

Women and Equalities Committee

Oral evidence: Tackling non-consensual intimate image abuse, HC 336

Wednesday 6 November 2024

Ordered by the House of Commons to be published on 6 November 2024.

[Watch the meeting](#)

Members present: Sarah Owen (Chair); Alex Brewer; David Burton-Sampson; Kirith Entwistle; Natalie Fleet; Catherine Fookes; Christine Jardine; Samantha Niblett; Rachel Taylor.

Questions 1 - 75

Witnesses

[I](#): David Wright, Chief Executive, SWGfL; Sophie Mortimer, Manager, Revenge Porn Helpline.

[II](#): Courtney Gregoire, VP and Chief Digital Safety Officer, Microsoft; Gail Kent, Director of Search for Government Affairs, Google.

Examination of witnesses

Witnesses: David Wright and Sophie Mortimer.

Chair: Good afternoon and welcome to the Women and Equalities Committee. Today we will be looking at non-consensual intimate imagery abuse, following a one-off session held by the previous Committee in the last session. I would like to welcome our first panel: David Wright, chief executive of the South West Grid for Learning and director of the UK Safer Internet Centre—that is a mouthful—and Sophie Mortimer, manager of the revenge porn helpline at the South West Grid for Learning. Thank you both very much for coming.

Q1 **Kirith Entwistle:** Thank you both for your time today. Sophie, how many cases has the helpline dealt with in this and previous years, if you do not mind me asking?

Sophie Mortimer: I do not mind at all. Total direct cases—those that have come by email or phone through to the helpline—is currently at just over 24,000. This year alone, it is at around 3,500. We also have a chatbot to manage capacity, and we have had about 35,000 sessions run through that chatbot.

Q2 **Kirith Entwistle:** What impact do you think non-consensual intimate image abuse has on victims?

Sophie Mortimer: It is enormous. The impact on victims reaches into every corner of their lives: their personal relationships; work relationships; their ability to go about their lives, go to the shop, go to work and engage with their family; and their emotional and psychological wellbeing. And it is ongoing, because if that content cannot be completely removed and continues to circulate online, it is there every day and they are re-traumatised every day.

Q3 **Kirith Entwistle:** My next question is for David. What is your assessment of the police response to non-consensual intimate image abuse?

David Wright: I would like to start by saying thank you for the invitation to come and share evidence and our experience. I am very grateful for that opportunity and, as the Chair says, it follows on from the session we had in May. It is such an important subject, and one that we deal with all day, every day, as you will discover. When I say we, I mean Sophie and the team in terms of the royal “we”.

In terms of police response, there is going to be no better person than Sophie to provide you with that from the experience we have because we spend a lot of time coaching victims who have reported to the police.

Sophie Mortimer: I can give you a statistic. We record people’s responses to how they have been dealt with by the police where they have already gone to the police before we speak to them. Of the people



who went to the police, four times as many reported a negative experience of reporting as opposed to a positive one, and that largely comes down to the officer or the call handler—the first person they have interacted with—around their understanding of the law as it is written in both its incarnations, how it might apply to their case, and what can be done to support them. The sense that many victims come away with is that there is no help, that their content cannot be removed from the internet, that perpetrators are rarely prosecuted and even more rarely convicted, and that there will be no just outcome for them.

Q4 Chair: Are there any forces that are very good and have had a more positive response to victims and others that would need a lot more improvement?

Sophie Mortimer: It would be really hard for us to say because the nature of the offence is that people do not often share demographic information that widely, so it might be unfair to name a particular force either way.

David Wright: In terms of the increase in caseload, Sophie has articulated the volume that we are dealing with at the moment, but something we have experienced over the last four years is a continued, massive increase in cases arriving or people reaching out. To give you some numbers, in 2019 we managed 1,600 cases. In 2020, largely fuelled by covid—I think that is fair to say—that rose to 3,200 cases, to 4,400 cases in 2021. In 2022 it was 8,900, and last year it was just under 19,000 cases. Over a tenfold increase in four years is terrifying.

As we project forward, we have no idea how big the problem is. We are just discovering more of this iceberg; we do not know how big this iceberg is, but at that rate, we are talking about hundreds of thousands of cases a year.

Q5 Kirith Entwistle: I would like to know a little more about what sextortion is and who it primarily affects, from your experience.

Sophie Mortimer: Sextortion is a very specific form of behaviour that comes through to the helpline, usually carried out by crime gangs based overseas, targeting predominantly men. This year, 91% of the victims of sextortion who have come to us are male. They meet somebody online, and it moves quite rapidly to sexual activity, which is recorded. Friends or contact lists have already been captured, and then there are threats to share unless they are paid money.

Q6 Kirith Entwistle: In your view, why is it a growing problem?

Sophie Mortimer: Because it works—it is as simple as that. It is very difficult for law enforcement to take any steps because the gangs are based overseas.

Q7 Catherine Fookes: David, can you explain how the StopNCII hashing tool works?



David Wright: Indeed I can. StopNCII.org is a website; it is a piece of technology that we built over a period. It stands as technology for victims who are being threatened to have their intimate images posted online. A great example is sextortion, as we have just been discussing. If somebody is threatening to post images online, we know the impact on those individuals of loss of control can be as impactful as actually having content shared as well.

StopNCII was created essentially for that. It supports adults anywhere in the world who are being threatened with having their intimate images shared online. As a victim—as an individual—you visit the website and create what is called a hash: a digital fingerprint; a unique code that identifies the image or the video you have on your device.

The unique thing about StopNCII, which just serves adults, is that in terms of privacy preservation, the image and the video always stay on your device. You do not share that image or video as an individual, which is clearly an important part here. When we were constructing the technology, we collected no data and no information. We did not want people to be second-guessing, “How much data am I giving away here?” So, there is no logging into the site and there are no cookies or collection of data. You create this hash—this fingerprint—that gets added on to the dataset and is then shared with participating platforms to enable them to identify and help to prevent anybody anywhere in the world who might try to post that particular image on their website.

We are using well-trodden technology. Hashing technology is nothing new; it has been used to identify and combat child abuse content for many years. PhotoDNA is 15 years old; you are going to hear from Courtney shortly on that. These are well-trodden technologies that we are trying to utilise and make it easier for platforms to engage.

Q8 Catherine Fookes: How effective has it been at preventing images appearing?

David Wright: Over 970,000 hashes have been created by individuals over the nearly three years since StopNCII began. When a platform identifies a match—that somebody is trying to upload that image—it invokes their normal moderation processes, whether automated or human. They then get to view the image or the video.

If it is NCII, they will then add a tag on to that hash, which then goes back into the system. The hash does two things: first, it provides information to all the other platforms that this is now a verified hash; and, secondly, it indicates to us how many times StopNCII has actively worked. Currently, we are working at about 22,000 instances.

Q9 Catherine Fookes: Has every platform signed up, though?

David Wright: Today, we have 12 platforms signed up, which are published on the website. They are platforms and companies with which we have data-sharing agreements, and which have integrated StopNCII.



HOUSE OF COMMONS

There are another 12 or 13 with which we have data-sharing agreements and are currently onboarding StopNCII.

It is important for us that there is no opportunity for misunderstanding when somebody uses StopNCII. We know the platforms listed on the website are taking the hashes and using them. We do not want there to be any opportunity for misunderstanding for somebody creating a hash, which is why we have this dual-step process.

Q10 Catherine Fookes: If the hashed image or video is edited in some way, will it still be recognised by your tool and stopped?

David Wright: There are two answers to that question. For image hashing, we use perceptual hashing; I will come back to that. For video hashing, we use cryptographic hashing. With perceptual hashing, there is a degree of tolerance. The perceptual hashing is PhotoDNA and PDQ, so two forms of hashing technologies. As I said, PhotoDNA is well-trodden from a child abuse content perspective and allows some form of tolerance in terms of image manipulation. With cryptographic hashing for video, there is no tolerance; if you change one pixel, it will create a different hash. We always recommend to everybody—particularly if you have a video—to create hashes of every single different version of the same video for exactly that reason.

