

Response to the CMA's call for information: algorithms, competition and consumer harm

16 MARCH 2021

FACEBOOK

Introduction

Facebook welcomes the opportunity to respond to the United Kingdom's Competition & Markets Authority's (CMA's) open consultation on algorithms, competition, and consumer harm.¹ The consultation involves critical issues deserving of careful consideration by regulators, companies, civil society, and academics. These important issues require broad participation and engagement because the questions are complex, nuanced, and context-dependent. We see great value in all stakeholders working together to co-create policy approaches that unlock the tremendous benefits from algorithmic technologies while limiting their potential consumer harms.

At Facebook we work hard to ensure that all of our technologies work together to help achieve our mission of giving people the power to build community and bring the world closer together. As we describe in the first section of the response, algorithms are a key part of how we fulfil that mission. First, they help us personalize what people see, whether that's the posts in their feeds, the Pages or Groups they might like, the ads they are shown, or the people that they might want to friend or follow. Second, they help us support the millions of businesses that use our services to reach people and grow, by offering affordable and effective advertising solutions. And third, they help us keep our community safe, by helping to detect and remove harmful content.

Offering a personalised experience is a central component of the value our services provide; and algorithms play a key role in delivering that personalised experience. We believe strongly that personalisation brings tremendous benefits. But we also agree that it is important for users to understand that personalization and how it is delivered. As described in the second section of this response, that is why we have created industry-leading tools that give our users transparency and control around the way algorithms are used to improve their experiences.

Notwithstanding the tremendous benefits they offer, algorithmic systems also raise important policy and legal issues, including how to ensure that our systems are fair, transparent, accountable, and privacy-protecting. We agree these are important questions for the CMA to prioritize. At the moment, however, there exist more questions than answers. In the third part of the response, we set out some of the challenges that we see, and the way forward to address these. One of the significant challenges is the lack of clearly defined best practices, practical standards, and technical frameworks for measuring and assessing algorithmic systems like AI, particularly in terms of their potential risks to individuals and society. Without those standards and best practices, the investigatory techniques described by the CMA are premature, risk being ineffective, and may be disproportionate to the harms they are intended to address. Nonetheless, we believe that by working with the CMA and other stakeholders, we can collaborate in shared efforts to establish meaningful standards and answer the hard questions about how best to govern algorithms in the 21st century.

Benefits of algorithmically enhanced services

Facebook users report again and again that they derive substantial value from the ways in which Facebook helps them connect and share with their friends, families and wider communities. An MIT economic study found that social networking tools like Facebook are worth ~\$48/month to the average user, and that online video services like Facebook Live and YouTube are worth ~\$100/month to the average user.² These services are valuable because they provide positive and engaging experiences for people. And many of those experiences are made possible through algorithmic tools.

Algorithms improve the user experience

Facebook's mission is to give people the power to build community and bring the world closer together. Helping build community is fundamentally about connecting people to the other users, businesses, causes, and information that they care about. To deliver on its mission, Facebook offers people a personalised experience where content, including ads, is tailored to an individual's interests. Algorithmic tools are essential to delivering this personalized experience.

People benefit from personalised content

Algorithmic tools are an essential part of how Facebook ensures that every user sees the content that matters to them. No two people have the same set of friends or set of Groups and Pages that they follow. Facebook by definition is personalised; for each person on Facebook, there is a different Facebook. It is that personalisation that allows us to show each user important updates and stories from their friends and families and do that in unique and meaningful ways for each of our more than 2 billion users around the world. And for each of those people, at any given moment there could be more than a thousand posts that could potentially appear in that person's feed. Our ability to deliver a positive, engaging, and delightful experience for people depends on Facebook's ability to show the handful of posts most meaningful and important to each user.³

This is possible through the use of algorithmic tools, including the News Feed ranking system, which consists of multiple layers of machine learning models and rankings that we apply in order to predict the content that's most relevant and meaningful for each user. At each stage in the process, the News Feed ranking system narrows down those thousands of candidate posts to the small subset that appear in someone's News Feed at any given time.

Put simply, the system determines which posts show up in someone's News Feed, and in what order, by predicting what content they're most likely to be interested in or meaningfully engage with. These predictions are based on a variety of factors, including what and whom they've followed, liked, or engaged with recently.⁴ All of the algorithmic tools that make up our ranking system are built around our primary objective: creating the most long-term value for people by showing them content that is meaningful and relevant to them.

In addition to using algorithmic systems to determine the content people see in News Feed, we also make personalized recommendations using machine learning and algorithms to help users discover new communities and content they may be interested in. Facebook may recommend content, accounts, and entities (such as Pages, Groups, or Events) that users do not already follow. Some examples of our recommendation experiences include Pages You May Like, “Suggested For You” posts in News Feed, and People You May Know.⁵ Our goal is to make recommendations that are relevant and valuable to each person who sees them, so that they can connect with the other users, communities, causes businesses, and information they care about.

Although algorithms are the essential tools that we use to create the most long-term value for our users, the process is ultimately driven by the experience of our users. Measuring the experience of our users is not a mysterious, black-box process. Facebook has hundreds of researchers that conduct surveys, focus groups, usability studies, one-on-one interviews, and more in order to learn what kinds of content users find meaningful.⁶ And we continuously feed the results of the research back into our systems to ensure that the News Feed reflects what our users enjoy and find meaningful. Algorithmic systems play a key role in personalizing content to deliver positive, beneficial experiences for the people who use our services.

Promoting meaningful interactions

As described above, algorithmic systems play an important role in highlighting and elevating the content that a user is most likely to find meaningful and engaging. But engaging content that drives long-term value for our users is not the same as addictive or sensational content. We design our algorithmic systems to prioritize long-term user enjoyment. For example, in 2018 we changed our ranking for News Feed to prioritize meaningful social interactions.⁷ The change led to a decrease of 50 million hours a day worth of time spent on Facebook. Similarly, although political content only makes up about 6% of what people see on Facebook,⁸ we know even a small percentage of political content can impact someone’s overall experience. That’s why starting in February 2021, we launched a multi-month study to better understand people’s varied preferences for political content including temporarily reducing the distribution of political content in News Feed for a small percentage of people in Canada, Brazil, Indonesia, and the US.⁹ We’re always trying to make News Feed better, and the results of these studies will help us improve for users around the world the balance of content people want to see.

Sensational content may be an easy way to manufacture low-quality, short-term engagement. But it undermines Facebook’s mission to bring our users high-quality content that they find meaningful and relevant. Here again, algorithmic systems make it possible for Facebook to try to identify things like political content and other content that is not as meaningful and relevant to our users in order to remove or reduce its prevalence in users’ News Feeds.

