Meta and WhatsApp response to Ofcom's Consultation on Fraudulent Advertising Codes of Practice

Consultation title:	Fraudulent Advertising Codes of Practice
To (Ofcom contact):	Ofcom Online Safety Team, FraudulentAdvertisingConsultation@ofcom.org.uk
Representing (self or organisation/s):	Meta Platforms, Inc. and WhatsApp LLC (joint response)

Introduction

We welcome the opportunity to respond to Ofcom's consultation on the Fraudulent Advertising Codes of Practice, and we support the shared objective of protecting UK users from fraudulent advertising. We recognise fraud as a serious and evolving harm, and we invest substantially in the systems, people and partnerships needed to detect and disrupt it.

We note that much of what the draft measures describe reflects practices we already operate at scale. Our feedback is therefore targeted: it is directed at ensuring the measures are proportionate, effective and workable in practice, having particular regard to the principles in Schedule 4 to the Online Safety Act, including that measures be technically feasible, proportionate to the size and capacity of the provider, and compatible with users' and advertisers' rights.

Ofcom's assessment of causes, impacts and risks

We broadly agree with Ofcom's characterisation of fraudulent advertising as diverse, adversarial and rapidly evolving, and with the observation that vulnerability is situational rather than confined to any single group. We would encourage Ofcom to give fuller weight to two points. First, that fraud is a cross-platform, cross-sector problem whose disruption depends on collaboration across the ecosystem (platforms, financial institutions, telecoms, law enforcement and regulators), such that platform-level measures alone cannot be expected to resolve it. Second, that the adversarial nature of the harm means overly prescriptive requirements can be counter-productive, as bad actors adapt to fixed rules.

We agree the characteristics relevant to risk that Ofcom has identified, including impersonation of public figures and brands, cloaked or unusual landing pages, generative AI and deepfakes, and suspicious account behaviours, are relevant and consistent with signals we use. We would stress that these are indicators rather than determinants: many of the same characteristics frequently appear in wholly legitimate advertising, so they are most effective when applied through risk-based, probabilistic systems rather than as fixed rules. We would caution against any approach that treats individual characteristics as dispositive, which risks over-removal of legitimate advertising.

Ofcom's proposed approach to designing the Codes

Taken as a package, we welcome the fact that most of Ofcom's measures appear to be appropriately targeted, workable and proportionate. We note Ofcom have drawn on their approach to tackling other harms in putting together this package and, broadly, we are supportive of this. A smaller number of measures, however, risk imposing a materially disproportionate burden relative to their incremental benefit, or raise real feasibility questions. Broadly, we would encourage Ofcom to preserve flexibility for providers to achieve the intended protective outcomes through the systems best suited to their service, in particular, in relation to measures regarding risk and governance (FAU B1 / FAS B1), verification (FAU H1 / FAS H1), and ad bans (FAU H5 / FAS H5). However, greater flexibility should not come at the cost of comprehensiveness: whilst the Codes must incorporate be proportionate in what they require, Schedule 4 also requires that the Codes must also be sufficiently clear and detailed that providers are able to understand what the measures entail in practice. We are concerned that certain measures suffer from this lack of detail, in particular, the measure relating to intermediaries (FAU K1 / FAS K1). We cover these points in more detail below.

Cross-cutting issue: transitory nature of categorisation

We note the measures generally treat services as if their categorised status is in permanent state. In reality, a provider's status as Category 1 or 2A may change as a result of changes to product features, fluctuations in user numbers relative to threshold conditions, or changes to the threshold conditions themselves. This issue is a particular feature of the Phase 3 additional duties requirements. This is not necessarily problematic where the measures are predominantly process based; however where substantial build work may be involved to bring providers into compliance, extra consideration must be given to the fact providers may only be categorised for a temporary period of time. We believe this is most relevant to the measures relating to advertising libraries (FAU E1 / FAS E1) and verification (FAU H1 / FAS H1). We suggest that the Code ensure these measures are calibrated to take account of the fact services may transition in and out of scope with relative frequency in some cases.

Risk, governance and control

Intermediaries measure (FAU K1 / FAS K1)

Harm arising in connection with advertising intermediaries is best understood not as a category of bad advertiser but as a structural problem in how advertisements reach the provider platform. Intermediaries, including (for instance): agencies, resellers, ad-tech platforms, affiliate networks and retail media networks run ads on behalf of other businesses, and sophisticated scam actors may sometimes deliberately route through them to scale their spending.

