

COPPAcalypse? The YouTube Settlement’s Impact on Kids

Content

Garrett A. Johnson, Tesary Lin, James C. Cooper, & Liang Zhong*

March 14, 2024

Abstract

We examine the tradeoff between privacy and personalization for online content by evaluating the impact of YouTube’s settlement with the Federal Trade Commission over violating the Children’s Online Privacy Protection Act (COPPA). Under the settlement, YouTube removed all forms of personalization for child-directed content starting in January 2020, which included personalized ads and platform features like personalized search and recommendations. We study the resulting impact on 5,066 top American YouTube channels by comparing the child-directed content creators to their non-child-directed counterparts using a difference-in-differences design. On the supply side, we find that child-directed content creators produce 18% less content and pivot towards producing non-child-directed content. Child-directed content creators also invest less in content quality: the proportion of original content falls by 11% and manual captioning falls by 27%, while user content ratings fall by 10%. On the demand side, views of child-directed channels fall by 20%. Consistent with the platform’s degraded capacity to match viewers to content, both content creation and content views become more concentrated among top child-directed YouTube channels.

Keywords: privacy, digital content, COPPA, YouTube

*Johnson: Questrom School of Business, Boston University <garjoh@bu.edu>; Lin: Questrom School of Business, Boston University <tesary@bu.edu>; Cooper: Antonin Scalia Law School, George Mason University <jcoope20@gmu.edu>, Zhong: Boston University: <samzl@bu.edu>. We gratefully acknowledge the financial support of University of Pennsylvania’s CTIC & Warren Center “Economics of Digital Services” initiative, Questrom’s Digital Business Institute at Boston University, and the Program on Economics & Privacy at George Mason University. We thank seminar participants at Columbia Business School, the Economics of Digital Services Research Symposium, and Simon Business School for their helpful comments. We thank Yansong Bai, Zhengrong Gu, Kushi Kiran, Yuqing Yang, and Kradon Zhao for their research assistance. We thank Jason Urgo at Social Blade and the YouTube Researcher Program for data assistance.

1 Introduction

Personalization plays a key role in online content platforms like YouTube. Personalized search results and recommendations help match users to content. Users can improve their experience by curating playlists and interacting with creators by commenting. Ad-funded platforms use personalized advertising to improve monetization. Past research finds that personalization increases ad revenue (Johnson et al., 2020; Ravichandran & Korula, 2019), and that ad revenue increases the supply of content (Shiller et al., 2018; Sun & Zhu, 2013). At the same time, personalization raises user privacy concerns, as it relies on persistent identifiers such as cookies. This suggests a privacy-for-content tradeoff, though the evidence for this intuitive relationship is mixed (Lefrere et al., 2022; Kircher & Foerderer, forthcoming).

Privacy regulations, like the General Data Protection Regulation (GDPR) in Europe, may reduce ad-funded content creation by restricting the use of persistent personal identifiers. In particular, regulators place a premium on children’s privacy: the Children’s Online Privacy Protection Act (COPPA) in the United States restricts the collection of data on children under 13; the EU’s Digital Services Act bans personalized advertising that targets children under 18; and American President Joe Biden has repeatedly called for a similar ban.¹ Nevertheless, these well-intentioned policies may negatively affect child-directed content creators and their audiences. YouTube’s 2019 COPPA settlement led the platform to deactivate all forms of personalization for child-directed content on its platform. We use this unique natural experiment to evaluate the consequences of strict privacy regulation for online content. We find the YouTube settlement reduced the supply of made-for-kids channel content by 18% and content views by 20%.

In September 2019, YouTube entered a consent decree with the Federal Trade Commission (FTC) to settle charges that it had collected persistent identifiers from children without explicit parental consent. YouTube’s parent company Alphabet paid \$170 million—by far the largest amount under COPPA at the time—and YouTube agreed to stop collecting personal information from made-for-kids (MFK) content viewers. After January 1, 2020, YouTube deactivated all forms of MFK-related personalization globally, including personalized searches, recommendations, playlists, and commenting functions for MFK content. Further, ads served with MFK con-

¹See <https://www.whitehouse.gov/briefing-room/speeches-remarks/2023/02/07/remarks-of-president-joe-biden-state-of-the-union-address-as-prepared-for-delivery>

tent could only target viewers based on context, which was expected to reduce ad prices. YouTube creators and viewers were concerned these policies would harm MFK content on the platform, and some termed this change the “COPPAcalypse.”