People benefit from personalised ads

In the same way that we use algorithmic systems to personalise the content that people see to show them the content that we think will be most interesting and relevant to them, we also use algorithmic systems to show them the ads that we think will be most interesting and relevant to them. Online advertising is what allows us to offer Facebook platforms for free to our users.¹⁰ As the Department for Digital, Culture, Media & Sport said, “Online advertising can also help consumers to discover valuable new goods, interests and services, and is creating more accessible and low cost routes for businesses to engage with customers.”¹¹ In fact, recent research from the UK’s Centre for Data Ethics and Innovation showed that most users have positive experiences with personalised content and ads.¹² We believe that good online advertising should be just as enjoyable and meaningful for our users as the rest of the content that they see on Facebook.

At their best, personalised ads deliver new, relevant information to consumers at the right time. People consistently tell us that if they’re going to see ads, they want them to be relevant. That is why Facebook determines which ads to show based on an auction model whose goal is to show people ads that will be additive to their experience.¹³ Whenever we are about to show an ad to a user, we use algorithmic systems to calculate a total value for that ad, which includes a measure of the quality of the ad for that particular user and the likelihood that the user will take the action that the business cares about (e.g., having people watch a video, install an app, or purchase a product). Ads with the highest bid don’t always win the auction. Ads with lower bids often win if our system predicts a person is more likely to respond to them or finds that they’re higher quality. This allows businesses of all sizes to compete in the auction and reach customers in their target audience on any budget.¹⁴ Without personalisation, we would not be able to pick the ads most meaningful and valuable to a user, and their overall experience on Facebook would be worse for it.¹⁵

Algorithmic systems drive value for small businesses

Just as Facebook uses algorithmic systems to improve the user experience, we also use algorithmic systems to create value for small businesses. As the UK Government’s Department for Digital, Culture, Media & Sport noted in 2020, “[o]nline advertising makes an important growing contribution to our economy, with the UK market now the third largest in the world. In 2018, it accounted for over half of UK advertising spend, contributing £13.4 billion to the economy—up from £3.5 billion in 2008. This reflects the significant shift in consumer consumption habits from traditional media, such as newspapers and TV, to online formats.”¹⁶ Algorithmic systems make it possible to unlock this value for businesses across the UK and the world.

Personalised online advertising, made possible by algorithmic systems, has created enormous value for small businesses. Prior to the internet, reaching customers through traditional channels was enormously expensive and inefficient, favouring dominant multinational corporations that could spend billions on research and development, market analysis, advertising, and shelf space.¹⁷ The internet, however, made it possible for

small businesses to find customers at a fraction of the cost. A study from the Progressive Policy Institute concluded that “the shift from print to digital advertising is being driven in large part by the relative (low) price of digital advertising. We calculate, based on several assumptions, that for every \$3 that an advertiser spends on digital advertising, they would have to spend \$5 on print advertising to get the same impact. In the economic sense, digital advertising is more productive than print advertising. The benefits of these lower prices flow directly to advertisers and consumers.”¹⁸

This democratisation of advertising is particularly important in the UK. According to analysis from the UK think tank Credos, “[t]here are an estimated 5.7 million SMEs in the UK, constituting 99.9% of all private sector businesses in the country.... These companies now employ 16.3 million people, or 60% of all private sector employees.”¹⁹ Even just a few years ago, effective advertising was simply not an option for large numbers of businesses in the UK, either because it was too expensive (for example, a commercial to air on prime-time TV) and/or too inefficient (for example, newspaper adverts which would likely only be relevant to a fraction of a newspaper’s readers). Personalised advertising makes it possible for SMEs to compete. Copenhagen Economics surveyed 7,000 businesses in 15 countries across Europe, including the UK and found that 58% find Facebook’s apps important in lowering their marketing costs.²⁰ We are proud of how we have democratised advertising; we have over 10 million active advertisers on Facebook. Personalised advertising through Facebook has helped SMEs find customers around the world benefiting both these businesses and their consumers.

Personalised advertising is cheaper and more efficient than TV or print, because rather than showing an ad to everyone watching a show or reading a paper, we can show an ad to only the users likely to find it relevant and meaningful. Matching users to the ads most relevant to them, however, is a complicated and difficult technological challenge.²¹ Millions of small businesses around the world rely on Facebook because we have invested in the technology that makes it possible for them to reach new customers efficiently and effectively. Ignoring the important pro-competitive benefits of Facebook’s personalised advertising offering to small businesses would be a short-sighted public policy that serves to protect incumbents from challenge by new entrants. Degrading personalised advertising will tend to benefit large brands and those other platforms, but not consumers or competition.

Algorithmic systems make Facebook platforms safer

The CMA’s call for information recognizes the important role that algorithmic systems play in making online spaces safer for everyone: “Businesses can also use algorithmic systems to mitigate online harms and prevent discrimination against particular groups [and] to provide consumers with better quality (and less harmful) products.”²² In particular, the CMA acknowledged that algorithmic systems help Facebook identify and remove harmful content in ways that would simply be impossible without such tools, citing how in “November 2019, Facebook announced that it was able to remove seven million instances of hate speech in the third quarter of 2019, 59 percent higher than the previous quarter. Most of this is being detected by algorithms.”²³

Algorithmic tools used to help detect and remove harmful content are necessary because of the volume of content posted to Facebook. Although such tools are not perfect, as our Community Standards Enforcement Report shows, our technology to detect violating content is improving and playing a larger role in content review. While there are still areas where it's critical for people to review (for example, discerning if someone is the target of bullying can be extremely nuanced and contextual; some AI systems rely on large amounts of training data from human review to develop them) algorithmic systems are a key part of a larger approach to keeping our communities safe. In the paragraphs which follow, we provide three key examples of how algorithms help us to do just that.

Identifying and removing child exploitation content

Algorithmic systems are a critical tool in our fight against child sexual exploitation, both to help us prioritize the most serious reports and to proactively find content and remove it. We have been using photo-matching technology since 2011 to thwart the sharing of known child sexual imagery on our platform, suspend violating accounts, and report apparent violations as required by law. We also use artificial intelligence and machine learning to proactively detect child nudity and previously unknown child exploitative content. The accuracy of these tools has enabled us to find problematic content faster than ever before.²⁴ We're also using technology to find accounts that engage in potentially inappropriate interactions with children on Facebook so that we can remove them and prevent additional harm. In August 2019, we announced we are open-sourcing two technologies that detect identical and nearly identical photos and videos, so our industry partners, smaller developers and non-profits can use them to more easily identify abusive content and share hashes — or digital fingerprints — of different types of harmful content.²⁵

Identifying and removing misinformation about COVID-19 and vaccines

Algorithmic systems are crucial tools to help prevent the spread of health-related misinformation, because it allows us to proactively detect, prioritize and in some cases automatically remove false content that could lead to imminent physical harm. Beginning in January 2020, we began applying to false claims about COVID-19 our 2018 policy of removing misinformation that can contribute to imminent physical harm. Under this policy we removed posts with false information about the existence or severity of the virus, how to cure or prevent the disease, how it is transmitted or who is immune, the access to or availability of essential health services, and discouraging of good health practices.

