

Algorithms: benefits, harms, and oversight Google's perspective

Google welcomes the opportunity to address the CMA's paper, '*Algorithms: How they can reduce competition and harm consumers*' (the **Algorithms Paper**). Our views, which are addressed in more detail below, are as follows:¹

Algorithms deliver enormous benefits to consumers and businesses. As the Algorithms Paper notes, algorithms can be used to save time, provide personalised recommendations, increase efficiency, and enhance product quality. Google has experience of using algorithms and automated systems to achieve these and a range of other benefits. In many contexts, algorithms are designed to promote consumers' long-term interests. Maintaining the long-term credibility of a search engine, for example, outweighs any short-term gains from exploiting users. In other contexts, though, service providers may use (or misuse) algorithms in ways that harm consumers and may be less concerned with preserving user trust.

Certain applications of algorithms are well-known to harm consumers, such as using algorithms to facilitate cartels, charge discriminatory prices on the basis of protected characteristics, and target vulnerable consumers with loyalty penalties. **Algorithm-based services that use manifestly harmful techniques are more likely to warrant scrutiny and intervention.** And the more significant an algorithm's impact on a person's livelihood, the greater the potential for harm. Algorithms that determine job offers, credit terms, access to housing, or exam scores are more impactful than algorithms for recommending music or films.

Of course, distinguishing beneficial, responsible applications of algorithms from harmful, exploitative services will not always be straightforward. But it is an important task to ensure that regulators can address harmful uses without undermining the benefits that algorithms deliver. The Algorithms Paper asks how, therefore, algorithms should be reviewed and audited. As the Algorithms Paper notes, firms are responsible for conducting impact assessments, monitoring, and evaluating how their algorithms work. Accordingly, **the results of firms' pre- and post-launch testing may help regulators focus their reviews efficiently and should be the starting point for analysis.**

Finally, the Algorithms Paper asks what measures are feasible, effective, and proportionate to remedy or prevent algorithmic harms from occurring.

- As regards transparency measures, a balance needs to be struck. On the one hand, website owners and business users can benefit from understanding the main parameters that ranking services use, as prescribed in the Platform-to-Business Regulation. On the other hand, **disclosing the precise signals that determine how rankings are calculated opens up algorithms to manipulation by bad actors.**

¹ See also Google's [written evidence](#) to the Science and Technology Committee's inquiry into the use of algorithms in decision-making (October 2017).

Disclosing user-level data risks disclosing the underlying algorithm itself, which in turn, reduces competition and dampens incentives to innovate.

- **Competition, consumer protection, and other laws can be enforced when algorithms are used unlawfully.** Examples of enforcement include the [Online Posters and Frames](#) and [Online Hotel Bookings](#) cases. Enforcement could be supplemented with participative frameworks to review and uphold standards of good conduct. For example, agencies could adopt mechanisms to let firms explain to them how particular algorithms work and seek agencies' views of those algorithms from a regulatory perspective. Reviews should start from a presumption of innocence and create opportunities for no-fault dialogue in what are likely to be complicated and fact-specific analyses.

These matters are addressed in detail in the following sections.

* * *

We hope our responses are helpful. Please let us know if you have any questions or would like to discuss these or other matters further.

A. Algorithms offer wide-ranging consumer benefits

The Algorithms Paper acknowledges the many and wide-ranging benefits that algorithms have for consumers, such as providing relevant recommendations, saving time, and allowing consumers to refocus their energies on what matters. Even when information is available in a hard copy encyclopedia, most people see the benefit in letting search algorithms do the work for them.² Algorithms have other benefits too that are not mentioned in the Algorithms Paper but that warrant discussion -- in further research and in the CMA's ongoing work:

- **Increased accuracy.** Techniques developed in one product area can be used in new, unpredictable ways, often to improve accuracy and performance. Machine learning and AI techniques that Google has developed, for example, have been deployed in medical diagnostics, leading to [significant improvements](#) in cancer screening, with fewer false positives and false negatives than human experts. Further research in this area is ongoing (see [here](#) and [here](#)).
- **Increased fairness.** With appropriate guardrails and thoughtful problem formulation, automated systems can base decisions on fair and consistent criteria, while excluding emotional appeal, subconscious prejudice, or inconsistent approaches that can characterise human decision making. Indeed, machine learning techniques can be used to rectify examples of human biases. For example, 2018 [updates to Google Translate](#) corrected for gender bias in language by offering both a feminine and masculine translation for a single word.
- **Improved safety and sustainability.** Algorithms can improve safety by delivering more reliable and higher quality performance. Unlike humans, algorithms do not lose concentration or get tired, and their greater processing power allows decisions to be taken more quickly and comprehensively. That's why self driving cars can be safer than human drivers. Algorithms can also find ways to make processes more sustainable. For example, Google's [Deepmind technology](#) has led to a 40% reduction in energy used for cooling and 15% reduction in overall energy overheads at certain Google data centres.
- **Improved productivity and efficiency.** Automated programs can improve the productivity and efficiency of human workers. They can save time on mundane processing tasks; help to prioritise the most pressing work; and boost the quality of decision making. For example the City of Memphis is using [Google Cloud's vision APIs](#) to review images of road surfaces from cameras on the front of buses to detect potholes that need repairing and more efficiently deploy work crews. This allowed the city to repair 63,000 potholes in one year -- a major improvement over

² For an indication of the value that consumers attach to certain algorithm-based services, see Brynjolfsson, Collis, and Eggers, [Using massive online choice experiments to measure changes in well-being](#) (2019) (estimating that users in 2017 would be willing to go without search engines for one year in return for approximately \$18,000).

previous manual efforts. Likewise, Google AI technology has been used to improve the process for [discovering new therapies](#) by enhancing ‘virtual screening’ methods that allow potential treatments to be evaluated computationally.