We study the impact of the YouTube settlement using a sample of 5,066 YouTube channels from July 2018 to December 2020. These channels are drawn from content creators in the United States within the top 100,000 global YouTube channels by subscription count. Our sample focuses on three content categories that represent about 80% of MFK content globally: education, entertainment, and film & animation. We collect rich video-level data including release date, MFK status, and transcripts (where available) on 1.8 million videos uploaded by these channels during our sample period. We also collect historical data on weekly views and subscriptions at the channel level. Our analysis divides channels into three groups based on their content types prior to the September 2019 announcement: MFK, non-MFK, and mixed content channels. Because the COPPA settlement only applies to MFK content, we use non-MFK channels as a control group within a difference-in-differences framework to identify the causal effect of these privacy measures on YouTube.

The settlement removed multiple elements at once, so we cannot distinguish between the impact of removing personalized advertising and of removing platform personalization in general. We believe our results are broadly applicable to settings where privacy regulation removes both elements at once such as can arise under Europe’s GDPR. Moreover, the FTC proposed in 2023 to strengthen COPPA by further limiting firms’ use of personalization to increase user engagement, which resembles YouTube’s strict approach to COPPA. Nevertheless, YouTube creators emphasized in 2019 that the loss of personalized advertising would most impact their business.² In particular, some YouTube creators that experimented with deactivating ads personalization saw large reductions in their ad prices between 60% and 90% (Katz & Fener, 2019).

First, we find a negative impact of YouTube settlement on MFK content production. MFK and mixed channels release 18% and 16% respectively fewer videos on average after January 2020. Moreover, MFK channels reduce their share of MFK videos by 2.7 percentage points, whereas mixed channels reduce their share of MFK videos by 36% (9.5 percentage points). Our findings are

²See comments to the FTC from Sockeye Media LLC, Skyship Entertainment, and J House Vlogs available at <https://www.regulations.gov/comment/FTC-2019-0054-117108>, <https://www.regulations.gov/comment/FTC-2019-0054-111477>, and <https://www.regulations.gov/comment/FTC-2019-0054-116904>.

consistent with lower ad revenue—i.e., lower ad prices potentially combined with fewer content views—moving content creation down the supply curve for YouTube creators. Mixed content channels are better able to pivot away from MFK content as ad prices for non-MFK content become relatively more lucrative. We also find that smaller channels make deeper cuts to their content supply.

Second, we find that the settlement reduces the quality of MFK content. We consider two vertical content attributes—the shares of original content and manual captioning—as well as user ratings as a subjective quality metric. All else equal, users benefit from having more original content on the platform, because this increases the possibility that they find content suited to their taste (Waldfogel, 2017). However, creators may cut costs by reusing old content, for instance, by making compilation videos. We measure duplicate content by identifying common passages in video transcripts with a channel’s prior released content. We find that the settlement reduces the original content share by 11% (7.7 percentage points) for MFK channels. This implies that our video release estimates understate the settlement’s impact on *original* content output by about a third. Video captions improves accessibility for users with hearing loss and can help children learn reading and vocabulary (Bird & Williams, 2002; Kothari & Bandyopadhyay, 2014; Linebarger, 2001). Creators can rely on YouTube’s free automatic captioning, which is prone to transcription errors, or creators can invest in manual captions to reduce transcript inaccuracies. The settlement lowers MFK channels’ manual captioning share by 27% (3.8 percentage points). Lastly, we find that the settlement reduces user content ratings by 10%. The decline in user ratings can reflect both creator’s reduced investment in content quality and YouTube’s degraded capacity to match users to content.

Third, the drop in content supply coincides with a reduction in content views and subscriptions. Relative to non-MFK channels, MFK channel views fall by 20% and mixed channel views fall by 13%. Moreover, new channel subscriptions fall by 25% for MFK channels and 24% for mixed channels. Consistent with our content creation results, we see larger relative reductions in both views and subscriptions among channels with fewer baseline subscribers. The asymmetric impact is more pronounced on the demand side, such that we see no significant change in views for top-quartile channels. This result is consistent with the effect of deactivating platform personalization: the platform can no longer pair users with long-tail content that matches their interests.