Q11 Chair: David, can I just come back to the question that Catherine asked quite rightly about the different platforms and search engines? You said they are listed on your site; they are also listed in our report. There is one notable exception from that list: Google. Google had previously told our Committee that it had not partnered with StopNCII, "Due to specific policy and practical concerns about the interoperability of the database, which have been communicated to SWGfL." We understand that dialogue between Google and SWGfL remains ongoing; I just wondered if there was any update or progress there with Google.

David Wright: We are making progress. It is not perhaps as quick as we might have liked, but there is progress and forward trajectory. We may have an update at the next witness event.

Q12 Samantha Niblett: I have a little additional questioning about the hashing so that I understand it correctly. It is a prevention tool, so somebody has to proactively be at threat of their image going online, and then they go on and create this hash? I am assuming they have already engaged with your organisation and are already feeling a sense of threat, or might have had it happen before and there is another photo that is about to go online.

David Wright: It is absolutely about prevention, yes. If the content is already online, there are likely to be quicker ways of getting it removed by directly reporting to the platform. If the content is already there, we would always suggest doing both because it would prevent further propagation of the same content. Certainly, if it has been posted, we



HOUSE OF COMMONS

would always suggest you directly report that content to the platform as well.

Q13 Rachel Taylor: What do you think are the barriers stopping those that have not signed up? Do you think we should move to a situation where Ofcom requires them to sign up?

David Wright: Thank you for that question. We are expending a lot of effort in engaging platforms. We have 12, plus another 12 or 13, which is great, but there should be thousands. The more platforms that engage with StopNCII, the more fear we can help to discharge from victims; there is a direct correlation.

Why is it hard? We have made every effort through the technology we are using to make it as easy as we possibly can by reducing the engineering lift in terms of onboarding and reducing the legal lift through our templates and data-sharing agreements. NCII of adults remains non-regulatory at the moment. Within companies, when this issue is competing with regulated issues, I fully expect that is what will take precedence in terms of their priority, and quite rightly, which is why we are trying to make it as easy as we possibly can.

At the last evidence hearing we were joined by Keily Blair, the CEO of OnlyFans. She very helpfully talked about their onboarding process and about it being a total of 80 hours of engineering time to integrate StopNCII, which is nothing in the grand scheme of things. It is helpful to give some confidence to platforms. We are trying to make it as easy as we can.

The answer to the second question about Ofcom is yes, and for the same reasons. That is the top and bottom of it. If it becomes an obligation, it becomes prioritised within platforms, and there is a re-orientation and prioritisation of resources.

Q14 Chair: Sophie, in practical terms, what difference would it make if Google signed up?

Sophie Mortimer: It would make the searching and locating of somebody's content so much more difficult because we know that Google is over 90% of the search traffic. Sadly, too many clients come to us whose content is shared directly under their full names and/or where they live— very identifiable information. If Google signed up, that would stop a lot of that in its tracks and would stop the continuous cycle of downloading and resharing it at another date.

Q15 Christine Jardine: I would just refer the panel to my entry in the Register of Members' Interests. David, are the Online Safety Act 2023 and Ofcom's related powers enough to require platforms with a UK presence to proactively remove NCII content?

David Wright: I have a number of responses to that. We are great advocates and supporters of the Online Safety Act 2023. However, for



this particular issue, from an NCII perspective, other than removing the need to prove intent, which clearly helps with prosecution and that is a positive thing, it does not actually do anything else in terms of classifying the content.

When we report content—again, the royal “we”—we have a 90-plus per cent. takedown rate, which is great, but the flip side is that just less than 10% of the content we report is still online. If we have platforms that will not take content down, usually because they are outside Europe, we have no other mechanism of disrupting or blocking access to that content like we do with other forms of illegal content. A big example is when a particular perpetrator was sentenced to 32 years in December 2021. He had over 200 victims, and we had reported 162,000 images that he had extorted from these victims. Today, there are still 4,000 of those images online and accessible within the UK. I am being very careful here to try not to provide too much detail around that for fear and thoughts of those victims.

We have made attempts to use the same mechanisms to block access to that content through ISPs, but the content is not illegal.

Q16 Christine Jardine: You have answered my next question to a certain extent, which was going to be to what extent do UK-based platforms take the content down? You are saying it is about 90%. Without naming names, are there platforms that are more effective than others that you can rely on more to take the content down?

David Wright: We have reported to over 1,000 different platforms over the years, to which we have had various responses. I am going to defer to Sophie, who has first-hand experience.

Sophie Mortimer: As David said, we report to over 1,000 different platforms. Much of this content that continues to circulate is on platforms beyond our reach, whose business model is entirely constructed around the sharing and resharing of this content. They have no interest in engaging with us because they do not have to and, sadly, they are also beyond the reach of Ofcom or any other authority to take action. The content is still there, and all it takes is a very small number of images to still be there and recirculate and then you are back to square one.

Q17 Christine Jardine: David, what difference would it make if adult non-consensual intimate images were classed as illegal content in the way that child sexual abuse material is?

David Wright: That is exactly what we want: for the NCII content to be treated in the same way as other illegal content—child abuse content, terrorist content—which will then enable a number of things, not least internet service providers to block access from a UK perspective to this content. Even in that regard, thinking about the victims whose content we have made every effort to take down but is still online, there is some comfort in knowing that it is not viewable or there are limitations on its



HOUSE OF COMMONS

viewability. We do not have that ability today. In fact, ISPs are telling us the content is not illegal and that it runs the risk of censoring the internet to non-illegal content if they do block access to it, which has to be wrong.

Q18 Christine Jardine: Would it be accurate to say that classing it as illegal content could be critical?

David Wright: It needs to be treated in the same way as child abuse and terrorist content. In terms of how that is articulated, again, there are forthcoming evidence hearings with lawyers, particularly Professors Lorna Woods and Clare McGlynn, I think. I would always defer to them in terms of how this would be classified from a legal and regulatory perspective, but we need it to be treated in the same way. The victims need access to the same instruments and protection to that content that children quite rightly enjoy from their image perspective.

Q19 Christine Jardine: Are perpetrators who are convicted of NCII abuse required to hand over or delete the content?

Sophie Mortimer: No, they are not. We had a case quite recently where somebody came to the helpline who had reported the sharing of her intimate images to the police. There had been a criminal process, he was convicted and given a suspended sentence, and the police rang her up and said, "Now that the case is concluded, we have to give him his devices back with all your content still on it." This was content that was not only non-consensually shared but was voyeuristically obtained. They said, "Our hands are tied," and she rang us to say, "This cannot possibly be right."

We did some investigation, but any moment that existed to do anything about it had passed because the court process had been completed. The orders available to magistrates and judges to deal with this are not designed for devices and the content on the device or additional forms of storage such as iCloud storage. The mechanisms have not been put in place or utilised. Again, I would refer to the legal experts who will give evidence in due course, but it is not happening routinely, and too often we have multiple cases of content that is still freely available to those perpetrators.

Q20 Chair: It will be shocking to most people to know that a perpetrator who has been convicted will be handed back all their material, possibly to do the same thing again. If an adult NCII was made illegal in the same way child sexual abuse material is, would we see that happen? Would perpetrators be given access to the same material? How would it impact the speed at which you are able to block content and take images offline?

Sophie Mortimer: To answer your first question, it would stop dead. I feel confident that would not happen if that individual who came to us was a minor.

In terms of going forward, it would be hugely powerful. We have partnerships; we are part of the UK Safer Internet Centre and we work



HOUSE OF COMMONS

closely with the Internet Watch Foundation, which is already doing this work. I have no doubt they would partner with us to make that a reality because they are already doing it and, ironically, we are in a position where the mechanism is the bit we can do, but we lack the legal ability to take that step.

Chair: There are many elements around NCII, one of which is cultural sensitivity, and Rachel is going to ask questions on this.

Q21 Rachel Taylor: David, to what extent do you think non-consensual images that are not necessarily sexual but may be culturally insensitive are a concern? For example, a woman not wearing a hijab when she ordinarily does.

