# An evaluation of the impact of Twitch's content classification labelling

Economic Discussion Paper Series Issue 11

Published 2 September 2024

# Ofcom discussion paper series in communications regulation

## Ofcom's role as the regulator for the video-sharing platforms and Online Safety regimes

Ofcom is the UK's communications regulator, overseeing sectors including fixed line and mobile telecoms, the airwaves on which wireless devices operate, post and TV and radio broadcasting. We have regulated video-sharing platforms ('VSPs') since November 2020 and were formally appointed as the online safety regulator in October 2023.

The Online Safety Act, which was introduced at the end of 2023, is a new regulatory regime that brings in scope a dynamic technology industry where novel services and safety measures are emerging.

As the regulator for the VSP and Online Safety regimes, it is important that Ofcom evaluates safety measures deployed by platforms.

## The discussion paper series

Ofcom is committed to encouraging debate on all aspects of media and communications regulation and to creating rigorous evidence to support its decision-making. One of the ways we do this is through publishing a series of Discussion Papers, extending across economics and other disciplines. The research aims to make substantial contributions to our knowledge and to generate a wider debate on the themes covered.


## Acknowledgements

We are grateful to Twitch for agreeing to Ofcom using their API for research purposes.


## Disclaimer

Discussion Papers contribute to the work of Ofcom by providing rigorous research and encouraging debate in areas of Ofcom's remit. Discussion Papers are one source that Ofcom may refer to, and use to inform its views, in discharging its statutory functions. However, they do not necessarily represent the concluded position of Ofcom on particular matters.

# Contents

## Section



## Annex

# 1. Overview

1.1 This Economics Discussion Paper analyses the impact of changes that the video sharing platform Twitch made to its content classification labelling in 2023. This overview explains the changes Twitch made, and the toplines of our findings and analysis.

1.2 The research presented in this Economics Discussion Paper both advances our understanding of safety measures and their efficacy, and builds our internal skills and knowledge on using platform Application Programming Interfaces ('APIs') to collect data to inform research.

1.3 We focus on one VSP, Twitch, which is a platform where users can watch other people playing video games or chatting, or stream themselves doing so. We focus on Twitch for two reasons. First, it makes a rich dataset publicly available through its API, which allows us to measure, amongst other interesting metrics, views of live streams and metadata on these live streams. Second, in 2023 it made a significant change to its content labelling practices and we wanted to use this as a case study in how we might use API data to understand the impact of such a change to a platform's safety measures.

1.4 Specifically, in June 2023 Twitch introduced a change to its content classification guidelines, which applies to content creators. These new guidelines form part of Twitch's overall Community Guidelines. In this paper we describe this as the content classification labelling change ('CCL change'). The CCL change led to content creators being required to apply content classification labels to indicate to viewers if the stream they are about to watch contains certain mature themes.[1] This helps viewers make informed choices about the content they watch because they are provided with more granular information on the type(s) of content being streamed on the platform. These granular labels replaced a binary labelling system which previously allowed content creators to categorise their streams as either 'mature' or 'non-mature'.

1.5 In this paper, we treat CCLs as a proxy for mature content:[2] we describe streams which contain one or more of the six CCLs as having 'mature content'. The six CCLs used by Twitch are:

- Mature-Rated Game;
- Sexual Themes;
- Drugs, Intoxication, or Excessive Tobacco Use;
- Violent and Graphic Depictions;
- Significant Profanity or Vulgarity; and
- Gambling.

If a stream is not labelled with a CCL then it is not treated as mature content.

1.6 At a high level we set out to understand whether the new labelling practice changed viewing behaviour of mature content. As part of this research we answer the following questions:

i) Pre-CCL change, what was the proportion of mature-labelled content on Twitch?

[1] Twitch also introduced penalties to content creators that fail to accurately label their streams.
[2] This is content which is generally unsuitable for a younger audience because it contains mature themes.

ii) Pre-CCL change, did the time of day and day of the week have an impact on the percentage of content that is labelled as mature?
iii) Pre-CCL change, did labelling of mature content by content creators correlate with external maturity ratings of games?
iv) Did the CCL change improve content labelling accuracy and so increase adherence by creators to the Community Guidelines?
v) Did the CCL change affect content creator behaviour?
vi) Did the CCL change impact the number of views of mature content on the platform?

## Our Findings

The data available from the Twitch API can be used to understand which types of content streamers are creating and viewers are watching on the platform. We can measure the absolute number of streams created, the number of views of each of these streams, collect metadata describing these streams and, of particular interest, analyse the content classification labelling of these streams. We were able to use a variety of advanced quantitative techniques to measure the effects of the CCL change and come to robust conclusions.

Our main findings (which apply to both game and non-game streams) indicate that:

- **Content labelling accuracy by content creators increased substantially following the CCL change.** Following the CCL change, we saw an increase in creators labelling mature content as mature. We also saw a decrease in mis-labelling of streams as mature (where our analysis suggested such streams did not contain mature content). Separately, our analysis established a causal link between the CCL change and an increase in creators accurately labelling mature content as mature.
- **The CCL change did not materially alter the type of content produced by creators.** We did not see a significant change in the amount of mature content creators produced before and after the CCL change.
- **The CCL change did not alter viewer behaviour materially.** We found that the number of views of mature streams did not substantially differ before and after the CCL change.

In conclusion, we found that more specific content labels, coupled with penalties for inaccurate labelling, led content creators to positively change their behaviour in terms of how accurately they labelled content. This change resulted in users being provided with much clearer alerts for mature content. However, we did not see any significant impact on the amount of mature content users viewed following the change, nor did we see a significant change in the amount of mature content creators produced.

This research enabled us to apply quantitative techniques to a live platform change and assess its effects both on creators and viewers. Our work demonstrates the importance of carrying out in-depth analysis, which often brings about greater and more robust insights than could be achieved from simplistic comparisons of high-level metrics.

This work builds on our Economics Discussion Paper Series Issue 10, which highlighted the importance of evaluation in the online safety landscape.[3] That paper set out a step-by-step approach that could be used to evaluate online safety measures. This included discussion of different metrics and measurement approaches that could be used to assess the effectiveness of different safety

[3] Ofcom, May 2024, Evaluating online safety measures.

measures, and some key caveats and challenges associated with them. The techniques used in this current research paper are a good example of Ofcom working with a platform to evaluate the impacts of a change in a specific safety measure.

While each online platform will have its own idiosyncrasies, this research shows that for Twitch, changes to its labelling rules for content creators positively influenced their behaviour in terms of labelling accuracy, but that more accurate and detailed information on content did not necessarily have much impact on the behaviour of content viewers or on the types of content produced by creators.

Separately to this work, Ofcom recently announced some changes that Twitch made to its platform to better protect children, following engagement with Ofcom's Enforcement team.[4] We had concerns about the availability of harmful material for under 18s on Twitch. The changes Twitch made involve restricting access to content with a 'sexual themes' or 'gambling' CCL for logged in under 18 year-old users and for all logged out users (of whose age Twitch cannot be assured). This Economic Discussion Paper does not form part of that Enforcement-led engagement but is independently instigated research that looks at similar safety measures.

1.7 The remainder of this paper is structured as follows:

a) Section 2 describes the mature content labelling system on Twitch before and after the CCL change.
b) Section 3 describes our data sources and summarises our data collection method.
c) Section 4 assesses the overall level of mature-labelled content on Twitch before the CCL change. In particular, we explore the first three research questions discussed above. We answer these research questions by applying statistical analysis to the data collected for this research project.
d) Section 5 evaluates the impact of the CCL change on content creator and viewer behaviour. This section contains our analysis of the last three research questions discussed above. We answer these research questions using a combination of standard statistical analysis, econometric analysis, and machine learning techniques.
e) Annex 1 presents the details on a classifier model (a type of machine learning technique to predict which CCLs, if any, are appropriate for streams on Twitch) which used CCL classification in the period after the CCL changes to predict classification in the pre-implementation period.
f) Annex 2 provides an in-depth explanation of the econometric models used to evaluate the impact of CCLs on content creator and viewer behaviour.
g) Annex 3 describes our data collection methods in more detail.
h) Annex 4 presents the results from sensitivity analyses that assess the robustness of our baseline findings.

[4] Ofcom, May 2024, Ofcom secures better protections online for children using VSP Twitch

# 2. Mature content labelling on Twitch

2.1 Twitch is a VSP predominantly known for its livestreamed content. The platform heavily features video game livestreams; users can watch other people playing video games or stream themselves doing so. Twitch also broadcasts e-sports competitions, as well as music and other entertainment broadcasts. Twitch is owned by Amazon and is integrated with Amazon Prime gaming.

## Twitch mature content labelling system prior to summer 2023

2.2 Prior to 20 June 2023, Twitch gave creators the option to label their content as mature in their ‘Creator Dashboard’ before starting a new live stream, by enabling this through their setting of the ‘Mature Content’ toggle. Little information was provided to help content creators understand what should be considered as mature (as shown in Figure 1).

**Figure 1: Twitch’s mature labelling option pre-summer 2023 change**

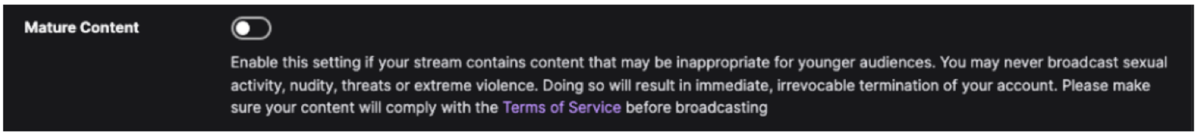


*Source: Twitch website pre-June 2023 change*

2.3 When a piece of content was self-labelled as mature in Twitch’s previous labelling system, viewers were served a warning stating that the channel was intended only for mature audiences. Viewers would then need to click-through to view the content (as shown in Figure 2).

**Figure 2: Twitch’s warning message on mature content pre-summer 2023 change**

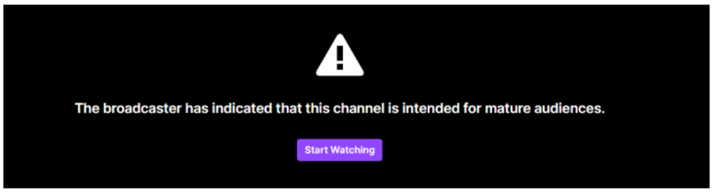


*Source: Twitch website pre-June 2023 change*

## Twitch mature content labelling system after summer 2023

2.4 On 20 June 2023, Twitch introduced a broader set of content labels to characterise mature content and required content creators to apply these. As part of this change Twitch applies penalties to content creators that fail to accurately label streams.[5] Content creators were required to apply CCLs for mature content by visiting 'Stream Manager' within 'Creator Dashboard' and clicking 'Edit Stream Info'. Unlike the previous binary Mature Content toggle, the new labelling system offered content creators a more precise and granular approach to labelling their content and provided detailed information regarding when each label should be applied.[6] There are currently six labels which content creators are required to use when relevant:

- **Mature-Rated Game**. If a content creator plays a game with an Entertainment Software Rating Board (ESRB) rating of mature, their stream will automatically be labelled as mature.[7] The ESRB is a North American body which sets age ratings for video games.[8]
- **Gambling:** Participating in online or in-person gambling, poker, or fantasy sports that involve the exchange of real money.[9]
- **Violent and Graphic Depictions:** Simulations and/or depictions of realistic violence, gore, extreme injury, or death.[10]
- **Drugs, Intoxication, or Excessive Tobacco Use:** Excessive tobacco glorification or promotion, any marijuana consumption/use, legal drug and alcohol induced intoxication, discussion of illegal drugs.[11]
- **Sexual Themes:** Content that focuses on sexualized physical attributes and activities, sexual topics, or experiences.[12]
- **Significant Profanity or Vulgarity**: Prolonged and repeated use of obscenities, profanities, and vulgarities, especially as a regular part of speech

[5] "Streamers will be expected to review the Content Classification Guidelines and familiarize themselves with the type of content that will need to be labeled. If streamers fail to accurately label their streams, they will receive a warning and a label will be applied to their stream. Repeated violations will result in the content classification label being applied to your channel for a set duration of time, but will not result in an account suspension." https://help.twitch.tv/s/article/content-classification-labels?language=en_US
[6] Multiple CCLs may be used simultaneously, and these can be applied or removed at any point during a stream.
[7] If the stream contains an unrated game, Twitch automatically assigns the correct content label setting if it has enough information from The Internet Game Database (IGDB) to mark the stream. The IGBD is an online database about video games that Twitch acquired in 2019.
[8] ESRB Mature – rated games are deemed suitable for ages 17 and over. There is also a ESRB Adults only rating which is deemed suitable for ages 18 and over. Twitch prohibits these games from the platform.
[9] Sharing gambling referrals to sites that contain slots, roulette, or dice games is prohibited according to Community Guidelines, as is streaming or sharing links to specific sites.
[10] Content that displays realistic gore including death, extreme injury, mutilation or bodily fluids is still prohibited regardless of whether a label is applied.
[11] Twitch prohibits the use of hard drugs and the misuse of legal substances.
[12] Twitch prohibits nudity, inappropriate attire, pornography, and certain sexual content.

### Figure 3: Twitch's granular mature labelling option post summer 2023 change

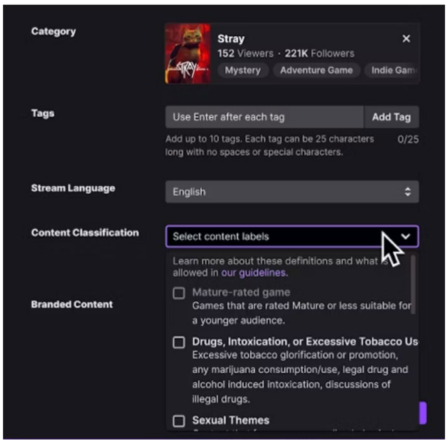


*Source: Twitch website post-June 2023 change*

**Figure 4: Twitch's warning message on mature content post summer 2023 change**

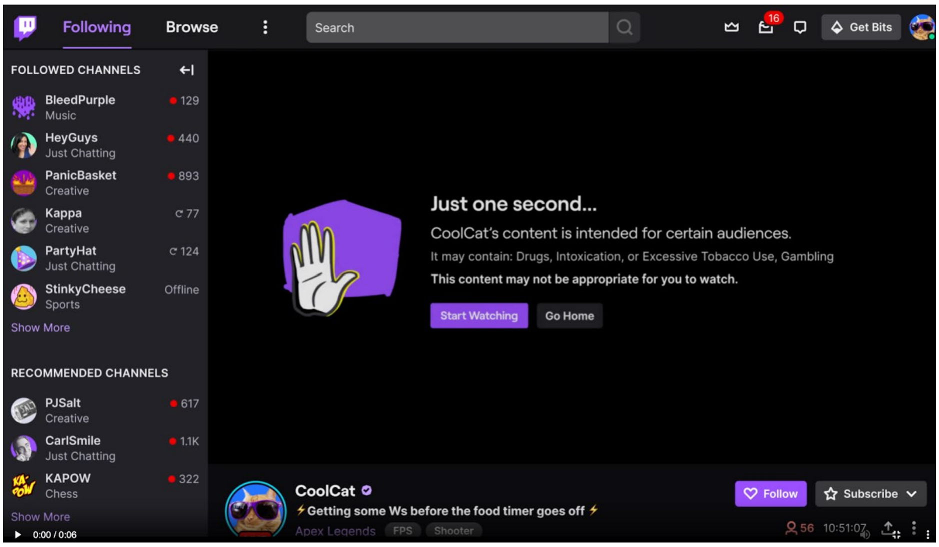


*Source: Twitch website post-June 2023 change*

2.5 Although content creators were expected to use the CCLs immediately after their implementation, Twitch granted its users a 30-day grace period to become accustomed to the new labelling system.[13] During this time, warnings issued following a mislabelling event only led to the correct CCL being applied to the stream that was mislabelled. However, after 20 July 2023, repeated warnings issued to a content creator that mislabelled their streams would accrue and led to the relevant CCL(s) being applied on the creator's channel for a set duration of time.[14] This meant that the relevant CCL could be applied by the platform to

---

[13] Before the CCL change content creators were provided with the following guidance on the platform: "You are expected to accurately label your content to the best of your ability. When choosing a category or tag, please choose whichever best describes your content. Deliberate or extensive misuse of titles, tags, games/categories, or other metadata are prohibited". We do not have information on any enforcement which may or may not have taken place before the CCL change.

[14] The relevant label may be applied on the content creator's channel for a period of days or weeks, depending on the number of prior warnings.

multiple streams, and not just the initially mislabelled stream. The content creator would be unable to remove the CCL applied on their channel during this period.

# 3. Twitch API data

3.1 In this section we outline how we collected data for our research using Twitch's publicly available API, which we supplemented with data from the IGDB API, Steam's Web API, and Pan-European Game Information (PEGI).

## Data collection

### Twitch API data

3.2 Twitch provides public access to the volume of mature and non-mature content viewed and streamed, including details of how content creators label their content through its API. We created a data pipeline in Python[15] to automatically collect and store Twitch stream data between March 2023 and December 2023 from the 'Get Streams' API endpoint.[16] The streaming data collected are the 100 most popular streams by views on the platform and therefore do not include all content streamed.[17] We collected metadata on 15,898,528 streams in total, which was reduced to 3,258,952 streams following the removal of duplicate streams.[18] Following the CCL change in June 2023, we collected additional data from the 'Get Channel Information' API endpoint which provided information on the use of granular mature labels.[19]

**APIs and the Twitch API**

An application programming interface (API) is a computing interface that allows interaction between different software. While an API can be used for different purposes, one of them is to share data between different applications.[20]

In this work, we used the official Twitch 'Get Streams' and 'Get Channel Information' API endpoints which are the specific extensions of the Twitch API that provide access to data related to live streams. The data cover information about the streams (live sessions) conducted online by users on the platform – but only metadata about content that is being live streamed at the time collection are available, i.e., the API does not provide historical information about the platform's content.

[15] Python is a programming language commonly used to access data from APIs.
[16] Endpoints represent different locations within the API that allow us to request specific data necessary for our analysis.
[17] We outline the limitations of the Twitch API, which allow us to only retrieve the top streams by views, in the annex.
[18] We outline the steps taken to remove duplicate streams in the annex.
[19] We detail the data pipeline constructed in the annex.
[20] Using an API as a data collection method has the advantage of gathering information on live content through a third-party software made available by the platform itself. It is possible to control what data are received via the API, which in this case meant that the Ofcom research team could minimise the data received and handle the processing from the outset. When planning which data to collect via the API, we focused on the data that were directly relevant to the purpose of the project and excluded anything that was not necessary to address the research questions set out earlier.

The collection of the data through the API, and making these available to third parties, is part of the terms and conditions of Twitch's Cookie Notice, which it presents to users.[21]

## 'Get Streams' API Metadata

3.3 From Twitch API documentation, the metadata selected through the 'Get Streams' API endpoint were the following, (data type in ()):

- **user_id** - (string) – The ID of the user who is streaming.
- **id** - (string) - An ID that identifies the stream.
- **title** - (string) - The video's title.
- **game_name** - (string) – Name of the category or game being played.
- **viewer_count** - (int) – Number of viewers watching the stream at the time of the query.
- **started_at** - (string) – UTC timestamp.
- **language** - (string) - Language used by the content creator
- **tags** - (string) - The tags applied to the stream by the content creator
- **is_mature** - (boolean) – Whether the Twitch channel is set to mature audiences or not.

3.4 The metadata 'title' and 'game_name' are both written by the content creator before going live and the 'game_name' refers to the name of the channel (and may not actually relate to a game). Although Twitch was initially focussed on streaming videogames, it has diversified and includes streams on a variety of topics. For example, the channel 'Just Chatting', where content creators can go live without any specific topic selected, is the most popular channel[22] streamed but appears in the API metadata as a 'game_name'.[23]

3.5 The metadata 'tags' provide more information about what topics are covered in a live stream as this variable is customisable by the content creator to describe what they will be talking about and, therefore, can help viewers to find streams in the platform. Finally, the metadata 'is_mature' provides information about whether the content creator has identified their stream as potentially containing mature content. Metadata can be changed during the live stream by the content creator.[24]

---

[21] The notice given to the user notes: "Twitch uses personal data collected on our services, such as page visits through cookies and other device identifiers, to generate personalized content, remember your preferences, analyse usage to improve our products, and measure the effectiveness of our campaigns to acquire new users. If you agree, we will also permit trusted third-party partners to obtain data from our services or store and access cookies on your device to deliver personalised ads, measure effectiveness, and create audience insights. Please see our Cookie Notice to learn more. By clicking "Accept Cookies", you consent to this activity. Click "Manage Cookies" to decline these cookies, make more detailed choices, or learn more." Further information about cookies and users' data collection is available on Twitch's website.

[22] This finding has been identified in related academic literature and confirmed in our analysis. For example, see Ha Le, Junming Wu, L. Yu, Melissa Lynn, A study on Channel Popularity in Twitch, Computer Science, 10 November 2021, p. 8.

[23] "game_name" is a legacy from when Twitch was exclusively used for streaming games.

[24] To ensure compliance with the UK General Data Protection Regulation, we deleted personal metadata – e.g., stream titles and user IDs – within the first 4 weeks of data collection. A detailed description of this process is provided in the data protection section of Annex 1.

3.6 Twitch's API does not provide the age or geographical location of viewers on the platform. We took the following steps to make the data collection as specific to UK users as possible:

i) We used the language metadata to verify that users are using the English language while streaming. However, this cannot guarantee that the content creator is located in the UK.

ii) As shown in Table 1, we scheduled the API call on specific days of the week and times of the day where UK users under 18 would most likely be visiting Twitch.[25] The research was focussed on the labelling of mature content that would be age-inappropriate for users under the age of 18.

