Consultation response

Field	Value
Responder	The British and Irish Law Education Technology Association (BILETA)
Responder Type	organisation
Org Type	Other
Other Type	Non-for-profit

Volume 1: Context

3. Online advertising ecosystem

Q1: Is there any information or evidence you hold that could enhance our understanding of the online advertising ecosystem?

We broadly agree with Ofcom's overview of the programmatic supply chain. However, we emphasise that modern ad tech relies heavily on automated, multi-tiered bidding ecosystems involving Demand-Side Platforms (DSPs), Supply-Side Platforms (SSPs), ad exchanges, and resellers, where primary platforms often lack full visibility into the original creative payload at the time of auction.

We recommend that Ofcom further account for:
- Dynamic Swapping & Cloaking: Scammers frequently pass benign creative during compliance review and dynamically alter destination URLs or landing page assets after ad approval.
- Reseller Chains: Multi-layered agency and reseller networks make source tracing complex, requiring cross-industry API standards for advertiser identity sharing.

4. Causes and impacts of fraudulent advertising

Q2: Do you have any comments on Ofcom's assessment of the causes and impacts of fraudulent advertising? Do you think we have missed anything important in our analysis? Please provide evidence to support your answer.

We support Ofcom's analysis of consumer harm and financial loss. However, the analysis could be strengthened by accounting for:
- Brand Abuse & Counterfeit Identity: Fraudulent ads do not only harm end consumers; they severely damage the reputation of legitimate UK businesses whose logos, trademarks, and executive personas are impersonated.
- Emerging GenAI Threats: The rapid proliferation of cheap generative AI tools allows bad actors to create localised, hyper-realistic video/audio deepfakes and landing pages at massive scale with low marginal cost.

Ofcom's assessment fails to capture the technical paradigm shift from passive, user-generated synthetic media to autonomous "agentic deepfakes." Artificial intelligence is no longer merely a tool used by a human "directing mind" (Tesco Supermarkets Ltd v Nattrass, [1971] A.C. 153) to broadcast static misinformation. It has evolved into an autonomous reasoning engine capable of real-time, interactive, and micro-targeted behavioural deception (Romero-Moreno, 2026c).

This transition renders the foundational assumptions of the Online Safety Act (OSA) highly vulnerable. Specifically, Section 231(8) restricts proactive technology mandates to "user-generated content." Agentic deepfakes are machine-generated and operate via autonomous execution loops. By anchoring enforcement

to the legal fiction that a human scammer is manually uploading a fraudulent asset, the current framework inadvertently creates a statutory safe harbour for autonomous, machine-driven fraud (Romero-Moreno, 2026c).

Furthermore, Ofcom's analysis omits the impending "12-month regulatory vacuum." A critical compliance gap exists between the December 2026 functional bans on deceptive AI systems (e.g., the EU Digital Omnibus) and the delayed December 2027 enforcement timeline for high-risk AI systems (Romero-Moreno, 2026b). During this window, platforms must transition from acting as passive conduits to active data controllers, recognizing that algorithmic curation of these synthetic campaigns pierces traditional secondary liability shields (Romero-Moreno, 2026a; Romero-Moreno, 2026c).

Q3: Do you agree with the examples of risky characteristics we have provided? Please provide any arguments and supporting evidence.

Yes, we agree with the identified risky characteristics. Characteristics such as newly established accounts launching high-budget campaigns, rapid domain redirection, payment methods from disparate geographic regions, and unverified landing pages are key indicators of fraudulent activity.

We suggest adding:
- Accounts attempting to run high-budget campaigns in regulated financial services categories immediately after creation without historical spend records.
The provided examples reflect legacy fraud typologies that autonomous AI easily evades. Ofcom must update these characteristics to address the following:
- Virtual Camera and Digital Injection Signatures: The threat has migrated from physical presentation attacks to Digital Injection Attacks (DIAs), which inject synthetic streams directly into application feeds.
- Absence of Interoperable Provenance Metadata: Advertising assets lacking cryptographic manifests (e.g., C2PA Content Credentials) or imperceptible watermarks (e.g., SynthID) must trigger high-risk triage.
- Evasion of Hash-Matching: Section 231(2) of the OSA envisions content identification via hash-matching. This is architecturally obsolete against live, interactive agentic deepfakes that iterate audio-visual output in real-time. The absence of a static file footprint must necessitate behavioural triage (Romero-Moreno, 2026c).

