[redacted]

Volume 1: Context

Online advertising ecosystem

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

  [*REDACTED*]

Causes and impacts of fraudulent advertising

- **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 broadly agree with Ofcom's assessment of the causes and impacts of fraudulent advertising. Data from the most recent publication of the Crime Survey for England and Wales, published after the consultation, further demonstrates the increasing scale of the threat; in the year ending March 2026 fraud accounted for 47% of headline crime measured by the survey. The estimated number of fraud victims also increased 10% from 3.4 million to 3.8 million.

  We consider that in your review of the causes and impacts of fraudulent advertising, greater emphasis should be placed on framing fraudulent advertising as part of a networked ecosystem, which increases complexity for its identification and detection. In Ofcom's assessment, evidence should be included to demonstrate how fraudulent campaigns operate at scale across interconnected services, infrastructure and actors, creating challenges for identification, attribution and disruption. This could include examples where activity spans multiple accounts and services, involves affiliates or other intermediaries, or relies on compromised or dormant accounts to sustain activity and evade detection. This would help illustrate how risk arises from the wider ecosystem in which fraudulent advertising operates, rather than solely from the characteristics of individual adverts or actors.

  We also encourage Ofcom to recognise that fraudulent advertising journeys do not always follow a route from advert to website. They can [*REDACTED*] redirect users across multiple URLs. Increasingly adverts have been observed which direct users to off-service communication channels, including messaging services, where the fraud develops through subsequent contact. Additionally, fraudulent adverts may serve to facilitate data harvesting. The impact of fraudulent advertising should therefore be considered across the full spectrum of fraud journeys, including where no conventional landing page is used and where financial loss or other harm occurs at a later date.

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

  We agree that the risky characteristics identified by Ofcom provide an appropriate starting point. However, fraudulent advertising methods evolve rapidly, and adverts may display few obvious indicators when considered in isolation. The characteristics set out in the consultation should therefore be treated as illustrative rather than exhaustive by providers. We also consider

that the absence of an obvious or directly observable characteristic should not in itself indicate that an advert or account presents a low risk of fraud.

The following additional characteristics may constitute risk indicators which, particularly when identified in combination, warrant further investigation by a provider:

- Adverts that direct users immediately to off-platform communication channels rather than to a website.
- Repeated or adapted creative material, branding or messaging across disparate accounts or adverts.
- Account compromise, changes in account administration, or renewed activity following a period of dormancy.
- Short-lived adverts, regular changes between active and in-active status, or edits to adverts or accounts following verification or moderation.
- Inconsistencies between an advertiser's stated identity or business purpose and the content being promoted (and any subsequent click-through landing page).
- Financial promotions that are not identified as such.
- Synthetic credibility signals, such as fabricated news content, review sites or testimonials
- Use of synthetic personas or AI-generated identities designed to circumvent platform controls, impersonate trusted individuals or adapt to mitigation measures deployed by providers
- Techniques used to create or manage multiple separate accounts (including the use of anti-detect browsers), conceal connections between accounts or facilitate re-entry following enforcement action.
- Common infrastructure or account information including payment information across seemingly disparate accounts and advertisers.

Approach to Codes

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

We support Ofcom's layered approach to tackling fraudulent advertisements, in particular its focus on upstream controls. This aligns with our priority to prevent fraudulent advertising from entering the ecosystem rather than relying on reactive reporting after user exposure and subsequent harm has already occurred.

As the regime matures, an important consideration will be how Ofcom assesses not only whether measures are implemented, but whether they are contributing to reductions in exposure to fraudulent advertising and associated harm.

