Online hate and abuse in sport

A report by Ofcom in partnership with Kick it Out

Published 16 May 2025

Contents

Section

Disclaimer

CONTENT WARNING: This report includes details of people's experiences of online hate and abuse. It also discusses the broader impact this content can have, including a brief discussion of disordered eating, depression and suicide.

Acknowledgments

We would like to thank the following organisations who contributed to this research:

- Kick It Out
- Governing bodies and clubs (including Premier League)
- Broadcasters (including Sky and Warner Bros. Discovery)

A special thank you to the following people who spent time with us sharing their stories:

- Andrew Wolfe Murray
- Azeem Rafiq
- Chloë Mitchell
- Claire Callaghan
- David McArdle
- Greig Mailer
- Jennie Blackmore
- Jo Tongue
- Joe Sayers
- Jonathan Hirshler
- Jonathan Sebire
- Kenny England
- Maddie Hudson
- Paul Struthers
- Ruby Malone
- Wayne Barnes

We would also like to thank everyone who participated but who chose to remain anonymous.

Ofcom Preface

Being online is an important part of modern life. The internet provides information, entertainment, and helps people stay connected.

For those working in the public eye – such as professional sportspeople, sports club staff, and individuals employed by broadcasters – an online presence can be an unavoidable part of their job and vital for their career. By using internet services, they can connect with fans, keep up with discussions about a variety of sports, and build and maintain a public profile that can lead to opportunities, like sponsorship deals. There are, however, risks with being online, and those working on issues in the public eye can face harmful online abuse and hate.

Ofcom is the UK's online safety regulator and has been given legal powers under the Online Safety Act 2023 ('the Act') to hold online services to account for protecting their UK users. This involves making sure that online services have systems and processes in place to protect UK users from illegal harms, such as inciting violence, and protecting children from content that is harmful to them.

Some of the online abuse people experience today is illegal under UK law, such as some types of threatening or abusive behaviour and harassment. Online services must assess the risk of users encountering this kind of material and then put in place measures to mitigate against these risks. Ofcom's Illegal Harms Codes of Practice set out ways service providers can comply with these duties.

However, much of the abuse people experience online, although distressing and harmful, is not illegal under UK laws. Platforms take different approaches to content that is harmful, but not illegal: some ban certain categories of harmful content in their terms and conditions; others provide users with more tools to protect and tailor their online experiences; and some may choose to do nothing at all.

Some platforms are subject to additional duties under the Act, such as those designated as 'Category 1' providers.[1] These services will need to provide adult users with features that enable them to reduce the likelihood of encountering certain types of abusive and hateful content, or to be alerted to the presence of this content on the platform. Ofcom must produce a Code of Practice setting out measures service providers can take to comply with these duties.

We collaborated with Kick It Out, the anti-discrimination body in sport, to engage directly with people who have lived experience of online hate and abuse, to listen and learn from these experiences while developing measures to be effective. We conducted seven individual interviews and nine group discussions with sportspeople, on-screen commentators, and professionals working in sport and broadcasting. They were recruited through Kick It Out's network.

We have previously commissioned The Alan Turing Institute to a report on online hateful content, and explored the volume of hate and abuse faced by footballers on Twitter, now X. Separately, we have researched first-hand the impact of online hate directed at members of the general public. We

[1] These are providers of user-to-user services that meet the threshold conditions specified in regulations made by the Secretary of State under Schedule 11 to the Act (known as 'Category 1 services'). Sections 15(2) and (3) of the Act set out these user empowerment duties. The type of content they apply to includes content that is abusive and targets people based on, or incites hatred against people with, particular characteristics, namely race, religion, sex, sexual orientation, disability or gender reassignment. Separate duties under the Act also require services likely to be accessed by children to protect children from content that is legal but harmful to them, including abusive and hate content as described above.

are therefore aware that online hate and abuse is a particularly challenging topic to research, especially since what people understand it to mean can be quite personal and not always linked to legal definitions. This means that it can be difficult for members of the public to establish whether content is illegal or legal under UK laws, or whether it breaches services' own terms and conditions.

Building on our existing understanding of online hate and abuse, this project enabled us to get a specific, in-depth account from professionals working in the sports and broadcasting sectors who have actively used or have a presence on social media as part of their job. This report focuses on those participants' perceptions about the harmful content they experience online and what they think platforms should be doing about it. This report reflects their views, and the findings should not be considered Ofcom's opinion. Where appropriate, these findings will be used as evidence to contribute to policy development. However, some suggestions may not be possible within Ofcom's powers under the Act.

Our participants recognised the inevitability of receiving criticism as high-profile figures, but believed more should be done to limit 'dehumanising' attacks about them based on their protected characteristics. Many of the people we spoke to told us that the harm they experience online can bleed into their offline lives and impact the lives of their families and friends. We have heard about the harm that online hate and abuse can cause, and we are talking to platforms about what more they should do to help protect their users.

This report shares users' personal experiences and views which we have heard first-hand. We recognise the importance and value of hearing directly from those impacted by online harm and listen to what changes they want to see to inform our evidence base. We will continue to learn directly from those with lived experience and share our findings in the months and years ahead.

1. Executive Summary

Ofcom collaborated with Kick It Out, the anti-discrimination body in sport, to engage directly with people who have lived experience of online hate and abuse. We conducted seven individual interviews and nine group discussions. Participants included sportspeople, on-screen commentators, and professionals working in sport and broadcasting.

Participants told us that, in their experience, online hate and abuse is:

- **Getting worse over time** as it is becoming more common, sophisticated, and normalised.
- **Rapidly evolving** as perpetrators are using different phrases, terms and emojis, which makes it more difficult to filter out.
- **Becoming bolder** as there is a perceived lack of consequences for accounts that post it, since some posts are made from accounts that are not anonymous. Where accounts with large followings encourage others to post hateful and abusive content, this can make it seem acceptable.
- **Being incentivised** by the business models of online services that monetise engagement. Hateful and abusive content can make people react so some perpetrators may post this type of content to increase engagement.
- **Outside of their control,** especially where the hate and abuse they saw was not a direct reaction to their own posts. Crucially, we were told that this limits their ability to protect themselves and others, especially when it is difficult to have the material removed.

Participants told us that online hate and abuse can have a serious impact on their lives. They said it lowers self-esteem, influences career choices, and limits what they feel able to say online and on-screen. It also affects their feelings about their ability to do their jobs and has consequences for their family and friends, both online and offline.

Participants said that they wanted platforms to enforce their terms of service and reduce online hate and abuse for all users, not just for those who choose to use specific tools. Specifically, they believed that the continued presence of abusive or hateful content was harmful even if they were not exposed to it. This is because they felt it contributed to its broader normalisation, which they said seeps into unavoidable offline behaviours.

Participants told us that the tools that exist today, such as blocking or muting, do not go far enough to help protect them and their families and friends against online hate and abuse. They said that existing content moderation systems often fail to detect harmful content and that they would like to see more human moderation to better address the issues that they raised.

