

Facebook response to the European Commission Communication on Covid-19 Disinformation

Report for July and August 2021

Executive Summary

User engagement with authoritative sources:

- **COVID-19 Information Center:** More than 280 million people globally, including over 29 million people in the EU, visited the COVID-19 Information Center during July 2021. More than 260 million people globally, including over 24 million people in the EU, visited the COVID-19 Information Center during August 2021.
- **Using AI to Help Health Experts Address the COVID-19 Pandemic:** The Facebook AI team is partnering with academic researchers and other experts on (i) open-sourcing our COVID-19 forecasting models so that response teams, governments, and researchers can use them to further help their communities and (ii) extending these models to the European Union.
- **Improving the Navigation into Facebook Settings:** We have redesigned the Facebook Settings page to make our tools easier to find. Settings are now grouped into six broad categories, each containing several related settings: Account, Preferences, Audience and Visibility, Permissions, Your Information, Community Standards and Legal Policies.
- **New Ways to Protect Our Community from Abuse on Instagram:** We have launched a set of new features to help protect people from abuse on Instagram: The ability for people to limit comments and direct messaging requests during spikes of increased attention; stronger warnings when people try to post potentially offensive comments and the global rollout of our Hidden Words feature, to filter abusive DM requests.
- **Introducing Sensitive Content Control on Instagram:** We have recently launched the “Sensitive Content Control,” which allows people to decide to leave things as they are, assuming they’re satisfied with their experience today, or to adjust the Sensitive Content Control to see more or less of some types of sensitive content.
- **Facebook Research on Preference Amplification in Recommender Systems:** Facebook recently published a paper on the interactions between users and recommender systems, exploring how these interactions may lead to potential harmful outcomes, and how understanding the long-term impact of ML systems helps build better safeguards and ensure that our models are optimized to serve the best interests of our users.

Actions on misinformation:

- **Applying Community Standards:** During the month of July, we removed over 880 thousand pieces of content on Facebook and Instagram globally, including over 110 thousand pieces of content in the EU, for violating our [COVID-19 and vaccine misinformation policies](#). During the month of August, we removed over 1.1 million pieces of content on Facebook and Instagram globally, including over 160 thousand pieces of content in the EU, for violating our [COVID-19 and vaccine misinformation policies](#).
- **Promoting Reliable Information and Reducing Harmful Misinformation About COVID-19:** Since the start of the pandemic through June:
 - We removed more than 20 million pieces of content from Facebook and Instagram globally for violating our [policies on COVID-19-related misinformation](#).
 - We removed over 3,000 accounts, pages, and groups for repeatedly violating our rules against spreading COVID-19 and vaccine misinformation.
 - We displayed warnings on more than 190 million pieces of COVID-related content on Facebook rated by our third-party fact-checking partners.
 - More than 18 million people globally have used Facebook profile frames supporting vaccines, and more than 25% of people globally on Facebook have already seen someone use the UNICEF COVID-19 vaccine profile frames.
 - Over 29% of people globally on Instagram have seen someone use a COVID-19 vaccine sticker. And Instagram COVID-19 vaccine stickers have been used by more than 7.6 million people globally.
- **Third-Party Fact-Checking:** We partner with over [80 fact-checking organizations](#) around the world, covering over 60 languages. In the EU and greater Europe, we partner with 36 fact-checking organizations, covering 33 languages.
- **Partnering with Meedan Digital Health Lab to Help Fact-Checkers Debunk Health Misinformation:** Through the Facebook Journalism Project, we have created a new partnership with the Digital Health Lab at Meedan, a global technology nonprofit, to provide Facebook's third-party fact-checking partners with free virtual training sessions and access to Meedan's Health Desk.
- **Supporting Media Literacy in Europe:** In July, we extended our partnership with fact-checking partner AFP by launching five new educational videos, intended to share tools, tips and methods with the public to help verify information, conduct online research, and use applications to verify an image or video.
- **Community Standards Enforcement Report:** In July, we published our Community Standards Enforcement Report on how we enforced our policies on Facebook and Instagram from April through June 2021.

Coordinated inauthentic behaviour: In July, we removed 144 Facebook accounts, 262 Instagram accounts, 13 Pages, and 8 Groups originating in Myanmar and Russia targeting both domestic audiences as well as audiences in other countries.

Advertising: As the COVID-19 crisis evolves, so do our policies. For a full list of our Advertising Policies about COVID-19, see [here](#).

