Consultation coversheet

Basic details

Consultation title: Fraudulent Advertising Codes Consultation
To (Ofcom contact):
Name of respondent: [redacted]
Representing (self or organisation/s): Advertising Association
Address (if not received by email):

Confidentiality

Please tick below what part of your response you consider is confidential, giving your reasons why

> Nothing ☒
> Name/contact details/job title ☐
> Whole response ☐
> Organisation ☐
> Part of the response ☐

If you selected 'Part of the response', please specify which parts:

If you want part of your response, your name or your organisation not to be published, can Ofcom still publish a reference to the contents of your response (including, for any confidential parts, a general summary that does not disclose the specific information or enable you to be identified)?

Yes ☐ No ☐

Declaration

I confirm that the correspondence supplied with this cover sheet is a formal consultation response that Ofcom can publish. However, in supplying this response, I understand that Ofcom may need to publish all responses, including those which are marked as confidential, in order to meet legal obligations. If I have sent my response by email, Ofcom can disregard any standard e-mail text about not disclosing email contents and attachments.

Ofcom aims to publish responses at regular intervals during and after the consultation period. If your response is non-confidential (in whole or in part), and you would prefer us to publish your response only once the consultation has ended, please tick here.

Name: [redacted]

AA response to Ofcom's consultation on the Fraudulent Advertising Codes of Practice

About the Advertising Association

The Advertising Association promotes the role and rights of responsible advertising and its value to people, society, businesses, and the economy. We bring together companies that advertise, their agencies, the media and relevant trade associations to seek consensus on the issues that affect them. We develop and communicate industry positions for politicians and opinion-formers, and publish industry research through advertising's think-tank, Credos, including the Advertising Pays series which has quantified the advertising industry's contribution to the economy, culture, jobs, and society.

The membership of the Advertising Association is very broad and includes the associations representing industry sectors, such as the advertisers (through the Incorporated Society of British Advertisers), the agencies and advertising production houses (through the Institute of Practitioners in Advertising and the Advertising Producers Association), all the media (from broadcasters and publishers, cinema, radio, outdoor and digital), advertising intermediaries and technology providers (which include platforms and the IAB UK), market research (through the Market Research Society) and marketing services (through the Chartered Institute of Marketing).

Table of Contents

Volume 1: Context 5

Online advertising ecosystem 5

Is there any information or evidence you hold that could enhance our understanding of the online advertising ecosystem? 5

Causes and impacts of fraudulent advertising 7

Do you have any comments on Ofcom's assessment of the causes and impacts of fraudulent advertising? 7

Do you agree with the examples of risky characteristics we have provided? 8

Approach to Codes 8

Do you agree with our proposed approach? Please provide any arguments and supporting evidence. 8

Combined impact assessment 8

Do you consider the overall burden on the service providers in scope of our proposed measures to be proportionate? 8

Volume 2: Risk, Governance and Control 10

Advertising intermediaries (FAU/FAS K1) 10

Do you agree with our proposal? Please provide any arguments and supporting evidence. 10

We have included two illustrative examples of when and how this measure may be applied, and what 'all reasonable endeavours' might look like. Do you have any comments on these examples, and do you have any additional examples of how providers could work with or require intermediaries to implement safety measures in circumstances where they have insufficient control? Please provide any arguments and supporting evidence. 12

Fraud indicator assessment (FAU/FAS B1) 13

Do you agree with our proposal? Please provide any arguments and supporting evidence. 13

Do you think there are other information sources providers should consider in order to understand how fraudulent advertising manifests on their service? If so, what information and why would they need to consider it? 13

Governance and accountability (FAU/FAS A1–A6) 14

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence. 14

Testing advertisement generation tools (FAU/FAS F1) 14

Do you agree with our proposal? Please provide any arguments and supporting evidence. 14

Volume 3: Ensuring account integrity 16

Account checks and actions (FAU/FAS H1) 16

Do you agree with our proposal? Please provide any arguments and supporting evidence. Where your feedback relates to a specific check we are proposing providers should carry out, please make this clear. 16

Are you aware of any other account checks and actions providers could do to protect users from fraudulent advertising? Please provide any arguments and supporting evidence. 17

Preventing fraudulent financial services advertising (FAU/FAS H2) 17

Do you agree with our proposal? Please provide any arguments and supporting evidence. 17

Countering account takeover (FAU/FAS H3, H4) 17
Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence. 17
Do you have any evidence about account takeover that could inform our assessment, including how frequently it occurs or its impact? 18

Advertising bans (FAU/FAS H5) 18
Do you agree with our proposal? Please provide any arguments and supporting evidence. 18
Are you aware of any additional exceptional circumstances in which it may be proportionate for providers to reverse advertising bans? 19

Account appeals (FAU/FAS J1–J4) 19
Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence. 19

Volume 4: Moderation 21

Advertising moderation (FAU/FAS C1–C7) 21
Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence. 21

Terms of service and publicly available statements (FAU G1–G2 / FAS G1–G2) 22
Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence. 22

Advertising complaints (FAU/FAS D1–D7) 22
Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence. 22

Ad libraries (FAU/FAS E1) 23
Do you agree with our proposal? Please provide any arguments and supporting evidence. 23
Do you agree with the principles, functionalities and information categories which we have included in the proposed measure? Please explain your reasoning and, if possible, provide supporting evidence. 23
Do you have any evidence or insights on the process of building an ad library from scratch or adjusting an ad library to align with the proposed measure, and how long that process would take? 24

Volume 1: Context

Online advertising ecosystem

Is there any information or evidence you hold that could enhance our understanding of the online advertising ecosystem?

Ofcom's account of advertising pathways (owned-and-operated, open-display and direct deals) is a reasonable description, and we note it draws substantially on the CMA's 2020 market study, the SparkNinety/DCMS work, and the AA's own AA/WARC expenditure data. We offer commentary on seven aspects.