- **Levelling the playing field.** Automated technologies can also help to level the playing field by enabling smaller firms to carry out tasks that previously required extensive resources. For example, [Google’s Cloud AutoML](#) tool enables developers who have little experience with machine learning to train machine learning models (e.g., a meteorologist wanting to train a model to detect different types of clouds to assist with weather forecasting).

Of course, in some contexts, algorithm-based services -- as with other services -- have the potential to cause consumers harm. Distinguishing beneficial from harmful applications of algorithms will be a central challenge for the CMA’s work. As the following sections explain, much turns on whether the service provider has a strong long-run incentive to improve product quality and retain customers. The answer will depend on the precise nature of an algorithm and the manner of its implementation.

B. In many contexts, providers have an interest in improving quality -- not exploiting consumers

The types of services that Google offers -- including information-providing services, such as Google Search -- are long-standing products that depend on maintaining user trust for their long-term viability. If trust breaks down, consumers are less likely to try out new products, less likely to look for answers or recommendations, and less likely to make purchases or click on ads. Sacrificing product quality, showing less relevant results, or depriving users of control could lead -- at best -- to fewer clicks on ads, which means less revenue. At worst, consumers who don’t trust Google could abandon our services altogether.

The following examples illustrate how Google has designed algorithms -- and offered associated controls -- to improve product quality.

i. Ads load

In 2015, a team of Google researchers considered the trade-off between short-term and long-term effects of changes to Google Search results, [noting that](#) “*Optimizing which ads show based on short-term revenue is the obvious and easy thing to do, but may be detrimental in the long-term if user experience is negatively impacted.*” To address this problem, the researchers developed metrics for measuring long-term user satisfaction, paying particular attention to ‘ads blindness’ -- how a users’ propensity to click on ads changes based on the quality of the ads and the user experience.

Applying these metrics, Google “*drastically reduced the ad load on the mobile interface*”; the resulting short-term loss of revenue was offset by increased satisfaction. Google was thereby able to reduce ad load without sacrificing long-term revenue; a net positive change

for businesses and users. This study illustrates that, for a product like Google Search, there is a clear business rationale to focus on the long-run user interest rather than short-term revenue generation.

ii. Search ranking

Algorithms play a fundamental role in ensuring that Google Search users see high quality results. Retrieving a set of quick, high quality results without algorithms would be virtually impossible. Google ranking systems sort through hundreds of billions of webpages in the Search index to find relevant results in a fraction of a second, and present them in a way that helps the user to find what they are looking for. In doing so, Google Search algorithms consider factors such as the query itself, the relevance and usability of web pages, the expertise of sources, and the freshness of the content.

Four examples illustrate how Google updates its Search algorithms to increase its functionality for users, including -- as the Algorithms Paper puts it -- to save users time and “*make effective improvements based on empirical evidence*” (para. 1.4):

- **Autocomplete.** The autocomplete feature in Google Search predicts the query that users want to run and enables users to enter those queries without having to type them out in full; a particularly useful feature on mobile or other small screen devices. On average, autocomplete reduces typing time by approximately 25 percent and is estimated, cumulatively, to [save more than 200 years of typing time per day](#). Contrary to the suggestion that autocomplete functionality is vulnerable to manipulation through adversarial attacks, it deploys extensive measures to protect against bad suggestions, including algorithmic protections against query spam and offensive predictions. These protections have, like other aspects of Google’s search product, been continually improved over time.
- **Hummingbird.** In 2013, the Hummingbird update changed how Google focused on and understood users’ queries -- a critical prerequisite for any search service to deliver relevant and useful search results. Specifically, Google noticed that search terms were becoming more conversational -- even more so for voice searches. The update focused on understanding natural language and how people communicate, improving Google’s ability to understand the *meaning* behind users’ search queries rather than just the words in the query. For ‘conversational’ queries, which are usually longer and made up of a complete set of questions, Hummingbird allowed Google to break the question into its component parts, and provide results that are nearly identical to the original keyword search result. Hummingbird thereby increased the quality of Google Search results.³
- **BERT.** One of the most important improvements in the quality of Google Search over the past few years has come from Google’s work on [Bidirectional Encoder](#)

³ Third parties have recognised this -- see for example [here](#) from Search Engine Journal.

[Representations from Transformers \(BERT\)](#), which enables search services to understand in even greater detail the context of words in a search query and how they fit together, rather than looking at words in isolation (Google has released BERT material on an [open source basis](#)). Therefore, Google Search now recognises that the query [2019 brazil traveler to usa need a visa] refers to Brazilian visitors to the US, not the reverse. With this better understanding of language, Google can show more relevant results.

- **RankBrain** is an AI-based system that Google began using in 2016 to understand how website pages are related to concepts. It enables Google Search to better return relevant pages, even if they do not contain the exact words used in a search query, by understanding instead that the page is related to other words and concepts.

iii. Ads personalisation

Another important component of maintaining user trust of algorithm-based services is offering users control over their experience. For example, beyond Google's commitment to [protecting user privacy](#) and providing [industry leading security](#), Google provides users with powerful controls to allow them to choose how their data is used, including for advertising. In partnership with ad industry initiatives such as [aboutads.info](#) and [YourOnlineChoices](#), Google offers users opt-out controls and information for almost every ad they see, including:

- The possibility to [opt out of seeing personalised ads](#) through a simple on/off toggle.
- "[Mute This Ad](#)" which is available for a large number of display ads, and which enables the user to "X" out of the specific ad (and other ads using the same web URL). These ads are not shown to the user again.
- The recently updated "[Why This Ad?](#)" page, which is reached by clicking on the "AdChoices" icon in the corner of most ads. This tool provides users with information about why they are seeing a particular ad and provides a link to the user's privacy controls to update their personalisation settings to avoid seeing similar ads in future. This page will also soon begin showing users the verified name of the advertiser behind each ad, giving users even more transparency into the ads they see.
- Like the CMA, Google supports (and is taking) measures to guard against discriminatory ad targeting on the basis of protected characteristics. Our [personalised ads policy](#) prohibits ads targeting based on identity and belief, including but not limited to sexual orientation, race, ethnicity, religious beliefs, marginalised groups, or transgender identification.