These results suggest that consumers did not find sufficient substitute content (new or existing) within YouTube’s MFK category to compensate for the post-settlement reduction in MFK content.

Our paper contributes to several strands of literature. First, by examining the tradeoff between content and privacy, our paper builds on prior work by Lefrere et al. (2022), Kircher & Foerderer (2023a), and Hui et al. (2023). Lefrere et al. examine the consequences of Europe’s General Data Protection Regulation (GDPR) for websites and find no impact on content provision. However, their result may simply reflect low regulatory compliance, as European websites made modest and transitory adjustments that fell short of the GDPR’s privacy requirements (Johnson et al., 2023; Lefrere et al., 2022; Peukert et al., 2022). In contrast, we see that regulation reduces content creation once YouTube imposed strict privacy compliance on its made-for-kids content. Another strand of literature examines how platform restrictions on personalized advertising affect mobile apps. Following a COPPA-related 2019 ban on targeted advertising, children’s mobile application developers invested less in their apps (Kircher & Foerderer, 2023a) and increased prices (Hui et al., 2023). Others find that Apple’s restrictions on personalized ads on its iOS platform imposed similar effects on apps (Cheyre et al., 2023; Kesler, 2023; Li & Tsai, 2023) and reduced ad prices (Cecere & Lemaire, 2023). These papers focus on supply-side responses,³ whereas we also examine demand-side responses for digital content.

We also contribute to the media economics literature by examining the role of personalized ads as a revenue source. Existing research shows that advertising supports content output (Shiller et al., 2018) and can affect content diversity (Sun & Zhu, 2013; Kerkhof, 2020). A growing literature finds that personalized ads generate greater revenue (Alcobendas et al., 2023; Cecere & Lemaire, 2023; Johnson et al., 2020; Laub et al., 2022; Ravichandran & Korula, 2019). We instead empirically examine how the loss of personalized ads affects both content output and quality.

Beyond eliminating personalized ads, the YouTube settlement also removed platform personalization. This change decreases the match quality between users and content with differential consequences for creators of different sizes. Several recent studies (see e.g., Sun et al., 2022; Donnelly et al., 2022; Korganbekova & Zuber, 2023) confirm that platform personalization generally favors smaller, more niche sellers. However, the impact of personalized ads on different-sized content creators is less understood. Bhargava (2022) theorizes that a platform that improves ad

³However, Li & Tsai (2023) also find negative demand effects for app downloads.

targeting technology would instead increase the concentration among its creators, because larger and more productive creators can better leverage the resulting higher ad prices. Given the opposing effects of eliminating platform personalization and personalized ads, the net effect on creator concentration is unclear. We show that both content production and consumption become more concentrated after the settlement, suggesting an important role for platform personalization.

Several scholars have studied YouTube. Early work modeled the supply (Tang et al., 2011, 2012) and propagation (Yoganarasimhan, 2012) of content. More recent work highlights secondary creator revenue streams like subscriptions (Panjwani & Xiong, 2023), affiliate marketing (Mathur et al., 2018), and influencer marketing (Li et al., 2023). Abou El-Komboz et al. (2022) is similar to ours in that they examine a platform change that raised the size threshold for creators to receive revenue-sharing from YouTube. This change also led to creator exit and reduced video production, though creators could grow their way out of this restriction, unlike made-for-kids content creators. Kerkhof (2020) finds that ad funding increases content differentiation on YouTube.

Our study connects to the broader economics literature on privacy (Acquisti et al., 2016; Goldfarb & Que, 2023) and the economic impact of privacy regulation in particular. Recent review articles consider the impact of health privacy regulation (Miller, 2022) as well as the GDPR (Johnson, 2022). However, the economic impact of COPPA has garnered less attention until the law’s recent application to YouTube and to mobile apps (Hui et al., 2023; Kircher & Foerderer, 2023a). In a law article, Beemsterboer (2020) provides additional background on the YouTube COPPA settlement and hypothesizes that it would hurt content creation. Our paper is closest to subsequent work by Kircher & Foerderer (2023b) who estimate the settlement’s impact on 1,676 educational channels from several English-speaking countries. Kircher & Foerderer (2023b) find a similar impact on the supply of MFK content, providing more evidence that our findings generalize outside the US. On the demand side, their study finds that subscriptions rise for less popular channels and fall for the rest. We find the opposite however, potentially because we consider all channels, whereas Kircher & Foerderer (2023b) exclude those that exit. Our results are more general in that we consider multiple content categories, separately examine mixed and pure-MFK channels, and consider multiple content quality measures. In particular, our content originality finding reveals that the settlement’s impact on content alone understates the policy’s full impact by about a third, because MFK channels also produce less original content.