Sophie Mortimer: I will start because I have some data about that. So far this year, 1% of our cases have been of culturally sensitive content. We have been tracking additional impacts this year, and 7.5% of the people who have told us about additional impacts have reported there is a cultural sensitivity for them. That tends to mean they are then open to further harassment and, of course, honour-based abuse, which is potentially life-threatening in some cases. It is an issue that is constantly raised in our work with StopNCII, where we work with over 100 NGOs around the world. Definitions of intimacy vary, and it is something we do need to be much more sensitive to.

David Wright: When letters were written by the Committee to 12 or 13 different companies about StopNCII back in March or April, one of the responses from Pinterest, which is clearly published in terms of the correspondence, suggested that it did not need to partner with StopNCII because it has perfectly adequate and effective sexually explicit filters.

This point gives us the opportunity to highlight that this is not necessarily about sexually explicit content. As Sophie said, the impact in some cases can be just as catastrophic if no nudity or explicit content is involved. It was a helpful response because it was a realisation from our perspective that platforms are connecting this issue with nudity, which is not necessarily the case.

Q22 Rachel Taylor: You are saying those concerns are not really being addressed by the major platforms. Do you think they are being addressed by legislation?

Sophie Mortimer: I would say they are not. To your example of a woman without a hijab where there will be an expectation that she will be wearing one, that image would not necessarily be classed as an intimate image, certainly not under the legislation that the Revenge Porn Helpline works under. You have to work with the tools you have, so we might be able to look at it in terms of harassment through our sister service, Report Harmful Content. However, that is more difficult and not necessarily as straightforward around that sort of content because it requires even more context and reporting platforms to industry



HOUSE OF COMMONS

platforms, so we need to give more information, which slows the process down.

Q23 **Chair:** One per cent. is still quite a large number of people. Is it growing?

Sophie Mortimer: Yes, I would say it is.

Q24 **Chair:** The other question I wanted to ask is about reporting. Many people who see that would probably take the same viewpoint as Pinterest has just displayed, which is wrong, but they would think, "Why would I go to the Revenge Porn Helpline if it is not an overly sexualised image?" Are you working with any other stakeholders that are having culturally intimate and sensitive image abuse reported to them and are not coming necessarily directly to you?

Sophie Mortimer: Yes. We work quite closely with both Karma Nirvana and the Muslim Women's Network, with which we have done reciprocal training and have good reporting routes. We support each other. We can do so much, but they can give that additional support and vice versa.

Q25 **Chair:** David, OnlyFans' CEO gave evidence to the Committee and said it was a lack of understanding about what intimacy is. Are you seeing any improvement with other companies and platforms about their cultural sensitivity?

David Wright: No. I would like to think this process will exactly highlight this as an issue and draw that out. Ultimately, I would like to be optimistic, but I have no evidence to suggest that it has changed at all.

Q26 **Chair:** Do you know any platforms that have policies specifically on culturally intimate images?

Sophie Mortimer: The Meta platforms do.

Q27 **David Burton-Sampson:** Sophie, how many cases of deepfakes do you encounter coming through the Revenge Porn Helpline? Maybe you can give us an idea over a period of time as well so we can see the trends.

Sophie Mortimer: I do not have specific numbers for this year, but I can certainly follow up. I can tell you that 72% of the reported cases were of women. My handy assistant has just handed me some statistics.

David Wright: As if by magic.

Sophie Mortimer: We have seen a 400% increase in cases between 2017 and 2024. The technology is certainly becoming much better and much more easily available, and it is much harder to identify. The behaviour around the creation of what we have been referring to as sexually synthetic content is somewhat different, and it is more in common with a behaviour for genuine content that we call collector culture, which is not about causing direct harm to a particular individual; it is about possessing, creating and sharing this sort of content peer to peer. So there is more similarity with synthetic content there.



Of the 72% of women who reported this content to us, 44% reported that the perpetrator was a known male, and 53% were completely unknown, which speaks to these different drivers where sometimes it is targeted at a particular individual, but very commonly it is not. So my point is they simply do not know.

Q28 David Burton-Sampson: What action would you like to see major platforms such as Google and Microsoft take to counter the threat posed by NCII, including deepfakes?

Sophie Mortimer: They need to be able to identify it and flag it because it is really important that people can see the difference. Somebody who has been affected by synthetic content can say until they are blue in the face that it is not them and people will not necessarily believe them. We need that creation of the registry so we can start to deal with this content in the same way we do where it is genuine content.

In terms of that identification, we had a case fairly recently of someone who was being harassed by being sent synthetic content of herself. When she reported it to the police, they said, "That's not you; that's someone who looks like you," and she said, "No, that is me; that is my face," and they said, "No, it's not. It doesn't meet the offence because it's not actually you." It looked like her, she said it was her and her husband said it was her, but the police said, "No, it's not you," which left her nowhere to go in terms of dealing with that content.

David Wright: StopNCII does not distinguish between whether an image is synthetic or genuine. As in that particular case, it is another mechanism where individual—victims—can and perhaps should be creating hashes to prevent the propagation. Again, it is just encouragement from platforms and companies to take StopNCII hashes in all these particular cases.

I would also like to mention nudification apps: the generative AI technologies and apps that have been built that we also encounter on a regular basis. This also brings app stores into scope. There should be policies and expectations—particularly on app stores—that these technologies should not be there. Back in March, we reported something like 29 different nudification app services to Apple, which then removed them, thanked us for reporting them and asked us to let them know if there were any more. Our question is: how did they get there in the first place?

Q29 David Burton-Sampson: Sophie, you mentioned the report to the police and the way that was dealt with. Do you think more needs to be done with police around how to handle these sorts of situations and, if so, what would you suggest is done?

Sophie Mortimer: Yes, there is a huge job of work in terms of training the police, all the way through that process from call handlers who take that first interaction with somebody right through to the end of the



criminal justice process. The first response somebody gets when they disclose to either a police officer or a call handler is going to be crucial to whether they take that further, and too often, we have people reporting to us that they were told just to come off that platform or that if they had not taken the pictures in the first place, this would not be happening. These victim-blaming narratives are still being perpetuated.

In fairness to the police—specifically frontline officers—they largely do not understand the law. They need the training to understand what the law says, how it would apply to somebody's real-life situation, and what overlapping behaviours—criminal or otherwise—might be going on around abusive behaviours, harassment or stalking, and domestic violence. There also needs to be an understanding of the impact of somebody's online life being targeted and abused in that way, because it is all part of people's lives now, and signposting to services such as ours and StopNCII can stop the sharing of that content.

What does evidence look like? We have no consistent approach to evidence around online offending. Sometimes we still get the police saying, "Can you leave that content live of this case because it needs to be live for the court case," but you could be talking months and months and months when we are being asked not to report content for takedown because the police say it needs to be done. Other forces will say, "No, we have screenshots, we have the links; take it down."

There is just no consistency, and this is not putting the victim at the heart of this process; they are almost forgotten. They are a means to give some evidence that a crime has been committed, but they are not getting justice because they are not getting their content removed, and it could be being reshared all the further while they are waiting. They are not being signposted to preventive services, and they are being made to feel worse than they did when they first turned up at the police. That cannot be right.

Q30 Catherine Fookes: Do you think we should criminalise the creation of sexually explicit deepfakes and introduce legislation to end this form of digital abuse?

Sophie Mortimer: Yes, we should, and we need to look quite comprehensively at what that looks like. It is the taking, making and solicitation of that content. Again, there are lawyer experts who will be able to support you with that much more than I can but, yes, there is no purpose to this, and it is widely available.

Q31 Catherine Fookes: What do you both think about how we can work in schools and around education and misogyny? This is part of the violence against women spectrum, is it not? We have talked a lot about the end results here, but how can we stop it happening in the first place?

Sophie Mortimer: There is a huge education piece, and yes, we always end up referring to schools, which is absolutely crucial, but we need to



target every generation. We need to be talking to the more vulnerable groups. It is difficult to reach adults, but we know this can be a form of elder abuse for older people and that people with additional needs can be more vulnerable to some forms of abuse.

So, yes, we do need to be talking to young people in schools about misogyny and the underlying attitudes and behaviours that are feeding into the sharing of intimate images, but I am always keen that we try to target everyone. Once you have turned 18, there is not a switch that has been flicked. We get a lot of reports of university students, particularly around sextortion. They have gone away from home for the first time, they have money in their pockets, and they are targeted. We need to be pinpointing those.