Beyond our comments on the specific approach taken to the design of the codes of practice described in this measure, our responses to individual measures identify a number of themes that we consider should be reflected across the code's framework:

1. Fraudulent advertising is a networked criminal activity and cannot be assessed solely through advert content.
2. Providers should be judged on the effectiveness of interventions, not just whether processes exist.
3. Trusted flagger reporting, intelligence sharing and transparent ad libraries are critical enablers of effective fraud disruption. However, these measures should complement,

rather than replace, providers' own systems and controls for proactively identifying and investigating fraudulent advertising activity. Intelligence sharing and engagement with law enforcement should be a two-way process, with providers both receiving and contributing intelligence to support wider disruption efforts.

Fraudulent advertising should be viewed as part of a wider criminal ecosystem rather than as isolated pieces of content. Effective regulation should therefore focus not only on individual adverts, but also on advertiser behaviour, account networks, associated infrastructure, click-through destinations and the organised criminal actors responsible for harm. We consider that moderation and account safety measures should more explicitly require services to take account of the full range of information available when identifying potentially fraudulent or harmful activity. Services should also be encouraged to identify connected accounts, coordinated behaviours and wider patterns of behaviour that may indicate organised criminal networks. Taking this broader approach should improve the identification and disruption of the underlying actors and networks driving harm, rather than focusing solely on the removal of individual accounts or adverts.

We also consider that the codes should place greater emphasis on measurable outcomes as demonstrations of compliance, such as how services monitor, evaluate and evidence the effectiveness of the interventions they deploy, rather than relying solely on process-based compliance requirements. Ofcom could assess providers against metrics, using minimum service standards as appropriate, such as time taken to remove confirmed fraudulent adverts, time taken to respond to trusted flagger reports, re-registration rates of banned accounts and overall exposure to fraudulent advertising.

In addition, effective fraud disruption requires the integration of intelligence from trusted flagger reporting, intelligence sharing and ad libraries. However, providers should not rely solely on intelligence received through these mechanisms and should also have effective systems and controls in place to proactively identify potentially fraudulent advertising, detect coordinated campaigns and generate actionable intelligence. We see value in Ofcom encouraging or facilitating structured intelligence-sharing arrangements between categorised services, trusted flaggers, financial institutions, cybersecurity organisations and relevant public sector bodies to enable identification of cross-platform fraudulent advertising campaigns and repeat offenders. Such arrangements should support the sharing of intelligence between providers, as well as the sharing of intelligence with law enforcement, to facilitate more effective intervention and wider disruption of fraudulent activity.

Combined impact assessment

- **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.**

We consider the overall burden on in-scope providers to be proportionate, given the scale and severity of harm associated with fraudulent advertising. Many of the proposed measures will build on providers' existing governance, account integrity and trust and safety systems, reducing the need to develop entirely new processes.

Given the complexities of identifying fraud (particularly given the contextual and intelligence assessment often needed), we do note that effective detection may require specialist expertise that is not captured by existing systems designed to address other harms. Some providers may need to invest in dedicated fraud capability, training and processes. However, we consider this

proportionate to the scale of the threat and the protections the measures intend to deliver for the public.

Volume 2: Risk, Governance and Control

Advertising intermediaries

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

  We support the proposal that where complex intermediary arrangements prevent a provider from implementing a measure directly, they should use all reasonable endeavours to implement a version as similar to the relevant measure as possible. We also welcome the requirement for providers to document where measures could not be implemented directly. This may help ensure providers can evidence why measures could not be implemented directly, the alternative approaches considered, and any resulting gaps in controls, creating a clearer basis for accountability and scrutiny.

  However, we encourage Ofcom to assess whether the different approaches adopted by providers result in sufficiently consistent and effective protection for users. As such Ofcom may consider requiring this information to be addressed explicitly within the provider's annual assessment. This is particularly important as Open-display advertising ecosystems may present particular challenges for assurance and enforcement where visibility, moderation responsibilities and advertiser relationships are distributed across multiple intermediaries.

- **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.**

  We consider the illustrative examples helpful in demonstrating how providers could work with, or place requirements on, intermediaries where they cannot implement a measure directly.