The defining difficulty is that the intermediary is often a genuinely legitimate business: its violating and non-violating spend are commingled, and the harmful client sits behind an account structure and may in some cases inherit the intermediary's accumulated trust.

Very often circumvention and recidivism rather than the first violation dominate the risk picture, the most effective approach is often one that targets entity and network-level accountability. In light of this, while we recognise the case for Ofcom looking to introduce measures to tackle these harms, we have some concerns around the way in which these have been framed and cast.

Our principal observation on Ofcom's intermediaries measure is one of **scoping**. We own and operate the advertising supply chain on our services, selling our own inventory directly, so advertisements placed through our proprietary interface are subject to a single set of standard terms, and systems and processes (associated with verification, modification among other things) whether or not an intermediary is involved in placing it. Ofcom's consultation material indicates that the measure should not apply in these circumstances, but its Codes of Practice do not necessarily reflect this, drawing no definitive line between walled-garden and open-display models expressly (noting that the definition of advertising intermediary in the Codes refers to actors "usually" operating within an open-display supply chain but not exclusively).

More generally, we have the following observations:

1. **It is overly broad.** Measure FAU/S K1 is cross-cutting in that it piggy-backs on Ofcom's suite of other measures for the most part and, as a consequence, is extremely broad in its potential impacts. It asks that providers, where they lack sufficient control and proximity in relation to the placement of the advert, implement a "version" of the relevant measure or "part" of it. It is not clear that such an approach, one that may lead to providers implementing a patchwork of partially implemented measures on grounds of having limited or partial "control" over ad placement, would lead to appropriate regulatory outcomes.
2. **It lacks sufficient legal certainty.** We are concerned that this measure does not express in sufficiently clear and detailed terms what providers must do in order to comply. We note that

the statutory duties in view (in sections 38 and 39) require “proportionate systems and processes”, in keeping with safety duties relating to illegal content and content harmful to children. It follows Ofcom’s recommended measures must specify, with sufficient particularity, what Ofcom considers proportionate for the purpose of meeting those requirements (noting, in particular, paragraph 2(b) of Schedule 4 to the OSA).

   a. As things stand, however, Measure FAU/S K1 does not do this and instead requires providers to use “all reasonable endeavours” to implement a version of other measures that is “as similar to the measure as possible”. This formulation effectively restates the general obligation to act proportionately without adding meaningful extra detail. A provider seeking to implement this measure cannot discern what steps would constitute compliance for the purposes of complying with the underlying statutory duty, which raises concerns around the lawfulness of such a measure.
   b. The measure appears to apply to in-scope providers where they have “insufficient control” over the placement of advertisements but it makes no attempt to set out what this means in practice or how it applies. Ofcom note that the lack of such detail is intended to confer flexibility but any such benefits are more than outweighed by the compliance risks associated with a measure that is quite so vague and broad.

3. **No account of alternative measures framework.** The Act permits providers to take alternative measures, recognising that Codes measures are one route to compliance, not the only one. Ofcom recognises this in its consultation proposals. The alternative measures framework depends on the Code setting a clear benchmark against which alternatives can be assessed. If the Code’s own recommendation is a nested chain of qualitative judgments - “all reasonable endeavours” to achieve “as similar as possible” - there is no fixed standard against which to measure an alternative. We appreciate that Ofcom has sought to offer flexibility in this regard but, in doing so, the approach raises questions about compatibility with foundational principles linked with the underlying statutory duties. More specifically, we note that a provider must also assess whether it may have only “partially” implemented that measure in every case. It is not clear whether an attempt to meet the underlying obligation by way of an equally effective alternative measure is itself to be treated as a partial attempt, with the record-keeping and reasonable-endeavours consequences that follow.

If Ofcom’s preference remains for something that is general, high level and far-reaching, we advocate for an approach that is less prescriptive in terms of measuring compliance to other measures, moving away from the concept of “sufficiency of control” that is overly nebulous and difficult to interpret. Instead we support an approach that confers the right on providers to *take account* of the other relevant controls within the Codes, and impose such of them that correspond to the risks reasonably apparent in the circumstances. Equally, where Ofcom considers that the measure should not catch providers that place ads on services solely via owned and operated models, Ofcom’s measure should be explicit about this, giving providers certainty for the purposes of judging compliance.

