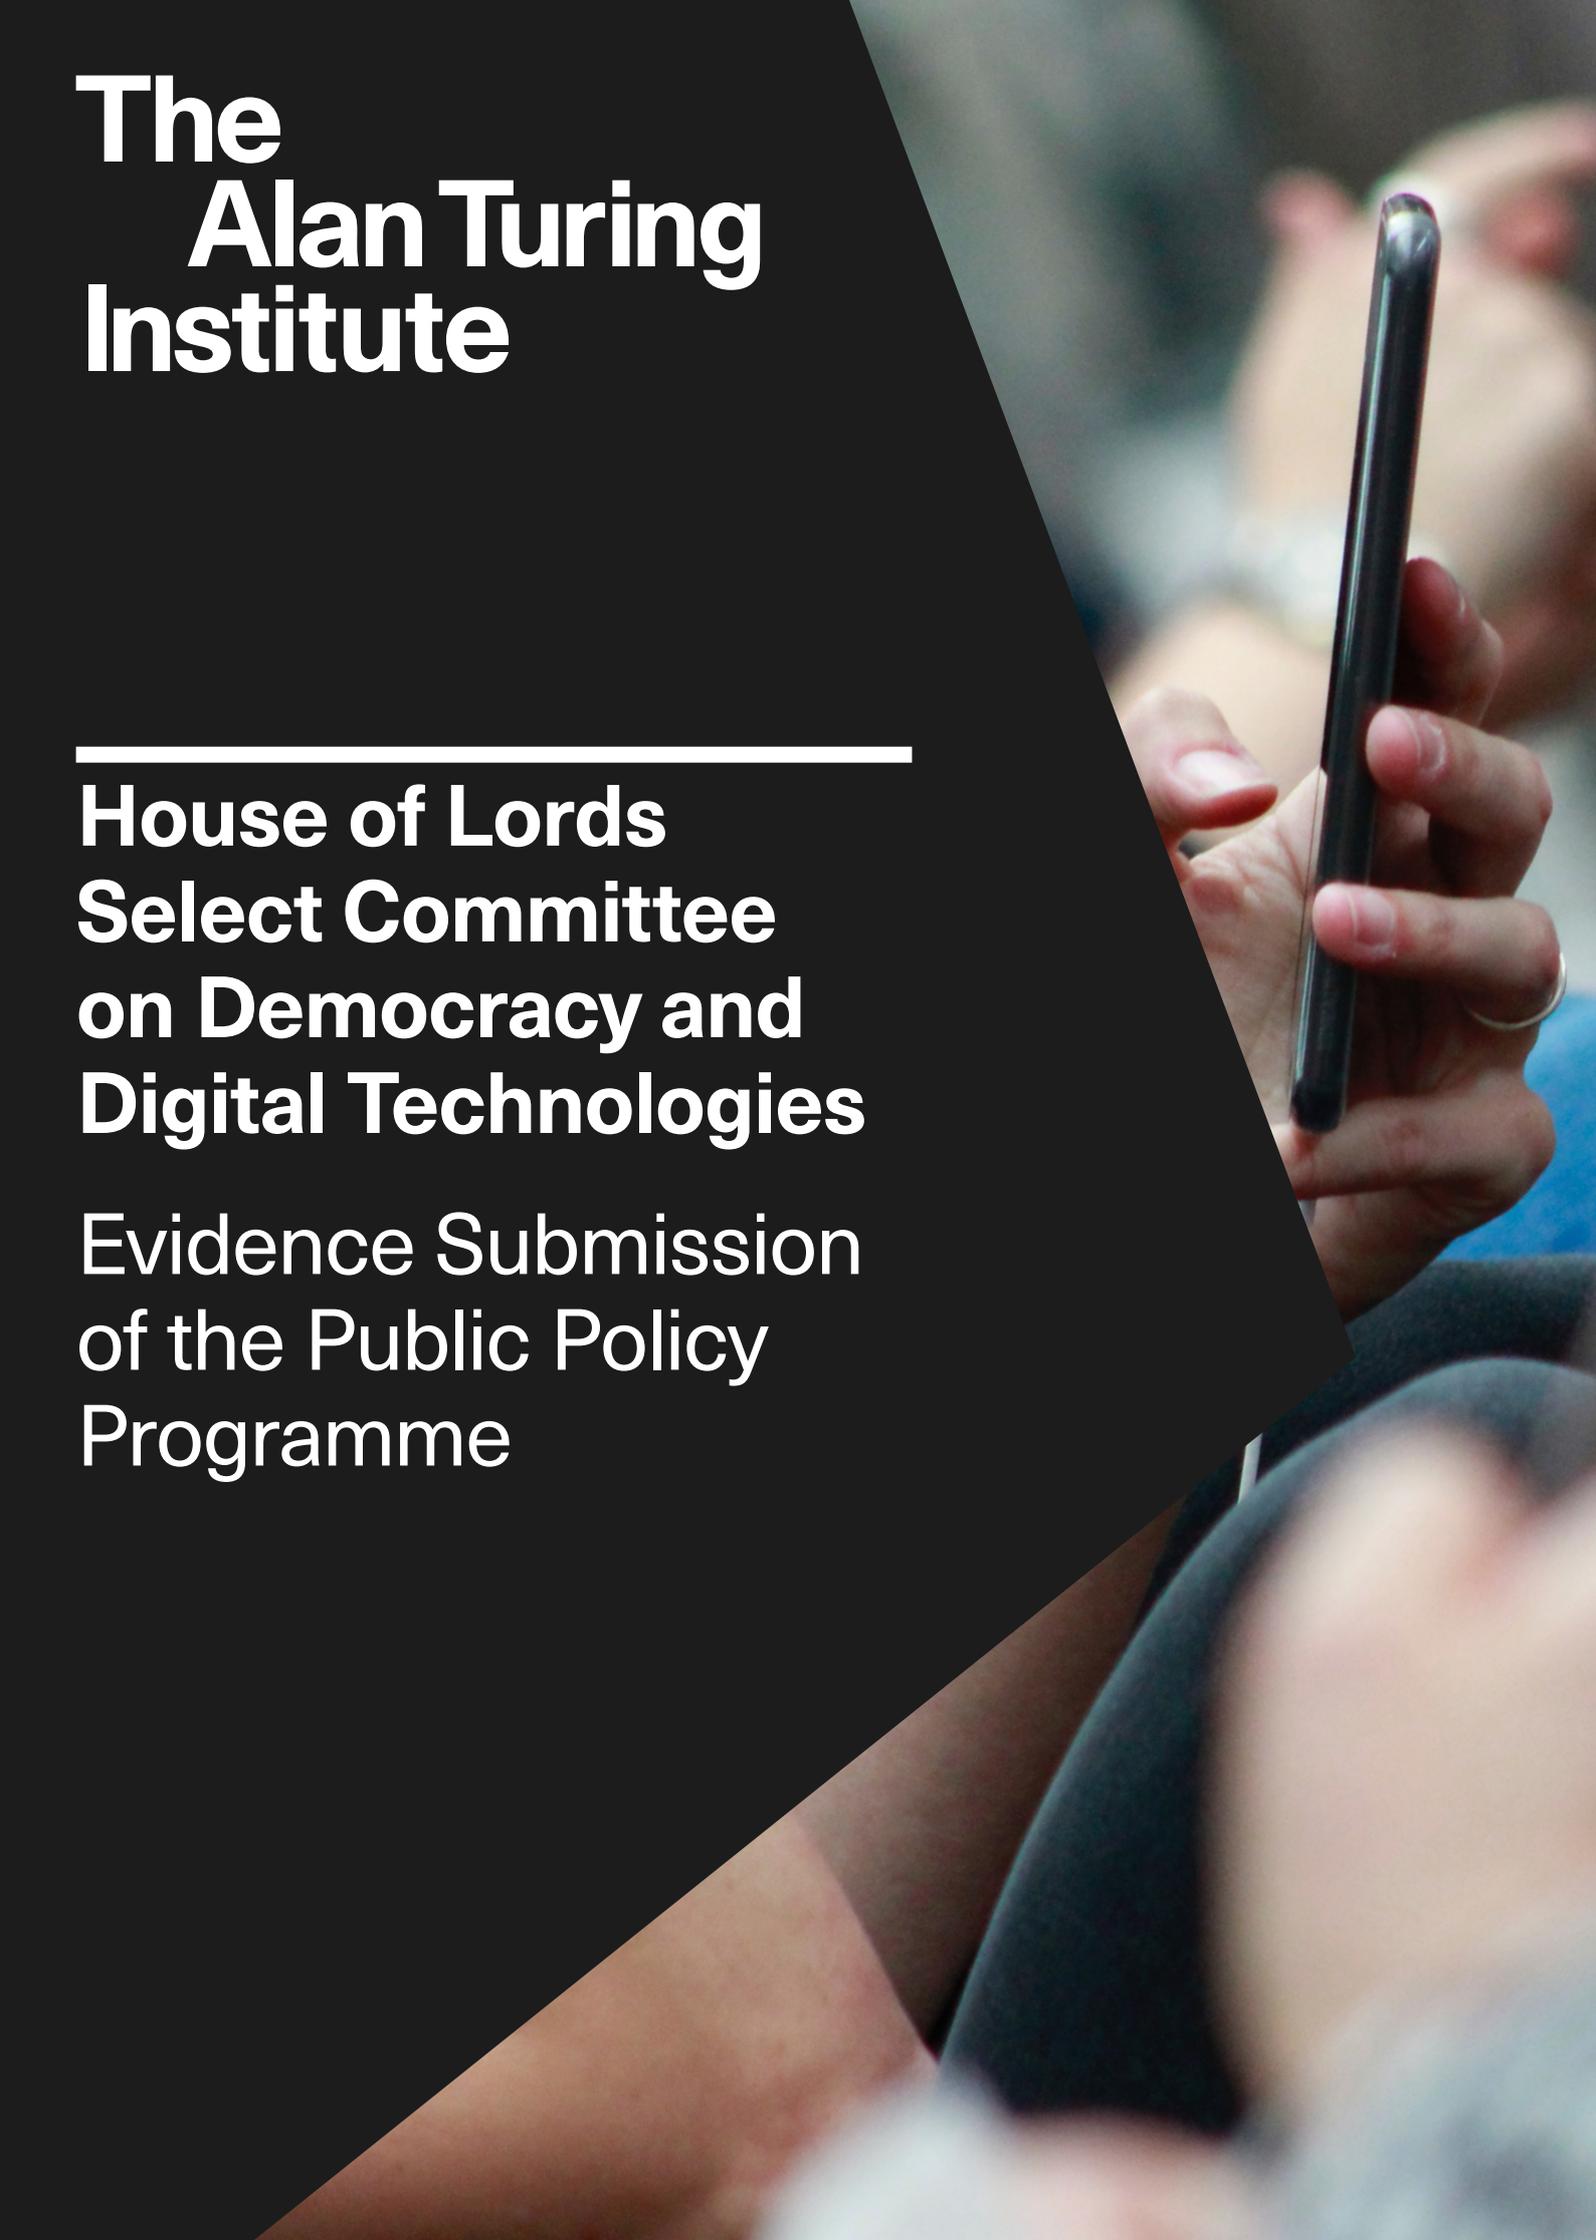


# The Alan Turing Institute



---

**House of Lords  
Select Committee  
on Democracy and  
Digital Technologies**

**Evidence Submission  
of the Public Policy  
Programme**

## **Evidence submission to the House of Lords Select Committee on Democracy and Digital Technologies: Questions 8 & 11**

