Consultation response form

Please complete this form in full and return to Part5Guidance@ofcom.org.uk by 5/3/24

Consultation title	Guidance for service providers publishing pornographic content
Full name	✂
Contact phone number	✂
Representing (delete as appropriate)	Organisation
Organisation name	Yoti
Email address	✂

Confidentiality

We ask for your contact details along with your response so that we can engage with you on this consultation. For further information about how Ofcom handles your personal information and your corresponding rights, see Ofcom's General Privacy Statement.

Your details: We will keep your contact number and email address confidential. Is there anything else you want to keep confidential? Delete as appropriate.	Nothing
Your response: Please indicate how much of your response you want to keep confidential. Delete as appropriate.	None
For confidential responses, can Ofcom publish a reference to the contents of your response?	Yes

Your response

Question	Your response
Question 1: Do you agree with our proposed guidance on scope? If not, please provide any information or evidence in support of your views, including descriptions of services or content where you consider it is unclear whether they fall within the scope of Part 5.	Confidential? – Y ✂

Question	Your response
Question 2: Do you have any comments on how our proposed guidance applies in respect of pornographic content created by generative-AI services within the scope of Part 5? Please provide any information or evidence in support of your views.	Confidential? – N We believe this question refers to *'user-generated content within the meaning of section 55(3) and (4) of the Act' ('Exemptions and exclusions from the scope of Part 5 of the Act',* 2.14, a)). This section of the guidance could be improved to clarify and match the wording of this question (*'content created by generative-AI services'*). We would question why content that has been generated by a user with help from artificial intelligence is being left out considering the speed at which generative AI is progressing. It is already well capable of generating pornographic content that resembles content produced by performers and professionals. An example of this is the endemic progression of pornographic deepfake content of singer Taylor Swift on X (formerly twitter), which has made headlines (*'Taylor Swift deepfakes spark calls in Congress for new legislation'*, BBC, 28 January 2024). The guidance should also be more explicit as to the criteria set out in 2.17 a), in reference to sections 80 (2) of the Act, to clarify what Ofcom will view as a *'significant number'* of adult users in proportion to the total UK population, and of potential underage users in proportion to a provider's total user base. We would point to the Information Commissioner's Office *'Likely to be accessed impact assessment'* document published in July 2023, which also references the 'significant number' threshold. We would like to see both regulator's definitions aligned.

Question	Your response
	We note the comments of the Canada Privacy Commissioner who conducted an investigation into user generated intimate image abuse and concluded that the adult content provider "had a legal obligation to obtain the complainant's consent and had failed to do so." However, we do not currently have confidence as to which regulator is leading on compliance in this sensitive area which spans both data and content systems and processes. We would encourage Ofcom to also consider artificially generated avatars; which can be built to resemble minors or hybrids (e.g. body of child or adult with face of child or adult.) Clearly there are several potential harms ensuing from such practises, of normalising AI generated child images in pornography or blended child/adult images. We would suggest that facial and body age estimation can be used to assess the ages of AI generated content and used to monitor all such content for underage performers.
Question 3: Do you have any comments on our proposed guidance in respect of the kinds of age assurance which could be **highly effective**? If you consider there are other kinds of age assurance which have not been listed that you consider could fulfil the proposed criteria, please identify these with any supporting information or evidence.	Confidential? – N *'Duties applying to providers within scope of Part 5.'* The clearest acid test as to the kinds of age assurance which are highly effective, is to assess the current market and current volumes of age assurance checks and feedback from commercial organisations undertaking age assurance. The popularity of facial age estimation amongst consumers has been echoed by Meta, 'we have found that 81% of people presented with our menu of options chose to use Yoti's video selfie to verify their age'. We would encourage Ofcom to engage with the German age regulatory bodies, KJM and FSM. The KJM and FSM have been engaged in reviewing age assurance approaches for a decade. The KJM has published a list of over 100 approaches, which it has reviewed and approved as effective, the 'KJM Raster'. The FSM engaged three independent tech experts to undertake an in depth review of Yoti's facial age estimation; copy available upon request, and subsequently awarded its Seal of Approval in 2020. We would also recommend Ofcom to review benchmarking data where that becomes available. NIST[1] (National Institute of Standards & Technology) has invested significantly in undertaking a global benchmarking of facial age estimation; given its market success and the

[1] https://pages.nist.gov/frvt/html/frvt_age_estimation.html

Question	Your response
	increasing accuracy rates. The benchmark was kicked off in September 2023 and the results should be published imminently. The test set includes several million of test images; from a wide global sample. We would refer OFCOM to the review undertaken in 2020 by ACCS to ascertain if Yoti Facial Age Estimation would meet the level required for Challenge 25 in the UK, for the purchase of alcohol. Yoti's facial age estimation has been certified since 2020 by the Age Check Certification Scheme for use in a Challenge 25 policy area, The intention of the test is to assess whether or not the Yoti Age Estimation System is fit for deployment by determining if an 18 year-old (the nominal age) would be incorrectly estimated as being over 25 (the Challenge Age policy). 'The report highlights how, subject to the exclusions mentioned in the report, our testing indicates that this version of the tool PASSES for deployment in a Challenge 25 policy area.' Even 4 years ago in 2020, the system was 'deemed fit for deployment in a Challenge 25 policy area and at least 98.89% reliable.'The Yoti AI Services Age API version 1.1.1 (Target of Evaluation) assessed on or before 17th November 2020 can be stated to accurately estimate the age of person of nominal age 18 as being under the age of 25 with 98.89% reliability where results are stated by the Yoti system to an uncertainty of less than 4.6 years. The mean absolute error, mean predicted age, upper and absolute tolerances were all within the permitted parameters as set out in ACCS 1:2020 Technical Requirements for Age Estimation Technologies. In addition, at the request of one of our clients, our May 2022 white paper was independently verified by the ACCS for our measurement methodology and accuracy of our results. The ACCS said that: "The training, testing and results reporting presented in the Yoti white paper have been independently validated by ACCS, who have certified that Yoti have deployed appropriate methodologies to analyse the performance of their Facial Age Estimation algorithm, including ensuring appropriate separation of machine learning training data, testing data and validation data." As we have stated in our response to Ofcom's previous consultation (*'Protecting people from illegal harms online'*), we believe that Ofcom should improve and uniformise its use of terminology across its documentation. We think it important that the guidance departs from the use of *'age verification or age estimation'* to instead uniformly use *'age assurance', 'age assurance techniques'*, *'age assurance technologies'*, or *'age assurance*