1. "Degree of control" is a gradient, not a binary, and hybrid pathways are the norm. Ofcom's two-pathway model presents owned-and-operated and open-display as relatively distinct, with a provider either able or unable to apply a measure. In practice, the largest categorised services run multiple pathways simultaneously for the same inventory, and the degree of control a provider has over placement varies continuously across a spectrum rather than falling into two separate boxes. This matters because the Codes are based on "the degree of control a provider has in relation to the placement" (the trigger for the FAU/FAS K1 intermediaries carve-out and for the "except where K1 applies" qualifier on most measures). A binary framing risks either under-applying measures where a provider has more control than the model assumes or requiring "reasonable endeavours" against intermediaries where meaningful control is genuinely absent. We would encourage Ofcom to articulate control as a gradient and to test the K1 boundary against real hybrid configurations. We outline further below why contractual mechanisms are the most effective and practical way for providers to fulfil their 'reasonable endeavours' obligations.

2. The buy-side and the agency/reseller layer are under-represented. Section 3 is largely sell-side (publisher/inventory) centric. It underplays the structure of the demand side: advertisers, media agencies (including holding-company networks, independents, and the growth of in-housing), agency trading desks, resellers and affiliate networks. The account-checks measure (FAU/FAS H1.3(a)) requires verifying that the advertising-account holder "works or acts for, or on behalf of" the advertised party, and the ecosystem section does not convey how routinely legitimate advertising is placed through layered agency and reseller relationships in which the account holder is several steps removed from the advertised brand. A fuller picture of the buy-side is necessary to assess the friction that H1 would impose on legitimate advertisers.

3. The advertiser base is heterogeneous and dominated by a long tail of SMEs. The population of advertisers using advertising services is vast and long-tailed, comprising very large numbers of small and micro-businesses self-serving through ad-manager tools. Any onboarding friction, verification requirement or moderation false-positive is therefore multiplied across a very large and disproportionately SME advertiser base. The ecosystem section would be strengthened by recognising this asymmetry between the sell-side and buy-side, because it is central to the proportionality assessment: the providers in scope are few and large, but the businesses that bear the friction are numerous and small.

4. The existing industry anti-fraud, brand-safety and verification infrastructure is significantly underweighted. Section 3 mentions blocklists, OpenRTB, Confiant and IAB Tech Lab's buyers.json, but it does not convey the breadth of the established cross-industry infrastructure already deployed against fraudulent and non-compliant advertising. A fuller account should recognise: the IAB Tech Lab transparency standards (ads.txt, app-ads.txt and sellers.json); the IAB UK Gold Standard; the Trustworthy Accountability Group (TAG); independent third-party verification vendors (such as IAS and DoubleVerify); the industry's brand-safety and trading-standards frameworks; the ICC/GASA Best Practice for combatting scams in advertising; and the ASA/CAP self-regulatory system for advertising content. We recognise that Ofcom does draw on parts of this infrastructure elsewhere, notably in the intermediaries measure (Volume 2, Section 2), which references the ASA Intermediary and Platform Principles, the IAB UK Gold Standard, the NCSC guidance and the Online Advertising Taskforce. The gap is therefore one of consistency given that the ecosystem overview in Section 3 does not map this infrastructure. This matters because Ofcom's understanding of what providers and intermediaries can already see and act on is more capable than what Section 3 conveys, which bears on what "reasonable endeavours" under K1 should require.

5. The scope of "paid-for advertisement" is narrower than a simple "paid versus organic" split suggests. Under s.236 of the Act, an advertisement is "paid-for" only where the provider of the service is paid for it and controls its placement under the advertising contract. Because the test is whether the platform is paid, most influencer and affiliate content falls outside it: where a brand pays a creator to post, or an affiliate is paid by an advertiser or network, the platform itself receives nothing, so the advertisement is not "paid-for". We would welcome Ofcom confirming this expressly. The exception to this would be paid amplification: where the platform is paid to boost a creator's post, that post comes back into scope.

Two boundary questions remain, on which guidance would help:

- Paid-for advertising versus user-generated content (Category 1 only). "Boosted" content can be both paid-for and user-generated, and Ofcom leaves that call to providers (3.11). This does not arise for Category 2A, where search advertisements are not user-generated content.
- Who the contracting parties are, in open-display chains. Where there is no single, clean contract, it can be unclear who "the parties entering into the contract" are (limb (b)) and how indirect payment is treated (limb (a)).

6. The incentives analysis (3.7) should be refined to distinguish the sell-side platform from the wider industry. Ofcom argues that providers have "weak incentives" to tackle fraudulent advertising because less fraud means fewer ads and less revenue. However, it does not describe the incentives of the wider advertising industry. Legitimate advertisers, agencies and responsible media owners have a strong and direct interest in eliminating fraudulent advertising, because it erodes consumer trust in advertising as a whole, diverts spend, and damages brand safety. The AA and its members are active on exactly this agenda, including

through the Trust Action Plan and the DCMS Online Advertising Taskforce. We would encourage Ofcom to frame the industry as an ally with aligned incentives, and to reserve the "weak incentives" analysis for the specific sell-side dynamic to which it applies.

7. Cross border nature of online advertising. The online advertising ecosystem is global, and that UK-specific obligations would interact with international frameworks (including the EU Digital Services Act ad-repository regime and cross-border industry standards). An accurate ecosystem picture is a cross-border one.

Causes and impacts of fraudulent advertising

Do you have any comments on Ofcom's assessment of the causes and impacts of fraudulent advertising?

We agree that fraudulent advertising is a serious harm with significant consumer and economic impact. At the same time, we do ask Ofcom to strengthen the evidential basis for a claim that is central to the proportionality case for the whole intervention: that advertising is the second most common way fraudsters reach victims online (see 1.6 and Section 4 introduction, and referenced at 6.4). Our chief concern is with what that evidence does and does not measure.

