

15 September 2020

September Update on Microsoft Corporation's Efforts to Tackle COVID-19 Disinformation

Microsoft Corporation welcomes the opportunity to provide this report on the steps we have taken, and continue to take, to combat COVID-19 disinformation during the month of August 2020. This report supplements, and should be read in conjunction with, our report to the Commission of August 15 entitled *Report of Microsoft Corporation on Tackling COVID-19 Disinformation* ("August report," available [here](#)).

Microsoft has taken substantial steps to promote the safety of our users and customers during the pandemic, and we are continually working to refine our efforts to combat COVID-19 disinformation on our services as these threats evolve. As explained in our August report, however, because many Microsoft services are offered primarily or exclusively to enterprise customers, or tend to be used primarily by individuals acting in a professional capacity (e.g., LinkedIn), our services tend to be targeted less frequently by COVID-19 disinformation than many other services. This fact necessarily affects how we have responded to the COVID-19 disinformation threat. In particular, Microsoft has tailored its efforts to combat COVID-19 disinformation to the unique attributes of our respective services. Also, because few of our actions to fight COVID-19 disinformation are limited to individual Member States, we generally track and report these efforts on a global or EU-wide basis. In those instances where our efforts are limited to a certain Member State, we have stated that below.

We would note again that the steps we are taking to address COVID-19 disinformation, while important, are only a small part of the efforts we are taking to help our customers, partners, employees, and the communities in which we operate respond effectively to the COVID-19 crisis. For more information on these efforts, please see our dedicated website, "[Responding to COVID-19 together](#)".

A. Initiatives to promote authoritative content

Point A of the Commission's Information Request asks platforms to "provide data on the actions taken to promote information from national and international health agencies, national and EU authorities, as well as professional media, including information on the use of trustworthiness indicators for news sources." The Request also asks platforms to provide "relevant output indicators to assess the impact of such initiatives."

Microsoft services continue to take actions responsive to these points based on the nature of the service and the type of COVID-19 disinformation threats that the respective service confronts.

1. Bing

In our August report, we discussed how Bing regularly refines its search algorithms, and the metrics it uses to measure them, in order to elevate the prominence of authoritative, high-quality content in search results and to prevent manipulation of search results by bad actors—including those seeking to spread disinformation about COVID-19 (more information on Bing’s organic search ranking practices can be found [here](#)). Our August report also discussed additional, proactive steps that Bing is taking to present users with authoritative information about COVID-19. This includes placing “Answers” and/or Public Service Announcements about COVID-19 at the top of search results for a number of COVID-19-related search queries, which provide information and statistics from authoritative sources such as the World Health Organization and relevant national health authorities. Bing is also showing “task panes” in prominent places on the first page of search results, which provide users with authoritative information about COVID-19, such as how to reduce the risk of catching or spreading the virus.

As we also noted in our August report, for specific queries, Bing will point users to special COVID-19 “[information hubs](#)” on Microsoft News. The articles in these information hubs, currently available in 39 markets globally, are sourced from over 4500 trusted news brands worldwide. These hubs also contain links to official tools and information sources on the crisis. The screenshot below provides an example of how Microsoft presents these various sources of information to users when they input a search term related to COVID-19 (in this case, the word “coronavirus”):

The screenshot shows a Bing search results page for the query "coronavirus". At the top, there is a search bar with the Microsoft Bing logo and navigation tabs for ALL, WORK, IMAGES, VIDEOS, MAPS, and NEWS. Below the search bar, it indicates 214,000,000 results and offers filters for Date, Language, and Region. A secondary navigation bar includes "Coronavirus" (selected), Statistics, Trends, Prevention, Symptoms, and Treatment.

The main content area is divided into several sections:

- Coronavirus (COVID-19) pandemic**: A header section with a menu icon.
- News about coronavirus**: A section featuring a large article from "La Libre" titled "Coronavirus: de l'Autriche à Israël, la deuxième vague devient réalité" with a sub-headline "La fameuse 'deuxième vague' de la pandémie, redoutée dans le monde entier, est devenu dimanche une réalité...". Below this are three smaller news snippets: "Coronavirus: les dépenses en ligne ont augmenté de 33% au premier semestre 2020" (Le Soir - 6m), "Coronavirus : comment nos voisins européens tentent de freiner l'épidémie" (France 3 Régions - 10m), and "Vaccin contre le coronavirus : comment devient-on cobaye pour la recherche" (RTBF - 1h).
- Top resources**: A section with four links to information hubs: "Coronavirus COVID-19 | Nederlands", "Coronavirus COVID-19 | Français", "Coronavirus COVID-19 | Deutsch", and "Coronavirus COVID-19 | English".
- North Rhine-Westphalia cases**: A statistics panel updated 13 sept. at 23:04 local, showing Confirmed (62,079, +259), Deaths (1,828), and Recovered (-).
- Germany cases**: A statistics panel updated 13 sept. at 23:04 local, showing Confirmed (261,255, +1,029), Deaths (9,428, +4), and Recovered (222,899).
- Global cases**: A statistics panel updated 13 sept. at 23:04 local, showing Confirmed (28,797,422, +212,963), Deaths (920,942, +3,759), and Recovered (19,491,050, +222,845).
- Confirmed cases**: A map showing the geographical distribution of confirmed cases, with labels for NETHERLANDS, Hamburg, Bremen, Hannover, Berlin, Poznań, and Warsaw.

At the bottom of the main content area, it states: "Data from: CDC · WHO · ECDC · Wikipedia · The New York Times · See full list".

As a result of these efforts, a person entering a COVID-19-related search query on Bing is almost certain to be presented with authoritative content from trustworthy sources. From 1 August to 31 August of this year, these Bing sources of authority on COVID-19 had over 30 million visitors, including over 4 million visitors from all EU markets combined.

In addition to these direct efforts, Bing also is taking several less direct but equally important steps to promote authoritative content on COVID-19. For instance, Bing recently [announced](#) that it would provide access to Bing's tracker of COVID-19 cases by geographic area for academic and research uses. This dataset is pulled from reputable sources such as the World Health Organization, the U.S. Centers for Disease Control, and others. By making these data available for non-commercial public use, Bing will make it easier for medical researchers, government agencies, and academic institutions to create and disseminate authoritative, fact-based content on the pandemic as it evolves over time.

Also, in July, we announced that we would make data about COVID-19-related search queries on Bing freely available on GitHub, with updates every month over the course of the pandemic. This dataset includes explicit COVID-19 queries (containing search terms such as "corona," "coronavirus," and "covid"), as well as related queries that access the same set of web page search results. To protect user privacy, infrequent queries and queries containing personal data are filtered from the dataset, while raw query counts are replaced by popularity scores between 1 and 100 reflecting normalized query counts for a given day and country. We are [already seeing](#) data scientists use these data to help people visualize the progress of the pandemic and to develop practical data tools for COVID-19 response efforts.

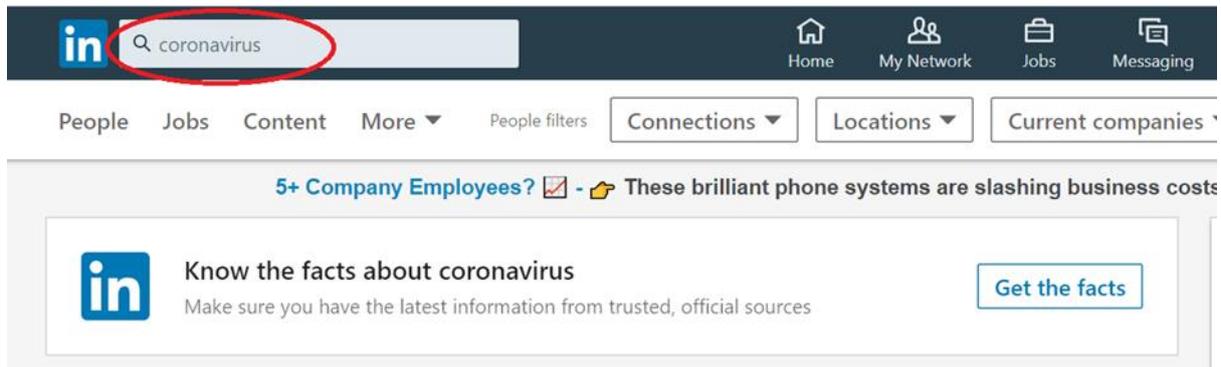
2. Microsoft Advertising

As described in our August report, Microsoft Advertising takes action against advertisements that contain disinformation about COVID-19 through our [Misleading Content policy](#), which prohibits advertising that can "reasonably [be] perceived as being deceptive, fraudulent or harmful to site visitors." On the basis of this policy, we have prohibited all advertising that seeks to exploit the COVID-19 crisis for commercial gain, spreads misinformation, or that may pose a danger to user health or safety. As a result of these efforts, Microsoft Advertising prevented a total of 3,871,425 advertiser submissions directly related to COVID-19 from being displayed to users globally in August 2020. This included 1,165,481 ads intended for display to users located in European markets.