The Algorithms Paper raises concerns about data being used "*in ways which consumers do not expect or have little control over.*" Enabling users to understand and control how their

data are used -- and which ads appear -- helps address this concern and guard against consumer harm.

C. Regulators should focus on applications of algorithms that are most likely to cause material consumer harm

Not all digital services are incentivised to maximise long-term user satisfaction. Some business models depend on one-shot interactions with consumers rather than looking for repeat business, so that short term revenue gains are prioritised. In such situations, some businesses may resort to exploitative, predatory techniques and use algorithms for this purpose. Two considerations can help regulators identify and focus on practices that are most likely to arise and most likely to cause harm. First, how impactful is the algorithm-based service? Second, is the algorithm used in a way that is well known to damage consumer welfare?

How impactful is the algorithm-based service? Certain applications of algorithms can have a much larger impact on users' lives and livelihoods than others -- either in positive or negative ways. For example, algorithms have been used to determine which candidates are offered jobs, on what terms borrowers will be offered credit, whether a particular individual will be given access to social housing, and, more recently, what exam grades students receive. These types of algorithms have a disproportionately greater impact on consumers and citizens than algorithms that recommend music, films, or places to eat. These more impactful algorithms should be the focus of regulators' attention.⁴

Is the algorithm used in a way that is known to harm consumers? Certain practices involving the use of algorithms are well-understood to damage consumer welfare. For example, the Algorithms Paper discusses services that target higher prices -- directly or indirectly -- to different ethnic groups; individualised perks that aim to maximise the amount users spend on addictive activities, such as gambling; increasing fees for insurance customers who are deemed unlikely to switch (including, potentially, vulnerable customers); gender and racial bias in rankings for workers who are listed on sharing economy platforms; scarcity messaging and pressure selling on hotel booking sites; and hotel rankings based on the commissions hotels pay (rather than their relevance to the user's search), without this fact being made clear to users. Algorithm-based services that use these techniques are more likely to warrant scrutiny and intervention.

D. A firms' own testing can help guide agencies' reviews

Of course, distinguishing clearly pro-consumer algorithms from those warranting further review, will not always be straightforward. But drawing this distinction correctly is important

⁴ By analogy, Article 22 of the GDPR establishes a right not to be subject to a decision based solely on automated processing. This right applies only if the data processing "produces legal effects" or "similarly significantly affects" the data subject. As the [ICO explains](#), "the decision must have a serious negative impact on an individual to be caught by this provision" such as "automatic refusal of an online credit application, and e-recruiting practices without human intervention."

to ensure that potentially harmful applications of algorithms can be addressed while pro-consumer algorithms are not inadvertently hindered.

The Algorithms Paper observes that “*firms are responsible for effective oversight of such systems, which should include... holistic impact assessments, monitoring and evaluation.*” Accordingly, in considering how to exercise their review functions, agencies could consider, among other factors, the results of firms’ own testing.

There are various types of evidence from testing that agencies might consider. By way of illustration, Google’s testing for its Search product includes pre-launch testing, long-term studies of user behaviour post-launch, and eye-tracking experiments.

i. Pre-launch testing

Google performs extensive user-focused testing of proposed changes to product design. Each time Google engineers have an idea that might improve the quality or relevance of search results, testing is carried out to verify that the contemplated changes do in fact lead to a better user experience. In 2019 Google [ran over 464,065 experiments, resulting in more than 3,620 improvements to Google Search](#). In general, launch testing includes the following steps:

- **Sandbox testing.** As a first step, Google tests the proposed change in an experimental ‘sandbox’. Google compares the pages with and without the proposed change and through multiple iterations to ensure the change is of the appropriate quality to proceed to the next stage of testing.
- **Rater testing.** If internal tests suggest that the proposed change is likely to improve user experience, [external ‘raters’](#) evaluate that change in blind side-by-side (‘SxS’) reviews. Raters are shown versions of Google’s search result pages generated with and without the change and are asked to rate each page, without knowing which version includes the change. They are asked to make their determinations based on [highly detailed rater guidelines](#). Google analysts then consolidate the feedback and determine whether the change does in fact improve the quality of search results for users. Google conducts each test with hundreds of individual raters and multiple different search result queries before proceeding to the next stage.
- **Live traffic experiments.** Only once the proposed change has passed rater testing is it run live. Google conducts [live traffic experiments](#) to see how users interact with a feature, before launching it to everyone. Google enables the feature in question to just a small percentage of people, usually starting at 0.1%, and looks at what people click on, how many queries were run, whether queries were abandoned, how long it took for people to click on a result, and so on. The results help measure whether engagement with the new feature is positive for the user, so as to ensure that the changes Google makes are increasing the relevance and usefulness of search results.

ii. Long-term studies of user behaviour post-launch

Google also continues to monitor and re-evaluate its results following launch. Examples of this practice are long-term user behaviour studies that Google conducted to assess the quality of its specialised product ads that it shows in Shopping Units. The long-term user studies in question analysed how users who experience Google's Shopping Unit react to it over time, compared to users who did not experience it. The studies find that users who experience the Unit return more often to Google to enter more queries and click more often on the Unit compared to users who did not previously experience the Unit. This provides strong evidence for the quality improvement that the Shopping Unit generates.

iii. Eye-tracking experiments

Google has also conducted eye-tracking experiments to assess whether adding images to specialised results improves quality and benefits users. In particular, Google has sought to understand the effect that displaying images has on the visual scanning of the search results, and whether displaying results with an image helps users find and choose between different results more easily.