If Ofcom have in mind particular behaviours that providers should adopt in respect of particular harms of entities, we think that these should be set out with a sufficient level of detail such that providers can either adopt the recommendation (being clearer on its face as to what is required) or meaningfully consider alternative options that are equally effective at addressing the harms in view. Together we believe this would bring about proportionate outcomes: ensuring providers adopt precise, explainable levers that reduce violating spend without penalising benign advertisers or the legitimate businesses that intermediaries also serve.

Fraud indicator assessment (FAU B1 / FAS B1)

We note Ofcom proposes providers undertaking assessments regarding the indicators of fraud on their services.

We maintain a detailed understanding of how fraudulent advertising manifests on our services.

Our concern with FAU B1 / FAS B1 is not its objective but its prescribed methodology. By defining a “fraud indicator” as a discrete characteristic capable of being catalogued, tracked and reported, the measure assumes a static, binary taxonomy that does not exist in a probabilistic system, where a signal’s significance is a function of the other signals accompanying it. A provider whose continuously-learning systems prevent very large volumes of fraudulent advertising from reaching users would nonetheless have to reverse-engineer them into indicator vocabulary for compliance purposes, with conceivably little or no gain in user protection.

That prescription also sits uneasily with the Act. Parliament imposed standalone risk assessment duties for illegal content and child safety but not for fraudulent advertising, a distinction a code of practice should be cautious about displacing. Key terms compound the difficulty: “material risk” is undefined, though Ofcom accepts at Volume 2 paragraph 3.76 that too low a threshold would implicate legitimate advertisers; “appropriate analytical method” is undefined and, if restrictive, leaves providers unable to know where the boundary lies.

We therefore propose reframing the measure around its purpose rather than weakening it: a general obligation to maintain an effective, evidenced understanding of how fraudulent advertising manifests, expressly recognising continuous probabilistic detection as compliant, documented in periodic snapshots; definition of “material risk” and definition or removal of “appropriate analytical method”; review at intervals appropriate to risk, with interim updates counting; evidence drawn on where available, relevant and proportionate rather than by fixed checklist; and consolidated assessments where services share advertising infrastructure.

Testing advertisement generation tools (FAU F1 / FAS F1)

Scope

We note that FAU F1.1 applies “regardless of whether the tool is made available on or off the service.” The statutory duties in sections 38(1) and 39(1) require providers to operate “the service” using proportionate systems and processes, preventing users encountering fraudulent advertising “by means of the service.” If a provider makes a generative AI tool available through a separate platform, API, or interface that is not itself part of the categorised service, the statutory basis for Ofcom to prescribe how that tool is tested is unclear. If Parliament had intended to regulate the tools used to create advertising content (as distinct from the systems for detecting and removing fraudulent content on the service), it could make provision for that effect. We would welcome clarity around the basis for Ofcom’s measure in this regard.

The definition of “advertisement generation tool” is overly broad

FAU F1.2 defines “advertisement generation tool” as a tool used by advertising account holders “for the purpose of creating a relevant paid-for advertisement, by generating content or significantly modifying existing content.” This captures text generation, image generation, video creation, translation, and format adaptation tools. Many are general-purpose creative tools not specific to advertising. The line between “significantly modifying” content and “minor or superficial edits” is unclear and will be difficult to apply consistently, incompatible with Schedule 4, paragraph 2(b)’s clarity requirement.

The measure does not reflect how large providers operate at scale

Large providers do not develop and deploy advertising tools on a fixed timeline that maps onto the “before launch / after deployment / at regular intervals” testing cycle that FAU F1.5 contemplates.

The proposed framework of discrete testing events with logged results and formal decisions implies a waterfall development model that does not exist in modern AI product development at scale. Providers would either need to treat every model update as triggering formal testing (creating unworkable volume) or draw arbitrary lines around what constitutes a “significant” update (introducing definitional uncertainty).

Overly granular logging requirements

FAU F1.7 requires logs of who tested the tool, methods used, metrics produced, results with analysis, and mitigation steps. FAU F1.8 requires this log to be available to development, testing, operations, annual review, and senior accountability. For providers conducting continuous safety evaluation, this transforms ongoing iterative safety engineering into a discrete compliance exercise. It incentivises fewer, larger, formal testing exercises over the continuous integrated testing that is actually more effective. The three-year retention requirement compounds this for a field where model architectures change fundamentally within that timeframe.