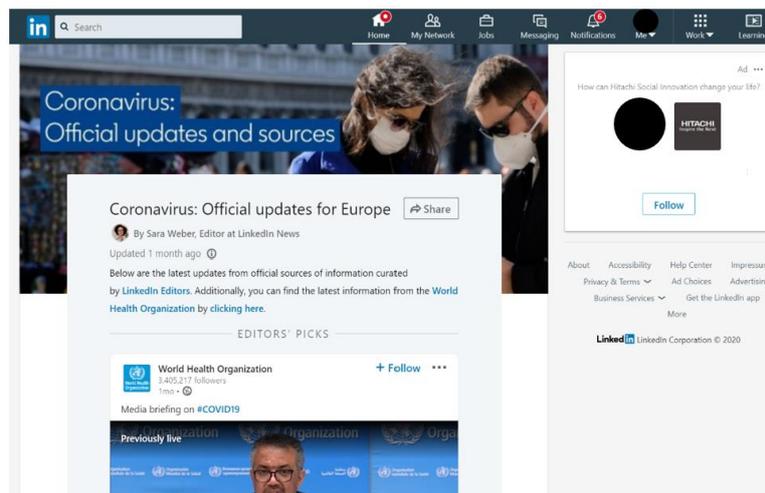
3. LinkedIn

As described in our August report, given LinkedIn's focus on enabling our members to meet other professionals in a network of trusted relationships, the service is an infrequent target of disinformation. This holds equally true with regard to COVID-19 disinformation. Accordingly, our efforts on COVID-19 have primarily focused on supporting the European Commission's efforts to combat disinformation by promoting trusted sources of information, and by extending the reach of authoritative government communications to local audiences.

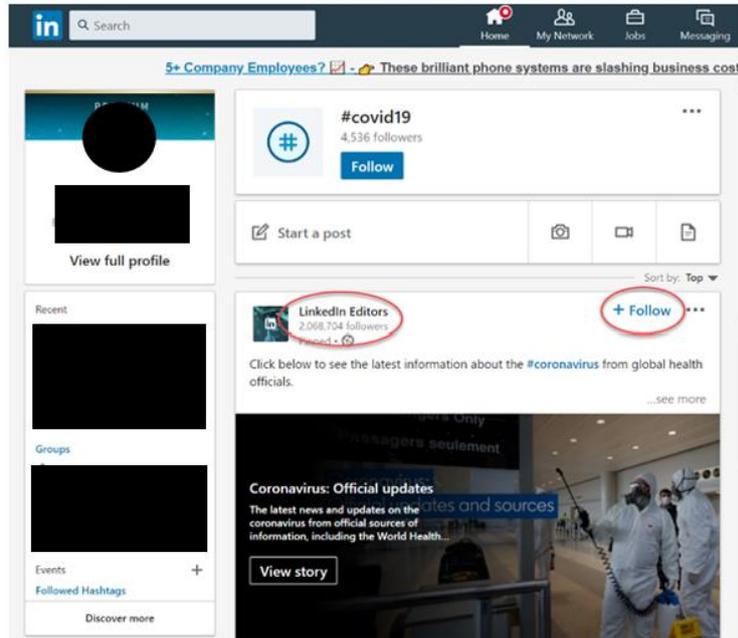
For instance, as illustrated in the screenshot below, we redirect any LinkedIn member that entered a search for the term "coronavirus" to a link labelled "*Know the facts about coronavirus*," which appears first in the list of search results:



By clicking on this link, members are directed to [LinkedIn's own official page on the coronavirus](#), which displays current information and broadcasts from verified international sources, primarily the World Health Organization (see screenshot below). The storylines on this page are available in 8 languages:



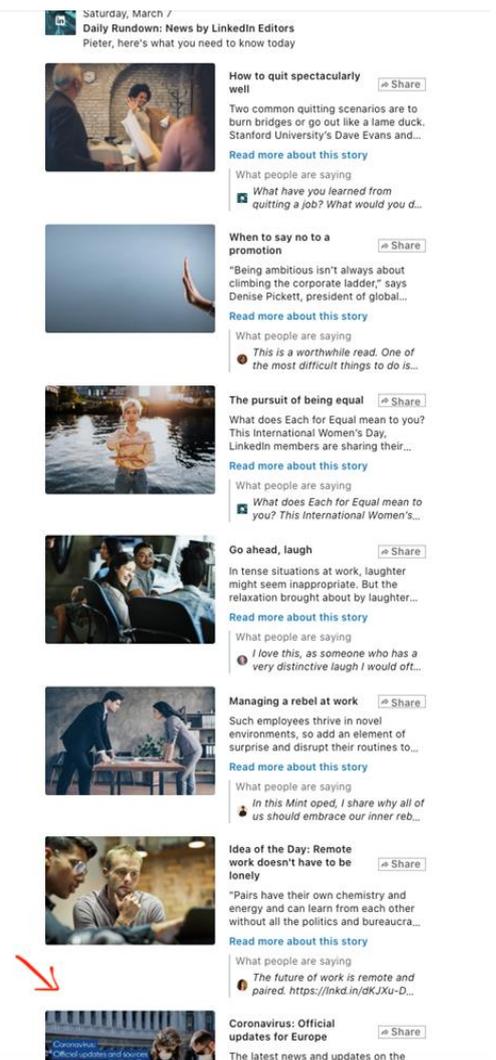
LinkedIn members can also undertake a search using #coronavirus or #COVID19 via the same search box on their LinkedIn profile page. This will produce a list of all posts that have been published on the LinkedIn platform using these hashtags (see screenshot below). The search result will also point members to the most recent post from LinkedIn News, which is generated by an in-house news team of 75+ journalists. By clicking on any of these posts, members are redirected to LinkedIn's official page on the coronavirus, as referenced above. Any member who wishes to receive regular updates from LinkedIn News can select the 'follow' option on any of the posts and all future updates will be added to the newsfeed of a member's own storyline.



LinkedIn also sends interested members a "European Daily Rundown," which is a summary of the day's news that is written and curated by experienced journalists. Members receive this in the form of a notification either in the news feed on their LinkedIn profile, or by email (see screenshot to the right). Since the start of the pandemic, and up until July, this notification included a reference to the "latest news on the coronavirus" and redirected members to [LinkedIn's official European page for the coronavirus](#). That page contains information and the most recent updates from official European sources such as the European Centre for Disease Prevention and Control, the European Public Health Association, the European Commission, and Member State governments.

The European Daily Rundown is distributed to LinkedIn members in all 27 EU member states and has a reach of approximately 9.7 million members. In addition, LinkedIn has local editions of the European Daily Rundown that are published in French, Italian, Spanish, Dutch and German. They are written and curated by a team of 18 experienced journalists based across the EU.

The chart below provides a list of the key trusted sources that the LinkedIn News team relies upon to source official information related to COVID-19. This list includes official pages of various Member State governments, as well as those of international organisations.



Company Page	Link	# of Followers on 31 August
Pan- European		
European Centre for Disease Prevention and Control (ECDC)	https://www.linkedin.com/company/ecdc/	28,494
European Public Health Association (EUPHA)	https://www.linkedin.com/company/eupha/	1,291
European Commission	https://www.linkedin.com/company/european-commission/	986,617
World Health Organization (WHO)	https://www.linkedin.com/company/world-health-organization/	3,362,231
Unicef	https://www.linkedin.com/company/unicef/	2,630,353
Red Cross	https://www.linkedin.com/company/ifrc/	131,844
United Nations	https://www.linkedin.com/company/united-nations/	3,090,151
France		
Government	https://www.linkedin.com/company/gouvernementfr/	67,814
Ministry of Health	https://www.linkedin.com/company/minist%C3%A8re-des-solidarites-de-la-sant%C3%A9/	95,560
Ministry of Labour	https://www.linkedin.com/company/minist%C3%A8re-du-travail/	176,190
Santé publique France	https://www.linkedin.com/company/sante-publique-france/	68,622
DACH		
Germany Ministry and minister of health (Bundesgesundheitsministerium)	https://www.linkedin.com/company/bundesgesundheitsministerium/	26,118
KfW Kreditanstalt für Wiederaufbau	https://www.linkedin.com/company/kfw/	35,946