- It does not isolate the regulated category. Ofcom's own footnote acknowledges that the underlying research relates to advertisements respondents saw on websites, apps, social media and video-sharing platforms, and that it "does not specifically relate to paid-for advertising as defined by the Act". The headline ranking therefore sweeps in organic and non-paid promotional content that falls outside s.236 and outside the Codes. This means it is likely to overstate the incidence attributable to the paid-for advertising the Codes actually regulate.
- It measures contact frequency, not harm. The 20% figure is the share of victims reporting contact via a channel, not a measure of financial loss. Proportionality should be tested against advertising's share of actual consumer losses, which is the metric that captures harm.
- The supporting figure is dated and not UK-specific. The 11% figure derives from a 2020 European Commission survey that is EU-wide rather than UK-specific and now five years old, on a different definition of "online advertisements" again.
- Potential over-attribution. The document cites the Home Office estimate that all fraud against individuals cost £9.2 billion. That figure is for all fraud, not advertising's share, and there is a risk that the reader associates the whole sum with advertising.

Accordingly, Ofcom should publish the share of fraud losses attributable to paid-for advertising as defined by the Act and calibrate the proportionality of the measures to that figure.

Data from the ASA's Scam Ad Alert system, kindly supplied to the AA, records the following:

Year	Scam reports received	Scam Ad Alerts issued
2023*	1,877	152
2024	1,691	177
2025	2,589	169

2026 (to date)	2,969	67

*These figures are for 12 months up to October 2023

These figures are shown to illustrate the limits of the current evidence base. It is important to note these are counts of reports received and referral alerts issued; they are not measures of confirmed fraud, of consumer losses, or of incidence on any categorised service, and they do not isolate paid-for advertising as defined by the Act. Nor do they move in a single clear direction. They cannot, in short, answer the question the proportionality case depends on, and that is precisely the evidence gap we are asking Ofcom to close.

Do you agree with the examples of risky characteristics we have provided?
No specific comments on this question.

Approach to Codes

Do you agree with our proposed approach? Please provide any arguments and supporting evidence.
We support a code-based approach in principle. We have two structural comments.

First, we note that the Category 1 and 2A Codes are substantively identical, differing almost entirely in the adaptation of user-to-user concepts to a search context. Hence, the drafting gives the impression that a search-specific assessment may not have been conducted for each measure. This matters because a measure's impact can differ sharply in the search context: the ad library's keyword-transparency requirement is the clearest example, since in paid search the keyword is the product. The presence of unamended Category 1 cross-references (there are two leftover "FAU H2.3(b)", instead of "FAS H2.3(b)", references inside limbs (c) and (d) in FAS H2.4) and an incomplete definition (FAS C1.2) reinforces this concern. We ask Ofcom to confirm, measure by measure, that the search application has been separately justified, in particular for the ad library, financial services verification, and account checks measures, where the operational realities of search advertising differ materially from user-to-user advertising.

Second, a service that is both Category 1 and 2A is subject to both Codes simultaneously (s.38 for its user-to-user service, s.39 for its search service). The two Codes appear to require parallel but separate compliance architectures. The impact assessment should therefore address the cumulative burden on combined services and whether a single, harmonised compliance route could satisfy both duties without duplication.

Combined impact assessment

Do you consider the overall burden on the service providers in scope of our proposed measures to be proportionate?
In several respects, no. While the providers directly in scope are large, the burden of many measures falls in practice on the entire advertiser base, including SMEs, and on intermediaries that are not themselves the regulated party. The specific proportionality concerns include ad

libraries; intermediaries; account checks; financial services verification; advertising bans; and moderation. Taken together, we are concerned that the cumulative effect is to impose friction, cost and disclosure obligations on legitimate advertising that are not proportionate to the marginal fraud-prevention benefit, and that some of that benefit is achievable through less intrusive means. The clearest example is the ad library, where the assessment weighs the fraud-detection benefit and providers' build costs but does not weigh the commercial cost to advertisers of publishing targeting, reach and keyword data publicly (see our answer on Ad libraries). We also note a proportionality concern in the other direction: because the duties attach only to categorised providers, fraudulent advertising served through content-recommendation and native-advertising networks onto non-categorised publisher inventory appears to fall outside the regime, risking displacement of harm onto the long tail of smaller publishers (see our answer on Advertising intermediaries).

Volume 2: Risk, Governance and Control

Advertising intermediaries (FAU/FAS K1)

Do you agree with our proposal? Please provide any arguments and supporting evidence.

The AA supports the underlying principle that safety outcomes should not vary by pathway, and we welcome that Ofcom has grounded this measure in existing industry and government initiatives, including the ASA's Intermediary and Platform Principles pilot, the IAB UK Gold Standard, the NCSC's malvertising guidance for brands, and the DCMS Online Advertising Taskforce (including the new ad-fraud-and-standards working group established in November 2025). This is the right foundation, and we address below how the measure could better build on it.

We accept Ofcom's evidence (2.3, 2.34) that most categorised services place advertising predominantly through owned-and-operated supply chains, and that this measure is therefore expected to apply only in limited circumstances. We do not, on that basis, maintain a broad claim that the Codes extend wholesale across the programmatic supply chain. Our residual concerns are these:

- The "sufficient/insufficient control" threshold is undefined. Ofcom expressly declines to define it (2.5), on the basis that it is context-dependent. The undefined threshold, combined with a detailed record-keeping obligation to justify insufficiency component-by-component and pathway-by-pathway (2.20 and footnote 13), creates real compliance uncertainty at exactly the boundary where providers must decide whether K1 or the underlying measure applies. We ask Ofcom to provide clearer indicative criteria or a wider set of worked examples for where the boundary falls.

