Your response – User empowerment duties

Questions 29 and 30: Features employed to enable greater control over content

For all respondents

Question 29: What features exist to enable adult users to have greater control over the type of content they encounter?

In your response to this request, please provide information relating to (a) – (d) where relevant.

Response: OneID is a UK certified digital ID service that enables anyone to prove who they are online by using their existing online banking credentials to securely share ID data.

In this example flow, an Instagram user could use OneID to simply verify their name, or just that they are a real person.

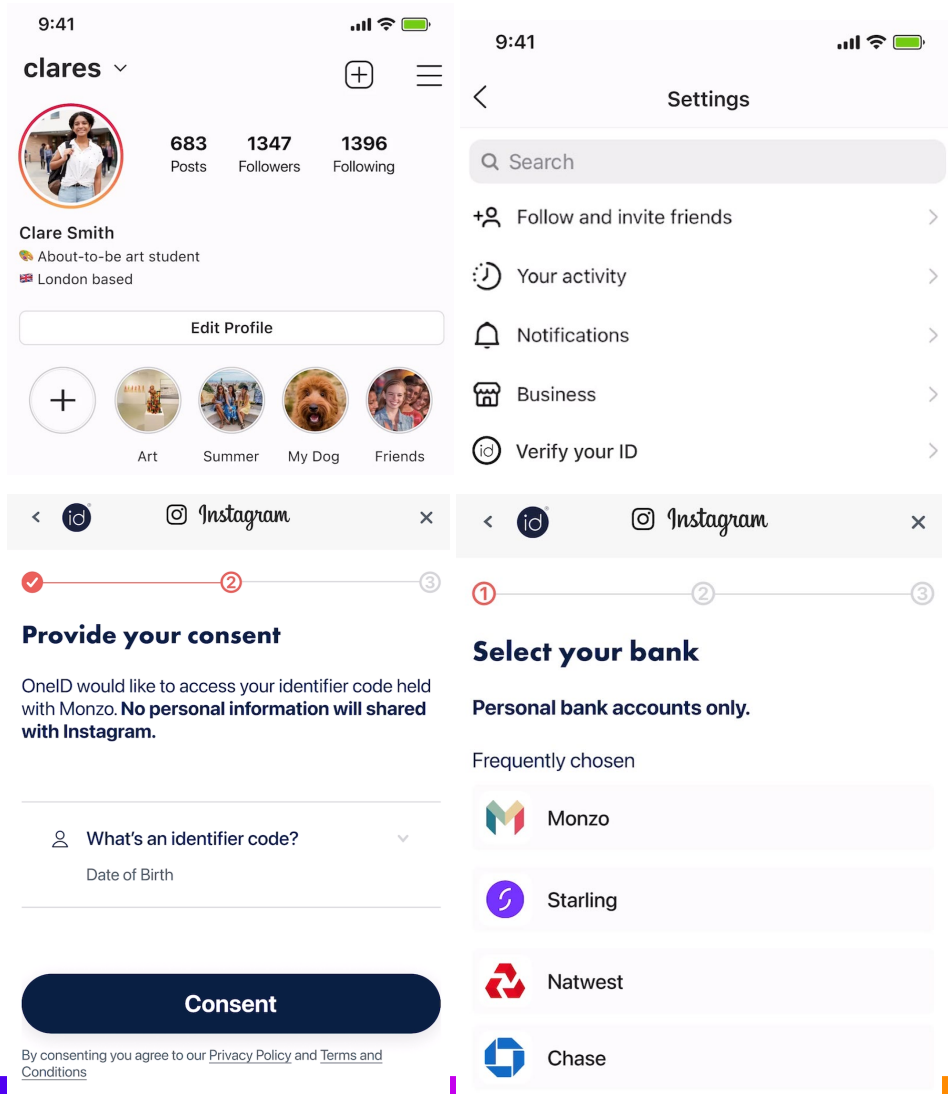

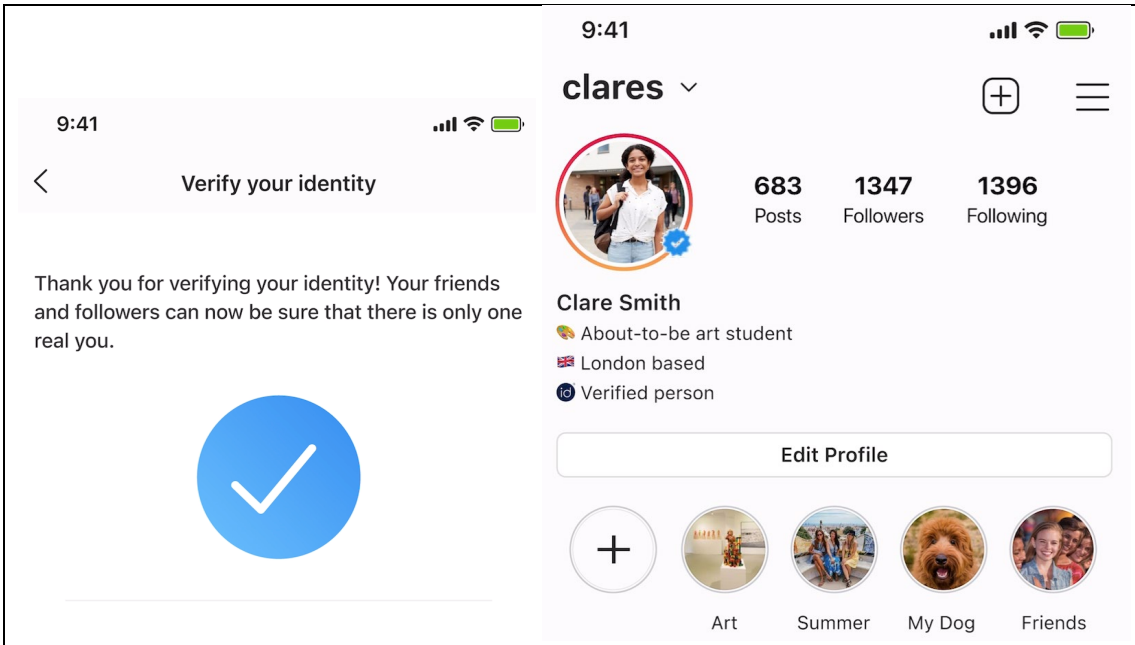

(a) features offered to users to reduce the likelihood of them encountering content they do not wish to see

Response: Social media platforms should provide options for their users to only see content from selected categories, e.g., adults should be able to turn off 'harmful but legal' content.

(b) features offered to users to alert them to the presence of certain categories of content

Response: Platforms sometimes alert users to harmful material, but with an option to click and view the material anyway. If the user is a child, the content should not be placed into the user's feed, and certainly shouldn't have any 'override' button.

(c) features offered to users to enable them to control their interactions with different types of users (e.g., non-verified)

Response: Social media platforms should provide options for their users to only see content from, and interact with, verified users. This would reduce the amount of 'trolling', harmful and abusive content, and if the user is making a payment to another one, this will reduce fraud.

(d) whether certain features are particularly valued or of use to users with protected characteristics, or by users likely to be affected by encountering relevant content

Response: None

Is this response confidential? (if yes, please specify which part(s) are confidential)

Response: No

Your response – User identity verification duties

Question 31 and 32: Circumstances where user identity verification is offered and how

For all respondents

