a new standard for transparency in Pages and political ads so people can see who is behind them.

In 2016, we were on the lookout for traditional cyber threats like hacking and stealing information. What happened was a much different kind of attack, meant to sow social discord around divisive political issues. We've learned lessons from 2016 and while we have seen threats evolve, we're working hard to stay ahead of those efforts so it's harder to use our platform for election interference.

We know that security is never finished and we can't do this alone—that's why we continue to work with policymakers and experts to make sure we are constantly improving.

02

II. Key Investments and Improvements

2016

DECEMBER 2016 Launched third-party fact-checking program

2017

APRIL 2017 Announced new tool to help people spot false news on Facebook

AUGUST 2017 Banned Pages that repeatedly share false news from running ads

SEPTEMBER 2017 Removed first accounts for violating our Coordinated Inauthentic Behavior policy

OCTOBER 2017 Announced new authorization requirements before advertisers can run ads about elections or politics on Facebook or Instagram

DECEMBER 2017 Launched 'Related Articles' to give people more information on content rated false by third-party fact-checkers

2018

APRIL 2018 Announced Election Research Commission to study social media's role in elections

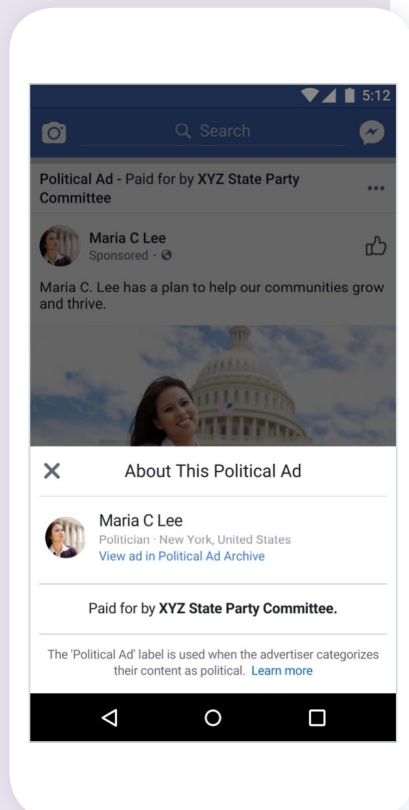
APRIL 2018 Updated our Community Standards to include internal guidelines used for decisions about content

APRIL 2018 Introduced appeals of content that we may have mistakenly removed

MAY 2018 Began partnership with the Atlantic Council's Digital Forensic Research Lab to study and investigate disinformation campaigns around the world, and continued to expand these external partnerships throughout 2018 and beyond

MAY 2018

Ads authorization policy expanded to cover issue and election-related ads



MAY 2018 Launched the Ad Library (with social issue/political ads archived for 7 years) and began requiring that advertisers running social issue/political ads confirm their location in the US and place “Paid for by” disclaimers on these ads

JUNE 2018 Allow people to see the ads a Page is running across Facebook, Instagram, and Messenger even if those ads aren't shown to you.

JUNE 2018 Announced new Pages transparency feature showing Page creation and name change dates.

AUGUST 2018 Launched new Pages authorizations requirements beginning with high reach Pages in the US

SEPTEMBER 2018 Started fact-checking photos and videos in addition to other content

SEPTEMBER 2018 First physical Elections Operations Center for monitoring ahead of key elections

SEPTEMBER 2018 Added “People Who Manage This Page” section to show country locations in Page Transparency

OCTOBER 2018 Expanded voter suppression policies and announced new way for users to report potential voter suppression

OCTOBER 2018 Rolled out the Ad Library report to make it easier to see who is spending money on social issue/political ads on Facebook

2019

MAY 2019 Launched fact-checking of Instagram content

JUNE 2019 Rolled out our ads authorization process for more advertisers globally

JUNE 2019 Launched policy banning paid advertising that suggests voting is useless or meaningless, or advises people not to vote

AUGUST 2019 Announced that organizations running social issue/political ads in the US now have to go through a stricter authorizations process, including providing information such as a Tax ID or FEC Committee ID number

SEPTEMBER 2019 Introduced the Deep Fake Detection Challenge to develop new ways of detecting and preventing manipulated media

OCTOBER 2019 Launched Facebook Protect to give campaigns, elected officials, their staffs, and others increased security protections

OCTOBER 2019 Announced stronger labeling of content rated false by third-party fact-checkers

NOVEMBER 2019 Started showing the “Confirmed Page Owner” of US pages with large followings

DECEMBER 2019 Reached over 50 international fact-checking partners covering over 40 languages around the world, and the program continues to expand

2020

JANUARY 2020 Announced policy to take down misleading manipulated media

JANUARY 2020 Announced new transparency features and controls for social issue/political ads

FEBRUARY 2020 Began requiring Pages to designate—and make transparent—a “Confirmed Page Owner” to continue running social issue/political ads

03

III. Fighting Influence Operations

COORDINATED INAUTHENTIC BEHAVIOR

Over the last three years, our teams of engineers and specialized threat investigators have worked to identify new and emerging threats and remove coordinated inauthentic behavior across our apps. In 2019 alone, we took down over 50 networks worldwide, many ahead of major elections.

We took down these networks based on their behavior, not the content they post. In each case, the people behind this activity coordinated with one another and used fake accounts to misrepresent themselves, and that was the basis for our action. We continue to share our findings with law enforcement and industry partners. You can find more information about each of the Coordinated Inauthentic Behavior (CIB) takedowns at newsroom.fb.com.

We continue to identify and take action against networks working together to mislead people about who they are and what they're doing. When we find domestic, non-government campaigns in which the use of fake accounts is central to the operation, we will remove all accounts, Pages and groups directly involved in this activity. Additionally, if the activity we remove is directly related to a civic event, poses imminent harm or involves a new technique or a new significant bad actor, we will share our findings at the time of enforcement.