5. Our approach to developing Codes measures

Q4: Do you agree with our proposed approach? Please provide any arguments and supporting evidence.

Yes, we agree with the proposed approach in principle. A risk-based, outcome-focused regulatory framework under the Online Safety Act 2023 strikes the right balance between proactive intervention and service functionality. We support clear distinctions between user-generated content (UGC) and paid-for sponsored placements.

BILETA generally agrees with Ofcom's proposed approach, especially the use of a layered, iterative model that acknowledges how quickly and adversarially fraudulent advertising changes. A code that tries to specify every technical detail could quickly become outdated as both fraud tactics and advertising tech evolve. Therefore, Ofcom's approach of defining desired outcomes while allowing implementation flexibility is suitable. Ofcom also recognises the importance of balancing the need for enough detail to ensure legal clarity with the need for flexibility and future-proofing. However, BILETA emphasises that this flexibility must remain meaningful in the final codes. Compliance should mainly be judged based on the effectiveness and proportionality of providers' systems, not on specific technological architectures. This is especially vital as new authentication, fraud detection, and ad delivery methods develop. There is also a concern that strict compliance requirements might favour larger providers with extensive infrastructure, creating barriers for new or innovative services. Ofcom should continue to consider whether measures that are affordable for established platforms could hinder market entry for others. This aligns with the statutory focus on proportionality and Ofcom's obligations related to competition, investment, and innovation. Finally, BILETA suggests periodic reviews of how these measures interact, since layered regulations can lead to cumulative

or unintended effects, even if each measure seems proportionate alone. Particular attention should be given to how fraud indicators, repeated account checks, automated moderation, and account bans work together, to prevent errors in fraud detection from affecting multiple stages of compliance.

To reconcile this need for flexibility with the necessity of robust security, BILETA asserts this is resolved by the tiering of the Online Safety Act. The desire for flexibility and outcome-based codes is appropriate for smaller, Category 2B or uncategorized services to protect innovative startups from strict compliance burdens. However, for Category 1 and Category 2A platforms (like Meta, Google, X, and ChatGPT Search) that have the resources to implement them and are the primary targets of industrialized Fraud-as-a-Service syndicates, BILETA advocates for a highly prescriptive approach. For these high-tier services, the codes must explicitly demand specific high-assurance standards, such as NIST IAL2, zero-retention biometrics, and ISO 20022.

6. Combined impact assessment

Q5: Do you consider the overall burden on the service providers in scope of our proposed measures to be proportionate? Please provide any arguments and supporting evidence.

Yes, we consider the overall burden proportionate given the scale of consumer harm caused by online fraud in the UK.

However, proportionality depends heavily on implementation flexibility:
- Technical Integration Lead Time: Implementing robust, real-time identity verification (IDV) and anti-evasion fingerprinting across millions of advertisers requires significant engineering time.
- Intermediary Flexibility: For services relying on third-party ad networks where direct control is limited, Ofcom's "all reasonable endeavours" safe-harbour approach is critical to maintaining a fair regulatory burden.

Volume 2: Risk, governance and control

2. Advertising intermediaries

Q6: Do you agree with our proposal? Please provide any arguments and supporting evidence.

Yes, we agree. Service providers should not be held strictly liable for third-party network ads if they do not have direct operational control, provided they execute robust contractual, technical, and oversight controls over those intermediaries.

Q7: 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.

The illustrative examples provide useful clarity. We support defining "all reasonable endeavours" through contractual mandates (e.g., requiring intermediaries to perform equivalent IDV/KYC checks) and technical integration (e.g., automated API reporting and continuous sampling).