### **About the contributors**

This contribution is made by The Alan Turing Institute's [Hate Speech: Measures & Counter-measures](#) project, which forms part of the Institute's Public Policy Programme. We are happy to work more closely with the House of Lords Select Committee and to provide any further information or guidance as needed. Please contact Dr. Bertie Vidgen, [bvidgen@turing.ac.uk](mailto:bvidgen@turing.ac.uk)

Dr. Bertie Vidgen is a post-doctoral researcher at The Alan Turing Institute, Research Associate at the Oxford Internet Institute and Visiting Researcher at the Open University. His research focuses on online abuse, hate speech, misinformation and the far right. He is a co-Investigator on The Alan Turing Institute's project, 'Methods for abusive content detection', lead researcher on The Alan Turing Institute's project 'Hate speech: measures and counter-measures' as well as an organiser of the 4<sup>th</sup> workshop on abusive language online (ACL). He received his PhD from the University of Oxford. In his work he combines advanced computational tools and social science theory, with the aim of contributing to both academic knowledge and the work of policymakers.

Professor Helen Margetts is Director of the Public Policy Programme at The Alan Turing Institute, and Professor of Society and the Internet at the University of Oxford and Professorial Fellow of Mansfield College. From 2011 to 2018, she was Director of the Oxford Internet Institute, a multi-disciplinary department of the University of Oxford dedicated to understanding the relationship between the Internet and society, before which she was UCL's first professor of Political Science and Director of the School of Public Policy (1999-2004). She sits on the UK government's Digital Economy Council, the Home Office Scientific Advisory Council, the WEF Global Agenda Council on Agile Government and the Ada Lovelace Institute for Data Ethics. She was a member of the UK government's Digital Advisory Board (2011-16). She is the founding Editor of the journal *Policy and Internet*, published by Wiley. In 2018 she was awarded the Friedrich Schiedel Prize by the Technical University of Munich for research and research leadership in technology and politics. In the 2019 New Years Honours List she was awarded an OBE for services to social and political science. In 2019 she was elected as a Fellow of the British Academy and also took up a visiting appointment as the John F Kluge Senior Chair in Technology and Society at the Library of Congress.

**Question 8: To what extent does social media negatively shape public debate, either through encouraging polarisation or through abuse deterring individuals from engaging in public life?**

Social media has transformed contemporary politics, from making huge audiences instantaneously accessible to sidestepping the role played by 'gatekeepers', such as the traditional broadcast media. Social media are where people find, consume and share political information and news, and participate, communicate and organise politically. In this way, social media enable users to engage in new 'tiny acts' of participation (through liking, following, viewing, sharing and so on) which have the positive effect of drawing people into politics (including groups where participation has traditionally been low) and sometimes this can scale up into large mobilizations and campaigns for policy change. Although most such mobilizations fail, the speed and scale of the ones that succeed have had dramatic effects on political life across the world, with almost every country in the world seeing a rise in collective action organised through social media, and offering people with no more resources than a mobile phone the opportunity to fight injustice and campaign for policy change.<sup>1</sup> Social media have had a number of positive effects on public life and political discourse, such as enabling like-minded people who are separated by time, space, social networks and culture, to connect with each other, and in those cases fostering respectable and constructive dialogues. Support networks and activist groups which would not previously be possible are now commonplace and it has never been easier to learn about competing perspectives, find out new information and to engage with people who have different views.

However, social media have also given new voice to otherwise marginalised actors across the political spectrum, such as extremist far right protestors and niche left-wing environmentalists. Some argue that the affordances of social media have fundamentally changed how people interact with each other. Anonymity and not being face-to-face with interlocutors mean that people are far more impolite and angrier. This line of reasoning suggests that although people troll, dox and 'cancel' opponents online, if they met face-to-face in the offline world then they would be civil and reasonable. As one study of comments below online news articles put it: 'Anyone can become a troll'.<sup>2</sup> For social science researchers, disentangling the effects – both positive and negative – of social media is a

---

<sup>1</sup> For a full discussion, see: Margetts et al. (2015), *Political Turbulence: how social media shape collective action*, Oxford: Princeton University Press.

<sup>2</sup> Cheng et al., 'Anyone can become a troll: causes of trolling behaviour in online discussions'. Available at: <https://dl.acm.org/citation.cfm?id=2998213>

considerable challenge: its ubiquity and the volatile nature of contemporary politics means that we cannot be sure exactly which political developments are due to social media and which are due to other causes. And, much as we are often pressed to ‘take a side’ when discussing whether social media has had a positive or negative impact, the truth is inevitably more complex: depending on how it is used, and by whom, social media can have both a positive and negative impact.

