A behavioural audit of online services in the UK

Technical report

BIT | Ofcom

June 2025

BIT is a global research and innovation consultancy which combines a deep understanding of human behaviour with evidence-led problem solving to improve people's lives. We work with all levels of government, nonprofits and the private sector, applying behavioural science expertise with robust evaluation and data to help clients achieve their goals.

Find out more at bi.team

If you'd like this publication in an alternative format such as Braille or large print, please contact us at: info@bi.team

Get in touch with Ofcom's Behavioural Insight team at behavioural.insights@ofcom.org.uk

Authors
Sujatha Krishnan-Barman, Isabel Kaldor, Laurence Fenn, Libby Woodhouse, Ruth Persian and Eva Kolker

Click the **BIT** logo on any page to return to the Contents

Contents

Introduction

Background and context

Under the Online Safety Act (OSA) 2023, Ofcom regulates a wide range of UK online services to ensure that users, particularly children, are protected from online harms. This research is intended to contribute to the evidence base used to inform the development of Codes of Practice and guidance relating to Illegal Harms, Protection of Children, and the additional duties on 'categorised' online services under the OSA. In particular, it helps to build Ofcom's understanding of design practices in relation to specific aspects of online safety measures.

In this project, Ofcom commissioned the Behavioural Insights Team (BIT) to conduct a behavioural audit of popular social media and video-sharing platforms (VSPs). A behavioural audit involves systematically mapping online design practices and evaluating their potential impact on user behaviour and outcomes.

This audit, conducted between December 2024 and January 2025, examined how platform design influences user behaviour, with a particular focus on four core areas of interest (AoIs) discussed in greater detail in Research objectives. Establishing a baseline of current practices in these areas offers important contextual evidence for evaluating changes in online services' features and functionalities related to Ofcom's codes of practices under OSA.

Research objectives

This audit explored how platform design influences user behaviour, focusing on the four AoIs and research aims (Figure 1), along with Help Centre design and discrepancies between child and adult accounts across all AoIs.