The remainder of this paper is organized as follows. Section 2 provides the regulatory background and the main policy change that enables our identification strategy. Section 3 describes our data and provides descriptive evidence. Section 4 presents our empirical approach. Section 5 provides our supply-side, content quality, and demand-side results. Section 6 discusses the welfare implications of our findings, and Section 7 concludes.

2 Background

Below, we provide some general background on YouTube as well as its COPPA-related settlement with the FTC.

2.1 YouTube

YouTube is the most important platform for free video content. With over 2 billion monthly logged-in users and localized versions in over 100 countries in 80 languages (YouTube Official Blog, 2023), YouTube is the second most popular website globally.⁴

YouTube is popular with children. 76% of American children ages 8-12 watch YouTube according to the non-profit Common Sense (Rideout & Robb, 2019), 53% report that they watch YouTube most often among online video platforms, and more than 80% of UK children ages 8-15 prefer YouTube to TV (PricewaterhouseCoopers, 2019). Reflecting its popularity among children, YouTube captured 23% of the global spending on child-directed digital ads in 2019 (PricewaterhouseCoopers, 2019). YouTube introduced a dedicated platform called YouTube Kids in 2015, which offers curated child-safe content free of personalized advertising. YouTube selects content for YouTube Kids from the universe of YouTube channels: i.e., creators do not determine this themselves. Still, children and especially older children seem to prefer YouTube to YouTube Kids.⁵

Independent creators generate YouTube’s content by uploading videos on their channels. Creators can monetize their content by allowing Google to serve display advertisements. These ads can take the form of video ads that run before or during the content, as well as banner ads next to

⁴<https://www.similarweb.com/top-websites/>, accessed on February 22, 2023.

⁵While YouTube Kids is more popular among younger children (39% of U.K. kids ages 5-7 use it exclusively: PricewaterhouseCoopers, 2019), children ages 8-12 prefer regular YouTube (only 23% have ever used YouTube Kids, and only 7% report using it the most: Rideout & Robb, 2019).

or below the video. YouTube shares a portion of ad revenue with creators whose channels have crossed certain size thresholds.

Creators of child-directed content play a key role on YouTube. In 2020, four of the top 10 YouTube channels by video views were directed at children: Cocomelon, Like Nastya, Kids Diana Show, and Ryan’s World.⁶ The largest MFK channels generate significant revenue: Ryan Kaji (age 8) was the highest earning YouTuber with an estimated \$26 million in 2019.⁷ Among the 37.9 million channels tracked by Social Blade, YouTube had over 170,000 channels classified as fully MFK as of January 2020 (Urgo, 2020). Collectively, these MFK channels represent over 6 billion subscribers and 2.6 trillion views (Urgo, 2020).

2.2 YouTube COPPA Settlement

The Children’s Online Privacy Protection Act (COPPA) prohibits operators of online services directed at children under 13 from collecting personal information without obtaining verifiable parental consent.⁸ In 2013, the FTC amended the COPPA rules so that the definition of personal information includes “persistent identifier that can be used to recognize a user over time and across different Web sites or online services,” such as a “customer number held in a cookie... or unique device identifier.”⁹ Because obtaining verifiable parental consent for free online services is difficult and rarely cost justified, COPPA acts as a *de facto* ban on the collection of personal information—and hence personalized advertising—by providers of free child-directed content.¹⁰

On September 4, 2019, YouTube entered a consent agreement with the FTC to settle charges that it had violated COPPA. The FTC’s allegations focused on YouTube’s practice of serving personalized advertising on child-directed content without obtaining verifiable parental consent. Although YouTube is a general audience website and users must be at least 13 years old to obtain a Google account ID (which makes personalized advertising possible), the FTC’s complaint alleged that YouTube knew many of its channels were popular among children under 13, citing YouTube’s

⁶<https://socialblade.com/youtube/top/100/mostviewed>.