This consultation response focuses on one particular aspect of online debate which has raised concern in recent years – abusive language. Our research in *Hate Speech: Measures and counter-measures* shows the myriad problems it creates for society: it can inflict harm on any victims who are targeted, create a sense of fear and exclusion amongst their communities, toxify public discourse and, even, motivate other forms of extremist and hateful behavior through a cycle of “cumulative extremism”. Online abuse is a deeply harmful part of online discourse<sup>3</sup> which, whilst it is often defended on the basis of free speech, we should take every effort to better understand, challenge and eradicate.

### **Prevalence of online abuse**

Surprisingly, relatively little is known about the prevalence of online abuse, which seriously constrains our ability to understand and tackle it. Through an ongoing review, we have identified several key insights about contemporary online abuse in the UK<sup>4</sup>:

1. The prevalence of legally defined abuse is incredibly low. 1,605 online hate crimes were recorded in England and Wales in 2017/2018. We estimate the number of offences for online harassment is around 1 per 1,000 people. Crime survey data from Scotland provides a higher estimate, indicating that ~1% of people have experienced online harassment. However, to fully understand this, we need more statistics to be made available, ideally collated in a single report.
2. The prevalence of online abuse on mainstream platforms which is serious enough for them to take down is also very low. We estimate that it is ~0.000001%. For instance, in Q1 of 2019, Facebook removed 4 million pieces of content for hate speech – but this figure relates to all of the content it hosts, including both new posts made during the period and the historical content it still hosts, going back to when Facebook was founded. For a

---

<sup>3</sup> As discussed. In the recent Online Harms White Paper from DCMS and the Home Office, available at: <https://www.gov.uk/government/consultations/online-harms-white-paper>

<sup>4</sup> A policy briefing report summarising the available evidence on online abuse is forthcoming from The Alan Turing Institute’s Hate Speech: Measures and counter-measures team.

platform which reported in 2013 that 4.75 billion posts are made each day, this means that a tiny fraction of its content is considered hateful.<sup>5</sup>

3. Measurement studies from academics and thinktanks indicate that prevalence of abuse is higher than the level indicated by mainstream social media platforms' content take downs. It is still typically less than 0.1%. However, some users and events generate more abuse, such as prominent figures (e.g. MPs) and terror attacks.
4. Niche online spaces contain far more abuse, and in some spaces between 5-8% of content is abusive. Most research in this domain has focused on analysing hate speech and there is a lack of research into other forms of abuse, such as harassment.
5. In strong contrast, the amount of people who report having observed, or being targeted by, online abuse is very high. Based on available survey data, we estimate that between 30-40% of people in the UK report having been exposed to online abuse. We also estimate that 10-20% of people in the UK have personally received abuse online. Note that this gives no insight to how many times people have experienced online abuse: in many cases, people will have only experienced online abuse once or twice or only seen one abusive posts out of thousands. Nonetheless, even limited experiences of abuse can be enough to seriously impact their wellbeing and health.

Overall, these results suggest that whilst the prevalence of online abuse is low, especially in terms of content which is illegal or contravenes platforms' guidelines, many people are still exposed to it. This somewhat counterintuitive finding (that online abuse is rare but still seen by many) is crucial for understanding the contemporary landscape of online abuse. It is a key reason for why online abuse is so hard to tackle: even just a few bits of abusive content online can have huge impact which, given the need to balance taking down abuse with protecting individuals' freedom of speech and right to privacy, creates a 'wicked problem' for policymakers.

### How abuse manifests

Abuse impacts on different people and groups very unevenly. For instance, MPs and public figures often receive far more online abuse than ordinary citizens.<sup>6</sup> Our research on how MPs used Twitter during the 2017 general election campaign, also shows that *how* they receive

---

<sup>5</sup> Reported by Facebook here: <https://www.facebook.com/FacebookSingapore/posts/did-you-know-that-475-billion-pieces-of-content-are-shared-daily/563468333703369/>

<sup>6</sup> See Gorrell et al. (2018), 'Twits, Twats and Twaddle', available at <https://www.aaii.org/ocs/index.php/ICWSM/ICWSM18/paper/viewPaper/17861>

Hate Speech: Measures & Counter-Measures

abuse is very different to other online actors.<sup>7</sup> We find that whilst there is always a background level of online abuse, infrequent but intense ‘events’ occur too: brief periods of time in which MPs are targeted by huge amounts of abuse from a wide array of Twitter users. Figure 1 shows the number of abusive tweets received by the Conservative MP Nadine Dorries. The turquoise sections show the occurrence of hate ‘events’ and the red sections show the background level of abuse she constantly experiences. Most of the abuse she receives happens during the short sharp event periods.