Figure 1 - AoIs and Research Aims

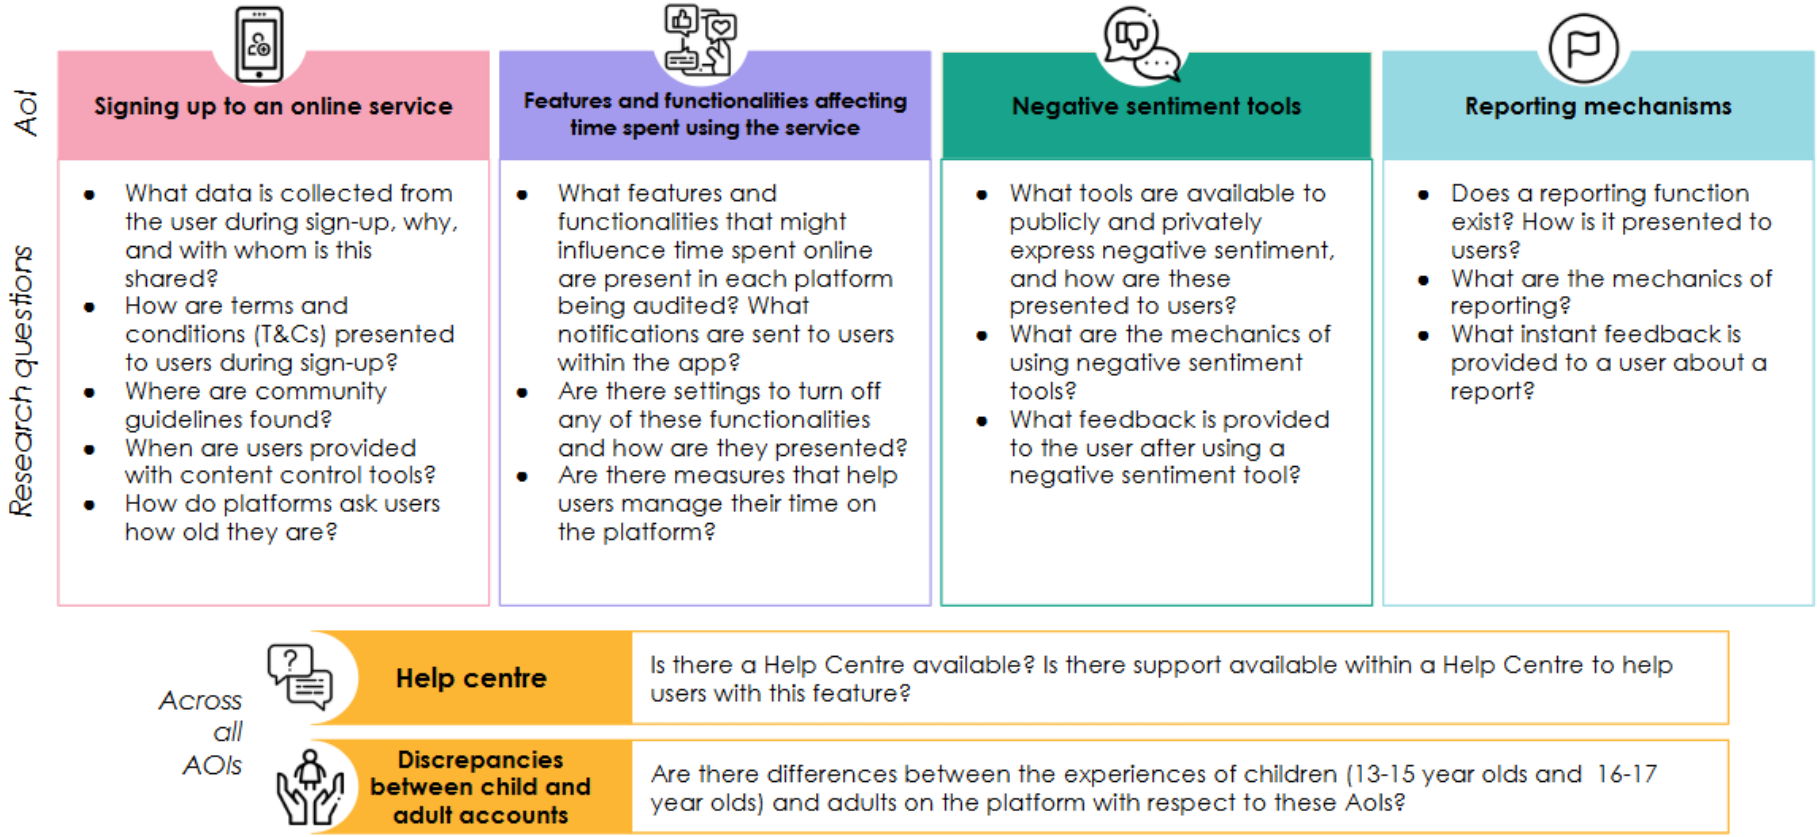

Audit methodology

Six online services were selected for the audit by Ofcom based on usage data[1] including the number of users and average time spent. For each platform, Ofcom also selected certain feeds to be included in the audit.

Feeds selected for the audit

Facebook

- **'Home'** – The main landing page of a user's Facebook profile, featuring a mix of posts from accounts the user follows and algorithmically recommended content based on their activity and engagement.
- **'Video for You'** – A personalised, algorithmically curated feed of videos selected based on the user's viewing history, interactions, and interests.
- **'Reels'** – A feed of short, engaging videos featuring music, effects, and filters. Content includes posts from accounts the user follows as well as algorithmic recommendations from Facebook.

Instagram

- **'Personal Feed'** – The main feed displaying posts from accounts the user follows, interspersed with algorithmically recommended posts based on their interactions and engagement.
- **'Reels'** – A dedicated section featuring short-form, vertical videos. This feed is algorithmically curated and includes a mix of videos from followed accounts and recommended content tailored to the user's interests and watch history.
- **'Search & Explore'** – While primarily a search function, this section also serves as a discovery feed, presenting algorithmically recommended posts and individual 'Reels' based on trending content and users' interests.

Snapchat

- **'Discover'** – A curated feed showcasing content from verified publishers, media companies, and creators. The selection is algorithmically determined based on user engagement and trending topics.
- **'Spotlight'** – A public feed of short, user-generated videos that are algorithmically selected based on engagement metrics, making them visible to all Snapchat users regardless of 'follow status'.

