Reset.

Algorithms, competition and consumer harm
March 2021

About Reset

Reset (www.reset.tech) was launched in March 2020 by Luminate in partnership with the Sandler Foundation. Reset seeks to improve the way in which digital information markets are governed, regulated and ultimately how they serve the public. We will do this through new public policy across a variety of areas – including data privacy, competition, elections, content moderation, security, taxation and education.

To achieve our mission, we make contracts and grants to accelerate activity in countries where specific opportunities for change arise. We hope to develop and support a network of partners that will inform the public and advocate for policy change. We are already working with a wide variety of organizations in government, philanthropy, civil society, industry and academia.

Foreword

Reset welcomes the work of the CMA in this important area. We have covered the problems posed by algorithms and options for regulatory oversight for a while and so are pleased to see progress in this area. Our recent paper on algorithmic audit, jointly penned with the Ada Lovelace Institute, can be downloaded [here](#).

Algorithms are at the heart of the digital market and so far have been much overlooked by regulators. In part, as the CMA's paper states, this is due to lack of empirical evidence about the harms to both consumers and markets. Of course as digital platforms share so little information about their data, systems and user experiences, this doesn't mean the harms aren't there - it's just incredibly difficult to identify them. However, as the evidence base grows, and as we understand quite how much data is behind closed doors, there is a compelling case for regulators to take active measures to counter harms before they arise. We are strongly supportive of the CMA's efforts and hope it is granted the appropriate tools and powers to deliver a robust regime.

While the CMA's focus is understandably on consumer protection, many of the harms explored in the paper have societal repercussions. Class, race and gender can be factors in determining what services and prices are made available to consumers. In addition, recommendation algorithms optimise for user attention at the expense of information quality. This, as the paper states, "lowers consumer welfare". How to differentiate the consumer from the citizen is a

difficult question, but one which regulators must consider collectively - perhaps as part of the Digital Regulation Cooperation Forum (DRCF). Many of the problems and solutions set out by the CMA apply to citizens and society just as much as to consumers and marketplaces. We hope other regulators, particularly Ofcom, benefit from the learnings in this exercise and can be awarded some of the powers, namely algorithmic inspection and audit, that the DMU is calling for.

Answers

1.   Are the potential harms set out in the review paper the right ones to focus on for our algorithms programme? Are there others that we have not covered that deserve attention?

As per our foreword, while we understand that the CMA's focus is on consumer detriment, where the consumer stops and the citizen begins is unclear. Harms witnessed as a result of market distortions may breach consumer law, but their effects are very personal. At scale, they have societal implications. So while it may not be for the CMA to regulate online harms to citizens, the problems set out in the paper also have ramifications for citizens at an individual level. On the basis that algorithmic distortion of the information market has a wider impact on society, as recognised in the paper, the CMA should consider societal implications as part of its investigations into consumer harms. Or - although ideally *and* - Ofcom, as the regulator for online harms, should be granted similar powers to the CMA. Without the two regulators working in step and with equitable enforcement powers, the UK's digital regulatory regime will be piecemeal and disjointed.

2.   Do you agree with how we have described each harm, and are there other examples that demonstrate them in addition to the examples we have included?

The harms are excellently set out in this paper. Recommendation and filtering algorithms are particularly harmful in distorting the information marketplace. They provide content which reinforces beliefs, polarises opinion and overextends our time online. The paper notes that they are not a focus for the CMA in this exercise, but we hope they will be properly interrogated by the CMA and Ofcom as part of the DRCF given the effects for consumers and citizens alike. This recent article in MIT Technology Review provides unprecedented insights into how Facebook trains and develops algorithms which aim to maintain and drive user engagement at any cost. It might be useful reading for the CMA.

3. How likely and impactful are the identified harms now, and how might they evolve in the next few years?

Without fast paced, robust regulation these harms will only steamroll over the coming years. Regulatory regimes should align globally to avoid the risk of regulatory arbitrage. The UK is slightly ahead of the rest of the world on this agenda, and a meaningful regulatory regime would have a waterfall effect on global policymaking. There is much resting on the UK's response to the challenges posed by digital platforms. The UK should be bold in setting the agenda and, critically, move fast.

4. Are there specific examples that we should investigate further to consider whether they are particularly harmful and potentially breaching consumer or competition law?

The interplay between consumer and citizen harm should be interrogated further. As the paper notes, algorithmic discrimation takes many forms. It is encouraging to see the CMA will work with the ICO and EHRC on this issue, and we would hope Ofcom would also be involved in those discussions.

Distortion of news content both as an economic harm and a democratic harm needs to be addressed. The paper touches on how publishers' content can be downranked unexpectedly and without explanation, causing them financial loss and expensive investigations. The knock-on effect to citizens - as consumers of news - is that they have inconsistent and unpredictable access to quality information. It also means news media outlets are weakened and distracted from doing what they do best. Coupled with recommendation algorithms which promote sensationalist content, often at the expense of trusted, quality news outlets, this equates to a broad societal harm. We know the CMA is looking at solutions to tackle this challenge, and we encourage HMG to take bold and concerted action to prop up a fundamental pillar of our democracy and digital markets.

5. Are there any examples of techniques that we should be aware of or that we should consider beyond those that we've outlined?

Reset.

In our recent paper on algorithmic audit, jointly penned with the Ada Lovelace Institute, we set out a range of different techniques for regulatory inspection of algorithms which we explore in the context of platform misinformation but which apply to other areas. The paper can be downloaded [here](#) and we have included below a table setting out the techniques for inspection.