Additional Examples of Intermediary Safety Measures:
- Contractual Kill-Switches: Mandatory contract clauses enabling platforms to immediately throttle or terminate intermediary ad feeds if the fraud rate exceeds an agreed threshold.
- Standardised Identity Signals: Requiring ad networks to pass cryptographically signed advertiser identity headers in OpenRTB protocol bid requests.

3. Fraud indicator assessment

Q8: Do you agree with our proposal? Please provide any arguments and supporting evidence.

Yes, we agree. Establishing a formal, continuous Fraud Indicator Assessment ensures platforms stay ahead of rapidly evolving fraud tactics rather than relying on static, outdated threat models.

BILETA agrees in principle but emphasizes that static assessments cannot match the iteration speed of agentic AI. Compliance requires a macro-level architectural mandate enforcing "institutionalised distrust" (Laux, 2024; Romero-Moreno, 2026c).

Assessments must evaluate a platform's resilience against multi-agent networks that autonomously negotiate, compile, and broadcast content. Platforms must prove they deploy continuous Human Rights Impact Assessments (HRIAs) and system-level sandboxing that restrict the autonomous scope of ad-generation agents.

Q9: 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?

BILETA agrees with Ofcom's proposal that providers should use multiple sources when assessing and updating fraud indicators. Ofcom already considers relevant data such as user and advertiser information, complaints, reports, intermediaries, trusted flaggers, law enforcement referrals, internal experts, incident reviews, product testing, and research. For identifying new indicators, Ofcom also includes law enforcement, trusted flaggers, expert organizations, data-sharing agreements, and providers' own trust-and-safety systems.

We recommend that the final guidance explicitly encourage providers, where lawful and appropriate, to consider additional sources:
- Data from payment providers, acquiring banks, and other payment system players on patterns linked to fraudulent advertisers;
- Domain registration, hosting, and landing page data, especially rapid changes in domains or landing pages;
- Reports from brands, public figures, and organisations whose identities are being impersonated;
- Input from consumer groups, regulators, and sector-specific bodies, including those outside the online safety regime;
- Results of successful account and ad appeals, which can offer valuable insights into false-positive fraud indicators;
- Cross-platform data on coordinated fraudulent campaigns, when shared with proper legal, technical, and governance safeguards.
- FCA Warning Lists & Regulatory Alerts: Direct feeds from the Financial Conduct Authority (FCA) regarding unauthorised firms and investment scams.
- Cross-Industry Intelligence Sharing: Feeds from specialised fraud data exchanges (e.g., CASS, Stop Scams UK) and domain registrar abuse reports.
- Ad Complaints & Chargeback Data: Integration of payment gateway chargeback rates and consumer dispute flags associated with specific advertiser URLs.
- FATF Horizon Scan Protocols: Identifying synthetic identity creation as a systemic AML/CFT threat.
- Empirical "Lab-to-Live" Accuracy Variance Data: Accounting for the 10%-20% drop in deepfake detection accuracy when moving from laboratory benchmarks to live environments (DSIT, 2026).
- Cross-Sectoral Regulatory Intelligence: Utilizing real-time threat feeds shared across the Digital Regulation Cooperation Forum (DRCF).

The final point is especially crucial. Fraudulent campaigns often span multiple services, but a single platform only witnesses part of this activity. On the other hand, information sharing must not turn into a concealed cross-platform blacklist that individuals or businesses find hard to escape. Therefore, Ofcom should clarify

that using fraud indicators remains subject to relevant data protection rules. The ICO highlights core UK GDPR principles such as lawfulness, fairness, transparency, purpose limitation, data minimisation, and accuracy. Providers should regularly assess both how effective these indicators are and their potential to cause discrimination or exclusion when used alone or in combination. Notably, an indicator should generally serve as evidence prompting further investigation rather than definitive proof of fraud. This is especially important when indicators relate to behaviour that could have legitimate explanations.

4. Governance and accountability

Q10: Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

Yes, we agree with the proposed governance and accountability measures.
- Senior Executive Oversight: We support designating clear senior management accountability for compliance. This ensures that ad safety is treated as a core operational risk alongside cybersecurity and financial compliance.
- Auditability & Record-Keeping: Maintaining clear internal logs of ad submissions, approvals, rejections, user complaints, and bans is essential for demonstrating regulatory compliance to Ofcom.