German Bundesagentur für Arbeit (Federal Employment Agency)	https://www.linkedin.com/company/bundesagentur-fuer-arbeit/	9,558
Switzerland Ministry of Health (Bundesamt für Gesundheit)	https://www.linkedin.com/company/bag-ofsp/	61,618
Netherlands		
Ministry of health (Ministerie van VWS)	https://www.linkedin.com/company/ministerie-van-vws/	95,703
Governmental institute RIVM	https://www.linkedin.com/company/rivm/	37,035
Italy		
Ministero della Salute	https://www.linkedin.com/company/ministerosalute/	12,216
Presidenza del Consiglio (government)	https://www.linkedin.com/company/presidenza-del-consiglio-dei-ministri/	49,689
Unicef Italia	https://www.linkedin.com/company/unicef-italia/	4,187
Spain		
Minister of Health	https://www.linkedin.com/company/sanidadgob/	11,971
Minister of Science	https://www.linkedin.com/company/cienciagob/	56,060
Minister of Social Security	https://www.linkedin.com/company/ministerio-de-inclusi%C3%B3n-seguridad-social-y-migraciones/	1,137
Unicef Comité Español	https://www.linkedin.com/company/unicef-spain/	18,023

As noted in our August report, we also provide opportunities for Member State governments to use “LinkedIn Live” to share critical information on COVID-19 to our members. For example, LinkedIn News conducted live interviews with [Bruno Le Maire](#), France’s Minister for the economy, about the reopening of the French economy, and with [Jens Spahn](#), Germany’s Minister for health, about the way out of the crisis. We have also provided free advertising to governments to enable them to extend the reach of their official communications to their citizens. For example, we provided the Belgian Health Ministry and the German Ministry of Health with sponsored marketing campaigns to promote their official COVID-19 websites.

As the pandemic took hold, LinkedIn identified a significant increase in hiring needs in certain countries, particularly in the healthcare and social services sectors, as well as recruitment difficulties for roles held by the frontline workers. In response, LinkedIn designed a support package for hospitals, supply chains, supermarkets, freight delivery services, and disaster relief non-profits (e.g., Red Cross and United Way) to support them with their critical hiring needs. Examples of entities we supported include the public health service of the Paris region (AP-HP) and the Irish national health services. We also scaled our Recruiting for Good program to meet urgent hiring needs. Our LinkedIn recruiting team has volunteered their time to help organizations that are looking to fill critical volunteer and full-time positions to support these organizations' COVID-19 responses.

B. Initiatives and tools to improve users' awareness

Point B of the Commission's Information Request asks platforms to "provide data about implementation of their policies to inform users when they interact with disinformation" about COVID-19.

An important element of combatting disinformation is the ability to maintain the integrity and authenticity of news articles and other media as they are shared across the Internet. As discussed in our August report, Microsoft has partnered with the BBC, CBC/Radio-Canada, and The New York Times to help meet this need through "[Project Origin](#)," which helps users authenticate the sources of content and identify when it has been manipulated.

Another important effort in Microsoft's fight against COVID-19 disinformation, also discussed in our August report, is our long-standing partnership with [NewsGuard](#). NewsGuard operates a coronavirus misinformation tracker that lists all of the news and information sites in France, Italy, Germany, the UK, and the United States that it has identified as publishing materially false information about the virus. When users of Microsoft's Edge browser who have installed the NewsGuard plug-in navigate to these sites, a label appears warning the user that the information on the site is unreliable (further details on the NewsGuard coronavirus misinformation tracker is available [here](#)). NewsGuard has also [published fact checks](#) of common misinformation and misconceptions surrounding the coronavirus pandemic, which are available for free on its website.

Microsoft's remains a committed partner with NewsGuard in its fight against COVID-19 disinformation. As an example of NewsGuard's ongoing efforts in this area, on 24 August, NewsGuard [announced](#) an agreement with the World Health Organization to provide the WHO, and the various technology platforms, with a variety of reports and data aimed at helping them fight COVID-19 misinformation online. This is essentially the same information that NewsGuard has been providing to Microsoft as part of our global partnership announced in August 2018. As a testament to the effectiveness of its efforts, NewsGuard recently [won](#) a contest sponsored jointly by the U.S. State Department and Pentagon for detecting COVID-19 misinformation and disinformation. Microsoft has sponsored the inclusion of NewsGuard in a number of public libraries in the EU, including the Cologne Public Library in Germany and the library system of Milan in Italy.