The cumulative procedural overhead creates a drag on innovation not justified by the marginal safety benefit. Providers already have strong commercial incentives to prevent tools being used for fraud. The measure does not identify any gap in existing practice that, in our view, justifies this level of prescription. The risk is that it slows development of improved tools, including tools designed to improve fraud detection and prevention.

Ensuring account integrity

Verification (FAU H1 / FAS H1)

We support the objective behind FAU H1 and already operate much of what the measure contemplates across our platforms. Our comments are directed at ensuring the verification limb of FAU H1 is proportionate, appropriately targeted at the underlying harm, and workable, consistent with the principles set out in Schedule 4 to the OSA.

Our central point is that the need for verification should be driven by fraud risk, not applied universally to every advertiser before any UK user can encounter their advertising. Verifying that a person works for or on behalf of an organisation answers a question of identity or authority, but it does not answer the question that matters, which is whether the advertising is likely to defraud users. A verified advertiser can still commit fraud, and the great majority of unverified small advertisers will not. Anchoring a mandatory, pre-publication gate to the relationship therefore misdirects effort away from the actual harm.

We firmly believe that this approach is more effective against evolving threats and less harmful to legitimate small businesses.

Three points reinforce why a risk-based design is the more effective response:

- **Fraud is dynamic.** An account can pass verification at onboarding and then pivot to fraudulent activity. A front-door “one-time” check gives false assurance and does little against this behaviour, which is best caught through continuous monitoring across the life of the account. Over-reliance on the gate can crowd out the monitoring that actually works.
- **Fraud is often detected after submission, not at onboarding.** Fraudulent advertising is often identified through automated and behavioural signals rather than identity checks at the point of entry. A universal verification requirement would divert finite detection and review capacity towards the front door and away from where marginal harm reduction is greatest.
- **A fixed, universal gate is gameable and its burden falls the wrong way.** A known and uniform verification step is precisely what a determined fraudster is resourced to defeat, for instance through rented or mule businesses or purchased verified accounts. Legitimate small advertisers, by contrast, bear the full friction. Probabilistic, risk-based systems are harder to circumvent because they do not present a single fixed target, so the universal approach tends to burden the compliant while the determined bad actor routes around it.

Ofcom’s own framework points to the risk-based approach

This reasoning is reinforced by the structure of Ofcom’s own proposals. The fraud indicator assessment (FAU B1) is expressly designed to identify the characteristics that indicate an account or advertisement poses a material fraud risk. Where a provider maintains such an assessment and applies checks and restrictions by reference to it, requiring that same provider also to verify every advertiser in every case before publication is duplicative and internally inconsistent. The coherent design is to let the fraud indicator assessment calibrate the intensity of verification, so that enhanced checks and restrictions are concentrated on higher-risk accounts, while low-risk small advertisers are not subjected to disproportionate onboarding friction.

A universal pre-publication gate would produce disproportionate outcomes

Beyond being poorly targeted, a universal requirement would carry real costs that cut against the Act’s wider aims.

- **Impact on legitimate small business and fairness.** Advertising is a form of commercial expression, and a universal pre-clearance step operates as a prior restraint on lawful advertising that should be reserved for demonstrably high-risk cases. The friction also falls hardest on those least able to navigate a verification process, including digitally excluded, disabled and non-native-speaking sole traders, which sits uneasily with the goal of a safe and thriving online economy.
- **Harm displacement.** Onboarding friction can push small advertisers towards less safe channels or off-platform, which risks increasing overall consumer exposure to fraud rather than reducing it.

Suggested approach

We would support FAU H1 being framed so that verification and restrictions are applied on a risk-based basis, calibrated by the fraud indicator assessment, rather than universally and pre-publication; so that a small business carve-out or proportionate threshold is expressly recognised; and so that relationship and organisational checks are treated as available risk signals rather than a mandatory determinant of whether an advertiser may reach UK users. This preserves the protective intent of the measure while keeping it proportionate, effective against the actual harm, and consistent with the Schedule 4 principles.

Financial services verification (FAU H2 / FAS H2)

We support the objective of the proposed measure. Fraudulent financial services advertising is among the most harmful categories of fraudulent advertising, it is disproportionately associated with actors operating without FCA authorisation, and verification against the Financial Services Register is an effective and appropriate control. Our comments relate principally to the warning list component of the measure. As we explain below, while we see the warning list as an important tool to which providers may have cause to rely in performing checks, we query whether a minimum requirement to check against it for universal cases is the right approach having regard to the statutory scheme and the nature of the information presently offered on the list itself.