Question 31: What kind of user-to-user services currently deploy identity verification and in what circumstances? In your response to this request, please provide information relating to (a) – (c) where relevant.
Response: Some dating platforms are starting to realise that user identity verification makes their platforms safer for their users, and are starting to deploy solutions.
(a) the ways in which these identity verification methods are beneficial, both to the user and to the service
Response: For the user, they feel safer using the service. For the service provider, the increased user protection demonstrates that the platform cares about their users, and wants to build a good reputation to enable them to win against competitors who do not protect users.
(b) what documentation you understand to be necessary for different types, or levels, of identity verification on user-to-user services
Response: The UK ID framework enables four 'levels of confidence' of a user's identity; low, medium, high and very high. The low and medium levels are probably the most relevant for platforms in scope of the OSA. Data minimisation can be applied at any level, e.g., to return the minimum amount of data for the use case, which could just be an identifier number (specific to the user on that platform). The ID framework caters for electronic evidence of ID as well as document-based solutions. OneID uses a bank account as the primary source of ID evidence, and this scores the same 'strength' score as a driving licence.
(c) whether you believe there are there any other circumstances where identity verification should be offered on user-to-user services.
Response: We believe that the scope of platforms that should offer identity verification should be broader that the current platforms identified, to maximise the customer protection that can be offered online, and establish 'verified user' as a widely known concept that citizens should look for online, and be able to check who is doing the verifying. This would help to communicate the Department of Science, Innovation and Technology's ID framework more broadly.
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Question 33: Cost and effectiveness of these methods