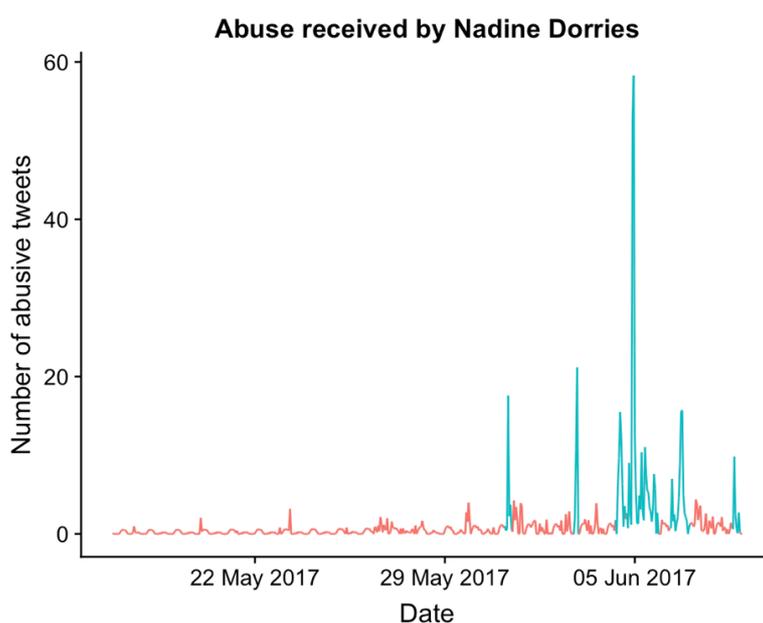


Figure 1, Abuse received by Nadine Dorries during the 2017 general election

Intense hate events are likely to have considerable emotional and mental health impact on MPs and may create the perception that everyone is opposed to them – even though, in many cases, it is just a small but vocal minority who are attacking them for a brief period. We need to develop ways of providing more support to figures who are in the public eye, and not just accept online abuse as part of the rough and tumble of politics. Otherwise, we risk excluding certain types of people from public spaces, in particular those who are either more vulnerable or have less support around them. This is not only unfair but will also reduce the diversity and quality of public discourse, negatively impacting all citizens and potentially alienating some groups from civic engagement.

---

<sup>7</sup> Publication is forthcoming.

Our research also shows that experiences of online abuse are also highly uneven in other ways, such as by peoples' demographics. This is based on original analysis of previously unpublished data from the Oxford Internet Institute's 2019 Oxford Internet Survey (OxIS). We are grateful to the authors for giving us permission to use this data, and to Dr. Grant Blank for making it available.<sup>8</sup> Our results show that younger people, non-Whites and more regular Internet users are more likely to experience online abuse. This unevenness in how online abuse manifests means that when we talk its impact on online democracy, participation and discourse we should not adopt a 'broad brush' and act as though all users' experiences are the same: experiences of online abuse vary systematically by background and activity.

### **Online abuse: OxIS 2019 survey results**

Two questions from the 2019 OxIS survey relate to online abuse

- (1) Have you seen cruel or hateful comments or images posted online?
- (2) Have you received obscene or abusive emails?

In total, 27% of respondents had seen cruel or hateful comments or images posted online. 10% of respondents had received obscene/abusive emails. However, these headline figures mask important differences based on respondents' demographics. Ethnicity impacted respondents' experiences of online abuse, with Black people experiencing far more online abuse than White people. Asian and 'Other' respondents fell in between these groups for both questions. This is shown in Figure 2.

Age also impacted respondents' experiences of online abuse. Younger people are more likely to experience abuse. 41.2% of 18-30 year olds had seen cruel/hateful content online, compared with 7.4% of 76+ year olds. Age also impacted whether respondents had received obscene/abusive emails but the relationship was far weaker, ranging only from 13.1% of 18-30 year olds to 6.77% of 76+ year olds. Differences in experiences of online abuse according to age are shown in Figure 3.

---

<sup>8</sup> Please see OxIS's website for more information about the survey: <https://oxis.oii.ox.ac.uk> and Dr. Grant Blank's online profile: <https://www.oii.ox.ac.uk/people/grant-blank/>

OxIS is the longest-running academic survey of Internet use in Britain and uses a multi-stage national probability sample of 2,000 people. Current internet users, non-users and ex-users are included, which allows the data to be used for representative estimates of Internet use, with a low margin of error.

Hate Speech: Measures & Counter-Measures

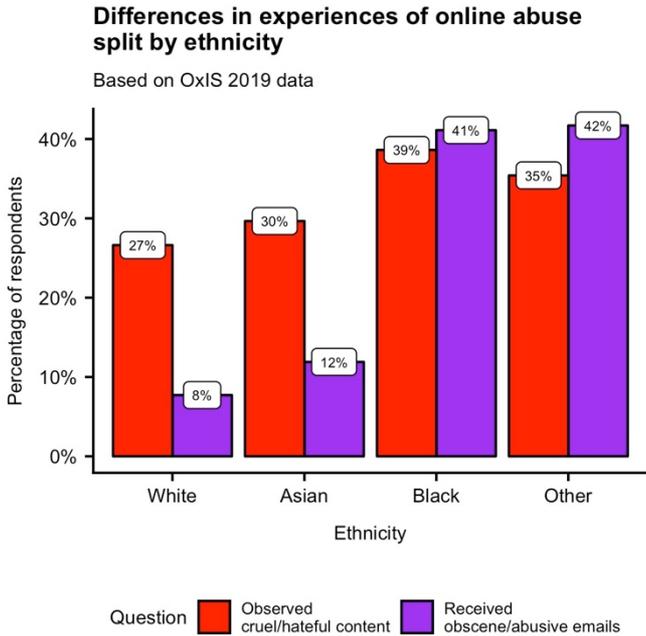


Figure 2, Experiences of online abuse, split by ethnicity (OxIS 2019 data)

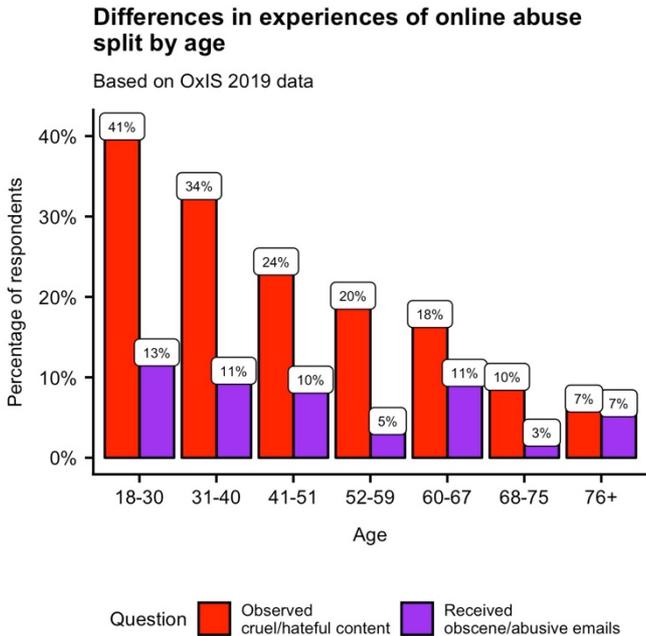


Figure 3, Experiences of online abuse, split by age (OxIS 2019 data)