[1] Ofcom. (2024). *Online Nation 2024 Report*. Retrieved March 6, 2025 Available here

TikTok

- **'For You'** – The primary, fully algorithmic feed showing personalised video recommendations based on a user's past interactions, watch history, and engagement patterns.
- **'Explore'** – A discovery feed highlighting trending and viral content, featuring a mix of videos from creators the user follows and algorithmically selected content from a broader pool of users.

X (formerly Twitter)

- **'For You'** – An algorithmically curated feed displaying a combination of posts from accounts the user follows and recommended content based on engagement patterns, interactions, and trending discussions. This feed prioritises personalisation and suggested content over chronological updates.

YouTube

- **'Home'** – The default landing page when a user opens YouTube, featuring a mix of video recommendations tailored to their watch history, subscriptions, and trending content. This feed is heavily algorithmic and includes both followed and suggested content.
- **'Shorts'** – A section dedicated to short-form, vertical videos (60 seconds or less). The feed is algorithmically curated, combining videos from subscribed channels and recommended content based on user preferences and trending topics.

Ofcom notified the platforms that the audit was being conducted, and a transparency notice was published on the Ofcom website to keep the general public informed about the research.

Developing a framework and a codebook for the audit

As a foundation for the audit, we conducted a rapid evidence review of the online choice architecture (OCA) of social media platforms, focusing on the four AoIs.

Box 1: What is online choice architecture?

Online choice architecture (OCA) refers to the design of digital environments that influences how individuals make decisions and interact with online platforms. Based on the Competition and Markets Authority (CMA) definition[2], this comprises three components:

- **Choice structure:** how options are designed and presented
- **Choice information:** how users receive information about their choices
- **Choice pressure:** indirect influences affecting decisions

OCA impacts user behaviour through design elements such as option order, default settings, and the complexity of accessing controls. It also includes how information is presented (e.g., clear vs. dense terms of service) and features that apply pressure, like time-limited offers.[2]

We drew on research from Ofcom, the CMA, and other research in the nascent field of OCA practices on social media. Based on this, along with a review of Ofcom's A-Sparc model, we developed a taxonomy of the OCA practices which may be encountered during the audit.[3] These included those that may be **dark**, **grey or bright** in nature. **Dark patterns** are intentional design choices that manipulate or deceive users into making decisions that may not be in their best interests. **Grey patterns** have a more ambiguous impact, as their effects depend on the user's preferences and context—sometimes nudging users in directions that may not align with their best interests, while in other cases enhancing their experience. **Bright patterns** are designs that foster trust, loyalty, and respect between users and platforms. These practices help users make informed, intentional choices without pressure, promoting a more balanced and user-friendly experience.

In addition, when it comes to the features that influence the time spent online, we have expanded our list of OCA practices to include design features of online services, such as the way algorithms function and the way content is presented on the platform. These features are not directly tied to user choices but relate to their experience on the platform.

Creating the codebook

This taxonomy informed the development of a structured codebook for data collection, which contained detailed questions for each Area of Interest (AoI). Researchers responded to these questions as they conducted the audit.

---

[2] *Competition & Markets Authority. (2022). Online Choice Architecture How digital design can harm competition and consumers.* Available here

[3] *Ofcom. (2021). The A-SPARC model of online platforms 2 Contents Section.* Available here

To ensure consistency in coding and maintain a systematic approach, BIT implemented several measures. These included detailed annotations within the codebook, multiple codebook workshops, and bi-weekly meetings where researchers cross-checked coding decisions. These discussions helped align interpretations of instructions and definitions across the team. However, given the qualitative nature of the research, some degree of individual interpretation was unavoidable. When researchers encountered uncertainties, they were addressed collaboratively to ensure a consistent approach to coding across all data collection efforts.