David Wright: StopNCII as a platform only supports adults. Back in 2022, we gave the technology to NCMEC—the National Center for Missing & Exploited Children—in the States to create Take It Down. It is the sister to StopNCII, but working for children. Now, there is clearly a reason for a demarcation because of all the regulations and, indeed, platform obligations if they do detect CSAM; there are mandatory reporting obligations if they identify child sex abuse material.

So StopNCII deals with adults, then NCMEC's Take It Down does exactly the same thing but for children. The first question you are presented with when you create a hash is, "What age were you in the image?" If you are at StopNCII and you say you were under 18, we just merely take you and transport you to Take It Down, and vice versa. So it helps support that this is not just about adults, particularly the gender split and how gender determines this issue. Whether it is relationship-based or extortion and the gender balance we talked about, it is the same for children as well as adults.

Q32 **Christine Jardine:** Listening to everything you have said, I wonder the extent to which the Online Safety Act 2023 is effective, but the legal system really has to catch up with the sophistication of what it is dealing with. Sophie, your answer about deepfakes especially made me think we have a lot to learn. Is that a fair assumption?

David Wright: There was an additional point I wanted to make. We know that DSIT and Ministers have said the Online Safety Act 2023 will cover this—again I am going to defer to Clare and Lorna about their interpretation of the Online Safety Act 2023 and how it does not determine NCII to be illegal—but Ofcom will have powers to provide service restriction requests and service access requests. So, being able to disrupt where platforms have not been compliant, and if that does not work, to submit service access requests to help block.

If DSIT does make that point, I would say Ofcom is not supporting individuals, and each of these processes is probably going to take weeks or months. As we have with child abuse and terrorist content, we need immediate intervention and immediate response. Again, when the team



are supporting victims, I am continually reminded that when they want help and support, they want help and support now. They want their content removed now, not in a few hours or a few days. So I would suggest all that mechanism is wholly unfit for purpose if that is going to be expected.

Q33 David Burton-Sampson: Jumping back to the sextortion figures, you mentioned 91% were male. Is there any disproportionality as to whether it is gay or straight males who are impacted, or do you have no data on that?

Sophie Mortimer: I do not think we have enough data to say. As I say, it is quite difficult to collect demographic data. The people who come to us are feeling very violated and tend to disclose the minimum, so I would not be able to speak to that.

Q34 David Burton-Sampson: Do you know if they have been targeted by people purporting to be women or men?

Sophie Mortimer: It appears predominantly they are purporting to be women.

Q35 Chair: We have touched on the international element of the work you do. I just wondered whether you could elaborate on the work NCII is undertaking with the UN, US, tech sharing, and so on, because this is an international problem. We can tighten up our laws in the UK, but we need to solve this problem internationally.

David Wright: That is a really pertinent question. I have worked in this space for well over 20 years and in terms of NCII, I am struck by the parallels with where children's harms were 15 to 20 years ago. We have an advantage as well in that it provides you with a blueprint or ideas and technologies that we can use.

It is an international issue. The Revenge Porn Helpline was the world's first helpline set up to support adults, so there is a lot of pride in that. Today, it is still one of the very few services that are available internationally; there are perhaps five or six. There are other NGOs around the world that do provide some support, but it is usually off the back of providing support to children.

As an example, a couple of years ago, a foundation was created in India to support children around intimate image abuse, and 40% of the cases in the first two months of its being live were from women because there is no other form of support available to them. Dedicated services are very few and far between for adults, but like I say, we are lucky in the UK in that victims have access to dedicated support.

StopNCII works with 104 amazing NGOs around the world that cover this. We convene at quarterly meetings—we have the next one next Tuesday—to try to share experience and regulation, and it is a very powerful way of



being able to understand exactly how international this is, as well as changing regulations.

StopNCII is always a big job. It is available in 22 languages; we are currently working on another 10 to try to make it as accessible as we possibly can, but that is really hard and we do rely on our NGOs to help us with that. When somebody arrives at StopNCII, they inevitably are going to need more help and support, so we are trying to direct them to where they can get more help and support in their own country. For example, in Pakistan, the Digital Rights Foundation provides help and support. It works in the opposite way as well; those NGOs will then use StopNCII or direct victims in their country to StopNCII.

I am really encouraged by the emergence of regulation. As well as European regulation, we have seen regulation in Australia, New Zealand and South Korea, the Digital Services Act, the Equality Act 2010 directed within Europe, and then more recently the UN cybercrime convention specifically having article 16 around non-consensual intimate image abuse and expecting members to have laws about NCII. We know that is going to help because there are big areas of the world that have no laws at all. I know we have been critical of our law and its not providing enough coverage and support for victims, but we have some form of experience and law, so we hope to be able to then provide that experience internationally back as to the gaps that may occur and what good law might look like.

Chair: On behalf of the new Committee—and the old Committee, I would guess—I want to thank you both for your patience in getting this far to the report and the inquiry, and thank you for all the work that you do and for the evidence you have given today.

Examination of witnesses

Witnesses: Courtney Gregoire and Gail Kent.

Q36 **Chair:** Welcome to our second panel. We have Courtney Gregoire, vice president and chief digital safety officer at Microsoft, and Gail Kent, director of government affairs and public policy (search, news and Gemini) at Google. Thank you both for coming and welcome to the Committee.

I want to give you both the opportunity, in under two minutes please, to explain and elaborate on what your respective platforms have been doing to combat NCII. I will start with Gail.

Gail Kent: Thank you very much for letting me come here today to explain what Google is doing about both synthetic and real non-consensual intimate imagery. As David said, the UK really is taking this seriously, and it is certainly really important to Google, so it is great to be able to continue the co-ordination.



HOUSE OF COMMONS

One of the things that I am responsible for at search is to balance that commitment to open access with protecting users, and that is something that we have done for the last 25 years. That has absolutely continued when we are talking about either real or synthetic non-consensual intimate imagery. One thing that is really important to us is to continue to have these conversations but also to talk to victim support charities, to victims themselves and to other tech companies.

To be really clear, Google does not create, attempt to host, distribute, or profit from non-consensual intimate imagery. We absolutely understand the profound distress that victims experience when they see that images have been shared non-consensually and they know that they are available online. We take this really seriously and we tackle it in three ways. First, we have clear, enforceable policies, and I can answer more questions on those later. Secondly, we want to make sure that victims have information about what they can do because we know that as search we are just providing a window on the world—a window on the internet. We are not actually hosting the data.

Thirdly, as you heard in the last panel, education is really important and that is something where we also put in a lot of effort. We know, however, that there is no single entity that can do this. This really is a whole societal problem, and online harm is something that has evolved over time. If I had been here two or three years ago, we would not have been talking about synthetic imagery; we would just have been talking about non-consensual intimate imagery. Keeping up and continuing to evolve with these nuanced problems and getting the balance right is really important to us.

We know that as a company that is responsible for helping users navigate trillions of websites and webpages, this is something that we really need to be talking about and working on, hence our continued focus on innovative solutions. We can talk more about that later, but thank you very much for inviting me and I look forward to answering your questions.

Courtney Gregoire: Thank you, Chair, and thank you, Committee, for drawing attention to this really important issue.

The new era of AI is truly upon us and it is being leveraged for positive forces like advancements in health care research. But as with any technology advancement, we know that there are those who are going to utilise it to do harm: the creation of currency yielded counterfeit; the creation of the telegram resulted in wire fraud. Creation and dissemination of synthetic NCII was actually around even before the advent of generative AI. If you look at 2019, it was Sensity AI that gave us the report that 96% of deep fakes were pornographic and 99% of those were impacting women.

I am really appreciative that the UK has been a global leader on this issue, criminalising the dissemination of intimate imagery in 2015 and, in



HOUSE OF COMMONS

2019, documenting in research evidence the very significant mental and physical impacts that this harm has on women. We look to those evidence reports as we think about our approach.

Microsoft has had a long-standing commitment to address image-based sexual abuse. We take this responsibility to contribute to a safer online ecosystem and protect our users, particularly children, from illegal and harmful content, very seriously. We do this by advancing technology, thinking deeply about important partnerships, and advocating for changes to the law and regulation.