FCA authorisation

We understand Ofcom wishes to target the communication of a financial promotion not issued or approved by an FCA-authorised firm contrary to section 21 FSMA 2000 (para 3.2 of Volume 3). This corresponds to the way fraudulent advertising is defined in the Act, by reference to section 25 of FSMA, which criminalises the section 21 prohibition. We note the application of these provisions depends principally on whether the entity in question has FCA authorisation.

It therefore follows that Ofcom's measure FAU/S H2 requires providers to verify entities seeking to promote financial services against the FCA's register for authorised firms. However, we have concerns with the additional requirement that providers should also, in every case, be expected to perform additional checks against the FCA's warning list. We believe the fact of FCA authorisation is the correct proxy to judge risk in the circumstances. Not only does this correspond neatly to the underlying FSMA provisions in view, but it also obviates the risk of unnecessary duplicative efforts because FCA authorised firms should, by definition, not appear on the warning list.

Value of warning list

We do not question that the warning list is a valuable tool against which providers should have regard but we believe this should be required only where the providers considers it appropriate in the circumstances (where FCA authorisation checks have taken place). This approach has the advantage of leveraging the real benefits whilst avoiding not only the risks of disproportionate outcomes outlined

above but also risks of unfair outcomes for advertisers associated with false positives. In this respect, we note that providers are not privy to the evidential basis on which a firm is placed on the warning list, nor to the processes governing addition, review or removal. As we understand it, entries may reflect a suspected unauthorised promotion or equally reflect a consumer complaint, a name resembling a clone, a historic matter since resolved, or conduct with no bearing on UK financial promotion at all. This contrasts with FCA authorisation decisions because (i) those correspond directly with the fraud offences (regarding promotion), (ii) the basis for entry on and removal from the register is clearly established in law and FCA's own guidance, and (iii) there is a clear UK nexus.

We are therefore concerned that treating appearance on the warning list as itself decisive has the effect of sub-delegating the decision on whether a firm may advertise to a determination whose basis we cannot see and which may bear no link to the underlying prohibition. In cases where the FCA may err in including entities on its list, it is not clear to us there is any right of appeal and it would be hard for providers to assess this evidence against their inclusion on the warning list without significant operational overhead. Separately, we also note that the information on the warning list does not typically leverage links to underlying accounts on relevant social media services. This makes adopting scalable solutions challenging in principle.

We therefore encourage Ofcom to recalibrate this measure so recourse to the warning list is permissive rather than compelled in every case.

Practical difficulties

In any case, we note the FCA's warning list is published as a set of alert pages and firm profiles intended mainly for human consumption. It is not, so far as we understand, available as a structured feed with stable identifiers equivalent to a Firm Reference Number. Matching advertiser-supplied names, addresses, telephone numbers and landing-page URLs against roughly 18,000 free-text entries will produce both misses and false positives, and the tolerances needed to catch deliberate obfuscation are the same tolerances that will catch legitimate advertisers with similar details. Ofcom's own estimate at paragraph 3.68, that a long list of possible matches may require further staff input to resolve, illustrates the point. In the event Ofcom decides to press ahead with a blanket requirement for providers to check against the list, we ask that the measure is reformulated to take account of the fact that, in its present form, the information and data offered is not apt for providers operating at scale to ingest.

What we suggest

We invite Ofcom to recast this element as risk-based rather than universal, and that recourse to the list is permissive rather than required in every case.

Advertising bans (FAU H5 / FAS H5)

We agree that advertising account holders who deliberately and repeatedly post fraudulent advertisements should be removed, and that providers should take reasonable steps to prevent their return. We are, however, concerned that Ofcom's proposals are too rigid and do not take proper account of content level controls.

Content-level action

We note that Ofcom concludes at paragraph 5.47 that warning and strike systems are insufficient, and that account-level action should follow the first instance.

The duty under sections 38(1) and 39(1) is to have proportionate systems designed to prevent individuals encountering fraudulent advertising. Where a provider (i) acts at content level so the advertisement can no longer be encountered by UK users, and (ii) issues a warning identifying the violation with enough specificity that the advertiser understands it and does not repeat it, users are protected and the conduct is corrected. That should meet the duty. Immediate account-level action adds no protection in that scenario and disregards the content-level controls built to deliver it.