One of the challenges of a general search service is that different users may be looking for different things with the same query. For example, a user entering the query [*Boris Johnson*] may be looking for news about the Prime Minister, videos, his biography, or his policies. Google may therefore show results for different kinds of information for a given query. Showing these different results with different formats, including, where appropriate, images, can make it easier for users to identify and select those results that are specifically relevant to them. Google tested this proposition with eye tracking studies.

For example, one study analysed the effect of including image thumbnails as part of specialised news and video results, with the objective of understanding whether the image thumbnails would (i) help users find the news or video result faster when the result is relevant to the query, or (ii) distract users when the news or video results are not relevant to the user's task.

In the study, participants were asked to complete a number of tasks by searching for information on Google. Participants were given a description of what information they needed to search for, and told to look for whatever the first result page said. The study found that "*the thumbnail image seemed to make results with thumbnails easy to notice when the users wanted them [...] and the thumbnails also seemed to make it easy for people to skip over the results with thumbnails when those results were not relevant to their search*". Google [publicly announced these findings on its blog](#).

In short, the eye-tracking studies demonstrated that the formats of Google's specialised results help users assess the relevance of results and select between them. Adding images does not automatically lead to increased clicks, but rather helps users select results and skip over a result if the result is not relevant for them.

E. Transparency and integrity need to be carefully balanced

There is often an assumption that transparency -- simply by making more information publicly available -- will lead to greater accountability and trust. By itself, however, transparency is no panacea. What is most appropriate will depend on many factors, including the audience and the wider technological and market context.

Google promotes transparency in a range of contexts, including in relation to algorithms. The [Model Cards](#) project, for example, enables people to understand how machine learning models work, including how well a particular model performs in some respects and what limitations it has in others. Likewise, in the context of search services and ads, Google recognises the benefits to website owners, businesses, and consumers from understanding how the underlying algorithms work:

- Transparency ensures that customers and webmasters benefit from understanding the criteria against which their products, services, or sites will be evaluated and ranked. Similarly, consumers benefit from the general parameters of search rankings being made public, including the key factors that help determine which results are returned for their queries. Specifically, Google publishes and maintains detailed [information](#) about how Google Search works, including information about how Google improves search quality and Google's approach to algorithmic [ranking](#), including publication of the [Search Quality Rater Guidelines](#) which define goals for Search algorithms.
- Transparency helps webmasters adapt to material changes in ranking or other issues. For example, when Google implemented the Speed Update, it [provided](#) webmasters with six months' notice of the change, giving them time to adapt. And Google has provided at least six months' notice of the [introduction](#) of the 'page experience' signal, which will further enhance users' search experience.
- Transparency helps publishers and advertisers understand the rules and processes of ads auctions. Google provides participants with explanations of the [key elements of Ad Manager's ad auction process](#) and the main parameters that influence it, such as [pricing](#) and [blocking](#) rules, Google's [relationship](#) with publishers and exchanges, and how [dynamic allocation](#) works.
- Transparency helps address questions about the fees charged when advertisers use Google's ad intermediation services. That is why Google recently published two blogs ([here](#) and [here](#)) showing that Ad Manager publishers keep over 69% of digital advertising revenues generated, and news publishers keep over 95% on average.
- Transparency ensures that consumers have access to clear information concerning which data are collected and how those data are used. Our [Privacy Policy](#) explains what data Google collects and why, Google explains [how data are used](#) in ads, and consumers are given the option to opt out of personalised advertising altogether.

At the same time, there are clear and well-established limits on how far certain types of transparency can go before they jeopardise the very services to which they relate.

i. Avoiding manipulation by bad actors

A careful balance is needed to ensure that ranking results are not manipulated by bad actors, which harms both legitimate businesses and consumers.

For example, while it may be helpful for a search service to provide guidance on the main parameters of a site that it takes into account in ranking, it would be prejudicial to the proper and safe operation of a search service to publish details of all the technical ‘proxy signals’ through which these parameters are assessed. Otherwise, websites could manipulate and improve their ranking in search results by optimising for the relevant proxy signal; not by increasing the quality or relevance of their site to users. This would have negative consequences for (i) consumers who will see more irrelevant or even harmful content, and (ii) genuine websites who play by the rules and will be pushed down in the search results to make room for websites that manipulate search signals.

A good example is the PageRank signal, which examines the number and quality of links that a website receives from other websites. It is unlikely that a user would notice the number of links that a website receives from other sites; however, if a website receives a lot of links from other websites, that may indicate that the website provides useful content. Google [published](#) the fact that it was using this signal as a proxy for relevance or quality. Because website operators knew that Google considered the number of incoming links as a signal, some websites engaged in practices to manipulate that signal, rather than genuinely improving their website. For example, they bought incoming links or engaged in link exchange schemes so that they appeared to Google’s algorithms to be of greater quality than they really were.

This serves to illustrate the importance of keeping proxy signals hidden. It is why the [Platform-to-Business Regulation](#) requires online platforms only to identify the “*main parameters*” that search services consider when ranking websites (Article 5(2)). It recognises in Recital 27 that the “*ability to act against bad faith manipulation of ranking by third parties, including in the interest of consumers, should [...] not be impaired.*”

ii. Avoiding collusive outcomes

A further objective is to avoid disclosures that displace innovation-driven competition and replace it with collusive outcomes.