In addition to these cross-company efforts, Microsoft services also have taken steps to alert users when they might be interacting with disinformation.

1. Bing

As discussed in more detail in our August report, another important way in which Bing helps users determine the reliability of content displayed in Bing search results is through its "[Fact Check](#)" feature, which gives users additional information with which to judge the trustworthiness of information online. Bing determines whether an article might contain fact checks by looking for the schema.org ClaimReview markup. In addition to this markup, Bing also looks for sites that follow commonly accepted criteria for fact checks, including of third-party fact checking organizations.

2. Microsoft Advertising

As described in response to Information Request 1, Microsoft Advertising takes action against COVID-19 disinformation through our policies that prohibit advertisements that can "reasonably [be] perceived as being deceptive, fraudulent or harmful to site visitors." When we learn that an advertisement on our system contains misinformation about COVID-19, we remove it (rather than leave it in place and merely alert users about the misinformation it contains). We provide more information on our practices in response to Point D, below.

C. Manipulative behavior

Point C of the Commission's Information Request asks platforms to "report all instances of social media manipulation and malign influence operations or coordinated inauthentic behavior detected on their services."

Because this section of the Commission's Request expressly seeks information on "*social media* manipulation and malign influence operations," our response below focuses primarily on actions LinkedIn has taken to identify and remove manipulative behavior on its services, including those that seek to exploit the COVID-19 pandemic. First, however, we wish to highlight again recent cross-company actions that Microsoft has taken to combat inauthentic and malicious actions in this area.

For example (and as discussed in more detail in our August report), Microsoft recently [initiated legal action](#) against web domains used in COVID-19 related cybercrime attacks. The cybercriminals at issue used COVID-19-related lures in their phishing emails to Microsoft customers in 62 countries in order to gain access to these customers' email, contact lists, sensitive documents, and other valuable information. Although 80% of victims compromised by this attack were US-based, some of our European customers were also targeted. Microsoft utilized technical means to block the criminals' activity and disable the malicious application used in the attack. In July, the court unsealed documents showing that our case resulted in a court order allowing Microsoft to seize control of key domains in the criminals' infrastructure, so that these domains can no longer be used to execute cyberattacks.

In addition, we are using our AI and human intelligence capabilities to stop attacks designed to exploit COVID-19. For instance, we been putting an emphasis on protecting critical services, especially hospitals, against [ransomware](#) attacks, and earlier this year extended our AccountGuard email protection service to healthcare workers, human rights organizations and humanitarian organizations in order to offer additional protections to those who are on the frontlines of responding to the pandemic.

We have also [published](#) extensive information and data, derived from our threat intelligence efforts, on how cybercriminals are seeking to capitalize on the pandemic and how their behaviors and strategies have evolved over the course of the pandemic. Because so many of these efforts involve the dissemination of disinformation or other deceitful practices, publishing this information can help users of all kinds more effectively identify and protect themselves against these harms.

In our August report, we provided information on the steps that Bing and Microsoft Advertising have taken to combat inauthentic behaviors on their services. In the interests of brevity, we do not repeat that information here. Instead, the balance of this response focuses on recent efforts by LinkedIn to address this problem.

Given the professional focus of LinkedIn and its members, it is rarely the target of social media manipulation, malign influence operations, or coordinated inauthentic behavior in relation to COVID-19. This might also be attributable in part to the fact that LinkedIn strictly enforces the LinkedIn User Agreement, which prohibits the use of “bots or other automated methods to access the Services, add or download contacts, send or redirect messages.” In addition, LinkedIn has systems and automated defenses in place for different types of content that may violate our User Agreement or Professional Community Policies. Our efforts in this regard can broadly be grouped into three categories: (1) Fake accounts; (2) content uploaded by members; and (3) advertisements.

1. Fake accounts

LinkedIn employs a “funnel” of defenses to detect and take down fake accounts. Our first line of defense involves the use of machine-learning models to detect groups of accounts that look or act similarly, which can imply they may have been created by a single bad actor. Any activity that registers a certain risk score is subject to a separate set of security measures designed to verify that the accounts belong to real individuals. For those accounts that meet this requirement (and have not been created in bulk by a single actor), we then run them through our activity-based models, which focus on types of behavior that are often associated with attempts to open a fake account.