Strict liability offences

Several of the fraud offences in view are strict liability, as Ofcom accepts at paragraph 5.76, including section 25 FSMA. An advertiser can therefore commit a listed offence with no intent to defraud and no awareness of having done so, for example by promoting a financial product without realising authorisation or approval was needed. We note that Ofcom accepts such cases may be technical breaches with materially less harm, and they are often easily remediated once explained. A mandatory first-instance ban imposes serious economic and expression consequences where a warning would achieve the same protective outcome.

We therefore suggest FAU H5 be framed so providers may issue warning(s) alongside content-level removal before banning. That preserves Ofcom's objective for deliberate and repeat offenders while avoiding disproportionate outcomes at the margin.

Extension to proxy content compounds the interference

The interference in service providers' business models and on advertisers is further compounded where Ofcom believes that the measure should apply to proxy fraudulent advertisement content as well as illegal fraud content. This creates a lower bar for enforcement where any violation of a service provider's fraud-related terms of service will lead to the proposed banning enforcement. This approach has no regard for the severity of the violation and low quality content that violates terms of service but is legal is treated in the same way as severe, illegal content.

Ofcom's answer at paragraph 5.71, that commercial incentives will keep proxies narrow, does not meet the point. The difficulty is not proxy breadth but the absence of gradation in the consequence once a proxy is breached. Ofcom has also not identified the categories of lawful activity falling within providers' fraud-related terms, or their volume or harm profile. Without that, we are concerned that the proportionality case is not made out.

We therefore invite Ofcom to recalibrate its measure (i) to ensure mandatory banning is reserved for the most serious or risky cases, and in any event (ii) where content assessed as a fraudulent advertisement under the fraudulent advertising judgements guidance (rather than to be proxy).

Moderation

Advertising libraries (FAU E1 / FAS E1)

We recognise the need for transparency in this field, illustrated by the fact we maintain an industry leading ad library which continues to be relied on and valued by a variety of different groups. We have some targeted feedback on certain aspects of the measures, insofar as they require the inclusion of previous advertiser account names and keyword transparency.

We note, in particular, the requirement for relevant keywords which (when entered by users) bring up the advertisements in view. We suggest introducing an effectiveness threshold. We urge Ofcom to

evaluate the true value of this requirement against risks to privacy and the competitive sensitivity of exposing proprietary advertiser targeting strategies.

We would welcome clarification about Ofcom's proposal in the draft Code of Practice to provide data about "the keywords which, when searched by United Kingdom users on the service, will prompt those users to encounter the relevant paid-for advertisement". We believe that Ofcom's intent is that an ad library must have the functionality to search by keyword (paragraphs 5.37 to 5.41 of Volume 4) rather than requiring that every relevant keyword is listed in the ad library. The latter interpretation would i) represent a risk to advertisers' privacy and present competitive sensitivity of exposing proprietary advertiser targeting strategies, ii) provide limited utility where specific keywords have not had a significant role in targeting an advertisement and iii) raise significant technical difficulties. Instead, we agree with Ofcom's evidence that keyword search functionality can better enable the identification of fraudulent advertisement and would welcome Ofcom reflecting that it is search functionality rather than a comprehensive list of keywords that should be required within an ad library in their Code of Practice.

Further, we would welcome Ofcom introducing a criteria of effectiveness / utility in their recommendation about which data should be included in relation to advertisements on the ad library. For example, it is unclear how the provision of previous advertising account names related to a specific advertisement helps the understanding of how fraud ads are placed. Where data attributes do not help with identifying trends/understand fraud ads, we do not think it is proportionate for service providers to include the attribute.

We would welcome Ofcom's clarification that either:

(i) such a user journey does not constitute a user "encountering content consisting of fraudulent advertisements by means of the service" and it is instead by means of the ad library which is a portal made available by the service provider; or

(ii) the fact the advertisement was not targeted at UK users represents a proportionate and effective process for preventing UK users encountering fraudulent advertising.

Were Ofcom to consider a UK user encountering a fraudulent advertisement targeted at international users via the ad library as falling short of its recommended measures, an alternative approach may be to limit the functionality of the ad library for UK users and researchers so it surfaces only advertisements targeted at UK users. This would reduce the effectiveness and transparency of the ad library which we do not believe to be the intent of Ofcom's measure.