- The measure still imposes obligations on non-regulated intermediaries, and the "reasonable even if revenue falls" position is significant. K1.3's "requiring… intermediaries" mechanism pushes measure-equivalents onto SSPs, ad servers and other intermediaries that are not the regulated party. Ofcom states that steps are "likely to be reasonable even if they result in a reduction in a provider's revenue" (2.13), which sets a demanding threshold. We ask Ofcom to clarify how contractual "requiring" is intended to work where an intermediary serves many platforms with potentially conflicting requirements, and how the accountability, appeal and safeguard architecture that protects advertisers travels down the chain when the acting party is an intermediary. In any case, we recommend that Ofcom clarify that a provider will have complied under this measure by negotiating and enforcing contractual obligations on the advertising intermediaries it engages with directly, including, where appropriate, obligations to flow equivalent requirements down to their own counterparties. Given that direct backend oversight of third-party compliance by service providers is practically impossible and disproportionate, we suggest that Ofcom recognise contractual agreements as the primary mechanism for providers to successfully

discharge this obligation. We agree with Example 2.1's focus on using contractual terms to obligate intermediaries to carry out necessary steps.

- The market-impact assessment rests on a contestable assumption. Ofcom acknowledges indirect costs on intermediaries and that some may pass these on as higher fees (2.51–2.52), but concludes that wider market effects are "likely to be limited" because the intermediary market serving Category 1 and 2A services is "small… and concentrated with larger intermediaries" (2.53). If that assumption does not hold, there are real competition and open-web risks.

We also make the following recommendations:

- Supply-chain accountability. Where an in-scope service's ads are delivered through third-party intermediaries, we recommend that Ofcom's guidance make clear that proportionate systems and processes include appropriate contractual safeguards with those partners. These could include audit rights, escalation and suspension rights, consequences for repeat failures, and timely provision of investigation data (such as buyer ID, advertiser name, domain, creative ID, deal ID, seat ID and reseller path), subject to appropriate data protection and competition law safeguards. The "reasonable procedures" defence under the failure to prevent fraud offence in the Economic Crime and Corporate Transparency Act 2023 offers a relevant model.

- Targeted supply-chain blocking. Once a fraudulent ad or buyer account has been confirmed, services should take prompt steps to suppress the associated domains and buyer IDs across the inventory they control. This should be based on clear evidence thresholds, with a route for appeal through the existing advertising appeals process, so that legitimate activity is not blocked where an issue is isolated, spoofed or misattributed.

- Repeat-offender intelligence sharing. Bad actors frequently rotate between domains, accounts, DSP seats and supply paths. Industry initiatives such as the OAT information-sharing pilot show how intelligence can flow faster across the ecosystem. We would welcome Ofcom recognising participation in such schemes as evidence of effective systems and processes.

- Mapping intermediaries by function. Our concerns above, and our comments on the ecosystem overview, point to a common gap. The Codes rely on providers determining which "equivalent measures" should apply to each intermediary in their supply chain, but neither the ecosystem overview (Section 3) nor the intermediaries measure (K1) sets out how intermediaries differ by function, what fraud risks each function presents, or which measures would be equivalent for each. Leaving that determination to individual providers means an intermediary serving several categorised services may face different, and potentially conflicting, requirements from each, with no common reference point for what is reasonable. We recommend that Ofcom map the main categories of intermediary (for example SSPs, DSPs, ad servers, ad networks, resellers

and content-recommendation networks) against the fraud risks each presents and the equivalent measures that would be proportionate for each, and publish this as guidance alongside the Codes. This would give providers and intermediaries a consistent baseline, reduce the compliance uncertainty at the K1 boundary, and support a more consistent level of protection across the supply chain. The AA would welcome the opportunity to support this work, drawing on the expertise of its members across the advertising supply chain.

We have included two illustrative examples of when and how this measure may be applied, and what 'all reasonable endeavours' might look like. Do you have any comments on these examples, and do you have any additional examples of how providers could work with or require intermediaries to implement safety measures in circumstances where they have insufficient control? Please provide any arguments and supporting evidence.

Ofcom's two worked examples (moderation, and account checks) are helpful and realistic. Two comments. First, in Example 2.2, the account-checks scenario assumes checks can be pushed to the demand side (media agency, advertiser ad server or DSP) via the provider's supply-side relationships. But the supply chain is often layered. A problem emerges if an SSP requires a DSP to act, when the provider has no relationship with that DSP and no means to compel it to act. We therefore ask Ofcom to confirm how far "requiring" is expected to propagate and what happens where it cannot. Second, both examples would be strengthened by referencing the existing transparency standards (ads.txt, sellers.json, the OpenRTB object, buyers.json) as concrete mechanisms through which "reasonable endeavours" can be delivered. This would anchor the measure to infrastructure the industry already operates.

It is currently unclear how the fraudulent advertising duties cover advertising served through content-recommendation and native-advertising intermediaries (e.g. "sponsored article" and "around the web" formats sometimes found on publisher and platform pages). This is a documented source of fraudulent and clickbait advertising.[1] Where such formats appear on a categorised service, we ask Ofcom to confirm that the intermediaries measure is expected to apply and how far "requiring the intermediary" is intended to propagate through the recommendation network. Additionally, the Which? evidence shows this harm arising through recommendation formats served programmatically, with different inventory shown to each user. Ofcom's own register of categorised services (published 30 June 2026, accessed 21 August 2026) confirms the coverage gap this creates. The fraudulent advertising duties attach only to Category 1 and Category 2A services, none of which is a content-recommendation or major native-advertising network. Where such advertising is served onto non-categorised inventory, no categorised provider sits in the chain, so this is a category of fraudulent advertising the duties do not reach at all. We do not suggest that any particular network is responsible for the underlying fraud (the networks concerned operate stated moderation and takedown policies),

[1] Scams, fake news and rip-offs lurking in sponsored ads (2022). Which? https://www.which.co.uk/news/article/scams-fake-news-and-rip-offs-lurking-in-sponsored-ads-aUWAn1M420Z3

and we recognise this partly engages the categorisation regime rather than the Codes alone. But the practical consequence is that the population of services carrying the fraudulent advertising duty is narrow, while a documented and significant vector of consumer fraud sits wholly outside it, and Ofcom should address this directly.

Fraud indicator assessment (FAU/FAS B1)

Do you agree with our proposal? Please provide any arguments and supporting evidence.