Since March 2020, these systems have helped us remove over 12 million pieces of harmful misinformation related to COVID-19 from Facebook and Instagram.²⁶ We're also using algorithmic systems to help us leverage and scale the work of our independent fact-checkers, who we work with to debunk false claims on our services that may not directly result in physical harm. In that same period of time, we've applied warning labels on about 167 million pieces of content related to COVID-19 on Facebook, based on thousands of articles by our independent fact-checking partners.²⁷

In addition to detecting misinformation, our algorithmic systems are helping us with other challenges related to the pandemic. For example, we have built new computer vision classifiers to help enforce our temporary ban of ads and commerce listings for medical face masks and other products. By using things like local feature-based instance matching we can more effectively find manipulated media at scale and take action proactively—before anyone has even flagged it to us. Using these algorithmic systems, we were able to remove more than 2.5 million pieces of content for the sale of masks, hand sanitizers, surface disinfecting wipes and COVID-19 test kits in just the first few months of the pandemic.²⁸

Identifying and removing increasingly complex harmful content

Keeping our platforms safe from increasingly complex and rapidly evolving forms of harmful content is important for ensuring that our users have positive and engaging experiences. This requires significant ongoing investment in algorithmic systems that work alongside our content reviewers. Two complex challenges are worth briefly mentioning: (1) hate speech and other contextually dependent harmful content; and (2) deepfakes and other synthetic media.

Detecting harmful content like hate speech requires complex human and algorithmic systems that can be flexible enough to try to take into account context, intent, and other factors that ultimately determine what content must come down. Our systems—both automated and human—are built and intended to detect and enforce against violations of our applicable policies. However, although our algorithmic systems are making a lot of progress, they are not able to detect all possible policy violations. For example, when review requires understanding of the context surrounding the content at issue, automated systems are particularly challenged. Indeed, there are many signals and characteristics of a piece of content that may, in isolation or on their face, appear benign. However, when placed in context, the content may be indicative of a policy violation, or vice versa. Similarly, our Community Standards do not, and are not intended to, reflect any particular local legal regime. This makes it challenging to detect things like personal rights violations which require additional information (such as the truth or falsity of allegedly defamatory remarks) which we often do not have and, as an online intermediary, cannot be expected to have.

Given these challenges, our measures to tackle policy-violating content—in particular our automated measures—are not, and cannot be, perfect. We, therefore, continue to work to improve our systems (both human and automated) and rely on stakeholders to report such content using our various reporting channels.

We also are working to protect our platforms from synthetic media like deepfakes that advance misinformation through realistic, but fake videos. For example, in collaboration with the Partnership on AI, Microsoft, and academics from University of Oxford, Cornell Tech, MIT, and UC Berkeley amongst others, Facebook designed and rolled out the Deepfake Detection Challenge (DFDC) aimed at the development of machine learning

models that everyone can use to better detect when AI has been used to alter a video in a misleading way.²⁹ To support that research, we even went so far as to create a brand new data set of deepfakes based on videos we collected from paid actors who agreed to having their images manipulated for this work in order to help researchers build algorithms to effectively detect malicious deepfakes in the future. Things like deepfakes are complex and challenging threats to consumers around the world. Algorithmic systems are critical tools for defending citizens and consumers and ensuring the platforms like Facebook remain safe and enjoyable spaces for our users.

Empowering people through transparency and control

The preceding sections have made it clear that algorithmic systems can bring tremendous benefits to people not only through making online experiences more meaningful, relevant, and engaging, but also through keeping our communities safe. Algorithmic systems are the tools we use to deliver positive experiences to our users and make contributions to the global good³⁰—and personalisation is a key part of delivering on those goals. The CMA in its call for information described how “[p]ersonalisation can be harmful because it is difficult to detect either by consumers or others, targets vulnerable consumers or has unfair distributive effects. These harms often occur through the manipulation of consumer choices, *without the awareness of the consumer*.”³¹ We agree that awareness and understanding of personalization is key and at Facebook we believe in empowering our users through transparency, control, and choice. That is why we have created industry-leading tools that give our users information about the data we are collecting and how our algorithms use that data to deliver them meaningful, relevant content. We then give our users a range of controls over how that information is collected and used.

Transparency around algorithms and personalisation

The CMA in its call for information worries that “[c]onsumers may have limited visibility or knowledge of the inferences the firms may draw about them and the categories in which they are placed.”³² At Facebook we agree that people should understand the inferences and categorizations that shape the content they see online. But just as importantly, we believe that transparency only empowers users when it is understandable. That is why we work to provide transparency to users in context through user-friendly tools like “Why Am I Seeing This?”, that helps users better understand and control the content that they see.

At the same time, there is a lot of policy debate around algorithmic transparency, and the regulatory response needs to be balanced in order to empower and not overwhelm the user. Facebook believes it is important that users know the principles behind our recommendation systems. We work to provide users with transparency into why we are making these recommendations in ways that are clear, understandable, and actionable. But developing clear, understandable, and actionable transparency is also challenging in practice. That is why Facebook has initiated and supported Trust, Transparency and Control Labs (TTC Labs)—a cross-industry effort dedicated to creating innovative design

solutions that put people in control of their privacy.³³ We believe that developing meaningful transparency tools requires thoughtful design, user experience research, and interdisciplinary collaboration. The insights that come out of TTC Labs help ensure that Facebook is giving our users meaningful transparency and not just transparency for transparency's sake.

“Why Am I Seeing This?” Tools

Facebook has built a range of “Why Am I Seeing This?” features to help users better understand some of the factors that influence the content that they see. Initially, “Why am I seeing this ad?” highlighted one or two of the most relevant reasons that an ad was targeted to a user, such as demographic information or that the user visited the advertiser's website. In 2019 we expanded the information available in this to show some of the interests or categories that an advertiser selected in targeting an ad. Right next to this information we highlight controls available to the user to adjust their experience.³⁴ We have evolved this tool to make it detailed, understandable, and actionable for our users.