Question	Your response
	solutions'. It is suggested in 2.18 to only use '*age assurance*' but we do not see this recommendation being followed throughout the document. As stated in our previous response, we believe this would help convey an understanding that fully verifying a user's age, rather than estimating whether a user is above or below an age threshold, should come as a last resort and only where it is proportional to do so. We note that Ofcom also uses the terminology *'age assurance method'* in its *'Guidance for service providers publishing pornographic content: Consultation on draft guidance on age assurance and other Part 5 duties'* (*'4. Guidance on age assurance duties'*, 4.3) document. Should this be made the preferred terminology, then we would like to see it used across all the documentation uniformly. We welcome the inclusion in the guidance of a mention that *'guidance may refer to industry or technical standards'* (Section 82(4) of the Act, and section 2.23 in this document) but regret that Ofcom has chosen not to do so (in 4.48). There are a number of standards that we would like to bring to the fore. PAS 1296:2018 is a Code of Practice for Online Age Verification service providers developed by the British Standards Institute and the Digital Policy Alliance. The PAS – a Publicly Available Specification – is intended to assist providers of age restricted products and services online with a means to adopt and demonstrate best practice and compliance when it comes to age checking. It helps businesses comply with regulation, and safeguard their reputation, by providing recommendations that help prove an online user's age. The standard fully addresses issues relating to privacy, security, safety, usability, accessibility and data protection online. Claims of conformity against PAS 1296:2018 should be verified by an independent third party. There is also an incoming ISO standard on age assurance. The document, ISO/IEC 27566-1:2025, presents a framework and core principles for age assurance systems used in age-related eligibility decisions. It is part of the Information security, cybersecurity, and privacy protection standards. The document covers topics such as the purpose of age-related eligibility decisions, the problem of inadequately defined age assurance processes, and the need for trust in terms of efficacy, acceptability, privacy, and security. The goal is to balance privacy outcomes for implementers, individuals, and policymakers. The document does not prescribe specific age assurance systems or methods but provides a framework for policymakers to specify applicable types and indicators of confidence based on their requirements. It also outlines terms and definitions related to age assurance, including concepts like

Question	Your response
	age verification, age estimation, age inference, and indicators of confidence. The scope includes privacy considerations, and normative references to other standards are provided. The document was prepared by the ISO/IEC JTC1 Information Technology, Subcommittee SC27, specifically Working Group WG5, focusing on Identity Management and Privacy Technologies. The incoming Institute of Electrical and Electronics Engineers (IEEE) 2089.1, Best Practice for Age Verification standard for an age appropriate digital services framework is based on the 5Rights Foundation principles (a) Recognition that the user is a child, b) Consideration for the capacity of and upholds the rights of children, c) Offers terms appropriate to children, d) Presents information in an age-appropriate way, e) Offers a level of validation for service design decisions.) in order to help build the digital world young people deserve. This Standard offers organisations the opportunity to create services that uphold children and young people's rights and support their evolving capacity. For the purposes of this standard, a child is any person under the age of 18. This standard provides a specific impact rating system and evaluation criteria and explains how vendors, public institutions, and the educational sector can meet the criteria. It also sets normative requirements for published terms, design, and delivery that can recognize and respond to the needs of children and young people. We would also point to a French working group in which Yoti participated in 2023. It was jointly led by social network Yubo and the *Association française de normalisation* (Afnor, the French Standardization Association which represents that country at the International Organisation for Standardisation). This working group recently published a document titled *'AFNOR SPEC 2305 Prévention des risques et protection des mineurs sur les réseaux sociaux'* (*'AFNOR SPEC 2305 Risk prevention and protection of minors on social networks'*). Yoti was responsible for drafting the sections on age assurance and identity verification, later approved by all participants. At the time of writing, an English translation is being finalised, and we would be delighted to share when it is available and to answer any questions Ofcom teams may have if this is of interest. Other participants and observers in this working group included Airmedia, Bodyguard, Dailymotion, E-Enfance, Internet sans craintes, Meta, Mym, Point de Contact, Respect Zone, Sorare, Tralalere, Université Toulouse Capitole and Université Paris VIII. Finally, the Proof of Age Standards Scheme (PASS) has developed a number of standards on general principles and definitions, requirements for identity and age verification, requirements for e-ID validation