  To illustrate how this approach could apply to other measures, we encourage Ofcom to include an additional example covering public ad libraries, given their importance in supporting reactive moderation and trusted flagger reporting. Outside a walled garden environment, adverts and the information required for an ad library may be distributed across several intermediaries. An example setting out how providers could deliver an accurate and functional ad library in these circumstances would provide useful clarity.

Fraud indicator assessment

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

  We are pleased to see the inclusion of a Fraud Indicator Assessment, which is one of the most strategically significant measures in the Codes. It aligns with our experience that fraudulent

advertising is complex, adaptive and not easily identifiable within the content of an advert alone. The Fraud Indicator Assessment proposal provides a valuable mechanism for providers to develop an ongoing understanding of how fraud manifests on their services and apply that across account and advert level controls.

We highlight that the value of the assessment will depend on how effectively its findings are translated into improved safety measures such as account controls and moderation. Fraudulent adverts may appear legitimate when viewed in isolation, with risk becoming apparent only through consideration of connected accounts, advertiser behaviour, click through destinations or patterns within wider infrastructure. Providers may require sufficient fraud expertise and analytical capability to identify these relationships and act on them.

The effectiveness of the Fraud Indicator Assessment may also be influenced by the threshold by which providers consider that a characteristic or combination of characteristics presents a risk. Overly cautious or differing thresholds could result in inconsistent levels of protection between services, even where similar evidence is available. We would welcome clarification on how Ofcom will assess whether the measure is consistently interpreted and how effective Fraud Indicator Assessments are in practice.

- **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?**

We agree with the range of information sources proposed by Ofcom. In particular, cross-sector intelligence products developed through data sharing arrangements between online platforms, financial institutions, cybersecurity organisations, relevant public-sector bodies and law enforcement could assist providers in identifying connections between adverts, accounts, infrastructure and wider fraud activity that may not be apparent from their service-level information alone.

Providers should also consider relevant malicious domain and URL intelligence, including appropriately maintained block lists, when assessing advert destinations and connections between accounts. This information should inform the assessment, recognising that newly created domains, compromised legitimate websites, redirect and cloaking may limit its coverage.

Cross-sector data sharing initiatives, including those being developed through the Online Crime Centre, may provide additional context on emerging fraud threats, repeat offender activity and broader criminal methodologies. Participation in such arrangements could be encouraged in the Codes and may help services strengthen their Fraud Indicator Assessments and resultingly design responses which extend beyond platform-specific signals.

Governance and accountability

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

We agree with the proposed governance and accountability measures and particularly support the requirement for the annual review to assess the effectiveness of the measures in practice. We consider this should be evidenced by reference to relevant operational outcomes and how

findings from the Fraud Indicator Assessment and internal assurance have informed improvements.

We would also welcome consideration of whether providers should make their annual assessments available to Ofcom, and where appropriate, relevant law enforcement bodies. This could support consistent supervision and earlier identification of areas where providers' measures may require improvement.

We note that Ofcom expects providers' governance procedures around the code of conduct and appropriate training to not cover advertising arrangements involving intermediaries. Clarification would be appreciated on how providers should assure themselves that relevant personnel working for intermediaries understand and apply the provider's standards; this includes whether appropriate training or equivalent requirements should form part of those arrangements.

Testing advertisement generation tools

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

We support the proposed testing of advert generation tools as a proportionate safety-by-design measure. Testing should help providers identify and address ways in which tools under their control could be used to create or modify fraudulent adverts.

We note that the proposal does not extend to external AI tools used independently by advertisers. We would welcome clarification how "effective control" will be determined where a provider enables a third-party AI tool to connect to its advertising service and create or modify adverts. Published examples indicate that this type of integration is emerging, creating potential ambiguity over responsibility for testing the tool and the functionality it allows.