⁷<https://www.forbes.com/sites/maddieberg/2019/12/18/the-highest-paid-youtube-stars-of-2019-the-kids-are-killing-it>.

⁸5 U.S.C. § 6501 et seq.; 16 C.F.R. § 312 et seq.

⁹16 C.F.R. § 312.2.

¹⁰Verifying parental consent for paid content is more straightforward, as a credit card payment signifies consent.

own claims to advertisers.¹¹ YouTube’s parent company Alphabet agreed to pay a civil penalty of \$170M, the largest amount under COPPA until 2023. As of 2023, this represents the FTC’s 10th highest penalty amount since 2000.¹²

As part of the settlement, YouTube also agreed to identify child-directed content and to stop collecting personal information from MFK content viewers. Beginning January 1, 2020 (hereinafter, the *post-settlement period*), YouTube required channel owners producing MFK content to designate either their entire channel or specific videos on their channel as MFK. YouTube augmented these self-designations with an automated classifier to identify content directed at children. The automated classifier looks for content that includes, for instance, child actors, child characters, games, toys, songs, and stories that children like.¹³ Between the announcement and implementation of the settlement (hereinafter, the *announcement period*), YouTube provided content creators details on its compliance plan.¹⁴ The FTC also provided legal guidance, explaining that MFK content creators could face civil penalties of up to \$42,530 per video if they fail to self-designate and YouTube’s classifier fails to correctly identify their content as MFK.¹⁵

Also beginning January 1, 2020, YouTube disabled all personalization elements to implement the settlement’s requirement to refrain from collecting children’s personal information¹⁶ and banned personalized advertising for MFK content on its platform. YouTube’s new MFK policies were broad in scope: they covered content creators worldwide, and include both new and existing

¹¹See *FTC v. Google, LLC*, at ¶¶ 49-50, Case No.: 1:19-cv-2642 (Dist. D.C. Sept. 4, 2019).

¹²https://violationtracker.goodjobsfirst.org/prog.php?agency_sum=FTC.

¹³<https://www.theverge.com/2019/11/13/20963459/youtube-google-coppa-ftc-fine-settlement-youtubers-new-rules>; <https://www.youtube.com/watch?v=KdIlQ9kq4F4>.

¹⁴See <https://www.youtube.com/watch?v=-JzXiSkoFKw>.

¹⁵<https://www.ftc.gov/news-events/blogs/business-blog/2019/11/youtube-channel-owners-your-content-directed-children>.

¹⁶These personalization elements include: video autoplay on home, video cards or end screens, video watermarks, channel memberships, comments, donate buttons, likes and dislikes on YouTube Music, live chat or live chat donations, merchandise and ticketing, notification bell, playback in the miniplayer, super chat or super stickers, save to playlist and save to watch later, posts, and stories. See <https://support.google.com/youtube/answer/9527654>. We note that the complete elimination of personalized content may go beyond what COPPA requires. 16 CFR § 314 exempts an operator from obtaining parental consent when it “collects a persistent identifier and no other personal information, and such identifier is used for the sole purpose of providing support for the internal operations of the Web site or online service.” Further, 16 CFR § 312.2 defines “internal operations” as including “personalizing content”, “so long as the information collected is not used or disclosed to contact a specific individual, including through behavioral advertising, to amass a profile on a specific individual, or for any other purpose.” Nevertheless, the FTC’s 2023 proposed COPPA rule changes (see <https://www.ftc.gov/news-events/news/press-releases/2023/12/ftc-proposes-strengthening-childrens-privacy-rule-further-limit-companies-ability-monetize-childrens>) would limit the “internal operations” exception by requiring operators to disclose related data-use purposes and by prohibiting the use of personal information in ways that would “encourage usage.” Thus, even if YouTube’s response went beyond what COPPA required in 2020, it potentially mirrors requirements under COPPA that may go into effect before the end of 2024.

MFK content. YouTube applied COPPA at the content level rather than the user level, so these policies affect all MFK content regardless of whether the viewer is an adult or a child.