Question	Your response
	technology, requirements for data protection, privacy and security, requirements for proof of age card design and construction, and requirements for digital presentation of proof of age - known as 'dPASS'. These standard documents can be found on its website. It would be useful for all departments to understand this approach and the accompanying audit process. (https://www.pass-scheme.org.uk/downloads/)/ *'Figure 4.1: Summary of our proposed approach to implement highly effective age assurance.'*, *'Examples of age assurance methods that are not capable of being highly effective.'* We have strong reservations about the current phrasing of this table, and in particular of the inclusion of *'Examples of age assurance methods that could be highly effective*'. We believe this should be more explicit by being rephrased to say '*methods that can be effective' or 'methods that are highly effective'*, and by amending the list accordingly. We are however surprised to see that credit cards and mobile network operator age (MNO) checks have been included in this list, whilst debit, solo and electro cards are instead rightly included in the *'not capable of being highly effective section'*. We would like to see published research published to support the basis on which credit card and mobile network operator checks are included. Has there been an exercise of regulatory freedom of information gathering that proves the efficacy of these approaches, with or without additional reauthentication? If this has occurred, it would be useful to publish this evidence. We think that it is also important to review the bias levels of all methods; including for instance credit cards. In the US, the Federal Reserve Board published a report (*'Economic Well-Being of U.S. Households in 2022'*, Board of Governors of the Federal Reserve System, May 2023) looking at the economic well-being of households, including data on credit card access. 82% of adults in the US own a credit card. But this falls to 62% of younger adults, aged between 18 and 29. In the United Kingdom, research suggests this number for the overall population is around 65-68%, and has declined in recent years. The Federal Reserve's report concluded that credit card usage also differs by race, ethnicity and disability status. Income data from employment, savings and investments are the most influential factors in determining whether a credit card is issued or not. But it's fair to say that more adults who are from certain racial or ethnic groups, or have a disability, face discrimination if a regulation or organisation requires evidence of credit card ownership to access a service. It is also important to note that adults

Question	Your response
	on lower incomes will find it harder to qualify for a credit card. These same adults may also struggle to buy other forms of identification, such as a passport, which can be expensive. It is crucial to know who is in possession of a document being presented. As with credit cards, there are concerns about the ease with which children can 'borrow' an adult's mobile phone device or use a hand me down device. This therefore would suggest that a reauthentication check would be proportionate to ascertain if the current user of a device or person using an account is an adult or a minor. In addition to this, the strength of the MNO-based assurance method will rely on the strength of the original age checks performed by third parties, such as mobile phone service providers in retail premises, which cannot be guaranteed. Therefore, we have strong reservations about photo-ID matching, credit card checks and MNO checks being featured in this list of highly effective approaches, unless further authentication is undertaken to ascertain who is using the card or mobile device. Without re-authentication we would suggest that they fall short of most of the four criterias as set out. However, we welcome the fact that Ofcom has included self-declaration in the *'Examples of age assurance methods that are **not capable** of being highly effective'* section of the table. Indeed that is in line with recommendations made by Ireland's national online safety regulator, the Coimisiún na Meán, which states in its latest *'Draft Online Safety Code'* consultation that *'mere self- declaration of age is not regarded as an effective age verification technique'*. This opinion is also shared by the Netherlands' national online safety regulator, the *Commissariaat voor de Media* (*'we think that self-declaration is not an appropriate age-verification tool'*, *'Responses to Coimisiún na Meán Call for Inputs: Online Safety Code'*), the Irish Safer Internet Centre and 5Rights Foundation in their responses (*'Responses to Coimisiún na Meán Call for Inputs: Online Safety Code'*, published by the Coimisiún na Meán). We would however question why, in spite of stating that *'research shows that self-declaration is not an adequate form of age-assurance'*, self-declaration is then included in section 18.78 of *'Volume 4: How to mitigate the risk of illegal harms - the illegal content Codes of Practice'* document which is part of Ofcom's *'Protecting people from illegal harms online'* consultation. We believe Ofcom's Online Safety Act documentation should be reviewed holistically to remove mentions of self-declaration, or where it is mentioned, to make it clear it is not an appropriate age assurance solution in the vast majority of instances where there are age inappropriate risks of content, conduct, contact, contract for minors.

Question	Your response
	We welcome the inclusion of facial age estimation and open banking, although in the case of the former, there should be requirements for independent review of bias and accuracy, transparency as to the origin of the data set, independent review that the images are instantly deleted, independent assessment of liveness detection. We welcome the inclusion of a statement in *'Examples of kinds of age assurance'* (4.8) that facial age estimation is included in the guidance because Ofcom believes it is *'capable of being highly effective*, [is a] *sufficiently mature technology'* and '[is] *being deployed at scale'* unlike other estimation techniques. We note the significant absence of mention of credit reference agency checks method which is currently highly widespread. It consists in a check to a credit reference agency of a name, date of birth and address. Indeed, if a child knows the name, date of birth and address of a family member and enters these, they would check out correctly with an electoral roll or credit agreement. We would therefore challenge the reason why this method is not listed as either effective or ineffective. If this is because Ofcom has undertaken research and concluded that it is too easy for children to enter an adult's details to pass the check, then this type of 'knowledge based check', which can easily be found out and easily shared, is not an effective method. This is a very widespread method deployed today, so we would welcome clarity on this point. We would welcome more clarity from Ofcom, as to whether they deem Open Banking checks to be effective. Technically, Open Banking does not disclose date or birth but banks could offer this additional information. A bank may state that an account is an adult bank account, because it is not a child account. However there can be joint, shared accounts and child accounts where a parent has access. It would also be useful for Ofcom to clarify a number of elements for this to be deemed a 'highly effective method': 1. Who holds the liability on incorrect information, if a bank says that a person is over 18 and the result is later contested; can the relying party or service provider go back to the bank and challenge the bank? 2 Is the bank obliged to act on information that account data is incorrect? 3. Does any party have any obligation to report incorrect information We welcome the use of a non-exhaustive list of examples as this facilitates innovation.