Volume 3: Ensuring account integrity

Account checks and actions

- **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:**

We support the proposed account checks and actions measure, which offers a welcome upstream control.

a. **verification that advertising account holders work for, or on behalf of, the individual or organisation they are advertising;**

Verification of the relationship between an account holder and the individual or organisation being advertised is an important safeguard. We would welcome further clarity on what constitutes satisfactory verification, particularly in more complex arrangements such as affiliate marketing where an account holder may promote multiple organisations without a direct relationship to them.

It will also be important to recognise that verification of the relationship does not establish the legitimacy of the organisation or advertising activity. Where other risk indicators are present, providers should undertake proportionate additional checks using relevant and reasonably available information, rather than treating successful verification as sufficient in itself.

b. **checks and actions to prevent banned advertising account holders returning; and**

We support this measure and the use of relevant indicators identified through the Fraud Indicator Assessment to prevent banned account holders from returning. Checks should not be limited to attempts to create new accounts though; once an account holder is banned, providers should consider whether available information identifies other active accounts operated or controlled by the same account holder.

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 support proportionate restrictions where indicators identified through the Fraud Indicator Assessment suggest that an account presents a material risk of posting fraudulent advertising. We would welcome clarification on the action expected where indicators warrant further scrutiny but are insufficient to support a determination that the account is fraudulent. Any restrictions could reasonably remain in place until the identified risk has been appropriate assessed or mitigated.

d. **repeat checks where there are relevant changes identified on an account.**

We support this measure, in particular proposals that checks should be repeated where relevant changes are identified on an account, or where the provider no longer has confidence that an earlier check remains accurate.

In addition to the circumstances identified in the consultation we encourage Ofcom to identify additional potential triggers for repeat checks, such as changes in account administration, significant changes in advertising activity, and renewed activity following a period of account dormancy.

- **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.**

  Our suggested additional account checks and actions are set out under parts a to d above. We also refer to our response on information sources for the Fraud Indicator Assessment, particularly regarding the use of cross-sector intelligence to identify connected accounts, infrastructure and repeat activity.

Preventing fraudulent financial services advertising

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

  We particularly welcome this targeted preventative measure, reflecting our team's prioritisation of investment fraud as a high harm and high value threat.

  In implementing this measure, however, it will be important that providers can accurately identify financial services advertising for them to go through the verification process. Although some services allow advertisers to declare that an advert relates to financial services, the relevant category may not be selected. A particular challenge arises when advert appear to be alternative content, with the financial promotion element apparent only via the

click-through journey [*REDACTED*]. We therefore encourage Ofcom to clarify whether providers should consider both the advert and its immediate destination when determining whether financial services are being promoted.

Checks against the FCA Register, or use of the Firm checker and Warning List are important, but absence from the Warning List should not in itself be treated as confirmation that an advertiser has the necessary permissions. Relevant entries may not capture every name or domain associated with an entity.

Countering account takeover

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

  We support the proposed measures to counter account takeover, which addresses a recognised enabler of fraudulent advertising. We would welcome further clarity on the minimum security outcomes providers should achieve and how the continuing effectiveness of their account security mechanisms will be assessed as technology and associated threats evolve.

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

  We do not hold quantitative evidence that can estimate the frequency of account takeover or the proportion of fraudulent adverts originating from compromised accounts. However, publicly available evidence does demonstrate the continuing significance of account takeover, although it is not specific to advertising accounts.

Advertising bans

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

  We support the proposed advertising ban measure, specifically its focus on preventing repeat activity at account level rather than addressing fraudulent adverts individually.

  We would welcome further clarity on how providers should determine the responsible account holder, particularly where more complex arrangements exist (including the involvement of affiliates) with multiple parties holding different permissions. This will ensure appropriate action is taken against those responsible without disproportionate consequences for legitimate parties within the same arrangement.

  We would also seek assurance that Ofcom will assess whether bans are applied consistently and effectively, including where the fraudulent nature of the advertising only becomes apparent from its click through destination.