1. Introduction

This report builds off our previous reports to the European Commission, in response to the [Joint Communication](#) for tackling COVID-19 disinformation, and provides an overview of the policies, products, and processes we have deployed to combat COVID-19 misinformation and disinformation across Facebook and Instagram in July and August 2021.

2. User Engagement with Authoritative Resources and Tools to Raise Awareness

We continue to find new ways to connect people with accurate, reliable and authoritative information. This is a core component of our strategy to combat misinformation because we want to be able to provide our users with the means to decide what to read, trust and share.

We believe that informing people with accurate and authoritative information, as well as more context, is an approach that can be more impactful than the alternative of just removing content. If we simply removed all posts flagged by fact-checkers as false, for example, the content would still be available elsewhere on the internet, other social media platforms, or even discussed around the dinner table. By leaving this content up and surfacing research from fact-checkers or pointing people to reliable information, we're providing people with important information and context.

As noted by an international group of [human rights experts](#) (in relation to COVID-19): "it is essential that governments and internet companies address disinformation in the first instance by themselves providing reliable information. Resorting to other measures, such as content take-downs and censorship, may result in limiting access to important information for public health and should only be undertaken where they meet the standards of necessity and proportionality."

During the coronavirus public health crisis, we have been supporting the global public health community's work to keep people safe and informed by connecting them to accurate, reliable, accessible and relevant sources of information about COVID-19. Our COVID-19 Information Center on Facebook provides people with the latest information from health authorities, news, resources, facts, and tips to stay healthy and safe. It is available globally, including all 27 EU member states. More than 280 million people globally, including over 29 million people in the EU, visited the COVID-19 Information Center during July 2021. More than 260 million people globally, including over 24 million people in the EU, visited the COVID-19 Information Center during August 2021.

Using AI to Help Health Experts Address the COVID-19 Pandemic

As part of the different tools that Facebook is putting forward, the Facebook AI team is partnering with academic researchers and other experts on a range of initiatives related to the COVID-19 pandemic. AI has been an important tool for public health experts around the world in their efforts to keep people safe and informed.

Open-sourcing our COVID-19 forecasting models: Since April of last year, we have created and shared high-quality, localized COVID-19 forecasting models to help healthcare providers and emergency responders determine how best to plan and allocate their resources in their particular area. We are now open-sourcing our entire stack of COVID-19 forecasting models so that response teams, governments, and researchers can use them to further help their communities. These AI models, developed by Facebook AI in collaboration with academic researchers at New York University's Courant Institute of Mathematical Sciences, the Universitat Politècnica de Catalunya (UPC), and the Faculty of Mathematics and the Data Science research platform at the University of Vienna, use publicly available, de-identified time series data about the spread of the disease. They have consistently been among the most accurate models since the beginning of the pandemic. Through this open research effort, we also hope to help advance epidemiological forecasting as it allows researchers to reuse, extend, and improve our methods.

Extension of our model to the European Union: As part of open-sourcing our code, we are also supporting UPC to extend our model to the European Union. Since early in the pandemic, the BIOCOS-SC team at the Universitat Politècnica de Catalunya has provided comprehensive reports and forecasts to the European Commission about the spread of COVID-19. Based on the success of our forecasts in the United States, BIOCOS-SC is now leading the effort to apply the model to similar scenarios in the European Union. First experiments by Enric Alvarez Lacalle, a professor of physics at the Universitat Politècnica de Catalunya, and collaborators have shown promising results, especially for longer forecast horizons and higher spatial resolution. We invite the global research community to explore our open source code for similar applications and forecasting tasks.

Improving the Navigation into Facebook Settings

As part of our ongoing efforts to improve people's experience on Facebook, we have redesigned the Facebook Settings page to make our tools easier to find. Whether it's managing the ads people see, adjusting sharing settings, or curating an audience for posts, people shouldn't have to think too hard about where to start. That's why we've reduced the number of categories, and renamed them: settings are now grouped into six broad categories, each containing several related settings: Account, Preferences, Audience and Visibility, Permissions, Your Information, and Community Standards and Legal Policies.

We want to design products and features that make information and tools accessible and easy to find. We're confident this new settings page will make it easier for people to visit their settings, find what they came for, and make the changes they want.