These defenses enable us to take down most fake accounts before they go live on our platform. However, in the rare event that a bad actor is successful in bypassing our automated defenses, we also have a team of investigators who proactively undertake regular quality assurance controls to detect any account that might have bypassed our automated defenses. These investigations can also yield valuable insights that can be incorporated into the various defense models. In addition, members always have the possibility to report any suspicious activity to us, which we investigate promptly.

2. Content uploaded by members

To help ensure compliance with our policies, LinkedIn passes all content uploaded by members through our Universal Content Filtering (UCF) system at the time of creation / uploading. UCF performs a series of proactive checks—through machine learning models, classifiers, hash matching, third-party service lookups, and rule-based detection methods—to detect possible violations of our policies. This includes checking for URL reputation, adult images, commercial scams, profanity, malware, and other types of objectionable content. In addition, we proactively perform velocity and pattern-match checks to detect abusive behavioral patterns. We impose these measures on all content formats, including text, images, video, and links. Our enforcement engine leverages both information about member reputation as well as signals from content classification to make enforcement decisions. If content is assessed to violate our terms, it is filtered from the view of all members except the author. Depending on the severity of the violation, the author account may also be subject to an enforcement action (e.g., account restriction).

3. Advertisements

As discussed in our August report, ads do not run on LinkedIn until they have gone through our ad review process. This involves a combination of automated scanning and/or manual review to check ads against LinkedIn's [Advertising Policies](#).

D. Data on flows of advertising linked to COVID-19 disinformation

Point D.1 of the Commission's Information Request asks platforms to "provide data on policies undertaken to limit advertising placements on their own services" relating to COVID-19 disinformation. Point D.2 asks platforms to "provide information on policies to limit advertising placements on third-party websites using disinformation around COVID-19 to attract advertising revenue."

As discussed in our August report, all Microsoft services that display advertising have adopted and vigorously enforce policies prohibiting disinformation. In addition, Microsoft Advertising works both with advertisers and online publishers, and employs a distinct set of policies and enforcement measures with respect to each, to prevent the spread of disinformation through the advertising that it serves.

1. Advertisers

As described above, Microsoft Advertising's [Misleading Content policy](#) prohibits advertising that can "reasonably [be] perceived as being deceptive, fraudulent or harmful to site visitors." On the basis of this policy, we have prohibited all advertising that seeks to exploit the COVID-19 crisis for commercial gain, spreads misinformation, or that may pose a danger to user health or safety. Moreover, and as explained in our August report, every ad loaded into the Microsoft Advertising system is subject to enforcement methods that leverage machine-learning techniques, automated screening, in-house expertise, and dedicated user safety experts.

Microsoft Advertising conducts a manual review of all advertisements flagged to its customer support team and removes advertisements that violate its policies.

As noted above, as a result of these systems and related efforts, Microsoft Advertising prevented a total of 3,871,425 advertiser submissions directly related to COVID-19 from being displayed to users globally in August 2020. This included 1,165,481 ads designated for display to users located in European markets.

2. Online publishers

Microsoft Advertising utilizes a distinct set of policies and measures to combat the display of advertising on—and thus disrupt the flow of advertising revenue to—sites involved in spreading disinformation. These policies and measures apply to all types of disinformation, including disinformation relating to COVID-19.

Specifically, Microsoft Advertising requires our publishing partners to abide by strict brand safety-oriented policies to avoid providing revenue streams to websites engaging in misleading, deceptive, harmful, or insensitive behaviors. As discussed in our August report, these policies include a comprehensive list of prohibited content that our ads cannot serve against, including disinformation. We require publishers to maintain a list of prohibited terms and provide us with information on their content management practices where applicable. In addition, we require publishers to abide by restrictions against engaging in business practices that are harmful to users.

Microsoft Advertising reviews publisher properties and domains for compliance with these restrictions and promptly notifies publishers of properties or domains that violate Microsoft Advertising's policies. Microsoft does not approve properties that violate our policies for live ad traffic; if a property or domain that violates our policies is already live, we remove it from our ad network until the publisher remedies the issue. We also give advertisers the option to block their ads from being displayed on particular web domains.