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

  A potential additional circumstance may arise where a legitimate advertiser's website or other advert destination has been compromised without its knowledge, causing users to encounter fraudulent content. Reversal in this scenario may be proportionate where the

provider is satisfied that the account holder was not responsible, the compromise has been resolved and appropriate steps have been taken to prevent further harm. Temporary restrictions may remain appropriate while those checks are completed.

Account appeals

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

J1 and J2: Appeals systems and accessibility
We support the measures proposed but welcome further clarity on how providers should assess information submitted through an appeal, particularly when responsibility may have been wrongly attributed to a party. Consideration should also be given to who should have appeal rights where multiple parties have different levels of permission or control over an account.

J3: Prompt determination of account appeals
While we support the requirement for appeals to be determined promptly, the term may be interpreted differently across providers. We encourage Ofcom to consider whether further guidance on appropriate timescales would support greater consistency, particularly where restrictions affect the ability of a legitimate business to advertise.

J4: Action following appeal
We support the proposed measure and would additionally encourage providers to give those that appeal clear information about the outcome and explain rationale behind any action. Ofcom should consider how the quality and consistency of appeal outcomes can be assessed, including whether providers should be required to undertake periodic reviews of appeal decisions. These reviews could include comparing similar cases to determine whether comparable facts and circumstances result in consistent outcomes, and monitoring whether accounts or advertisers that successfully appeal subsequently become the subject of further complaints, enforcement action or other indicators of fraudulent activity.

Volume 4: Moderation

Advertising moderation

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

Unless otherwise stated, our comments apply to the equivalent moderation measures for both Category 1 and Category 2A services.

Fraudulent advertising may appear legitimate when viewed in isolation. Effective identification increasingly requires consideration of advertiser behaviour, associated infrastructure, connected accounts and known fraud methodologies. Providers should therefore ensure moderation arrangements remain capable of identifying fraud that may not be detectable through content-level assessment or automated review alone.

C1: Review and assessment of suspected fraudulent adverts

We particularly support Ofcom's recognition that providers should consider relevant and reasonably available information beyond the content displayed within an advert. Effective assessment will require fraud knowledge and consideration of advertiser behaviour, connected accounts, destinations, and existence of redirect chains or cloaking. It should also consider whether material changes have been made to an advert after its initial approval, including changes to content or click through destination.

[*REDACTED*]

We therefore particularly encourage Ofcom to ensure moderation approaches are applied consistently so that identifying (or having reports on) one fraudulent advert prompts proportionate consideration of related activity and accounts to determine whether further account level action is required.

C2: Swift removal of fraudulent adverts

Whilst we support this measure, the meaning of "swiftly" may be interpreted differently across providers. We would welcome clarity on how Ofcom will assess whether action taken over content has been sufficiently timely to reduce user exposure and harm. For example, providers could be required to retain and, where appropriate, report information on the number of impressions, clicks or other engagements that occurred between the point at which content was identified for moderation and the point at which action was taken. This could provide a useful indication of the potential harm arising from delays in enforcement and help inform assessments of whether moderation actions are being taken sufficiently quickly.

C3: Internal advertising policies

The requirement for providers to set and record internal advertising policies is a useful foundation for consistent decision making. Where moderation is undertaken by a third party we would welcome clarity that the provider should ensure those carrying out the function have access to the relevant policy standards and enforcement guidance and apply them consistently. While it is appropriate for services to develop policies tailored to their platforms, Ofcom could consider setting expectations for key components that should be addressed within those policies. This should include clear policies and enforcement guidance relating to false impersonation and the unauthorised use of an individual's image, likeness or identity, including through deepfakes and other synthetic media. Where providers identify advertisements that appear to involve false impersonation or the use of deepfakes, we would expect these to be subject to appropriate moderation action, and subsequent removal if verified.

C4: Moderation performance targets

We welcome the inclusion of performance targets for advert moderation but request clarity on how Ofcom will assess whether provider-defined targets are meaningful and appropriately reflect moderation decisions in practice. Additionally, where moderation relies on automated or other non-human controls, performance targets should enable providers to assess their continuing effectiveness, including false positive and false negative outcomes, escalation to human review, and performance as fraud techniques evolve.