Question	Your response
	We would ask that in each instance where approaches are deemed effective, that there is a requirement for publicly available, transparent materials, independently assessed to justify their inclusion.
Question 4: Do you agree that service providers should use the proposed criteria to determine whether the age assurance they implement which is highly effective at correctly determining whether or not a user is a child? Please provide any information or evidence in support of your views.	Confidential? – N *'Overview'* We would suggest a rewording of the introductory text, for it to be clearer and more assertive as to what it would like to achieve. For instance, where the guidance states '*Service providers <u>should</u> implement an age assurance process that fulfils each of the criteria of technical accuracy, robustness, reliability, and fairness to ensure that it is highly effective at correctly determining whether or not a particular user is a child'. W*e believe that it should instead read *'Service providers <u>must</u>* (...)' to match the wording of other sections of the document (such as in 2.5, 2.17, etc.). *'Figure 4.1: Summary of our proposed approach to implement highly effective age assurance.*', *'Criteria that the age assurance should fulfil to be highly effective.'* We welcome the introduction of *'Criteria for ensuring age assurance is highly effective.*' (4.1). We would challenge the statement in 4.12 (*'Currently, we do not have sufficient evidence as to the effectiveness and potential risks of different age assurance methods to recommend specific metrics'*). We could suggest two initial key metrics: mean absolute error (MAE) and levels of circumvention of the method. As we have suggested previously, we think that it would be helpful if Ofcom were to undertake research into what a person can for instance undertake with £10 under 1 hour and low skill, £100 under 1 day and medium skill, £1000 in a 1 month and with high skill. Ofcom could then consider what level of assurance is deemed proportionate for a given use case, and what levels of circumvention are deemed acceptable. We agree that *'new solutions [are] likely to emerge over time*' (4.13) and think it would be appropriate for Ofcom to *'set a base level'* for different types of age assurance technologies. We think it should be Ofcom's role and responsibility to periodically conduct reviews and academic research, scan the horizon and update these base lines as time goes in line with technological advancements. We would also like to point to 4.54, with which we disagree where it is assumed that *'it is reasonable*' to assume

Question	Your response
	that children would only have access to *'most basic forms of falsified documents'*. Sophisticated fake identity documents can be purchased at a very low price (often for less than it would cost to procure official documentation), and many children have digital literacy skills. Therefore, we believe it should be assumed children will indeed have access to *'sophisticated'* or *'elaborate' forms of falsified documents*', and that there is a need for Ofcom to assess circumvention. We would highlight that statements in 4.12 (*'we do not have sufficient evidence (...) to recommend specific metrics'*) conflict with statements in 4.15 (*'a service provider should (...) ensure the method has been evaluated against appropriate metrics'*) and in 4.17 (*'We welcome stakeholders' views on the suitability of metrics we have suggested'*). We think there is a further opportunity to reword the language, as the words *'metrics'* and *'criteria'* seem to be used interchangeably and create confusion. It is unclear whether the word *'metrics'* in 4.15-4.17 is meant to refer to the four *'criterias'* listed in the table. Whilst we understand that Ofcom may not want to set base levels (we think it should), it should at least be clear what metrics should be taken into consideration (we have previously suggested MAE (mean absolute error) and circumvention levels), otherwise the risk is that all assessments are made on the basis of radically distinct metrics and base levels, risking a very fragmented and uneven delivery of the Act's policy aims across the Part 5 providers landscape. We would challenge the statement in 4.13 (*'the age assurance industry is still nascent'*), particularly considering it conflicts with the earlier statement in *'Examples of kinds of age assurance*' (4.8) that facial age estimation is included in the guidance because Ofcom believes it is *'capable of being highly effective*, [is] *sufficiently mature technology*', and '[is] *being deployed at scale*' unlike other estimation techniques. Age checks via transactional or reusable digital identity have been possible and undertaken for many years. A sensible set of questions to distinguish a *'nascent'* versus a mature industry might include: ● Is there a healthy ecosystem of providers? ● Are there standards in place? ● Has it been adopted by global organisations? ● Are there independent audits with consistent measurement? ● Has there been transparent benchmarking at scale? ● Is there an established trade body?