Persona Non Grata: If, in the course of a CIB investigation, we determine that a particular organization is primarily organized to conduct manipulation campaigns, we will remove it from our platforms in its entirety and work to ensure that the organization can never again have a presence across Facebook products.

There are two types of CIB that are particularly egregious:

- Foreign-led efforts to manipulate public debate in another country
- Operations run by a government to target its own citizens. These can be particularly concerning when they combine deceptive techniques with the real-world power of a state.

If we see any instances of CIB conducted on behalf of a government entity or by a foreign actor, we will apply the broadest enforcement measures including the removal of every on-platform property that is part of the operation itself and the people and organizations behind it. We will also announce the removal of this activity at the time of enforcement.

Identifying and disrupting influence operations is a whole-of-society challenge, and requires multi-stakeholder solutions to be effective. For each of our CIB takedowns, we work with third party independent researchers, like the Atlantic Council's DFR Lab, Graphika, or Stanford's Internet Observatory, who conduct deep cross-platform investigations to understand the broad nature of these operations and dissect their core narrative themes. We also partner closely with our industry partners, academic researchers, investigative journalists, and law enforcement partners to ensure that we are bringing every tool available to bear against these threats.

INAUTHENTIC BEHAVIOR POLICY

As we've improved our ability to disrupt these foreign influence operations, we've also built a deeper understanding of different threats and how best to counter them. We investigate and enforce against many types of inauthentic behavior. However, the most appropriate way to respond to someone boosting the popularity of their posts in their own country may not be the best way to counter foreign interference. While significant public attention has been on foreign governments engaging in these types of violations, over the past two years, we have increasingly seen non-state actors, domestic groups and commercial companies engaging in this behavior. And we've seen financially-motivated campaigns relying on fake accounts and other inauthentic tactics to drive clicks and mislead people. That's why we've updated our inauthentic behavior policy to clarify how we deal with the range of deceptive practices we see on our platforms, whether foreign or domestic, state or non-state.

In line with our commitment to authenticity, we don't allow people to misrepresent themselves on Facebook, use fake accounts, artificially boost the popularity of content, or engage in behaviors designed to enable other violations under our Community Standards. We routinely take action against other inauthentic behaviors, including financially motivated activity like spam or fake engagement tactics that rely on inauthentic amplification or evading enforcement, rather than a core use of fake accounts. We enforce against Inauthentic Behavior (IB) based on specific protocols that are reviewed and approved by

multiple product and policy teams, to make sure we are catching the right things and not impacting authentic behavior. The consequences of engaging in inauthentic behavior include temporary restrictions, warnings, down-ranking or removal. You can find our policy [here](#).

PROTECTING ACCOUNTS THROUGH FACEBOOK PROTECT

In October, we [launched](#) Facebook Protect to further secure the accounts of elected officials, candidates, their staff and others who may be particularly vulnerable to targeting by hackers and foreign adversaries. As we've seen in past elections, they can be targets of malicious activity. However, because campaigns are generally run for a short period of time, we don't always know who these campaign-affiliated people are, making it harder to help protect them.

Page admins can enroll their organization's Facebook and Instagram accounts in Facebook Protect and invite members of their organization to participate in the program as well. Participants will be required to turn on two-factor authentication, and their accounts will be monitored for hacking, such as login attempts from unusual locations or unverified devices. And, if we discover an attack against one account, we can review and protect other accounts affiliated with that same organization that are enrolled in our program. Read more about Facebook Protect and enroll [here](#).

04

IV. Increasing Transparency

MAKING ADS MORE TRANSPARENT

In 2016, Russia tried to use Facebook ads to sow social discord and interfere in the US election. In addition to going after coordinated influence operations, we've introduced an unprecedented level of transparency and authenticity around political and issue ads in the US so people can see who's trying to influence their vote:

- In 2018, we started requiring advertisers to get authorized in the US before running ads about social issues, elections or politics. Advertisers that want to run these ads must prove who they are and where they are located, by providing an ID issued in the US and verifying they have a residential address in the US
- Ads about social issues, elections or politics must include "Paid for by" disclaimers which show the organization or person behind the ad and advertisers must provide more information for Facebook to confirm this, like their Federal Elections Commission or Tax-ID number. Read more [here](#).
- We put political and social issue ads in a publicly searchable Ad Library for seven years, and provide people with information such as a range of how much was spent on the ad and how many people saw the ad.