We note potential challenges when interpreting removal volumes. Lower volumes could indicate reduced prevalence but could equally reflect ineffective detection or moderation. Performance should therefore be assessed through a balanced set of measures rather than removal figures alone.

C5: Prioritisation of suspected fraudulent adverts

We support Ofcom's inclusion of trusted flagger intelligence and findings from the Fraud Indicator Assessment within prioritisation decisions. This should mean that adverts identified by trusted flaggers as likely fraudulent are prioritised by providers for review and moderation and the risk of managing the harm caused from a flagged ad is managed by the provider. The expectation should be that the provider consider the desirability of minimising users encountering the advert and the likelihood that an advert is fraudulent when making prioritisation decisions should lead to consideration of the severity and immediacy of the potential harm including likely financial loss and targeting of vulnerable users (such as older demographics).

We would encourage consideration that when a trusted flagger reports multiple connected adverts, accounts or click through destinations, providers should be able to prioritise and assess the activity collectively rather than treating each reported item as an unrelated case.

C6: Resourcing moderation function

We support the focus on adequate resourcing for content moderation. The approach should be consistent with the standards applied to other illegal online harms, whilst recognising the specialist knowledge required to identify fraudulent advertising.

Providers should ensure that suspected fraudulent advertising can be escalated to appropriately trained human reviews where contextual indicators cannot be assessed reliably through automated systems alone. Where moderation is outsourced, providers should retain effective oversight and ensure that third-party teams have access to relevant information, sufficient fraud expertise and appropriate quality assurance processes.

C7: Training and materials

We welcome the requirement for fraud-specific training and supporting materials but seek assurance that these are expected to apply consistently to directly employed and contracted moderation teams.

Terms of service and publicly available statements

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

We support the proposed measures, but note that clear and accessible published information, whilst useful for law enforcement and service users, will not in itself demonstrate that the relevant systems are operating effectively or being applied consistently.

Advertising complaints

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

  Unless otherwise stated, out comments apply to the equivalent measures for both Category 1 and Category 2A services.

  D1 and D2: Complaints systems and accessibility
  We support the proposed measures and would encourage Ofcom to ensure that complaints systems can accommodate reports concerning multiple connected adverts or accounts, and support bulk submission with references then assigned to individual reports. Our operational experience indicates that reporting routes can be difficult to navigate, particularly where services have multiple routes for different types of content, or allow only individual or small batch URLs to be reported at one time with key information requiring repeated input. These limitations affect the ability of law enforcement to report at scale.

  D3: Action in response to advertising complaints
  We welcome the requirement that a complaint should give providers a reason to suspect an advert or account may have fraudulent intent. We would request inclusion that the subsequent review should consider all relevant information on wider context provided by the complainant, and not just determine an outcome from the individual entity reported.

  [*REDACTED*]

  D4: Prompt determination of advertising appeals
  We support this measure but note that the term "prompt" may be interpreted differently across providers and would welcome clarification on how Ofcom will assess whether appeals have been resolved in an appropriate timeframe, particularly where the decision restricts a legitimate advertiser's activity.

  D5: Action following a successful advertising appeal
  We support the proposed actions and have no additional comments on this measure.

  D6: Unfounded complaints
  We support the proposed high threshold for treating complaints as manifestly unfounded and particularly welcome the exclusion of reports submitted through dedicated trusted flagger channels.

  D7: Dedicated reporting channels for trusted flaggers
  We welcome the proposed measures around dedicated trusted flagger reporting channels. Our experience shows that existing arrangements can be difficult to use at scale, including limits on batch sizes, fragmented reporting routes and limited visibility of report status and outcomes.