Question	Your response
	● Has any sectoral research been done? ● Have there been any regulatory reviews in other jurisdictions? To each of these questions, the answer is yes, these are in place, so we would argue that *'nascent'* is not an accurate description. Yoti as a company has entered its tenth year having been founded in 2014. One of the bodies which certified Yoti's age technologies, the Age Check Certification Scheme (ACCS, which is quoted in this document) and the professional organisation for the age assurance sector, the Age Verification Providers Association (AVPA) were both founded in 2018 and are well-established and well-respected organisations. Our feedback on the criteria specifically is that we feel the *'technical accuracy'*, *'robustness'* and *'reliability'* sections feel confluent. We think they largely describe the same aim, which is that the age assurance method should be good at determining the age of a person. Therefore, it could be merged into one single '*precision'* criteria. As we have also said in our response to the previous consultation (*'Protecting people from illegal harms online'*), we believe that accounting for the ease of circumvention of measures and the evolution of circumvention techniques (for example virtual private networks), and users' literacy levels in that field seems to have been left out of the documentation aside from a duty on Part 5 providers not to promote them on their sites. We think these are important factors to consider when assessing and implementing an age assurance solution and believe they should form the basis of a criteria. We would welcome formal studies in this field. We fully support Ofcom's definition of the risk of using facial age estimation without implementing liveness check technology (4.54). We believe this should be a mandatory requirement for any facial age estimation technology provider. We would also suggest the *'fairness'* criteria should be renamed *'equity'*, as 'fairness' does not feel like the best term to use in this context. For this criteria, we would suggest that Ofcom should encourage the use of the Fitzpatrick scale, as currently used in Yoti's *'Facial age estimation white paper'* published in December 2023 (and available at https://www.yoti.com/wp-content/uploads/2023/12/Yoti-Age-Estimation-White-Paper-December-2023.pdf). The Fitzpatrick scale is a dermatological test that involves grading skin tone at two different points in time, one before exposure to sun and then after a week's exposure to sun. We would encourage Ofcom to be more thorough in its guidance in

Question	Your response
	this section, particularly in order to mitigate the potential harms identified in the *'Equality Impact Assessment'* section of this document (especially A1.27 b)). This is important as providers will need to assess whether the datasets used by age assurance technology providers to train their algorithms will have been ethically sourced, and representative of the broader United Kingdom population. In addition to this, we would suggest that Ofcom should add two further considerations that service providers should have when reviewing solutions (in 4.37). The first should be to have regards to whether the solution poses any barriers to users because it is reliant on the possession of an object, document or device, and therefore outright excludes a segment of the population. This is relevant considering the cost and complexity of obtaining identity documents, or modern devices. It is an important consideration for providers at the point of deciding whether to implement a second age assurance technology on their site to provide users with an element of choice to help mitigate this risk. We welcome the inclusion of this point in the *'Equality Impact Assessment'* section of this guidance (A1.28) We have advocated throughout our responses so far that users should always have a choice of what age assurance methods to use. We would welcome more detail as to how Ofcom will ensure it stays *'up to date with developments'* (4.41) in the age assurance space, and more crucially how it intends to disseminate this knowledge within the online safety ecosystem to encourage best practices and the improvement of standards over time.
Question 5: Do you have any information or evidence on the extent of circumvention risk affecting different age assurance methods and/or on any steps that providers might take to manage different circumvention risks for different methods?	Confidential? – N We have provided information to support our views about the importance of considering the ease of circumvention for each age assurance technology in our responses to previous questions. As we have suggested previously, we think that it would be helpful if Ofcom were to undertake research into what a person can for instance undertake with £10 under 1 hour and low skill, £100 under 1 day and medium skill, £1000 in a 1 month and with high skill. Ofcom could then consider what level of assurance is deemed proportionate for a given use case, and what levels of circumvention are deemed acceptable. We would suggest that Ofcom also reviews where it considers it proportionate for reauthentication to be required. For instance this could

Question	Your response
	include reauthentication of the live user of an account in relation to data from mobile phone operators, credit reference agency data or any knowledge based data which can be shared and indeed in terms of access to a credit card - which may be borrowed or a shared card between parent and child.
Question 6: Do you agree with our proposed guidance that providers should consider accessibility and interoperability when implementing age assurance? Please provide any information or evidence in support of your views.	Confidential? – N Whilst not strictly feedback for Ofcom, we are supportive of the work undertaken by Ofcom in cooperation with other regulators within the Digital Regulation Cooperation Forum (DRCF). Considering the potential that digital identity can play in helping deliver the policy objectives of the regime, we continue to consistently feedback that the proposed Office for Digital Identities and Attributes (OfDIA) should have a status enabling it to join the DRCF. This is also relevant considering the latest suggestions by the Home Office in its latest consultation (*'Alcohol licensing: age verification'*, published 24 January 2024) that certification to the United Kingdom Digital Identity & Attributes Trust Framework (UKDIATF), which will be eventually be overseen by OfDIA, could be a requirement for the employment of age assurance technology providers in retail premises. We welcome the introduction of the additional criteria of *'accessibility and interoperability'* (4.32). We think that, in addition to a section on *'accessibility'*, there should be a section on *'inclusivity'* too. These two sections should be part of the core *'criteria that the age assurance should fulfil to be highly effective*' (*'Figure 4.1: Summary of our proposed approach to implementing highly effective age assurance*.'). Our general feedback to the 'accessibility' section is that it is currently lacking detail. 24% of the population having a disability (*'UK disability statistics: Prevalence and life experiences'*, House of Commons Library, 23 August 2023) should also warrant the use of the word *'must'* rather than *'could'* for this section of the guidance. We would like to see references to accessibility standards, principles and techniques such as the Children's Code, the Hemingway system of 'grade level' judging in terms of the review of language used, or the Web Content Accessibility Guidelines (WCAG). We believe it should be a goal for age assurance technology providers to achieve a minimum level of WCAG 2.2 (Yoti achieved this in July 2023), and that WCAG assessments should be carried out by an independent third party (the true independence of which is guaranteed by the fact that assessors will get paid whether the firm hiring them fails or passes the assessment) rather than Part 5 providers or age assurance technology providers themselves.