Audit process

To conduct a comprehensive audit of social media platforms, we created four distinct Researcher Accounts, each designed to simulate different user experiences and test platform functionalities across various user types. These accounts were systematically used to explore differences in platform design, browsing experiences, safety features, and engagement strategies, and included:

- **13–15 year old account** – simulated the experience of a younger teen user
- **16–17 year old account** – represented an older teen user
- **Adult account** – examined the experience of a standard registered user
- **User without an account** – used to analyse what content and functionalities are accessible without an account

In addition, we created a second adult account on each platform to upload a piece of content. This allowed the other Researcher Accounts to report content that did not originate from a real user, enabling us to evaluate the reporting process without engaging with genuine user-generated content. This process is outlined in Figure 2.

Figure 2 - Outline of the audit process

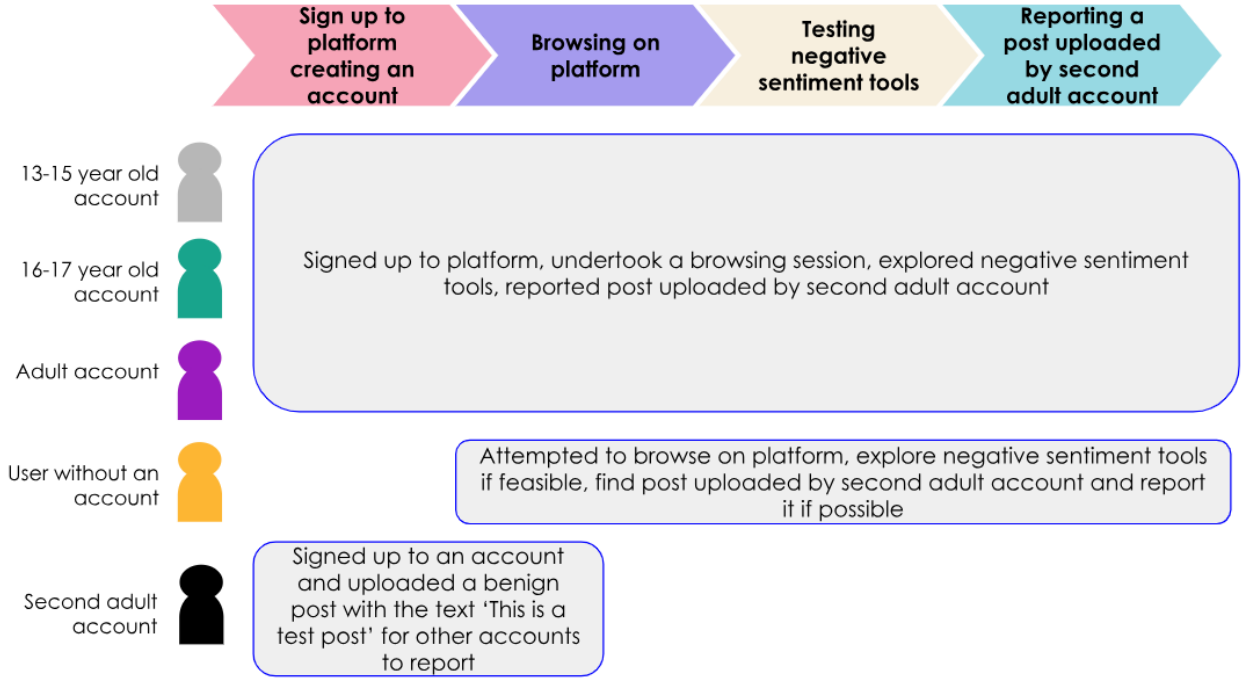

For the purposes of this audit, we define a *journey* as one Researcher Account going through the processes needed for one AoI on one feed/platform. Below in Figure 3, we outline all of the journeys completed as part of the audit, with supporting explanations given in Box 2.