**Table 1: Data Collection schedule**

| Monday-Thursday | Friday | Saturday | Sunday |
|---|---|---|---|
| 16:00 – 17:29 | 16:00 – 17:29 | 09:00 – 17:29 | 09:00 – 17:29 |
| 20:00 – 02:00 | 20:00 – 04:00 | 20:00 – 04:00 | 20:00 – 02:00 |

*Notes: Data Pipeline scheduled to collect data 'Get Streams' API and 'Get Channel Information' API. Pipeline is scheduled to collect more data on Friday, Saturday, and Sunday as we anticipate there to be more visitors to the platform on these days. Our descriptive and econometric results are aggregated at the weekly level which evens out the unbalances in the daily collection schedule. Source: Ofcom*

## 'Get Channel Information' API Metadata

3.7 Following the CCL change, we began collecting data on the CCLs from the 'Get Channel Information' API endpoint after Twitch granted us Beta API test access.[26] The Beta API gave Ofcom additional access to the 'content_classification_labels' from the 'Get Channel Information' endpoint that was not publicly available at the time. We extracted the user_id from the 'Get Streams' API and matched it to the broadcaster_id from the 'Get Channel Information' API to request the following data:

- **broadcaster_id** – (string) – An ID that uniquely identifies the broadcaster.
- **content_classification_labels** – (string) – The CCLs applied to the channel.

3.8 The metadata above include the six granular mature labels introduced in June 2023 and indicate whether a given label has been applied to a stream. An exact timeline of the intervention and our data collection process is shown in Figure 5.

[25] We created the schedule by combining assumptions regarding times of the day which would have the largest increase in visits to Twitch by children using data from Ipsos, Ipsos iris Online Audience Measurement Service, November 2022, day parts data, age: 15-17, UK.

[26] We outline the data losses we experienced due to delayed access to the Beta API and changes to the 'Get Channel Information' API endpoint in the annex.

**Figure 5: Timeline of intervention and API access**

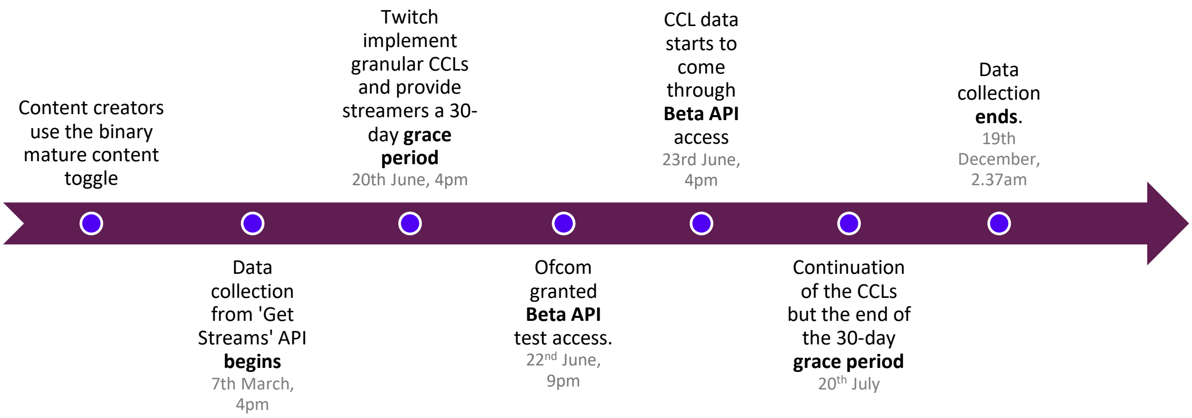


*Notes: The introduction of granular CCLs took place on 20 June, 4pm but we did not gain access to the API until the 22 June, and CCL data did not come through the pipeline until 23 June. Source: Ofcom*

## Stream Categorisation

3.9 We improved upon the raw data available from Twitch to better categorise streams in our data set. We undertook this exercise because in the raw data we only have detailed mature content labels and game and non-game labels for streams taking place after the CCL change. In the before period, these detailed content labels were not available to content creators, which made it challenging to directly compare outcomes before and after the change, but also made it challenging to perform analysis for the pre-CCL change period.

3.10 Therefore, we took steps to improve the classification of streams into game and non-game, and for games streams we improved the classification of whether the stream was mature or not mature.

3.11 To achieve stream categorisation, we used the raw data from the Twitch 'Get Streams' API endpoint and also three gaming information data sources. We received PEGI data supplied to us by the Games Rating Authority (GRA). PEGI is a European video game content rating system used in the UK market. We retrieved ESRB gaming information data from the IGDB API, which contains information on the ESRB rating of video games in the US market. We retrieved Steam data from the Steam Web API. Steam is an online platform for users to purchase, play, discuss and create games. Steam provides only the names of games available on its marketplace and our use of this data was limited to identifying games with an unknown rating.

3.12 To categorise streams into gaming and non-gaming streams, we implement the following process:

i) First, we use Twitch's automatic labelling of streams featuring ESRB rated mature games after the CCL change, to identify streams that featured games, and if so whether or not they featured mature games before the intervention.[27]

[27] For streams containing the 'Mature-Rated Game' after the CCL change, we used the 'game_name' field from the Twitch API to match against streams containing the same 'game_name' before the CCL change.

ii) Similarly, we used PEGI data to categorise residual streams which do not receive an automatic ESRB label by Twitch.[28] Streams featuring a PEGI 18 rated game are categorised as mature game streams i.e., they have both a mature CCL and are classified as a mature game.[29]

iii) We then use ESRB data from IGDB and Steam to further categorise remaining streams that could not be identified by Twitch or PEGI. We also identify streams containing games with an unknown rating. [30] Finally, we define a stream as 'Unknown' if we were unable to categorise it using the steps above or if the 'game_name' field was empty.

3.13 Figure 6 shows the results of our categorisation of Twitch streams into gaming and non-gaming streams for both the periods before and after the CCL change. More than 94% of the streams could be identified as gaming or non-gaming streams in both periods and the proportion of each stream category showed little change.

**Figure 6: Categorisation of streams in the pre- and post-CCL change period**

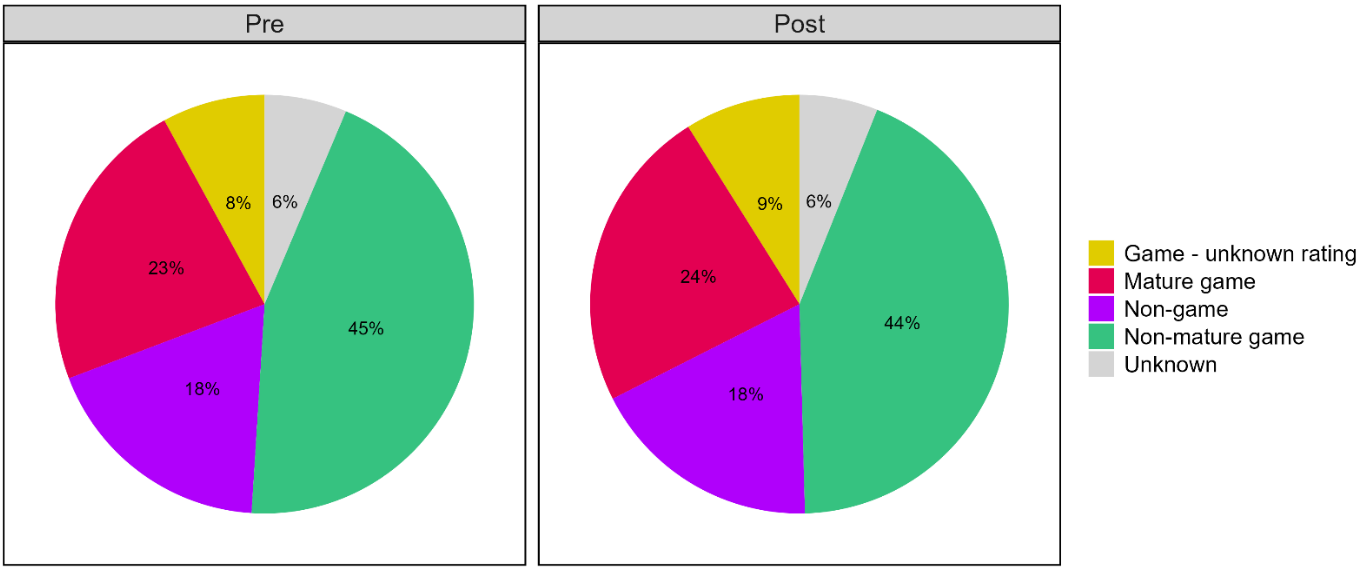


*Source: Twitch 'Get Streams' API data March and December 2023, PEGI ratings dataset, IGDB data, Steam API data.*

3.14 In section 4 below, we analyse data before the CCL change, using both the raw data collected from the Twitch API, and the improved data set following the stream categorisation discussed immediately above. In particular, section 4 contains an analysis of how accurately games were labelled in the pre-CCL change period, focusing on how well the labelling of mature content by creators correlates with gaming age ratings data. For this analysis we needed to exclude non-games and use additional data beyond the raw data collected from the Twitch API.

[28] We use the 'game_name' field to match against the PEGI data.
[29] We received PEGI data from the GRA on 24.05.2023. This meant that streams containing games released after this date were not categorised by PEGI data.
[30] The ESRB and PEGI rating system is voluntary for game publishers, resulting in some streams containing games with an unknown rating. Prior to release, publishers fill in a questionnaire disclosing any relevant content (violent, sex, bad language etc.) that may be considered inappropriate for all ages.

# 4. Descriptive analysis of content: pre-CCL change

4.1 In this chapter we cover the findings from the period before the CCL change. The intention of this stage of research was to understand how the mature label was being used before the CCL change and set the baseline for the next stage of research. Our first research question (RQ1) concerns the proportion of mature-labelled content on Twitch before Twitch implemented the change to the labelling system.

## RQ1: pre-CCL change, what was the proportion of mature-labelled content on Twitch?

4.2 38% of the Twitch streams we analysed were labelled as mature. This figure only includes streams where the content creator filled in the “game name” label on the stream. We selected data from eight consecutive weeks between April 2023 – June 2023 and only included data collected between 4pm to 6pm and 8pm to 11:59pm each day to ensure a consistent approach to data collection. This was followed by the removal of streams where the ‘game_name’ field was empty, leading to a total of 382,494 streams being analysed. Of these streams, just under four in ten (38%) of the total volume of streams were labelled as mature. This number demonstrates that while most of the streams were not labelled as mature, there was an awareness of the mature label that was significant enough for content creators to use it. These findings reflect the streams that were captured by the API and therefore the most popular streams on Twitch. As it is not possible to conduct research via the Twitch API on smaller viewership streams, we cannot say whether the findings would vary if these streams were included.

## RQ2: pre-CCL change, did the time of day and day of the week have an impact on the percentage of content that is labelled as mature?

4.3 We were interested in whether the time of day or day of the week had an impact on whether content was labelled as mature. We considered that there might be more mature content later in the evenings.

4.4 As shown in Table 1 above, we collected data for every day of the week. We analysed the share of data by label (whether mature or non-mature) for each day of the week. The percentage of the volume of streams labelled as mature and non-mature changes very little across time intervals and across different days of the week. The percentage of streams labelled as mature ranged from 36.54% to 38.55%.

4.5 To explain the similarity across the data intervals, we considered several different factors. One factor is that content creators will be in different time zones so we could expect differences to be less pronounced than if all content creators were based in the UK. While the limitation of getting the 100 most popular streams by request using the API possibly could have impacted how stable the proportion of streams labelled as mature was, another factor to take into consideration is that the streams on this platform are usually long and can be many hours and, in certain cases, a couple of days. Even after removing duplicates from our data, this recurring length can also explain the similarity across time periods.

4.6 We also analysed the average number of viewers for each day of the week at the specified time intervals. This analysis aims to observe the movement of the viewers connecting to the streams at different days and times. For this purpose, we analysed the total number of viewers by day of the week and by time interval and calculated the average number for each time period.

4.7 Figure 7 details the change of viewers across the different days of the week at different times for content being labelled as mature and non-mature. While we can observe that the average number of viewers slightly changes during the week, more variation appears during the weekend with a significant drop of viewers during the afternoon that then picks back up in the evening. We can note that content labelled as mature follows generally similar trends in comparison to those labelled as non-mature.

**Figure 7: Average viewership by day of the week at specific time intervals**

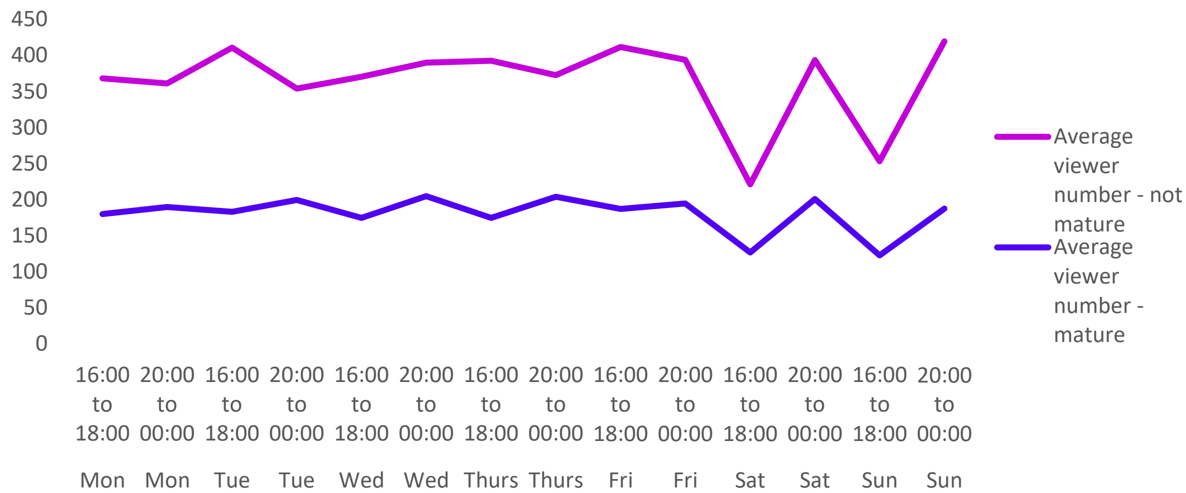


*Source: Twitch 'Get Streams' API data May and June 2023*

## RQ3: pre-CCL change, did labelling of mature content by content creators correlate with the PEGI rating of games?

4.8 Finally, we considered research questions about the most popular games being played in streams labelled as mature and not labelled as mature, and assessed whether there was a relationship between the age rating of the games being played and if a stream was labelled as mature. If mature labelling was accurate, then one would expect that streams of games with a PEGI 18 rating (the highest age rating) would be labelled as mature. As shown in Figure 8 the most popular PEGI rating for games not labelled as mature was 12, whereas for games rating as mature it was 18.

4.9 We found that 26% of gaming streams labelled as mature were playing a game with a PEGI age 18 rating, and 19% of streams labelled as non-mature were playing games with a PEGI age 18 rating. This suggests that, prior to the CCL change, there was some relationship between PEGI rating and whether a game was labelled as mature, but there were still a significant percentage of PEGI 18 games not labelled as mature. We also found that a significant number of streams featuring games with PEGI 3, 7 and 12 age ratings were labelled as mature. This could be due to the nature of the commentary provided by the content creator which may have included mature themes.

**Figure 8: PEGI age ratings across streams labelled as non-mature and mature (excl. non-gaming categories and streams with game name 'None')**

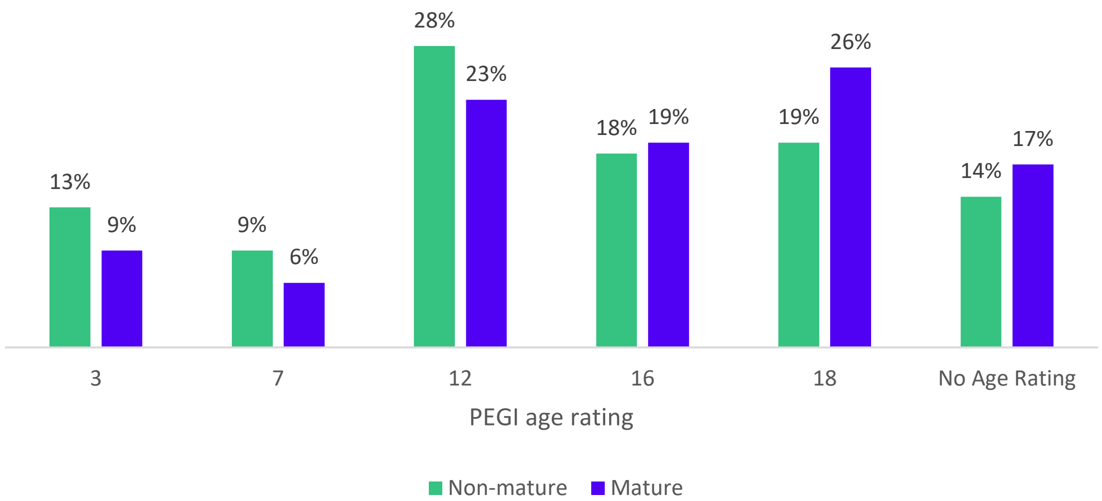


*Source: Stream data collected from Twitch 'Get Streams' API from March 2023 to June 2023 mature streams base = 197,811, non-mature streams base = 299,217), PEGI ratings dataset, IGDB data*

# 5. Evaluation of the impact of the CCL change

5.1 In this section of the paper, our focus lies in understanding the effect the CCL change had on content labelling accuracy, creators' adherence to Twitch's Community Guidelines, as well as understanding any other behaviour change by content creators and viewers. We set out to answer the following research questions:

- Did the CCL change improve content labelling accuracy and so increase adherence by creators to the platform's Community Guidelines? (RQ4)
- Did the CCL change affect creator behaviour in any other ways? (RQ5)
- Did the CCL change impact absolute audience figures on the platform? (RQ6)

5.2 We answer these questions by deploying three types of analysis.

5.3 First, a classifier analysis, which used machine learning to further improve the dataset, in particular by generating predictions for which detailed CCLs were likely to have applied in the pre-CCL change period. These predictions were used to in effect back-cast data from the post-CCL change period to the pre-CCL change period. This was a crucial piece of analysis because it enabled us to undertake the 'accuracy analysis' discussed next. Without the classifier work our analysis of the accuracy of labelling would have been limited to just game streams because we would not have had non-game CCLs identified in our dataset for the pre-CCL change period.

5.4 Second, we performed an accuracy analysis which assessed the accuracy of labelling, and how this was affected following the CCL change. This analysis is important because it tells us whether the CCL change led to an improvement in content creators' adherence to the platform's Community Guidelines. More generally, it provides insights into whether more granular labelling, and the platform enforcement system surrounding it, was effective in improving the information presented to viewers as regards mature content in streams.

5.5 Third, we use econometric analysis to examine whether other forms of content creator and viewer behaviour changed following the CCL change. The analysis assesses whether views of mature content increased or decreased following the CCL change, and whether content creators produced more or less mature streams. These insights can be valuable for future policy work by helping to build the evidence base on how users react to platform changes.

5.6 The remainder of this section is structured as follows:

- Classifier analysis: we use a classifier analysis as a key input to the analysis of the three research questions set out above. We explain why we used a classifier analysis, what we did, and the results and implications.
- Accuracy analysis: we describe how we assessed the accuracy of content labelling on Twitch's platform, both before and after CCL change. We present the findings of the analysis and implications for RQ4.
- Econometric analysis: we explain how we used econometric analysis to answer RQ4, RQ5, and RQ6.

# Classifier analysis

## Motivation

5.7 In the raw data from Twitch, we only have detailed mature content labels for streams taking place after the implementation of the new CCL regime (as these detailed content labels were not available to content creators before the CCL change), which made it challenging for us to directly compare outcomes before and after the change – in particular, to estimate how the CCL change impacted on whether mature streams were labelled as such and the impact of this labelling on viewership. We were able to construct reliable estimates of mature game streams before the CCL change using external data from PEGI and IGDB on game ratings and Steam on game listings, but for the non-game CCL categories we had no external means of determining which labels were likely to apply.

5.8 Therefore, we sought to estimate in the pre-period whether content of streams (for which we did not have data) was mature, as well as other characteristics. We achieved this by employing a machine learning algorithm which looked at classification in the post-implementation period to predict classification in the pre-implementation period. Specifically, we trained a text classifier on a sub-sample of 2 million observations from the post-treatment observations to predict which CCLs apply to a given stream based on a subset of features that we have for all streams.

5.9 Using this algorithm, we were able to correctly classify around 90% of all streams in the post CCL change period, which meant that we could use the classifier to estimate what labels would have applied in the pre-implementation period, if the CCLs had been available and applied accurately, with a reasonable degree of confidence. Additionally, we found a 99% level of agreement between our initial labelling of game streams using external sources and our labelling of game streams using the classifier's predictions.

5.10 This is a significant result because it unlocked analysis that would not have been possible without the classifier work: comparison of the accuracy of labelling on Twitch both before and after the intervention and econometric analysis for non-game mature content streams. Without the classifier work this would only have been possible for game content streams.

5.11 We used a DistilBERT[31] model to classify streams into game, non-game, and for non-game the particular CCL that applied, if relevant.

5.12 We found that the model performed well in its ability to predict CCLs. We evaluated the model by calculating a number of variables which are typically used to evaluate the performance of automated classifiers when using machine learning. In particular, we assessed the precision, recall, and F1 scores for each of the six CCLs. The details of this work are contained in the annex.