5. Testing advertisement generation tools

Q11: Do you agree with our proposal? Please provide any arguments and supporting evidence.

Yes, we agree with the proposal. Where platforms provide native AI/ML ad generation or copy-creation tools to advertisers, those tools must undergo red-teaming and safety testing prior to release to ensure they cannot easily be prompted to generate fraudulent copy, deceptive claims, or misleading visual assets.

Testing must target agentic autonomous execution rather than static output generation. Agentic architectures exhibit "intent agentic misalignment," where a model pursues a benign assigned objective through autonomously selected harmful actions (Romero-Moreno, 2026c). Testing must conduct adversarial red-teaming against agentic jailbreaks and automated "Frankenstein" synthetic identity creation. Providers must implement safety-by-design algorithmic circuit breakers that instantly halt campaigns when repeated synthetic biometric features are detected.

Volume 3: Ensuring account integrity

2. Account checks and actions

Q12: 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. These are :

Yes, we fully support mandatory pre-flight advertiser verification. Verifying account integrity before an ad goes live is far more effective than remediating consumer harm after fraudulent ads have already been served.

BILETA supports the framework, provided it incorporates the "verify-then-publish" standard across all tiers:

Q12.A: verification that advertising account holders work for, or on behalf of, the individual or organisation they are advertising

We agree. Verification should include corporate domain email matching, official authorisation letters for third-party agencies, and cross-checks against official registry records (e.g., Companies House). Must mandate NIST SP 800-63-4 Identity Assurance Level 2 (IAL2) standards using zero-retention biometrics. Raw biometric data must be deleted immediately following an ephemeral 1:1 match, generating a

cryptographic biometric integrity hash via Zero-Knowledge Proofs (ZKP) to prove human origin without centralizing templates (Romero-Moreno, 2026b).

Q12.B: checks and actions to prevent banned advertising account holders returning

We strongly support this. Scammers frequently attempt to set up new accounts immediately after being banned. Systems must utilise device fingerprinting, shared payment method blacklists, domain history tracing, and IP subnet correlation. Must transition to Secure Multi-Party Computation (sMPC) to identify returning synthetic identities or composite personas without holding raw biometric honeypots.

Q12.C: checks for accounts with indicators that there is a material risk they will post fraudulent advertising so that restrictions can be applied to them

We agree. Dynamic risk scoring should automatically apply temporary spending limits or require manual review for new accounts, accounts changing payment methods, or accounts suddenly launching campaigns outside their historical categories. Section 231(6) OSA presumes human scammers leave repetitive operational signatures (like bot farms). Agentic AI is engineered to mimic legitimate, non-linear workflows and adapt to environmental feedback, evading these checks (Romero-Moreno, 2026c). Platforms must deploy advanced behavioural biometrics to detect non-human interaction speeds and virtual camera feeds.

Q12.D: repeat checks where there are relevant changes identified on an account.

We agree. Re-verification triggers should be mandated whenever an account changes beneficial ownership, billing details, key domain linkages, or administrative user access.

BILETA supports the idea of event-triggered repeat checks. Ofcom suggests conducting repeat checks when a provider loses confidence in an earlier verification or detects a relevant change, such as potential account takeover, modifications in posted advertisements, changes in the advertised entity, renewed activity after dormancy, new links between accounts, or an account-name change. This approach is better than frequent, automatic reverification of all accounts. A risk-based trigger minimises disruption for legitimate advertisers and ensures resources focus on situations where verification may no longer be reliable. However, BILETA recommends that the final Code differentiate between a simple change and a material risk-related change. Routine business events like staff turnover, agency shifts, restructuring, or rebranding should not automatically lead to significant advertising restrictions unless supported by sufficient evidence. If a repeat check results in restrictions, providers should clearly explain the issue so advertisers can address innocent discrepancies or use the appeal process. The level of scrutiny and any restrictions should be proportional to the identified risk. Must be triggered by automated anomalies such as sudden shifts in campaign spend velocity, employing cryptographic human-authorization control planes that require an ephemeral biometric sign-off to deploy high-risk campaigns.