Without personalized advertising, creators expected ad prices to fall significantly. Creators expressed greatest concern for the resulting loss in ad revenue and its corresponding business impact.¹⁷ Personalized advertising generates value for advertisers because it helps them to target, measure, and optimize ad effectiveness. YouTube’s ad price data are not public, so we do not know the extent of the price drop. However, some channels found that ad revenue fell by 60-90% when they experimented with deactivating personalized ads during the announcement period (Katz & Fener, 2019). Consistent with this anecdote, we obtained data from one YouTube MFK creator showing that their ad prices fell 73% after the YouTube settlement. For comparison, Ravichandran & Korula (2019) and Johnson et al. (2020) both find that ad prices fell 52% without third-party cookie identifiers on the open web.

Advertisers who serve ads on MFK content can still target ads based on context, such as the video topic (as inferred by YouTube’s content labels) or specific channels.¹⁸ The FTC’s COPPA rule allows the use of persistent identifiers for “support for the internal operations of the Web site or online service,” which includes serving contextual ads and ad frequency capping.¹⁹ However, advertisers could not deploy behavioral targeting on MFK content.²⁰ Advertisers cannot target consumer segments, based on their own first-party data, which precludes retargeted advertising. Nor can they use YouTube’s user behavioral segment information, including detailed demographics, in-market users, affinity, or custom segments.²¹ Google cautions advertisers that target children or MFK content against collecting data using “third-party trackers,” which limits advertisers’ ability to measure ad conversions.²² We discuss online ads on YouTube’s MFK content in greater detail in Appendix G, which includes evidence from our own ad survey.

YouTube also rebuilt its MFK-related video search and recommendation engine to exclude personal data. Research in e-commerce settings suggests that personalized recommendations improve sales—particularly for niche sellers (Sun et al., 2022)—and increase consumer surplus

¹⁷See creator comments to the FTC in footnote 2.

¹⁸<https://support.google.com/google-ads/answer/2497832>; <https://support.google.com/google-ads/answer/2470108>.

¹⁹16 C.F.R. §312.2.

²⁰<https://support.google.com/adspolicy/answer/9683742>.

²¹<https://support.google.com/google-ads/answer/2497941>.

²²<https://support.google.com/adspolicy/answer/9683742>.

(Donnelly et al., 2022). Korganbekova & Zuber (2023) also find that personalized search results improve e-commerce sales. YouTube deactivated many features that help content creators engage their audience but rely on personal data, including end screens (a YouTube feature for channels to promote other content at the end of a video), subscriber notifications (e.g., new content alerts), and adding videos to playlists.²³

YouTube creators and users were alarmed by these changes. Multiple related online petitions collected over a million signatures in total.²⁴ The FTC received 119 thousand public comments on its 2019 review of COPPA, and 71% of these referenced “YouTube”. Then-FTC Commissioner Noah Phillips (2019) worried the “order may reduce content creators’ incentives to develop child-directed programming” so that children may “see less and lower quality content.” Anecdotal evidence supports this concern: popular MFK channels claimed they would cut back on content creation, pivot towards adult-oriented content, or cease content creation altogether.²⁵ We consider large-scale descriptive evidence for these claims in the next section.

3 Data

3.1 Sample

Our sample consists of 5,066 top American YouTube channels drawn from the three most common content categories for MFK videos—“film & animation”, “education,” and “entertainment” categories—which comprise about 80% of MFK videos globally. To construct our sample, we acquired a list of the top 100,000 YouTube channels by subscriber counts as of June 2021 from Social Blade, a social media analytics firm. Although this list reflects some survivor bias in that it excludes channels that dropped out of the top channel list due the settlement’s impact on subscriptions, we were unable to obtain a comparable channel list pre-2020.²⁶ We select channels

²³YouTube also deactivated commenting, which was an essential channel for creators to get feedback on their content. However, YouTube had already disabled comments on many MFK videos in 2019 to prevent predatory comments (see <https://www.theverge.com/2019/2/28/18244954/youtube-comments-minor-children-exploitation-monetization-creators>), thus attenuating the impact of disabling comments.

²⁴See <https://www.change.org/p/youtubers-and-viewers-unite-against-ftc-regulation>, and <https://www.change.org/p/youtube-reconsider-the-new-rules-regarding-children-family-videos-on-youtube>.

²⁵<https://twitter.com/SocraticaKids/status/1195156921256251392>; <https://variety.com/2020/digital/news/ftc-rules-child-directed-content-youtube-1203454167/>; and <https://www.c-span.org/video/?468062-1/childrens-online-privacy-protection-act&playEvent#!>.