2. About this report

Background

The user empowerment duties in the Act require Category 1 service providers to offer tools that help adult users manage the content they see.[2] These tools should reduce the likelihood of users encountering certain categories of content or alert the user to its presence.[3] This applies to legal content that incites hatred or is abusive. It includes content targeted at the characteristics of race, religion, sex, sexual orientation, disability, and gender reassignment.[4]

To help service providers comply with their duties, Ofcom will publish:

- guidance explaining the types of content that fall under these rules and those that do not, giving examples of each;[5]
- content assessment guidance, that in part explains how service providers should assess the likelihood of adult users with a certain characteristic or who are members of a certain group encountering content that could affect them;[6] and
- a Code of Practice recommending measures that providers should follow to comply with their duties.[7]

The user empowerment duties were, in part, influenced by the abuse and hate experienced by people in the public eye, including sports professionals.[8] Kick It Out campaigned for the inclusion of duties to reduce online hate and abuse in the Act. To better understand the reasons for Kick It Out's campaign and the harms that prompted it, we partnered with Kick It Out to hear about these issues first-hand.

Our objectives were to understand:

a) the impact of online hate and abuse on the sports sector;
b) why some online hate and abuse is particularly harmful; and
c) participants' experiences of tools designed to mitigate harm.

Interviews were carried out from August to November 2024. During interviews, we did not ask participants to differentiate between legal and illegal online content as they may have found it difficult to distinguish between them. In addition, some participants described broader experiences of being abused online where the abuse didn't involve reference to any specific characteristic.

The Act also requires providers of internet services that are likely to be accessed by children to protect them from certain categories of content. Some of the categories of content that children must be protected from are described similarly to the categories in the user empowerment duties.[9]

---

[2] Category 1 service providers refers to providers of user-to-user services that meet the threshold conditions specified in regulations made by the Secretary of State under Schedule 11 to the Act.
[3] Section 15(2) and (3) of the Act.
[4] Section 16(4) of the Act.
[5] Section 53(2) of the Act.
[6] Section 52(1) of the Act.
[7] Section 41(3) and (10)(c) of the Act.
[8] New plans to protect people from anonymous trolls online - GOV.UK, 25 February 2022.
[9] The Act sets out that content promoting suicide, self-harm or eating disorders as primary priority content harmful to children. It also sets out that content that is abusive or incites hatred against listed characteristics is priority content harmful to children.

Ofcom has already consulted on and published its Statement for the protection of children online, which includes one volume containing examples of the categories of content that must be restricted for child users.[10] Therefore, during our engagement with participants, we discussed examples of both in-scope and out-of-scope content that was included in the guidance on protecting children from harms online.

Our approach

This report is based on seven individual interviews and nine group discussions. Participants were introduced to the project by Kick It Out and came from a range of backgrounds, including:

- professional sportspeople who are now retired
- on-screen and off-screen individuals from broadcasters[11]
- management and support functions
- sports clubs and governing bodies
- online safety specialists.

The sample was purposely selected rather than representative. People were included because they had experiences of online hate and abuse that they wanted to share. Therefore, this report focuses on different things that have happened or that have been experienced, but we are unable to comment on the prevalence of these among the wider population.

Experienced researchers, working as independent contractors for Ofcom, conducted individual interviews and group discussions. Their expertise ensured they understood the relevant provisions in the Act. The research sessions followed a semi-structured discussion guide and were audio recorded and transcribed to ensure accurate analysis.

We hosted hybrid workshops with participants before and after the main research sessions. The first workshop introduced the project, ensured that people were comfortable with the process, and explained what to expect. The second workshop shared early findings and gathered feedback to shape this report.

Please note that the remainder of report discusses experiences and impacts of online hate and abuse, including a brief discussion of disordered eating, depression, and suicide.

[10] Protecting children from harms online - Volume 3: The causes and impacts of online harms to children, Chapter 8.6 Guidance on abuse and hate content, pg. 324

[11] The broadcasters approached by Ofcom and Kick It Out wanted to widen the discussion to include professionals who deal with other on-screen talent (including news and reality TV), so some participants talked about their broader perspective.

3. How online hate and abuse leads to harm

What people are seeing online

This section explains the types of harmful content participants encounter, and then discusses its impact. The most frequently mentioned examples were:

- Sexism, misogyny, sexualisation, and double standards for women in sport

> *"A female TV sports presenter will get horrendous amounts of abuse [online], often just about what she's wearing. People will say things like call her a 'slut' and that she's wearing her outfit to try and attract male attention...that kind of stuff. All just based on what she's wearing on TV."*

- Racism

> *"Racist abuse online has only continued to increase because people know the N-word is almost certainly going to get removed [by platforms]. But people find other ways, like if they put a monkey emoji on there then it will get through."*

- Religious intolerance of several different religions in various forms, mainly:
    - Personal attacks based on religion – comments that target an individual's religious beliefs to personally attack or discredit them.
    - Assumptions about a person's religion leading to abuse – sending abusive content (such as Islamophobic comments) to an individual based on perceived religious beliefs.
    - False accusations and online 'pile-ons' – spreading false or provocative accusations that incite others to accuse an individual of religious intolerance – which escalate and prolong the negativity directed at the individual online.

- Transphobic and homophobic content

> *"So, a female [sportsperson] receives a message early evening on a Friday [before she is to compete the next day]. It says, 'You're f*cking useless, you ugly f*cking bitch. Can you not just go and do some f*cking lesbian parade, you bitch?! I wish you get injured and never compete again. You and that stupid lesbian partner of yours will die!'. And then it [the abuse] goes on with more clear and implicit threats."*

- Ableist content such as dismissing an individual's disability as exaggerated, assuming an individual is not working, or describing 'inspiration porn'.[12]

[12] Inspiration porn objectifies disabled people for the benefit of nondisabled people. The purpose is to inspire the nondisabled person or to motivate them – for example, "Well, however bad my life is, it could be worse. I could be that person." (definition abbreviated from What is Inspiration Porn?).

The factors that increase the likelihood of harm

This section explains how several factors contribute to risk of harm. Participants discussed a range of factors that can increase the impact of online hate and abuse. These can be grouped into five themes:

1. The volume and cumulative impact of online hate and abuse
2. Who online hate and abuse comes from
3. The form online hate and abuse takes
4. The timing and duration of online hate and abuse
5. The target or subject of online hate and abuse

Each of these factors is discussed further in this section. What is striking is that every story is different. While some factors are almost always likely to contribute to harm, the harm is often context dependent. For some people, anonymous trolls cause the most harm, whereas for others, knowing who sent the hate or abuse can be more concerning. This means that every person's experience of online harm was unique. We will now examine each of these factors in more depth.

The volume and cumulative impact of online hate and abuse

Once a hateful or abusive comment is posted, it can often be followed by hundreds or thousands of additional harmful comments. From these scenarios, participants said that they felt like everyone was 'piling on', which they found particularly stressful. Several said that the sheer volume of abusive comments was sometimes overwhelming. In addition to abusive comments themselves, seeing the number of likes or support that such a comment received could add to participants' perceptions that everyone agreed.