In 2019 we also extended this technology with “Why am I seeing this post?” to help users better understand how algorithmic content ranking works in News Feed. “Why am I seeing this post?” provides just-in-time context to users, with the tool available in the controls for News Feed posts. When a user clicks on “Why am I seeing this post?” they get a clear explanation for how their past interactions impact the ranking of posts in their News Feed. They are able to see:

- Why they are seeing a certain post—for example, if the post is from a friend, a Group they joined, or a Page they followed.
- What information generally has the largest influence over the order of posts, including: (a) how often they interact with posts from people, Pages or Groups; (b) how often they interact with a specific type of post, for example, videos, photos or links; and (c) the popularity of the posts shared by the people, Pages and Groups they follow.
- Shortcuts to controls, such as See First, Unfollow, News Feed Preferences and Privacy Shortcuts, to help shape what they see.

Meaningful transparency requires clear, understandable, and actionable information. The “Why Am I Seeing This?” tools provide users information in context, specific to the relevant post or ad. And based on that information, users can make changes to their settings and preferences with shortcuts that make it easy to manage what users see in News Feed right from this feature.³⁵

Ads transparency

At Facebook we believe transparency helps ensure that our systems—like those for ads—are fair for everyone. That is why we created the Ad Library, which provides advertising transparency by offering a comprehensive, searchable collection of all ads currently running from across Facebook apps and services, including Instagram.³⁶ Where the “Why am I seeing this ad?” tool provides actionable, user-level transparency into some of the reasons an individual sees a particular ad, the Ad Library allows users and stakeholders to have a broader view into advertising on Facebook and Instagram, beyond what they might encounter in their own personalised experiences. This wider lens into advertising on our services is one mechanism through which users and other stakeholders can hold us accountable for fairness in our personalised ads system.

We’ve added features to the Ad Library for particularly sensitive kinds of ads (e.g., housing, employment, credit, and politics) for certain countries. For example, in over 90 countries and territories, including the UK,³⁷ we require authorizations and proactively detect or reactively review possible social issues, electoral, or political ads.³⁸ Our Ads Library contains seven years’ worth of political and social issue ads, which are identified on Facebook by the “Paid for by” disclaimer. Displayed with the ad in context, this disclaimer is followed by information about who paid for the ad. The Ads Library includes detailed information about these ads, including the total spent by particular Pages on ads about social issues, elections, or politics. We also make an API available to third parties so that they can independently search and analyse these ads. We believe that through powerful transparency tools like the Ad Library, we can help ensure that all users are treated fairly.

Finally, in a few specific areas we simply block the use of certain targeting options available to advertisers. For example, for any advertisement for housing, employment or credit, advertisers are not permitted to target based on gender, age, or any interests that could be viewed as relating to people of a certain race, religion, ethnicity, sexual orientation, disability status, or other protected class. We also prevent advertisers from targeting ads based on narrow location options, including US ZIP code or from using lookalike targeting.³⁹ We use a mix of automated and human review in order to ensure that advertisers are not circumventing these rules.

Algorithmic transparency for businesses

Transparency is not only important for our users but also for the businesses that rely on Facebook’s algorithmic systems as a sometimes-substantial driver of their customers and revenue. That is why we provide a variety of tools to advertisers to help them better understand the relationship between their specific ads and our algorithmic systems. In particular, we offer three measures that collectively can provide important insight to advertisers:

- Quality Ranking: Estimate of whether a particular ad is likely to be interesting to users as compared to ads competing for the same audience.
- Engagement Rate Ranking: Estimate of a particular ad's expected engagement rate as compared to ads competing for the same audience.
- Conversion Rate Ranking: Estimate of a particular ad's expected conversion rate as compared to ads with the same optimization goal competing for the same audience.⁴⁰

The ad quality ranking, for example, is measured through various signals, such as feedback from people viewing or hiding the ad and assessments of low-quality attributes in the ad.⁴¹ Collectively these tools help advertisers understand whether the ads they are running are relevant to the audience they reached through our systems. They can use this data to make changes to their creative assets, post-click experience, or audience targeting so that they get better results through our algorithmic systems. We want users to have the best experience on our apps and services, and that includes providing transparency that helps businesses create and deliver the most high-quality and relevant ads to the right user at the right time.

Control around algorithms and personalisation

In addition to transparency around algorithms and personalisation, we also provide meaningful controls to empower users. To that end, we provide multiple ways for users to adjust the ads they see and how their data is used to personalise those ads. In fact, we provide more ways for users to exercise meaningful control than exist for other kinds of advertising like TV and print.

We provide to users a number of built-in controls⁴² to adjust the kinds of ads and content they see and what information influences those ads and content. These include:

- **Ad Preferences:** In Ad Preferences, users can see (i) their “interests,” which are keywords associated with a user based on actions such as liking “Pages” and clicking ads; and (ii) the advertisers that are currently showing them ads based on information provided by the advertiser, including contact information the user may have previously provided to the advertisers or previous use of the advertiser’s website or app. The Ad Preferences page allows users to update and edit these interests, influence what information Facebook uses to show them ads and adjust their general ad settings.
- **Content Controls:** Users can choose to hide or see less of particular posts, advertisements or content from particular people or advertisers. This goes beyond what is possible on many other ad-supported services. As noted above, whether users choose to hide particular ads informs our ad quality ranking. Users can also choose to hide some recommendations, including suggestions from People You May Know, which informs future recommendations.

- **“Data about your activity from partners” control:** Users can choose whether Facebook is permitted to use data it receives from advertisers and other partners about activity on non-Facebook service products to display personalised ads to them.
- **Off-Facebook Activity (OFA):** OFA is a control that gives people the ability to see a summary of activity that businesses and organisations have shared with Facebook about the user’s interactions off-Facebook through tools like the Facebook Pixel, Conversions API, App Events, and Offline Conversions (such as visiting that business or organisation’s apps or websites, viewing content, adding an item to a shopping cart) and to disconnect this activity from their Facebook account. Facebook began rolling out OFA worldwide (previously known as “Clear History”) in the second half of 2019. As part of the testing for OFA, Facebook conducted months of research to get input from stakeholders, including privacy advocates, policymakers, advertisers and industry groups, and we made important changes in response to what we learned.
- **Political Ads Control:** Last year we launched a control that enables people to turn off all social issue, electoral, or political ads from candidates or other organizations that have the “Paid for by” political disclaimer on them. Users can do this on Facebook or Instagram directly from any political or social issue ad or through each platform’s ad settings. This control is now available in more than 90 countries, including the UK.⁴³

Our controls extend beyond advertising, as we think consumers should have a choice in the content they see. That’s why we also provide controls for the News Feed itself. We recently made it easier to access alternative ways to browse your News Feed and tailor it to your liking, including “Favourites,” which is a way to choose the friends, family, and Pages whose content you’d like to see more frequently;⁴⁴ and Most Recent, a way to see all News Feed posts from the people and Pages you follow, sorted with the newest posts first, without our ranking algorithm.⁴⁵