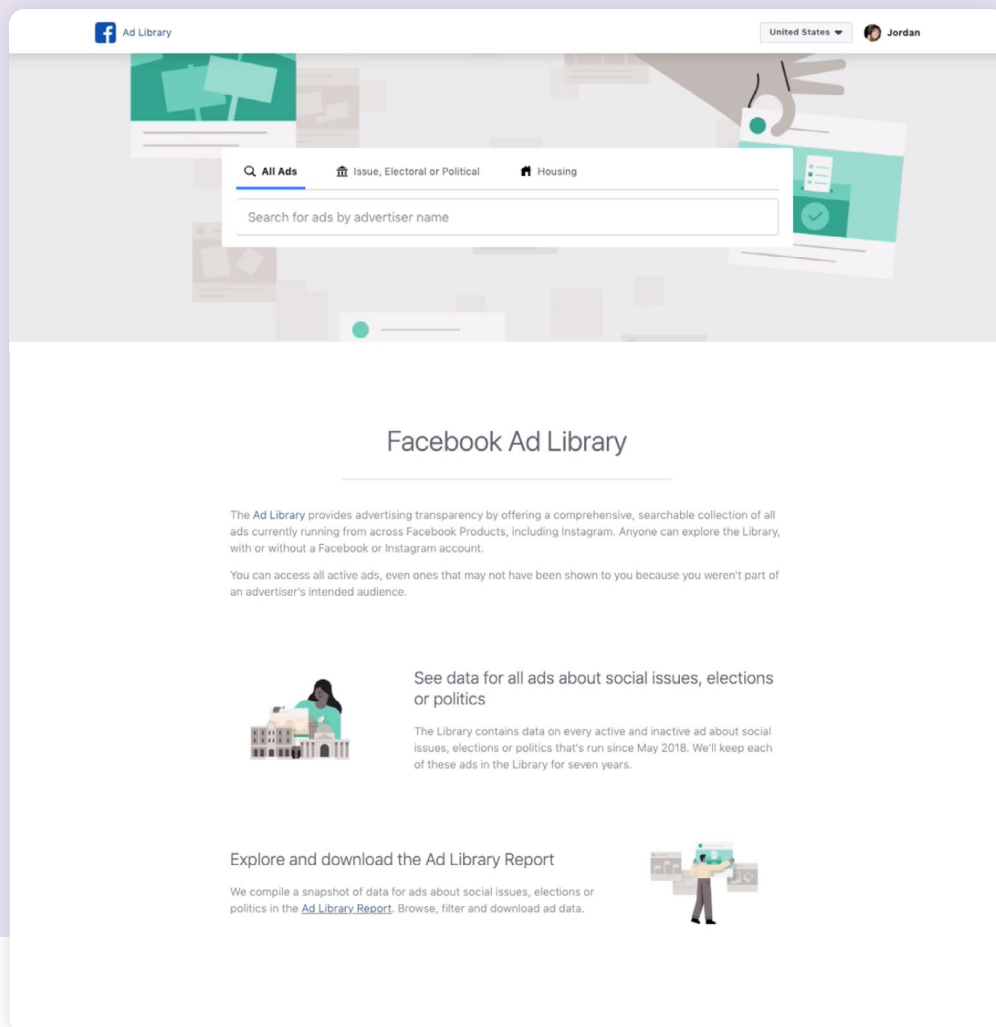
BRINGING GREATER TRANSPARENCY TO POLITICAL ADS AROUND THE WORLD

Elections are happening all over the world and we are committed to requiring authorizations and "Paid for by" disclaimers in more places, to help prevent interference in elections globally. As of January 2020, we proactively enforce on these requirements and review issues, electoral or political ads in 36 countries, which include the US, in addition to the EU, UK, India, Canada, Israel, Ukraine, Singapore, Argentina, Taiwan and more.

We also believe that transparency leads to increased accountability and responsibility over time—not just for Facebook but advertisers as well. It's 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

AD LIBRARY

A transparency hub that provides a searchable view of ads across our products.



who's behind them. In June 2019, we rolled out these tools globally for advertisers wanting to place ads about social issues, elections or politics. For a full list of all 196 countries and territories where these tools are now available, visit our [Help Center](#).

MAKING IT EASIER TO UNDERSTAND POLITICAL ADS

The [Ad Library](#) is Facebook's comprehensive ads transparency hub, providing a searchable view of ads across Facebook products. It gives people more information about the ads they see. For ads about social issues, elections, or politics, we keep these ads in the library for seven years. In October 2019, we updated the Ad Library, Ad Library Report and Ad Library API to help journalists, lawmakers, researchers and others learn more about the ads they see. These updates include:

- A US presidential candidate spend tracker, so that people can see how much candidates have spent on ads

- Additional spend details at the state or regional level to help people analyze advertiser and candidate efforts to reach voters geographically
- Making it clear if an ad ran on Facebook, Instagram, Messenger or Audience Network
- Adding useful API filters, providing programmatic access to download ad creatives and a repository of frequently used API scripts

EXPANDED TRANSPARENCY AND MORE CONTROL FOR POLITICAL ADS

In January 2020, we [announced](#) updates to our Ad Library to increase the level of transparency it provides for people and give them more control over the ads they see.

Throughout the first half of 2020, we'll add new features so voters, the media and researchers can better understand the ads candidates and political organizations are running, and people can have more control over the ads they see:

View audience size in the Ad Library

We are adding ranges for [Potential Reach](#), which is the estimated target audience size for each political, electoral or social issue ad so you can see how many people an advertiser wanted to reach with every ad.

Better Ad Library search and filtering

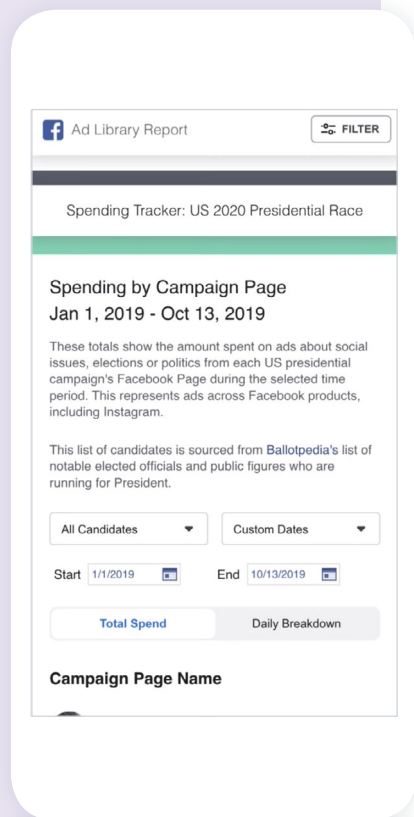
We are adding the ability to search for ads with exact phrases, better grouping of similar ads, and adding several new filters to better analyze results—e.g. audience size, dates and regions reached. This will allow for more efficient and effective research for voters, academics or journalists using these features.

Control over [Custom Audiences from a list](#)

We are rolling out a control to let people choose how an advertiser can reach them with a Custom Audience from a list. These Custom Audiences are built when an advertiser uploads a hashed list of people's information, such as emails or phone numbers, to help target ads. This control will be available to all people on Facebook and will apply to all advertisers, not just those running ads about social issues, elections, or politics. People have always been able to hide all ads from a specific advertiser in their [Ad Preferences](#) or directly in an ad. But now they will be able to stop seeing ads based on an

ADS LIBRARY

The US presidential candidate spend tracker displays how much candidates have spent on ads.