²⁶Survivor bias would make our impact estimates conservative.

based in the United States that mostly produce videos in the three categories above (based on Social Blade’s channel classification), and released content between the start of our sampling period and the settlement announcement date (September 4, 2019). Applying these filters results in 5,066 YouTube channels. We collect data for these channels from July 1, 2018 to December 31, 2020.

We then classify channels based on their video composition between July 1, 2018 and the announcement date of the YouTube settlement. Following the settlement, YouTube requires that channels either indicate that their whole channel is child-directed (MFK) or classify each video’s MFK status separately. We therefore separate channels into three MFK types: 1) *MFK channels* released only MFK videos, 2) *non-MFK channels* released only non-MFK videos, and 3) *mixed channels* released a mixture of both MFK and non-MFK videos.²⁷ In our sample, 3,772 (74.5%) channels are non-MFK, 697 (13.8%) are MFK, and 597 (11.7%) are mixed. We also reclassify channels by the content category that applies to the majority of their videos prior to the settlement announcement date: 3,011 (59%) channels are majority “entertainment”, 928 (18%) channels are majority “education”, and 738 (15%) channels are majority “film & animation.”²⁸ The remaining “other” 389 channels (8%) produce a mix of content categories.

We use YouTube’s Data API to identify the 1.8 million videos that these channels released during our sample period. For each video, we extract its attributes including release date, MFK label, content category, and the cumulative number of likes and views as of July 2022. Appendix B considers generalizability to an all category, all country sample. We discuss our data gathering in greater detail in Appendix H. We use the video-level data to construct a panel at the channel-week level for our sample period. Our panel is unbalanced because some of our channels launch during the pre-announcement period and some outcome variables (e.g., MFK share) are missing in weeks that a channel releases no content. We then use the YouTube Data API to collect video caption data for English videos—when available—containing video transcripts and an indicator for whether the captions were generated manually.

²⁷Figure I.1 of Appendix I presents a histogram of the MFK content share among mixed channels. The distribution is somewhat bimodal and skewed toward zero. Examples of mixed channels include Warner Brothers and Disney, both of which have cartoons that are marked as MFK (e.g., Looney Tunes or Mickey Mouse) as well as videos aimed at those older than 12 (e.g., Marvel or Star Wars). Similarly, “FishFam,” a reality channel based on a family named Fish, has mixed content because some of the videos feature vignettes involving the young daughters in the family. Educational channels, such as “makemegenious,” include animated videos that explain scientific or historical topics, which may be marked as MFK depending on the subject’s complexity.

²⁸Social Blade’s assigns content category to each channel based on the channel’s ten most recent public videos, which may be less accurate than our classification here.

We use these variables to construct three measures of content quality—content originality, share of manual transcripts, and the like/view ratio—which we elaborate on in Section 5.2. In particular, we use the transcript data to construct a measure of original content by identifying duplicate text passages reused from the channel’s prior video releases. As detailed in E, we construct the originality measure for the subset of videos with English-language transcripts after removing the top 5% of channels by video uploads.

We also use Social Blade data to investigate the demand-side impact of the YouTube settlement on viewers. Unlike YouTube’s cumulative views data, Social Blade provides historical data on the number of total views and total subscriptions by channel and by day. Social Blade is able to offer only three years of historical data; since we collected these data in October 2022, our panel of new views and new subscriptions at the channel-week level runs from October 6th, 2019 to December 31, 2020. Note that YouTube rounds its subscription figures to three significant digits, so that our new subscription figures are less precise than new views.

Table 1: Summary Statistics for YouTube Channels

	Obs.	Mean	Conditional Means			Min.	Max.
			MFK	Mixed	Non-MFK		
<i>Creator Supply Metrics</i>							
Weekly video release	650,830	2.739	1.728	3.806	2.753	0	1,978
MFK share of release	380,270	0.163	0.987	0.267	0.002	0	1
<i>Creator Quality Metrics</i>							
Original Content (%)	251,163	0.907	0.629	0.900	0.960	0	1
Manual Caption (%)	273,982	0.140	0.117	0.131	0.145	0	1
Like/view Ratio	380,225	0.028	0.005	0.018	0.034	0	1
<i>Viewer Demand Metrics</i>							
Weekly views (in millions)	325,258	3.474	12.642	4.996	1.554	0	3,492
Weekly subscriptions (in thousands)	325,258	7.531	18.212	9.147	5.318	0	14,900

Notes: Summary statistics for the main sample of 5,066 American top YouTube channels where the unit of observation is a channel-week. The supply-side data covers the full July 1, 2018 to December 31, 2020 period whereas the demand-side data begins after October 6, 2019. All variables related to video content (e.g., MFK share and like/view ratio) omits weeks in which the channel releases no videos and represents the average across videos when the channel releases multiple videos in a week. Original content is defined as videos with an originality score over 0.5 (see Appendix E for details).