Hate Speech: Measures & Counter-Measures

Internet use impacted respondents' experiences of online abuse. 38.9% of users who 'Constantly' go online had seen cruel/hateful content online, compared with just 5.2% of users who go online less often than once per week. More time spent online was also associated with being less likely to receive obscene/abusive emails, from 16.6% of users who 'Constantly' go online to 2.2% of users who go online once per week. However, one exception is users who go online less often than once per week, of whom 21.5% had received obscene/abusive emails. Differences in experiences of online abuse according to Internet use are shown in Figure 4.

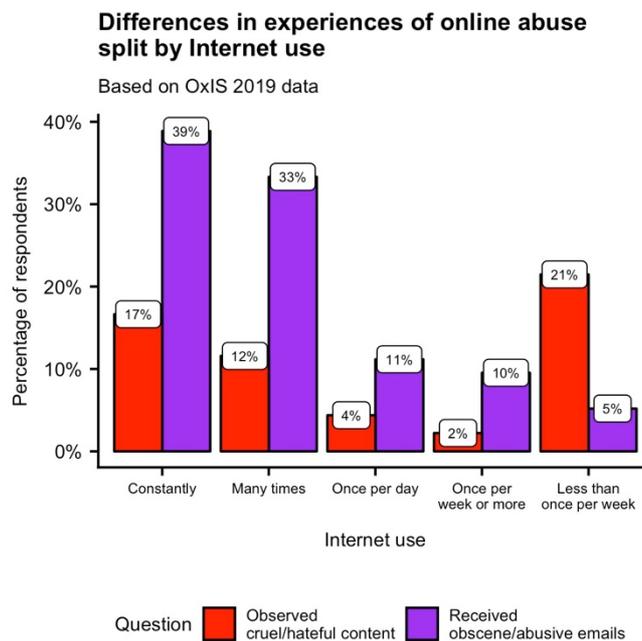


Figure 4, Experiences of online abuse, split by Internet use (OxIS 2019 data)

Finally, our results show that gender played a far smaller role in experiences of online abuse. For both variables, a similar proportion of Males as Females experienced online abuse. See Figure 5. However, we note that other research indicates that gender is a crucial factor in experiences of online abuse, and so advise caution when interpreting this result.

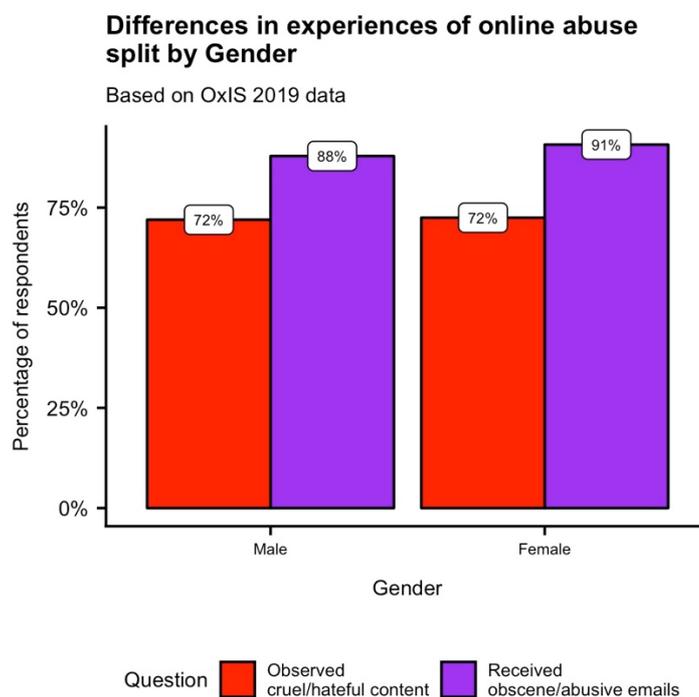


Figure 5, Experiences of online abuse, split by Gender (OxIS 2019 data)

**Question 11: How could the moderation processes of large technology companies be improved to better tackle abuse and misinformation, as well as helping public debate flourish?**

Content moderation is crucial for ensuring a safe and fair digital ecosystem in which all users are free to reap the benefits of communication technologies. There is no such thing as a non-moderated online space: all platforms are ‘designed’ and users’ experiences curated. Platform design effects everything that users experience, including what features are available to them, what content they are exposed to, and how they can connect/engage with other people. Thus, the question facing society is not *whether* we want content to be moderated but *how* we want it to be moderated.

In this evidence submission, four key aspects of content moderation are discussed: (1) Scrutiny, transparency and collaboration, (2) Identifying harmful content: the role of technology, (3) How should platforms intervene? and (4) Time sensitive interventions in content moderation.