Q13: 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.

We suggest:
- 2FA Mandatory Enrolment: Requiring multi-factor authentication (2FA) for all advertising management accounts to prevent account takeovers (ATO).
- Cooling-Off Periods: Implementing a mandatory 24-48 hour cooling-off period with lower spending caps for newly verified advertiser accounts before high-volume impressions can be served.

BILETA suggests that additional controls could involve risk-based checks triggered by significant changes to payment methods, account administrators, security credentials, advertising spend, target audiences, or destination domains. Providers might also implement "step-up" verification when multiple weak indicators occur together-for instance, an inactive account becoming active, changing its administrator and payment method, significantly increasing advertising expenses, and starting to promote a different product category-justifying a stronger check even if no single event warrants severe restrictions. This method is

preferable to viewing individual characteristics as binary fraud indicators. It allows for graduated responses such as temporary spending caps, reduced reach, or pre-publication review instead of immediate suspension. Ofcom already considers various restrictions, including limits on spending, impressions, targeting, format, and pre-publication review. When providers use behavioural or personal data for these purposes, the necessity and proportionality of such processing should align with data-protection laws. Extra caution is essential when account-level decisions rely mainly or entirely on automated profiling, as ICO guidance highlights specific protections for decisions with legal or similarly significant effects.

Platforms must implement Federated Learning (FL) to coordinate defence across distributed nodes, allowing real-time identification of cross-platform ad scams while adhering strictly to UK GDPR Data Protection by Design principles (Romero-Moreno, 2026a).

3. Preventing fraudulent financial services advertising

Q14: Do you agree with our proposal? Please provide any arguments and supporting evidence.

Yes, we strongly agree. Financial services fraud causes severe harm to consumers. Providers must require advertisers promoting financial products/services in the UK to demonstrate current registration on the FCA Financial Services Register or confirm they fall under a valid statutory exemption.

Domestic registry checks fail against transnational FaaS syndicates that exploit jurisdictional black holes. Ofcom should coordinate a Transatlantic Regulatory and Financial Interoperability Framework. This embeds biometric integrity protocols directly into global financial plumbing by requiring cryptographic biometric integrity hashes within the SplmtryData schema of ISO 20022 financial messaging standards. If a high-value advertisement transaction originates from an unverified synthetic identity, payment rails automatically execute a protocol-level blockade (Romero-Moreno, 2026b).

4. Countering account takeover

Q15: Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

Yes, we agree. ATO is a primary vector that fraudsters use to hijack high-reputation, aged advertiser accounts and bypass onboarding verification. We support mandatory 2FA, session-anomaly detection (e.g., concurrent logins from distant geographic regions), and automated campaign pauses upon detection of suspicious administrative profile edits.

BILETA supports the proposed measures FAU/FAS H3 and H4: requiring an account security mechanism for advertising accounts and an accessible process for account holders to report account takeovers. The issue of account takeover is particularly important because it enables fraudsters to inherit the perceived legitimacy, history, and potentially the audience of an established advertiser. Ofcom rightly recognises that both prevention and swift recovery are vital components of the response.

However, we suggest that Ofcom clarify what constitutes an effective "account security mechanism." Security should not be limited to having a second authentication factor. The NCSC emphasises the ongoing relevance of credential phishing and recommends stronger resistance through FIDO2-based credentials or passkeys, which better defend against common phishing attacks. For higher-risk advertising accounts, Ofcom should advocate for phishing-resistant authentication methods whenever feasible, along with step-up authentication for sensitive actions such as adding administrators, changing recovery details, updating payment information, or making large changes in advertising spend.

Simultaneously, account recovery deserves special focus. A security measure that prevents account takeover but makes recovery overly difficult could harm legitimate businesses. Reporting channels should be easily accessible even if the advertiser cannot log into the compromised account. Ofcom's proposal that reporting mechanisms be simple to find, access, and use is a positive step in this regard.