The Algorithms Paper refers to the risk of collusive outcomes resulting from firms using “*the same algorithmic system to set prices, including by using the same software or services supplied by a third-party, or by delegating their pricing decisions to a common intermediary.*” In this regard, the Algorithms Paper is too narrow. Collusive outcomes can affect other parameters of competition too, such as quality and innovation (not only price) and can result from regulatory actions (not only market failure).

For example, there has been discussion about whether Google should disclose user-level click and query data to rival search engines as a means of helping those rivals compete. But evidence suggests that sharing this data would not enhance competition to find the best results⁵; rather, it would inform rivals as to how Google answers a particular query. It would therefore enable rivals to clone Google's search results in a systematic way, reducing product diversity and chilling incentives of Google and its rivals to invest in product improvements. In other words, disclosing click and query data would lead to collusive outcomes on the core parameter of competition in search; namely, competition to find the best answer to a user's query.

This risk is borne out in the [comments of one of Google's search rivals, Mojeek](#), which noted that this measure would force Google to:

“open up what is essentially their product and share it with others, or to open up search query and click data they have obtained by way of that product... If these steps are made in the name of positive competition, it will actually just result in multiple search engines all offering the same service but under different banners. And whilst it's important that metasearch engines like DuckDuckGo and Startpage exist to offer users better privacy than mainstream search engines, they are not offering any new innovation with regards to improving the core element of search... instead we call for more search engines with independent search indexes and algorithms.”

As the CMA's [Final Report](#) commented, *“there is a risk, if such a remedy included a requirement to disclose the outputs of proprietary search algorithms, which are the result*

⁵ One way to look at the question is what happens when rivals get more data. For example, the [Microsoft/Yahoo! deal](#) doubled Bing's query volume overnight but failed to improve the relevance or monetisation of Bing's search results. In fact, it was [publicly reported](#) that *“Yahoo's revenue per search has been worse under the Microsoft deal than when it operated its own Web-search technology and advertising system.”* In other words, having more data did not lead to an improvement in rivals' performance.

Another way to look at the question is what happens when Google gets less data? When breaking news events occur, an important question for search engines is how quickly they learn to associate related queries to results that discuss the breaking news event. Of course, in the immediate aftermath of the event, search engines cannot fall back on associations of prior queries and clicks to tell them what the 'right' or 'relevant' results are; after all, a breaking news event is -- by definition -- something new. Comparing how Google Search and rivals perform when confronted with searches related to breaking news events allows Google to test whether its superior quality comes from its greater access to query and click data, or because it is simply better at figuring out what results users are interested in.

When the Japanese Tsunami occurred on 11 March 2011 at 06:46 CET, Google was able to identify, promote, and surface relevant responses to the query “japan tsunami” by 07:10 CET; just 24 minutes later and well before Google had received a large volume of queries relating to the tsunami. In other words, Google was able to identify and show relevant results quickly, even without a large trove of query and click data. By contrast, 650 hours after the event, Bing was still showing irrelevant results, despite having (presumably) received a large number of Tsunami-related queries by that time (top Bing results for “japan tsunami” related to tsunamis from 1993, 2005, and 2006, as well as the Wikipedia page for tsunamis).

of investments in search and associated infrastructure, that this could dampen incentives for Google to innovate and improve its algorithm by enabling free riding” (para. 8.40).

Moreover, sharing such granular data could expose users to privacy violations, as borne out in both historical examples and [a 2019 paper in Nature](#) by an author of the EC Special Advisers’ Report on digital competition. Accordingly, transparency and disclosure requirements need to be carefully confined to avoid damaging algorithms’ integrity.

F. Clear principles can guide appropriate monitoring of algorithms

Regulators have a legitimate interest in understanding what impact algorithm-based services may have on competition and consumer welfare. While many -- and perhaps most -- uses of these types of systems benefit consumers and businesses, harmful or exploitative uses of algorithms undermine the reputation of algorithm-based services as a whole. In determining how regulators should carry out review functions, two issues emerge: (i) how to prioritise algorithms or services for review, and (ii) how to carry out those reviews.

i. Prioritisation principles

Ever more industries are becoming digitised, and as the CMA’s paper points out, algorithms are used in a wide range of contexts; not only by large digital platforms such as Amazon, Apple, Facebook, Google, and Microsoft, but also other firms in a variety of sectors, ranging from transportation (Uber) to freelancing (TaskRabbit/ Fiverr), and from stationery suppliers (Staples) to casinos (MGM) and hotel booking sites (Booking.com, Expedia, Hotels.com). Accordingly, criteria will be needed to prioritise areas for review, consistent with the CMA’s proposal to apply its [prioritisation principles](#) in a way that focuses on the risk or impact of harm. The following considerations appear most relevant to prioritising regulators’ work:

- ***Prioritise interventions based on the severity of consumer impact.*** As discussed above, some automated systems have a greater impact on consumers’ lives than others -- both the capacity to do good when they function well, and the capacity to harm when they entail predatory design. An algorithm that makes a poor song recommendation is less impactful than an algorithm that determines which job offers a user receives. Likewise, some applications of algorithms are well-established as harming consumer welfare (e.g., supporting cartel agreements with algorithms, or imposing loyalty penalties). Regulators should prioritise accordingly.
- ***Focus on the harms that appear most likely to transpire.*** The Algorithms Paper discusses a wide range of potential harms arising from the use of algorithms. Some of these concerns are accompanied by indications that they are transpiring in practice (e.g., pressure selling; undisclosed payments in ranking in purportedly organic search results; and setting service conditions, directly or indirectly, on the

basis of protected characteristics). Other concerns appear theoretical and lack indications of consumer harm (e.g., personalised rankings).⁶