We broadly support a structured, evidence-led fraud indicator assessment. It is a sensible foundation for risk-based action. However, it is important to note a potential downstream effect. The assessment feeds the "reason to suspect" trigger that runs through moderation, complaints and account checks, so its breadth directly affects the volume of mandated reviews and the risk of over-removal (see Moderation, below). We ask that the assessment be designed to support accurate targeting rather than blanket suspicion. We recommend that Ofcom make clear that fraud indicator assessments should be reviewed outside the regular review cycle where there is a material change in risk, such as the launch of new AI ad-generation tools or the emergence of new fraud tactics. Generative AI and deepfake techniques continue to evolve quickly, and assessments should not be left out of date between scheduled reviews.

With respect to obligations concerning requesting information from third parties, we urge Ofcom to confirm that providers are only expected to initiate these requests when the information is not already held by the provider and is known to be possessed by the third party. Consistent with UK GDPR data minimisation principles, Ofcom should clarify that such requests must only occur when the data is necessary and proportionate to the specified safety duty.

To promote operational efficiency and align with Ofcom's commitment to a principles-based, outcomes-focused regulatory regime, the Code should clearly state that providers can use their existing OSA or DSA risk assessments, including their established illegal content and children's risk assessment frameworks, to inform their fraudulent ads risk assessments. The requirement to keep a separate record of the fraud indicator assessment should be removed.

Do you think there are other information sources providers should consider in order to understand how fraudulent advertising manifests on their service? If so, what information and why would they need to consider it?

The AA is currently leading a pilot which establishes a collaborative advertising fraud intelligence-sharing mechanism between Google, Amazon and Microsoft (via The Media Trust) using the Global Signal Exchange (GSE) platform operated by Oxford Information Labs. A neutral, third-party platform of this kind is needed for major participants in the online advertising ecosystem to overcome competition issues.

The logic is straightforward: no single participant has complete visibility across the entire ecosystem, so sharing information on threats they have each encountered builds collective

defence. The pilot enables participants to share signals, initially URLs associated with advertising campaigns, allowing them to identify new and overlapping threats and enhance their malicious URL detection and blocking capabilities. Eventually the pilot will be expanded to bring in more participants and increase the volume and range of data. On reaching that scale, the potential of the GSE can be leveraged through its capability to enrich data, identify trends, support law enforcement in tracking down the organisations behind these threats, and help companies strengthen their own defences.

We would encourage Ofcom to recognise cross-industry intelligence-sharing of this kind as a source providers should draw on in their fraud indicator assessments, and to support its development by providing a conducive regulatory environment.

Governance and accountability (FAU/FAS A1–A6)

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.
No specific comments on this question.

Testing advertisement generation tools (FAU/FAS F1)

Do you agree with our proposal? Please provide any arguments and supporting evidence.
We recognise the growing role of generative AI in ad creation and the potential for these tools to be misused, particularly to impersonate brands, public figures and institutions. We also note that fraudulent advertisements are driven by the intent of the bad actor and can readily be created without such tools, so testing advertisement generation tools addresses one route to fraud rather than the problem as a whole. The measure should therefore be proportionate to the incremental risk these tools present, and should not discourage providers from offering native creative tools whose output remains subject to their existing review processes.

Where platforms offer AI ad-generation tools (FAU F1 / FAS F1), we recommend that Ofcom's guidance make clear that proportionate systems and processes should include safeguards against those tools being used to generate content that impersonates brands, public figures or institutions. AI-generated creative should be subject to at least the same review as other ads, and should not become a route around existing checks. We recognise that safeguards will need to distinguish impersonation from legitimate uses such as authorised resellers and comparative advertising.

We note that the measure applies "regardless of whether the tool is made available on or off the service" (footnote 151). We ask Ofcom to explain how this is consistent with s.38(7) and s.39(7) of the Act, which confine the duties to the design, operation and use of a Category 1 or 2A service, and in particular whether, and on what basis, a tool the provider makes available off the service is caught. We understand that the measure does not capture third-party creative tools that an advertiser independently chooses (5.7–5.8), and would welcome confirmation of this.

Many providers of advertisement generation tools already test, log and evaluate model performance as part of their wider model governance. We ask Ofcom to confirm that providers can rely on these existing processes to meet the measure, to the extent that they address the misuse risks it targets, rather than building a separate UK-specific testing pipeline.

Where a provider makes available a tool built or operated by a third-party supplier, the testing and logging burden (including logging the supplier's design decisions) appears to fall on the provider. We ask Ofcom to confirm this is indeed workable in practice.

Volume 3: Ensuring account integrity

Account checks and actions (FAU/FAS H1)

Do you agree with our proposal? Please provide any arguments and supporting evidence. Where your feedback relates to a specific check we are proposing providers should carry out, please make this clear.

We support measures that keep known fraudsters and impersonators out of advertising accounts. However, there are some concerns about the additional friction this causes for legitimate advertisers, namely:

a) **Verifying the account holder works for, or on behalf of, the advertised party.** We welcome that Ofcom offers a flexible menu of verification methods (2.18–2.22). But the most concrete of them presuppose certain infrastructure that legitimate advertisers may not have. The Companies House check is unavailable to sole traders and unincorporated businesses, and the domain-email method would not work for many small advertisers that use free email services like Gmail or Outlook. Other jurisdictions may not have an equivalent to a searchable or reliable database such as Companies House, making it difficult to verify overseas advertisers targeting UK users in this manner. Applied uniformly, the check therefore risks excluding a wholly legitimate segment of the advertiser base and adding friction disproportionate to the risk many advertisers present. We ask Ofcom to make clear that the methods listed are illustrative rather than exhaustive, that providers may rely on other reasonable evidence, and that the depth of verification may be calibrated to risk, for example by reference to the provider's fraud indicator assessment. The most demanding checks should be reserved for higher-risk advertisers and categories, including where the advertised party is a brand, public figure or institution that is commonly impersonated. We also ask Ofcom to set out how sole traders and overseas advertisers are to be verified proportionately.