> *"I was very, very scared, you know, I didn't leave my house for a week because the impact of online abuse, the sort of wave [of intensity] and the amount of people that are abusing you. And then the media writes about it and then it becomes this sort of overwhelming feeling of just dread that so many people are saying such horrible things about you, without you actually having done anything, you know?"*

While some participants described the volume of comments spiking after a particular event or point in time, others described the cumulative impact of seeing a lot of hate and abuse over an extended period. This often led to participants feeling numb or desensitised to the severity of the content. For some participants, this was brought into sharp resolution when family members or friends commented on what they were seeing, not least because the content that was public was sometimes far less extreme than what was being sent to the participant privately.

> *"[Family member] would see comments on like a [broadcaster] video on Instagram about me, and [they said], 'Shall I message this person, or shall I find his wife and message his wife?'. And I was like...'don't do that'. But those weren't really the bad comments either. They didn't see the really crazy messages that I was getting…kind of like pervy comments, but they weren't really disgusting, whereas the ones in my DMs were disgusting."*

Participants also reported having the sense that once something was posted, it would always be there, especially if it went viral at the time of posting. Once the initial flurry had died down, we were told that hateful or abusive content can continue to resurface months or years later and lead to further harm.

However, it is important to note that a single message or post can also have a significant impact, particularly if it intentionally or unintentionally plays on a specific vulnerability of the individual. We were told examples that included making reference to sensitive aspects of a person's family life.

Case Study: *On-screen talent experience high volumes of abuse*

People working for several UK broadcasters took part in the research. This case study builds on the comments of on-screen talent and people who support them.

Broadcasters see an increasing volume of online abuse directed at on-screen talent, some of whom they say are particularly vulnerable to the impact of such content because of being new to being in the public eye. In some sports, new female presenters may initially be the target of increased abuse before they 'prove' themselves in the role.

Some potential presenters and pundits have actively chosen not to go on screen due to fears about potential online abuse. For the people supporting talent, the job of managing huge volumes of abusive content across social media can lead to feelings of frustration and helplessness. They have to take steps such as reporting content, blocking or muting people, and removing comments where that function is available – which they see as a necessary job. They do this as they do not want their accounts to be associated with hate and abuse. Both people on screen and people in supporting roles described developing coping mechanisms, such as making jokes about the worst comments they receive. However, they realise this is not necessarily healthy.

"Then I kind of, it sounds horrific, but I kind of got used to it. By the end of, like, the Olympics, it was kind of like a funny joke, with me and some of my friends in the team, which is probably a coping mechanism, which is kind of depressing."

Seeing people on the television makes the public more likely to comment. For example, after major live sporting events such as the Olympics or the UEFA Champions League there is a spike in the volume of comments, including hate and abuse. Typical comments include racist, Islamophobic and sexist abuse, such as claiming that a person is unqualified to do their job because they are a woman or have a foreign accent. Female presenters describe receiving large quantities of 'disgusting' comments about how they look, while threats of sexual violence are common.

People supporting on-screen talent in sports report challenges that are echoed by those we spoke to working with on-screen talent in reality TV and entertainment. This highlights the widespread and serious nature of online abuse. Typically, sportspeople and on-screen talent in sports have personal social media accounts instead of, or alongside, corporate accounts. This can make it difficult for the broadcaster to provide them with support, and there is limited information available about how to do this effectively.

"People that want to take part in...like a panel show. They come in, they showcase their knowledge. They are very, very clever people. They are subjected to a lot of online abuse and threats. And they have only been in a 60-minute programme."

Equally, the people responsible for operating corporate accounts do not want to make the situation worse for these individual public figures in sport. Therefore, they intentionally do not link to or tag on-screen talent or sportspeople in stories which might be controversial, even though they would like to (for example, a story about a gay couple having a baby). They describe having a sense for which stories might be most likely to lead to hate and abuse.

Some on-screen talent tried to use platform tools to address abuse but typically found that they did not noticeably reduce the volume of hate and abuse directed towards them, either in general or from a specific individual. For example, participants found that people intent on sending abusive messages would find ways to bypass restrictions, making tools like blocking ineffective. Furthermore, they received so many comments that it was not possible to report everything they saw.

"One of the other issues with blocking I know we haven't really talked about is verification. There's so many [unverified accounts], like people will just make new accounts and carry on the abuse. And there's kind of multiple accounts and also, you know, fake."

Who online hate and abuse comes from

Participants told us that the hate and abuse they receive came from a range of different people, from people they know, to complete strangers, to unknown individuals using anonymous accounts. Participants had different views on anonymous accounts. For some, hate and abuse from anonymous accounts was worse because it was impossible to hold the person to account. This was typically the view of those who were being proactive in trying to reduce the volume of hate and abuse they were receiving. However, other participants were more concerned when the accounts were not anonymous. Some felt that hate and abuse from non-anonymous accounts was in some ways worse because it demonstrated the normalisation of hateful and abusive behaviour. They believed that a person who does not hide their identity must either think their conduct was acceptable or feel no shame in being associated with it.

Participants said that certain people online can have a disproportionate impact on the volume and tone of hate and abuse. For example, they were concerned that abusive comments directed towards them from accounts with large numbers of followers could reach a much wider audience. This had the potential to fuel the hate and abuse they received. Some of these large accounts were owned by private individuals posting in a personal capacity. Others were individuals posting in a professional capacity, such as sportspeople, managers, coaches, commentators, or journalists. Sometimes participants believed that such accounts were intentionally stirring up hate. They thought this should result in sanctions, such as the person being banned from the platform.

"What happens is when somebody, like, [an individual with lots of followers] starts to say these things it green lights many, many other accounts to do the same."

On other occasions, participants thought that misinformation and/or inaccurate reporting around contentious topics could act as a further trigger for online abuse, escalating the volume and intensity of it. Participants believed that people with large online followings should be required to fact-check more carefully to reduce hate caused by misinformation. They thought this would help prevent large accounts from spreading and legitimising hate. They also pointed out that it would align with longstanding broadcasting sector rules on verifying information before publication.

Many participants mentioned the growing impact of online hate and abuse they received from gamblers seeking to influence the outcome of a competition by distracting or intimidating sportspeople. They said that such tactics had become more finely tuned, strategic, and premeditated to maximise personal impact. Other participants expressed the hurt they felt in cases where most of the hate and abuse they received came from their own fans.

Case Study - Azeem Rafiq: *We have had to change the way we live. It's like you're not in control of your life anymore.*

Azeem Rafiq is a former English cricketer who played professionally for Yorkshire County Cricket Club.

I played county cricket, so I was used to the odd bit of criticism, but nothing prepared me for the volume and nature of the abuse I received on social media when I spoke up about the racism I'd suffered as a British Muslim playing for Yorkshire Cricket Club. The impact of this experience on me as a human being and on my mental health has damaged my life to such an extent, I'm not sure I'll ever be able to quantify it.