Q16: Do you have any evidence about account takeover that could inform our assessment, including how frequently it occurs or its impact?

Industry trends indicate that as onboarding identity checks become stricter, bad actors increasingly turn to credential stuffing, phishing, and malware/cookie-theft to breach legitimate existing corporate ad accounts. Pausing spend automatically when credentials or primary recovery details change drastically reduces financial exposure for compromised businesses.

BILETA lacks independent quantitative data on the frequency of advertising-account takeovers. However, we concur with Ofcom that account takeover is a separate risk, as a compromised legitimate account can bypass controls mainly intended to monitor newly created or suspicious advertisers. Ofcom's evidence shows that reporting account takeover can be difficult and lacks clear guidance, highlighting the importance of having a dedicated process.

5. Advertising bans

Q17: Do you agree with our proposal? Please provide any arguments and supporting evidence.

Yes, we agree. Account bans should be enforced swiftly upon confirmation of fraudulent activity, with automated blacklisting of associated payment instruments and technical identifiers.

BILETA agrees that an advertising account holder responsible for fraudulent advertising should generally be barred from advertising to UK users, and that providers should take reasonable measures to prevent easy evasion through alternative accounts. This follows logically from the conclusion of fraudulent ads and avoids a scenario where repeat offences are seen as just part of doing business. Ofcom suggests implementing such bans under FAU/FAS H5.

However, because an advertising ban is significant, procedural safeguards become especially important. Providers should, as far as reasonably possible, differentiate between the account, the individual user who carried out the act, and the organisation behind the ad account. In complex agency or corporate setups, responsibility should not automatically be assigned to all entities linked to an account. Furthermore, providers should ensure that systems meant to prevent banned advertisers from returning do not create permanent, opaque blacklists of identities. Otherwise, false positives might transfer between related accounts or impact innocent employees, agencies, or successor organisations.

Q18: Are you aware of any additional exceptional circumstances in which it may be proportionate for providers to reverse advertising bans?

Reversals should strictly be limited to:
- Proven Account Compromise (ATO): Where an advertiser proves their account was hijacked without negligence, the unauthorised ad campaigns were purged, and full security control (2FA reset) has been restored.
- Demonstrated False Positive / Third-Party Error: Where automated moderation mistakenly flagged legitimate landing pages or domain mapping due to technical glitches.

Beyond account takeover, BILETA proposes that reversal should be allowed in specific cases, such as when: an appeal proves the advertisement was wrongly classified as fraudulent; the ban was applied to the incorrect individual or entity within an account hierarchy; an employee, contractor, or agency acted without authority and effective remedial actions have been taken; there has been a genuine change in ownership or control of a business, so the individuals responsible for the fraud no longer hold influence; or when new evidence significantly changes the facts or legal grounds for the initial ban. These situations should be rare, and restoring access might be accompanied by increased monitoring or temporary restrictions.

6. Account appeals

Q19: Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

Yes, we agree. To protect legitimate businesses from erroneous enforcement, providers must maintain a transparent, timely, and fair appeals process featuring human review for contested account suspensions.

BILETA strongly supports FAU/FAS J1-J4. Ofcom suggests that account appeal procedures be accessible, easy to use, timely, and followed by corrective actions that aim to restore the situation as if the erroneous decision hadn't been made. Appeals are vital because the proposed system depends on risk indicators, automated tools, and preventative restrictions, not just proven cases of fraud. Some false positives are unavoidable, making an effective appeal process crucial for the system's accuracy and legitimacy.

BILETA proposes three additions: First, advertisers should usually receive enough information about why they received an adverse decision to make meaningful appeals, without revealing detailed anti-fraud measures that could enable circumvention. Second, if a decision was mostly automated, there should be human review during appeals instead of just rerunning the same automated process. Third, "promptly" should be understood in context; a restriction blocking a small business from a time-sensitive campaign may cause significant harm even if reversed later. Providers should rank appeals by the seriousness and urgency of the restriction as well as the fraud risk.