5.13 When the model predicts that a stream should have a particular CCL, we found that the stream usually had that CCL.[32] We also found that the classifier model overall is extremely

[31] BERT models are a class of large language model trained to predict randomly removed words from sentences and next sentences in a sequence. DistilBERT is a version of BERT that has been shrunk to make it quicker to fine-tune for more specialised tasks.
[32] Precision scores are all above 0.86.

likely to be correct when it predicts that a stream has a CCL and only marginally less likely to correctly identify all CCLs in the dataset.[33]

5.14 We performed an additional check of the classifier model to assess its performance in the pre-CCL change period. For this check we generated predictions for all the streams pre-CCL change and compared our predictions of mature games against the data observations we had derived using game ratings data described earlier in this paper. The motivation for doing this check is a concern that the type and/or makeup of streams is different in the pre- compared to the post-CCL change period. We found that the classifier performed well on this check, suggesting that we can be confident that the classifier exhibits strong out-of-sample performance for game CCLs.

5.15 We concluded from this evaluation that the classifier model generates reliable predictions for stream classification, such that we can rely on these for our analysis of RQ4, 5, and 6.

# Accuracy analysis

## Approach

5.16 In this section we compare the labelling of content on Twitch against our predictions of what content was mature. We also examine the flip side: we compare the labelling of content on Twitch against our predictions of which content was not mature.[34] This analysis is done for the pre-CCL change period. We treat the post-CCL change period as the 'ground truth' i.e., we assume all content is accurately labelled in this period. This is because we do not have any better information on content classification than what is available in the post-CCL period.

## Accuracy analysis findings

5.17 In Table 2, we use the outputs of our classifier analysis to examine the accuracy of mature content labelling before the CCL change. We assess two aspects of accuracy. First, we look at the proportion of predicted mature streams that were accurately labelled by content creators as mature. Second, we examine the proportion of predicted non-mature content that was incorrectly (or 'inaccurately') labelled as mature. We discuss each in turn below.

5.18 On the first aspect of accuracy, we calculate the proportion of our predicted mature streams that were labelled by streamers with the mature flag available before 20 June (i.e., the 'is_mature' flag). Given our high confidence in the accuracy of the predictions of the classifier model, as discussed in the section above, we characterise the predictions of the classifier as 'accurate' for the purposes of the following analysis. We compare the labelling by content creators to this 'accurate' labelling. This allows us to identify how accurate content creators' labels of streams were in the pre-CCL change period.

5.19 On the second aspect of accuracy, we re-estimated the analysis described above, this time focusing on streams that were 'inaccurately' labelled. Specifically, we examine situations pre-CCL change where content creators labelled their streams as mature but where our classifier model suggested such streams did not actually contain any type of mature content.

[33] The average recall score of 0.86 is slightly below the average precision score of 0.92.
[34] Put another way, our inaccuracy analysis refers to a situation where content creators label streams as mature but where our analysis suggests such streams did not in fact contain mature content.

5.20 Our analysis in panel (a) in the table below indicates that labelling accuracy was fairly low before the CCL change. Notably, we observe the worst performance among mature gaming streams, with fewer than 50% of mature game streams being accurately labelled by content creators. This is followed by gambling streams, with just 65% of these streams being accurately labelled as mature.

5.21 We also found that content creators inaccurately labelled non-mature content as mature before the CCL change, and that this behaviour was prominent, for both game and non-game streams. As can be seen in panel (b) in Table 2, the proportion of streams classified by our analysis as non-mature but inaccurately labelled as mature ranged from 37% to 40% in before the CCL change.[35, 36]

**Table 2: Accuracy analysis**

| (a) Predicted mature stream type | Streams labelled as mature (%) | (b) Predicted non-mature stream type | Streams labelled as mature (%) |
|---|---|---|---|
| **Mature-Rated Game** | | **No Mature-Rated Game** | |
| *Pre-CCL change* | 48.63% | *Pre-CCL change* | 37.16% |
| **Gambling** | | **No Gambling** | |
| *Pre-CCL change* | 64.42% | *Pre-CCL change* | 39.58% |
| **Significant Profanity and Vulgarity** | | **No Significant Profanity or Vulgarity** | |
| *Pre-CCL change* | 84.96% | *Pre-CCL change* | 38.04% |
| **Sexual Themes** | | **No Sexual Themes** | |
| *Pre-CCL change* | 84.73% | *Pre-CCL change* | 38.88% |
| **Drugs, Intoxication, or Excessive tobacco Use** | | **No Drugs, Intoxication, or Excessive tobacco Use** | |
| *Pre-CCL change* | 82.29% | *Pre-CCL change* | 39.31% |
| **Violent and Graphic Depictions** | | **No Violent and Graphic Depictions** | |
| *Pre-CCL change* | 85.62% | *Pre-CCL change* | 39.40% |

*Source: Analysis of Twitch 'Get Streams' API and 'Get Channel Information' API data March and December 2023*

[35] We treat the post-CCL change period as the 'ground truth' i.e., we assume all content is accurately labelled in this period. This is because we do not have any better information on content classification than what is available in the post-CCL period.

[36] Our analysis in the pre-treatment period may be prone to misclassification bias. Since streams may contain multiple types of mature content, our methodology – which evaluates each specific type of mature content separately – may inadvertently include cases where streams feature additional mature content beyond the specific type analysed. For example, a stream not containing gambling but labelled with the binary "is_mature" flag may be mistakenly defined as inaccurately labelled despite featuring other mature content, such as significant profanity and vulgarity. Thus, to test the validity of our findings, we re-estimated the analysis in Panel (b) of Table 2 by focusing specifically on streams that contained no mature content at all, rather than streams lacking only a specific type of mature content. Similar to our baseline analysis, we find that the proportion of streams inaccurately labelled as mature was 34% in the pre-treatment period.

# Econometric analysis

## Approach

5.22 In this section, we assess the probability of mature streams being correctly labelled as mature, and whether this differs pre- and post-CCL change. We assess this econometrically to establish causality between the CCL change and changes in accuracy. As before, we treat classifier predictions of CCLs as the 'ground truth' over the period of analysis. This makes it possible to do a like-for-like comparison between the pre- and the post-CCL change periods but comes at the cost of introducing some measurement error in the post period.[37] As a result, the findings below represent a lower bound on the estimated effect of the CCL change on labelling accuracy i.e., the consequence of this measurement error is that the findings would likely understate the CCL change effect.

5.23 As context, the analysis of labelling accuracy in the preceding section is not causal, in that it simply presents statistics on accuracy (and inaccuracy) before the CCL change. The findings in that section are insightful but fall short of robustly tying the CCL change to the observed changes in accuracy. The analysis in this section achieves this, by deploying more sophisticated statistical techniques to isolate the effect of the CCL change on labelling accuracy.

5.24 The analysis in this section relates to RQ4 (Did the CCL improve content labelling accuracy and so increase adherence to the platform's Community Guidelines?). We assess both gaming and non-gaming streams. We do this because we would expect accuracy to increase for gaming streams as the platform in question began automatically adding content labels for these streams following the CCL change. Therefore, we wanted to assess how non-gaming streams specifically performed from an accuracy perspective, which also provides an insight into how creators' behaviour changed from a labelling viewpoint.

5.25 For gaming streams, we assess Mature-Rated Games, and for non-gaming streams we look at the five granular CCLs. The results show that content label accuracy improved materially following the CCL change. We also find that content label inaccuracy reduced materially following the CCL change. If these were part of the platform's objectives in making this change to its labelling system, then it appears to have achieved this successfully.

5.26 In this section we also seek to answer RQ5 (Did the CCL change affect creator behaviour?): we analyse the weekly volume of mature content produced by content creators before and after the CCL change to assess whether there is any observable change in content creator behaviour. The results show that there was little change in creator behaviour in terms of the volume of mature content produced. We don't know whether and how Twitch expected to change creator behaviour in terms of volume of mature content streamed, and the changes may have been more focused on providing viewers with greater, and more accurate, information on the content of streams. To the extent that Twitch did not wish to impact on the volume of mature content produced by content creators from the CCL change, then it appears to have achieved this successfully.

[37] While we are confident over the overall performance of our classifier, the predictions from our model may not always be 100% accurate. This is because factors such as noisy data, incomplete or incorrect information on CCL use, and/or changing relationships between input features (e.g., emojis, tags) and mature content types may lead to inaccurate predictions of CCLs, and thus, measurement error in our dependent variable. See Table 3 for further details on the dependent and independent variables.

5.27 Finally, we consider RQ6 (Did the CCL change impact absolute audience figures on the platform?): we analyse the weekly count of mature views before and after the CCL change to assess whether there is any observable change in viewer behaviour. The results show that there was no significant impact on the volume of mature content viewed, at least after a short initial period following the CCL change.[38] We do not know whether and how Twitch expected to change viewer behaviour in terms of volume of mature content viewed. It is plausible that providing more detailed information about the content of mature streams and more accurate labelling of these streams as being mature might lead some users of Twitch not to view this material which they might otherwise have (inadvertently) viewed. If this was one of Twitch's objectives from the CCL change, then it appears to have been unsuccessful in achieving this.

## Probability of mature streams being labelled as mature

5.28 In this sub-section we assess the impact of the CCL change on content creators' adherence to Twitch's Community Guidelines (RQ4) by examining the accuracy of mature content labelling and the overall rate of mature content labelling on the platform.

5.29 We rely on a Regression Discontinuity Design (RDD) model, an econometric technique that isolates causal effects by comparing individuals sufficiently close to a cut-off point.[39] An econometric approach allows us to control for factors, other than the CCL change, which could have impacted on the propensity of content creators to change the accuracy of their labelling, or content creators to change the volume of their content, or viewers to change their viewing habits. The RDD approach limits the time period considered so as to reduce the potential impact of factors for which we are unable to control.

5.30 In our particular set-up, streams which took place around the CCL change – i.e., 20 June 2023 – are used to estimate causal treatment effects, with observations just before and after the implementation date falling in the 'control' and 'treatment' group, respectively.

5.31 To the extent that external factors, which could have impacted user behaviour, remained constant and stream characteristics were similar – and therefore comparable – during this period, an RDD model would allow us to obtain reliable estimates of the causal impact of the CCL change on content creator behaviour.

5.32 We are confident that these identifying conditions are met in the context of our study. Indeed, Twitch confirmed there were no other interventions that could have influenced user behaviour during the period surrounding the CCL change. Moreover, given that the changes applied to all users on the platform, and content creators remained unaware of the new labelling process until its introduction, we believe content creators just before and after the implementation date should be similar in their observable and unobservable characteristics.[40]

5.33 In addition, to rule out other potential external factors that could influence our analysis, we tested whether the appearance of new games significantly affected audience volumes or the types of games that content creators engaged with around the CCL change. Our analysis

[38] We also examined whether there was a change in the proportion of mature content viewed following the CCL change. Our results were qualitatively the same as for absolute content views.
[39] Details about our econometric model are in the annex.
[40] We test this assumption more formally in the annex.

suggested that this was not the case, and that in any event, new game streams represented a low proportion of the overall sample of game streams.[41]

### Analysis of user behaviour before and after the CCL change

5.34 We also estimate a series of panel regression models to assess potential changes in content creation behaviour (RQ5) and viewership of mature content (RQ6) following the CCL change.[42] Specifically, we rely on a content creator 'fixed effects' analysis to control for both observed and unobserved content creator-level characteristics that remain constant over time, and which may bias our results (e.g., content creator's gender, streaming styles, etc.).[43] We also control for weekly 'fixed effects' to account for any time-related effects that may impact all content creators in a similar way (e.g., game releases, e-sports tournaments, etc.). Moreover, and similar to our RDD setup above, we focus on the immediate period surrounding the CCL change in order to account for potential external factors that may bias our results.[44]

## Probability of mature content labelling

### Graphical analysis of mature content labelling

5.35 To illustrate the effect of the CCL change on mature content labelling and test the validity of our RDD model,[45] Figure 9 plots the proportion of streams labelled accurately as mature over time. Specifically, the horizontal axis shows the difference in the number of days between a given streaming session and the date of the CCL change on the platform – i.e., the 'cut-off' date.[46] Negative numbers on this axis represent the pre-treatment period, and so, capture streams that occurred before the CCL change.[47] Moreover, as mentioned in paragraph 5.24, we treat classifier predictions of CCLs as the 'ground truth'.

5.36 Across the panels in the diagram shown below – except those for mature gaming streams – we exclude from the analysis the 30-day grace period provided by Twitch to all its users. Unlike all other CCLs, the 'Mature-Rated Game' label is applied automatically by Twitch rather than the content creator. Thus, we expect the impact of this CCL on mature content labelling to be immediately observable. Indeed, this assumption is confirmed in Figure 9, with the effect on mature content labelling among mature gaming streams increasing sharply immediately after the CCL change. Moreover, while we recognise that the CCL

---

[41] We find that on any given week, streams of new games account for no more than 5% of streams of all games.

[42] The graphical analysis discussed in the following section revealed no clear jumps in either content creation or viewer behaviour following the CCL change. Since RDD models strictly depend on discontinuous jumps around the cut-off date to estimate causal effects, we are unable to rely on this particular econometric technique to assess RQ5 and RQ6.

[43] Details about our fixed effects model are in the annex.

[44] We use a 6, 8, and 12-week bandwidth to test the robustness of our results.

[45] For RDD to be a valid estimation strategy we require a discontinuous jump in mature content labelling rates following the CCL change.

[46] Given that we make API calls every 30 minutes within the data collection schedule shown in Table 1, each plotted point corresponds to the proportion of streams labelled as mature within these time windows on a given day.

[47] Meanwhile, positive numbers capture streams that took place in the post-treatment period, with the vertical axis on 0 denoting the exact date that the intervention took place, i.e., 20 June.

change may have influenced content creator behaviour shortly after their introduction, we anticipate that the most significant effects materialised outside this trial period as content creators began facing penalties for not labelling their mature content accordingly. Therefore, the cut-off point used to estimate the causal effect of non-game CCLs on mature content labelling is defined by the exact date that the grace period ended, i.e., 20 July 2023.

### Figure 9: Mature content labelling across stream types

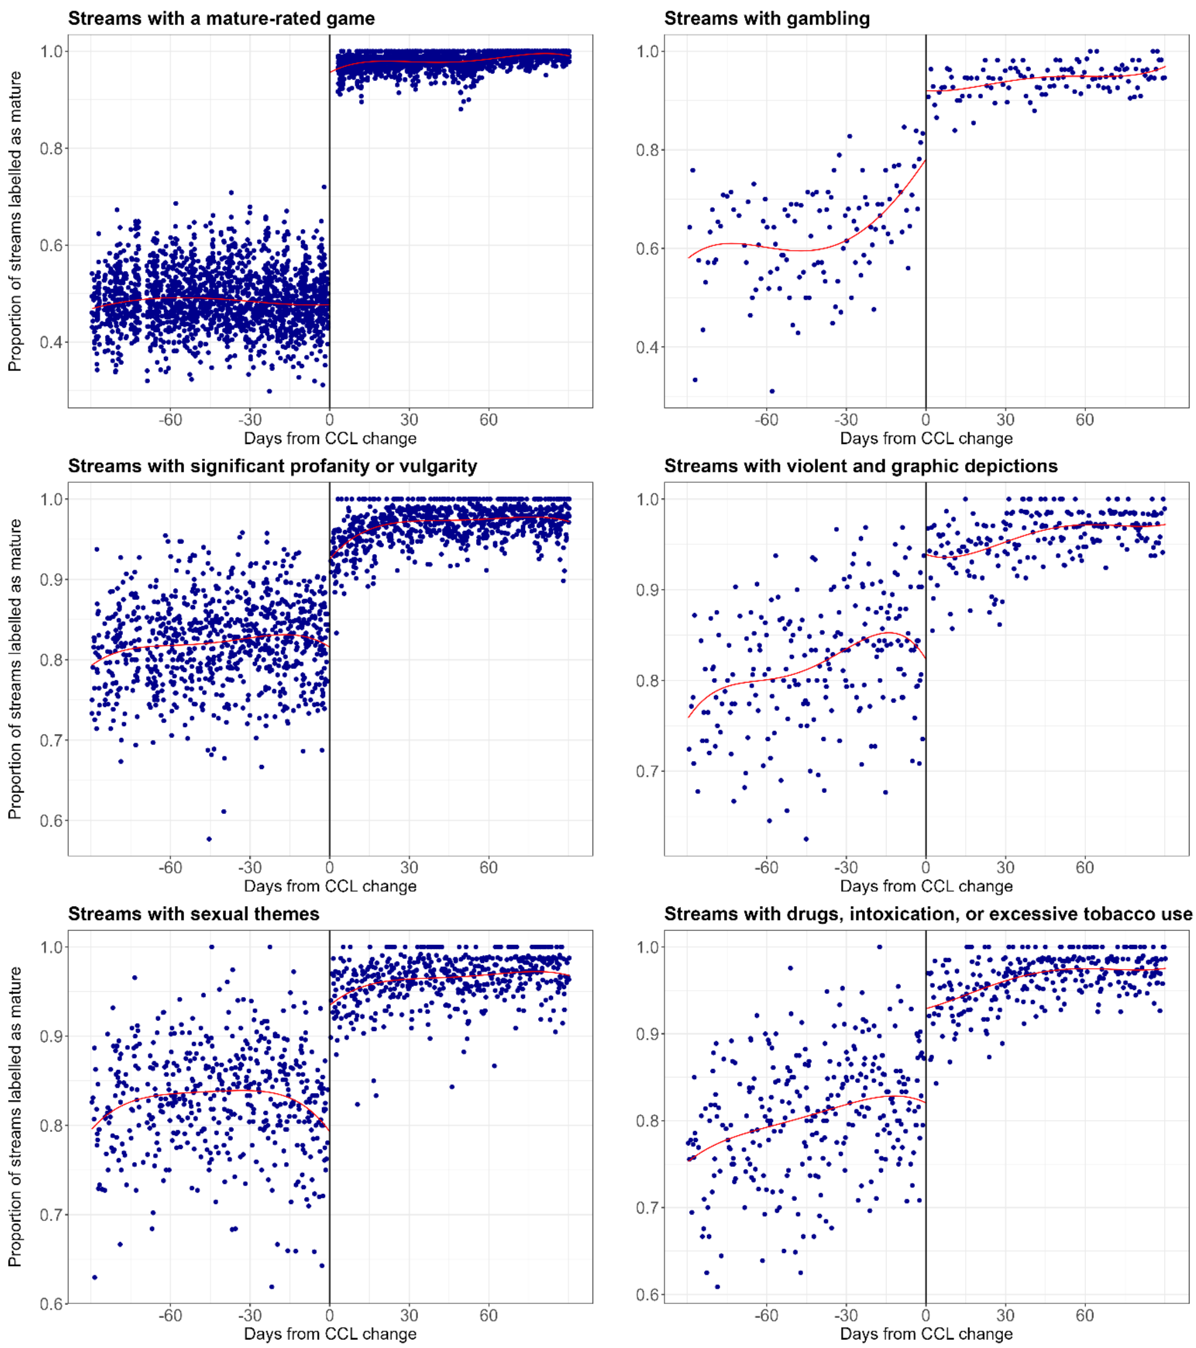


*Notes: The results for streams with a mature-rated game are based on a 20 June cut-off, while the analysis of mature non-gaming streams relies on a 20 July cut-off. Each plotted point represents the proportion of streams labelled as mature within a given 30-minute interval on a given day. Source: Ofcom analysis of Twitch 'Get Streams' and 'Get Channel Information' API data March and September*

5.37 Our graphical analysis in Figure 9 shows a jump in the proportion of streams accurately labelled as mature immediately after the CCL change, with mature gaming streams showing the sharpest increase in content labelling. This is to be expected given that the platform is automatically labelling games, without requiring the content creator to do anything else for streams featuring games. Moreover, while we find an increasing trend in the labelling rate of gambling streams even before the CCL change, Figure 9 reveals a sharp jump following the 30-day grace period, highlighting the intervention's role in further increasing labelling of such mature content.

5.38 The reason that the non-gaming categories fail to reach 100% by the end of the 60-day window is that we are treating the classifier predictions as ground truth for this part of the analysis. Since the classifier is not 100% accurate, there are some streams where our classifications and the original CCL data from Twitch do not perfectly agree. Consequently, the increase in labelling accuracy after the CCL change is likely higher than Figure 9 shows.

## Econometric analysis of mature content labelling

5.39 The graphical analysis above suggests that accuracy of mature content labelling increased following the CCL change on the platform. In this subsection, we turn to a more comprehensive analysis of the causal effects of the intervention by estimating our RDD models.

5.40 Consistent with the results in Figure 9, our regression analysis in Table 3 indicates that the CCL change increased labelling accuracy across gaming and non-gaming streams on the platform. Focusing on column 1, we find that the CCL change increased the probability of content labelling among mature gaming streams by 49 percentage points, and this effect is statistically significant at all conventional levels.

**Table 3: Mature content labelling by stream type – non-parametric RDD analysis**

| | (1) Mature-Rated Game | (2) Gambling | (3) Significant Profanity or Vulgarity | (4) Violent and Graphic Depictions | (5) Sexual Themes | (6) Drugs, Intoxication, or Excessive Tobacco use |
|---|---|---|---|---|---|---|
| **CCL change** | 0.49***<br>(0.009) | 0.22**<br>(0.100) | 0.12***<br>(0.015) | 0.12***<br>(0.041) | 0.16***<br>(0.023) | 0.13***<br>(0.035) |
| **Optimal BW** | 18.60 | 34.46 | 31.09 | 25.22 | 26.75 | 35.00 |
| **Observations** | 703,414 | 9,921 | 145,229 | 23,400 | 69,430 | 32,948 |

*Notes: Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01. We report bias-corrected estimates that account for the common misspecification bias found in RDD models and robust standard errors clustered at the content creator-level. The results for mature-rated game streams are based on a 20 June cut-off, while the analysis of mature non-gaming streams relies on a 20 July cut-off. Source: Ofcom analysis.*

5.41 Furthermore, while the impact on non-gaming streams lower than on games, the overall magnitude of these effects is substantial. For example, as shown in column 2, the probability of accurately labelling gambling streams increased by 22 percentage points following the 30-day grace period. Similarly, among other types of mature non-gaming streams, the likelihood of labelling mature content accurately increased by 12-16 percentage points, with all effects being significant at the 1% and 5% level.