People would post on social media that I was a *"Dirty P-word"*; that I should *"F*ck off back to Pakistan"* and that *"All Muslims are bombers"*. And then there were posts that linked me to negative stories that had nothing to do with me, simply because I was a Muslim with brown skin, like the stories in the news at the time about grooming gangs in the North of England. I felt overwhelmed, as if the abuse was coming at me from all sides – and when people with their own large online followings piled on to criticise me, it led to me receiving even more racist and Islamophobic posts.

A key thing that made it worse was the constant drip-drip [of abuse] over days, weeks and months. It was all I could think about; it chipped away at me constantly. And over time, that wears you down. The abuse left me feeling incredibly paranoid, at times, and often made me question my sanity. I have had my wife ring me saying someone was watching our house; and that my shop had been attacked. I've also been abused on the street. So, we've suffered real world impacts from what people post online and I've had to change the way we [as a family] live. It got to the point that we no longer felt safe in the UK, so I made the difficult decision to resettle with my family in Dubai.

At first, I did not block anyone because I felt that would make it look like they were winning. But it got to the point where I had to start engaging with tools with the help of a third-party firm. I soon discovered that the tools were limiting my engagement with online spaces on all levels, which made me worry I might be cutting myself off from positive interactions and opportunities. It's left me in two minds: I think it's good that people have access to tools, but they alone won't stop people posting abuse online.

The form online hate and abuse takes

Participants said that the specific format of content – such as text, images, or video – did not automatically impact their perceived severity of the harm. Many participants thought that those posting hateful and abusive content were becoming increasingly creative in terms of the formats they used, which included emojis, nicknames, and imagery. They said that it is often harder to report such content and get it removed due to the absence of specific words or phrases. However, they felt that the impact of this kind of content was no less damaging than other forms of abusive content for the person seeing it.

Participants reported receiving hate and abuse through both public and private channels. Both were perceived to be impactful in different ways. Some felt private abuse was worse as it was often severe, while others felt public abuse was worse, as others could see it. They were also concerned such content could have the potential to go viral through re-shares. However, participants also found private messages harmful, especially those containing direct threats to them or their family. Some described their experience of how information flows between online and offline worlds. They

said that abusive or hateful content directed towards them, or towards people with a particular characteristic, online would then be repeated in an offline setting, and vice versa.

Participants said that the impact of hate and abuse can feel different depending on whether it is either: a response to something that the individual posted – and therefore has some control over; or whether it has been instigated by a third party using platform tools like tagging other users, hashtags, or features which allow creating and sharing lists of other users – over which the individual has little to no control.

Loss of control

Online hate and abuse can often feel inescapable and overwhelming, particularly if the individual being targeted did not start the discussion.

People can trend without posting content (such as when a third party creates a hashtag), or can receive unsolicited content (for example, if they are added to a list). They can also lose control of original content if it is copied or screenshotted and shared by a third party, or if a clip or photo becomes a meme. In some cases, this could result in hate and abuse related to the original content. The content may also be subject to manipulation, which can potentially make it more harmful. Some participants talked about the impact of false accusations, or secrets that are guessed and shared, which can be particularly difficult for the individual involved.

Participants contrasted this with their experiences of online hate and abuse in response to their own online content. While the comments or replies might be equally hateful or abusive, users had the option to turn off comments or to take the content down. As a result, they found that content over which they had no control, had the potential to be more harmful.

The timing and duration of online hate and abuse

Online hate and abuse can range from an isolated post, comment, or message to a more prolonged campaign. It may be unprompted or may be in response to a specific incident or scheduled sports event. We were told that the duration of hate or abuse related to a specific incident or event could vary, but whether it lasts a few hours or a few weeks, it can feel all-encompassing for the individual. Furthermore, the impact of the content can be lasting, whether or not it remains online. Some participants said that the impact could be particularly harmful if the content involved a meme or clip that people would continue to refer back to, as it would endure long past the original incident and they would continue to be reminded of it.

Some participants felt they could predict times when they might see a higher volume of hate and abuse. This was often linked with times when they are in the public eye – for example, before, during and immediately after a match, or when they appeared on television. Generally, participants believed that a larger number of viewers was likely to lead to more hate and abuse.

However, participants noted that there was sometimes an unanticipated reason for an increase in abuse or hate – for example, a high-profile person with a large following saying something controversial about a sportsperson. Equally, there could be another spike if someone then spoke up in that person's defence. It could be particularly difficult for on-screen talent to predict when someone might decide to post something that could lead to a spike in harmful content.

Those working in broadcasting were aware that posting on certain topics would be likely to increase the volume of hate and abuse, and they took this into consideration when deciding what to post and the timing of it.

A couple of participants noticed that late night content could be particularly bad and speculated that this reflected the state of mind of the person posting the content.

The target or subject of online hate and abuse

The content of the hate and abuse and the motivation behind it could also influence how much impact it had on the individual.

Participants listed a wide range of motivations beyond hatred that could have influenced the sender. Specifically, they mentioned frustration, intimidation, the desire to monetise, or to seek some sense of fun or enjoyment. The perceived motive did not always change the impact that the content could have, but participants thought this context was important in illustrating the worsening problem of online hate and abuse over time.

Content often questioned the sportsperson's ability to do their job. In some cases, this had a negative impact on the person's performance, leading to further abuse. In other cases, hate or abuse was the result of misattribution – for example, incorrectly guessing something about the individual, or creating links between stories that had no link. The impact was similar whether this occurred by accident or on purpose and whether the individual had the protected characteristic being targeted or not. For example, a participant mentioned that some people sent them hate and abuse because they were Muslim, and some sent them hate and abuse because they were not. The person's religion was not public knowledge, and they found both types of abuse equally impactful.

For those participants that have chosen to be in the public eye, their families have not. As such, participants considered instances where family members, friends, or colleagues received hateful or abusive content to be particularly harmful. Sometimes this content was completely unprovoked, while at other times it occurred because a member of the participant's family, a friend, or a colleague, sought to defend them. Participants told us of other scenarios where seeing a 'pile-on' of hate and abuse directed towards supportive individuals who had stood up for them online was particularly difficult. They were concerned that these individuals might not have any support in place, particularly if they were members of the public and not high-profile individuals.

Case Study – Wayne Barnes OBE: *A shocking experience that has hurt our family relationships.*

Wayne Barnes OBE is a retired English international rugby union referee.

As an international rugby referee, I didn't expect to always be popular. I accepted that there was always going to be a certain level of criticism towards me on social media about my performance, because it is the job of a referee to make big decisions. I could handle that. The trigger for the abuse online would usually be an incident in a big international game – and especially if a coach had said something that cut through into the media. Fans would jump on that bandwagon and there would be forum posts saying things like *"Wayne Barnes should die!"*, *"Wayne Barnes should be hung!"* or images of me as an effigy being attacked.

But ultimately it wasn't about the impact on me – I felt I could take it. What made it worse was the direct impact over time on my family. I wasn't active on social media, but my wife was, so she would read the bile posted about me first-hand. She then became the subject of the abuse, with people attacking her personally via direct messages to her social media accounts and work email address, or by posting fake and offensive friend requests. The abuse went on for some time - misogynistic language aimed at my wife because of her association with me, including comments like *"You f*cking bitch! You slut! Tell your husband he's sh*t"* and, at times, threats of sexual violence.