New Ways to Protect Our Community from Abuse on Instagram

We have launched a set of new features to help protect people from abuse on Instagram:

- The ability for people to limit comments and direct messaging (DM) requests during spikes of increased attention;
- Stronger warnings when people try to post potentially offensive comments;
- The global rollout of our Hidden Words feature, which allows people to filter abusive DM requests.

We have a responsibility to make sure everyone feels safe when they come to Instagram. We don't allow hate speech or bullying on Instagram, and we remove it whenever we find it. We also want to protect people from having to experience this abuse in the first place, which is why we're constantly listening to feedback from experts and our community, and developing new features to give people more control over their experience on Instagram, and help protect them from abuse.

Introducing Sensitive Content Control on Instagram

People have very different views about the experience we offer, and if it's the right one for them. We believe people should be able to shape Instagram into the experience that they want. We've started to move in this direction with tools like the ability for people to turn off comments, or restrict someone from interacting with them on Instagram. We have recently launched the "Sensitive Content Control," which allows people to decide how much sensitive content shows up in Explore.

This new feature gives people control over sensitive content. People can decide to leave things as they are, assuming they're satisfied with their experience today, or they can adjust the Sensitive Content Control to see more or less of some types of sensitive content. We recognize that everybody has different preferences for what they want to see on Explore, and this control will give people more choice over what they see.

Facebook Research on Preference Amplification in Recommender Systems

Facebook recently published a paper on the interactions between users and recommender systems, exploring how these interactions may lead to potential harmful outcomes.

The repeated interplay between people and algorithms creates a feedback loop that results in recommendations that are increasingly customized to people's tastes. Ideally, these feedback loops ought to be virtuous all the time and the recommender system provides users with recommendations that enhance the quality of their lives. However, the question is what happens when the system overindexes and amplifies interactions that do not necessarily capture the user's true preferences? Or if the user's preferences have drifted toward recommended items that could be considered harmful or detrimental to their long-term well-being? Under what conditions would recommender systems respond to these changes and amplify preferences leading to a higher prevalence of harmful recommendations?

Our main assumption is that users have a slight inclination to reinforce their opinion (or drift), i.e. a situation in which the user preferences do not change in response to the system's recommendations. We show that even under a mild drift and absent any external intervention, no such fixed point exists. That is, even a slight preference by a user for recommendations in a given category can lead to increasingly higher concentrations of item recommendations from that category. We refer to this phenomenon as preference amplification.

Several strategies to mitigate the effects of preference amplification of problematic entities were tested in this paper, including an experiment we conducted using a real-world large-scale video recommender system. In the experiment, we downrank videos considered to include borderline nudity (the platform already filters out videos that violate community standards) for users who have a high level of exposure to them consistently. The results of the experiment show that in addition to reducing exposure of this content in the impacted population, we saw that overall engagement go up by +2%. These results are highly encouraging, as not only we can prevent exposure to problematic content, we also have an overall positive effect on the user experience.

Understanding the long-term impact of ML systems helps us, as practitioners, to build better safeguards and ensure that our models are optimized to serve the best interests of our users.

The full paper can be accessed [here](#).

3. Actions on Misinformation

Our goal is to create a place for expression and give people a voice. Building community and bringing the world closer together depends on people's ability to share diverse views, experiences, ideas and information. Our commitment to expression is paramount, but we recognise that the internet creates new and increased opportunities for abuse. When considering whether to provide more context, allow, reduce distribution, or remove misinformation, we do it in service of one or more of our [Community Standards Values](#): voice, authenticity, safety, privacy, and dignity.

We define misinformation as content that is false or misleading. We enforce on misinformation by looking at content or behaviors that violate our Community Standards or content that may be reviewed by our third-party fact-checking partners. We define disinformation as coordinated efforts to manipulate public debate for a strategic goal.

Applying Community Standards to COVID-19 Content

As people around the world confront this unprecedented public health emergency, we want to make sure that our [Community Standards](#) (Facebook) and [Community Guidelines](#) (Instagram) protect people from harmful content and new types of abuse related to COVID-19. We're working to remove content that has the potential to contribute to offline harm, including through our policies prohibiting the coordination of harm, hate speech, bullying and harassment, and misinformation that contributes to the risk of imminent violence or physical harm. Oftentimes, misinformation can cut across different types of abuse areas; for example, a racial slur could be

coupled with a false claim about a group of people and we'd remove it for violating our hate speech policy. So in addition to our misinformation policies, we have a number of other ways we might combat COVID-19 misinformation. The following are updated metrics:

- During the month of July, we removed over 880 thousand pieces of content on Facebook and Instagram globally, including over 110 thousand pieces of content in the EU, for violating our [COVID-19 and vaccine misinformation policies](#). During the month of August, we removed over 1.1 million pieces of content on Facebook and Instagram globally, including over 160 thousand pieces of content in the EU, for violating our [COVID-19 and vaccine misinformation policies](#).
- Under our **Hate Speech** policy, we are removing content that states that people who share a protected characteristic such as race or religion have the virus, created the virus or are spreading the virus. This does not apply to claims about people based on national origin because we want to allow discussion focused on national-level responses and effects (e.g., "X number of Italians have COVID-19"). We also remove content that mocks people who share a protected characteristic such as race or religion for having COVID-19. We report on Hate Speech content removals every quarter in our [Community Standards Enforcement Report](#).
- Under our **Bullying and Harassment** policy, we remove content that targets people maliciously, including content that claims that a private individual has COVID-19, unless that person has self-declared or information about their health status is publicly available. We report on Bullying and Harassment content removals every quarter in our [Community Standards Enforcement Report](#).

As the situation evolves, we are continuing to look at content on the platform, assess speech trends and engage with experts, and will provide additional policy guidance to our [Community Standards](#) (Facebook) and our [Community Guidelines](#) (Instagram) when appropriate to keep the members of our community safe during this crisis.

Promoting Reliable Information and Reducing Harmful Misinformation About COVID-19

COVID-19 is still a major public health issue, and we are committed to helping people get authoritative information, including vaccine information. We continue to remove harmful COVID-19 misinformation and prohibit ads that try to exploit the pandemic for financial gain. Since the start of the pandemic through June:

- We removed more than 20 million pieces of content from Facebook and Instagram globally for violating our [policies on COVID-19-related misinformation](#).
- We removed over 3,000 accounts, pages, and groups for repeatedly violating our rules against [spreading COVID-19 and vaccine misinformation](#).
- We displayed warnings on more than 190 million pieces of COVID-related content on Facebook that our third-party fact-checking partners rated as false, partly false, altered

or missing context, collaborating with 80 fact-checking organizations in more than 60 languages around the world. When they rate a piece of content with one of these ratings, we add a prominent label warning people before they share it and show it lower in people's feed.

We've provided authoritative information to help improve vaccine acceptance, connecting 2 billion people to resources from health experts through our COVID-19 Information Center and educational pop-ups on Facebook and Instagram and helping 4 million people in the US alone access vaccines through our vaccine finder tool.

We know from [public health research](#) that people are more likely to get vaccinated if they see others in their community doing so. In countries where vaccines are available to most people, we ramped up our efforts to show when friends and neighbors share their support for vaccines through profile frames and stickers.

- Since the beginning of the pandemic, more than 18 million people globally have used Facebook profile frames supporting vaccines.
- More than 25% of people globally on Facebook have already seen someone use the UNICEF COVID-19 vaccine profile frames.
- Over 29% of people globally on Instagram have seen someone use a COVID-19 vaccine sticker.
- Instagram COVID-19 vaccine stickers have been used by more than 7.6 million people globally.

Globally, we have seen vaccine acceptance rising. For example, our COVID-19 Trends and Impact Survey data which we conduct in partnership with Carnegie-Mellon and University of Maryland has since the beginning of the year shown vaccine acceptance rising by 35% in France, 25% in Indonesia, and 20% in Nigeria. For people in the US on Facebook, vaccine hesitancy has declined by 50%.

Our Third-Party Fact-Checking Program in Europe

For misinformation that does not lead to real world harm, but undermines the authenticity and integrity of our platform, we continue to work with our growing [network of independent third party fact-checking partners](#). We partner with over [80 fact-checking organizations](#) around the world, covering over 60 languages. In the EU and greater Europe, we partner with 36 fact-checking organizations, covering 33 languages.

Partnering with Meedan Digital Health Lab to Help Fact-Checkers Debunk Health Misinformation

Through the Facebook Journalism Project, we have created a new partnership with the Digital Health Lab at Meedan, a global technology nonprofit, to support fact-checkers globally in

fighting health misinformation online. This is the latest step in our ongoing effort to [support fact-checkers](#) and [connect them to the resources](#) they need to address health misinformation.