Collaborating on the Future of Algorithmic Governance

Facebook CEO Mark Zuckerberg has called for a “more active role for governments and regulators” so that “by updating the rules for the internet, we can preserve what’s best about it – the freedom for people to express themselves and for entrepreneurs to build new things—while also protecting society from broader harms.”⁴⁶ Notwithstanding the tremendous benefits they offer, algorithmic systems raise important policy and legal challenges around how to ensure algorithmic systems are fair, transparent, accountable, and privacy-protective. We agree these are important questions for the CMA to prioritize. That said, we believe the investigatory techniques described by the CMA are premature, risk being ineffective and may be disproportionate to the harms they are intended to address. Before auditing and investigating algorithmic systems, there need to be clearly

defined best practices, practical standards, and technical frameworks for measuring and assessing algorithmic systems, particularly in terms of their potential risks to individuals and society. At the moment, however, there exist more questions than answers. The CMA should prioritize addressing these conceptual gaps in close collaboration with stakeholders that can collectively advance the knowledge around how algorithmic systems impact consumers. Facebook is eager to work with all stakeholders to collaborate in shared efforts to establish meaningful standards and answer the hard questions about how best to govern algorithms in the 21st century.

Creating best practices, standards, and technical frameworks

While there are many collaborative efforts underway, because this is such a new field, clearly defined best practices, practical standards, and technical frameworks for measuring and assessing algorithmic systems are still lacking. There exist established mechanisms for assessing the environmental or human rights impacts of technologies which can be drawn on, but more work is needed to usefully or effectively apply these approaches to algorithmic systems.

All stakeholders in this space—companies big and small, together with policymakers and civil society—need to collaborate on how to make algorithmic assessment work at scale, based on clear and reasonable standards and with at least the initial steps automated wherever possible, so that we can effectively protect people from unanticipated consequences without diminishing the many societal benefits of quickly developing and deploying new technologies. Similar collaboration is needed to develop basic, practical standards around other aspects of algorithmic fairness, accountability, and transparency.

For example, ensuring fairness in machine learning models is much more challenging than coming up with the right mathematical equation or tweaking a statistical formula — it’s a process that involves working hard to understand what it means for a product or system to work well for all users while fairly trading off between values that may be in tension. What “fairness” means will often be unclear and contested, and fairness goals for a given system will need to turn on the context of the particular product at issue or the particular communities it impacts. There is no one-size-fits-all approach to addressing algorithmic fairness concerns, pointing to a distinct need for collaborative work to develop more tailored approaches that can address specific types of fairness concerns for specific sectors, use cases, and user populations.

Supporting responsible algorithmic systems through engagement and practice

We believe that building responsible algorithmic systems will require collaboration with stakeholders across industry, government, and civil society. That is why Facebook has been engaging with regulatory efforts through responses to calls for information, such as this one from the CMA. We’ve also participated alongside experts from the United Kingdom and many EU Member States in the expert group⁴⁷ that helped develop the OECD’s AI Principles,⁴⁸ and we are now working through OECD’s AI Policy Observatory⁴⁹

to help put those principles into practice. Similarly, our global Facebook AI Research (FAIR) team has contributed to, and advanced, both fundamental AI research and broader discussion of issues like explainability through academic publication and open-source AI frameworks, tools, libraries, and models. One such library is Captum, a model interpretability library for PyTorch that makes state-of-the-art algorithms for interpretability readily accessible, so researchers and developers can better understand their models' predictions.⁵⁰

Even in the absence of clear standards and best practice, Facebook is taking a proactive approach by surfacing important questions about the responsible use of algorithmic systems early and often in the building of products and services. To drive this work, we have created a dedicated, multidisciplinary Responsible AI (RAI) team of ethicists, social and political scientists, policy experts, AI researchers and engineers focused on understanding concerns associated with the deployment of AI in Facebook products. That team's overall goal is to develop guidelines, tools and processes to tackle issues of AI responsibility and help ensure these resources are widely available across the entire company to encourage a systematic approach to these hard questions. One such tool is Fairness Flow, a technical toolkit developed in consultation with external experts that enables our teams to analyse how certain AI models and labels perform across different groups. We are continuing to improve and expand Fairness Flow, scaling our tooling and building infrastructure to support measurement for additional, more complex model types as well as enabling long-term, recurring automated measurement across models.

Facebook, however, cannot address such challenges on its own. In areas where clear standards and best practices do not yet exist, regulatory support of policy prototyping, regulatory sandboxes and other flexible approaches to policy development will help identify reasonable, evidence-based approaches to minimizing harms and maximizing benefit.

Tensions and challenges in the proposed approach to investigation and remediation

The CMA's report, although acknowledging the important role that algorithms can play, focused almost exclusively on mitigating risks, not enhancing benefits. Effective regulatory engagement will require a focus not only on the potential harms but also on the benefits of the algorithmic systems in question, and a recognition that *not* using algorithms will also create risks in certain contexts. We have described throughout this response how algorithmic systems can bring crucial benefits and mitigate existing harms. Accordingly, the CMA's ultimate approach should reflect a sensible balancing of the risks and the benefits of such systems.

We are glad that the CMA recognized that "harms from an algorithmic service are best understood with the full and proper context" of how a system was designed, developed, and deployed.⁵¹ However, many of the techniques for investigation that the CMA proposes present practical challenges, create new risks that may be disproportionate to

the harms they are intended to address, and fail to fully reflect the context in which the systems are developed and deployed. We highlight a few of these below:

Proposed Investigatory Technique	Practical Challenges
<p>Access to “the full and proper context of how such a service or automated system was designed, developed and trained, the data used as inputs, and how consumers respond to and make use of the output of the system”⁵²</p>	<p>There are emerging approaches to documenting machine learning systems, such as “model cards”,⁵³ that can potentially provide a snapshot of how a system was developed and how it performs. These emerging approaches show real promise, but there are still no clear standards for such documentation efforts, nor have they yet been demonstrated to be practical at a meaningful scale.</p> <p>Before requiring algorithmic audits, we need to address important gaps in standards and expectations. For example, we need to develop approaches that scale with algorithmic systems that may learn, adapt, and operate dynamically as users respond. We also need security protocols to ensure that sharing details of how algorithms operate does not enable adversarial attacks that allow bad actors to more effectively circumvent detection mechanisms or otherwise exploit algorithms in ways that could cause harm to others.</p>
<p>Access to “relevant data from the algorithm owner, such as inputs and outputs from live deployment or A/B testing”⁵⁴</p>	<p>In addition to raising the same security concerns mentioned above, providing inputs, outputs, and A/B testing results when this data would not otherwise be retained would create serious tensions with data protection regulations like the GDPR as well as with companies’ data minimisation and data retention policies, which generally err on the side of reducing data retention and giving people control over their data when it is stored. Sharing training data and real-world inputs would require additional data capture and retention and could create substantial new risks to the privacy of consumers’ data.</p> <p>Furthermore, a requirement to retain such data could harm privacy in another way—by unintentionally prohibiting privacy-preserving approaches like federated learning, where the data used to train an algorithm is not held in a</p>