At Microsoft, we have comprehensive policies to address non-consensual intimate imagery, whether real or synthetic. Those apply across our range of services. In addition, we have to think about how we address this on specific services. So, for example, on GitHub we have a policy that prohibits the sharing of software tools that could create non-consensual intimate imagery and we prohibit explicit sexual content being created through our Microsoft Generative AI services. As was mentioned in the last panel, we were very proud to partner with StopNCII through our Bing search engine and believe we are the first search engine in the world to do so. How we approached this was first and foremost donating the technology PhotoDNA that we created 15 years ago to tackle child sexual exploitation and abuse imagery.

It has been explained why hashing technology is so important in putting a victim's rights first. What we did was to update the PhotoDNA technology so that it could be utilised on a device, and a victim would not ever have to send that image even to an NGO they trusted: they could hash it within their device and send the hash of the content that they wanted to stop being disseminated, or, hopefully, prevent from being disseminated. We have utilised NCII hashes in our search index and that has resulted in the taking down of over 350,000 images of NCII in the past six months.

Yesterday, we brought together some stakeholders to talk about this very critical issue and to put forth our proposals to address this from a policy and legislative perspective.

I am happy to talk about that more, but as we think about a whole society solution, you can guarantee that Microsoft will continue to advance technology to ensure we can use AI for good, such as tackling this very sincere harm and detecting it across the ecosystem, making sure we have deep partnerships with those who understand and listen to the rights of survivors first, and advocating for laws that will get us up to speed to ensure that AI works for humans, not against us.

Q37 Chair: Thank you both. There is clearly a contrast in activity, and I want to dig down a bit further into that. Gail, I am going to come to you first, and Courtney, I am going to ask you about how your partnership with StopNCII is working and how it is progressing.

Gail, you said that Google is taking three steps to increase security on



HOUSE OF COMMONS

this: first is policies, secondly is victim information and thirdly is education. How are you assessing the effectiveness of this? Have you assessed that your approach is more effective than partnering with StopNCII, or do you have plans in the future to partner with them?

Gail Kent: Perhaps it is easier to answer the StopNCII question first. As David said, we absolutely continue to be in conversation with them, and when I explain what we do I can also explain where the one problem has been and the conversation we are having on that.

It is not just about having policies; it is about having enforceable policies. The first thing we do is to enable victims to report to us, and we have spent a long time developing a system that does that while taking into account the needs of a search engine. The most important thing to remember is that we are providing a window on to the internet, we are not hosting the content itself. But obviously we know, as was previously discussed, that it is not comfortable for users—it is deeply difficult for someone to come across their images on search. So, what we want to do is to give them the ability to not only ask us to remove those images from search, but also to go further and to look at the terms that were used to search for those images and the URLs.

When we do that, we do two things: we not only look at the images themselves and remove them, but we also look at duplicates. That is really important because we know it is re-victimising victims if they have to go and report again and again and again. We know that 90% of NCII images that are reported to us have duplicates. We do not differentiate between synthetic and real duplicates, and there are usually about 30 duplicates. So, we are not talking about one or two; we are talking about a number of images that were perhaps created when the one image was created.

The reason that context is really important to us is that we treat this in a very similar way to how society treats sex. We have talked a lot about consent, which is really important. When we ask people to report to us, we also want them to tell us the circumstances to give us more context about how those images were shared. We make no judgment on whether the images are nude or intimate; that is really important to us. If this is something that a victim feels is intimate, as we heard in the previous conversation, then we are relying on what they are telling us to make that call.

Talking about consent is really important because we enable them to say that they might have shared it once consensually and then removed that consent. They might even have commercialised it and then removed that consent. But if, right now, they do not want those images to be available in our search, then we will remove them.

We also allow them to tell us the specific search terms that were used to find that image so that we can make sure that they do not come up with those images, but we also look at similar ones. So, for example, if the



HOUSE OF COMMONS

images came up when searching for news images of Jane Doe, you do not want those images to come up again if you search for explicit images of Jane Doe. We want to take all that action. Then we make sure that we are down-ranking any sites we see that are repeatedly sharing non-consensual intimate imaging.

That is the enforceable policy stuff; I can come on to the information.

Q38 Chair: It might be useful for the Committee if this is sent in written evidence, because obviously you have extensive policies. What I wanted to ask is: if NCII content was made illegal, would Google's current policies be robust enough to deal with that change?

Gail Kent: We already go further than if this content was illegal. The issue with StopNCII is around the context that is shared with us. When an image is uploaded to StopNCII, we do not know what the context is and we do not know whether it is non-consensual intimate imagery. As you pointed out, as one of the largest search engines in the world, we really need to get the balance right between open access to information and protecting our users. We believe that, with slightly more context about why the user is sharing these images, we can make sure we are taking the right action.

What we do not want—Courtney pointed to this—is people who misuse systems to further abuse, and we really want to make sure that we are doing exactly what is right for the user.

Q39 Chair: A final question before I come on to Courtney: I appreciate that there is progress and constant communication between yourselves and StopNCII, which is obviously welcome, but what is stopping Google officially partnering?

Gail Kent: It is about that context of when an image is shared. StopNCII was set up to focus on social media and as you heard, Meta was the main collaborator. We think there are different requirements for search engines. We are absolutely delighted that Bing has managed to work through some issues and that is absolutely our intent as well, but it is around the context of when an image is shared.

Q40 Chair: Thank you, Gail. Courtney, perhaps this is a good point for you to come in and explain how it is that Microsoft Bing has been able to get round this, where Google perhaps has not just yet.

Courtney Gregoire: Our approach with StopNCII was to recognise that as a setup hotline we respected the content being reported by individuals and treated it as non-consensual report. We do have the capability as we implement this to think about what quality controls we should have in case there is content outside the bounds of intimate imagery. Let us call it a false positive that has somehow got into the hash database. In this context, we have erred on trusting this NGO to be having that conversation with victims.



You ask a really important question, Chair: this would be clean if this content was illegal and this was a repository of illegal content. That would set the standard to hopefully raise all bars across the online ecosystem. In other content areas, it has unlocked the capability to address it across platforms. Ideally, we would be moving towards a global standard in that context, but it does help raise the bar that we need.

Chair: Thank you. Samantha, you look like you want to come in?

- Q41 **Samantha Niblett:** I am just a little confused about the reason to not do the partnership. Everything you are describing is after the fact, so the information is already out there. Forgive me if I have misunderstood, but the hash is about if there is a threat of a photo; you put a hash against that particular image. When you say it is the context, that feels very much like Google is deciding to be the judge and jury as to whether a victim's—I use that word deliberately—story matches up. But you are then saying that the context you look at is broader.

Gail Kent: Just to be really clear, we do not decide who are victims and who are not victims, and that is why it is important to us to be really clear that it is intimate imagery and not nudity; we want to make sure that users are in control of what happens. If we have the context then it enables us to understand what exactly the search terms were; it is less the context in which it was shared, more the search context. We are not asking, "Was this an image that was taken in these circumstances?" I should have been clearer on that. We want to know the search terms used to find it, the URLs where it was found, and then the image itself. That enables us to make sure that we can deal with duplicates, because, as David explained, hash technology is really miraculous; it had a huge impact on tackling child sexual abuse, but it can be manipulated.

We also know that there is going to be more than one image, so we want to look at duplicates as well, and having that additional information helps us do that. It also helps us downrank sites that are regularly sharing this information because it is not just the one URL that we want to remove. If there is a site that is regularly sharing non-consensual intimate imagery, we want to make sure that is coming nowhere near the top of search results.

- Q42 **Chair:** Thank you. Courtney, just so we can dig down into how your partnership works with StopNCII, if Bing produces a search result that matches a hash you have received, are you able to feed that URL back to StopNCII so they can make efforts to have that material removed as well? How does that information sharing work?

Courtney Gregoire: We are still on that step of what is the right information sharing mechanism back to StopNCII to help it understand where this content might be in a broader context.

Chair: Thank you. Rachel.