Figure 3 - Audit journey breakdown

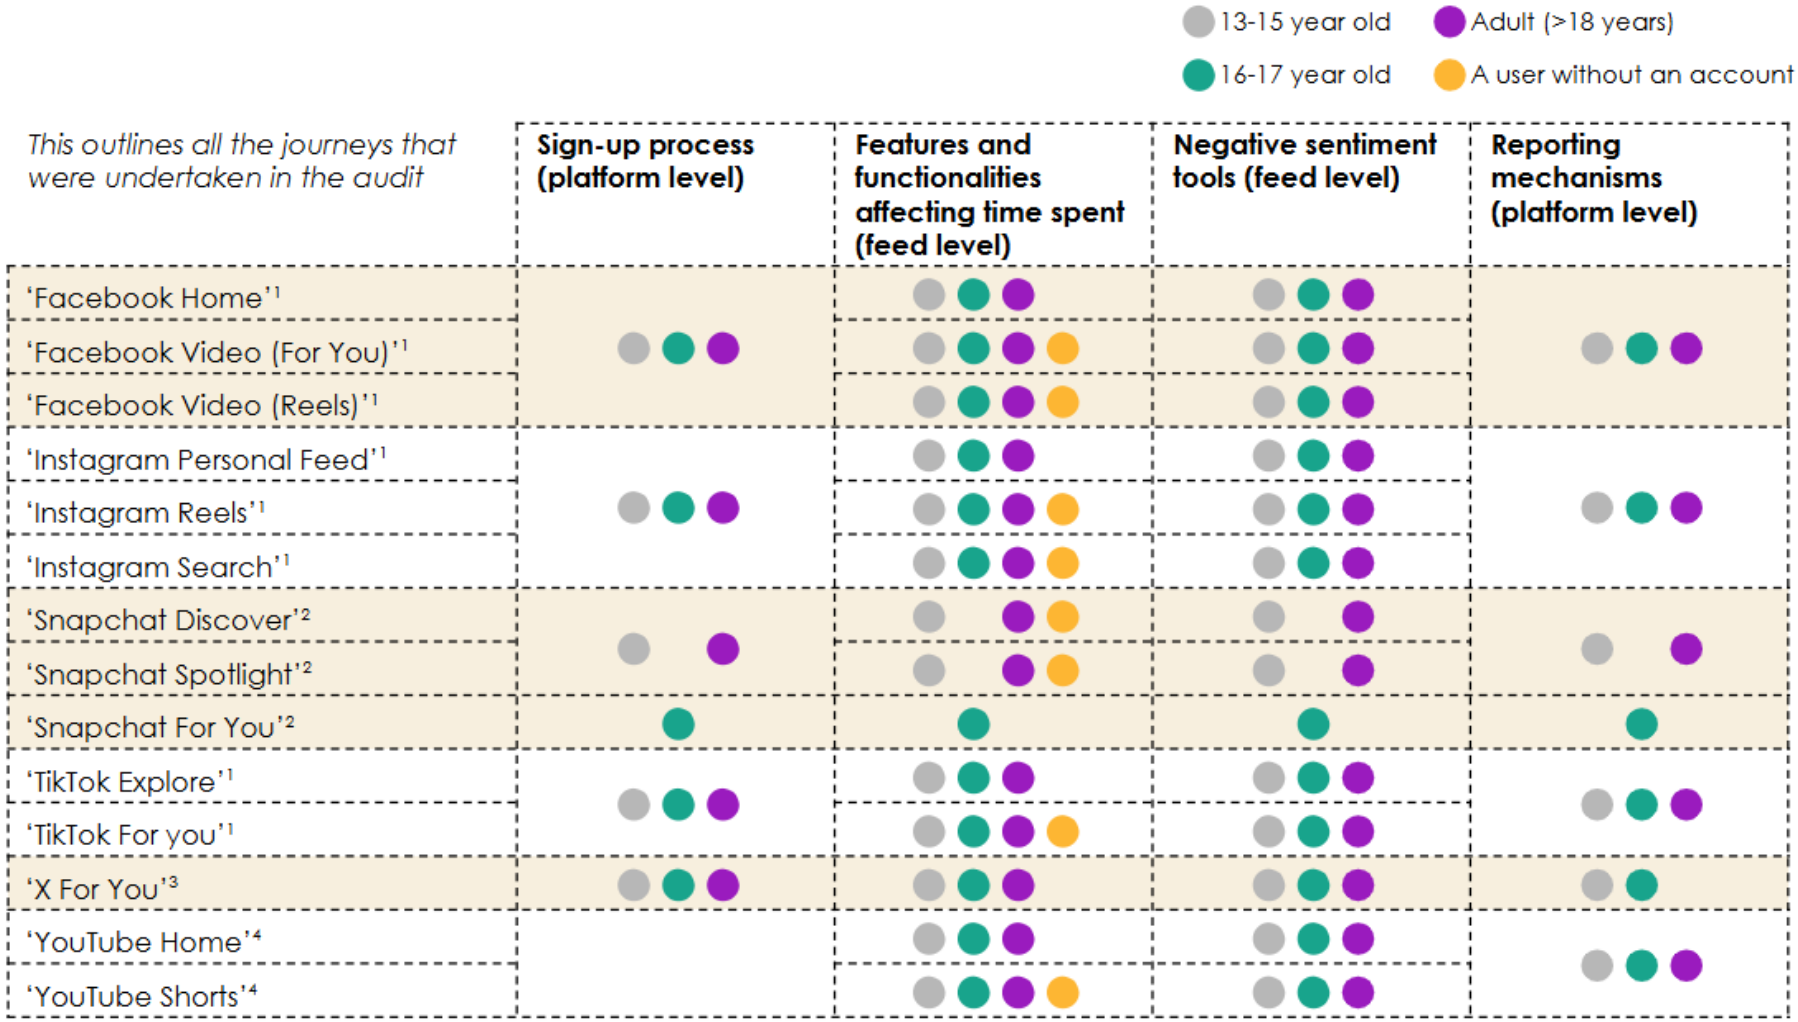

Box 2 - Notes on the audit journey breakdown

[1] Certain journeys could not be undertaken for a user without an account since it was not possible to access the feed or the tools without creating an account.

[2] During the period when the audit was being undertaken, the Snapchat app had a major update, moving to 'Simple Snapchat' which includes only three tabs (down from five previously). This meant that the 'Discover' and 'Spotlight' feeds were combined into one 'For You' feed. This affected the 16-17-year-old account, which only had access to the new 'Simple Snapchat'. However, this is unlikely to have impacted comparisons between child and adult accounts, as we were still able to compare the 13–15-year-old account with the adult account. Likewise, cross-platform comparisons are unlikely to have been impacted, as insights from the updated 'For You' feed were considered alongside findings from the earlier 'Discover' and 'Spotlight' feeds, helping to maintain consistency in analysis.

[3] Our Researcher Accounts were suspended so rapidly and repeatedly that it was not possible to complete the Reporting journey for the Adult account. We had two researchers instead use the 16-17 year old account to conduct two Reporting journeys independently to collect more data for analysis. From other journeys, we are confident that there are no differences in the Reporting journey between adult and child accounts.