	<p>centralised location but is instead distributed across multiple devices such as a user’s smartphone.</p> <p>Finally, it is rarely the case that training data sets are static. Given the growing number of applications of AI, the size and dynamic nature of the relevant data sets, and the increasing sophistication and complexity of training procedures, it may be technically impossible for organisations to keep all versions and iterations of training data sets ever used in designing and developing their AI systems.</p>
<p>Access to “the organisation’s documentation, pseudo-code and general explanations [including] communications and internal documents about the business context, objectives, design, architectural diagrams, training (including relevant function(s) that has been maximised during an algorithm’s training stage), key performance indicators (KPIs), and monitoring of algorithmic systems.”⁵⁵</p>	<p>Sharing details of sensitive commercial products, whether under this provision or those mentioned above, will also create substantial risks to valuable trade secrets if not permitted to be kept confidential during and after an investigation.</p>
<p>“[O]rder firms to make certain changes to the design and operation of key algorithmic systems”</p>	<p>There may be subsets of critical tasks that are only accomplishable using certain types of algorithms. It may be particularly important to prioritize model function in cases where integrity or safety are at risk, such as when enforcing our community guidelines around harmful content. Mandatory re-architecting of such systems could risk impeding such work, so serious caution is warranted prior to recommending this remedy.</p>

The risks highlighted above—including risks to privacy, trade secrets, and user security—are substantial, and when combined with a lack of clear or practicable standards in this area, further highlight the need for collaborative policymaking in this area to ensure an appropriately balanced approach. The systems at issue are profoundly complex, often the result of the combined effort of thousands of engineers around the world, and the question of how to provide meaningful external auditability while protecting against those risks is still unanswered. Quickly proceeding to require broad storage of or broad access to wide ranges of technical information as the CMA proposes risks creating substantial risks for companies and their users, while potentially leaving the CMA no better informed about the harms it seeks to address.

A multi-stakeholder path forward

We believe that there is great merit in multi-stakeholder collaborations to clarify practical expectations before other measures are sought. Facebook has demonstrated a commitment to working collaboratively with policymakers around the world to address the critical questions that arise when developing and deploying products and services that use algorithmic systems.

In addition to our work with the OECD, we have supported a diverse range of academic initiatives focused on AI and governance, such as our partnership with the Technical University of Munich to support the creation of an independent Institute for Ethics in AI,⁵⁶ as well as AI-focused academic projects in India, the Asia-Pacific region, and Latin America; and created AI for Good partnerships like our work with the Digital Ethics Lab of the University of Oxford's AIxSDGs project to explore how AI can help meet the United Nations' Sustainable Development Goals. We've also contributed to the Partnership on AI's multistakeholder ABOUT ML project to help identify best practices and approaches to AI documentation.

Alongside the initiatives above, we are also piloting innovative new approaches to regulatory sandboxes and policy prototyping. We recently launched Open Loop⁵⁷, a global initiative that promotes and deploys experimental regulatory efforts for algorithmic systems, with the express purpose of co-creating and testing new and existing governance frameworks through hands-on policy prototyping or sandbox activities. Open Loop programs are based on a consortium of diverse actors, from governments to industry, academia to civil society. Our first Open Loop report, based on findings from prototyping an AI Impact Assessment with 10 European companies, holds promising results.⁵⁸

Other Open Loop projects involve collaborations in the Asia-Pacific region and Latin America. For example, we worked with Singapore's Infocomms Media Development Authority (IMDA) and Personal Data Protection Commission in running an accelerator to help start-ups provide more transparency and control to users about data and develop new models for AI explainability that can support Singapore's AI Model Governance Framework.⁵⁹ Similarly, we are working with the Mexican National Institute for Transparency, Access to Information and Personal Data Protection (INAI) in Latin America in order to prototype policy approaches that would advance transparency and explainability in AI systems.⁶⁰ We are exploring other ways to similarly engage with policymakers on emerging algorithmic policy questions in the EU and around the world.

We believe these types of co-development processes are essential to understand and address the complex questions that algorithmic tools present, and we welcome similar public-private collaborations with the CMA and UK policymakers.

Conclusion

Algorithmic systems have delivered enormous benefits. As shown by the tensions and concerns highlighted above, however, they also pose novel and complex challenges, and deciding what an effective, proportionate, and technically feasible set of regulations for algorithms will not be easy. There are many difficult questions ahead of us, and we are eager to collaborate on how to answer them.

As noted above, Facebook is especially interested in the potential opportunity to participate in regulatory sandboxes and policy prototyping programmes to develop and test policy ideas around emerging technology like AI in collaboration with local regulators, as we are currently doing with Singapore. We welcome such collaborations with the CMA in order to experiment with and assess the impact of different evidence-based policy and regulatory approaches to algorithmic systems. In the meantime, we look forward to engaging in conversation with the CMA as we work together to ensure that algorithmic systems benefit all consumers and would be happy to participate in the roundtables planned for later this year.