- Q43 **Rachel Taylor:** Thank you both for coming along today. The speed with

which NCII is taken down is really critical to reducing its spread and impact. Why then do you both review content reported as NCII before taking it down or delisting it from searches?

Courtney Gregoire: One of the challenges as we have approached digital safety for over 20 years is misuse of reporting mechanisms. The healthiest use of technology, as we think about safety all up, is how we get to a place where we reduce false positives. There are additional complications in this content area, obviously, because we are talking about a wide range of intimate imagery and, as noted, consent is really important. But in this context we do want to take the step of understanding whether this meets the boundaries of the policy to ensure we are addressing it. And in fact that is a core element of regulation in this space: how are we going to prove that we are accurately enforcing our policies and accurately and transparently reporting that we are actioning that content?

Even in the hash context, we know there is a space to ensure product quality. Your point, however, is very well taken. Speed is critical here. It is speed and addressing the duplicates and addressing where else this goes across the online ecosystem. It is a place where we humbly have to continue to get better and where, honestly, AI can help us get better.

Q44 **Rachel Taylor:** On average, how long does your review process take?

Courtney Gregoire: It varies by content harm type, if that makes sense, because, as you have articulated, we prioritise both child sexual exploitation and abuse imagery and non-consensual intimate imagery to be addressed rapidly. The challenge—this is one of the places where we want to get better—comes when users sometimes, understandably, do not report it under the right category and so it may be NCII that they have reported elsewhere.

So, it varies, but it is a priority.

Q45 **Rachel Taylor:** On average, how long would it take?

Courtney Gregoire: Our SLA is we endeavour it to be in a space of hours.

Gail Kent: I would give exactly the same answer. Our aim is to take images down as quickly as possible because we understand the harm that it causes. As David clearly said, and as we heard from previous witnesses, we can all understand what it is like to know that images are out there. We want to take them down as quickly as possible, but we also want to make the right decision.

Q46 **Rachel Taylor:** So, you leave it there to review rather than taking it down to review it and then putting it back; is that correct?

Courtney Gregoire: It varies by our service. We can suspend access to content while reviewing depending on the search context; that is much harder to do given that it is not our content. So, it does vary by our



HOUSE OF COMMONS

service, if that makes sense. The ability to suspend content from being accessed by other eyes while a review process is happening obviously happens in our gaming environment or LinkedIn environment. In a search index that step is not possible.

Gail Kent: This is where getting it right is so important to us because we are committed to open access. We want people to be able to go to Google and find the information they need. What we do not want is for information to be missing because of the sort of false positives that Courtney has talked about, where we have taken something down without checking that it is indeed non-consensual intimate imagery.

Chair: One last question, Rachel.

Q47 **Rachel Taylor:** I was going to push you on that point, but I will move on. What action is taken upon a user reporting a non-consensual culturally insensitive image of themselves, such as a woman pictured without her headscarf or hijab?

Gail Kent: We leave it for the user to define what is an intimate image. We do not make any decision on that; we treat it in the same way as we would any request of this type.

Courtney Gregoire: We also let the user tell us what is intimate imagery. Our approach in this context is often to start looking at other harm types, whether that means harassment or others. This is usually the space that it comes in. But I do want to be honest, this is a space that we need to learn how to get better at. Right now, if an individual reports such content, we would likely action it as intimate imagery or access of private information, but the question is, how do they know where to go? How do we make them understand where to report this? That is the space where we are seeing a delta in the ecosystem.

Gail Kent: If I can add one thing: this is where, when thinking about harms, Google really takes its global responsibility seriously. When we are looking at people who are responsible for reviewing this type of information, we make sure that we are drawing from a wide range of languages, cultures and nationalities, so there is that experience and knowledge.

Q48 **Rachel Taylor:** Would it be fair to say that your primary objective is to make sure that the data that people are searching for is paramount, as opposed to there being harm with some images?

Gail Kent: No, I would say Google's goal or mission is to organise the world's information and make it accessible and useful. We really are trying to balance both access to information that people are looking for with protecting users and complying with laws. I do not think it is an either/or.

Q49 **Rachel Taylor:** Do you have that balance right?



Gail Kent: It continues to be something that we need to think seriously about as harms develop and how we use the internet, including with AI, changes.

Q50 **Chair:** I want to pick up on where Rachel left off about this balance. You both talked about the false positives and the fear around those false positives. There must be quite a large number of false positives generated, surely? Do you have the statistics and figures for that, either one of you?

Gail Kent: I do not.

Q51 **Chair:** If you do not have them to hand now it would be great if you could write into the Committee. Because obviously, if there is such an imbalance, I would be interested to see the percentage of false positives. If you are preventing yourselves from deleting or delisting an image or video because you are worried it might be a false positive, surely there must be quite a high percentage of false positives.

Gail Kent: We know that people, as Courtney said, misuse these systems and they misuse these systems globally as well. You could imagine the impact if we took action on an image that was maybe related to current events and did that automatically, rather than making sure that it was something we should be enabling people to see. The impact of not getting this right is really significant.

Courtney Gregoire: Chair, if I may, we would be happy to follow up on this, but I acknowledge that one of the real risks is in the reported concern flow for our users. Honestly, on average, we see 50% to 60% of that content being unactionable. Sometimes a nefarious actor is using a bot to report the same piece of content about a real-world event that they would like suppressed off the internet. We need to get better at detecting when it is a bot submitting a report of abuse into a space. As you might imagine, these actors will of course use a prioritised harm as they are doing our workflow because they know we are going to prioritise looking at child exploitation content or gender-based online violence.

So, honestly, tech has to continue to evolve and understand how those are abusing even the systems we have designed. When I say false positive, I want to acknowledge it is this whole system—how do you ensure that the reports that are coming in are actually from users and get the noise out? How do you then ensure that users understand how to report actionable content? How do we address that? That is where we have leaned into partnerships, where there is an opportunity for a single moment to that report. Our quality control on that looks a little different than when I say a bot is attacking in a report abuse scenario. But I give that context.

Chair: Thank you. I have some follow-up questions from Kirith, then David and then from Catherine.

Q52 **Kirith Entwistle:** That was going to be my question, Chair: I wanted to



know if you might have information you are able to share with the Committee about how many false positives you are receiving and how grave a threat it really is.

Courtney Gregoire: It is a difficult term, if that makes sense. Sometimes we say false positives as if the technology to detect has caught something it should not. In this case we know we can get better at everything. How do we leverage the right report abuse scenario so a person knows that they are telling us what they need to tell us as minimally as possible so we can action? How do we make sure we get other people out of that system? How do we not put artificial barriers into the reporting system? It is a hard thing to unpack, but it is a place we know we continually have to invest in.

Q53 **Chair:** On culturally sensitive image usage, do you both have policies and training in place?

Gail Kent: For Google employees, yes. We really are about putting the users first.

Courtney Gregoire: There is an intimate imagery policy. I suspect we can get better at training.

Q54 **David Burton-Sampson:** I want come back on something Rachel brought up around the review process and how long it takes. This is really important because the longer that image stays online, the longer that person is suffering harm and potentially further harm. Courtney, your answer was, "We endeavour within a couple of hours," and your answer, Gail, was, "As soon as possible." That is not really an answer, so I am interested to know what your actual SLA is. If you cannot provide that now, that is fine; maybe again you could write into the Committee to let us know. But what is the actual SLA, and how often do you meet it? Because with all respect, Gail, as soon as possible could be a week, and in a week, a lot of harm could be done to that individual.

Gail Kent: It is definitely quicker than a week, but we can come back to you with precise data. I mentioned briefly the information part. When we reply to the user that has reported to us, we also provide them with additional resources—we know that just having your image taken off search results is not enough; that image is likely still hosted somewhere—in terms of other things they need to do, including reporting to StopNCII, which means that the information is no longer going to be available. And that is really important to us.

Chair: Thank you. Natalie, we will move on to you.

Q55 **Natalie Fleet:** Earlier this year, the Committee heard from NCII survivor and campaigner, Georgia Harrison. She described the positive difference it would make for victims if platforms had a phone line that they could call to report this. Is this something you have and if not, is it something you could have?



Gail Kent: This comes on to the information that we have. As I said, when we reply to a user we point them to helplines that they can speak to so that they can speak to someone who is an expert in dealing with this sort of thing.