Question	Your response
	We would welcome more co-production (*'Co-production'*, National Health Service website, https://www.england.nhs.uk/always-events/co-production/) of standards and guidance by Ofcom involving disabled persons at the earliest possible stage of conception of policies and guidance, such as by creating citizen and user-led co-production groups that can co-write the guidance, and harnessing the expertise of Ofcom's *'colleagues networks'*. We believe the guidance also would benefit from the inclusion of a mention of age tokens, and the definition of standards. Age tokens and other tokens-based age assurance technologies are currently widely available and employed around the world. Age tokens serve as digital proof of age verification, allowing users to access various integrated websites without repeatedly proving their age. These tokens, devoid of personal information, only contain the age result and details of the verification process. Users, after verifying their age once, can create a free age account with an anonymous username and password, enabling them to transfer age tokens between browsers. Despite relying on first-party cookies, the system prioritises user privacy, sharing information only during direct interactions with the portal. Stored within the infrastructure, age tokens avoid passing on to relying parties, ensuring no linkage to personal identifiers. Data stored includes a shared ID, verification method, liveness and authenticity check details, check time, and a unique ID for auditing decisions. Age tokens are live, stored securely, and do not entail personal data sharing during age verification. Anti-spoofing measures involve token signatures and controls for detection if shared. While working in incognito mode, tokens are session-specific, emphasising the creation of age accounts for cross-browser access. Businesses are charged a volume-based subscription fee for age tokens, with an unlimited generation capability meeting specified criteria. Quick to generate, the age token payload is less than 1KB. Yoti currently employs age tokens in the euCONSENT project (https://euconsent.eu/), with potential for other providers to issue and receive tokens. We would welcome rules in this guidance that would set the frequency at which an age token should be reverified, and what cybersecurity standards Ofcom would view as suitable in order to ensure the safety of tokens-based age assurance technologies.

Question	Your response
Question 7: Do you have comments on the illustrative case study we have set out in the guidance? Do you have any supporting information or evidence relating to additional examples of how the criteria and principles might apply to different age assurance processes?	Confidential? – N The video-sharing platform regime (2.28 and 2.29) Whilst strictly not a case study, we welcome the introduction of a mention of the video-sharing platforms (VSP) regime, which Ofcom has been responsible for since 2021, and really is a precursor to the Online Safety Act regime. However, we believe there was a missed opportunity to publish a report on lessons learned, observations, and the successes and challenges of Ofcom's enforcement efforts. This would have enabled better responses to this consultation. We would nevertheless encourage Ofcom to publish such a report ahead of its wider consultation on age assurance. *'Overview'* We would challenge the wording of a *'non-exhaustive list of kinds of age assurance that could be highly effective at correctly determining whether or not a user is a child'* in this guidance. Ofcom in its guidance should ensure it gives Part 5 providers the maximum certainty that the solutions named by Ofcom would ensure compliance with the policy objectives of the regime, all the while, we agree, leaving flexibility for providers to adapt to their unique circumstances. We would recommend that Ofcom liaises with the KJM, in terms of its list of over 100 approved methods for age assurance for the German marketplace ('KJM Raster' webpage available at https://www.kjm-online.de/aufsicht/technischer-jugendmedienschutz/unzulaessige-angebote/altersverifikationssysteme). And also with the FSM to understand the approach it offers for independent expert review of new approaches; it would be interesting to review this approach with NCSC or other such technically competent bodies. We think the paragraph *'Service providers should implement an age assurance process that fulfils each of the criteria of technical accuracy, robustness, reliability, and fairness to ensure that it is highly effective at correctly determining whether or not a particular user is a child.'* should be rephrased to say *'Service providers <u>must</u> implement an age assurance process that fulfils each of the criteria (...).'* This is to ensure the document stands firm in its commitments to deliver the aims of the regime from its onset. Similarly, we think the following paragraphs should read *'Service providers <u>must</u> also consider the principles of accessibility and interoperability (...)'* and *'Service providers <u>must</u> ensure access controls.'*

Question	Your response
	As we have said previously, we disagree with 4.54 where it is assumed that *'it is reasonable'* to assume that children would only have access to *'most basic forms of falsified documents'*. Older teens are making money by producing or procuring fake documents for younger teens. This is why, considering the ease with which elaborate fake identity documents[2] can be obtained at a very low price (often for less than it would cost to procure official documentation), and that children are often digitally literate, we have suggested that circumvention rates should be a factor in assessing the effectiveness of all age assurance solutions. Currently there has only been a review of a few of the age assurance methods in terms of mean absolute error, false positive, false negative rates. The same scientific comparison should be available across all methods.
Question 8: Do you agree with our proposed guidance on the record-keeping duties? Please provide any information or evidence in support of your views.	Confidential? – part (sources to be redacted) As we have stated in our response to Ofcom's previous consultation (*'Protecting people from illegal harms online'*), the guidance on record-keeping duties as currently drafted has the potential to confuse providers. The main risk to avoid is that providers believe they have to keep a record of the actual result of an age assurance technology, such as a user's age or full date of birth. We will repeat that we believe providers should not retain any information beyond where a user is situated in relation to an age threshold. Where the guidance states *'This includes details of any third-party supplier contracted to provide an age assurance process and what kind of age assurance the process uses, whether made up of one or multiple age assurance methods*', we think Ofcom should be more precise as to what details it would like included by providers. Where the initial text states '*Service providers should keep a durable written record of the age assurance process in use. The record must be up-to-date and easy to understand.',* we would like more detail set out as to what *'durable' and 'easy to understand'* mean to the regulator. We would like to see more information about what format it expects the *'details of any third-party supplier contracted to provide an age assurance process and what kind of age assurance the process uses'* (5.4) should be in. We welcome the provisions in 5.6, but believe they could be expanded to include considerations of the privacy trade-offs that a provider's choice