  We would encourage Ofcom to place requirements on trusted flagger channels to improve both the efficiency of reporting processes and the visibility of moderation outcomes. This would include functionality such as bulk reporting, unique case references, submission tracking, outcome notifications, clear reasons for decisions, information on any associated enforcement decisions and accessible escalation routes for any disputed decisions. Trusted

flaggers should be able to submit intelligence about connected adverts, accounts or infrastructure. We would also welcome efforts to establish greater standardisation in the trusted flagger experience across services, which could include minimum expectations for functionality, information requirements, feedback and escalation routes.

Trusted flagger DRCs should also support more bidirectional intelligence sharing, enabling providers to share relevant information on emerging threats, repeat offenders, associated infrastructure or moderation outcomes where appropriate, to support wider disruption efforts and improve future reporting.

[*REDACTED*]

Ad libraries

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

We strongly welcome the proposed ad library measure. Ad libraries are an important operational resource for law enforcement and other organisations, enabling scrutiny of adverts and identification of patterns of behaviour by bad actors online. Having advert and advertiser information available to a wider audience should facilitate the identification of fraudulent adverts that have incorrectly passed any moderation or proactive technology processes, so that the provider can be alerted to their presence and they can be swiftly taken down. We have found ad libraries to be a particularly useful tool for understanding how fraudulent advertising activity manifests in practice. The wider visibility of advert content and context greatly enhances the information we are able to submit to providers in trusted flagger reports, and has allowed us to develop our working relationship with these services. Additionally, ad library content can be particularly valuable when combined with other data sources, such as fraud reports. This can be used to contextualise individual reports, identify connections between separate incidents across platforms and build a more complete picture of the behaviour of the actors involved.

We do, however, welcome clarification on how equivalent transparency and visibility of advert content will be achieved where adverts are delivered through intermediaries and the provider does not hold all the information required for a complete ad library.

- **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 support Ofcom's proposed minimum requirements for the accessibility, visibility and searchability of public ad libraries. An API alongside keyword, exact-phrase and advertising account searches, with consistent advertiser information will support our cross-platform analysis and help identify common content and wider campaign patterns.

In addition, we would welcome some additional functionalities, such as the ability to search by destination domains and by payment identifiers. Further features, such as similarity searching across creative assets, advertiser or product names, and click through destinations could support the identification of related adverts and accounts. Additionally, historical versions of adverts, including changes to their content or destination, as well as information

on historical advertiser account changes would also assist in identifying evasion techniques and wider campaign patterns.

As a further useful feature, Ofcom could consider whether trusted flaggers should have access to additional information that may not be appropriate for a fully public library, such as relevant moderation and enforcement outcomes, identified fraud indicators, evidence of material edits or indicators of cloaking. Any enhanced access would of course require appropriate safeguards.

- **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 do not hold evidence on the technical process, costs or time required to build an ad library or adapt an existing library to meet the proposed requirements.

Annexes

Equality Impact Assessment and Welsh language assessment

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

  We agree with Ofcom's Equality Impact Assessment and have no additional comments.

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

  We agree with Ofcom's Welsh language assessment and have no additional comments.

Guidance on making fraudulent advertising judgements: updates to the Illegal Content Judgements Guidance

- **Do you agree with our proposals? Please provide any arguments and supporting evidence.**

  We support the proposed guidance and strongly welcome its recognition that information about an advert destination can be essential to determining whether it is fraudulent.

  We would, however, welcome greater clarity on the different treatment of landing pages and other destinations for Category 1 and Category 2A services, particularly where a visible advert appears legitimate and indicators of fraud become apparent only after interaction.

  The guidance should make clear that destination information may extend beyond the URL declared when the advert is submitted. Relevant additional information may include the final destination following redirects, destinations concealed through cloaking, app-store links and direct links to messaging services where there is no conventional landing page.

We would also welcome clarity on what destination information providers should regard as reasonably available, including the extent to which URL scanning or other technical capabilities should be used to identify redirects, cloaking and the destination presented to UK users.