Consultation response

Field	Value
Responder	Moonshot
Responder Type	organisation
Org Role	Senior Communications Manager
Org Type	Other
Other Type	Private Company

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?

No response

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.

Moonshot's investigative work provides evidence to support Ofcom's assessment that fraudulent advertising is becoming more sophisticated, more adaptive, and harder to disrupt through platform reporting alone. The characterisation in Volume 1 of fraudulent advertising as a "highly adversarial harm type" where perpetrators "continuously adapt and evolve their tactics to evade detection" has been observed by Moonshot in a sustained, live investigation.

The Brain Honey operation uses AI-manipulated footage of a globally well-known public figure to fabricate personal endorsements of fraudulent memory loss supplements. It operates through a network of paid advertising accounts across Meta platforms, routing users through spoofed health news websites impersonating CNN, CBS News, and NBC's Today.
More than 4,000 fraudulent paid adverts tracked. Since March 2026, Moonshot has identified more than 4,000 fraudulent paid advertisements across Facebook, Instagram, and Threads, posted by hundreds of unique accounts across multiple reporting periods. The same infrastructure and creative assets are used to target multiple high-profile individuals simultaneously, consistent with a professional, scalable operation.

Ofcom's list of risky characteristics, including the impersonation of public figures, use of deepfakes, cloaked landing pages, and suspicious account behaviours, accurately reflects what Moonshot has documented. We would add one further characteristic: operator-level recycling of creative assets across campaigns targeting multiple individuals simultaneously. In the Brain Honey investigation, the same AI-manipulated video assets, spoofed domain templates, and shell company administrator accounts were used concurrently across scam campaigns targeting at least four separate public figures. This cross-campaign reuse is a meaningful fraud indicator that current platform detection frameworks do not appear to capture.

Consumer reports submitted to the Better Business Bureau, a US consumer protection organisation, document financial losses incurred by victims of this scam in the United States. Given the international infrastructure of this operation, with indicators pointing to operators based in Vietnam and Brazil and confirmed UK ad targeting, the financial harm documented in the US is consistent with harm being caused to consumers across jurisdictions, including in the UK. The scam specifically targets older demographics through health-related claims about Alzheimer's and memory loss, an audience that may be less equipped to identify deepfake content or spoofed news sites.

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

No response

5. Our approach to developing Codes measures

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

No response

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.

No response

Volume 2: Risk, governance and control

2. Advertising intermediaries

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

No response

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.

No response

3. Fraud indicator assessment

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

No response

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?

Moonshot's investigation identified several fraud indicators that Meta's existing systems failed to flag prior to our escalation. These indicators were visible at the account level before fraudulent content was even posted. We offer them as direct evidence of information sources providers should incorporate into their fraud indicator assessments (FAU B1/FAS B1):

Bulk-purchased aged accounts. The advertising accounts behind the Brain Honey campaign were almost certainly rented or purchased from services that create and sell aged social media accounts in bulk. Nearly all were originally registered under Vietnamese names before being changed to anglicised names. They contained fabricated biographical data - non-existent Instagram profiles, false addresses, fictitious phone numbers - designed to pass account review processes. Multiple pages were administered by two Brazilian shell companies, providing an appearance of business legitimacy that assisted in clearing Meta's ad review processes.

Shared infrastructure signals across accounts. Across hundreds of apparently independent accounts, Moonshot identified:
a single shared US phone number associated in open-source checks with wider fraudulent activity
recycled profile images (15 of 29 early-stage accounts sharing a single group photograph)
parent-page relationships revealing that what appeared to be organic user activity was in fact centrally coordinated.

These signals are detectable before fraudulent content is posted.

Domain registration patterns. The scam landing pages used in this campaign were predominantly registered through low-cost registrars in the same month the advertising campaign launched, with WHOIS privacy protection and Cloudflare hosting obscuring true ownership. Several domains hosted multiple concurrent scam campaigns - confirming they function as shared multi-campaign infrastructure. Platforms should consider incorporating domain registration signals - age, registrar, hosting provider, and WHOIS privacy status - into their fraud indicator assessments.

Simultaneous multi-target campaigns. Moonshot's investigation found that the same advertising accounts and domain infrastructure used in the Brain Honey campaign were simultaneously running scam campaigns using the likenesses of at least three other public figures, including a weight loss scam using the likeness of a prominent US television personality. The ability to identify accounts running multi-target campaigns would be a significant addition to platform fraud indicator frameworks.

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.

No response

5. Testing advertisement generation tools

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

No response

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 :

Moonshot strongly supports the proposed account checks measure, and can offer direct evidence in support of each of the three specific checks Ofcom proposes:

a) Verification that account holders work for who they say they do. The Brain Honey campaign used accounts administered by Brazilian shell companies that had no apparent legitimate business relationship with the individuals or organisations being advertised. In several cases, the administrator name associated with the account differed entirely from the page name and the contact details listed, consistent with fabricated business identities designed to pass verification. Robust verification of the advertiser-administrator relationship would have been a material barrier to this campaign operating at scale.

b) Preventing banned advertising account holders from returning (FAU H5). This is the most critical measure Ofcom is proposing, and Moonshot's evidence directly demonstrates why current approaches are failing. Following Meta's removal of identified Brain Honey advertising accounts, the same network continued operating through replacement accounts with near-identical characteristics - same shell company administrators, same shared phone numbers, same recycled creative assets.