I'm in shock that my wife had to deal with that. You don't mind people abusing you, but you want to protect your family.

4. The impact of online hate and abuse in sport

How online hate and abuse shapes people's lives

The impact of online hate and abuse can be significant and affects people in different ways,[13] so it was of no surprise that our participants described their experience of both emotional and physical impacts. However, we also learned that the impact can be more broadly experienced across participants' family, friends, and even colleagues, whether they were personally targeted because of their association with the person with the public profile or not.

The impact of online hate and abuse on how people feel

Participants described a range of different emotional reactions to online hate and abuse. Some participants reflected that while they sometimes claim that online content does not impact them, this may be because they do not want to admit its effects to themselves or others. However, they thought it was only human to be affected in some way if subjected to a barrage of hate or abuse.

The feelings participants described experiencing because of online hate and abuse included:

- **Numbness, resignation and tiredness** – Participants felt that the volume of content they experienced meant they became desensitised to it over time and increasingly accepted it as the norm. However, occasionally something would happen to make them see the content with fresh eyes (such as a family member drawing their attention to it) and they would again realise how unacceptable the content was.
- **Helplessness and depression** – Some participants felt a sense of helplessness or depression because of the relentlessness of encountering hateful and abusive content, and because of the perceived ineffectiveness of current tools to deal with this content.

> *"I kind of feel like I'm numb to it now. But it's just still really depressing when any little opportunity, without any real facts or tangible knowledge, is just used to drive a wedge or make money or whatever it is that they're trying to do; it is hugely depressing."*

- **Questioning self-worth, feelings of self-doubt** – Some participants questioned themselves and their skills after seeing a high quantity of content questioning their abilities. This content could refer to certain characteristics or assumed characteristics, such as their ethnicity or gender. Even if participants knew that they were performing well in their jobs, exposure to a large number of critics undermined their confidence. Equally, if they had a bad day some told us they might view content as a form of self-punishment.
- **Fear, anxiety, paranoia** – Depending on the nature of the content, some participants were concerned that online threats could potentially result in negative real-world outcomes. This could range from fear that random people might follow through when encouraged to harm them, to concerns that a specific person might move their harassment from online to real life.

[13] For example, see The impact of online hate - Ofcom.

> *“Most importantly, I feel incredibly scared of actually what's around me."*

- **Isolation and loneliness** – The quantity of hateful and abusive comments could sometimes drown out positive and supportive online content. In some cases, participants decided to stop using their accounts completely, which led to some of them feeling even more cut off from family, friends, sponsors and supporters.
- **Guilt** – A few participants said that they grappled with feelings of guilt generated by online hate and abuse, typically due to the impact on family and friends. Participants noted that that if they had not posted or if they had not been in that line of work, their family and friends (who did not have a choice) would not have been so exposed.

> *"You want to protect your family. You don't mind people abusing you [a level of criticism comes with the territory] - I can cope with that. I'm sitting there not realising the kind of the extent of what's happened. Like in shock that me and my [partner] had to deal with that."*

- **Overwhelmed** – Some participants said that the volume of online hate and abuse they received felt overwhelming.
- **Anger** – Some participants remained angry at the content they had seen, while others described moving past anger to feeling numb, as referenced previously.

Some participants described using therapy to help support them. Others talked to their employer. However, participants in management positions felt that they were not well-equipped to support people facing these issues.

Participants reflected that, while hate and abuse in sport is not new, the pervasive nature of social media meant it was ‘always on’ and that ‘turning off’ was not perceived to be a viable option.

The impact of online hate and abuse on how people behave

In addition to affecting how people feel, online hate and abuse also had an impact on how they behaved, both online and offline. These impacts were wide-ranging, but broadly related to four areas:

- how people do their jobs
- how people spend time online
- disordered eating, self-harm, and depression
- offline behaviour

The impact of online hate and abuse on how people do their jobs

One of the most notable impacts related to the participants’ approach to their jobs. They told us that the level of online hate and abuse a job might attract could be a deciding factor in whether to take it. Specifically, some participants employed by broadcasters said they had actively chosen not to apply for on-screen roles due to the hate and abuse they feared or expected they would receive.

> *"Some members of my team have turned down opportunities to do things that should be hugely beneficial for them and would be progressive for them in their careers, for fear of the hate that they would receive if they did it… either as young women largely or as people of colour."*

Sportspeople and on-screen presenters also mentioned examples of the different ways experiences with online hate and abuse could negatively impact their performance. Presenters said that they sometimes found themselves over-preparing or self-censoring to try to avoid too much online hate

and abuse. Sportspeople said that online hate and abuse had the potential to undermine their confidence, affecting their performance and/or the outcome of competitions.

> *"So, there's the pre-game abuse targeted on making you make a mistake, lose a point. And then when that has happened and you've lost money for a punter, it comes afterwards because [they say] 'I've lost my money on you.'"*

The impacts of online hate and abuse extend beyond high-profile people in sport and sports broadcasting. Participants who worked on the social media accounts of major sports broadcasters also reported being exposed to high volumes of hate and abuse in response to the content they created as part of their job. Some said this made them question whether they wanted to go to work each day.

The impact of online hate and abuse on time spent online

Participants typically described one of two behaviours when subjected to online hate and abuse: stepping back or engaging more.

Some people would choose to disengage. In the most severe cases they would come off social media entirely or reduce their online presence by not actively posting content. However, being online and being visible is perceived to be a significant part of the job – sponsors and advertisers want to see high levels of online engagement from sportspeople and can even make offers of sponsorship deals through social media channels. Consequently, stepping back was typically considered to be a temporary solution, but not one that could work in the long term for these professions. This was particularly the case for sportspeople and individuals working in broadcasting, both on-screen and off-screen, as they were expected to connect with audiences directly.

> *"Filtering out messages from unverified people has helped, but it has also filtered out these types of interaction that I hold precious, like someone asking if I could help their child with their homework and of those sorts…"*

By contrast, other people described engaging more than they felt was appropriate because of online hate and abuse. This could be because they felt compelled to read and to potentially respond to or challenge the content. Although some believed that engaging with the content was not encouraged, they were not sure why and felt uncomfortable leaving it unchallenged. For others, the compulsion to view more came from a need to either seek positive or negative validation, or a preference for knowing what people were saying about them.

Some people described other impacts, such as reduced engagement with their content and therefore a reduction in how many people might see it, because they felt compelled to turn off comments to reduce the hateful or abusive responses. However, with content that they knew would result in a lot of negative comments (such as a broadcaster posting a story about an LGBTQ+ sportsperson) they might turn off comments by default to prevent this from occurring.