> 85%

of spend by US presidential candidates on Facebook is for ad campaigns targeted to audiences estimated to be greater than 250,000

advertiser's Custom Audience from a list—or make themselves eligible to see ads if an advertiser used a list to exclude them. For example, if a candidate has chosen to exclude you from seeing certain fundraising ads because they don't think you will donate again, but you still want a chance to see those ads, you can stop yourself from being excluded.

See fewer political ads

Seeing fewer ads about social issues, elections, or politics is a common request we hear from people. That's why we plan to add a new control that will allow people to see fewer ads about social issues, elections, or politics on Facebook and Instagram. This feature builds on other controls in Ad Preferences we've released in the past, like allowing people to see fewer ads about certain topics or remove interests.

OUR POLITICAL ADS POLICY

There has been much debate about political advertising online and the different approaches that companies have chosen to take. While Twitter has chosen to block political ads and Google has chosen to limit the targeting of political ads, we are choosing to expand transparency and give more controls to people when it comes to political ads.

Unlike Google, we have chosen not to limit targeting of these ads. We considered doing so, but through extensive outreach and consultations we heard about the importance of these tools for reaching key audiences from a wide range of NGOs, non-profits, local candidates, elected officials, political groups and campaigns, including both Republican and Democratic committees in the US. And when it comes to targeting, our data actually indicates over 85% of spend by US presidential candidates on Facebook is for ad campaigns targeted to audiences estimated to be greater than 250,000.

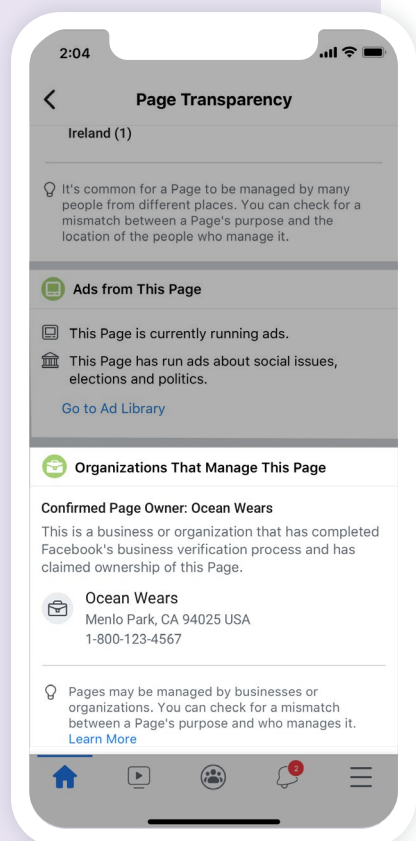
Ultimately, we don't think decisions about political ads should be made by private companies, which is why we are arguing for regulation that would apply across the industry. The Honest Ads Act is a good example—legislation that we endorse and many parts of which we've already implemented—and we are engaging with policy makers in the European Union and elsewhere to press the case for regulation too. Frankly, we believe the sooner Facebook and other companies are subject to democratically accountable rules on this the better.

In the absence of regulation, Facebook and other companies are left to design their own policies. We have based ours on the principle that people should be able to hear from those who wish to lead them, and that what they say should be scrutinized and debated in public. This does not mean that politicians can say whatever they like in advertisements on Facebook. All advertisers, including politicians, political parties, PACs and advocacy groups, must abide by our Community Standards, which apply to ads and include policies that, for example, ban hate speech, harmful content and voter interference and suppression. We regularly disallow ads from politicians that break our rules.

We recognize this is an issue that has provoked much public discussion—including much criticism of Facebook’s position. We will continue to work with regulators and policy makers in our ongoing efforts to help protect elections.

PAGE TRANSPARENCY

We’ve added more transparency to help people understand more about Pages they follow, including the primary countries they are managed from and confirmed Page owners.



MAKING PAGES MORE TRANSPARENT

We want to make sure people are using Facebook authentically, and that they understand who is speaking to them. We’ve taken steps to ensure Pages are authentic and more transparent by showing people the Page’s primary country location and whether the Page has merged with other Pages. This gives people more context on the Page and makes it easier to understand who’s behind it.

Increasingly, we’ve seen people failing to disclose the organization behind their Page as a way to make people think that a Page is run independently. To address this, we’ve added more information about who is behind a Page, including a new “Organizations That Manage This Page” tab that will feature the Page’s “Confirmed Page Owner,” including the organization’s legal name and verified city, phone number or website.

LABELING STATE-CONTROLLED MEDIA

We want to help people better understand the sources of news content they see on Facebook so they can make informed decisions about what they’re reading. Soon we’ll begin labeling media outlets that are state-controlled media. This label will be on both their Page and in our Ad Library, and we plan to apply these labels on Instagram as well. We will hold these Pages to a higher standard of transparency because they combine the opinion-making influence of a media organization with the strategic backing of a state.

We are developing our own definition and standards for state-controlled media with input from more than 60 experts around the world specializing in media, governance, human rights and development. Those consulted represent leading academic institutions, nonprofits and international organizations in this field, including Reporters Without Borders, Center for International Media Assistance, European Journalism Center, Oxford Internet Institute, Center for Media, Data and Society (CMDs) at the Central European University, the Council of Europe, UNESCO and others.

It's important to note that our policy draws an intentional distinction between state-controlled media and public media, which we define as any entity that is publicly financed, retains a public service mission and can demonstrate its independent editorial control. At this time, we're focusing our labeling efforts only on state-controlled media.

We will update the list of state-controlled media on a rolling basis. For any organization that believes we have applied the label in error, there will be an appeals process.