For all respondents

Question 33: Please share any information about the costs and the effectiveness of identity verification methods In your response to this request, please provide information relating to:

- (a) – (d) where relevant for all respondents, and
- f) and g) where relevant for providers of user-to-user services that provide some types of identity verification for individual adult users.

Response: OneID's solution is both cost-effective, and highly effective in proving a user's age is over a particular boundary.

(a) any insight into the cost of identity verification methods, including set-up and on-going costs, in terms of employee time and any other material costs, as well as any intended and unintended impacts on services more broadly

Response: OneID has no set up fee, is a simple API integration based on open standards. Our pricing is a 'software as a service' pay as you go model with flexible options for bundles, and typically is cheaper than document-scanning solutions, whilst also being more effective.

(b) how effective these identity verification methods are in verifying the identity of a user for the particular purpose for which verification is carried out

Response: We believe that OneID is the most effective solution in the market across all four criteria of 'highly effective' in the Ofcom framework, scoring 99.99% or 100% across each:

Accurate

The criterion of technical accuracy refers specifically to how an age assurance method can correctly determine the age of a user under test lab conditions.

Bank-based age verification is the most accurate, 'strict', which means it is at least 99.99% accurate. It also means that anyone can prove they are 18 on their actual birthday, so no need for a tolerance level or 'challenge 25' scheme; it's a binary 'over 18' result based on a verified date of birth.

Robust

The criterion of robustness describes the degree to which an age assurance method can correctly determine the age of a user in unexpected or real-world conditions.

As bank-based age verification does not vary between test lab and real-world conditions, it also scores a 'strict' level of accuracy, which means it is at least 99.99% accurate. OneID uses an API, so it doesn't need to capture audio or video, so does not have problems relating to environment.

Reliable

The criterion of reliability describes the degree to which the age output from an age assurance method is reproducible and derived from trustworthy evidence.

OneID is both reproducible (user can access their online banking each time, and this is not dependent on an AI decision) and derived from trustworthy evidence (the users bank). This would align with the top category in reliability. OneID also has a 100% uptime status https://status.oneid.uk/

Fair

The criterion of fairness describes the extent to which an age assurance method avoids or minimises bias and discriminatory outcomes.

Bank-based age verification is 100% fair – anyone who has online banking can access it. OneID uses bank-supplied Strong Customer Authentication, rather than using facial biometrics directly, so issues such as skin-tone bias for AI image analysis, are not an issue for us. OneID would score in the top range for all 4 criteria for 'high effectiveness'. Accuracy is independently assessed in the DRCF-commissioned ACCS report, the other metrics are self-assessed.
(c) any other benefits or unintended consequences from these schemes existing
Response: Benefits for corporates using digital ID include; better customer experience on onboarding (leading to more users/sales), no form-filling, can remove passwords/improve account security, better data quality (no typo errors), lower fraud.
(d) the safeguards necessary to ensure users' privacy is protected
Response: Each digital ID service provider is assessed against the ID framework requirements to ensure data privacy. OneID goes even further by using the user's bank as the storage and protection layer in our model – we do not store or monetise customer data.
For providers of user-to-user services that provide some types of identity verification for individual adult users
(e) any unintended consequences of implementing identity verification, such as the impact this may have on your site's ecosystem
Response: Providers to answer
(f) how you envisage your service operating in the digital identity market, bearing in mind moves towards cross-industry and federated identity schemes
Response: Providers to answer
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Question 34 and 35: User attitudes and demand for identity verification on user-to-user services