Meedan's Digital Health Lab will facilitate a series of virtual training sessions between its team of doctors, scientists and health experts and Facebook's [third-party fact-checking partners](#). We know that strong partnerships and sharing expert information on best practices and approaches play a big role in effectively addressing misinformation. That's why the Lab's team of experts and our third-party fact-checking partners will work together to design the training sessions, identify the needs of fact-checkers, and share potential solutions. The trainings will be available at no cost to all of Facebook's fact-checking partners.

In addition to the training opportunities, Facebook's fact-checking partners will also have access to Meedan's [Health Desk](#). This will provide them with timely answers to key public health questions to support their health-related fact-checking efforts.

Supporting Media Literacy in Europe

We want to give people the tools to make more informed decisions about the information they choose to consume and share online. We regularly run media literacy campaigns to help people identify and detect false news. For example, at the start of the Covid-19 crisis in France, we collaborated with our fact-checking partner AFP to produce a series of videos in which fact-checkers deconstructed viral fake news on Covid-19.

In July, we extended our partnership in France to provide educational video to combat covid-19 misinformation by launching five new educational videos, intended to share tools, tips and methods with the public to help verify information, conduct online research, and use applications to verify an image or video. More information can be found [here](#).

Community Standards Enforcement Report, Second Quarter 2021

We're committed to sharing meaningful data so we can be held accountable for our progress, even if the data shows areas where we need to do better. We publish our [Community Standards Enforcement Report](#) on a quarterly basis to more effectively track our progress and demonstrate our continued commitment to making Facebook and Instagram safe and inclusive. The report provides metrics on how we enforced our policies from April through June 2021. This is our 10th report and some of our long-term trends include:

- Prevalence of hate speech has decreased for three quarters in a row since we first began reporting it. This is due to improvements in proactively detecting hate speech and ranking changes in News Feed.
- Hate speech content removal has increased over 15X on Facebook and Instagram since we first began reporting it.

- Our proactive rate (the percentage of content we took action on that we found before a user reported it to us) is over 90% for 12 out of 13 policy areas on Facebook and nine out of 11 on Instagram.
- We now include 13 policy areas on Facebook and 11 on Instagram, and have added new metrics on appeals, restores, and prevalence.

Some other highlights from the report on our content moderation efforts include:

- **Hate Speech:** We continue to make progress on Hate Speech, as prevalence continued to decrease for the third quarter in a row. In Q2, it was 0.05%, or 5 views per 10,000 views, down from 0.05-0.06%, or 5 to 6 views per 10,000 views in Q1. This is due to continued improvement in our proactive detection. Our investments in AI enable us to detect more kinds of hate speech violations on Facebook and Instagram. This technology helps us enforce our policies across billions of users and multiple languages. Steady, continuous AI improvements and advancements, such as the Reinforcement Integrity Optimizer (RIO), enable our AI models to spot hate speech using real-world data and improve over time. And in terms of content removed for violating our Hate Speech policy, on Facebook, we took action on 31.5 million pieces of content in Q2 2021, 97.6% of which were found and actioned on before it was reported by people. While on Instagram, we took action on 9.8 million pieces of content in Q2 2021, 95.1% of which were found and actioned on before it was reported by people.
- **Bullying and Harassment:** Because bullying and harassment are highly personal by nature, in many instances, we need a person to report this behavior to us before we can identify or remove it. This means using technology to proactively detect bullying and harassment can be more challenging than other violation types. However, we continue to invest in our proactive detection technology to ensure we are tackling the problem and protecting our community. On Facebook, we took action on 7.9 million pieces of content in Q2 2021, 54.1% of which were found and actioned on before it was reported by people. On Instagram, we took action on 4.5 million pieces of content in Q2 2021, 71.5% of which were found and actioned on before it was reported by people.
- **Fake Accounts:** Our goal is to remove as many fake accounts on Facebook as we can. These include accounts created with malicious intent to violate our policies and personal profiles created to represent a business, organization or non-human entity, such as a pet. We prioritize enforcement against fake accounts that seek to cause harm. Many of these accounts are used in spam campaigns and are financially motivated. We expect the number of accounts we action to vary over time due to the unpredictable nature of adversarial account creation. Our detection technology helps us block millions of attempts to create fake accounts every day and detect millions more, often within minutes after creation. We do not include blocked attempts in the metrics we report here. On Facebook, we took action on 1.7 billion accounts in Q2 2021, 99.8% of which were found and actioned on before it was reported by people.

The full Community Standards Enforcement Report can be found [here](#).