Method	Examples	Benefits	Challenges
Documentation	**Policy documentation,** including definitions of misinformation or harmful content, related platform rules and actions, and reasoning behind them	Provides evidence of the company's (claimed) expected behaviour Enables initial scrutiny of policy stance	Without details of company processes and systems, risk of being a high-level understanding of policy intent (and not of realities on the platform)
	Process documentation, including instructions given to manual content moderators	Provides evidence of the company's (claimed) expected behaviour Enables initial scrutiny of process design	If made public, risks making it easier to 'game' moderation system
	Technical system documentation, including: - tools used to identify and moderate information - content recommendation and sharing systems	Provides evidence of the company's (claimed) expected behaviour Enables initial scrutiny of technical design	If made public, risks making it easier to 'game' moderation systems Concerns about intellectual property
Self-reported metrics	Self-reported metrics on misinformation and harmful content, such as: - Model performance for recommender and moderation systems (including false positives and false negatives) - Commercial data for promoted content that's later moderated. - Engagement metrics for content that's later moderated	Provides evidence of the extent to which company believes it is meeting standards	Lacks independent verification Platforms can selectively choose what to report

Reset.

Interviews	Interviews with staff beyond the typical policy and legal teams who interface with regulators, such as: - Technical staff on product teams focused on moderation and recommendation software (product managers, engineers, data scientists) - Moderation teams implementing policies	Direct access to those who design and implement systems will more quickly reveal the principles underpinning the system, and design and engineering decisions and trade-offs	The power dynamic of employer-employee relationship may pressure interviewees Technical staff themselves may not fully understand algorithm behaviour and output
Dataset provision	Datasets shared with inspectors could include samples of moderated and unmoderated content and/or training data to develop moderation or recommendation models	Enables independent scrutiny of system, and provides inputs and outputs to verify function and impact	Datasets provide a snapshot of a single point in time - they may become out of date as user behaviour or system algorithms change Datasets may be selective Privacy concerns for users
API access	Access to new or extended APIs for an inspector, such as access to live platform data	Enables real time/rolling scrutiny of a system's inputs and outputs to verify function and impact	Ongoing access must be agreed upon Companies could manipulate data available through the API
Code access	Access to code that underpins moderation or recommendation systems	Allows interrogation of algorithms and verification of system function	Code changes over time; access would need to be ongoing to be meaningful Security threat of ongoing access to systems Privacy concerns for users Understanding the code would require technical expertise (which may vary by platform). This would likely be slow and would benefit from support of engineers working at the social media platform Concerns about

			intellectual property
Inspector-set test results	A test or dataset for companies to run on their platforms (or for the inspector to run through a private API), in order to collect test results This could include benchmark datasets for different types of harms (which could be used to compare performance across platforms, or for a single platform over time)	Allows access to information and systems that are not public without direct access to systems	Results are not independently verifiable; concerns raised about reliability Hard to set universal tests for different platforms due to different content formats or processes, and it's challenging to keep them up to date as platforms develop

6.   Are there other examples where competition or consumer agencies have interrogated algorithms that we have not included?

We would recommend the work being carried out by Rebekah Tromble at George Washington University who is a leading expert on algorithms and associated regulatory interventions. Some of her work can be found here: https://www.rebekahtromble.net/overview. We would be happy to make a connection to Rebekah is helpful.

If the CMA is unaware of the Citizen Browser, it is worth exploring as an insight into how algorithmic targeting and filtering impacts social media feeds. https://themarkup.org/citizen-browser

The Citizen Browser Project is a custom web browser designed to audit the algorithms that social media platforms use to determine what information they serve their users, what news and narratives are amplified or suppressed, and which online communities those users are encouraged to join. Initially, the browser will be implemented to glean data from Facebook and YouTube

The Citizen Browser publishes the data behind its finding on Github. Here is the data showing how different ethnic groups in America are shown varying levels of public health information

about Covid. Minority communities were shown much less official public health information than white communities.

The Browser also looks at how recommendation algorithms affect visibility of news content. For example, in the run up to the elections for the Senator of Georgia, USA, The Browser looked as users' Facebook feeds before and after Facebook allowed political posts to be reinstated on the site, following a ban on such posts during the Presidential election. While Facebook's controls *were* in place, they found that links to traditional news sites were present in almost all election-related posts that appeared on our Georgia panelists' feeds. After they were removed and Facebook flipped the switch to turn on political advertising for the Georgia election, they noticed that partisan content quickly elbowed out news sites, replacing a significant proportion of mentions of the election in users' feeds.

While issues of focus for The Citizen Brower are slightly out of scope for the CMA, given the limited evidence available to regulators and researchers it is worth the CMA being aware of their findings.

7.   Is the role of regulators in addressing the harms we set out in the paper feasible, effective and proportionate?

Yes. As the paper notes, regulatory oversight of these issues have been slow and ineffective. Regulators have their work cut out in getting ahead of these issues. The focus on transparency is absolutely right one - CMA and other regulators need to be able to access and interrogate algorithms, which run the digital markets. Without this access, the regulatory agenda will fail. Reset and the Ada Lovelace Institute have published a paper on algorithmic audit which can be downloaded here: https://www.adalovelaceinstitute.org/blog/algorithms-social-media-realistic-routes-to-regulatory-inspection/

The recommendations in this paper align with those in the CMA paper, calling for greater regulatory oversight of algorithms. In the paper we note how the FCA and ICO have similar powers already, although these should be bolstered by learnings from the CMA's consultation.

Reset.

8. **Are there other ideas or approaches that we should consider as part of our role?**

No answer.

END