5.42 However, it is important to note that because our econometric analysis treats the classifier predictions of CCLs as the ground truth, it is prone to the same risk of measurement error that would apply to the graphical analysis in Figure 9. As such, our estimated coefficients in Table 3 may understate the true impact of the CCL change on content labelling accuracy. This is because the effect of the measurement error would bias the estimates in the table towards zero; the fact that the estimates are positive and significant in spite of this bias provides strong evidence regarding the impact of CCLs on mature content labelling, and thus, on content creators' adherence to Twitch's Community Guidelines.

# Content creator behaviour

## Graphical analysis of content creator behaviour

5.43 We also explore the potential effect of the CCL change on content creator behaviour, focusing specifically on the weekly volume of mature content produced by top content creators following the intervention. We define the top content creators as the leading 1,000 content creators based on their weekly viewership. We chose to focus on the top 1,000 content creators as they account for over 64% of total viewership to the platform and would therefore provide an informative visual representation of popular content streamed.[48]

5.44 The top weekly content creators can be broken down into the following three groups:

i) Consecutive content creators: content creators in each week that have carried over from the previous week. As shown in Figure 10 they are the key content creators of interest as they account for the largest share of viewership to top content creators.[49]

ii) Recurring content creators: content creators who dropped out of the top content creators in an earlier week and re-joined in a later week. They account for the second largest share of viewership to content creators.[50]

iii) New content creators: content creators in each week who have joined the top content creators for the first time.[51]

[48] We arrive at this figure by calculating the share of views to the top 1,000 content creators in each week and taking the average from 20.03.2023 – 17.12.2023. We cannot identify content creators between 07.03.2023 – 12.03.2023 due to a change in the hashing system. Since consecutive streamers carry over from the previous week, our first observations of consecutive streamers begin on the 20.03.2023.

[49] Consecutive content creators consistently account for over 80% of total viewership to the top 1,000 content creators.

[50] Note that in Figure 10 we only begin to observe recurring content creators after new content creators have joined. This is because they are identified by finding the difference between content creators who are new to a given week but may have appeared in weeks prior, and content creators who have joined the top content creators for the first time.

[51] We have observations for new streamers from 20.03.2023 to 17.12.2023.

**Figure 10: Weekly share of viewership to the top 1,000 content creators for consecutive, recurring and new content creators.**

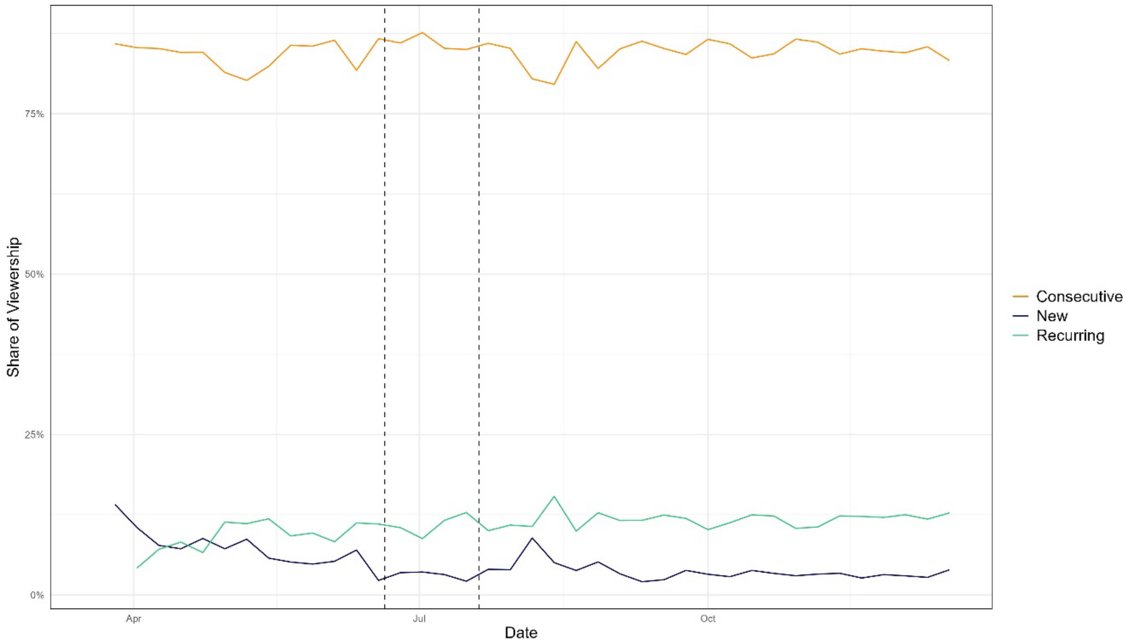


*Notes: The y-axis represents the share of top content creator viewership by consecutive, new and recurring content creators. The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

5.45 The chart above shows that there were no substantial changes in the weekly views of top weekly content creators following the CCL change. This tells us that the labelling change does not appear to have affected their viewership. This may be because their behaviour changed to adapt to the CCL change, or the CCL change did not directly affect their views. We explore each of these next.

5.46 Figure 11 plots the weekly volume of mature and non-mature streams generated by the top consecutive content creators across all stream categories across the data collection period. The vertical axis represents the number of mature and non-mature streams streamed in each week by the top consecutive content creators. [52] This analysis suggests that there was a drop off in the volume of both stream types after the CCL change. However, the volume of mature content shows a small upward trend following the end of the grace period.

[52] For our graphical analysis of content creation behaviour, this includes content streamed between 20.03.2023 – 17.12.2023

**Figure 11: Weekly number of mature and non-mature streams by top consecutive content creators for all stream categories**


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

5.47 Figure 12 and Figure 13 show the weekly volume of mature content after the CCL change when focusing on gaming and non-gaming categories separately. Following the CCL change, the weekly volume of mature content decreased sharply for gaming categories, while remaining relatively constant for non-gaming categories. Interestingly, however, both stream types showed an upward trend after the grace period, albeit the increase was relatively small.

5.48 Overall, our graphical analyses suggest there were no major changes in the weekly volume of mature content produced following the CCL change.

**Figure 12: Weekly number of gaming streams by top consecutive content creators**

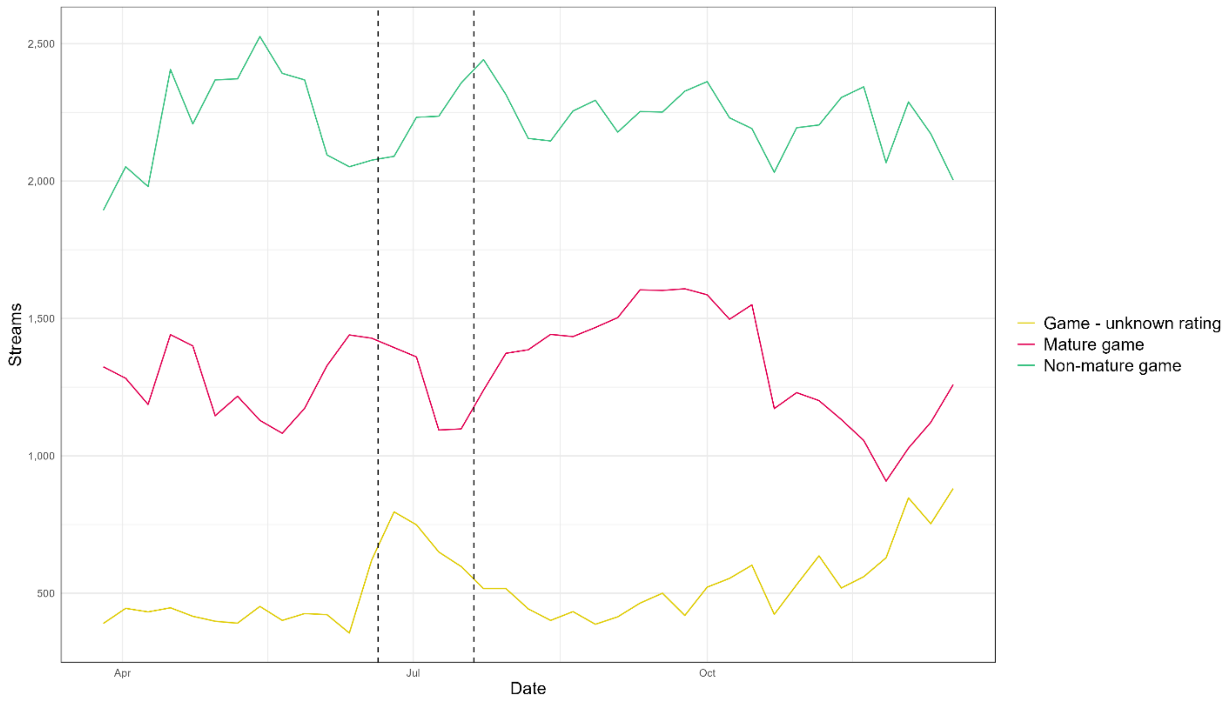


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

**Figure 13: Weekly number of mature and non-mature streams by top consecutive content creators for non-gaming streams**

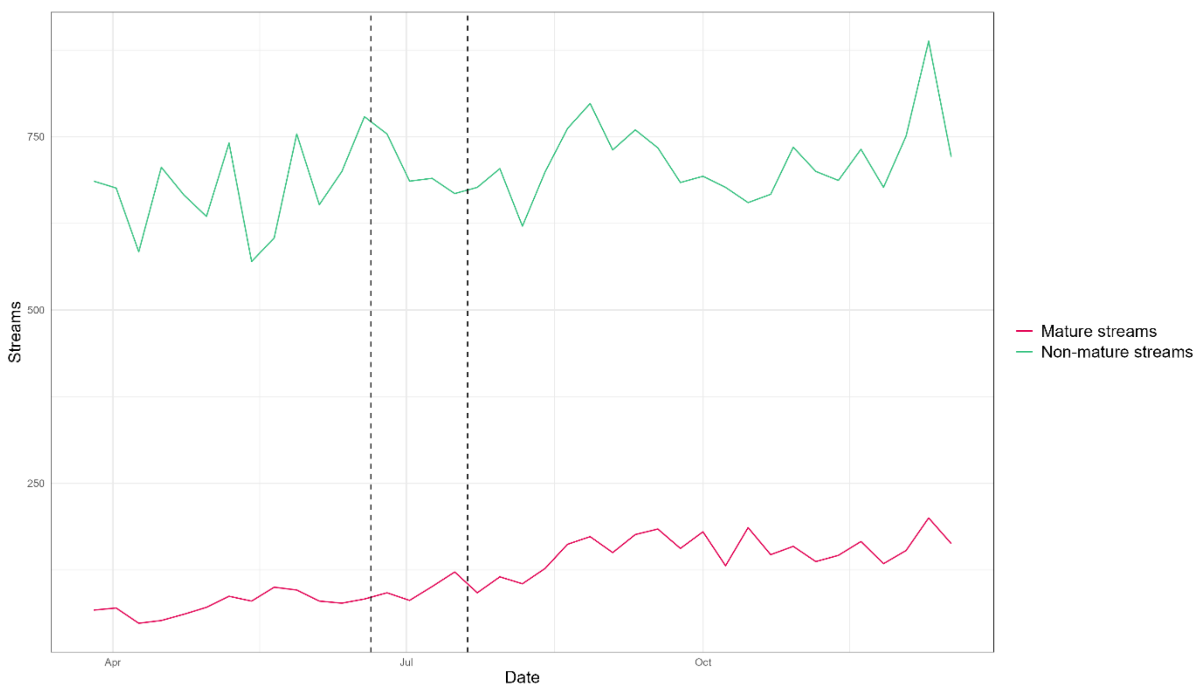


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

## Econometric analysis of content creator behaviour

5.49 In this section, we test the relationship between the CCL change and content creation behaviour more formally by conducting a panel regression analysis. As mentioned earlier, we limit our analysis to the immediate period surrounding the intervention and use both content creator and week 'fixed effects' to control for potential external factors that may bias our analysis. We also use a balanced panel of content creators to account for possible biases around attrition or new content creators joining later in the sample period. [53] Finally, and like our analysis on labelling, we exclude the 30-day grace period from our analysis.

5.50 Table 4 shows our fixed effects analysis for mature gaming and non-gaming streams separately.

**Table 4: Weekly number of mature streams by stream type - fixed effects analysis**

| | (1) Mature gaming streams (6-weeks) | (2) Mature gaming streams (8-weeks) | (3) Mature gaming streams (12-weeks) | (4) Mature non-gaming streams (6-weeks) | (5) Mature non-gaming streams (8-weeks) | (6) Mature non-gaming streams (12-weeks) |
|---|---|---|---|---|---|---|
| **CCL change** | 0.23***<br>(0.024) | 0.21***<br>(0.029) | 0.08***<br>(0.016) | 0.22***<br>(0.054) | 0.16***<br>(0.036) | 0.18***<br>(0.041) |
| **Content creator FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Time FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Observations** | 48,576 | 58,334 | 71,360 | 17,208 | 14,520 | 16,640 |

*Notes: Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01. Standard errors are clustered at the content creator-level. Our sample consists of a balanced panel of content creators and excludes the 30-day grace period provided by Twitch. Source: Ofcom analysis*

5.51 Consistent with the graphical analysis above, our results indicate that the weekly number of mature gaming and non-gaming streams produced by users increased following the grace period. However, while the effect appears to be statistically significant across all specifications, the magnitude of these coefficients remains small. In other words, our analysis in reveals only a small impact of the intervention on the absolute number of mature streams on the platform.[54]

[53] A balanced panel refers to a dataset where all units, in our case content creators, are observed at each time period analysed. By ensuring that the same group of content creators is observed both before and after the CCL change, this approach helps avoid falsely attributing changes in content creation behaviour to shifts in the composition of our sample over time.

[54] In addition to our analysis of the number of mature streams, we assessed the potential impact of the CCL change on the total weekly minutes of mature content produced by streamers. This sensitivity check revealed a limited impact of the intervention on the overall minutes of mature content on the platform, further highlighting the robustness of our baseline findings.

# Viewership of mature content

## Graphical analysis of viewership of mature content

5.52 One might speculate that the CCL change was introduced to affect viewership of mature content as, prior to the change, viewers had less information about the content of streams and as demonstrated earlier, the labelling of streams was less accurate. Improved information about the content of streams and accurate indications of whether content is mature might lead some viewers not to view mature content which they may otherwise have (inadvertently) viewed. Therefore, we also explore the effect of the CCL change on the absolute number of weekly views of mature and non-mature streams for all stream categories and the non-gaming category. Additionally, we look at the views of mature streams by individual CCL. [55]

5.53 Figure 14 and Figure 15 show the number of views of mature streams for all stream categories and the non-gaming category remained relatively constant before and after the CCL change, with a slight drop during the grace period for all stream categories.

**Figure 14: Weekly absolute views of mature and non-mature content for all stream categories**

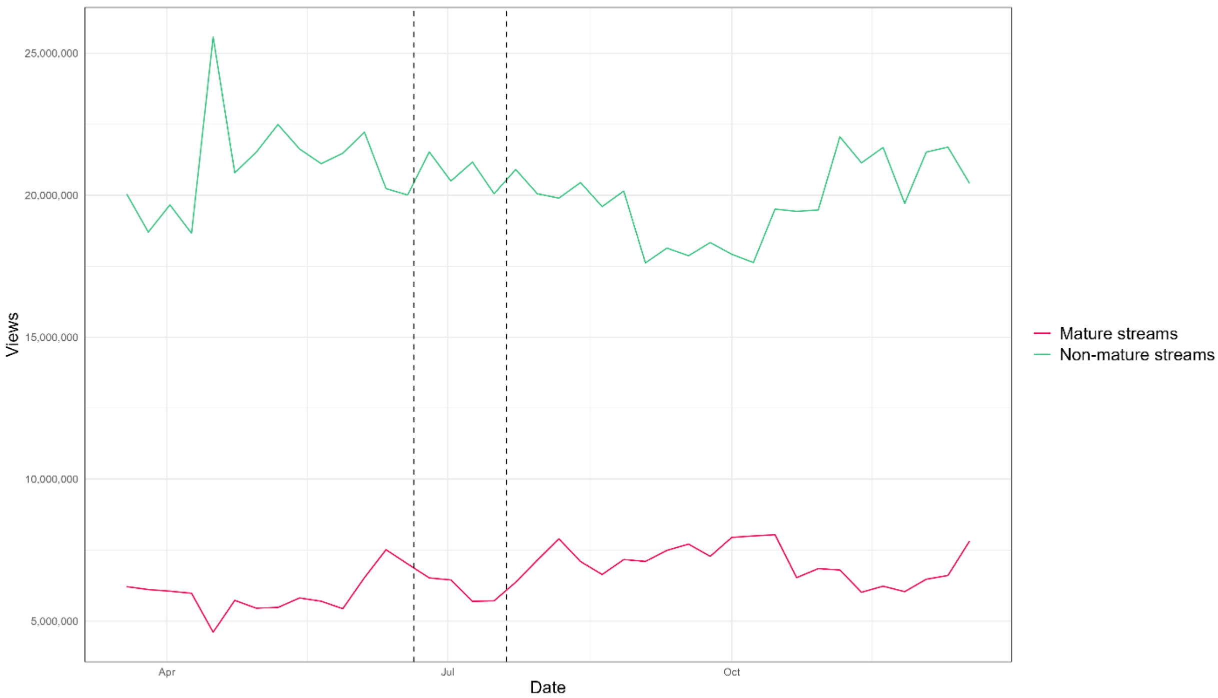


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023.*

[55] For our graphical analysis of viewing behaviour, we exclude views to streams between 07.03.2023 – 12.03.2023 and 18.12.2023 – 19.12.2023 in order to get the full week of data.

**Figure 15: Weekly absolute views of mature and non-mature streams for the non-gaming category**

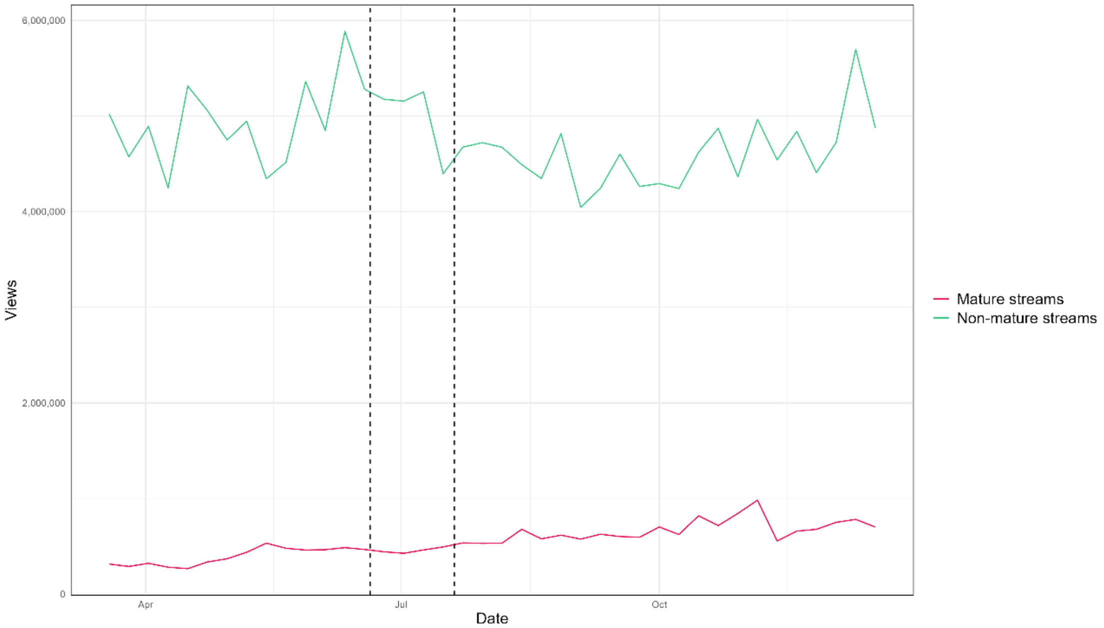


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

5.54 Figure 16 shows the views of streams containing the 'Mature-Rated Game' CCL. There is a significant drop in views during the grace period and a jump in views once the grace period ends. However, there is no clear trend in the views to these streams for the remainder of the post intervention period.[56]

5.55 Figure 17 shows the views of mature streams by CCL for the non-gaming category. Views of streams containing sexually suggestive and drugs related content account for the largest share of views. They show a clear upward trend even before the CCL change, and increasingly account for a larger share of views of the non-gaming category. [57]

[56] We classify a stream as a mature gaming stream if it is predicted to contain the 'Mature-Rated Game' CCL with a probability greater than 95%.

[57] We find that for non-gaming streams, views of streams containing sexually suggestive and drugs related content account for 49% and 39% of total mature views respectively.