INCREASING TRANSPARENCY WITH CROWDTANGLE

We've provided Secretaries of State and election boards in all 50 states and Puerto Rico access to [CrowdTangle](#), a free social media monitoring tool owned by Facebook to help them quickly identify misinformation, voter interference and suppression. CrowdTangle has worked with local officials to customize public Live Displays for each state.

As a result, state officials can follow the conversations related to voting and the upcoming elections from influential accounts across all 50 states. Each live display surfaces content from elected officials, government agencies, colleges and universities, as well as local media. Alongside this work, CrowdTangle has partnered with the majority of news outlets at the state level to provide access and training for journalists for the upcoming 2020 elections and beyond to better understand and analyze voter interference and misinformation at the state and national level.

We've already seen this work pay off. In 2019, because of their access to CrowdTangle, Louisiana state officials identified election misinformation (for example, inaccurate poll hours) that local elected officials had inadvertently posted. They were then able to reach out directly to the officials to have them correct the information on their posts.

This work isn't just limited to state elections officials. CrowdTangle has also created several [public live displays](#) that anyone can use to monitor what US presidential candidates are posting on Facebook and Instagram in one dashboard. In addition to the 2020 presidential candidates, CrowdTangle has also created live displays that allow anyone to track what senators, representatives in the House and governors are posting on both their official and campaign Pages. This includes live displays for 2020 US Presidential Candidates, US Senate, US House GOP, and more.

Our elections operations center teams also use CrowdTangle as one potential signal of content that may violate our policies, and the tool has been particularly useful for helping us proactively identify voter interference and suppression.

05

V. Reducing the Spread of Misinformation

We've cut the amount of engagement with fake news on Facebook by more than half according to [multiple independent studies](#) so we know people are seeing much less of it. But misinformation is a complex and evolving problem, and we have much more work to do.

We fight the spread of false news by following a three-part framework:

Remove

We remove content that violates our [Community Standards](#), which we use to protect authenticity, safety, privacy, and dignity on the platform.

Reduce

We partner with more than 50 partners around the world, fact-checking content in more than 40 languages. When content is rated false, we dramatically reduce its distribution so less people see it.

Inform

We also believe it's important to inform users when they encounter misinformation, so they can decide for themselves what to read, trust, and share. If content is rated false by a fact-checker, people who see it, try to share it, or already have, will see warnings alerting them that it's false.

REMOVE: ENFORCING AGAINST MANIPULATED MEDIA AND VOTER SUPPRESSION

There are some extreme forms of misinformation that we may remove. This includes claims or content that contribute to manipulative interference in the voting process through misrepresentations about the participation in democratic processes, and misleading manipulated videos.

VOTER SUPPRESSION AND INTIMIDATION

Attempts to interfere with or suppress voting undermine our core values as a company, and we work proactively to remove this type of harmful content. Ahead of the 2018 midterm elections, we expanded our voter suppression and interference policies to prohibit:

> 45K

pieces of content that violated our voter suppression and intimidation policies removed by our Elections Operations Center ahead of the 2018 midterm elections

> 90%

of that violating content was detected by our systems before anyone reported the content to us

- Misrepresentation of the dates, locations, times and methods for voting or voter registration (e.g. "Vote by text!")
- Misrepresentation of who can vote, qualifications for voting, whether a vote will be counted and what information and/or materials must be provided in order to vote (e.g. "If you voted in the primary, your vote in the general election won't count.")
- Threats of violence relating to voting, voter registration or the outcome of an election.

We remove this type of content regardless of who it's coming from, and our systems are now more effective at proactively detecting and removing it. Ahead of the midterm elections, our Elections Operations Center removed more than 45,000 pieces of content that violated these policies—more than 90% of which our systems detected before anyone reported the content to us.

We also recognize that there are certain types of content, such as hate speech, that may impact people's likelihood to vote. That's why our [hate speech policies](#) ban efforts to exclude people from political participation on the basis of things like race, ethnicity or religion (e.g., telling people not to vote for a candidate because of the candidate's race, or indicating that people of a certain religion should not be allowed to hold office).

In advance of the US 2020 elections, we're implementing additional policies and expanding our technical capabilities on Facebook and Instagram to protect the integrity of the election. Following up on a commitment we made in the [civil rights audit report](#) released in June 2019, we have now implemented our policy banning paid advertising that suggests voting is useless or meaningless, or advises people not to vote.

We are also continuing to expand and develop our partnerships to provide expertise on trends in voter suppression and intimidation, as well as early detection of violating content. This includes working directly with secretaries of state and election directors to address localized voter suppression that may only be occurring in a single state or district. This work is being supported by our Elections Operations Center during both the primary and general elections.

MISLEADING MANIPULATED VIDEOS

People share photos and videos on Facebook every day, creating compelling and creative visuals on our platform. Some are

manipulated, often for benign reasons, like making a video sharper or audio more clear. But there are people who engage in media manipulation in order to mislead, including through “deep learning” techniques to create videos that distort reality—usually called “deepfakes.” These videos present a significant challenge for our industry and society as their use increases.

In January 2020, we [announced](#) how we are addressing both deepfakes and all types of manipulated media. Across the world, we’ve been driving conversations with more than 50 global experts with technical, policy, media, legal, civic and academic backgrounds to inform our policy development and improve the science of detecting manipulated media.

As a result of these, we have strengthened our policy toward misleading manipulated videos that have been identified as deepfakes. We will remove misleading manipulated media if it meets the following criteria:

- It has been edited or synthesized—beyond adjustments for clarity or quality—in ways that aren’t apparent to an average person and would likely mislead someone into thinking that a subject of the video said words that they did not actually say. And:
- It is the product of artificial intelligence or machine learning that merges, replaces or superimposes content onto a video, making it appear to be authentic.