> *"It's altering our behaviour and actions, which is a bigger problem. Um, because realistically we shouldn't be thinking about not posting something because we're worried about negative comments if we think it's the right thing that we should be getting out there."*

The impact of online hate and abuse on disordered eating, self-harm and depression

A couple of participants mentioned having a difficult relationship with food, which they felt could be exacerbated by online hate and abuse. For example, posts which commented on their weight could

exacerbate their stress or influence their decisions about what to eat. Equally some people talked about how they would compulsively look for criticism, which was described by one participant as akin to a form of self-harm. Participants raised concerns that it might only be a matter of time before online hate and abuse becomes a major factor contributing to the suicide of someone in the sector.

The offline impact of online hate and abuse

Some participants described changes to their behaviour offline because of online hate and abuse. For example, some felt that they were now more careful about when they went out. Some people who had experienced peaks of online hate and abuse described feeling the need to barricade themselves inside for fear of who they might meet outside, or to spend time abroad to avoid offline repercussions. There were also incidences of people arranging additional security or approaching the police with requests for protection when online hate and abuse was perceived to be likely to lead to physical violence.

Other participants expressed the belief that what people see online impacts what they say offline. Specifically, they noted that the normalisation of hate and abuse online was increasingly visible in real-world conversations.

> *"I'm starting to hear echoes of those views in conversations I have with people, even in this building, uh, people who are very reasonable people. But because it's so prevalent online, sometimes you do feel like you're having conversations, especially over the [transgender] boxing, and I'm hearing yeah, echoes of what you hear online."*

The cost of legal representation to bring civil or criminal cases against people posting online hate and abuse was another significant offline impact.

How online hate and abuse is changing

The overarching view of participants was that the online hate and abuse they see is getting worse. They described an increase in the quantity of this type of content and how far it went in terms of pushing boundaries of what was permissible.

Most participants named a specific platform which they perceived as having particularly high levels of hate and abuse. They expressed the belief that the site rewarded content that received more engagement and that hateful and abusive content got more engagement. However, they also noted that hate and abuse appeared across most platforms, and many different user-to-user services were named over the course of discussions. While some platforms were perceived to be ‘less bad’, this was typically attributed to the type of users the site attracted rather than to any actions or tools being provided by the platform.

5. Reducing harm from online hate and abuse

After discussing the impact of online hate and abuse, and the factors that could lead to harm, participants in the interviews and groups discussed the tools that could be used to reduce the risk of harm on the platforms they use. This discussion was intentionally wide-ranging and covered everything that participants wanted to see in place to reduce harm. As such, some of the platform tools discussed are outside the scope of the user empowerment duties in the Act.

Participants' experiences of different tools

Participants felt that being online, and being aware of what is being said online, is part of having a successful career in sport, either as a sportsperson or a commentator. As such, they did not feel that they had the luxury of stopping their use of online platforms. They also noted that even if they did step back, the hate and abuse would continue online and their friends and family would be exposed to it.

> *"I get quite frustrated that people are just like, oh, don't read the comments, don't look at your inbox, and I feel like that's not helpful for anyone, really. It's easy for people to say that who haven't experienced it and don't get those messages, but, I'll always read what's in my inbox because sometimes it'll be like a great opportunity from someone or they 'like' a story."*

Participants discussed their experiences using tools to reduce their exposure to online hate and abuse. **Figure 1** shows the tools that were discussed, and each one is explored in more detail in the paragraphs that follow.

Figure 1: Impact of different tools

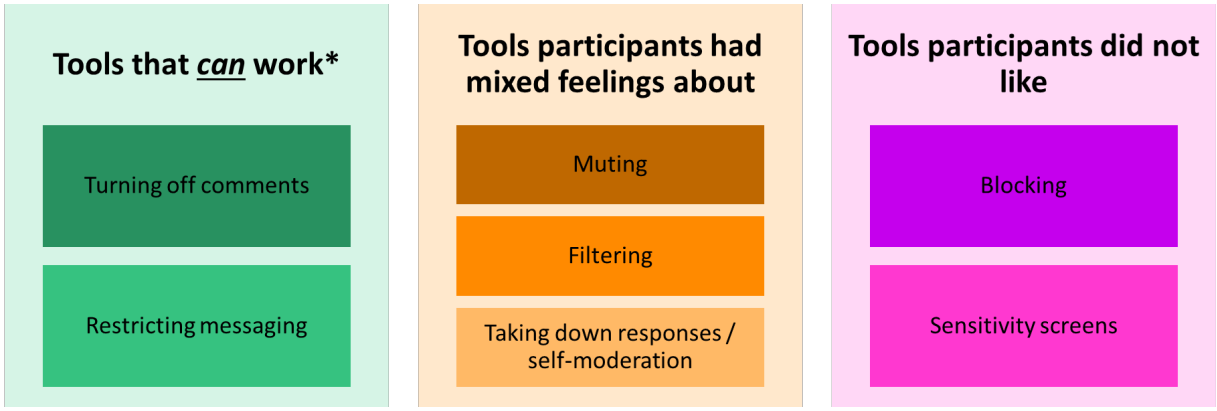

** Tools that are perceived to have the potential to work but do not necessarily do so consistently.*

Tools that participants believed had the potential to work in limited situations

Participants discussed two tools they believed could help manage online hate and abuse. However, many said that these tools were not suited to all situations and negatively impacted a platform's functionality, making them less desirable.

- **Turning off comments** – Some participants reported using this either pre-emptively, if they knew a post was likely to be controversial, such as a broadcaster posting a story about a trans athlete, or later in response to a barrage of negative comments. However, this tool only works where the target of the abuse posted the original content. Importantly, it also reduces engagement and thus the likely reach of the content. Therefore, it was seen to be an essential tool but one to be used sparingly.
- **Restricting messaging** – Participants said that private messages could be more extreme than those posted in public, and some therefore chose to restrict direct messaging. However, while this stopped some messages, participants reported that people would typically find workarounds such as sending emails or targeting other family members. In addition, restricting messaging risks stopping important messages from getting through – one participant mentioned that they had received important sponsorship offers through their social media direct messages.

Some participants mentioned wanting to receive positive messages from family, friends, and fans. They found these tools to be overly blunt because they stopped encouraging and supportive messages coming through.

Case Study: *Online hate and abuse pose huge challenges for broadcasters' show accounts on social media*

People working for several UK broadcasters took part in the research. This case study builds on the comments of those managing TV channels' show accounts on social media.

Show accounts on social media are an important channel for broadcasters. The people looking after these accounts commented on how demoralising the constant exposure to online abuse and hate could be, especially when they were committed to communicating the inclusive values of the broadcaster but knew that some of their audience would push back. The knowledge of how sections of the audience would respond led to difficult decisions having to be made about what to post or when to turn off comments. They thought this limited engagement with some of their messages and stories, a notable example being content relating to LGBTQI+ issues in sport. They were also mindful of the fact that, by enabling fans to express opinions without repercussions, they were potentially encouraging these views to flourish.

"It's like the tribalism. It's the opinions that will kind of increase the level of hate because that then gets served more and more."

Ongoing efforts by individual users or broadcaster social media teams to keep up with and moderate the huge volumes of abusive content posted on the broadcaster's different social channels is an increasingly significant drain on time, resources, and sometimes morale.