b) **Preventing banned account holders returning**. We support this in principle; our concerns about the underlying ban mechanism are under "Advertising bans" below. Ofcom should clarify that FAU/FAS H1.3(b) requires providers to take proportionate and reasonable steps to screen for banned individuals, rather than holding them to an unfeasible standard of guaranteeing exhaustive elimination.

c) **Risk-characteristic checks and restrictions**. Supported in principle, provided restrictions are proportionate and the false-positive rate on legitimate advertisers is monitored and minimised. Providers should have flexibility over which risk characteristics they check and how, calibrated to the risks identified in their fraud indicator assessment, rather than being expected to apply every listed check to every account. This avoids unnecessary operational burden and limits the collection of advertisers' personal data to what the risk justifies.

d) **Repeat checks on relevant change.** Reasonable in principle; we ask that "significant change" and re-check triggers be defined to avoid unnecessary re-verification of established advertisers. As with point (c), providers should have flexibility over the form a re-check takes.

Are you aware of any other account checks and actions providers could do to protect users from fraudulent advertising? Please provide any arguments and supporting evidence.
No specific comments on this question.

Preventing fraudulent financial services advertising (FAU/FAS H2)

Do you agree with our proposal? Please provide any arguments and supporting evidence.
The AA supports the objective of keeping unauthorised financial promotions out of advertising. We raise two points. First, there is duplication with the FCA regime. Section 21 FSMA already restricts who may communicate a financial promotion, and the FCA's financial promotion gateway already narrows the types of firm permitted to approve promotions for unauthorised businesses. Moreover, s.40 of the Act, which is a closed list of specified criminal offences, makes breach of the financial promotion restriction (FSMA s.25) and carrying on a regulated activity while unauthorised (s.23) predicate offences for the fraudulent advertising duty, as Ofcom itself confirms at footnote 81. An unauthorised financial promotion is already capable of being a fraudulent advertisement caught by the general moderation duty (C1–C2). H2 then asks platforms to run a separate, front-loaded verification of substantially the same authorisation status before an ad is encounterable. We acknowledge the rationale: H2 is preventive (verify before the ad is seen) where the general duty is reactive (detect and take down), and the platform check catches impersonation of an authorised firm, which the FCA gateway does not. Our concern is proportionality: front-loaded verification applies to all financial services advertisers, including the many legitimate authorised firms. We therefore ask Ofcom to set out the marginal benefit of the front-loaded check over the combination of the FCA regime and the reactive s.40 duty, and to ensure the measure is calibrated to the residual impersonation risk rather than imposing blanket friction on compliant advertisers.

Second, the definition of "financial services advertisement" as any ad that "promotes, or is capable of promoting, financial products or services" risks catching adjacent or borderline content that is not in substance a financial promotion, creating friction for legitimate advertisers. We therefore ask that the phrase "capable of promoting" be tightened.

Countering account takeover (FAU/FAS H3, H4)

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.
The AA supports account security and account-takeover reporting measures; protecting legitimate advertisers from account compromise is directly in members' interest. We note that there appears to be drafting error at FAU/FAS H4.3(c) and H4.4, which cross-reference "H2.2" where "H4.2" is evidently intended. We kindly ask Ofcom to correct this.

However, we recommend that Ofcom clarify that providers need only make security mechanisms available for all advertising accounts. Mandating their adoption is unnecessary and risks interfering with advertisers' independence in managing their own accounts.

Do you have any evidence about account takeover that could inform our assessment, including how frequently it occurs or its impact?
No specific comments on this question.

Advertising bans (FAU/FAS H5)

Do you agree with our proposal? Please provide any arguments and supporting evidence.
We support the banning of genuine fraudsters, and we welcome Ofcom's acceptance that fraud is an adversarial space, and its setting of the standard at "reasonable steps" that "materially reduce" return rather than eliminate it (5.32, 5.39). Particularly helpful is the acknowledgement of phoenixing, lifeboating and account takeover as return routes (5.35).

We also support Ofcom's principles-based approach to the technical means providers use to prevent banned individuals returning, rather than the imposition of prescriptive technical mandates. This allows providers to keep pace with the evolving tactics of bad actors and tailor their approach to the service. Providers that take reasonable steps to prevent banned individuals from creating and using new advertising accounts should not be held to a more prescriptive standard, provided their methods are effective both in preventing return and in avoiding the wrongful identification of legitimate advertisers, and the safeguards set out below are in place.

We do have concerns about effectiveness of the proposed measure, particularly against the determined actor, as well as the collateral impact on legitimate advertisers, and a gap the measure leaves open.

- Effectiveness against clean re-entry is asserted but not evidenced. The steps Ofcom proposes for preventing return rely on matching a new account to a banned one by shared characteristics: account names, registration details, content and behavioural similarity (5.38, 5.42, and the provider practice cited at footnote 193). That is effective against simple re-entry that reuses identifiers, but it is easily defeated by a determined fraudster who sets up a fresh legal entity, new registration details, new payment information and new creative, which leaves nothing to match against. Setting up such an entity is cheap and fast, particularly in jurisdictions with light company-formation checks, which also defeats the H1 verification the ban relies on to detect a returning actor. The effectiveness of the whole account-integrity system (H1 and H5 together) is therefore limited by the weakest global incorporation regime, which the Codes cannot reach. We ask Ofcom for its evidence that the "material reduction" it expects holds against re-entry via freshly incorporated entities specifically, as opposed to simple re-entry that reuses identifiers.
- The identification net creates a false-positive asymmetry. To catch phoenixes, the characteristic-matching net must be wide. But the wider it is, the more it catches legitimate

advertisers who happen to share characteristics with a banned actor, the same payment processor or company-formation agent, or an agency acting for many clients, while the sophisticated fraudster with a genuinely fresh entity still evades it. Ofcom does not address this collateral risk in the preventing-return sub-section, and the only safeguard is the account appeals route, which (as we note under Account appeals) does not appear to carry a defined determination or restoration timeframe. The Code should require explicit false-positive safeguards and a stated evidential threshold for attributing a new account to a banned person. Where identification is wrong, restoration through the account appeals route should be subject to the indicative timeframes we propose under Account appeals.