2 ✂

Question	Your response
	of age assurance solution may mean for a user. As we have stressed throughout our responses, this would help bring in an element of proportionality of the measure in relation to the amount of personal information disclosure that it requires. In similar consultations, Ofcom has given an indication of how many sites it expects will come under the scope of the Online Safety Act. In order to determine whether the procedure described in 5.10 (*'Where we have concerns that a provider, based on its written record, has not complied with its obligations under data protection law, we may refer the matter to the ICO'*), we would first want to get a better understanding of how many sites are expected to be in scope of the '*Part 5*' regime. This is also a unique situation where the scope of this consultation technically overflows onto another regulator, whose very structure is due to change following passage of the re-introduced Data Protection and Digital Information Bill. Therefore, there are many variables and unknown parameters, in addition to questions around the Information Commissioner's Office future form, such as whether it will be able to manage such an assuredly high number of sites and referrals from Ofcom. We are supportive of the guidance as drafted in *'Our proposed guidance for summarising the written record in a publicly available statement*.' However, we would encourage Ofcom to develop its guidance further to include some accessibility criteria, such as requirements to follow the Children's Code to ensure that these statements are easy to understand by an average citizen, taking into account the national average literacy and digital skills levels and allowing for wider inclusion[3].
Question 9: Do you have any comments on our proposed approach to assessing compliance with the duties on service providers who publish or display pornographic content, including on the proposed examples of non-compliance? Please provide any information or evidence in support of your views.	Confidential? – N Similarly to the feedback we have provided about record-keeping duties and the processes described in 5.10, we have reservations about the process described in 6.11 regarding *'Principles for assessing compliance*'. The effectiveness of this process can only be assessed on the basis of a knowledge of the number of providers that will fall in scope of this guidance against the amount of human resources that will be assigned to the enforcement of this specific regime that sits in the Online Safety Act. We would welcome more information on whether there will be a specific Part 5 providers team, and enforcement unit within the Online Safety Group. We would reiterate, that there needs to be a level playing field in terms

[3] ✂

Question	Your response
	of enforcement. If, either as in the case of the VSP regime, the perceived risk of enforcement remains zero after legislation comes into force or there is only a focus on the largest players, then large organisations will focus their energy on resisting changes. We fear then that the business case will not meet the threshold for action at platforms and the intended positive impact of the Act will be at best delayed at worst unsuccessful.
Question 10: Do you have any comments on the impact assessment set out in Annex 1? Please provide any information or evidence in support of your views	Confidential? – N We are generally very supportive of the approach taken in *'Figure A1.1: Summary of our assessment of the impact of our proposed guidance on regulated service providers'*, as well as the assessment of proportionality of the measures required throughout. Whilst we agree that there may be staff costs in the form of a very limited amount of hours taken to research the internet and find information about products, this is likely to be very limited (*'there may be staff costs associated with understanding and considering the criteria / the principles when implementing age assurance'*). This will likely not be significant because Part 5 providers will be able to rely on a certain number of existing registers such as the UKDIATF's list of certified providers (which will be updated by the Department of Science, Innovation & Technology following passage of the Data Protection & Digital Information Bill), as well as the very easy to use *'find a provider'* page on the Age Verification Providers Association (AVPA) website (https://avpassociation.com/find-an-av-provider/). Further, it is likely that age assurance technology providers will adapt their literature in a way that is easily legible against the requirements of the Part 5 regime. We would argue against the assessment in A1.10 (*'Impact on service providers who are small and micro businesses'*) that *'the overall direct costs relating to our proposed guidance are likely to be a greater proportion of the total costs/revenues for smaller firms'* in the case of the implementation of an age assurance technology. Indeed, Yoti provides specific offers and packages to help small platforms manage their costs, including through the use of high assurance methods such as the Yoti App that are free to use. Yoti also offers flexible tokens which were built and deployed previously in the context of the Digital Economy Act 2017. Similarly, the onboarding cost and experience can be made to be favourable both to large and small operators. In order to ensure that adults are not *'unduly excluded'* from accessing legal content, we would re-emphasise our earlier points about the need for

Question	Your response
	Ofcom to make proportionality a central criteria, and to make clear to Part 5 providers that users should be given a choice of measures. We believe this risk could also be mitigated by ensuring that there are steps where a person could move to a different method of age assurance should they wish to challenge an outcome. We also believe that Ofcom correctly identifies that there is a risk that the regime widens inequalities should it be over-reliant on technologies that require users to possess a device, an object or an identity document. This is why we have made recommendations in previous sections that the guidance on criterias should be enlarged to include an *'inclusivity'* criteria, as well as the other considerations we have suggested. We are supportive of the approach taken in the *'Other impacts'* section. We agree that the mitigation measures as proposed would not *'unduly affect competition'*, however this will also be dependent on Ofcom thoroughly and firmly enforcing the regime, and ensuring a level playing field across the industry that does not mean that some providers are put at a disadvantage and therefore losing traffic and commercial revenue. See link to recent certification from Google[4] for facial age estimation. We would request that fair competition in the age and identity assurance marketplace, be tabled as a standing agenda item, for discussion at the DRCF, along with counterparts from ICO, FCA and CMA Hence our previous points about the importance of resourcing of Ofcom's enforcement team, but also of the ICO as per the system of referrals described in the guidance (5.10). *'Ex-post evaluation of the impact of our guidance*.' As we have said previously in this response, we have concerns about the proposed timelines for implementing the Part 5 regime. We welcome the inclusion in this text of mentions of a *'report on regulated services' use of age assurance for the purpose of compliance with their duties set out in the Act, and how effective the use of age assurance has been for that purpose*' (A1.19), which we take as being the same document referred to elsewhere in the guidance as '*a statement setting out the final guidance*' expected *'in 2025'*, after which *'Government will bring these duties into force*.' As we have said previously, we would welcome a firmer commitment on the part of Ofcom as to when in 2025 it would publish such a statement, considering that if done on 31 December 2025, it would potentially not be until late 2026 or even early 2027 that we see this guidance being