Q56 **Natalie Fleet:** This was specifically about reporting. There are helplines available to talk through it, but is there anything to report it?

Gail Kent: We do not have a helpline. We have an online form, and we have spent a lot of time making it as simple as possible and then staying in online communication with the user. But we do not have somebody that you can speak to.

Courtney Gregoire: We have a similar approach. I watched that testimony with interest because I thought to myself that we are trying to make sure we scale and take the minimal amount of information we need from a user in order to action swiftly. But, as indicated, we really want these individuals to have the right support system that they need.

You asked why we sort the content. In the UK, if someone says, “Non-consensual intimate image,”—what StopNCII does—what our search engine does through our authoritative principles, is, we hope, to help them to understand the nature of this problem through useful news stories, and then they are directed to people who can actually help them through all the elements, both in terms of how to take down the content, but also the legal and other victim supports that should happen. Where we do not have that—I say that because our customer support agents would not be trained, and I think people should be experts in this space to do that—we want to ensure that we get the reporting content for what we can address on our platform and direct them to the right resources for the full range of supports they need.

Q57 **Chair:** Does anyone have any follow-up questions on that? No? Okay, I am going to move on to the legal status of NCII content. Both of you will have heard similar questions in the previous panel. Are the Online Safety Act and Ofcom’s related powers enough to require platforms like yours, with a UK presence, to proactively remove NCII content? Are you prepared for greater powers from both?

Gail Kent: Obviously, we have followed very closely how NCII is being treated under UK law. As I said, we believe at the moment we already go further than if this was illegal, but as you have heard from other witnesses, and as Microsoft said in their statement, clarity is always helpful.

Q58 **Chair:** Can I just push back a little, because you have said a couple of times now that you go further than if the content was illegal? If that is the case, surely there would be no availability of NCII for adult content, and it would be the same as child sexual abuse material. What level are you are holding yourself to there?



Gail Kent: That is such an important question. It comes down to consent. I have spent a lot of my life dealing with child sexual abuse material. It is not easy, but we are good at identifying child sexual abuse material. As David said, we also have got much better as a society about understanding the spectrum of material that exists. Because it concerns somebody under the age of consent, it is easier for us to define what that looks like. Once you get to adult intimate imagery, there are legitimate times when intimate imagery is produced by consenting adults who may also wish to share that material and do that with consent. Understanding that this is without consent is the difficult part.

Q59 **Chair:** Do you not see some parallel with where we were maybe 20 to 30 years ago when talking about rape and consent that we have hopefully got to a position now where we believe a victim? That is where I would hope most organisations and most public services would strive to be, whereas at the moment it seems very much that your search engine goes, "We need to find out the context first, before we believe the victim."

Gail Kent: First, as I said to Catherine, we make no judgment. The context is definitely not, "Are you a victim, or are you not a victim?" The context is: what search terms turned up this intimate imagery? The context is not trying to understand when something was shared or when it was not. But we know there are consensual sexual images and intimate images online that users would not want us to take down. So, we need somebody to tell us—whether that is the victim, law enforcement, or somebody acting on behalf of the victim—that this is non-consensual imagery. That is not the context we are asking for.

Q60 **Chair:** The context you are asking is whether this is consensual or not.

Gail Kent: No. First, in terms of reporting it, if somebody has ticked a box saying, "This is non-consensual," we do not make a judgment on that; we absolutely believe them. That is not what the false positive is about. The false positive is about somebody sharing an image of a portcullis rather than something that would be in the context of non-consensual intimate imagery. That is not the context.

Q61 **Chair:** Why do you not take it down from the point of reporting then?

Gail Kent: This is the conversation we have just had where we absolutely do want to take it down; we just need to make sure that it is intimate imagery and not the sort of false positive that Courtney was talking about.

Q62 **Chair:** I appreciate this is a conversation we have gone over, but I feel like we have not got down to the real reason as to why that is a priority over something else. Courtney, how are you treating NCII right now? If it was to be made illegal, what changes would that make to your processes?



Courtney Gregoire: I am going to be clear on this: I think making this content illegal would be better for victims, and that is who we should be putting first. It would unlock the capability of hopefully an Ofcom or a similar model to create a registry so that there was a single place for a victim to report that content, making it illegal across the ecosystem, which would unlock faster action. From my perspective, it is about: what is the best legal regime to support victims and survivors in this? This legal clarity would help with that.

Chair: Thank you. Catherine, did you have something to add?

Q63 **Catherine Fookes:** Following on from the Chair, actually, I am really struggling here. Surely, we should just use the precautionary principle: if there is a danger, why not take the things down immediately? I am also struggling with the comprehension in this matter; if you were signed up to a partnership with the organisations we heard from before, and the imagery that had been reported to them was hashed, presumably all those places on the internet where they appear in Google searches, everything would be hashed. If you are in a partnership, is that what would happen? So, why not just get into the partnership and get rid of these online images asap on your platforms?

Gail Kent: First, we do not host the images. We are already a member of numerous organisations that share hashes, so what we would do with any of our hash databases is remove them from search results. We do not remove them from where they are hosted; we only remove them from being seen in search results, which is why it is also important to provide that extra information to users so they and StopNCII can go to the actual hosters and have them removed. So, we are only showing that window on the world.

We are absolutely committed to making sure that users have a safe experience. I am going to try to explain again what we are looking for and how we help victims. When a victim reports something to us, we ask them to confirm that either they, or the person they are representing, are in the image, that it is an intimate image, and that they are not sharing this on this website or anywhere else for commercial purposes.

We then ask them what search terms resulted in the image being shared and the URLs that came up. I just want to emphasise that we do this as quickly as possible. We do not spend a long time making a judgment as to whether this is a picture of a dog instead of what they say it is. We look to do two things as quickly as possible: we remove the image from search terms and duplicates—just as a reminder, we normally see 30 images for 90% of cases that are reported on this, so a lot of images—we then remove the URL from search terms, and we look to make sure that similar search terms do not result in those images as well, if that is what the user wants. Sometimes they just want them removed from individual search terms.



I am happy to follow up in writing, Chair, with more detail and maybe the pages that have our reporting flow so that you can see exactly what processes the victims go through. We have spent a long time trying to make this as simple, as straightforward and as human as possible.

Q64 Chair: Thank you. It would be really useful in that information to find out the number of drop-offs in terms of what you have decided are false positives. I know you do not like the phrase, but it would be really intriguing to know what the numbers are.

Gail Kent: Yes.

Chair: I will hand over to Alex who is going to ask questions around the Online Safety Act.

Q65 Alex Brewer: I have a question for both of you; Gail, maybe you would like to go first. Sharing non-consensual intimate images is a priority offence in the Online Safety Act. What difference does that make to your respective approaches in tackling this?

Gail Kent: I have covered what our approach is to policies and take-down and then information. We also do a lot of educational work. The difference that this will make, or the OSA makes in general, is in our obligation to report to Ofcom and to clearly explain what we are doing. I know previous witnesses have said this as well, but the UK really is leading the way when it comes to online safety. You see that in the many conversations we have with Ofcom and the requests that it makes to us. What we do not know is specifically what reporting will look like when it comes to non-consensual intimate imagery, and we are in conversations with Ofcom about that.

Courtney Gregoire: My perspective is, as a priority offence, I do not think it would change our approach. I believe we treat it as a priority offence under the OSA across our ecosystem. My hope would be that Ofcom approaches that by focussing attention on those who are hosting this content. Search engines have a responsibility to think about how our search engine looks, but these victims deserve to get to the place that is hosting this content. If we move that towards a priority offence, potentially coupled with the legality, there should be clearly a responsibility to the underlying hosts of this content to action this in a much more meaningful way: give victims rights and a reporting mechanism which should be actioned. And to be perfectly honest, get rid of what we see as one of the most egregious practices: some hosting providers requiring victims to pay money to report their content to get it taken down. My hope would be that this would change in the ecosystem. We are treating it as a priority offence, but of course Ofcom could tell us if it would take a different approach.