This policy does not extend to content that is parody or satire, or video that has been edited solely to omit or change the order of words.

Consistent with our existing policies, audio, photos or videos, whether a deepfake or not, will be removed from Facebook if they violate any of our other Community Standards including those governing nudity, graphic violence, voter suppression and hate speech.

Videos that don’t meet these standards for removal are still eligible for review by one of our independent third-party fact-checkers, which include over 50 partners worldwide fact-checking in over 40 languages. If a photo or video is rated false or partly false by a fact-checker, we significantly reduce its distribution in News Feed and reject it if it’s being run as an ad. And critically, people who see it, try to share it, or have already shared it, will see warnings alerting them that it’s false.

This approach is critical to our strategy and one we heard specifically from our conversations with experts. If we simply removed all

manipulated videos flagged by fact-checkers as false, the videos would still be available elsewhere on the internet or social media ecosystem. By leaving them up and labelling them as false, we're providing people with important information and context.

Our enforcement strategy against misleading manipulated media also benefits from our efforts to root out the people behind these efforts. In December 2019, we identified and removed a network using [AI-generated photos](#) to conceal their fake accounts. Our teams continue to proactively hunt for fake accounts and other coordinated inauthentic behavior.

We are also engaged in the identification of manipulated content, of which deepfakes are the most challenging to detect. That's why we launched the [Deep Fake Detection Challenge](#), which has spurred people from all over the world to produce more research and open source tools to detect deepfakes. This project, supported by \$10 million in grants, includes a cross-sector coalition of organizations including the Partnership on AI, Cornell Tech, the University of California Berkeley, MIT, WITNESS, Microsoft, the BBC and AWS, among several others in civil society and the technology, media and academic communities.

In a separate effort, we've [partnered with Reuters](#), the world's largest multimedia news provider, to help newsrooms worldwide to identify deepfakes and manipulated media through a free online training course. News organizations increasingly rely on third parties for large volumes of images and video, and identifying manipulated visuals is a significant challenge. This program aims to support newsrooms trying to do this work.

As these partnerships and our own insights evolve, so too will our policies toward manipulated media. In the meantime, we're committed to investing within Facebook and working with other stakeholders in this area to find solutions with real impact.

REDUCE: PREVENTING THE SPREAD OF VIRAL MISINFORMATION

For misinformation that does not directly violate our [Community Standards](#) but nonetheless is spammy or misleading, we take a number of steps to reduce its distribution, including demoting or down-ranking the content in News Feed, not recommending it, and more. In other words, we allow people to post false information as a form of expression if it doesn't violate our Community Standards, but we're not going to show it at the top of News Feed or Instagram

Explore. This significantly reduces the number of people who see the content at issue and helps us limit virality.

We use various signals to help us identify stories that might be false or misleading. Comments expressing disbelief (e.g. “no way this is real!”) are one signal that helps inform the process. We also use [feedback](#) from our community such as when people report something as false news. In many countries, including in the US, if we have signals that a piece of content is false, we temporarily reduce its distribution pending review by a third-party fact-checker.

FACT-CHECKING PROGRAM

We partner with independent third-party fact-checking organizations [certified](#) through the non-partisan International Fact-Checking Network (IFCN). These fact-checkers review videos, images, articles and text posts across our platform, check their facts, and [rate their accuracy](#). Fact-checkers can review content identified by Facebook's algorithms as potentially false, and may also identify stories to review on their own. We now have more than 50 partners around the world fact checking content in more than 40 languages, and are investing in ways to scale these efforts further. We recently brought the global fact-checking program to Instagram.

When a fact-checker rates a story as false, we reduce its distribution so fewer people see it—on Instagram, we remove it from Explore and hashtags, and on Facebook, we reduce its distribution by showing it lower in News Feed. This helps stop the hoax from spreading and reduces the number of people who see it.

On Instagram, we also make content from accounts that repeatedly post misinformation harder to find by filtering content from that account from Explore and hashtag pages for example.

And on Facebook, if Pages, domains or Groups repeatedly share misinformation, we'll continue to reduce their overall distribution and we'll place restrictions on the Page's ability to advertise and monetize.

In December 2019, we launched a [pilot program](#) that leverages Facebook's community as researchers to find information that contradicts the most obvious online hoaxes or corroborates other claims. Their findings will be shared with the third-party fact-checkers as additional context as they do their own official review. This is another signal to help fact-checkers identify false content faster.