Endnotes

- ¹ Competition & Markets Authority (2021) *Algorithms, competition and consumer harm: call for information*, GOV.UK. Available at: <https://www.gov.uk/government/consultations/algorithms-competition-and-consumer-harm-call-for-information> (Accessed: 17 February 2021).
- ² Brynjolfsson, E., Eggers, F. and Gannamaneni, A. (2018) ‘Measuring Welfare with Massive Online Choice Experiments: A Brief Introduction’, *AEA Papers and Proceedings*, 108, pp. 473–476. doi: [10.1257/pandp.20181035](https://doi.org/10.1257/pandp.20181035).
- ³ Lada, A., Wang, M. and Yan, T. (2021) *How does News Feed predict what you want to see?*, *Facebook Technology*. Available at: <https://tech.fb.com/news-feed-ranking/> (Accessed: 16 February 2021).
- ⁴ *Id.*
- ⁵ Facebook, *What are recommendations on Facebook?* / Facebook Help Center. Available at: https://www.facebook.com/help/1257205004624246/?helpref=search&query=recommendation_s&search_session_id=5d5bca7781cbb1308fa641cf4818cf2a&sr=0 (Accessed: 15 March 2021).
- ⁶ Facebook (2018) ‘How Users Help Shape Facebook’, *About Facebook*, 13 July. Available at: <https://about.fb.com/news/2018/07/how-users-help-shape-facebook/> (Accessed: 18 February 2021).
- ⁷ Mosseri, A. (2018) *News Feed FYI: Bringing People Closer Together*, Facebook for Business. Available at: <https://www.facebook.com/business/news/news-feed-fyi-bringing-people-closer-together> (Accessed: 2 March 2021).
- ⁸ Schultz, A. (2020) ‘What Do People Actually See on Facebook in the US?’, *About Facebook*, 10 November. Available at: <https://about.fb.com/news/2020/11/what-do-people-actually-see-on-facebook-in-the-us/> (Accessed: 18 February 2021).
- ⁹ Gupta, A. (2021) ‘Reducing Political Content in News Feed’, *About Facebook*, 10 February. Available at: <https://about.fb.com/news/2021/02/reducing-political-content-in-news-feed/> (Accessed: 17 February 2021).
- ¹⁰ Department for Digital, Culture, Media & Sport (2020) *Policy Paper: Online advertising - call for evidence*, GOV.UK. Available at: <https://www.gov.uk/government/publications/online-advertising-call-for-evidence/online-advertising-call-for-evidence> (Accessed: 16 February 2021) (noting “[o]nline advertising is at the heart of the digital economy. As the primary source of revenue for major online platforms, it underpins the provision of key online services”).
- ¹¹ Department for Digital, Culture, Media & Sport (2020) *Policy Paper: Online advertising - call for evidence*, GOV.UK. Available at: <https://www.gov.uk/government/publications/online-advertising-call-for-evidence/online-advertising-call-for-evidence> (Accessed: 16 February 2021).
- ¹² Centre for Data Ethics and Innovation (2020) *Online targeting: Final report and recommendations*, GOV.UK. Available at: <https://www.gov.uk/government/publications/cdei-review-of-online-targeting/online-targeting-final-report-and-recommendations> (Accessed: 16 February 2021) (finding that people did not want to stop personalization of content and advertisements, but instead wanted more meaningful transparency and control over those systems).
- ¹³ Metz, C. (2015) ‘Facebook Doesn’t Make as Much Money as It Could—On Purpose’, *Wired*, 21 September. Available at: <https://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose/> (Accessed: 23 February 2021).
- ¹⁴ Facebook for Business (2020) *How Does Facebook Use Machine Learning to Deliver Ads?*, Facebook for Business. Available at: <https://www.facebook.com/business/news/good-questions->

[real-answers-how-does-facebook-use-machine-learning-to-deliver-ads](#) (Accessed: 5 March 2021).

¹⁵ For more about how we use machine learning to deliver ads, see Facebook for Business (2020) *How Does Facebook Use Machine Learning to Deliver Ads?*, Facebook for Business. Available at: <https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads> (Accessed: 5 March 2021).

¹⁶ Department for Digital, Culture, Media & Sport (2020) Policy Paper: Online advertising – call for evidence, GOV.UK. Available at: <https://www.gov.uk/government/publications/online-advertising-call-for-evidence/online-advertising-call-for-evidence> (Accessed: 16 February 2021).

¹⁷ Thompson, B. (2020) *Email Addresses and Razor Blades*, Stratechery by Ben Thompson. Available at: <https://stratechery.com/2020/email-addresses-and-razor-blades/> (Accessed: 19 February 2021) (“Being good at advertising, whether it be coupons in the Sunday paper or television ads during the evening news, was mostly a matter of the ability to spend, which was itself a matter of scale.”).

¹⁸ Mandel, M. (2019) *The Declining Price of Advertising: Policy Implications – Progressive Policy Institute*. Available at: <https://www.progressivepolicy.org/issues/regulatory-reform/the-declining-price-of-advertising-policy-implications-2/> (Accessed: 16 February 2021).

¹⁹ Credos (2019) *Advertising Pays 7: UK Advertising’s Digital Revolution*. Available at: <https://www.adassoc.org.uk/wp-content/uploads/2019/06/Ad-Pays-7-UK-Advertisings-Digital-Revolution-compressed.pdf> (Accessed: 19 February 2021).

²⁰ Copenhagen Economics (2020) *Empowering the European Business Ecosystem – an impact study of business using Facebook apps and technologies*. Available at: https://www.copenhageneconomics.com/dyn/resources/Publication/publicationPDF/2/522/1579535391/empowering-the-european-business-ecosystem_copenhagen-economics.pdf (Accessed: 16 February 2021).

²¹ Thompson, B. (2020) *Email Addresses and Razor Blades*, Stratechery by Ben Thompson. Available at: <https://stratechery.com/2020/email-addresses-and-razor-blades/> (Accessed: 19 February 2021).

²² Competition & Markets Authority (2021) Algorithms, competition and consumer harm: call for information, GOV.UK. Available at: <https://www.gov.uk/government/consultations/algorithms-competition-and-consumer-harm-call-for-information> (Accessed: 17 February 2021).

²³ *Id.*

²⁴ Facebook, *Safety Center*, Safety Center. Available at: <https://www.facebook.com/safetyv2> (Accessed: 9 March 2021).

²⁵ Davis, A. (2020) ‘Facebook Joins Industry Effort to Fight Child Exploitation Online’, *About Facebook*, 11 June. Available at: <https://about.fb.com/news/2020/06/fighting-child-exploitation-online/> (Accessed: 9 March 2021).

²⁶ Jin, K.-X. (2021) ‘Reaching Billions of People With COVID-19 Vaccine Information’, *About Facebook*, 8 February. Available at: <https://about.fb.com/news/2021/02/reaching-billions-of-people-with-covid-19-vaccine-information/> (Accessed: 2 March 2021).

²⁷ *Id.*

²⁸ Facebook AI (2020) *Using AI to detect COVID-19 misinformation and exploitative content*. Available at: <https://ai.facebook.com/blog/using-ai-to-detect-covid-19-misinformation-and-exploitative-content/> (Accessed: 16 February 2021).

²⁹ Schroepfer, M. (2019) *Creating a dataset and a challenge for deepfakes*. Available at: <https://ai.facebook.com/blog/deepfake-detection-challenge/> (Accessed: 19 February 2021).

³⁰ As discussed above, we also use algorithms to prevent negative experiences. Rosen, G. (2018) ‘F8 2018: Using Technology to Remove the Bad Stuff Before It’s Even Reported’, *About*

Facebook, 2 May. Available at: <https://about.fb.com/news/2018/05/removing-content-using-ai/> (Accessed: 5 March 2021); see also Nuno, B. S.-A. and Krambeck, H., *Addressing COVID-19 through Public-Private Data Partnerships—Where Do We Put New Testing Facilities?* Available at: <https://datapartnership.org/updates/covid19-and-public-private-data-partnerships/> (Accessed: 19 February 2021).