[4] Signing up is only possible through an existing Google account or by creating a new one, which was outside the scope of this audit. There is no separate sign-up process for YouTube. Instead, we reviewed some of YouTube's policies related to the research themes for the sign-up section.

Setting up and conducting the audits

The audits were conducted using an Android smartphone using the relevant app for the 13-15 year old child, the 16-17 year old child, and the adult user. For the user without an account, the audit was conducted using a normal Chrome web browser. All the audits were conducted using alias email addresses (e.g. bitaudit3@bi.team). All Researcher Accounts followed the Rules of Engagement when engaging on the platform.

For each journey, researchers signed up or signed in using the relevant Researcher Account details, navigated the platform or feed, and completed the codebook. They also captured screen recordings and screenshots of the journey. To ensure consistency and minimise subjectivity, a senior researcher reviewed just under one-quarter (24.4%) of the journeys, providing quality assurance and an objective standard of analysis.

Rules of engagement for Researcher Accounts

The Rules of Engagement outlined in Table 1 were developed to address the ethical considerations for this audit to minimise the likelihood of our Researcher Accounts interacting with other users (particularly children) on the online services and to minimise the impact of any such interactions. These rules governed how the Researcher Accounts could behave on the platforms, outlining what they were allowed to do and proscribed from doing, and are in line with the Rules of Engagement used in previous work commissioned by Ofcom in this area.[4]