For all respondents

Question 34: What are user attitudes and demand for identity verification on user-to-user services? In your response to this request, please provide information relating to (a) – (d) where relevant.
Response: There is a large demand for identity verification to prevent problems arising from online platforms being anonymous. There are stories in the press every day about online safety issues and fraud cases which could have been prevented by having identity verification for online accounts and user empowerment features to enable unverified accounts to be filtered out.
(a) whether they value verification being offered on a service

Response: Many users would value verification, and this doesn't necessarily mean sharing more data with the platforms; it could be a 'verified human' identifier, with an IDSP holding the actual identity in case enforcement becomes necessary.
(b) whether verification influences user behaviour, such as whether they perceive identity verification to signify authenticity
Response: If verification is done properly to a defined set of rules, e.g., the DSIT ID framework, users would have confidence that a verified user was authentic. Where 'verification' is done only to the platform's rules, e.g., X's blue tick for making a card payment, this leads to no confidence and problems with fraudsters and impersonators having verified accounts. We believe that user verification should also be offered to all users for free, with the platform payment the cost. Using verification as a paid 'premium' feature leads to minimal use, and also the costs for the user are typically not in proportion to the costs of verification.
(c) attitudes towards non-verified, anonymous or pseudonymous users and the willingness to engage with them
Response: OneID endorses an approach to maintain user choice and access to anonymous accounts where this adds value, e.g., for journalists, whistle-blowers, authoritarian states etc. Users should be free to choose who they interact with, but if making payments, users should be advised by the platform to only trade with verified users. Platforms could also monetise this approach by enabling customer protection and safe trades (and ultimately take some liability for the trade, the cost of which would be covered by the business growth).
(d) who you deem to be 'vulnerable' in terms of verifying their identity online – for example, whether this includes users unable to access or less likely to hold identification documentation, and those who may become vulnerable by displaying their identity to other users.
Response: See above – users should be able to remain anonymous. Platforms should also have a duty of care to protect those who are identified as more vulnerable (such as children).
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Your response – Fraudulent advertising

Questions 36 – 42: Overarching considerations

For all respondents

Question 36: Please provide evidence of the following:
(a) The most prevalent kinds of fraudulent advertising activity on user-to-user and search services (e.g. illegal financial promotions, misleading statements, malvertising)
Response: We do not have empirical evidence of the prevalence, but these fraud types are frequently reported and fraudulent content is readily available on some online platforms.

(b) The harms associated with different kinds of fraudulent advertisements, the severity of such harms, and, if relevant, how this varies by user group
Response: UK Finance publishes fraud statistics every 6 months, that show for instance that investment fraud can lead to the biggest losses per case.
(c) The key challenges to successfully detecting different types of fraudulent paid-for advertising, and how these challenges can be minimised or resolved
Response: Platforms have started to check that advertisers have the required permissions to offer the financial product that are being advertised, but this is done manually and therefore can be error-prone or 'gamed' by the users placing adverts. A digital ID would make this process more secure and effective, by removing weaknesses in the manual human approach.
(d) The prioritisation of suspected fraudulent advertising within all categories of harmful advertising queues, e.g. account verification, user reports, appeals
Response: No comment
(e) The proportion of fraudulent advertisements that are currently estimated to remain undetected by services' systems.
Response: No comment
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Question 37: What technological developments aiding the prevention/detection of fraudulent advertisements do you anticipate in the coming years, and how costly and effective do you expect them to be? What are the challenges/barriers to their development?
Response: Better availability of trusted data from secure, 'authoritative' data sources. E.g., if the FCA were to provide an Open Banking-like feature for an individual or company to digitally share their FCA register permissions via a secure API, this could be connected to a digital ID to enable more secure verification of permissions to make the platform checking process more robust. This would also make it harder for fraudsters to 'clone' firms by impersonation. The technology (OIDC, FAPI) and process of Open Banking (regulated/certified 3rd parties, trust registries, digital ID of prividers) should be applied to enable more 'smart data' to be usable in secure processes.
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Question 43: Proactive technology

For all respondents

Question 43: Please provide any evidence you have regarding proactive technologies which could be used to identify fraudulent advertising activity.