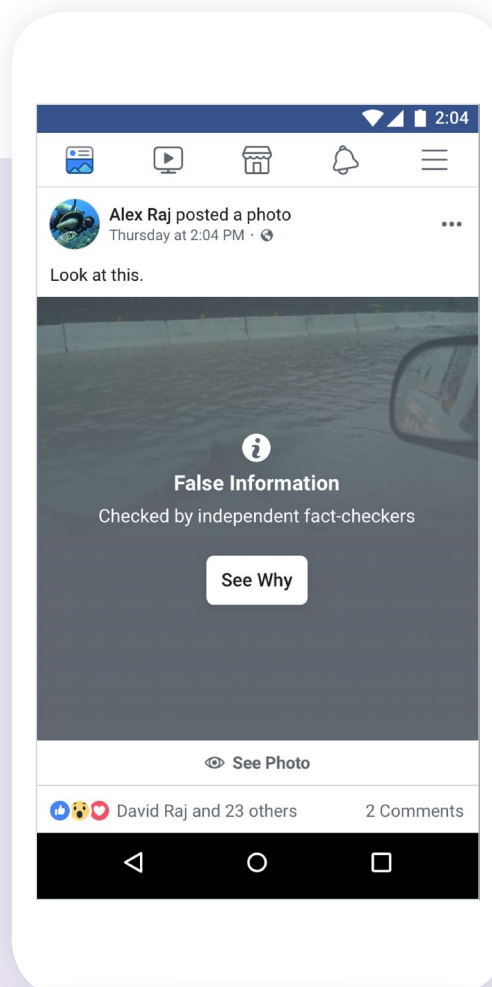
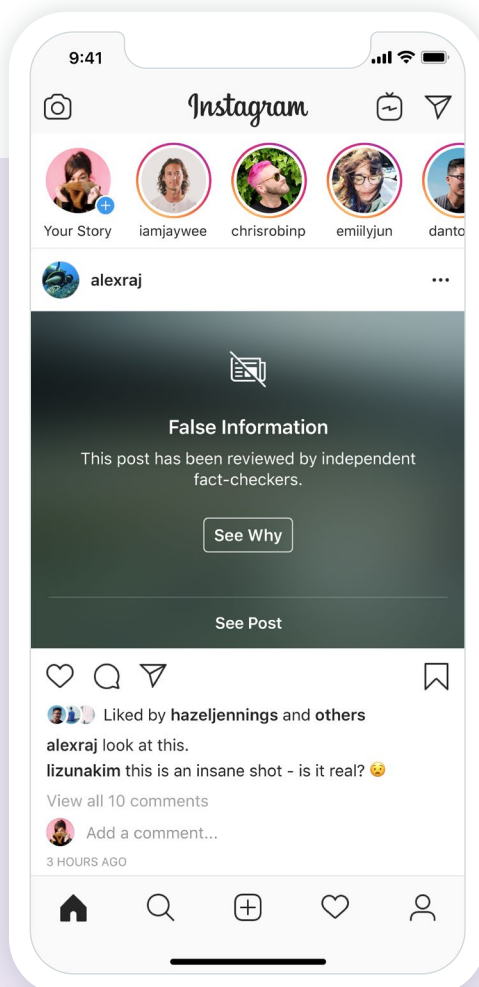
We are always looking to expand our fact-checking program, and continue to bring on new partners ahead of the election. We keep

an open line of communication with these partners about critical topics such as voter interference and suppression, and are providing them with public tools including CrowdTangle Live Displays, that help them better track voting-related keywords and content online.

INFORM: HELPING PEOPLE BETTER UNDERSTAND WHAT THEY SEE ONLINE

The answer to misinformation isn't less information—it's more context. We've heard from experts that if we simply removed misinformation from Facebook, the content would still exist elsewhere on the internet and broader social media ecosystem. So, we provide more context and information from fact-checkers on false news so people can better decide for themselves what to read, trust, and share.

In October 2019, we introduced more prominent warning labels across Facebook and Instagram content rated false or partly false by a third-party fact-checker. The labels are shown on top of false and



LABELING MISINFORMATION

Content across Facebook and Instagram that has been rated false or partly false by a third-party fact-checker is more prominently labeled so that people can better decide for themselves what to read, trust and share.

partly false photos and videos, including on top of Stories content on Instagram, and will link out to the assessment from the fact-checker.

Additionally, a pop-up will appear when people attempt to share posts on Facebook and Instagram that include content that has been debunked by third-party fact-checkers.

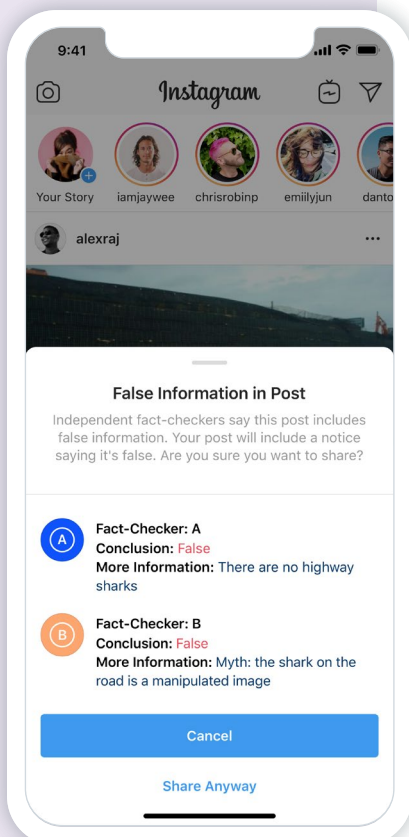
We also partner with organizations and experts in media literacy. In October 2019, we announced an initial investment of \$2 million to support projects that empower people to determine what to read and share—both on Facebook and elsewhere. In January 2020, we provided an update on our partnerships for these new [media literacy initiatives](#) including:

- **Poynter’s MediaWise** to launch a new national media literacy program called the [MediaWise Voter Project](#) to reach students who will be voting for the first time in the 2020 election. The MediaWise Voter Project kicked-off with events on college campuses in January where the MediaWise team trained hundreds of students.
- **PEN America**, a literary and free expression organization, to launch a new program called ‘Knowing the News’, a series of activities in ten cities across the country that will help give news consumers a better understanding of the journalistic process and equip communities with tools to assess the trustworthiness of information they receive through workshops and visits to local newsrooms.
- And **Arizona State University News Co/Lab**, that will develop a multi-platform suite of media literacy resources for adults, specifically for senior citizens, first-time voters, and the general voting public. The project will include three videos, co-produced with Arizona PBS.

In addition, we’ve included a new series of media literacy lessons in our Digital Literacy Library. These lessons are drawn from the Youth and Media team at the Berkman Klein Center for Internet & Society at Harvard University, which has made them available for free worldwide under a Creative Commons license. The lessons, created for middle and high school educators, are designed to be interactive and cover topics ranging from assessing the quality of the information online to more technical skills like reverse image search.

PREVENTING MISINFORMATION SHARING

A pop-up appears when people attempt to share posts on Instagram that include content that has been debunked by third-party fact-checkers.



06

VI. Supporting Political Speech

Facebook's approach toward political speech on the platform is grounded in the fundamental belief in free expression and respect for the democratic process, as well as the fact that, in mature democracies with a free press, political speech is already arguably the most scrutinized speech there is.