Table 1 - Rules of engagement

Type of engagement	Rules to be followed by Researcher Accounts
Account setup	All Researcher Accounts were set up as private accounts if that was allowed on the platform to minimise the chances of interacting with other real users' accounts. All demographic information (age, gender, location) was kept as private as possible within the constraints of the platform.
Following other accounts	Researcher Accounts aimed to not follow any accounts, and only engage with public content recommended by the platform. If this was not feasible on a particular platform, the Researcher Account only followed neutral public accounts with a minimum of 1000 followers so as to be following accounts where the impact of our Researcher Accounts following/unfollowing was likely to be minimal; this is in line with the parameters followed in previous Ofcom research.[4] Neutral accounts included accounts related to non-controversial topics such as sports, music and travel. Researcher Accounts aimed to follow a maximum of 10 accounts during the browsing phase if needed for the purposes of the audit. The only private account Researcher Accounts followed is the Audit upload account. Figure 2 outlines the different Researcher Accounts. The second adult account (which uploaded a piece of content for other accounts to report) did not follow any accounts.
Engagement with children online	Researcher Accounts only followed adult accounts. If there were any doubts as to whether the account holder was a child, Researcher Accounts took a risk-averse approach and applied a 'challenge 25' principle – if the account holder appeared (from the information and content present) under 25, the Researcher Accounts ceased further engagement. If, after applying this principle, there were still any doubts, Researcher Accounts also stopped any engagement in line with being risk averse. When considering whether an account holder was under 25, researchers took into account the following

[4] Avatar Methodology Pilot Study (2024) Available here

	indicators: references to recent birthdays, references to school or university activities, handles that hinted at birth or graduation years.
Engaging with content	Researcher Accounts only reacted to or commented on neutral content with a visible reaction (e.g., like, dislike) count of at least 1000 reactions on platforms where this was feasible. Where this was not feasible, Researcher Accounts relied on the number of followers of an account to serve as a proxy for popularity, and only interacted with creators who had at least 1000 followers. Neutral content included content related to non-controversial topics such as sports, music, and travel. Reactions and comments were standard positive emojis (e.g., thumbs up, heart) or short neutral statements ("Interesting"). The Researcher Account expressed negative sentiment on one post that already had a high level of engagement to test the engagement journeys. High engagement was at least 1000 reactions if this metric was available and displayed on platforms. Otherwise, researchers used their judgment to select a post where negative sentiment was unlikely to affect the creator. This negative sentiment took the form of a dislike (if the feature was available on the platform) or a "not interested" signal.
Direct messages	Researcher Accounts did not initiate direct messages with other users and did not respond to direct messages received from other users.
Reporting content	Researcher Accounts planned to only report content uploaded by the Audit upload account. For other content encountered by Researcher Accounts, they referred to the safeguarding actions table to decide actions to take as appropriate.
Uploading content	Other than the Audit upload account, none of the Researcher Accounts uploaded any content. The Audit upload account uploaded one piece of neutral content (a piece of text on a plain background saying “This is for testing purposes”) for the purposes of testing the reporting mechanisms. This account did not engage with any comments left on the content they had uploaded.

While these rules served as important ethical safeguards, they also imposed certain limitations on our research. For instance, we were unable to engage in specific types of interactions, such as messaging users or interacting with accounts with lower follower counts. That said, our audit primarily focused on structural patterns and platform design rather than user-to-user interactions. Therefore, while this constraint is worth noting for future research, we do not consider it a significant limitation for this study.

All information reflects the platform features and settings available at the time of the audit. Subsequent updates or changes by platforms may mean that some details are no longer current.

Ethical considerations and safeguarding procedures

BIT is committed to conducting research ethically and to the highest standard. This project was subject to our research ethics process, which meets the criteria set out by the UK Government's Social Research (GSR) Unit[5], the Market Research Society (MRS) Code of Conduct[6] and the Economic and Social Research Council's (ESRC) guidance on governance arrangements for research ethics committees.[7] Our ethics policies are regularly updated ensuring alignment with GSR. We have a research ethics panel consisting of trained staff who conduct project reviews. To ensure the independence of the ethics panel, only panel members who are not involved in the project in question can review it.