In particular, we are interested in information related to the following points:
(a) The kinds of proactive technology which are/could be applied to identify or prevent fraudulent advertising
Response: Digital ID checks using certified vendors from the DSIT 'Digital Verification Services' DVS register.
(b) A brief description of how these technologies are/could be integrated into the service
Response: OneID has a certified service, that is a simple API integration, based on the global open standard 'OpenID Connect' (OIDC). This typically takes a developer only a few hours to connect. We also have an SDK – code to enable the OneID button to appear on any digital journey.
(c) The effectiveness, accuracy and lack of bias of such technology (including compared to alternative proactive and non-proactive methods) in relation to detecting fraudulent advertising and accounts which post fraudulent advertising material
Response: Bank-based ID does not have any bias issues; anyone who has access to online banking can use it to access their bank account.
(d) How proactive technologies are maintained and kept up to date
Response: OneID is maintained constantly, and recertified annually against the DSIT framework.
e) Information related to the associated time and/or costs for set-up, operation, and human review
Response: Approx. 2 hours to integrate.
f) The cost of integrating such technologies: (a) for the first time; and (b) when updating these technologies over time
Response: OneID has no onboarding cost, and a market-leading 'pay as you go' pricing model.
g) Whether there are cost savings associated with these technologies
Response: Yes (vs. other document-scanning solutions)
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Question 44: Advertising onboarding and verification

For all respondents

Question 44: Please provide any evidence you have regarding the processes for advertiser onboarding and verification related to protections against fraudulent advertising. In your response, please indicate whether these processes are currently implemented in respect of services which are in scope of the Act or whether they stem from another sector In particular, we are interested in information related to the following points:
(a) The criteria which advertisers are verified against, including documentation/evidence used to support verification, and what advertisers are required to declare

Response: Advertisers should be able to use a digital source of ID evidence to prove their FCA permissions. This would increase security.
(b) The role of (a) automated processing and (b) human processing in the verification process, and how they interact
Response: Digital ID enables a more automated process, and is less prone to errors.
(c) The costs associated with advertiser verification and how those costs vary as scale increases
Response: Pay as you go, can lead to decreased costs with increased volumes.
(d) The percentage of advertiser accounts that are verified
Response: No comment
e) Whether advertisers are permitted to publish advertisements on the service while the verification process is ongoing
Response: Advertisers should be verified first, but if a digital ID is used this could be done at the time of placing the advert, leading to no delays in verification.
f) Whether there are additional/specific verification checks for advertisers placing adverts of certain kinds or targeting certain audiences, such as about specific products or services, or targeting users under the age of 18
Response: Advertisers advertising FS products should have the relevant FCA permissions to offer those products.
g) Whether the verification of an advertiser account expires after a certain amount of time or certain activity, such as when advertisers make changes to their account or profile
Response: A digital ID service could notify the platform if/when an advertisers permission changes, e.g., they lose FCA permissions to offer a product.
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No

Question 49: Generative AI and deepfakes

For all respondents

Question 49: Please provide any evidence you have regarding the impact of generative AI developments and deepfakes on the incidence and detection of fraudulent advertisements on services in scope of the Act. In particular, we are interested in information related to the following points:
(a) The frequency of deepfake fraudulent advertisements' occurrence, in absolute terms and/or as a proportion of all fraudulent advertisements, and how you expect this to evolve in the future
Response: A bank-based digital ID does not have the attack vector of deepfake faces or document ID evidence, so is safer than document scanning and selfie techniques to verify advertiser's ID.

Content authenticity techniques such as the standard C2PA can be applied to verify the content's authenticity.
(b) What methodologies/technologies are currently employed to detect fraudulent advertisements which include deepfake or otherwise AI-generated content, and the effectiveness of these tools
Response: No comment
(c) Whether detection technologies are developed in-house or acquired from a third-party, and how long it takes to develop and/or integrate those tools into wider systems
Response: No comment
(d) The accuracy of detection methods, including true positive and false positive rates
Response: No comment
(e) The costs associated with the development/acquisition and deployment of these detection mechanisms
Response: No comment
(f) The types of deepfake or AI-generated content (in terms of either media type or subject) in fraudulent advertisements that are most difficult to detect i) via automated processes, ii) by human moderators, iii) by service users
Response: No comment
Is this response confidential? (if yes, please specify which part(s) are confidential)
Response: No