- **Prioritise straightforward solutions.** In certain cases, interventions can correct algorithmic harm without wide-ranging unintended consequences. For example, ranking purportedly organic results on the basis of undisclosed payments to the platform provider can be resolved through a straightforward disclosure requirement. Other possible interventions could lead to substantial risks to innovation, competition and consumer welfare, such as requirements to disclose user-level click and query data. As discussed above, this type of disclosure would undermine incentives to invest in algorithm improvements, lead to a collusive outcome where search engines align their results rather than competing on quality, and jeopardise user privacy. Accordingly, it makes sense to prioritise straightforward solutions where they exist.
- **De-prioritise cases where there are safeguards against manipulation.** The Algorithms Paper discusses harms that occur through user manipulation and dark patterns that impede user choice. Algorithms and choice architecture are separate issues. Nonetheless, in circumstances where firms have put in place measures to protect user choice, it may be easier to exclude concerns that users suffer from a lack of options -- whether those concerns arise from the operation of an algorithm or something else. For example, choice screens and choice carousels (described at **Annex 1**) preserve consumer choice and guard against perceived default bias.
- **Focus on concerns that are directly related to algorithms.** Many of the illustrations of consumer harms that the Algorithms Paper mentions are core examples of algorithm-driven harms. An algorithm that decides the cost of a loan or the prices charged based on protected characteristics is one such example. Other issues, such as self-preferencing or choice architecture, are legitimate issues for discussion and debate (and were recently considered by the Digital Markets

⁶ The section of the Algorithms Paper on personalised rankings (pp.14-15) does not include tangible examples of personalised ranking outside e-commerce or of examples of consumers being harmed.

Taskforce). But they are separate from harms that might be caused by uses of algorithms themselves.⁷ They are, at most, tangentially related.

ii. Oversight principles

Algorithms exist in a fast-moving area of technological development, generating substantial consumer benefits. Oversight should therefore be evidence-based, proportionate, and targeted to ensure that it does not dampen incentives to innovate or slow down the deployment of new product features that benefit consumers. With these considerations in mind, we propose the following principles for carrying out oversight of algorithm-based services:

- **Deploy participative approaches.** Participative antitrust looks for opportunities for agencies and companies to address concerns -- or promote competition -- more quickly, effectively, and creatively than traditional enforcement or centralised rule-making (although enforcement of competition, consumer protection, or other areas of law remains as a possibility). Under a participative approach, agencies set the objectives or direction, and companies make proposals to achieve those objectives, making use of industry players' technical expertise.⁸ This model was discussed as an attractive option in the Digital Markets Taskforce's recent report⁹ and other reports recognise the value of participative approaches for

⁷ The Algorithms Paper refers to the European Commission's Google Shopping decision. In that case, the EC did **not** challenge Google's search algorithm design. The EC acknowledged, as the CMA's paper notes, that Google's algorithms "were designed to stop sites with poor quality content appearing in Google's top search results." The Shopping Decision finds that Google's demotion algorithms "improve the relevance of the generic search results on its general search results pages" (Decision, recital 16). The Shopping Decision states that the EC does not prevent Google from applying these algorithms (Decision, recital 661). Rather, the alleged abuse in Shopping concerned unequal access as between Google Shopping and rival comparison shopping services to an attractive design on Google's results page called the Shopping Unit. Whether that constitutes an abuse is a matter before the General Court. But this is separate from the question of whether a demotion algorithm in and of itself harms competition or excludes rivals, which was not part of the case. For further discussion on the case, see Graf and Mostyn, [Do We Need to Regulate Equal Treatment? The Google Shopping Case and the Implications of its Equal Treatment Principle for New Legislative Initiatives](#), *Journal of European Competition Law & Practice*, 2020.

⁸ See Quartz, [A Nobel-winning economist's guide to taming tech monopolies](#), 27 June 2018, citing Professor Tirole ("*Drawbacks of classical approaches are well-known: self-regulation tends to be self-serving; competition policy is often too slow; public utility regulation, as we discussed, is mostly infeasible (and it is sometimes captured). We must develop what I would call 'participative antitrust' in which the industry or other parties propose possible regulations and the antitrust authorities issue some opinion, creating some legal certainty without casting the rules in stone*").

⁹ CMA, [A new pro-competition regime for digital markets](#), Advice of the Digital Markets Taskforce (CMA 135), December 2020, Overview, para.14 (envisaging "a participative approach, whereby the DMU seeks to engage constructively with all affected parties to achieve fast and effective results").

innovation-intensive sectors.¹⁰ In the context of algorithms, a non-adversarial context might provide an option for firms voluntarily to discuss material changes to algorithms with agency officials (including the results of firms' own testing related to the change) and receive an indication of how the agency views the proposed changes from a regulatory perspective. The CMA's proposal to provide firms with guidance fits well within this framework.