Gail Kent: If I can add one point to that, what we hear in the Ofcom conversations is not just that it is focusing on what more we should be doing, but also what industry best practice is, which just as Courtney



said, could be then obligated to other platforms who are not following these rules.

Q66 **Alex Brewer:** With that in mind, do you think the content itself rather than just the act of sharing it should be made illegal?

Gail Kent: From what we have heard today that clarification would be useful, but obviously that is a decision for Parliament.

Q67 **Alex Brewer:** But what impact would that kind of legislative change have on your ability to deal with this problem?

Gail Kent: I can very clearly answer that question. At the moment, because the content is not illegal, we downrank it; we do not remove it, which is why the search term part is important. It is still available on the internet; if somebody searched for that information, perhaps because they themselves are the victim, it is way down the ranking but it still exists. If it is illegal, we then remove that altogether.

Alex Brewer: That has filled in a whole lot of gaps from earlier. Thank you.

Q68 **Samantha Niblett:** The phrase, “We go further” is interesting, because I have heard several times, “We go further than if it was illegal already.” It feels like we do not, because you have just said, “If it was illegal, we would remove it rather than downrank it.”

Gail Kent: We would still need somebody to report it to us to know that it was non-consensual.

Q69 **Samantha Niblett:** Slightly moving on but staying in and around the same topic, I know your Microsoft report was talking about some of the extra steps Government need to take, particularly when it comes to resourcing police. We have touched on some things that you think would be good from Government, but just to summarise—I will come to you first, Courtney, if that is okay—what action do you think is required from Government, both legislative and otherwise, to effectively address the emerging challenges of NCII, including deepfake stuff?

Courtney Gregoire: First, one clarification that is really important is the understanding that previous law has focused on dissemination because that was the nature of the harm before the advent of generative AI. Let us be honest: the creation of synthetic sexualized imagery needs to be considered now as illegal because of how generative AI is being used to cause this abuse. The previous witnesses have been very clear about how it may be different, and so the intent of that person is different, but at the end of the day the creation of it is causing the harm and needs to be legally addressed.

Legal clarity would unlock help and support, and we have seen the mobilisation that can happen around a registry for a centralised place like StopNCII to be resourced. Every user, every citizen needs to understand that it exists, that it is where they should go, that it has the resources



and the support to counsel. We spend a significant amount of time hearing directly from those who have been impacted by this harm, and thanks to a session here two weeks ago, actually, we heard frontline experience from survivors who often times approach law enforcement that does not understand the intersectionality between cyber-harassment, stalking, image-based sexual abuse, and how quickly those things overlap. I can see a potential to unlock that toolbox for frontline law enforcement to know how to help when someone comes in its front door.

We have partnered previously in the nature of cyber-crime to do just that. How do you help train frontline law enforcement as cyber-crime evolves? There are experts who really deeply understand how that online harassment evolves in many ways, and frontline law enforcement could have those skills. We have seen that has been a priority of the Biden White House and something that it was considering supporting in the Department of Justice, and I hope our new leadership would continue that.

Q70 Chair: Courtney, you have talked about multiple partnerships in solving NCII as a problem. Do Google have any partnerships, and if so, with whom, if they are not partnering with StopNCII?

Gail Kent: We have not signed up to StopNCII, but I would say that we are partnering with it. The reason I say that is we have funded the Revenge Porn Helpline in the past, and there are conversations that we continue to have with it. We also speak to numerous UK bodies such as Refuge to make sure that we really understand this issue. Also, as David, Sophie and you, Chair, have said, there is a huge overlap with what we can learn from child abuse, and that is another area where we spend a lot of time understanding exactly what we can do.

Q71 Chair: You said that you are partnering, but it is not really a partnership if it is not official. Which other organisations and stakeholders are you officially partnering with to combat the problem of NCII through Google?

Gail Kent: I can come back to you with that answer.

Chair: Thank you very much.

Q72 Samantha Niblett: Generative AI is a topic of conversation for so many reasons, but one example that probably hit everybody's radar earlier this year was deepfake pornographic images of Taylor Swift, which were created with a Microsoft tool. A question for you both: to what extent is it possible for your companies to prevent your generative AI systems from producing NCII content?

Gail Kent: Google has been leading the way on responsible AI. We put out our first responsible AI principles in 2018, where safety and privacy are key for Gemini, which is the Google generative AI tool and image creator. We do not allow sexually explicit images to either be asked for or



created, and we have numerous checks to make sure that does not happen.

Courtney Gregoire: Samantha, we take a six-pillar approach as we think about how the challenge of abusive AI can happen. I will be quick about this. First and foremost, consistent with our responsible AI principles, which are about a decade old now, is: what is the safety architecture at the point of model creation and model application? You raised an example where our safety architecture failed. We recognise that we had tested it. We have an explicit ban and prohibition on sexually explicit content being generated through our generative AI. We took the critical step of investigating, through the spaces of the internet of those who had taken upon themselves to do everything they could to break our safety architecture, what search terms they used, how they did it, and how they repeatedly attack it so that we could learn. We know that safety is an adversarial space. It has always been, much as the cyber-criminal world has been in malware for a while. We are going to have to constantly monitor those who seek to do harm and get around safety architecture systems to continue to make ours even better. But we focus on that safety architecture first.

Secondly, the role of content provenance and watermarking is really important. Average users want to understand, "Is the content I am coming across synthetic?" That is why we create services that say, "You deserve, as a Member of Parliament, to put a stamp on your content so that you own that content in the ecosystem." We hope that standard continues to evolve and give credit to Adobe for advocating for that. That is a space in which we advocate for this.

Those are two important things but, thirdly, we are going to have to continue getting better at detecting the harmful and illegal content that will be disseminated. While we make these commitments that we do not want that content, we know there are applications out there right now that are enabling it. So, we are going to have to get better at detection. I have talked a lot about our detection, which is technology, but it is definitely partnerships, and then when the law can help and support that.

Lastly, and I say this humbly, we cannot do this without learning directly from how users are experiencing the internet, so bringing that voice in constantly and directly to our engineers as they build this system to find out what happened, and what went wrong, so that they can understand it deeply. Those are our approaches, and we believe you have to do them all at the same time to make the flywheel work.

Q73 **Samantha Niblett:** What actions are you taking to prevent the promotion of tools for creating deepfake stuff? As you have mentioned, the search engines are a set of curtains to the world. I appreciate that you can pay for optimisation tools or sponsored prominence on the search result, but how can you stop products being presented that should not even be out there, such as applications that make people look nude?



Gail Kent: This is still a work in progress. As we understand harms and how they develop, these are things that we collectively need to consider. Our biggest defence is to make sure that we are providing high-quality, useful information, and I think we would all agree that these apps do not necessarily provide that.

Q74 **Samantha Niblett:** One thing I have had brought to my attention is that the top three results for “deepfake porn” on Microsoft Bing are deepfake porn websites. So, if you put in “deepfake porn”, it brings up deepfake porn websites, despite their content being against Microsoft policy. Google has downranked them; why has Microsoft not?

Courtney Gregoire: It is very much our intent to do so. I would gladly take that and make sure we improve the algorithm so that it responds to that query with authoritative content, as intended. Authoritative content from our perspective would be news articles explaining the pervasiveness of this problem and challenge, or directing people to resources that they need.

Humbly, I say we can do better and that should not be the results. We also know these search queries, and how people will seek this content, will continue to evolve. That was a simple one; we should get it right. But we need to get better as the user queries for this content evolve, because what “authoritative” means in this context should be shining a light on the problem and the challenge as it is covered by news articles and reports of authority such as StopNCII.

Q75 **Chair:** Thank you. Before I come to our last question, I just wanted to check around the Committee and see if anybody had any follow-up questions to any of these, or anything that they wanted a bit more information on. No? Okay.

I want to thank you both for coming today and answering our questions. We want to get to a solution and a better place together. I hear loud and clear that you both say you want to work together, to learn together, and to learn from partner organisations. In that vein, we will be sending a report with our recommendations from the cross-party Committee. Would you both be prepared to respond to the relevant recommendations that we put in this report?

Courtney Gregoire: Yes.

Chair: Thank you. Thank you both very much for your time today and for all that information. We will look forward to some written answers as well.