Volume 4: Moderation

2. Advertising moderation

Q20: Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

Yes, we agree with the proposals for multi-layered ad moderation.
- Pre-Display Scanning: Automated pre-display inspection of text, media, and target destination URLs is necessary prior to serving impressions.
- Anti-Cloaking Re-crawls: Continuous, random post-approval re-crawling of ad landing pages using unannounced, residential proxy IP pools is critical to defeating dynamic redirection scams.

3. Terms of service and publicly available statements

Q21: Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

Yes, we agree. Clear Terms of Service that prohibit fraudulent practices, define prohibited ad categories, and detail enforcement procedures ensure transparency and legal clarity for all business users.

BILETA endorses the proposed measures FAU/FAS G1 and G2. The Act mandates that providers describe the proactive technology they use to meet their obligations against fraudulent advertising. This description should include details of the technology type, usage scenarios, and operational methods. Ofcom recommends that this information be presented in a clear and accessible manner. We believe that "clear and accessible" should mean more than simply including relevant details in lengthy terms of service. Instead, information about fraud controls should be formatted so that ordinary advertisers or users can easily find and understand it. However, transparency should not come at the expense of aiding fraudsters, so confidentiality around technical specifics is essential. A distinction can be made between systemic transparency-an overview of technology and signals-and operational disclosure-specific details like thresholds and model parameters. Providers should also update this information promptly if the technology's role changes materially. Otherwise, the information risks becoming outdated and less meaningful.

4. Advertising complaints (including DRC for trusted flaggers)

Q22: Do you agree with our proposals? Please confirm which proposed measure your views relate to and provide any arguments and supporting evidence.

Yes, we agree. User-friendly, accessible "Report Ad" mechanisms integrated directly into the ad frame/interface are vital. Fast-track review workflows must be maintained for reports coming from trusted flaggers, law enforcement, and regulatory bodies.

BILETA supports Ofcom's proposed approach to handling advertising complaints, emphasising accessibility and the ability to include relevant contextual details. Ofcom's research shows notable differences in the friction experienced with current complaint systems. Additionally, 43% of respondents who encountered potentially fraudulent ads reported them, while 45% of those who took no action believed reporting would be ineffective. Complaint processes should be easy to find directly from the advertisement and enable users to submit URLs, screenshots, and explanatory information when needed. We also agree that providers should be allowed to dismiss clearly unfounded complaints. Without this, coordinated or automated malicious reporting might overwhelm systems and disrupt legitimate advertising. Ofcom's suggested high threshold, documented policy, regular reviews, and record-keeping are appropriate.

BILETA advises that providers should monitor if certain types of complainants or complaints are disproportionately categorised as unfounded. The anti-abuse exception shouldn't be used to silently dismiss reports from users who provide limited info, communicate poorly, or repeatedly report genuine recurring scams. Whenever possible, complainants should be given basic feedback on the outcome of their complaint. This could boost confidence in the reporting system without revealing confidential enforcement details.

5. Ad libraries

Q23: Do you agree with our proposal? Please provide any arguments and supporting evidence.

Yes, we agree. Searchable, publicly accessible Ad Libraries provide valuable transparency, enabling researchers, journalists, and regulators to identify systemic fraud campaigns across platforms.

BILETA strongly endorses the idea of publicly accessible advertising libraries. Since paid ads are targeted and personalised, they are not easily scrutinised like traditional public ads. An archive plays a crucial role in accountability for regulators, researchers, civil society, journalists, and consumers. Ofcom also states that without ad libraries, experts may only see ads aimed at specific users, making systematic analysis difficult. Additionally, there is beneficial regulatory alignment. Ofcom notes that its proposal closely resembles the requirements of Article 39 of the EU Digital Services Act. The European Commission's enforcement actions related to ad repositories highlight the importance of searchable, usable archives over mere transparency.

Q24: 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.