4. Coordinated Inauthentic Behavior (CIB) and Influence Operations

We know that people looking to mislead others - whether through phishing, scams, or influence operations - try to leverage crises in order to advance their goals, and the COVID-19 pandemic is no different. As the situation evolves, we are actively working to find and stop coordinated campaigns that seek to manipulate public debate across our platforms.

Historically, influence operations have manifested in different forms: from covert campaigns that rely on fake identities to overt state media efforts that use authentic and influential voices to promote messages that may or may not be false on their face. As we've studied emerging adversarial tactics and actors, our understanding of influence operations has evolved. Today, we define **influence operations** as coordinated efforts to manipulate or corrupt public debate for a strategic goal.

Tackling the many tactics that make up influence operations requires multiple approaches, which is one reason why, when designing policies intended to counter IO, it is important to distinguish deceptive content from deceptive behavior.

The most egregious form of this type of deception is **Coordinated Inauthentic Behavior (CIB)**: CIB is any coordinated network of accounts, Pages and Groups on our platforms that centrally relies on fake accounts to mislead Facebook and people using our services about who is behind the operation and what they are doing. Since 2017, Facebook's Coordinated Inauthentic Behavior policy has been a primary vehicle for enforcing against these covert deception campaigns.

Our approach to Coordinated Inauthentic Behavior (CIB), and Influence Operations (IO) more broadly, is grounded on behavior- and actor-based enforcement. This means that we are looking for specific violating behaviours exhibited by violating actors, rather than violating content (which is predicated on other specific violations of our Community Standards, such as [misinformation](#) and hate speech). Therefore, when CIB networks are taken down, it is based on their behavior, not the content they posted. For a comprehensive overview of our approach, see [here](#).

To date, we have not found evidence of influence operations created to focus specifically on COVID-19. What we have seen is that people behind campaigns opportunistically use coronavirus-related posts among many other topics to build an audience and drive people to their Pages or off-platform sites.

July 2021 Coordinated Inauthentic Behavior Report

Our teams continue to focus on finding and removing deceptive campaigns around the world — whether they are foreign or domestic. In July, we removed two networks from Russia and Myanmar. In this report, we're also sharing an in-depth analysis by our threat intelligence team into one of the operations — a network from Russia linked to Fazze, a marketing firm registered in the UK — to add to the public reporting on this network's activity across over a dozen different platforms. We have shared information about our findings with industry partners, researchers, law enforcement and policymakers.

We know that influence operations will keep evolving in response to our enforcement, and new deceptive behaviors will emerge. We will continue to refine our enforcement and share our findings publicly. We are making progress rooting out this abuse, but as we've said before — it's an ongoing effort and we're committed to continually improving to stay ahead. That means building better technology, hiring more people and working closely with law enforcement, security experts and other companies.

Here are the numbers for the new CIB networks we removed in July:

- Total number of Facebook accounts removed: 144
- Total number of Instagram accounts removed: 262
- Total number of Pages removed: 13
- Total number of Groups removed: 8

Networks removed in July 2021:

- **Myanmar:** We removed 79 Facebook accounts, 13 Pages, eight Groups, and 19 Instagram accounts in Myanmar that targeted domestic audiences and were linked to individuals associated with the Myanmar military. We found this activity after reviewing information about a portion of it shared by a member of civil society in Myanmar. Our investigation revealed some links between this operation and the activity we removed in 2018.
- **Russia:** We removed 65 Facebook accounts and 243 Instagram accounts from Russia that we linked to Fazze, a subsidiary of a UK-registered marketing firm, whose operations were primarily conducted from Russia. Fazze is now banned from our platform. This cross-platform operation targeted audiences primarily in India, Latin America, and to a much lesser extent the United States. We found this network after reviewing public reporting about an off-platform portion of this activity.

A detailed report of the networks taken down and examples of content they posted can be found [here](#). Previous reports can be found [here](#).

5. Advertising

As the COVID-19 situation develops, we have implemented a variety of measures to prevent ads from being used to spread misinformation; to prevent ads from promoting content that could contribute to physical harm; to prohibit exploitative or deceptive ads; and provide transparency on ads about health issues. We have applied our [Advertising Policies](#) to new types of abuse that we're seeing on the platform. We have made adjustments to our enforcement protocols to prevent people from exploiting the COVID-19 pandemic, and continue adapting or removing temporary bans on specific products as the situation stabilizes.

For a full list of our Advertising Policies about COVID-19, see [here](#).