## 1. Scrutiny, transparency and collaboration

Evaluating the moderation processes of large (and small) tech companies is challenging for one overarching reason: they provide little information and only rarely open up their processes to outside scrutiny. For instance, according to Pew Research, the most popular social media platforms in the USA in 2019 are Facebook, Twitter, YouTube, Snapchat, Instagram, WhatsApp, LinkedIn, Pinterest and Reddit (Pew Research, 2019).<sup>9</sup> Of these 9 platforms, only 4 (Facebook, Twitter, YouTube, Reddit) provide statistics on the amount of hate speech and bullying content they remove globally. At present, transparency reporting is entirely optional and there are no clear frameworks for what should be reported on. Platforms each report different statistics at different levels of granularity and use different categories. This makes it very difficult to compare their processes and to build an evidence base about what works and what needs to be improved.

Part of the reason for the lack of transparency, and inconsistency in how platforms report on content moderation, is that it is a commercially sensitive area. Nonetheless, the impact of ineffective content moderation is felt by all in society, particularly by marginalised and vulnerable groups. As such, whilst respecting their right to protect their commercial interests, platforms should aim to share more information with researchers, civil society and policymakers through cross-sector collaborations. We make three proposals for improving collaboration:

1. Industry wide standards and frameworks for determining what content is considered harmful should be agreed. Clear definitions, with examples, should be provided. This should address more than just what is prohibited by law (as most platforms' guidelines do currently).
2. Opportunities to share resources, data and frameworks should be supported, publicised and given funding. For example, in the Turing's *Hate speech: measures and counter-measures* project, we maintain a list of datasets for other researchers to use, available at: <http://hatespeechdata.com>. More endeavours like this should be pursued.
3. A forum should be established to regularly bring together industry, civil society and academics to discuss how to tackle online harms. This should be organized as a threat assessment exercise, with participants discussing which challenges in content moderation are growing and receding, any emergent problems, and how issues can be better tackled.

---

<sup>9</sup> Figures from the USA are used as comparable survey data for the UK is not currently available.

## 2. Identifying harmful content: the role of technology

Detecting harmful content at scale in a timely, robust and fair manner is remarkably difficult.<sup>10</sup> Despite the increasing sophistication of available computational tools, most companies still rely on humans to review individual posts – an approach which is both expensive and necessarily reactive (human moderators do not proactively find content, they only review content reported to them). To improve content moderation, industry practitioners need to approach it as not only an *engineering task* (requiring ever more computational power, big datasets and machine learning) but also a *social challenge*. This is because content moderation is fundamentally a question of fairness: who is given protection from unsolicited harmful content? Who is allowed/enabled to spread such content? Which groups systematically have their content over-moderated? Which groups are left feeling unsafe and, in some cases, feel they need to withdraw from online spaces? If we keep approaching content moderation as only an engineering problem, or a financial imposition, we will end up with coarse and ineffective moderation strategies which do not challenge discrimination and injustice but perpetuate them.

Based on our research at The Alan Turing Institute, we identify five challenges that content detection methodologies must address.<sup>11</sup> These are based on our review of hate speech detection technologies but are also relevant for other types of harm. Given existing technologies it is highly unlikely that these challenges would be solved solely with automated computational methods:

6. *Variations in spelling.* Language use online is very varied and spelling varies hugely. There are many reasons for this, including: (1) different communities develop their own unique lexicons, (2) users change words for emphasis (e.g. elongating 'no' to 'nooooo'), (3) platforms' affordances and norms drive changes in spelling (such as with Twitter's constraint on number of characters) and – most problematically with harmful content – (4) users obfuscate their language to avoid detection by existing technologies (e.g. changing 'bitch' to '31tch').
7. *Humour, irony, sarcasm.* Many researchers suggest that content which is ironic, humorous or sarcastic is not intrinsically harmful and as such should not be moderated – even if the content purports to make aggressive or highly prejudiced statements. This remains an open research question as, at the same time, others argue that people

---

<sup>10</sup> See: <https://www.aclweb.org/anthology/W17-1101>, <https://arxiv.org/pdf/1705.09899> & <https://www.aclweb.org/anthology/P19-1163>

<sup>11</sup> Vidgen et al., 'Frontiers and Challenges in abusive content detection', *Proceedings of the 3<sup>rd</sup> workshop on abusive language online*, Available at: <https://www.aclweb.org/anthology/W19-3509>

who *view* such content may not know the authors' intention and could still experience it as harmful. Furthermore, even humorous content can speak a certain truth: ironic hate speech may still reproduce negative stereotypes or hurtful ideas, even if it aims to lampoon them. This issue should be tackled head-on and platforms should make clear where, and for what reason, they draw the line on humour, irony and sarcasm.

8. *Polysemy and context.* Words with a single spelling can have multiple meanings; which meaning is elicited depends on the context. For instance, some far right activists use seemingly neutral terms, such as 'Skype' or 'Bing'. to derogate particular groups (in these cases, Jewish and Asian people).<sup>12</sup> Whether these words are harmful depends entirely on the context in which they are used. Making these distinctions remains a huge challenge for computational systems. This problem can only be addressed through engagement with victims of harmful content to understand their experiences and what they perceive to be the most hurtful.
9. *Long range dependencies.* Longer pieces of online content are far harder to monitor and detect than short pieces because they are far more complex, introducing uncertainty, different lines of reasoning and, in some cases, conversational dynamics. Platforms need to keep developing ways of capturing long range dependencies and following conversational threads so that harmful content (such as expressing support for violent terrorism) can be identified even if it only emerges over a chain of posts or several pages of text.
10. *Language change.* The syntax, grammar, and lexicons of language change over time, often in unexpected and uneven ways. This is particularly true with informal forms of 'everyday' language, which proliferate in most online spaces. Systems need to be constantly updated to account for this, and to understand the evolving nature of online harm. Platforms need to think about how meaning changes and how what was acceptable previously may no longer be acceptable today. They should also consider whether they want to do retrospective 'clean up' missions, removing content that was once allowed but now falls foul of moderation guidelines, or only focus on newly uploaded content.