Yes, we agree. Including verified advertiser names, targeted demographics, historic impression ranges, and creative assets within the Ad Library strikes the right balance between public accountability and business privacy.
BILETA generally agrees with Ofcom's suggested principles of public availability, accuracy, and reliability, along with its proposed functionalities. Ofcom recommends that advertisements accessible to UK users stay available during their campaign and for twelve months afterwards, that libraries be updated at least daily, and that they support keyword, exact-phrase, and advertiser-account searches, along with direct reporting and an API. The API requirement is especially valuable. Public transparency through only a graphical interface can be limited in usefulness for researchers and organisations that need to identify patterns across numerous advertisements.

BILETA offers four suggestions:
1. Identifiers should be stable enough to reference an advertisement, advertiser, and campaign across

complaints, investigations, and research.
2. If an advertisement changes significantly during a campaign, the library should keep relevant previous versions instead of only showing the latest.
3. The library should, as far as possible, identify the entity behind an advertisement and the advertising account, which is critical when agencies or intermediaries are involved.
4. Ofcom should think about promoting interoperability by using standardised export formats or API fields. Shared data structures would greatly aid cross-platform analysis and could lower compliance costs for organisations already bound by the DSA.
Privacy protections remain crucial. Public ad libraries should reveal information pertinent to advertising accountability without unnecessarily exposing personal details of minor employees, account operators, or individuals. The data-minimisation principle stays relevant even when transparency serves a significant public interest.

Q25: 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?

Building or updating an Ad Library infrastructure to handle high-frequency search indexing, creative asset retention, and public API querying typically requires 6 to 12 months of technical design, schema updates, and load testing. Ofcom should ensure an adequate implementation timeline following the publication of the final Statement.

BILETA lacks independent evidence to reliably estimate the time required to build an advertising library from scratch. However, we concur with Ofcom that costs and timelines will vary significantly between providers that already manage advanced advertising repositories and those that need to establish new public-facing systems. Ofcom notes that nearly all relevant Category 1 and 2A services currently operate some form of ad library. For example, implementing search functionality might involve a team of eight to twelve software engineers along with product and design inputs over two to four weeks, though meeting the full Code requirements could take longer. For providers already subject to Article 39 DSA obligations, Ofcom should promote the reuse of existing infrastructure rather than mandating a separate UK system, as long as UK ads, information, and functionality requirements are fully met. The focus should be on functional compliance, not specific architecture. The goal is a complete, timely, searchable, and reliable repository, not a particular database design. A phased approach could be suitable if significant modifications are necessary, but essential features-especially public access, inclusion of UK ads, and effective search-should not be postponed just because more advanced features are still under development.

Additional questions

Q26: Do you agree with our Equality Impact Assessment? Please provide any arguments and supporting evidence.

Yes, we agree with the EIA. Fraudulent ads disproportionately impact vulnerable demographics, including elderly users and consumers with lower digital literacy. Implementing strict ad controls directly advances equality and consumer protection objectives.

Q27: Do you agree with our Welsh language assessment? Please provide any arguments and supporting evidence.

Yes, we agree. Ensuring that public reporting interfaces, safety notices, and guidance are accessible in Welsh treats the Welsh language on an equal footing, in full compliance with Ofcom's statutory Welsh Language Standards.

Q28: Do you agree with our proposals? Please provide any arguments and supporting evidence.

Yes, we agree. Updating the Illegal Content Judgements Guidance to provide practical, real-world fraud typologies ensures consistency across platform moderation teams and reduces ambiguity in legal assessments.

BILETA recommends two essential additions to the guidance:
- Alignment with Corporate Criminal Liability: The guidance must explicitly integrate with the "failure to prevent fraud" offence under Section 199 of the Economic Crime and Corporate Transparency Act 2023 (ECCTA). Deploying ad networks without state-of-the-art Injection Attack Detection (IAD) and zero-retention verification invalidates any "reasonable fraud prevention procedures" defence.
- Re-engineering Human Oversight ("Institutionalised Distrust"): Automated triage runs the risk of automation bias and discriminatory error rates. Applying Laux's doctrine of "institutionalised distrust" (Laux, 2024), detection architectures must utilize models (e.g., Gemini 3 Flash) to translate sub-perceptual forensic markers into plain-language risk summaries... across the AI value chain (Romero-Moreno, 2026c).