³¹ Competition & Markets Authority (2021) Algorithms, competition and consumer harm: call for information, GOV.UK. Available at: <https://www.gov.uk/government/consultations/algorithms-competition-and-consumer-harm-call-for-information> (Accessed: 17 February 2021) (emphasis added).

³² *Id.*

³³ TTC Labs, *About TTC Labs*. Available at: <https://www.ttclabs.net/about> (Accessed: 20 February 2021).

³⁴ Thulasi, S. (2019) 'Understand Why You're Seeing Certain Ads and How You Can Adjust Your Ad Experience', *About Facebook*, 11 July. Available at: <https://about.fb.com/news/2019/07/understand-why-youre-seeing-ads/> (Accessed: 20 February 2021).

³⁵ Sethuraman, R. (2019) 'Why Am I Seeing This? We Have an Answer for You', *About Facebook*, 1 April. Available at: <https://about.fb.com/news/2019/03/why-am-i-seeing-this/> (Accessed: 16 February 2021).

³⁶ FACEBOOK, *Ad Library*. Available at: [https://www.facebook.com/ads/library/?active_status=all&ad_type=political_and_issue_ads&country=ALL&sort_data\[direction\]=desc&sort_data\[mode\]=relevancy_monthly_grouped](https://www.facebook.com/ads/library/?active_status=all&ad_type=political_and_issue_ads&country=ALL&sort_data[direction]=desc&sort_data[mode]=relevancy_monthly_grouped) (Accessed: 16 February 2021).

³⁷ Facebook for Business, *Ads about social issues, elections or politics in the UK*, *Facebook for Business*. Available at: <https://www.facebook.com/business/m/one-sheeters/ads-with-political-content-uk> (Accessed: 15 March 2021).

³⁸ Facebook for Business, *Availability for Ads About Social Issues, Elections or Politics*, *Facebook Business Help Center*. Available at: <https://www.facebook.com/business/help/adstransparencycountries> (Accessed: 15 March 2021).

³⁹ Murphy, L. W. (2020) *Facebook's Civil Rights Audit - Final Report*, at 72. Available at: <https://about.fb.com/wp-content/uploads/2020/07/Civil-Rights-Audit-Final-Report.pdf> (Accessed: 16 February 2021).

⁴⁰ Facebook, *About Ad Relevance Diagnostics*, *Facebook Business Help Center*. Available at: <https://www.facebook.com/business/help/403110480493160> (Accessed: 10 March 2021).

⁴¹ Facebook, *About Quality Ranking*, *Facebook Business Help Center*. Available at: <https://www.facebook.com/business/help/303639570334185> (Accessed: 10 March 2021).

⁴² Facebook, *Control the Ads You See | Facebook Help Center*. Available at: https://www.facebook.com/help/1075880512458213/?helpref=hc_fnav (Accessed: 16 February 2021).

⁴³ Gleit, N. (2020) 'Launching The Largest Voting Information Effort in US History,' *About Facebook*, 16 June. Available at: <https://about.fb.com/news/2020/06/voting-information-center/> (Accessed: 10 March 2021).

⁴⁴ Facebook, *How do I set up and manage Favorites on my Facebook News Feed? | Facebook Help Center*. Available at: https://www.facebook.com/help/1634545223376778?helpref=faq_content (Accessed: 20 February 2021).

⁴⁵ Facebook, *How can I view and adjust my Facebook News Feed preferences? | Facebook Help Center*. Available at: <https://www.facebook.com/help/371675846332829> (Accessed: 17 February 2021).

-
- ⁴⁶ Zuckerberg, M. (2019) 'Four Ideas to Regulate the Internet', *About Facebook*, 30 March. Available at: <https://about.fb.com/news/2019/03/four-ideas-regulate-internet/> (Accessed: 20 February 2021).
- ⁴⁷ OECD, *List of participants in the OECD Expert Group on AI (AIGO)*, OECD. Available at: https://www.oecd.org/going-digital/ai/oecd-aigo-membership-list.pdf?fbclid=IwAR3jUalKeK9O6EKu9Z5f2uEZ5zufXSiulpIKIS0b89im_4kPtMMeUwaS17c (Accessed: 17 February 2021).
- ⁴⁸ OECD, *OECD Principles on Artificial Intelligence - Organisation for Economic Co-operation and Development*. Available at: <https://www.oecd.org/going-digital/ai/principles/?fbclid=IwAR0xJph4cqLWUvazLpUoS8iHrVXzimXWOr9IS-P8HF6a4JFvktakSMcLw> (Accessed: 17 February 2021).
- ⁴⁹ OECD, *The OECD Artificial Intelligence Policy Observatory - OECD.AI*. Available at: https://oecd.ai/?fbclid=IwAR0_NFdHYf-kqL3zoU1s26HoVEHaJDUKJgkrhyy4w-g37FTa9oOl9f32nss (Accessed: 17 February 2021).
- ⁵⁰ Facebook AI (2019) *Open-sourcing Captum: A model interpretability library for PyTorch*. Available at: <https://ai.facebook.com/blog/open-sourcing-captum-a-model-interpretability-library-for-pytorch/> (Accessed: 16 March 2021).
- ⁵¹ Competition & Markets Authority (2021) Algorithms, competition and consumer harm: call for information, GOV.UK. Available at: <https://www.gov.uk/government/consultations/algorithms-competition-and-consumer-harm-call-for-information> (Accessed: 17 February 2021).
- ⁵² *Id.*
- ⁵³ Mitchell, M. *et al.* (2019) 'Model Cards for Model Reporting', *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 220–229. doi: [10.1145/3287560.3287596](https://doi.org/10.1145/3287560.3287596).
- ⁵⁴ *Id.*
- ⁵⁵ *Id.*
- ⁵⁶ Joaquin, Q. C. (2019) 'Facebook and the Technical University of Munich Announce New Independent TUM Institute for Ethics in Artificial Intelligence', *About Facebook*, 20 January. Available at: <https://about.fb.com/news/2019/01/tum-institute-for-ethics-in-ai/> (Accessed: 10 March 2021).
- ⁵⁷ Open Loop, *A Global Experimental Governance Program, Open Loop*. Available at: <https://openloop.org/> (Accessed: 16 March 2021).
- ⁵⁸ Open Loop, *Let's Unlock, Open Loop*. Available at: <https://openloop.org/lets-unlock/> (Accessed: 10 March 2021).
- ⁵⁹ Naidoo, T., *TTC Labs - Our Learnings: FB Accelerator Singapore, Season 2*. Available at: <https://www.ttclabs.net/insight/Our-Learnings-FB-Accelerator-Singapore-Season-2> (Accessed: 10 March 2021).
- ⁶⁰ Open Loop, *Let's Unlock, Open Loop*. Available at: <https://openloop.org/lets-unlock/> (Accessed: 10 March 2021).