- The proxy route lowers the threshold for a ban. Because the measure applies to fraudulent advertising proxies as well as to full fraudulent-advertisement judgements (5.19–5.21), a ban can be triggered on the lower "suspicion plus terms-of-service breach" basis without a s.192 judgement. A legitimate advertiser removed through the over-cautious proxy route can therefore be banned on that basis, with the ban then propagating across the provider's services (5.3, 5.18). We ask that bans be reserved for full fraudulent-advertisement determinations, or that proxy-based bans carry enhanced appeal safeguards.
- A single determination triggers a ban, with no proportionality gradient. The measure bans whoever is determined responsible for a fraudulent advertisement, with no distinction between a serious repeat fraudster and an isolated or borderline case, and the ban then propagates across the provider's other Category 1 and 2A services. We believe that proportionality should be built into the trigger.
- The measure would be materially more effective if paired with cross-industry signal-sharing. Bans propagate only across a single provider's own services (5.3, 5.18). The Codes do not propose a mechanism for sharing known-fraudster signals between providers. A fraudster banned by one provider can re-enter another fresh, and each provider must rediscover the same actors independently. The single most effective response to phoenixing is cross-service and cross-industry intelligence sharing, so that a banned actor's fingerprints are recognised across the ecosystem. Bans siloed per-provider are inherently limited, and the cross-industry GSE signal-sharing pilot we describe in our answer on the fraud indicator assessment is a potential mechanism that gives the ban duty real effect.

Are you aware of any additional exceptional circumstances in which it may be proportionate for providers to reverse advertising bans?

No specific comments on this question.

Account appeals (FAU/FAS J1–J4)

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

We support an accessible appeals route, but the Codes do not have a defined timeframe. Account appeals must be determined "promptly" (6.20), and Ofcom cross-refers the meaning of "promptly" to the same internal-process examples as advertising appeals (6.22); restoration then follows determination (6.23–6.24). Neither is time-bound. Ofcom itself recognises that undue delay in reversing account restrictions after a successful appeal can have a

disproportionate impact, and that for many advertisers the ability to post paid-for advertisements to UK users is an essential part of their business (6.21). For a legitimate advertiser wrongly restricted or removed during a live campaign, downtime is unrecoverable commercial harm. Since Ofcom accepts this harm, it should introduce indicative timeframes for appeal determination and for restoration where an appeal succeeds. Providers should be able to exceed these where an appeal requires complex investigation, provided they record the reason and keep the advertiser informed of progress. Moreover, it should ensure the accuracy limb of the moderation performance targets (FAU/FAS C4.2(b)) carries genuine weight against the speed limb.

Ofcom should clarify that "appropriate action" for upheld appeals is limited to platform-level remedies (e.g. account and ad restoration) and excludes any form of financial compensation.

Volume 4: Moderation

Advertising moderation (FAU/FAS C1–C7)

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

We support effective moderation of fraudulent advertising. However, there is a structural risk of over-removal of legitimate advertising. This arises from the interaction of the low "reason to suspect" trigger (met by any complaint per D3.2 and any in-scope trusted-flagger report per D7.9), the "fraudulent advertising proxy" route (C1.4(b), Vol 4 2.10–2.11), which lets a provider act on suspicion plus a breach of its own terms of service without the fuller s.192 judgement, and the resulting reliance on automated moderation at scale (2.14). Ofcom engages this risk in its freedom-of-expression assessment (2.36): it accepts there is a risk of erroneous moderation. However, we believe that Ofcom should reconsider how it has discounted those risks.

- The "primarily commercial" justification understates the harm. Ofcom concludes the error risk is "justified" because paid-for advertising is primarily commercial (2.36). Commercial expression nonetheless attracts ECHR Article 10 protection, and wrongful removal inflicts real and sometimes unrecoverable economic harm on legitimate businesses, with a chilling effect on lawful advertising. The commercial character of the speech lowers neither the harm nor the standard of justification as much as the assessment implies.
- Ofcom's answer addresses the wrong case. Ofcom notes that fraudulent advertisements do not attract Article 10 protection, so their removal involves no interference (2.36). That is true but beside the point. The concern is the legitimate advertisement wrongly removed, to which that answer offers no protection.
- Ofcom cannot fully disclaim an incentive its measure creates. Ofcom attributes any over-restriction to providers' own discretion in how they set their proxies rather than to the measure (2.35, footnote 11). But the Codes structure the proxy route as the lower-cost compliance path, so the incentive is a feature of the design, not merely of provider choice.

Ofcom contends (footnote 11) that providers have commercial incentives not to over-remove, since doing so forgoes advertising revenue and generates moderation and appeal costs. This may be true for higher-value advertisers, but it will be less so for the long tail of lower-value advertisers, where the volume of ads and per-ad revenue does not justify careful individual review. The over-removal risk therefore concentrates precisely on the smaller and SME advertisers least able to absorb it.

We are not arguing against swift action on genuine fraud. We ask Ofcom to (i) treat the false-positive risk to legitimate advertisers as a substantive harm rather than one discounted by the "commercial" character of the speech; (ii) give real weight to decision accuracy (C4.2(b)) against the speed target; and (iii) introduce the defined appeal-determination and restoration timeframes we press for under Account appeals, since Ofcom currently leaves "promptly" entirely to providers' internal timescales (4.119) with no external standard. In the search context the remedy is de-listing rather than take-down, so a wrongly actioned ad is removed from

results while remaining live on the advertiser's own site; the commercial harm and the case for prompt restoration are the same.

Terms of service and publicly available statements (FAU G1–G2 / FAS G1–G2)

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.
On the terms of service and publicly available statement measures (FAU/FAS G1–G2), we do not object to transparency about the proactive technology used for fraud prevention. We note the sensible distinction the Codes draw between "terms of service" (Category 1) and "publicly available statement" (Category 2A), reflecting that search users are not in a terms-of-service relationship.