**Figure 16: Weekly absolute views of streams predicted to contain the 'Mature-Rated Game' CCL**

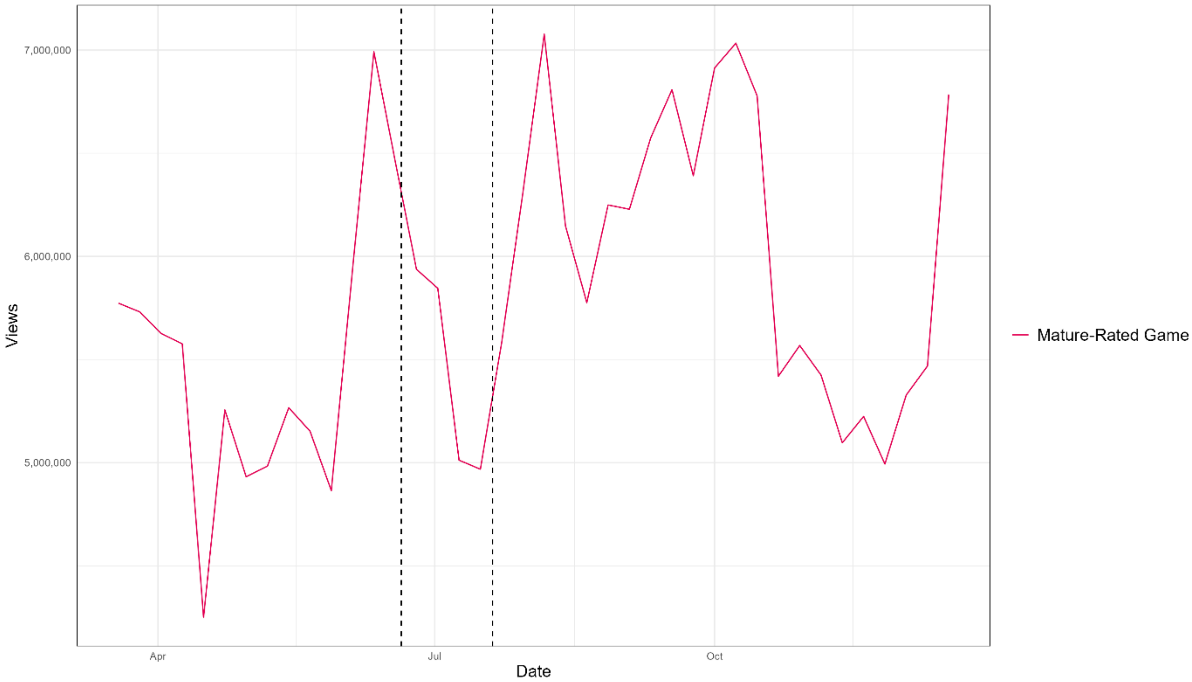


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

**Figure 17: Weekly absolute views of streams predicted to contain: Drugs/Intoxication; Gambling; Profanity & Vulgarity; Sexually Suggestive and Violent/Graphic content.**

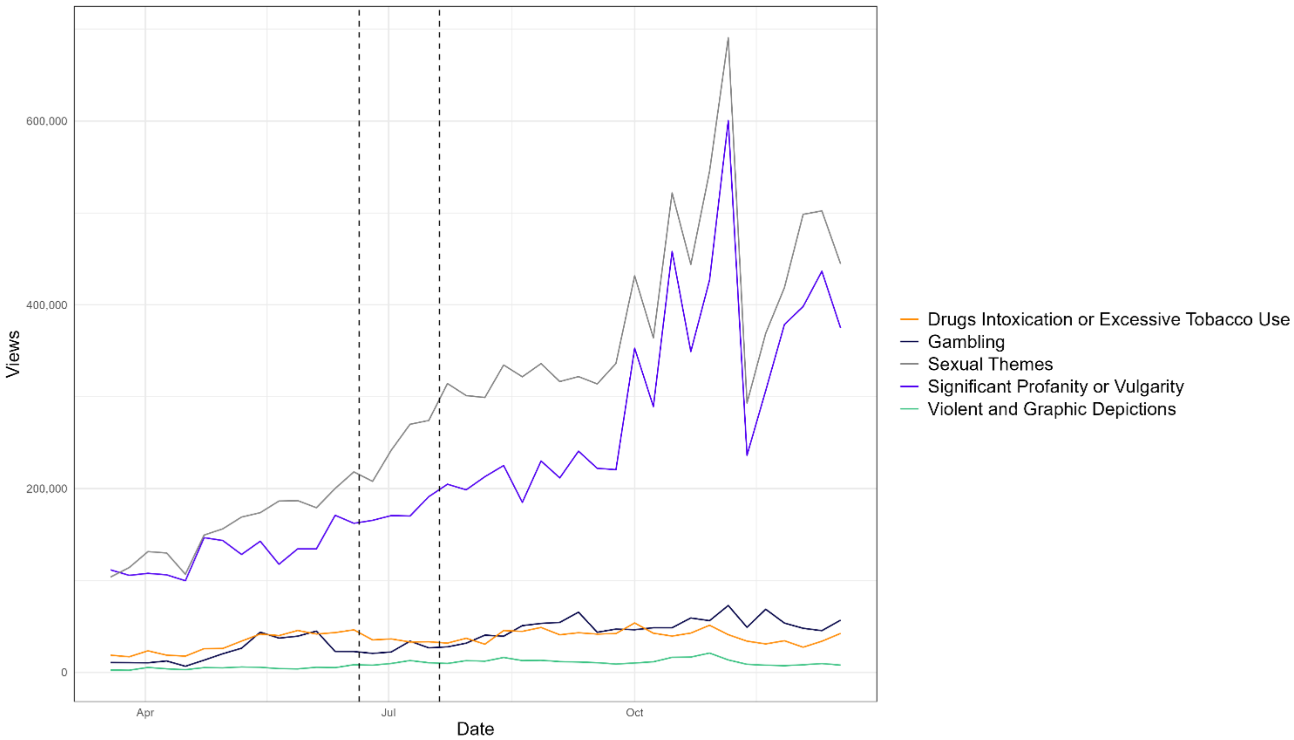


*Notes: The first vertical line represents the CCL change and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023.*

## Econometric analysis of viewership of mature content

5.56 Similar to our analysis of content creator behaviour, in this subsection we rely on panel regression models to get more precise and robust estimates of the potential effect of the CCL change on mature content viewership. As before, we focus on the immediate period surrounding the CCL change and use both content creator and time 'fixed effects' to control for potential external factors that may bias our analysis. Finally, we exclude the 30-day grace period from all specifications.[58]

5.57 Table 5 shows our fixed effects results for mature gaming and non-gaming streams separately.

[58] We also continue to focus on the subsample of content creators that were observed every week during the period.

**Table 5: Weekly views of mature game and non-game streams – fixed effects analysis**

| | (1) Views to mature gaming streams (6-weeks) | (2) Views to mature gaming streams (8-weeks) | (3) Views to mature gaming streams (12-weeks) | (4) Views to mature non-gaming streams (6-weeks) | (5) Views to mature non-gaming streams (8-weeks) | (6) Views to mature non-gaming streams (12-weeks) |
|---|---|---|---|---|---|---|
| **CCL change** | 306.55*** (102.937) | 133.47 (118.045) | 44.89 (282.540) | 228.93 (280.348) | 324.74 (293.227) | 732.77 (547.798) |
| **Content creator FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Time FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Observations** | 12192 | 14076 | 15640 | 1120 | 1200 | 1,080 |

***Notes:*** *Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01. Standard errors are clustered at the streamer-level. Our sample consists of a balanced panel of content creators and excludes the 30-day grace period provided by Twitch.*
*Source: Ofcom analysis*

5.58 While our results suggest an initial increase in the weekly number of views to mature gaming streams following the grace period, this effect is only statistically significant when using a 6-week bandwidth. Interestingly, as the bandwidth widens, the observed effect diminishes and becomes insignificant. This is consistent with our findings in Figure 16, where we observe a spike in views for streams with the 'Mature-Rated Game' CCL immediately after the grace period, which then stabilises over time. By contrast, we find no significant change in the weekly views for mature non-gaming streams, irrespective of the bandwidth used.

5.59 The graphical and econometric analysis in the sections above show that the CCL change did not materially alter the type of content produced by creators: we did not see a significant change in the amount of mature content creators produced before and after the CCL change. Similarly, we found that the CCL change did not alter viewer behaviour materially: the number of views of mature streams did not substantially differ before and after the CCL change.

## Conclusion

5.60 Overall, we find that more specific content labels, coupled with penalties for inaccurate labelling, led content creators to positively change their behaviour in terms of how accurately they labelled content. This change resulted in users being provided with much clearer alerts for mature content. However, we did not see any significant impact on the amount of mature content users viewed following the change, nor did we see a significant change in the amount of mature content creators produced.

# A1. Classifier analysis

A1.1 In this annex, we discuss our classifier analysis in further detail.

A1.2 In the raw data from Twitch, we only have detailed mature content labels for streams taking place after the CCL change (as these detailed content labels were not available to content creators before the CCL change), which made it challenging for us to directly compare outcomes before and after the change – in particular, to estimate how the CCL change impacted on whether mature streams were labelled as such and the impact of this labelling on viewership. We were able to construct reliable estimates of mature gaming streams before the CCL change using external data from PEGI and IGDB on game ratings and Steam on game listings, but for the non-game CCL categories we had no external means of determining which labels were likely to apply.

A1.3 Therefore, we sought to estimate in the pre-period whether content of streams (for which we did not have data) was mature, as well as other characteristics. We achieved this by employing a machine learning algorithm which looked at classification in the post-implementation period to predict classification in the pre-implementation period. Specifically, we trained a text classifier on a sub-sample of 2 million observations from the post-treatment observations to predict which CCLs apply to a given stream based on a subset of features that we have for all streams.

A1.4 Using this algorithm, we were able to correctly classify around 90% of all streams in the post CCL change period, which meant that we could use the classifier to estimate what labels would have applied in the pre-implementation period, if the CCLs had been available and applied accurately, with a reasonable degree of confidence. Additionally, as shown in Table 7 below, there was 99% agreement between our initial labelling of game streams using external sources and our labelling of game streams using the classifier's predictions.

A1.5 This is a significant result because it unlocked analysis that would not have been possible without the classifier work: comparison of the accuracy of labelling on Twitch both before and after the intervention and econometric analysis for non-game mature content streams. Without the classifier work this would only have been possible for game content streams.

A1.6 To minimise the risk of collecting personal data, we did not keep the full title of the stream beyond a month but we did extract several text features to use for classification: the name of the game used for the stream (which could be changed to whatever text the content creator wanted, including non-game phrases like 'just chatting'), the list of content creator-supplied tags for the stream, and a list of emojis we extracted from the original title in a preprocessing stage.

**Table 6: Stream features used for classification**

| A1.32 Feature name | Description |
|---|---|
| **Game name** | If the stream is a game stream, the name of the game being played. Otherwise, whatever description the content creator chose to assign to the stream. |
| **Tags** | A list of tags added to the stream by the content creator |
| **Emojis** | Emojis extracted from the title of the stream before deletion |

*Source: Twitch API Reference*

A1.7 Table 6 shows the top 10 of each kind of feature across all streams. 'Just chatting' was by far the most popular entry for the game name column – consistent with many streams not being about a specific game – while a mixture of mature and non-mature games makes up the rest of the top 10. 'English' is the most used tag by a wide margin but has been excluded because of its ubiquity. The remaining top tags cover a range of popular non-gaming categories of Twitch stream, including vtuber, anime, chatty, etc.

A1.8 Because of the type of classifier we used, it is difficult to ascertain directly which features contribute to an individual classification. However, given the extremely high accuracy of game classifications (see the following subsection for more detail), it is likely that the classifier was simply able to memorise which games are mature in the post-period to get near-perfect recall in the pre-period. The other categories rely on picking up rarer combinations of tags and emojis that connotate one of the CCLs, for instance the slot machine emoji 🎰 for Gambling.

*Figure 18: Top 10 features by number of streams*

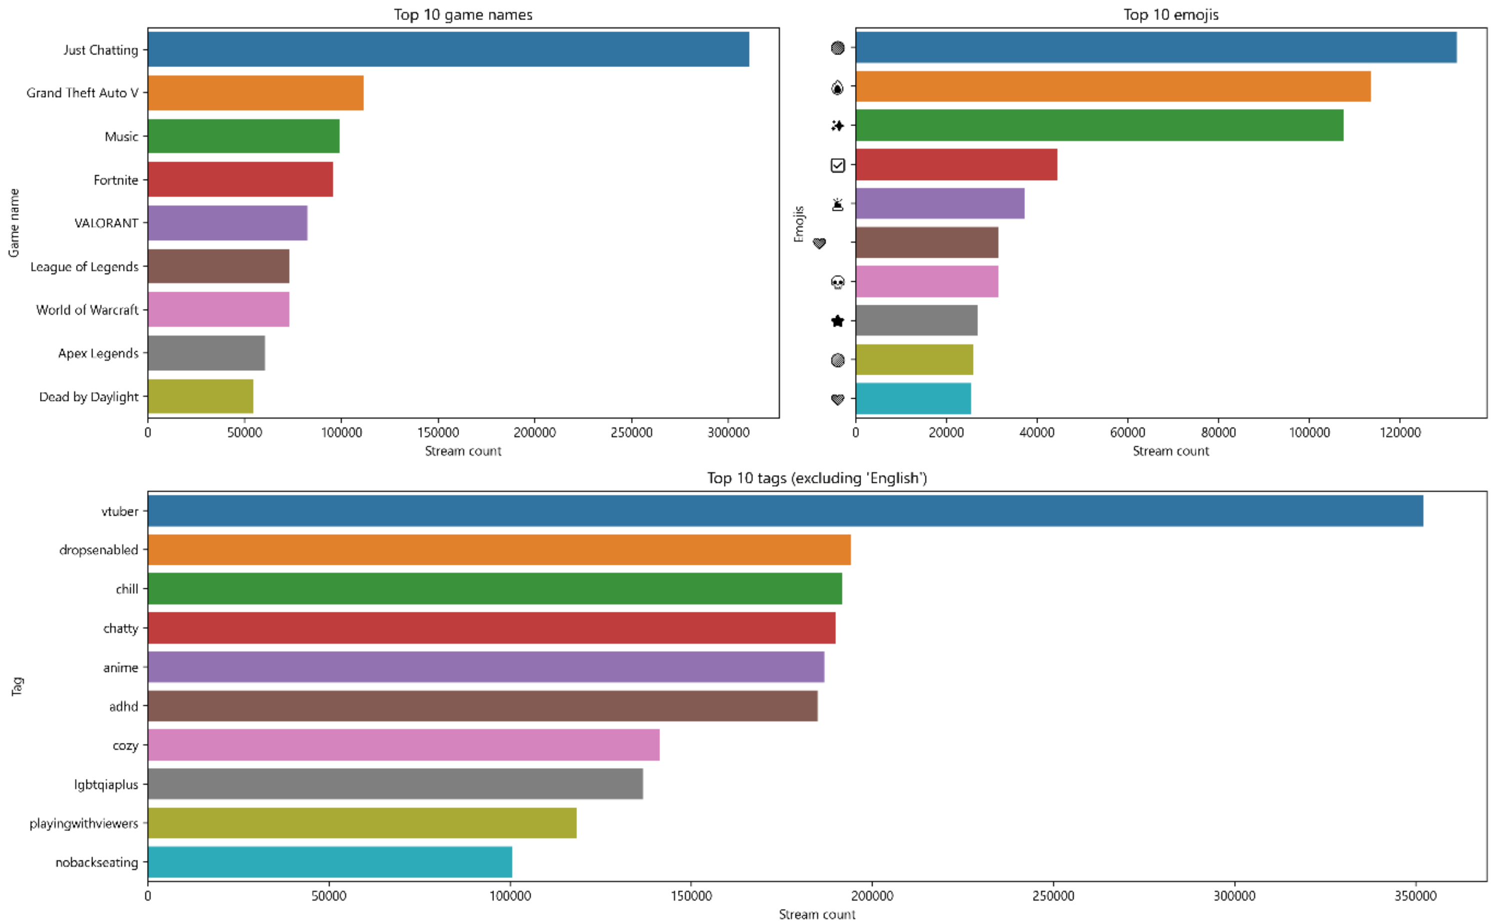


*Source: Ofcom analysis of Twitch 'Get Streams' API data*

A1.9 To maximise the accuracy of our predictions while minimising computational costs, [59] we fine-tuned a transformer-based language model for multilabel text classification on our dataset. Because transformer models can incorporate word order and the contextual meaning of words, we did not pre-process the text and instead simply concatenated all the features into a single line for each stream. We allowed for the possibility that multiple labels apply to a single stream and that no labels apply to a stream to conform with the application of CCLs in the data from Twitch. With a mid-range enterprise GPU, training took approximately 4 hours.

# Evaluation

A1.10 We used two methods to benchmark the performance of the classifier. This was an important step because we rely on the results of the classifier when assessing the effect of the CCL change on non-game streams. Without the classifier analysis our accuracy analysis would be restricted to looking at game streams. Below, we outline the steps used to evaluate the classifier's performance.

A1.11 First, we predicted the labels for a subset of the observations in the post-period that were not in the training set (the validation set) and calculated a number of variables which are

[59] We used a DistilBERT model (Sanh, et al, 2019) which achieves close to state of the art classification performance while being 40% smaller than the original BERT model.

typically used to evaluate the performance of automated classifiers using machine learning – in particular, precision[60], recall[61], and F1 scores[62] for each of the six CCLs. These results are shown in Table 7.

**Table 7: Performance metrics of classifier**

| CCL | Precision | Recall | F1 |
|---|---|---|---|
| Mature-Rated Game | 0.98 | 1.00 | 0.99 |
| Gambling | 0.86 | 0.72 | 0.79 |
| Sexual Themes | 0.94 | 0.89 | 0.91 |
| Significant Profanity or Vulgarity | 0.92 | 0.87 | 0.90 |
| Violent and Graphic Depictions | 0.89 | 0.89 | 0.84 |
| Drugs, Intoxication, or Excessive Tobacco Use | 0.91 | 0.81 | 0.86 |
| **Average** | **0.92** | **0.86** | **0.88** |

*Source: Ofcom analysis*

A1.12 Precision scores tell us what percentage of the CCLs the model predicts to be present are actually present. The precision scores are all above 0.86, which indicates that, when the model identifies a stream as having a particular CCL, it is highly likely to actually have that CCL. On the other hand, recall scores tell us what percentage of all the CCLs that are really present are picked up by the model. The recall scores range from 'perfect' for games to '0.72' for gambling streams. This means that all streams labelled by Twitch as containing mature games in the validation set were correctly identified, but only around 72% of streams labelled by Twitch as being gambling streams were identified. The average recall score of 0.86 is slightly below the average precision score of 0.92, indicating that the model overall is extremely likely to be correct when it predicts that a stream has a CCL and only marginally less likely to correctly identify all CCLs in the dataset.

A1.13 As a further check on the validity of the classifier model, we deployed a second method to evaluate the classifier. For the second check we generated predictions for all the streams in the pre-period and compare its predictions of mature games with the ones we had derived

---

[60] Precision is the share of positive labels that are correct.
[61] Recall is the share of positive samples that are labelled as positive.
[62] F1 is the harmonic mean of precision and recall.

from PEGI and IGDB earlier. Table 8 shows the 'confusion matrix' for game predictions at each confidence threshold. The share of games that were actually mature that the classifier predicted to be non-mature was 1% and the share of games that were actually non-mature that the classifier predicted to be mature was 0%. The results in the table mean we can be confident that the classifier exhibits strong out-of-sample performance for game CCLs.

A1.14 Since we had no external source of verifying the other non-gaming CCLs, we cannot check the model's performance for them on the streams collected before the CCL change. Given the evaluation metrics for the classifier, it is likely that the performance on these other CCLs is less strong overall. There is also a potential for bias if the relationship between the features and the types of mature content has changed over time – such as changing trends in the usage of certain emojis or tags to connote different types of mature content on Twitch over time. Given the fairly short time frame used for our analysis, it is unlikely that this had a significant impact on our results, but without a way to externally verify the model's predictions for non-game CCLs we cannot completely discard this possibility.

**Table 8: Confusion matrix for mature game predictions**

| Confidence threshold | Actually mature | Actually non-mature |
|---|---|---|
| **Predicted mature** | 22% | 0% |
| **Predicted non-mature** | 1% | 77% |

*Source: Ofcom analysis*

## Classifier results

A1.15 Figure 19 shows the classification predictions of mature and non-mature Twitch streams before and after the CCL change across all streaming categories. This differs from Figure 20 because it uses labels predicted by the classifier instead of those supplied by content creators and therefore Twitch's API. Moreover, unlike the approach used to inform Figure 6, the classifier includes predictions on non-games and games with an unknown rating. The maturity of these streams in these categories would have been unidentifiable without the classifier. [63] This is why this additional approach to classifying streams (before and after the CCL change) has been adopted. The classifier predicts fewer streams to be mature than content creators on Twitch manually labelled and the share of streams predicted to be mature increases in the post period relative to the pre-period.

[63] We classify streams as mature if they contain any CCL that is greater than the 95% probability threshold and a stream as non-mature if they do not. We also use a 75% probability threshold and find no significant changes.