"There are other tools [other than reporting], like banned words and the like. And you can create lists of words which can't appear in your comments and emojis as well. But I think the problem with it is they're not foolproof, so obviously people can find creative ways to get around it... just by using spaces and other different things...."

One broadcaster said its sports accounts were appealing to children and young people and that it was concerned about the nature of the content that they were being exposed to in the comments beneath its posts. Some participants believed they had observed a reduction in platforms' efforts to help broadcasters to keep their social media accounts and related comments clear of hate and abuse, with many commenting on how they had noticed processes becoming increasingly automated and unresponsive.

"Nothing seems to be changing and you know you wanna be able to, kind of, call out individuals or see individuals who you know are the worst perpetrators [being] shut down for… you know, spewing this bile or whatever, but it doesn't seem to happen."

Broadcasters spend money on both internal staff and third-party services to help them manage the volume and increasingly varied nature of hate and abuse posted to their social media accounts. However, third-party services take time to set up, are expensive to run, and require effort to maintain due to the constantly evolving language used by those posting abuse and hate online. Ultimately, broadcasters questioned having to pay for such services and believed that platforms should be taking accountability for these issues directly rather than putting the onus on individual users.

"I think our view is that the platform should be taking this responsibility rather than it being on us. And I think there is a quite large cost associated with those, we're talking six figures… And the more comments you get, the more successful we are [as a broadcaster], the more [show] accounts you have, the more expensive it becomes."

Tools that received mixed opinions from participants

Participants told us that being able to turn off comments and restrict messaging on social media profiles could be beneficial because it prevents harmful content from being created in that way. While participants perceived these tools to be very blunt and to have drawbacks, they said that the tools did achieve the ultimate goal of reducing online hate and abuse content.

Participants had mixed feelings about the efficacy of muting, filtering, and moderation tools:

- **Muting** allows a user to completely stop seeing any content from other specific users without them being notified, and without needing to block that user. Some liked it because the sender did not know their content was not being seen. However, others disliked it for the same reason as they were concerned that threats could escalate without them being aware.
- **Filtering** allows a user to proactively block content containing a specific set of words set by the user. Participants had varying degrees of awareness of this tool. Those who were aware of it had mixed views for two reasons. First, although filtering stops the user being exposed to abusive or hateful content, they still know the content is there. It is still being seen by others (such as family, friends, and fans) and could therefore be potentially causing harm to other people. Second, filtering was not seen to be particularly effective as it can be bypassed using simple workarounds. These include misspelling words or using apparently innocuous words or phrases that are hard to block. Participants perceived that filtering worked better when they wanted to avoid content on a particular topic for a time-limited period – such as content about exams when revising – rather than hate and abuse.

"I've found those tools helpful when I'm trying to hide something like people talking about an exam that I'm taking and not wanting to think about it... But… from an online hate perspective, it's just that there are so many things you would have to put in there."

- **Taking down responses or self-moderating** is available on some platforms and allows a user to remove other people's posts related to their content, such as comments on a live stream. Some participants reported doing this directly or using a third party to do so. However, they

felt it put a lot of pressure on the user and could be very resource intensive. Consequently, while relatively effective as a tool, it was not seen as a practical option.

Ultimately, most participants did not want to hide content – they wanted it to not be posted in the first place or for platforms to be accountable for taking it down quickly with minimal burden on the user. Overall, participants did not consider any of the options given here to be particularly desirable.

> *"I don't know how to do it [mute/filter], and I can't be bothered to like learn a new skill, but also because I think it's actually really important, especially as a journalist, to sometimes not sit within your own echo chamber."*

Tools that most participants disliked

These tools were mostly perceived to not be effective. While some participants were using them, most participants said they would not choose to use these tools.

- **Blocking** results in the blocking user and the blocked user being unable to send direct messages to each other or encounter each other's content and severs their connection if they were previously connected. Typically, participants described having tried blocking in the past, but most had stopped using it. This was because they found that blocking somebody was perceived to be an indication that the sender had 'got to' the target. Therefore, being blocked was sometimes perceived to be the goal of certain perpetrators. Participants did not want to give perpetrators the satisfaction of being blocked, especially as they felt platforms were not doing enough to prevent repeat offenders. In other words, it appeared to participants that it was easy for perpetrators to set up a new account or use another platform to continue sending hate and abuse even after being blocked. Therefore, blocking was perceived to be an ineffective measure. Some participants used blocking and had found it to work in some cases. However, these participants were not aware that being blocked might be seen as being an achievement by people posting hate and abuse.
- **Sensitivity screens and content labelling** are tools that typically hide potentially harmful content and give users the option to look at it or label it so that the user knows what to expect. Participants were not necessarily aware of these tools. However, most participants thought that a sensitivity screen would arouse curiosity and increase rather than decrease the likelihood of them engaging with the content. As such, they thought it would not achieve its objective.

In addition to formal tools, some participants described taking time to challenge hateful and abusive content. While this allyship was appreciated by the targets of such content, it was acknowledged that it could also result in more hate and abuse. As well as provoking a reaction which could reignite the original debate, such engagement could potentially teach an algorithm to send more hateful and abusive content for the ally to engage with. Participants were particularly concerned when other people – normally outside of the public eye – came to their defence and experienced a high volume of hate and abuse as a result.

6. What participants wanted to see

Platforms reducing the volume of hate and abuse

Participants were very keen to emphasise the importance of service providers taking more responsibility for the content on their platforms. Currently, they felt that the onus was on them to protect themselves against this kind of content, whereas they believed that platforms should be more proactive in ensuring that they are not enabling, encouraging, or rewarding people who post hate and abuse.

Participants felt that offering content control tools did not necessarily go far enough to address the root cause of the issue – which, in their opinion, is people posting hate and abuse.[14]

In the interviews and groups, participants told us that what they wanted were effective reporting tools to send a clear message that online hate and abuse were unacceptable. They also discussed a range of tools which could empower users in these circumstances.

Participants said that it was important for the default to be that protections were turned on. They believed this would mean that people were automatically protected and therefore would only be exposed to potentially harmful content if they made an active choice. There was limited awareness of some control tools, but participants did not feel that it should be up to the victim of hate or abuse to learn about what was available. They thought that greater priority should be given to the design and delivery of processes for mitigating online harms, as they did not feel that any of the existing tools were particularly effective.

Platforms providing the services currently offered by third-party tools

Some organisations, such as broadcasters or clubs, and individuals, such as sportspeople, had paid for third-party tools or services to help them stay safe online and reduce the online hate and abuse visible to people viewing their social media accounts. These tools or services worked in different ways. Some hid or reported content on the user's behalf, either through manual review by content moderators or automated filtering. Others took an offline route, tracking down perpetrators and sending them letters warning that further activity would result in criminal or civil action against them.