In October 2019, Mark Zuckerberg spoke at Georgetown University about the importance of protecting free expression. While we're increasingly seeing laws and regulations around the world that undermine free expression and human rights, he underscored his strong belief that giving everyone a voice empowers the powerless and free expression is essential for driving progress and building a more inclusive society around the world. You can read Mark's full speech [here](#).

These are some of the principles he mentioned that remain core to Facebook's values:

- Like the other major internet platforms and media, Facebook doesn't fact-check content or ads from politicians. By limiting this speech we would leave people less informed about what their elected officials are saying and leave politicians less accountable for their words.
- We don't think it's right for a private company to censor politicians or the news in a democracy.
- As a principle, in a democracy, people should decide what is credible, not technology companies.
- While we may disagree on exactly where to draw the line on specific issues, we at least can disagree. That's what free expression is.
- Democracy depends on the idea that we hold each others' right to express ourselves and be heard above our own desire to always get the outcomes we want.

We ask our third party fact-checkers not to review organic content or ads from politicians. It is consistent with our view that it's not an appropriate role for us to referee political debates and prevent a politician's speech from reaching its audience, where it will be subject to public debate. Our policies don't mean that politicians

can say whatever they want on Facebook. They can't, for example, spread misinformation about where, when or how to vote or post content that may incite violence. And when a politician shares vaccine hoaxes or otherwise debunked content, including links, videos and photos, we reduce the distribution of that content, and reject its inclusion in advertisements.

NEWSWORTHINESS EXEMPTION

Facebook has had a newsworthiness exemption [since 2016](#). This means that we may allow organic content that otherwise violates our Community Standards when the public interest value outweighs the risk of harm. In making this assessment, our determination is holistic and accounts for international human rights standards. In the case of politicians' speech, for example, we presume a public interest value in people hearing from and responding to politicians, but the speech is still evaluated against the risk of harm.

In making a determination as to the public interest value, we consider a number of factors, including:

- Country-specific circumstances like whether there is an election underway or the country is at war
- The subject matter of the speech, including whether it relates to governance or politics and who the target is; and, among other things,
- The political structure of the country

In determining the risk of harm, we have established a robust framework that looks at factors, including:

- Potential to incite violence
- Efforts to silence or intimidate others
- Civic harms, including voter interference
- Impact from bullying or harassment

The newsworthiness exemption only applies to organic content; all ads, including those posted by politicians, must still comply with our Community Standards and our Advertising Policies, and will not be deemed as newsworthy no matter who runs them.

07

VII. Global Partnerships and Coordination

We know we can't do this work alone, which is why we partner with outside experts to increase our effectiveness. We have made partnerships with organizations focused on civic engagement and civic tech a priority. One of the biggest challenges we have as we build tools is making sure those tools are a positive force for citizen participation and democracy. Through strategic partnerships, we are able to create tools and inform people who use Facebook about elections, candidates and voting rules so they are prepared and excited to vote on Election Day. These partnerships help drive civic engagement and encourage civic participation everywhere.

Some of our key partners include:

[Atlantic Council](#) is an organization we work with on fighting foreign interference through the Atlantic Council's Digital Forensics Research Lab. They are key in providing more context into who and what the networks we take down are.

[Democracy Works](#), a nonpartisan organization that supplies data for hundreds of elections at every level across the country. Democracy Works also maintains the Voting Information Project, a partnership with states to provide official information to voters about where to vote and what's on their ballots where they look for it most—online. Democracy Works provides elections information and voter registration resources to civic organizations across the political spectrum, making it possible for underserved communities to easily access accurate, current information about elections and voting.

[The Newseum Foundation](#), which offers free resources on media literacy and education.

[Nonprofit Vote](#), who partners with America's nonprofits to help the people they serve participate and vote. Nonprofit Vote is a leading source of nonpartisan resources to help nonprofits integrate voter engagement into their ongoing activities and services. Nonprofit Vote organizes and runs National Voter Registration Day.

[VoteRiders](#), a national nonpartisan, nonprofit organization providing free information and help to voters to make sure they have the right kind of ID to vote in their state. VoteRiders is a leading organization focused exclusively on voter ID.

We also work with a variety of organizations that work creatively to promote civic engagement in communities across the country, including the Civic Alliance, HeadCount, the Alliance for Youth Action, the National Archives, and the National Conference on Citizenship.

This partnership work also includes our Election Research Commission, to help provide independent, credible research about the role of social media in elections, as well as democracy more generally. This is an important new model for partnerships between industry and academia.

We continue to improve our coordination and cooperation with law enforcement, including DNI, DHS, FBI, as well as federal officials, state election officials, and other technology companies, to allow for better information sharing and threat detection.

We are also working with academics, civil society groups, and researchers, including the Atlantic Council's Digital Forensic Research Laboratory and the Stanford Internet Research Laboratory to get the best thinking on these issues.

VII. Conclusion

Since 2016, we have fundamentally changed our approach to elections integrity. We have invested in both people and technology. We have more than 35,000 people dedicated specifically to safety and security, with about 40 teams contributing to our work on elections. We have dedicated unprecedented resources to finding and removing malicious activity on our platform.

And we have made progress. Through improved technology, we've increased our efforts towards finding and removing bad actors, blocking fake accounts, limiting the spread of false news and misinformation, and bringing industry-leading transparency to political advertising and Pages.

We've learned that stronger coordination and cooperation is key. We continue to improve coordination with law enforcement, federal officials, and other technology companies. We are also working with researchers, academics, and civil society groups around the world to get the best thinking on these issues. We know we can't do this alone, and these partnerships serve as a force multiplier to prevent both foreign and domestic meddling.

We take our role in elections seriously, and we are dedicated to doing everything we can to prevent bad actors from interfering in the US 2020 election.