We recognise the increased level of risk associated with this research owing to the novelty of the research method, which makes it harder to anticipate challenges. Unlike traditional methods, where risks are well-documented, new approaches can come with unforeseen risks and vulnerabilities. In addition, there is a (limited, but present) possibility of interacting with children in a non-standard setting on online services. Finally, there is a risk that researchers working on the project may encounter sensitive or harmful content online. The following ethics and safeguarding measures were put in place to mitigate these risks.

Prior to the audit

BIT and Ofcom developed Researcher Account Rules of Engagement prior to the audit to minimise the risk of interacting with children on online services.

During the audit

Researchers strictly adhered to the Rules of Engagement and were briefed on the safeguarding actions required when encountering harmful or illegal content of low, moderate or high severity (see table below). While our classification was informed by the Online Safety Act's classification of harms,[8] our tiers were designed to address the risks faced by adult researchers, rather than the harms that children may experience when exposed to such content. The tiers were structured around the specific safeguarding actions required of researchers, ensuring they had clear, practical guidelines while conducting the audit. Additionally, our classification includes a broader range of potential content than the OSA, reflecting the unique focus of our

[5] *Ethical Assurance Guidance for Social Research in government*. (2011). GOV.UK. Available here.
[6] MRS Code of Conduct (Mat 2023). Available here.
[7] *Governance arrangements for research ethics committees* (n.d.). UK Research and Innovation. Available here.

[8] *Protecting people from illegal harms online* (2025). Ofcom. Available here.

work on researcher exposure and response rather than regulatory compliance for platform safety.

Safeguarding actions

Safeguarding actions were designed in response to the risks involved in the audit, given in Table 2.

Table 2 - Safeguarding actions

Risk	Safeguarding actions
Researchers encounter **low-severity harmful content** online such as: • Fake news or misinformation • Scams or phishing • Exposure to content that may negatively shape body image (e.g. heavily edited or filtered content)	No safeguarding actions required.
Researchers encounter **moderate-severity harmful content** online such as: • Exposure to hateful content online • Exposure to pornography	Report this content to the platform using the standard procedures offered by the platform.
Researchers encounter **high-severity harmful or illegal content** online such as: • Terrorist content • Child abuse material • Extreme pornography	Report this content to the platform using the standard procedures offered by the platform. Take all necessary steps to report the content to the appropriate agency, according to Ofcom's published guidance.[9]

During the audit, Researcher Accounts encountered some instances of moderately harmful content, which were reported using the standard procedures available on each platform. However, they did not encounter any high-severity harmful or illegal content.

When sharing outputs of the audit with Ofcom colleagues, we removed personally-identifiable information from screenshots and shared only a limited number of screen recordings, specifically where they were relevant and where it was possible to minimise any personally-identifiable information.

[9] *Harmful online content: how to report it and where to get help* (2023). Ofcom. Available here; *What should you do if you've seen harmful content online?* (2023). Ofcom. Available here.

Avoidance of harm to researchers and Ofcom colleagues

Key risks for BIT and Ofcom staff were exposure to potentially harmful materials. BIT staff joined the project voluntarily after receiving a full briefing on potential risks and were free to withdraw at any time without penalties. All BIT staff could avail themselves of mental health support offered by BIT, e.g., trained Mental Health First Aiders or the Employee Assistance Programme. Should any project team member experience harm and withdraw, the project leads convened to understand the situation and determine how best to avoid future instances. To reduce the risk to Ofcom staff, BIT reviewed all materials and screen recordings that were shared with Ofcom and provided any necessary warnings or labels for potentially harmful content in advance. In addition, Ofcom has its colleague safeguarding and wellbeing processes which were followed as relevant.

Analysis and reporting

A mixed-methods approach was used to analyse the data. For textual data and images, the qualitative data analysis method of thematic analysis was used. This involved reviewing the codebooks in detail both within each platform and for each AoI across platforms to identify patterns and trends that emerged from the qualitative data. Where relevant, descriptive statistics were also used to describe trends across the audit sample. All outputs were reviewed by a senior researcher to provide robust quality assurance. The outputs from this analysis are available here.

58 Victoria Embankment
London EC4Y 0DS
+44 (0)20 7438 2500
info@bi.team
𝕏 @B_I_Team
www.bi.team

Registered office: 58 Victoria Embankment, London EC4Y 0DS.