There were mixed views on the effectiveness of third-party moderation services, especially those relying on automated tools, which needed to be regularly updated as trends change quickly. Furthermore, third-party tools and services were viewed as taking time and money to set up, and participants were unclear why they should be responsible to foot this cost. Furthermore, they

[14] The interviewers explained that *"the Online Safety Act requires service providers to remove illegal content from their platforms to protect users from significant harm. Under the 'user empowerment' duties it also requires certain service providers to offer content control tools to adult users which, if applied by the user, will alert them to or reduce the likelihood of them encountering certain categories of content, including legal hate and abuse. Some of these tools exist already and some might be new."*

considered that the existence of these tools proved it was possible to significantly reduce online hate and abuse, but felt that platforms should be offering the service, rather than them having to pay for it.

Case Study: *Protecting followers*

Third-party services reduced harm for the individuals targeted but were also seen as a useful way to protect account followers. For example, a broadcaster mentioned that their accounts often attract a lot of porn bot accounts, which post comments unrelated to the original content in order to advertise different porn sites. They felt this was particularly inappropriate as sports content is often viewed by children, who may then be tempted to view the sites being advertised. They used third-party tools to ensure this content was taken down where possible, but it was difficult to ensure this always worked effectively as the content was constantly changing.

Simpler, more effective reporting tools

Reporting, when combined with sanctions for the person posting the hate or abuse, was frequently mentioned as a way to reduce the prevalence or impact of online hate and abuse. Effective reporting was considered as potentially one of the most powerful tools to stop hate and abuse. In reality, however, participants felt that reporting often did not work for two main reasons:

- **Ineffective classification of content** – participants reported systems not identifying content as hateful or abusive where the individual believed that it was (for example, because they felt the full context or implicit meaning of the content had not been sufficiently considered)
- **Ineffective or non-existent sanctions** – participants believed that reporting people posting hate and abuse did not stop them doing so.

Participants said it was relatively straightforward for people to post content consisting of seemingly innocuous words or phrases, which, when used together, can be perceived as hateful or abusive – for example, suggesting or inferring that women belong in the kitchen. They noted that systems, which they viewed as being increasingly automated, appeared to struggle to identify these cases because they require an understanding of the wider cultural context and history to determine the nature of the content. Consequently, they felt that human moderators were vital to ensure that decisions were appropriate and fair.

Participants' views on what should happen after reporting varied in scope from taking down the single post that was reported to barring the user from the platform. However, the consistent message was that there should be repercussions from posting online hate or abuse so that people would learn that it is not acceptable. Currently, due to perceived issues with classification of content, participants found that even reporting content did not necessarily lead to it being removed.

Participants discussed the different reporting options available to them:

- **Reporting to service provider as an individual user, in a personal capacity** – Participants saw this as a good tool in theory, but in almost all cases perceived not as quick, simple, or effective as they would like. Participants said that it could currently take too many clicks to report a single piece of content and described often getting an immediate or very quick response (that they suspected was automated) saying 'no breach'. Ideally, participants thought it should be possible to report multiple pieces of content in one go where they related to the same issue. Sometimes responses would take longer – but even when they did, communication was not necessarily compassionate. Responses from the service did not

recognise that those who took the time to report content were likely to feel strongly about it and would be upset if the outcome of a complaint is 'no breach'.

- **Reporting to social media platforms by a sports club or company** – Sometimes systems were in place for the social media team at a football club or at a broadcaster to report the abusive or hateful content, rather than the targeted individual. Participants working in these organisations told us that, in the past, they had named contacts at the major platforms who were helpful and would assess and take down content quickly. However, they reported that this was increasingly no longer the case. As a result, employees at these organisations said they were now struggling with time-consuming processes that they perceived to be ineffective.
- **Reporting to the police** – Some participants had tried reporting online hate and abuse to the police, with mixed results. Participants said they had no clear idea why the police did not respond as hoped, but assumed they did not know what to do, or they did not have the time or will to do anything. They also noted that it could be hard to collect the evidence the police required and said they would prefer an option to report content through the platform so that the evidence could be easily submitted directly to law enforcement. Some participants hoped that if the police started taking action against hate and abuse online, this would send a message and reduce the number of people posting such content (as seen in summer 2024 with the riots following the events in Southport, where there was a clear connection between online activity and violent disorder seen on UK streets, resulting in arrests for online offences).[15]

Some participants who were responsible for managing social media accounts for broadcasters expressed frustration with what they saw as a lack of clarity and support from platforms. They said this made the task of supporting individuals by reporting online hate and abuse more difficult. This was illustrated by at least one participant who mentioned having their ability to report content blocked by a service provider for submitting too may reports.

[15] Our immediate action after Southport attacks - GOV.UK: "There have already been hundreds of arrests, including for online offences."

7. Conclusions

It is helpful to bring together the two key elements of this report – the nature of online hate and abuse and the experiences of using tools to navigate this content – to summarise what we have learned. **Figure 2** sets out the factors that participants perceived as exacerbating harm and the tools they knew of that could mitigate harm.[16]

No single tool is likely to be enough on its own due to the many factors and impacts involved. Participants preferred that hateful and abusive content be removed and discouraged, rather than the onus being on individuals to protect themselves. Moreover, they emphasised that tools need to be easy, quick, and effective to use.

Figure 2: Factors in online hate and abuse and the perceived effectiveness of mitigation tools

Factor	Tools
Volume and cumulative impact of online hate and abuse	Turning off comments and restricting messaging can help reduce volume but also has an impact on engagement, which in turn impacts reach. Participants did not describe any user empowerment tools that could mitigate against cumulative impact.
Who online hate and abuse comes from	Blocking or muting users were both discussed as potential tools, but blocking in particular was seen as having limitations because it could be perceived as giving the individuals posting the abusive comments the recognition that they want. Some tools were viewed as being more effective if it meant that other users, in addition to the user being targeted, would not see the abusive comments posted and the senders would not be aware that other users could not see their comments.
The form online hate and abuse takes	The main ask was a tool which adapts quickly to evolving tactics as people try to get around blunt tools such as banning the use of specific words or emojis. Currently, creating and updating lists of banned words or emojis was perceived to be too time intensive. Reporting content was considered the main way to take back control but was seen as being rarely effective or taking too much time, meaning harmful material could become viral before it was removed. Similarly, features that allow creation and sharing of lists of other users can be misused. Users may be added without their consent, and it can be difficult to be removed. Typically, participants felt they had limited control over hashtags relating to them.
The timing and duration of online hate and abuse	Apart from completely stepping back from social media at times when hate and abuse were likely to be particularly bad, participants did not identify any tools that could address the issue of hate and abuse happening at particular times, except possibly using a time-limited individual mute or block function on a live stream or turning off comments entirely, depending on the circumstances.

[16] This information is based on participants' current awareness and does not necessarily reflect the full range of tools that might be available now or in the future.

Factor	Tools
The target or subject of online hate and abuse	Participants said that filtering and sensitivity screens could potentially help but would not protect family and friends from viewing the content or from becoming a target.

This report is a deep dive into the experiences of a particular group of people who have experienced online hate and abuse. It will add to our evidence base alongside previous pieces of work including:

- Ofcom commissioned report on online hateful content:
- Ofcom research exploring the volume of hate and abuse received by footballers; and
- Ofcom research into the impact of exposure to online hate on people with protected characteristics.