[4] https://www.accscheme.com/registry/google-inc-llc

Question	Your response
	seriously enforced. As we have set out in our response to Ofcom's previous consultation ('*Protecting people from illegal harms online*'), we believe that Ofcom should bring forward its roadmap on age assurance ('*Ofcom's approach to implementing the Online Safety Act*') especially considering the gravity of what is recognised as a priority harm. We stand ready to support the work of the *'evaluation workstream'*, the composition and work of which we would like to see more clarity and transparency about. *'Equality Impact Assessment'* We agree with the assessments that existing inequalities may be reinforced if age assurance technology is not responsibly sourced, implemented correctly, and accurately assessed for compliance with performance objectives. In this response, we have made recommendations for Ofcom to further develop its guidance in order to address this. We have also made a number of recommendations in this consultation response in order to address the potential negative impacts identified in this section, such as by encouraging Ofcom to revisit its *'criteria'* section. We are supportive of the assessment as made in A1.28. False positives A 'False positive' is when we ask a question with a yes/no answer, and the answer comes back as *'yes'* when it should have been *'no'*. So for example, when dealing with age-restricted goods or services, if we ask *'Is this person old enough to buy alcohol*?' and facial age estimation tells us *'yes they are'*, but actually they are not, then we have a *'false positive'*. In this kind of use case, we can regard false positives as a measure of facial age estimation being too lenient. When dealing with age-restricted goods and services, the age of interest is what we call the age stipulated in the relevant law or regulation. For Part 5 providers, 18 years of age. In many use cases, we will ask '*is this person above the age of interest?*' (e.g. *'are they over 18?*'). We can then configure facial age estimation technology to simply return a response of *'yes, they are over 18*' or '*no, they are not*'. Facial age estimation has a margin of error, and we would expect some false positive replies when asking if a person was above the age of interest (particularly if their true age is close to it). For this reason, where it is illegal to access a certain good below a given age, we offer the

Question	Your response
	possibility to configure a threshold age above the age of interest, to create a safety buffer. Instead of asking facial age estimation if the person is above the age of interest, we ask if they are above the threshold age. For an age of interest of 18, we might choose a threshold age of 23. We ask facial age estimation whether or not people are over the age of 23. If the answer is 'yes, they are', we can accept with confidence that they are over 18. In contrast, a regulator may deem for an age of interest of 13, where it is part of the terms of service of an organisation, and not illegal for a service to be offered to a 12 year old, they may decide that a buffer is not required. Therefore, the challenge is picking an appropriate threshold for the given use case which delivers an acceptably low false positive rate, and we believe Ofcom's guidance should address this. At Yoti we publish the accuracy and bias levels, transparently so that our relying parties and regulators can make the decision as to whether a buffer is appropriate for a given use case, and can review the buffers periodically, as the technology improves. However, when considering the acceptability of false positive rates for any given use case, the relative risk involved should be considered too in terms of the potential relative harm in an 11 year old accessing an age restricted good or service versus an 18 year old. To conclude, false negatives are an annoyance to those trying to access an age-restricted service or purchase age-restricted goods. They can cause friction and conflict between customers and retail staff, with assaults and abuse being a growing problem, and customers having to carry physical ID documents. These documents, such as passports and driving licences, can be expensive to obtain and a significant proportion of young people do not have them. Additionally large numbers of physical ID documents are lost every year, (circa 1 million driving licenses[5]) increasing the risk of identity fraud as well as incurring a replacement cost. It is important, when discussing the choice of a threshold age and safety buffer for use with facial age estimation, to consider both where facial age estimation is too lenient and where it is being too cautious. Higher thresholds will likely decrease false positives at the expense of causing more false negatives. We think Ofcom should provide more clarity in this guidance as to its risk tolerance for any given deployment of facial age estimation, and guide Part 5 providers in choosing a threshold which is likely to deliver an acceptable balance between false positives and false

[5] https://www.gov.uk/government/news/drivers-lose-almost-a-million-licences-in-the-last-year

Question	Your response
	negatives. Regular review will be needed as the technology continues to improve[6].
Question 11: Do you agree that our proposed guidance is likely to have positive effects on opportunities to use Welsh and treating Welsh no less favourably than English? If you disagree, please explain why, including how you consider the proposed guidance could be revised to have positive effects or more positive effects, or no adverse effects or fewer adverse effects on opportunities to use Welsh and treating Welsh no less favourably than English.	Confidential? – N We do not have any comments.

Please complete this form in full and return to Part5Guidance@ofcom.org.uk.

[6] https://pages.nist.gov/frvt/html/frvt_age_estimation.html