In a single month following enforcement actions, over 1,000 new advertisements were identified from more than 330 unique accounts. In the following month, this figure exceeded 3,000 advertisements.

Account-level banning without network-level disruption is insufficient. Ofcom's proposed measure should be explicit that "reasonable steps to prevent banned account holders from returning" includes identifying and acting on network-level signals - shared billing entities, administrator relationships, device fingerprints, and domain infrastructure - not only individual account identifiers.

c) Checks for accounts with indicators of material risk. The indicators described above - bulk-purchased aged accounts, shared phone numbers, shell company administrators simultaneously running unrelated scam campaigns - were all present before a single fraudulent advertisement was posted. Platforms should be required to assess these pre-posting indicators as part of their account checks and actions policy.

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

Moonshot found that the Brain Honey campaign used accounts administered by Brazilian shell companies that had no apparent legitimate business relationship with the individuals or organisations being advertised. In several cases, the administrator name associated with the account differed entirely from the page name and the contact details listed, consistent with fabricated business identities designed to pass verification. Robust verification of the advertiser-administrator relationship would have been a material barrier to this campaign operating at scale.

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

Moonshot believes this is the most critical measure OFcom is proposing and our evidence directly demonstrates why current approaches are failing. Following Meta's removal of identified Brain Honey advertising accounts, the same network continued operating through replacement accounts with near-identical characteristics - same shell company administrators, same shared phone numbers, same recycled creative assets.

In a single month following enforcement actions, over 1,000 new advertisements were identified from more than 330 unique accounts. In the following month, this figure exceeded 3,000 advertisements.

Account-level banning without network-level disruption is insufficient. Ofcom's proposed measure should be explicit that "reasonable steps to prevent banned account holders from returning" includes identifying and acting on network-level signals - shared billing entities, administrator relationships, device fingerprints, and domain infrastructure - not only individual account identifiers.

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

Moonshot evidence shows that the indicators described in a and b - bulk-purchased aged accounts, shared phone numbers, shell company administrators simultaneously running unrelated scam campaigns - were all present before a single fraudulent advertisement was posted. Platforms should be required to assess these pre-posting indicators as part of their account checks and actions policy.

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

No response

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.

Additional account check Moonshot recommends: Platforms should check whether a new advertising account's administrator has a relationship with accounts that have previously been banned or flagged for fraudulent activity. In the Brain Honey investigation, parent-page relationships connected newly-created accounts to previously banned operators in a way that would have been identifiable through administrator network analysis.

3. Preventing fraudulent financial services advertising

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

No response

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.

No response

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

No response

5. Advertising bans

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

No response

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

Moonshot has no evidence that advertising bans are being over-applied in this context, and no circumstances where reversal would be proportionate. Our concern is the opposite: that bans are applied too narrowly (account-level only) and too slowly to prevent network regeneration.

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.

No response

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.

Moonshot strongly supports the proposed trusted flagger framework and dedicated reporting channel (FAU D7/FAS D7). This proposal addresses one of the most significant structural gaps in the current enforcement ecosystem.

Moonshot's experience is that standard platform reporting mechanisms are systematically insufficient for sophisticated campaigns. Content that does not fall cleanly within a defined policy category typically survives routine reporting. In the Brain Honey investigation, Moonshot secured the removal of a fraudulent video not through standard reporting but through direct escalation to a platform contact with a targeted evidence package making the case under an Intellectual Property policy; the content did not meet the threshold for removal under AI-generated content or fraud policies. This approach is effective but depends on prior relationship access that most organisations and individuals cannot obtain.

The formalisation of a trusted flagger framework would address this dependency on informal relationships. Moonshot would welcome clarity from Ofcom on the criteria and process for organisations to be recognised as trusted flaggers, and the minimum response timeframes that platforms would be required to meet for trusted flagger submissions. Defined SLAs for trusted flagger reports would be a significant improvement over current informal arrangements.

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.

No response

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.

No response

5. Ad libraries

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

Moonshot has used Meta's existing ad library as a core investigative tool and strongly supports the proposed ad library measure (FAU E1/FAS E1). The ad library was essential in identifying the scale and coordination of the Brain Honey campaign, allowing Moonshot to cross-reference accounts, identify shared creative assets, and map the network of shell company administrators.

Moonshot specifically supports the inclusion of "previous advertising account names" in the required information categories, as proposed. This field is directly relevant to detecting network regeneration: accounts that cycle through name changes to evade detection can be identified through historical name records in a way that current search functionality does not support.
However, significant limitations in current ad library provision constrained our work, and we offer these as

specific improvements Ofcom should require:

Geographic data gaps. The majority of Brain Honey advertisements were not served in the EU, meaning Meta's transparency data - which is most complete for EU-served content - did not capture audience reach figures for this campaign. Ofcom's proposals should address this gap for UK-targeted advertising that falls outside the EU's transparency framework.

Billing entity visibility. The ability to identify shared billing entities and shell company relationships across advertising accounts would significantly enhance investigators' ability to detect coordinated inauthentic advertising. Current ad library data does not surface this information. Moonshot recommends that Ofcom considers adding billing entity data, even in aggregated or anonymised form, to the required information categories for expert access.

Creative asset search. Ad library search currently requires specific account or page identifiers. The ability to search by creative asset fingerprint, domain, or phone number would substantially accelerate the identification of coordinated networks. Moonshot recommends this be added to the required functionalities.

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.

No response

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?

No response

Additional questions

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

No response

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

No response

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

No response