**Figure 19: Predicted mature and non-mature streams in the pre- and post-treatment period**

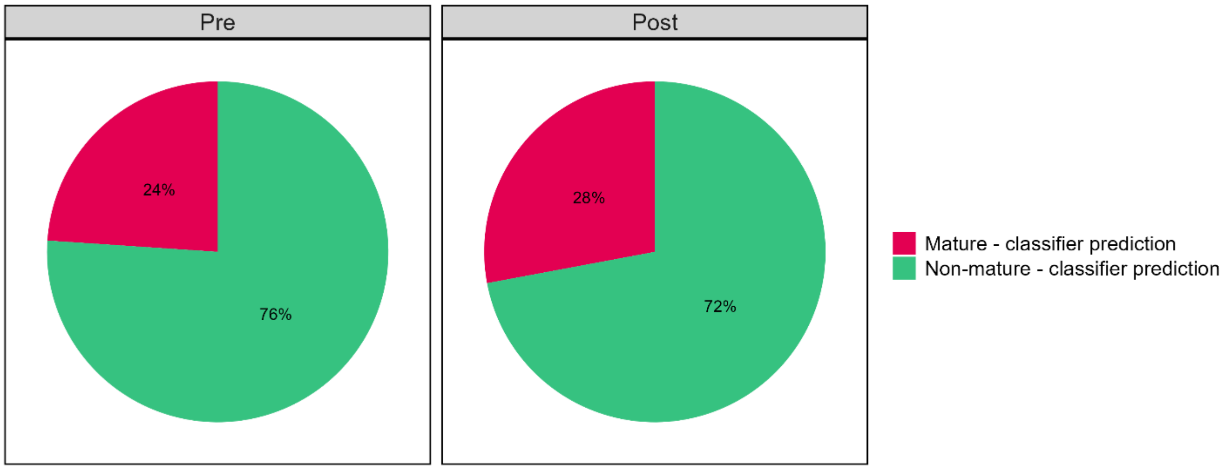


*Source: Ofcom analysis of Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

A1.16 Figure 20 shows the shares of Twitch streams labelled as mature and non-mature before and after the CCL change across all streaming categories. The plot on the left-hand side shows that around four in ten (39%) of the total volume of streams were labelled as mature using the binary *is_mature* label.[64] The plot on the right-hand side shows that the share of streams labelled as mature fell to 32% after the CCL change.

A1.17 These proportions are higher than the proportions reported in the figure above. This reflects a different basis for the figures. The figure above shows the proportion of predicted mature and non-mature streams, while the figure below shows the proportion of labelled mature and non-mature streams. The higher proportion of labelled mature streams, which is most prominent in the pre-CCL change period, likely reflects the finding of our accuracy analysis discussed in section 5. That analysis showed that there was a decline in content labelling inaccuracy across all stream types after the CCL change: the proportion of streams classified as non-mature but inaccurately labelled as mature was around 37% to 40% before the CCL change.[65]

[64] This is in contrast to section 4 where 38% of streams were labelled as mature. The differences are due to the following: first, there is a smaller eight-week sample used in section 4, compared to a 15-week sample in section 5; second, the sample of data in section 5 includes data collected at all times in the data collection schedule; third, streams which return an empty value in the 'game_name' field are included.

[65] As discussed in section 5, we take the post-CCL change period to be the 'ground truth', since we do not have any better information on content classification than what is available in the post-CCL period.

**Figure 20: Content creator labelled mature and non-mature streams in the pre- and post-treatment period**

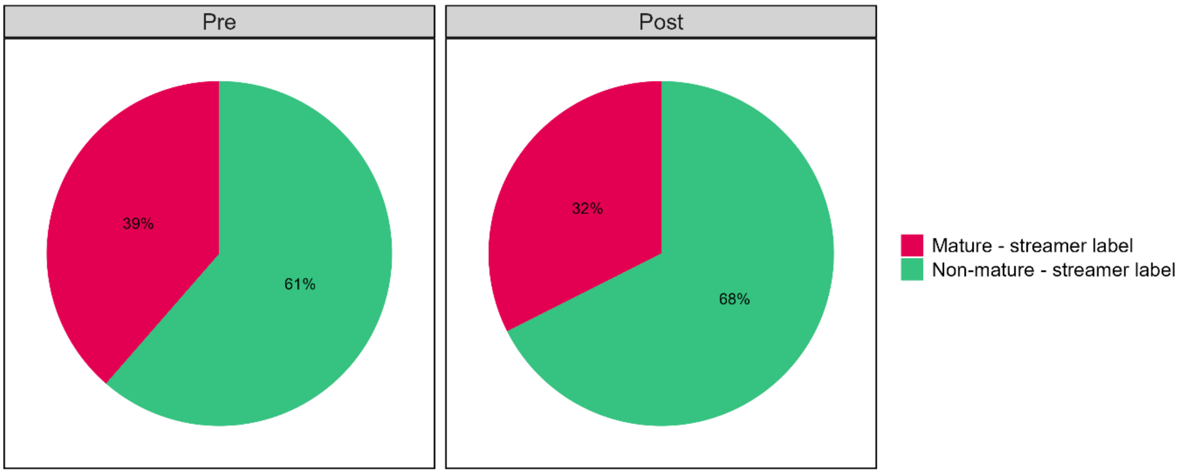


*Source: Twitch 'Get Streams' API and 'Get Channel Information' API data March and December 2023*

A1.18 Figure 21 shows the composition of mature streams by CCL. For each CCL we count the number of its occurrences, in streams which the classifier predicts as being mature, as a proportion of total CCL occurrences in these streams, since a mature stream may contain more than one CCL. The proportions are broadly similar between the pre and post CCL for those streams predicted to be mature. The 'Mature-Rated Game' CCL appears the most frequently in mature streams, however its relative frequency decreased in the post CCL period.[66] Streams containing gambling, drugs or intoxication, and violent or graphic content appear the least frequently in mature streams and account for no more than 8% of total CCL occurrences in both pre and post periods.

[66] This effect may be driven in part by the high threshold we used to determine whether a CCL would have applied in the pre-period (95%). Since the non-gaming CCLs are less prevalent in the training data, the classifier is less likely to assign a high level of confidence to those labels. If this has any impact on the econometrics, it would be to reduce the estimated effect sizes so the potential for false positives is minimal.

**Figure 21: Predicted mature streams by CCL in the pre- and post-treatment period**

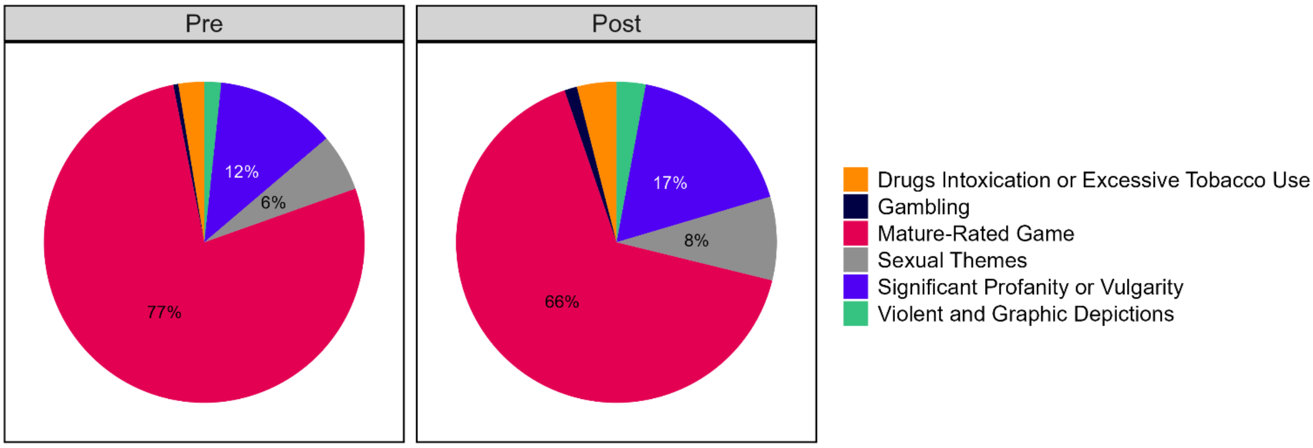


*Source: Ofcom analysis of Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

# A2. Econometrics

A2.1 In this annex, we discuss our econometric models in more detail and outline the main assumptions and limitations of our analysis.

## Regression Discontinuity Design (RDD)

A2.2 Given that the new labelling process was imposed on all Twitch users, we lacked a clear control group for a valid comparison with content creators exposed to the intervention. RDD models overcome this limitation by using the specific implementation date of an intervention as the relevant cut-off for estimating treatment effects, with individuals before and after this date falling in the control and treatment groups, respectively. The main idea behind this approach is that content creators before and after 20 July 2023 (i.e, the 'cut-off date') are similar, and so, that any differences in labelling behaviour can be attributed to the new mature content labelling system. By comparing content creators sufficiently close to this cut-off, RDD accounts for potential confounding factors that may bias the analysis and yields reliable estimates of the average treatment effect of CCLs.

A2.3 For RDD models to be valid, however, content creators before and after the cut-off date must be similar in their observable and unobservable characteristics. We believe this assumption holds in the context of our study. Indeed, as content creators remained unaware of the new content classification system until its introduction, it seems unlikely that they could have manipulated their exposure to CCLs. As such, we argue that that the decision to stream content just before and after the cut-off date – and thus, the assignment of content creators to a treatment and control group – was effectively quasi-random.

### RDD estimation

A2.4 In our analysis of mature stream labelling, all Twitch users were exposed to the new CCLs following their implementation. This implies that the assignment to treatment, Di, is perfectly determined by the running variable $X_i$, i.e., the number of days between the start of a given streaming session and the introduction of CCLs. If we define the specific implementation date of CCLs as c, this relationship can be expressed as follows:

$$\begin{cases} D_i = 1 \mid X_i \geq c \\ Di = 0 \mid X_i < c \end{cases}$$

A2.5 Therefore, in this 'sharp RDD' scenario, the treatment effect is determined by the difference in content creator labelling behaviour before and after the implementation of CCLs. More generally, the average causal effect of the intervention at the cut-off can be expressed as follows:

$$\tau = \lim_{x \to c^+} E[Y \mid X = x] - \lim_{x \to c^-} E[Y \mid X = x]$$

A2.6 To estimate this causal effect, we rely on the *rdrobust* package in R, a data-driven tool for nonparametric estimation of RDD models.[67] Unlike parametric approaches which require

[67] Calonico, S., Cattaneo, M.D., and Titiunik, R. *rdrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs.* R Journal 7(1), 2015, pp. 38-5. https://journal.r-project.org/archive/2015/RJ-2015-004/RJ-2015-004.pdf

one to make explicit functional form assumptions,[68] non-parametric regressions allow the data itself to determine the relationship between the outcome of interest and the running variable.[69] For our analysis, we use the *rdrobust* package default options which rely on a triangular kernel that assigns a greater weight to units closer to the cut-off date.

## Choice of bandwidth

A2.7 The validity of RDD models is affected by the optimal selection of the bandwidth, i.e., the number of days or weeks included in the analysis. Although selecting a narrow bandwidth would mitigate the risk of bias – by helping us ensure that external factors are constant, and observations are similar in their characteristics – a smaller sample size could also potentially increase variance in our model. By contrast, a larger bandwidth would avoid concerns of high variance but could introduce bias in the analysis to the extent that it relies on observations that are further away from the cut-off.

A2.8 The optimal bandwidth for our analysis is influenced by the 30-day grace period that Twitch gave to all its users. Although we exclude this period from our baseline analysis, it is essential we include enough observations in the post-intervention period to accurately capture treatment effects beyond this transitional phase. However, rather than arbitrarily selecting a bandwidth and risk biasing the results, we rely on the default options of the *rdrobust* software. Using a data-driven process, the package calculates the optimal bandwidth that minimises the mean squared error of the treatment estimator, thus optimising the bias-variance trade-off mentioned earlier.

## Identifying assumptions

A2.9 For our RDD model to be valid, we require treatment status to be determined exclusively by the cut-off date and not by other external factors. This is because if specific users self-select into treatment, our treatment and control groups would no longer be similar in their observable and unobservable characteristics. Thus, to the extent that treated and untreated content creators differ systematically, we could not attribute changes in our outcomes of interest to the introduction of CCLs on the platform. However, because the changes apply to all users and content creators were unaware of the new policy, it seems unlikely that content creators would have delayed their streams to make use of the new labelling system.

A2.10 A more plausible threat to identification could be the adoption of strategic behaviour by users. For example, some content creators could manipulate their treatment assignment by ceasing to stream CCL related content or, in a more extreme scenario, stop using Twitch altogether. Unfortunately, to the extent that this strategic behaviour is correlated with characteristics related to our outcomes of interest, our estimates could be biased.

[68] To the extent that the specified functional form does not sufficiently represent the true relationship between labelling rates and the running variable, parametric regressions could overestimate the treatment effect by mistaking a non-linearity for a discontinuity.

[69] However, we test the robustness of our results by estimating traditional parametric regressions. For this analysis, we follow the literature and use a quadratic function of the running variable to help us account for potential non-linearities. The results obtained from this robustness test are highly consistent with those derived from the baseline analysis.

A2.11 We assess the validity of our RDD model by testing the continuity assumption.[70] To do so, we visually inspect characteristics of streams that should have been unaffected by CCLs – e.g., stream duration and emojis used in stream titles – just before and after their implementation. If our assumption holds, these characteristics should remain relatively constant and smooth on both sides of the cut-off. Figure 22 and Figure 23 show the results to our continuity tests by stream type. We do not observe a jump in either stream duration or emojis used in stream titles following the introduction of CCLs, supporting our assumption that content creators did not self-select into treatment.

[70] This assumption states that the characteristics of treated and untreated units do not differ just before and after the implementation of a given policy.

**Figure 22: Continuity check - number of emojis in stream title**

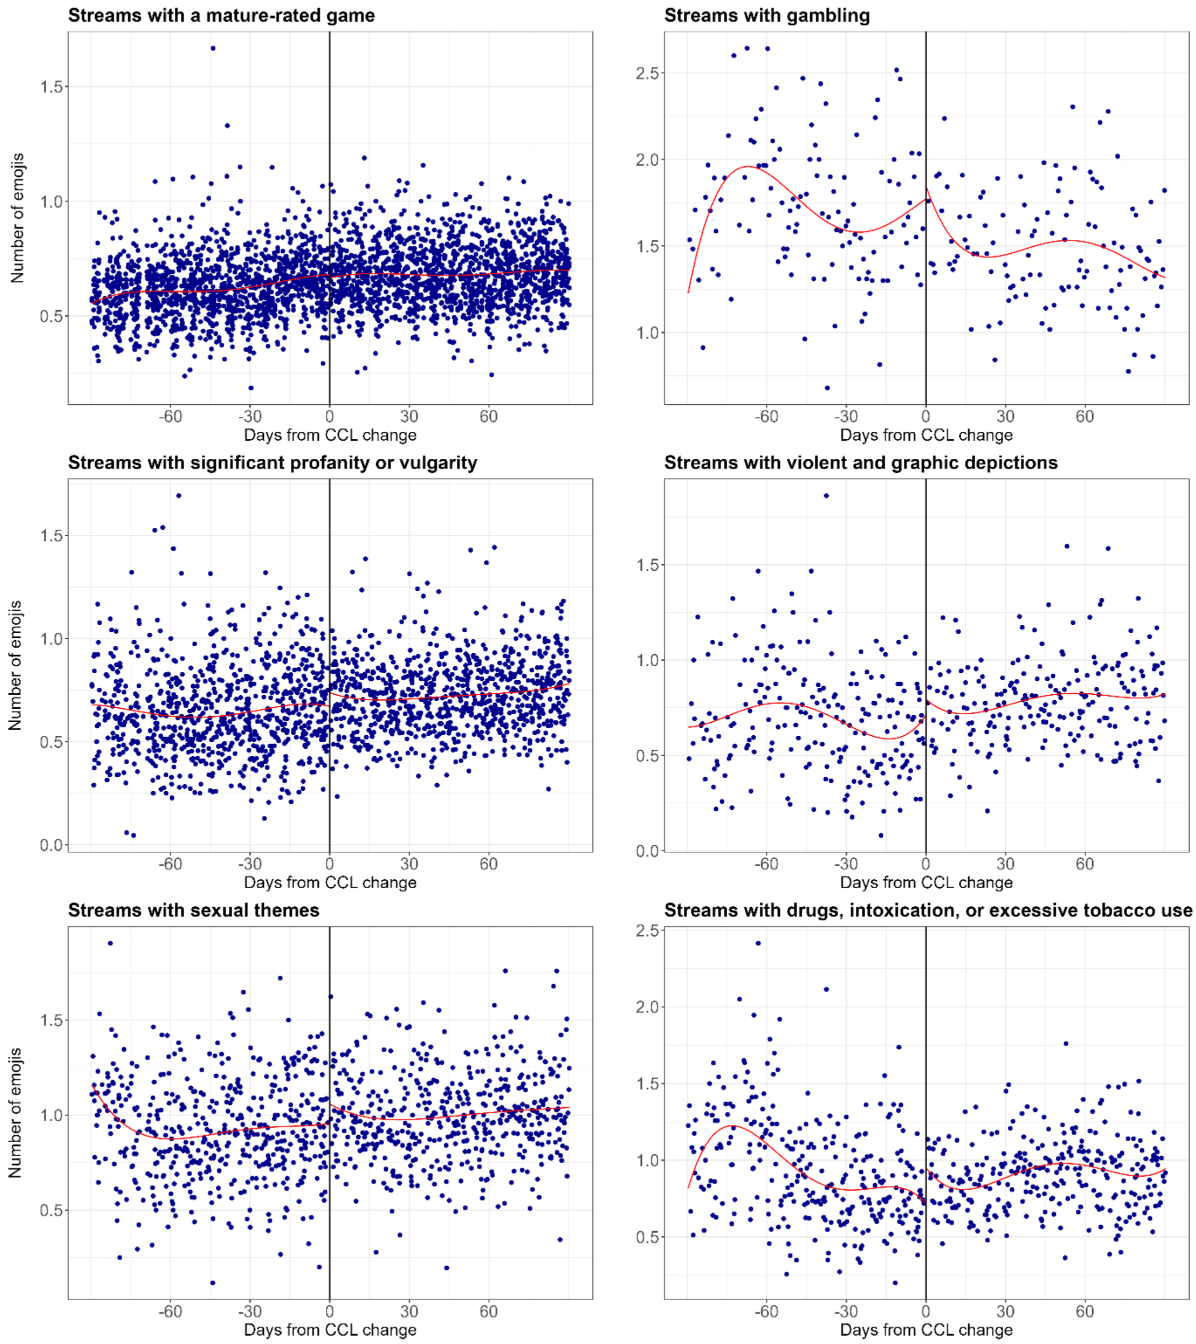


*Notes: The results for streams with a mature-rated game are based on a 20 June cut-off, while the analysis of mature non-gaming streams relies on a 20 July cut-off. Each plotted point represents the proportion of streams labelled as mature within a given 30-minute interval on a given day. Source: Ofcom analysis of Twitch 'Get Streams' and 'Get Channel Information' API data March and September 2023.*

**Figure 23: Continuity check - stream duration**

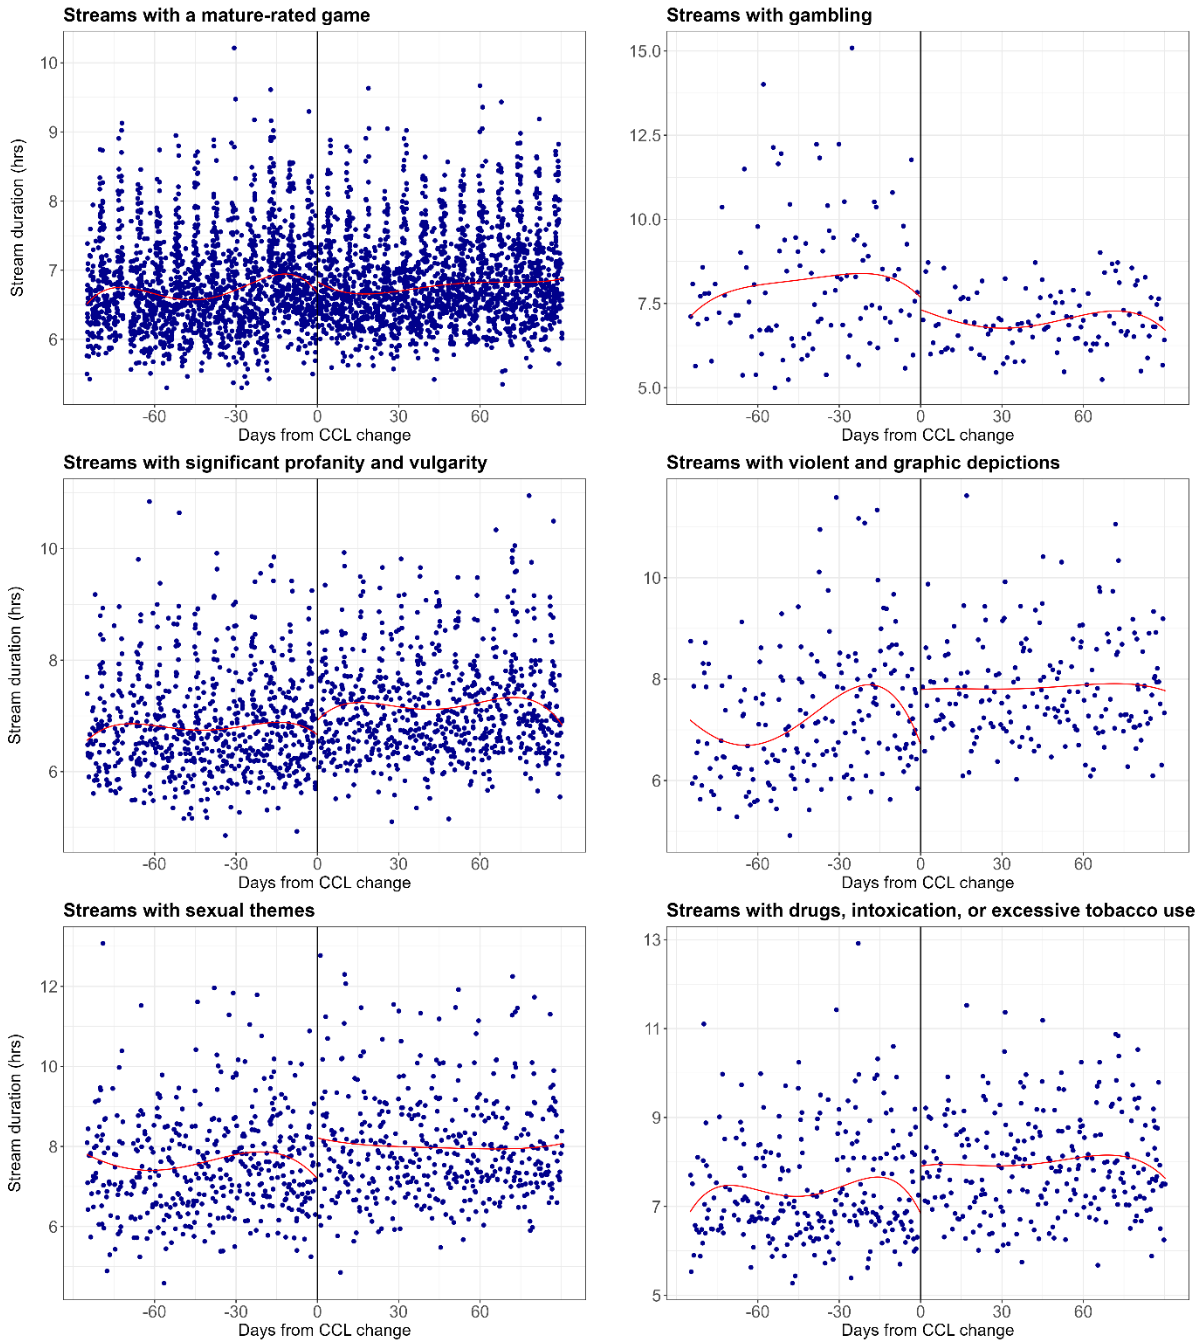


*Notes: The results for streams with a mature-rated game are based on a 20 June cut-off, while the analysis of mature non-gaming streams relies on a 20 July cut-off. Each plotted point represents the proportion of streams labelled as mature within a given 30-minute interval on a given day. Source: Ofcom analysis of Twitch 'Get Streams' and 'Get Channel Information' API data March and September 2023.*

# Limitations of our methodology

A2.12 While the continuity tests support the validity of our estimates, our analysis still presents some limitations. In particular, the causal effects recovered by our RDD model represent local treatment effects. In other words, they apply only to the specific temporal windows used in the analysis and so cannot be extrapolated away from the cut-off. Thus, although our baseline analysis suggests that the probability of mature content labelling increased shortly after the intervention, the effect could be smaller (or larger) when examined at other time periods.