- **Use proportionate evidence gathering.** Agencies have wide-ranging information gathering tools -- and where particular algorithms are suspected of causing harm, it is reasonable for regulators to use those tools. At the same time, information gathering should be conducted proportionately and in a way that respects trade secrets. For example, the Algorithms Paper recognises the “*significant value in investigating automated systems without direct access to the underlying code*”; the principle that audited companies should not be subjected to “*undue burdens*”; and the opportunity to ask companies to provide “*general explanations*” of how their algorithms work, as well as observing the outputs of those algorithms.
- **Ensure regulatory coherence.** Oversight of algorithms cannot be treated in isolation; many of the issues addressed in the Algorithms Paper may fall within the scope of other pre-existing legislation and regulatory initiatives, such as the GDPR, Platform-to-Business Regulation, existing competition legislation, as well as contemplated legislation like the CMA's proposed Digital Codes of Conduct. Regulatory coherence militates against layering on new rules that overlap with recently implemented (or not-yet implemented) regulations that address the same issues. Before deciding whether to intervene, the prior question is whether there is a regulatory gap at all. Likewise, regulatory coherence requires that agencies with overlapping interests in these matters take consistent and coordinated approaches. The [Digital Regulation Cooperation Forum](#) has been set up to support cooperation and coordination on online regulatory matters, to enable coherent, informed and responsive regulation of the UK digital economy to serve both citizens and consumers alike. And the UK Regulators Network can further support consistency among agencies.

¹⁰ World Economic Forum, Agile Regulation for the Fourth Industrial Revolution, December 2020, p.20 (“Clear and timely regulatory advice is vital for innovators who are developing new ideas, products and business models. Where businesses face uncertainty about whether their ideas will be considered compliant with regulation, they are less likely to be able to persuade potential investors or consumers of the merits of their innovation – and less likely to innovate... many governments have introduced advice services for innovators to help reduce uncertainty about the regulatory implications of their ideas”).

Annex 1

Choice Screens and Choice Carousels

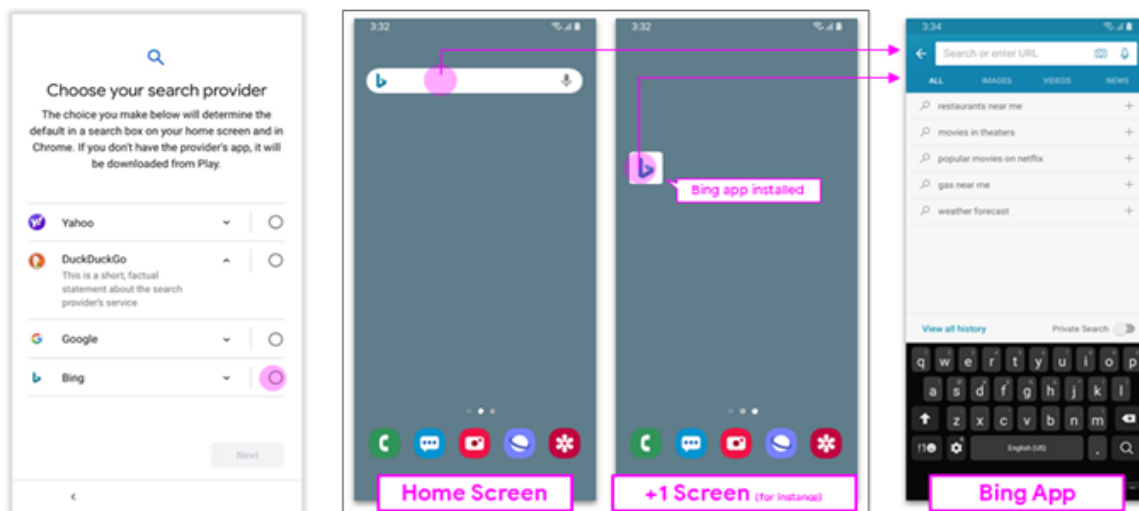
Consumer choice is essential to the functioning of a free market economy. When consumers shop around and switch, they get the best service possible at the best price. Businesses are rewarded or punished by consumers depending on the value of the service they provide. Offering tools such as choice screens and choice carousels could provide a suitable basis for addressing concerns about algorithms limiting the choice available to consumers.

A. Choice screens

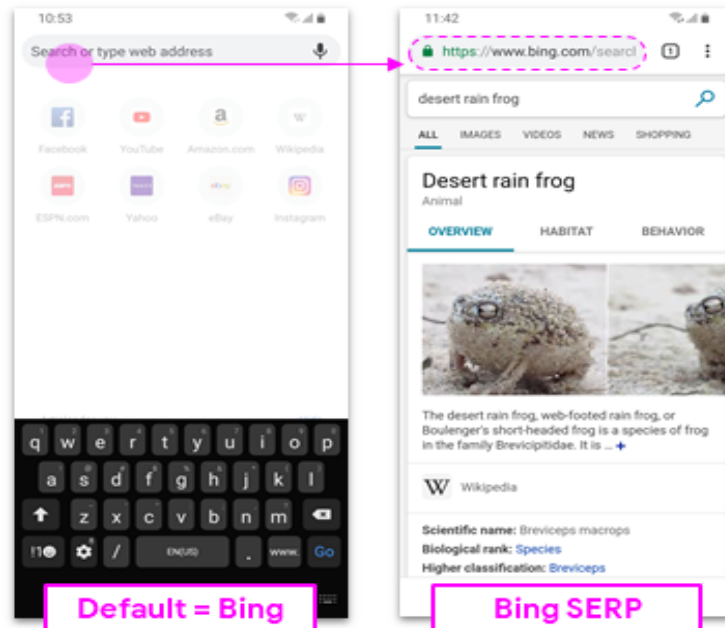
Choice screens can be an effective way to address concerns about behavioural biases identified by the CMA in the Algorithms Paper, including ‘default bias’. Furthermore they are proven, effective, can be straightforward to implement, and require little ongoing monitoring or enforcement.