Advertising complaints (FAU/FAS D1–D7)

Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.
The AA supports accessible complaints routes and welcomes Ofcom listing the ASA as a recommended trusted flagger (4.143) recognising the ASA's existing role in receiving scam-ad reports from the public. The ASA's Scam Ad Alerts are issued on the basis that an advertisement is likely to be fraudulent, not merely misleading or irresponsible, and the substantial majority concern uncontentious impersonation of celebrities and companies; such referrals sit close to the s.40 fraud standard. We note that consumers can already report fraudulent advertisements to both the ASA and platforms.

Note: FAU D6/FAS D6 allows providers to disregard complaints that fall within their own definition of "manifestly unfounded". We recognise that providers need this flexibility, including to deal with coordinated or malicious complaint flooding, and we do not suggest a single industry-wide definition. However, sophisticated tactics such as cloaking, where reviewers and automated systems are shown benign creative while targeted users are redirected to fake news websites, mean that a well-founded complaint can appear unsubstantiated to a system that only sees the benign version. We would welcome guidance from Ofcom that a complaint should not be treated as manifestly unfounded solely because the provider's own review of the advertisement did not reproduce the harm reported, particularly where the complaint concerns redirection or landing-page behaviour.

We also urge Ofcom to revisit the scope of FAU/FAS D5.2(c) due to the risks it introduces. We support protecting legitimate advertisers from repeated removal of an advertisement that has already been reviewed and restored, consistent with our concerns about over-removal above. However, an exemption from automated re-flagging should not become a blind spot. Bad actors are known to use dormant accounts and account takeovers to inject malicious payloads or fraudulent redirects into previously approved advertisements. The protection should therefore apply only while the advertisement remains materially unchanged, and providers should remain

able to re-evaluate it where its content, destination or account control changes, or where new evidence of fraud emerges.

Ad libraries (FAU/FAS E1)

Do you agree with our proposal? Please provide any arguments and supporting evidence.
Do you agree with the principles, functionalities and information categories which we have included in the proposed measure? Please explain your reasoning and, if possible, provide supporting evidence.
We recognise that ad libraries, and granular search functionality, can help expert bodies (the FCA, ASA, UK Finance and others) detect and report fraudulent advertising. Our concern is that the proportionality analysis supporting the measure is one-sided, and that the measure is not calibrated to reflect the costs it does not assess.

- The analysis does not consider the commercial-sensitivity cost to advertisers. Ofcom's assessment sets out the fraud-detection benefits and estimates providers' build and maintenance costs (5.100–5.104), but it does not weigh the commercial-sensitivity or competitive-intelligence cost to advertisers of publishing targeting, reach and keyword data. Several mandated fields are commercially sensitive: target audience by demographic (E1.6(e)), target or actual reach (E1.6(f)), and the content keywords used to target users' searches (E1.6(g)). Published daily, to non-registered users, via an open API, these fields could allow a competitor to reconstruct an advertiser's media strategy. This is of particular concern for the keyword-targeting field in the search context, where the keyword is the core of the paid-search product. We ask Ofcom to assess the advertiser-side commercial cost of public disclosure, which its current analysis omits entirely.

- Less-intrusive access models have not been genuinely weighed. Ofcom rejects limiting access to accredited experts on the basis that maintaining an expert list would be time-consuming and costly, and that wider access yields more reports (5.30). That reasoning weighs administrative cost and marginal reporting benefit, but not the commercial harm of fully public disclosure. Given that the detection benefit accrues principally to expert bodies (5.30 concedes "most of the benefits" come from experts), we ask Ofcom to consider authenticated or accredited access to the commercially sensitive fields, with a narrower public tier, rather than a single fully public API. This would preserve the detection benefit while reducing the competitive-intelligence exposure.

- Scope. The library applies to all paid advertising, not a risk-tiered subset. We ask Ofcom to explain why universal coverage of the sensitive fields is proportionate to a concentrated fraud problem, and whether higher-risk categories could be prioritised for the more sensitive fields.

- Personal data. E1.6(b) requires publishing the advertiser's name, defined as the name of the person the account is registered to (footnote 180), which for sole traders and individuals is the public disclosure of personal data. Ofcom offloads this to a one-line instruction to comply with data protection law (5.10).

Taken together, we ask Ofcom to revisit E1.6(e), (f) and (g) in particular. These fields carry technical build and maintenance costs for providers as well as the greatest commercial-sensitivity cost for advertisers, while Ofcom's own analysis indicates that most of the detection benefit accrues to expert bodies (5.30). Our preferred approach is for these fields to be available only through authenticated or accredited access. If Ofcom cannot demonstrate that publishing them to the general public adds materially to fraud detection beyond what such access would deliver, they should not form part of the public library.

More broadly, we note that implementation engages significant personal-data processing (the ad library, account checks, banned-person identification), and that the Codes hand this to a general instruction to comply with data protection law. Given the live tension between several measures and data minimisation, we ask Ofcom and the ICO to publish a joint position on how these obligations are to be reconciled, so that providers and advertisers are not left holding conflicting duties.

- DSA alignment is not a substitute for a UK proportionality assessment. Ofcom relies on alignment with Article 39 of the EU Digital Services Act (5.90, 5.96–5.98), including the one-year retention period. We accept that much of the field set broadly tracks the DSA ad-repository regime, and we withdraw any suggestion that the measure wholesale exceeds it. But alignment with an EU regime does not discharge Ofcom's duty to assess proportionality under the Act and UK law, and at least the "content keywords used to target" field appears more granular than the DSA's "main parameters" disclosure.

Do you have any evidence or insights on the process of building an ad library from scratch or adjusting an ad library to align with the proposed measure, and how long that process would take?

We note Ofcom's own evidence that most in-scope providers already operate ad libraries in the UK or EU, so that for the majority the exercise is one of adaptation rather than building from scratch (5.100). However, the commercial cost this creates for advertisers is not accounted for.

Advertising Association
2 October 2026