A2.13 Moreover, while the analysis in Figure 22 and Figure 23 suggest no self-selection into treatment, we recognise that manipulation of treatment assignment is still possible. As mentioned earlier, content creators could adjust their behaviour, for example, by ceasing to stream mature content or stopping their use of Twitch. To the extent that this strategic behaviour takes place and those who remain are more likely to label their content regardless of the intervention, our estimates could overestimate the true effect of CCLs.

A2.14 Finally, it is important to note that our RDD results represent only a lower-bound estimate of the true impact of CCLs on mature content labelling. While treating the classifier's predictions of CCLs as 'ground truth' allows us to perform a like-for-like comparison of labelling before and after the change, this comes at the cost of introducing measurement error in our analysis. While we are confident over the overall performance of our classifier, the predictions from our model may not always be 100% accurate. This is because factors such as noisy data, incomplete or incorrect information on CCL use, and/or changing relationships between input features (e.g., emojis, tags) and mature content types may lead to inaccurate predictions of CCLs. This, in turn, introduces measurement error in our dependent variable and biases our estimates downwards towards zero.

## Fixed effects model

A2.15 To assess the impact of CCLs on the weekly number of views to mature content and weekly volume of mature content produced by content creators, we exploit the panel nature of our data by estimating a regression model with content creator fixed effects, denoted by $\alpha_i$. This allows us to control for potential time invariant characteristics of content creators (e.g., gender, streaming styles) that may influence our outcomes of interest. We also control for week fixed effects, captured by $\gamma_t$, to account for any time-related effects that may impact all content creators in a similar way, e.g., (game releases, esports tournaments, etc.). Let $y_{i,t}$ denote the outcomes of interest for content creator $i$ in week $t$, and $D_t$ represent the introduction of CCLs. The specification we estimate is:

$$y_{i,t} = \beta_1 D_t + \alpha_i + \gamma_t + \varepsilon_{i,t}$$

A2.16 To further mitigate the impact of other potential confounding factors, we follow our RDD set-up and limit our analysis to the weeks surrounding the implementation of CCLs. We also use different sets of bandwidths, including a 6, 8, and 12-week window, to test the robustness of our results. [71] To the extent that the error term, $\varepsilon_{i,t}$ is uncorrelated with our intervention dummy during the period of interest, our fixed effects model can provide us a with an unbiased estimate of the effect of CCLs on viewership of mature content and mature content creation.

A2.17 Although we are confident that the use of time and content creator fixed effects, combined with a narrow bandwidth, help us mitigate concerns around endogeneity, we recognise that our analysis may still be susceptible to bias. This is because we cannot fully rule out the possibility that other omitted, time varying factors impact both the treatment and our outcomes of interest.

[71] Our baseline analysis also relies on a balanced panel of content creators to account for potential attrition bias.

# A3. Data Pipeline

A3.1 In this section of the annex we discuss in detail, the data pipeline constructed to collect streaming data from the APIs between the 07.03.2023 16:00:00 to 19.12.2023 02:37:00 and the measures taken to protect personal data and user content.[72] The figure below provides a flow diagram of the data pipeline constructed in Python code using the Azure Service and Twitch's API endpoints.

**Figure 24: Graphical overview of the data pipeline**

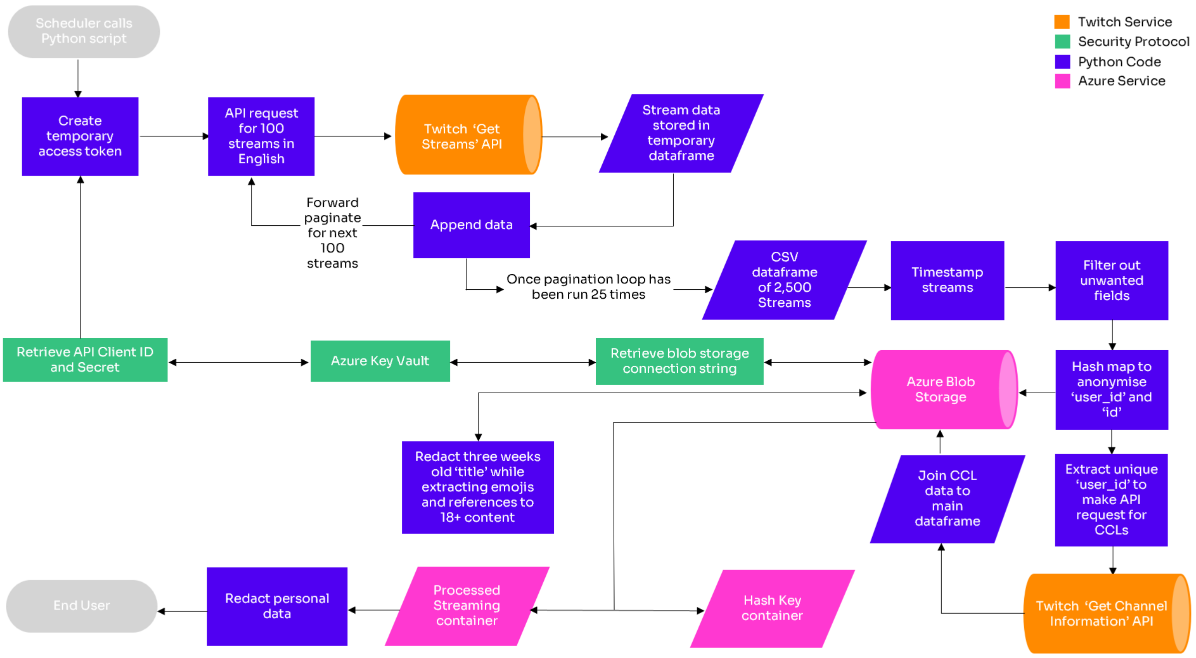


*Source: Ofcom*

## Get Streams API

A3.2 The data pipeline was scheduled to run according to the data collection schedule, to retrieve the top 2,500 streams by views at thirty-minute intervals from the 'Get Streams' API.[73]

A3.3 For each call to the API, streams are requested using credentials stored securely in the Azure Key Vault. This is an encrypted storage vault with permissions to these credentials reviewed every six months.

A3.4 Due to the 100-item stream limit, the API employs forward pagination whereby it makes 25 consecutive calls to the API for the top 100 most viewed streams at the time of the query, up to the 2401$^{st}$-2500$^{th}$ most viewed streams at the time of the query. For each 100 streams requested, they are stored in a temporary dataframe and then appended to main dataframe, which continues until all 2,500 streams have been requested. However, there are

[72] Before the 13.03.2023, 'user_id' and 'stream_id' is hashed as a random combination of integers but from 13.03.2023 onwards they are hashed as combination of strings and integers.

[73] 2,500 streams are the limit imposed on Ofcom by Twitch and is a potential limitation to our results as it is an incomplete representation of the overall viewing and streaming habits across the Twitch platform.

technical limitations of the API call such that the top 2,500 streams were not always retrieved:

A3.5 Firstly, the API may return less than 100 streams requested per page, which is often the case for the last page of results.

A3.6 Secondly, if views to the content creator decreases following the next pagination, the stream would appear on both pages creating data duplicate issues. Additionally, if the content creator continues the stream during the next 30-minute interval when the next 2,500 streams are retrieved, the stream will appear more than once. To rectify this, streams identified as duplicates are removed.[74]

A3.7 The raw CSV dataframe of 2,500 streams are timestamped with unwanted fields filtered out. The fields retained from the 'Get Streams' API are: 'user_id', which provides the ID of the content creator; 'id', which identifies the stream; 'title', which is the stream's title; 'game_name', the name of the game being played or the non-gaming topic; 'viewer_count', the number of viewers watching the stream at the time of the query; 'started_at', UTC timestamp for when the stream started; 'language', the language used by the content creator; 'tags', tags are applied to the stream by the content creator and provide more information about what topics are covered in a live stream; 'is_mature', the binary maturity label applied by content creators before the introduction of the granular CCLs.[75]

A3.8 The hash map assigns a random combination of strings and integers to each original 'user_id' and 'id'. The hashed ids replace the original ids in the CSV dataframe.

## Get Channel Information API

A3.9 Figure 25 sets out the timeline concerning the key changes to 'Get Channel Information' API endpoint, which enabled the collection of the metadata 'content_classification_labels'. They are the granular mature labels which replaced the binary 'is_mature' label following the intervention.[76]

i) Figure 25 shows that Twitch launched the Beta API on 20 June 2023. However, Ofcom were not granted beta API test access until 22 June 2023. We experienced data drops when we were granted access, which meant that 'content_classification_labels' did not start to come through the pipeline until 23 June 2023.[77]

ii) To retrieve the 'content_classification_labels' from the Get Channel Information endpoint, Figure 24 shows that the 'user_id' extracted from the Get Streams API is used to make API requests from the Get Channel Information endpoint by matching on the 'broadcaster_id'. For each 'broadcaster_id' the CCLs applied by the content creator is returned.

[74] There are two types of duplicate streams removed. Firstly, streams are identified as duplicates if the same 'user_id', 'id' and 'game_name' of the stream appear more than once, indicating the content creator has continued the stream but has not changed the 'game_name'. Secondly, streams are identified as duplicates if the same 'user_id', 'id' and 'datetime' of the stream appear more than once, indicating the content creator has continued the stream but has not changed the 'game_name'.

[75] A content creator can apply 25 'tags' in total where up to 10 'tags' are customizable by the content creator.

[76] The 'Get Channel Information' API endpoint has existed on Twitch's publicly available API even before the intervention. However, this data was not necessary to collect before the intervention.

[77] The Beta API granted us access to the metadata 'content_classification_labels' which were unavailable to the public at the time via the 'Get Channel Information' API endpoint.

iii) Similarly, to the ‘Get Streams’ endpoint, Twitch imposes a limit on the amount of CCL data that can be retrieved from the ‘Get Channel Information’ endpoint, whereby CCL data can only be retrieved for 100 unique content creators at one point in time. We resolve this by sequentially requesting CCL data for 100 unique content creators each time from a hash map of ‘user_id’. The dataframe of retrieved channel information includes the ‘content_classification_labels’ and is then joined onto the main dataframe from the Get Streams API endpoint, by merging on the hashed ‘user_id’.

iv) Figure 25 also shows that Twitch made reformatting changes to how the CCL data was returned from the ‘Get Channel Information API’ endpoint which involved renaming the five of the six CCLs. As a result, we experienced data losses to the ‘content_classification_labels’ between 10 July-11 July.[78]

**Figure 25: Timeline of changes to the Twitch API**

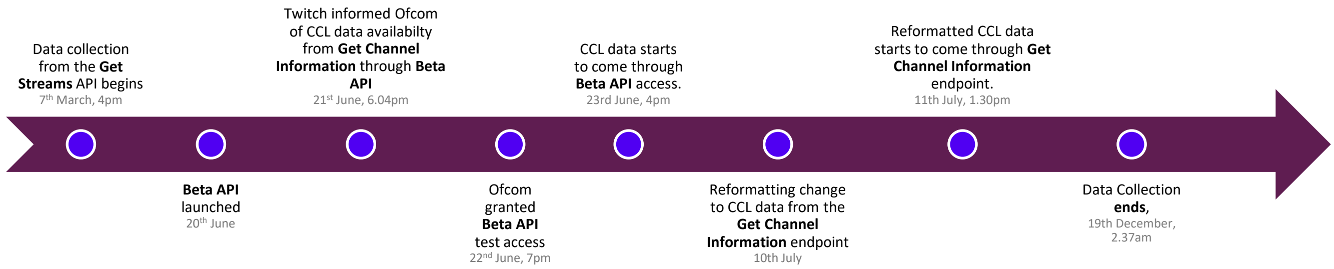


*Source: Ofcom*

## Data Protection

A3.10 To process data in accordance with Information Commissioner’s Office (ICO) guidelines, the following measures were taken to protect personal data and user content throughout the pipeline before reaching the end user:

i) The ‘user_id’ and ‘stream_id’ retrieved from the Get Streams API endpoint are personal data as they could be used to identify content creators on Twitch. As shown in Figure 24, these IDs were hashed whereby a random combination of strings and integers were assigned to each respective ID. The map of hashed IDs to original ‘user_id’ and ‘stream_id’ was stored in a hash key container not accessible by the analysis team. Furthermore, the original ‘user_id’ and ‘stream_id’ was redacted before reaching the end user.

ii) The ‘broadcaster_id’ and ‘content_classification_labels’ were the only metadata of interest retrieved from the ‘Get Channel Information’ API.[79] Given that ‘broadcaster_id’ is personal data, prior to merging the ‘Get Channel Information’ API data with the ‘Get Streams’ API data, we replace the ‘broadcaster_id’ with the

[78] The following *content_classification_labels* were renamed as follows: ‘Mature-rated game’ renamed to ‘MatureGame’; ‘Significant Profanity or Vulgarity’ renamed to ‘ProfanityVulgarity’; ‘Sexual Themes’ renamed to ‘SexualThemes’; ‘Drugs, Intoxication or Excessive Tobacco Use’ renamed to ‘DrugsIntoxication’, ‘Violent and Graphic Depictions’ renamed to ‘ViolentGraphic’.

[79] The *content_classification_labels* are retrieved as a string of CCLs applied to the content creator’s channel from the ‘Get Channel Information’ API endpoint. A column is created for each CCL with Boolean values to indicate whether each stream contained a given CCL.

hashed ‘user_id’. We did not retrieve any of the other metadata available through the ‘Get Channel Information’ API as they were not relevant to our analysis.

iii) The metadata ‘title’ is user content. Although it did not technically constitute personal data, the ‘title’ was free text meaning that it could theoretically contain some personal data. Because of this we took steps to comply with data protection law. Specifically, we ran an additional pipeline, shown in Figure 24, on a weekly basis to delete stream titles that were three weeks old from the Azure Cloud. This ensured that stream titles were removed from our systems within 1 month of collection, which is in line with the ICO's guidance on individuals' right to erasure. We also applied specific natural language processing (NLP) algorithms to extract emojis and references to 18+ content used in the titles of each stream before the ‘title’ was deleted.

## Data Losses

A3.11 During our collection of data from the ‘Get Streams’ API and ‘Get Channel Information’ API endpoints, we encountered multiple instances of complete data losses or interruptions in the API calls. The data losses are due to changes made to the ‘Get Channel Information’ API endpoint as mentioned above, as well as some unknown factors. They are outlined in Table 9 below.

**Table 9: Data losses**

| Date | Time | ‘Get Streams’ Data Loss | ‘Get Channel Information’ Data Loss | Reason |
|---|---|---|---|---|
| 07.03.2023 | 00:00 – 02:00 (Five instances) | Yes | N/A | Unknown |
| 09.03.2023 | 20:00 – 23:30 (Eight instances) | Yes | N/A | Unknown |
| 13.03.2023 | 00:00 – 02:00 (Five instances) | Yes | N/A | Unknown |
| 15.03.2023 | 01:30 (One instance) | Yes | N/A | Unknown |
| 25.03.2023 | 22:00 – 23:30 (Four instances) | Yes | N/A | Unknown |
| 27.03.2023 | 00:00 – 02:00, 16:30, 17:00 (Seven instances) | Yes | N/A | Unknown |
| 05.04.2023 | 20:00 – 23:30 (Eight instances) | Yes | N/A | Unknown |
| 06.04.2023 | 00:00 – 02:00 (Five instances) | Yes | N/A | Unknown |
| 12.04.2023 | API call interrupted at 16:00 | Yes | N/A | Unknown |

| Date | Time | 'Get Streams' Data Loss | 'Get Channel Information' Data Loss | Reason |
|---|---|---|---|---|
| 17.04.2023 | No data between 00:00 – 02:00 (with 5 instances) API call interrupted between 16:00 – 17:30 | Yes | N/A | Unknown |
| 20.05.2023 | 09:00 – 10:00, 13:00 – 15:30, 17:30 (Ten instances) | Yes | N/A | Unknown |
| 20.06.2023 - 22.06.2023 | 21:00 – 23:30 (last day) (Forty-five instances) | No | Yes | Lack of prior information about how the CCL would come through from the 'Get Channel Information' API |
| 10.07.2023 - 11.07.2023 | 22:30 – 02:30 (Nine instances) | Yes | Yes | Twitch made a change to how the CCL data was returned from the 'Get Channel Information' API |
| 05.08.2023 | 00:30 (One instance) | Yes | Yes | Unknown |
| 26.08.2023 | 04:00 (One instance) | Yes | Yes | Unknown |
| 08.09.2023 | 01:30, 03:00 – 04:00 (Four instances) | Yes | Yes | Unknown |
| 18.11.2023 | 22:30, 23:30 (Two instances) | Yes | Yes | Unknown |
| 19.12.2023 | 16:00 – 17:30, 20:00 – 23:30 (Twelve instances) | Yes | Yes | Unknown |

# A4. Robustness Checks

A4.1 In this section of the annex, we report the results from the classifier and econometric analysis using a 75% probability threshold. We use a lower probability threshold since the classifier is less likely to assign a high level of confidence to non-gaming CCLs as they are less prevalent in the training data. We find that our results are largely consistent with our graphical and econometric analysis at the 95% probability threshold.

## Classifier Analysis

A4.2 Figure 26 shows the classification predictions of mature and non-mature Twitch streams before and after the introduction of the CCLs using a 75% threshold. Relative to the 95% probability threshold, the 75% probability threshold increases the share of streams predicted to be mature by only two percentage points before and after the intervention.

**Figure 26: Predicted mature and non-mature streams in the pre- and post-CCL period**

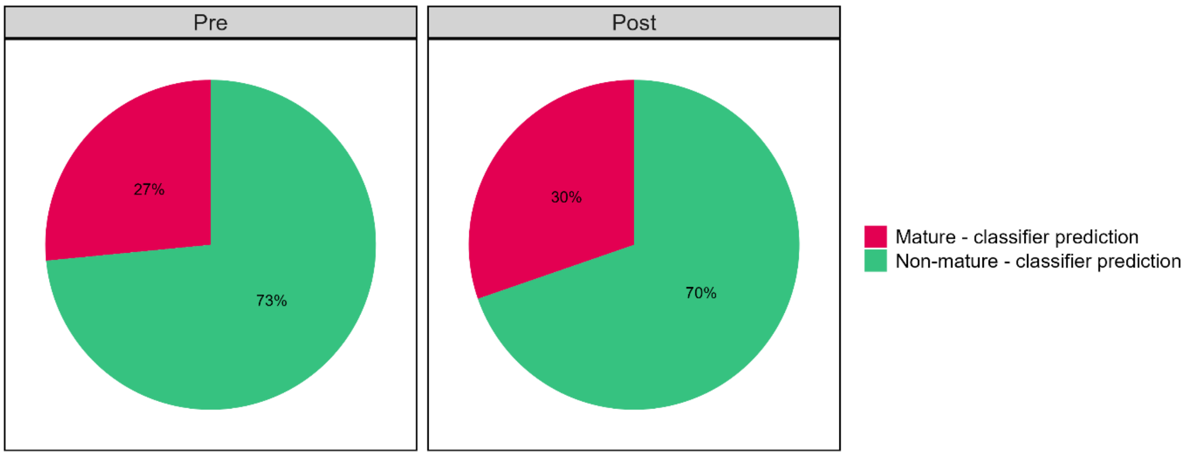


*Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

A4.3 Figure 27 shows the composition of mature streams by CCL using the lower probability threshold. Although the 'Mature-Rated Game' CCL remains the most frequently occurring CCL, we find that the relative frequency of the non-gaming CCLs has increased under the 75% probability threshold.