---

<sup>12</sup> See: <https://www.aclweb.org/anthology/W18-5112>

### 3. How should platforms intervene?

Once harmful content has been detected, platforms have to decide how it should be handled. This is an area in which platforms could both provide considerably more information about their existing processes and do far more to strength public debate and ensure users' safety online. Most public discourse around content moderation focuses on bans, such as suspending and blocking users or removing their content – but this is simply the most overt end of a wide spectrum of interventions, which include:

1. De-monetising content (as used by YouTube)
2. Stopping content from being promoted (potentially used by Facebook)
3. Making content unsearchable (as used by Reddit)
4. Making content less visible (as used by Twitter)
5. Attaching a 'warning' to content (as used by Twitter)
6. Constraining how many times content is shared (as used by Whatsapp)
7. Showing content with a competing viewpoint (No well-known examples exist but it has been discussed in civil society and by academics).

There are clear cases when suspending individual social media accounts may be viewed as entirely appropriate – such as with violent threats, paedophilic content and revenge porn. But, for many other types of content, platforms should consider drawing from the far wider array of options available to them. Unless we open up this space of debate, the ethical, political and social implications of these different options will not be fully explored and more nuanced and appropriate responses for dealing with harmful content will not be developed. Developing a wide range of moderation interventions will also expand the range and types of harms that can be moderated; if bans are your only option then you can only tackle the most extreme forms of content. If you have a wider range of interventions then you can deal with a wider range of problems.

Different interventions may work better for tackling certain types of online harm: some forms of misinformation are shared in order to drive traffic to certain websites and then let the hosts derive financial benefit from advertising (in which case demonetising could work well). In contrast, political figures might share false, inflammatory or hateful content to gain public attention and increase their notoriety (in which case constraining the number of shares could work) and some porn stars share shocking and sensitive content to drive attention, but are only interested in appealing to certain target markets (in which case a sensitivity warning could be effective). However, it is important to draw attention here to the fact that *there is little evidence about what interventions work to tackle harmful online content*. To remedy this, the

Government should utilise its unique position to evaluate different approaches. More broadly, it should adopt a joint-up approach which considers how different types of harmful content intersect and considers the impact of content moderation on the entire digital ecosystem rather than a single platform. It is possible that when mainstream platforms aggressively moderate content they push some users to migrate to niche alternative platforms, which could, in turn, lead to greater radicalization – we need to build an evidence base that addresses these complex issues and not only think about platforms in isolation.<sup>13</sup>

#### 4. Time sensitive interventions in content moderation

The 2019 Christchurch terrorist attack showed the need for platforms to develop time sensitive responses to harmful online content: in just a few hours the New Zealand terrorist's homemade and livestreamed video of the attack had spread across social media. The challenge here is that during these highly volatile, unpredictable periods the amount of harmful online content sharply accelerates: tools, process and systems which are sufficient in normal periods start to break. To address this challenge, we propose four strategies for short intense periods where there is a spike in the amount of harmful content:<sup>14</sup>

1. *Adjust the sensitivity of content detection tools:* All tools for detecting harmful content have a margin of error. The designers have to decide how many false negatives and false positives they are happy with. False negatives are bits of content which are allowed online even though they shouldn't be and false positives are bits of content which are blocked even though they should be allowed. There is always a trade-off between the two when implementing any content detection system. During times of increased harmful content, tools should be adjusted so that platforms catch more harm (reducing the number of false negatives), even if that means temporarily also having more false positives.
2. *Enable easier takedowns:* At present, moderated content is flagged by users and then sent for manual review by a content moderator, who checks it using predefined guidelines. During times of increased harmful content, platforms could introduce special procedures so that staff can quickly work through content without fear of low

---

<sup>13</sup> For more information about regulation of online harms, see the Turing's Public Policy Programme's response to the 'Online Harms' white paper: <https://www.turing.ac.uk/research/research-programmes/public-policy/programme-articles/turings-public-policy-programme-responds-online-harms-white-paper>

<sup>14</sup> Based on a previous blog post by this submission's author for *The Conversation*: <https://theconversation.com/four-ways-social-media-platforms-could-stop-the-spread-of-hateful-content-in-aftermath-of-terror-attacks-113785>

performance evaluation. They could also introduce temporary quarantines so that content is temporarily removed but then re-examined at a later date.

3. *Limit the ability of users to share:* In normal times, platforms encourage users to share as much content as possible. However, research shows that extreme and hateful content is imported from niche far-right sites and dumped into the mainstream<sup>15</sup> where it can quickly spread to large audiences. During times of increased harmful content, platforms should limit the number of times that content can be shared within their site and potentially stop sharing across platforms.
4. *Create shared databases of content:* All of the big platforms have very similar guidelines on what constitutes harm and will be trying to take down largely the same content during times of increased harmful content. Creating a shared database of harmful content would ensure that content removed from one site is automatically removed from another. This would not only avoid needless duplication but enable the platforms to quickly devote resources to the most challenging content which is the hardest to detect.

---

<sup>15</sup> See: <https://conferences.sigcomm.org/imc/2017/papers/imc17-final145.pdf>



---

**turing.ac.uk**  
**@turinginst**