In March 2020, following a consultation with the European Commission, Google introduced a choice screen for Android mobile devices shipped to the EEA. This choice screen has further expanded opportunities for consumers to make active choices about the services they use. When a mobile device is set up for the first time, the consumer is given a choice of search providers, each showing a description of its service. Selecting a search service from the choice screen causes the device to (i) configure the home screen search box to the selected service, (ii) install the search app of the selected provider (if not already installed), and (iii) set the selected search service as the default in Chrome. The effect of a user selecting Bing is illustrated below.

Choice screen configures the device to the user’s selection



Searches in the Chrome URL bar are directed to user-selected service



At present, choice screens are shown only on certain platforms, including Google’s Android operating system. However, there have been calls to adopt a more systematic approach. The Digital Markets Taskforce has recognised choice screens as a “*long-term solution to delivering consumer choice*”,¹¹ and the CMA has stated that “[c]hoice screens can help improve consumers’ access to alternative search engines”.¹² The recent Penrose report noted that “*services such as ‘choice screens’ will reduce the power of default settings, so customers can switch (for example) to a different internet browser from the preset one on their laptop, tablet or phone simply and easily*” (p.31). And the CMA’s Algorithms Paper recognises that “*well-designed choice architecture including default options and rankings can help consumers make decisions efficiently. If there is sufficient competition, informed and active consumers can switch to other platforms if they are unsatisfied with the results of one platform*”.¹³

B. Choice carousels

In addition to choice screens, other tools have also emerged to facilitate and further enhance consumer choice. Google is implementing a series of ‘choice carousels’ in search verticals to present users with additional options. Alongside the shopping vertical with which the European Commission’s Google Shopping decision was concerned, Google has been working on ‘carousels’ that provide links to alternative vertical search providers alongside Google’s own specialised results boxes for jobs, local amenities, flights, and

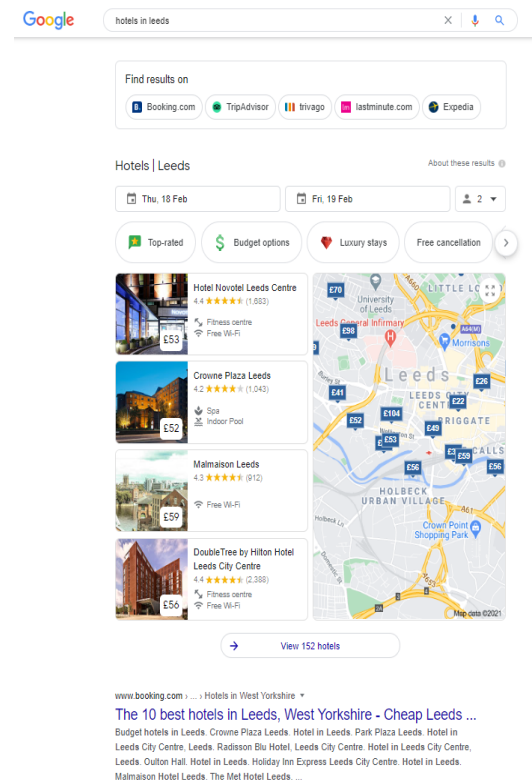
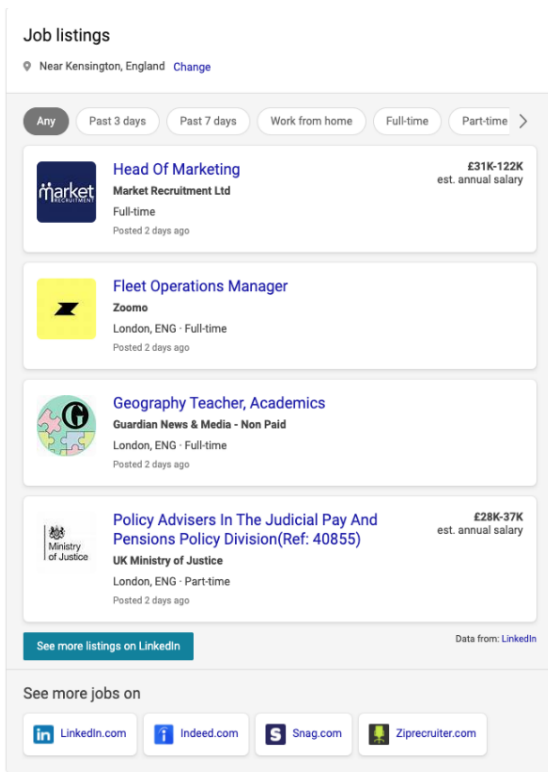
¹¹ Digital Markets Taskforce, A new pro-competition regime for digital markets, Advice of the Digital Markets Taskforce, Appendix D, para. 10.

¹² CMA, Final Report, Market Study into Digital Advertising and Online Platforms, Appendix V, para. 56.

¹³ CMA Algorithms Paper (para 2.45).

hotels. As shown in the example below, a search for [hotels in Leeds] takes users to a carousel -- above Google's own set of specialised results -- with links to Booking.com, TripAdvisor, Trivago, Lastminute.com and Expedia.

Likewise, other search services -- such as Microsoft's Bing -- appear to be linking to third party jobs sites when users search for job listings. Alongside choice screens, these carousel-type solutions merit consideration as a tool to help consumers make proactive choices and address concern about firms exploiting ranking effects (or 'self-prefencing') identified by the CMA in the Algorithms Paper.¹⁴



¹⁴ CMA Algorithms Paper, “Firms can exploit default effects and ranking effects by placing options that are more profitable in prominent positions. This can be done in a way that is not transparent or understood and accepted by consumers, and potentially at the expense of the consumer if he or she would have chosen a superior alternative under a more neutral presentation of options. Where the favoured options belong to the same entity controlling the platform, this is a form of self-prefencing” (para 2.51).