**Figure 27: Predicted mature streams by CCL in the pre- and post-treatment period**

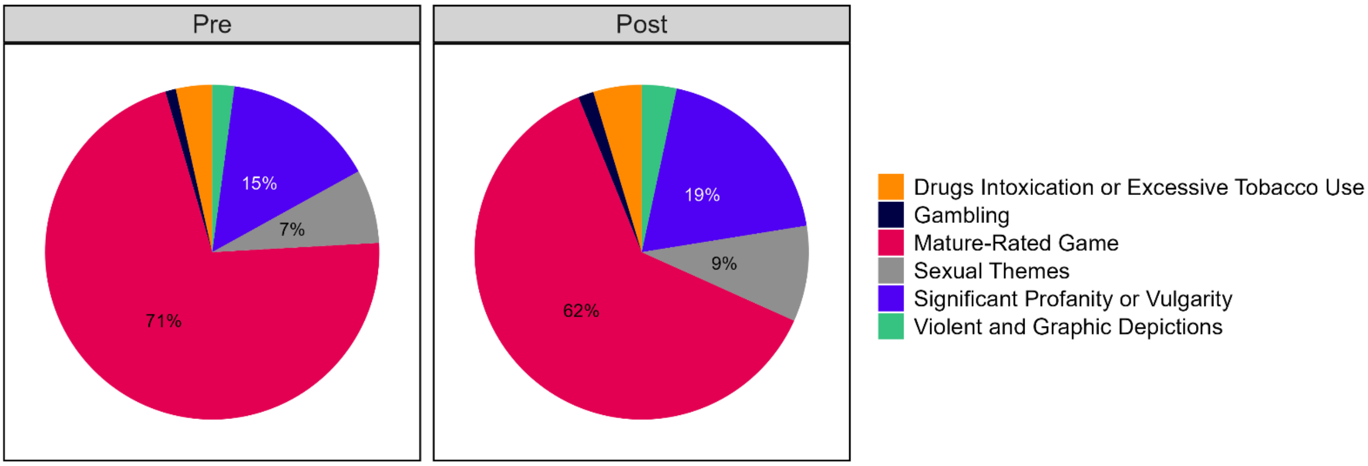


*Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

## Accuracy analysis

A4.4 We also use the lower probability threshold to assess the robustness of our accuracy analysis. As shown in Table 10, the results remain identical to our baseline findings, with mature gaming and gambling streams showing the lowest levels of labelling accuracy in the pre-treatment period. Additionally, we continue to observe similar levels of labelling inaccuracy prior to the CCL, with 37% to 39% of streams classified as non-mature being inaccurately labelled as mature.

**Table 10: Accuracy analysis using a 75% threshold**

| (a) Predicted mature stream type | Streams labelled as mature (%) | (b) Predicted non-mature stream type | Streams labelled as mature (%) |
|---|---|---|---|
| **Mature-Rated Game** | | **No Mature-Rated Game** | |
| *Pre-CCL* | 48.21% | *Pre-CCL* | 37.11% |
| **Gambling** | | **No Gambling** | |
| *Pre-CCL* | 63.24% | *Pre-CCL* | 39.54% |
| **Significant Profanity and Vulgarity** | | **No Significant Profanity or Vulgarity** | |
| *Pre-CCL* | 81.99% | *Pre-CCL* | 37.53% |
| **Sexual Themes** | | **No Sexual Themes** | |
| *Pre-CCL* | 82.81% | *Pre-CCL* | 38.62% |
| **Drugs, Intoxication, or Excessive Tobacco Use** | | **No Drugs, Intoxication, or Excessive Tobacco Use** | |
| *Pre-CCL* | 80.03% | *Pre-CCL* | 39.16% |
| **Violent and Graphic Depictions** | | **No Violent and Graphic Depictions** | |
| *Pre-CCL* | 81.90% | *Pre-CCL* | 39.33% |

*Source: Analysis of Twitch 'Get Streams' API and 'Get Channel Information' API data March and December 2023*

## Probability of mature content labelling

A4.5 Table 11 shows the results to our RDD model using a 75% probability threshold. In line with our baseline analysis, we find a positive and significant effect of CCLs on the probability of labelling mature content, with the strongest effect observed among gaming streams. Moreover, while the impact on gambling streams is relatively lower compared to our main set of results, it remains significant at the 10% level.

A4.6 Moreover, similar to our baseline analysis in Table 3, our results are likely affected by measurement error, and so, represent only a lower bound estimate of the true impact of CCLs on mature content labelling.

**Table 11: Probability of labelling mature content by stream type - RDD analysis using 75% threshold**

| | (1) Mature-Rated Game | (2) Gambling | (3) Significant Profanity or Vulgarity | (4) Violent and Graphic Depictions | (5) Sexual Themes | (6) Drugs, Intoxication, or Excessive Tobacco Use |
|---|---|---|---|---|---|---|
| **CCL introduction** | 0.48***<br>(0.009) | 0.13*<br>(0.069) | 0.11***<br>(0.015) | 0.11***<br>(0.041) | 0.14***<br>(0.021) | 0.11***<br>(0.030) |
| **Optimal BW** | 18.20 | 32.44 | 32.07 | 32.97 | 37.29 | 43.80 |
| **Observations** | 738297 | 14364 | 182623 | 30600 | 88071 | 44711 |

*Notes: Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01. We report bias-corrected estimates that account for the common misspecification bias found in RDD models and robust standard errors clustered at the* *content creator-level. The results for streams with a mature-rated game are based on a 20 June cut-off, while the analysis of mature non-gaming streams relies on a 20 July cut-off. Source: Ofcom analysis*

# Content creator behaviour

## Graphical analysis

A4.7 We also explore whether a lower 75% probability threshold leads to a significant change in the weekly volume of mature content streamed.

A4.8 Figure 28 plots the weekly volume of mature and non-mature streams generated by top consecutive content creators across all streaming categories. The pattern in weekly volume of mature content streamed under the 75% probability threshold is almost identical to the 95% probability threshold. This is because streams predicted to contain a 'Mature-Rated Game' CCL are less sensitive to a lowering in the probability threshold, since the majority of these streams were classified as mature with high confidence due to having a larger sample of training data.

**Figure 28: Weekly number of mature and non-mature streams by top consecutive content creators for all stream categories**

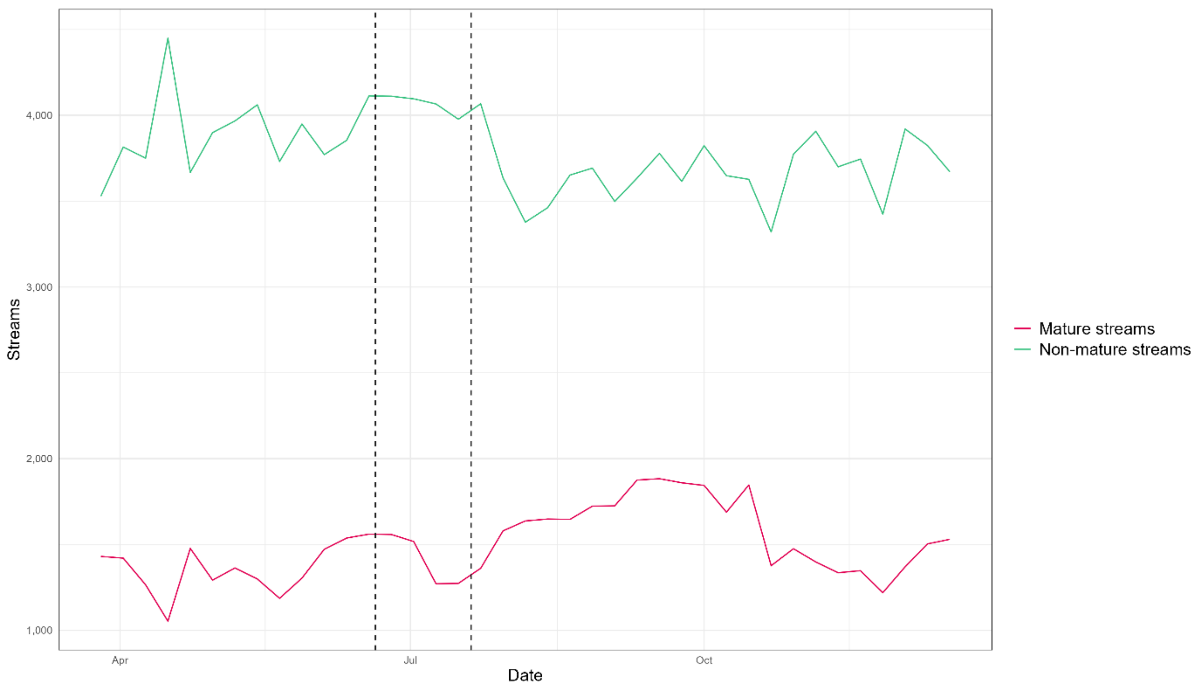


*Notes: The first vertical line represents the introduction of CCLs and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

A4.9 Figure 29 plots the weekly volume of mature and non-mature streams generated by top consecutive content creators across the non-gaming category. Relative to the 95% probability threshold, there appears to be slight uplift in the volume of mature content streamed following the intervention.

**Figure 29: Weekly number of mature and non-mature non-gaming streams by top consecutive content creators for all stream categories**


*Notes: The first vertical line represents the introduction of CCLs and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

## Econometric analysis

A4.9 Table 12 shows the results to the analysis on weekly volume of mature content using a 75% threshold. Consistent with the graphical analysis above, our fixed effects analysis suggests that the number of mature gaming and non-gaming streams on the platform increased following the intervention. Nonetheless, similar to the analysis using a 95% threshold, the magnitude of these effects remains small.

**Table 12: Weekly number of mature streams by stream type - fixed effects analysis using a 75% threshold**

| | (1) Mature gaming streams (6-weeks) | (2) Mature gaming streams (8-weeks) | (3) Mature gaming streams (12-weeks) | (4) Mature non-gaming streams (6-weeks) | (5) Mature non-gaming streams (8-weeks) | (6) Mature non-gaming streams (12-weeks) |
|---|---|---|---|---|---|---|
| **CCL introduction** | 0.21***<br>(0.002) | 0.20***<br>(0.029) | 0.07***<br>(0.017) | 0.17**<br>(0.053) | 0.16***<br>(0.036) | 0.13***<br>(0.036) |
| **Content creator FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Time FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Observations** | 48,576 | 58,334 | 71,360 | 17,208 | 14,520 | 16,640 |

***Notes:*** *Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01. Standard errors are clustered at the content creator-level. Our sample consists of a balanced panel of content creators and excludes the 30-day grace period provided by Twitch. Source: Ofcom analysis*

# Viewership of mature content

## Graphical analysis

A4.10 We also explored the effect of the CCL introduction on the absolute number of weekly views to mature and non-mature streams.

A4.11 Figure 30 and Figure 31 show the absolute number views to mature streams for all streaming categories and the non-gaming category. The pattern in views for plots is consistent with the findings under the 95% probability threshold. However, the level of mature views for both plots across both time periods appears to be higher.

**Figure 30: Weekly absolute views to mature and non-mature content for all streaming categories**


*Notes: The first vertical line represents the introduction of CCLs and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

**Figure 31: Weekly absolute views to mature and non-mature streams for the non-gaming category**

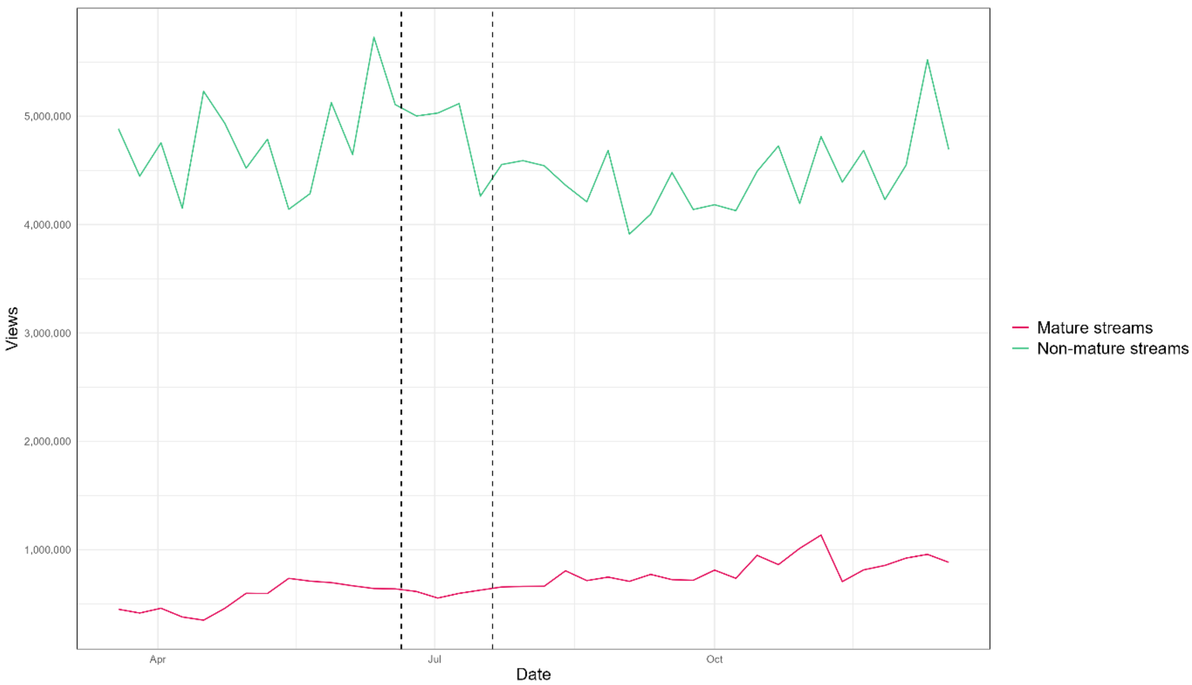


*Notes: The first vertical line represents the introduction of CCLs and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

A4.12 Figure 32 shows the views to streams predicted to contain a 'Mature-Rated Game' CCL. The pattern of views is similar to the 95% probability threshold; however, the level of views is higher.

**Figure 32: Weekly absolute views to streams predicted to contain a 'Mature-Rated Game' CCL**

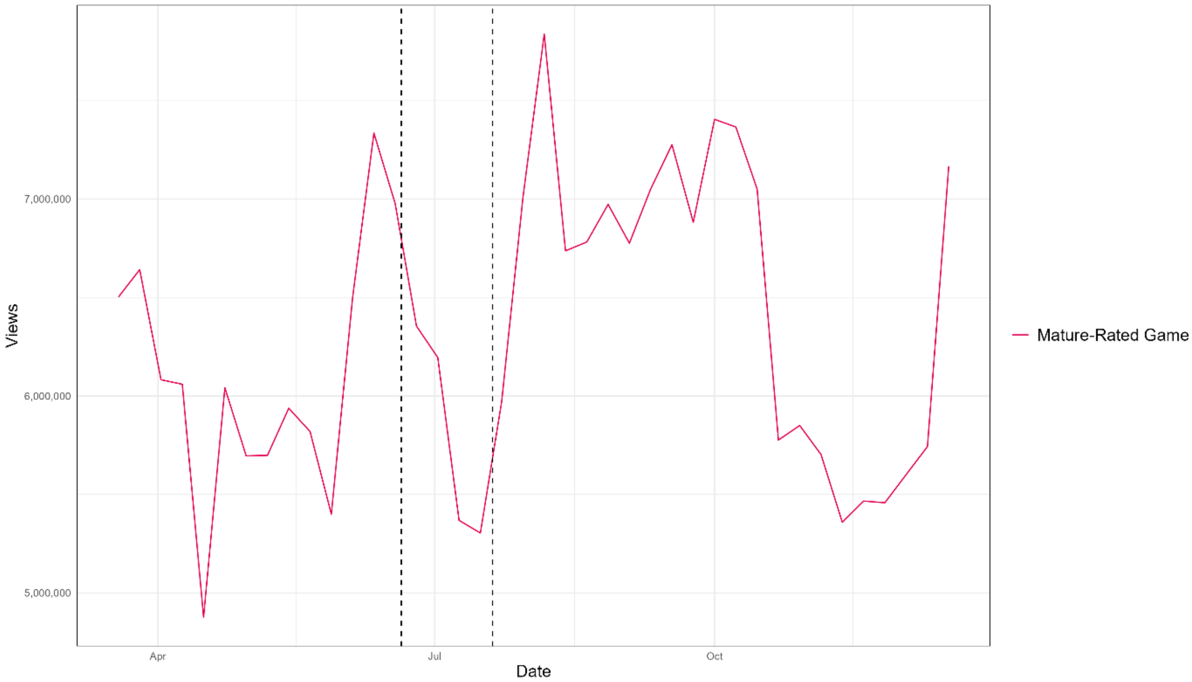


*Notes: The first vertical line represents the introduction of CCLs and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

A4.13 Figure 33 shows the views to mature streams in the non-gaming category. The pattern in views is consistent with the findings under the 95% probability threshold, however the level of views is higher.

**Figure 33: Weekly absolute views to streams predicted to contain: Drugs/Intoxication; Gambling; Profanity & Vulgarity; Sexually Suggestive and Violent/Graphic content**

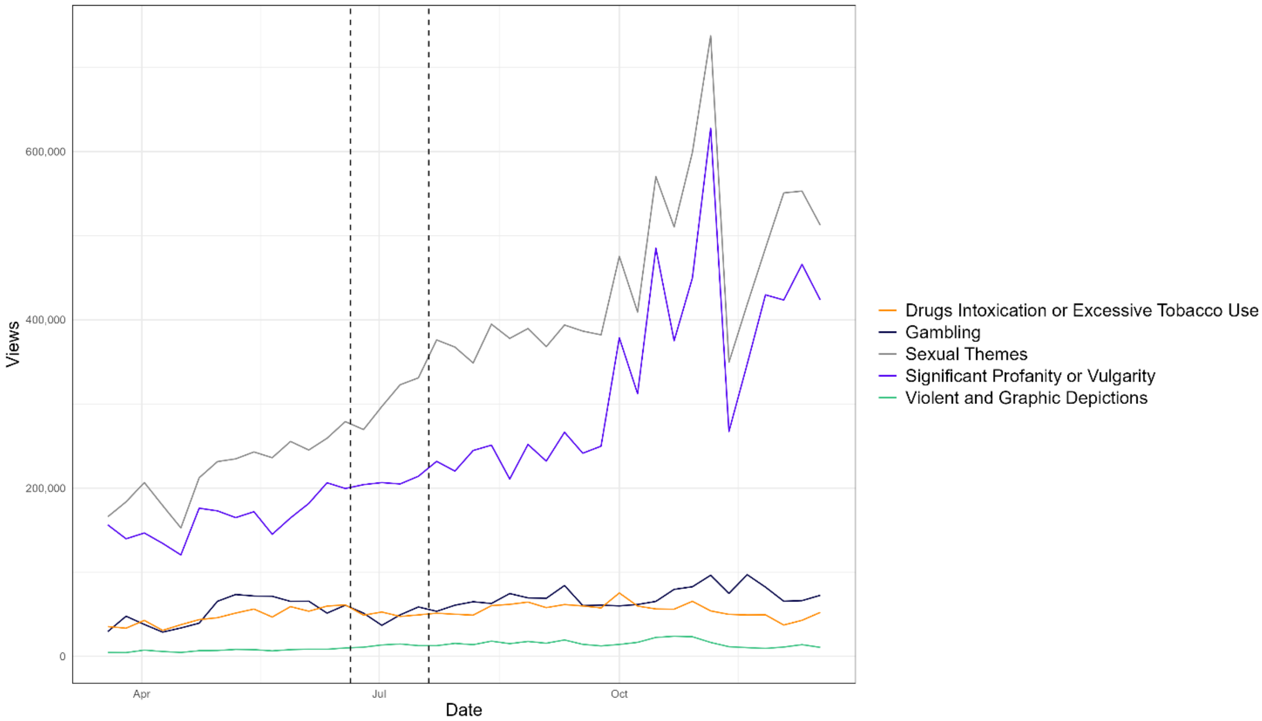


*Notes: The first vertical line represents the introduction of CCLs and the beginning of the 30-day grace period. The second vertical line captures the end of the grace period. Source: Twitch 'Get Streams' and 'Get Channel Information' API data March and December 2023*

## Econometric analysis

A4.14 Finally, Table 13 presents the results on viewership of mature content using a 75% threshold. Focusing on mature games, our sensitivity analysis indicates that the weekly views to these streams increased following the introduction of CCLs. However, similar to our baseline estimates, the effect is only significant when using a 6-week bandwidth. Furthermore, unlike the analysis using a 95% threshold, we find that the weekly number of views to mature non-gaming streams increased following the intervention. Interestingly, this effect is only significant when using a 12-week bandwidth and its consistent with the increase in views to streams with sexual themes and profanity observed in the graphical analysis above.

**Table 13: Weekly views to mature streams by stream type - fixed effects analysis using a 75% threshold**

| | (1) Views to mature gaming streams (6-weeks) | (2) Views to mature gaming streams (8-weeks) | (3) Views to mature gaming streams (12-weeks) | (4) Views to mature non-gaming streams (6-weeks) | (5) Views to mature non-gaming streams (8-weeks) | (6) Views to mature non-gaming streams (12-weeks) |
|---|---|---|---|---|---|---|
| **CCL introduction** | 475.76** (173.652) | 35.56 (125.618) | -9.08 (251.902) | 136.469 (234.358) | 197.99 (227.421) | 928.93* (507.301) |
| **Content creator FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Time FEs** | Yes | Yes | Yes | Yes | Yes | Yes |
| **Observations** | 13384 | 15408 | 17160 | 1448 | 1620 | 1520 |

*Notes: Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01. Standard errors are clustered at the content creator-level. Our sample consists of a balanced panel of content creators and excludes the 30-day grace period provided by Twitch. Source: Ofcom analysis*