Table 1 shows summary statistics for our sample of YouTube channels, including conditional means by our three channel MFK types. Table I.1 of Appendix I provides pre-announcement and

post-settlement means by channel MFK type. On average, MFK channels release fewer videos per week (1.73) than both mixed (3.81) and non-MFK (2.75) channels. The mixed channels' full sample MFK share is 26.7% compared to 31.8% pre-announcement. For our quality metrics, MFK channels feature less original content on average (62.9%) than non-MFK channels (96.0%). MFK channels also have fewer manual captions (11.7% versus 14.5%) and a much smaller like/view ratio (0.005 versus 0.034). On the demand side, MFK have 12.6 million new weekly views on average, whereas mixed channels have 5.0 million and non-MFK have 1.6 million. MFK channels again have the most average, new subscribers per week (18.2 thousand) relative to mixed (9.1 thousand) and non-MFK (5.3 thousand).


3.2 Descriptive Evidence

We first consider descriptive evidence on the impact of the YouTube settlement on MFK content creation. Figure 1 illustrates the total number of MFK and non-MFK videos released by our sampled channels over time. Non-MFK content releases show a modest upward trend prior to January 2020, whereas MFK content releases are relatively flat over that period. In the first couple weeks of 2020, both series drop, though this seems to be a seasonal decline based on the prior year's data. Afterwards, non-MFK video releases (Figure 1a) rebound and continue along an upward trend similar to the pre-settlement period. By contrast, the MFK releases (Figure 2b) never recover, and the post-settlement level remains below the pre-settlement average through the end of 2020.

Figure 2 shows average video releases in log form broken out by the channel MFK types. MFK and non-MFK channels are similar in levels and show a similar upward trend prior to the settlement announcement, whereas mixed channels show a weak downward trend. After a brief seasonal decline that affects all channel types at the beginning of 2020, non-MFK channel releases rebound to or exceed pre-settlement levels, whereas mixed and especially MFK channels' average releases stay below pre-settlement levels.

YouTube channels also appear to shift away from releasing MFK content after the YouTube settlement. Figure 3 shows Sankey diagrams that chart the evolution of channel types between the pre-announcement and the post-settlement periods. Figure 3a shows the largest change comes from mixed channels, with 36.0% shifting to exclusively produce non-MFK content, compared to

Figure 1: Total Video Releases by Made-for-Kids Status



only 2.3% of MFK channels. Further, channel exit rates—defined as a channel releasing no videos in 2020—are higher for MFK channels (17.5%) than mixed (4.9%) and non-MFK (5.1%) channels.

Figure 3b illustrates shifts in production patterns for MFK and mixed channels only, to highlight the differences in how these channels react to the YouTube settlement. This pattern is consistent with what one may expect: mixed channels pivot away from MFK content more readily than pure MFK channels after MFK content became less lucrative. The observed reduction in MFK programming does not appear to merely result from channel owners relabeling their content as more lucrative non-MFK videos. Again, 17.5% of MFK channels exit for all of 2020. Moreover, only a small share of MFK channels become mixed channels (10.9%), despite the greater ad prices for non-MFK content.

Figure 4 shows histograms of our continuous content originality scores (see Appendix E) averaged at the channel-week level, separately for MFK and non-MFK channels. Prior to the FTC's

Figure 2: Average Log Weekly Releases by Channel by Made-for-Kids Type




Note: Video releases variable is transformed using $\ln(1 + y)$.

announcement, 89.0% of non-MFK channels have video originality scores over 90% whereas the comparable figure is half as high (44.5%) for MFK channels. This pattern is consistent with the notion that children better tolerate duplicate content.²⁹ The distribution of originality scores moves to the left for MFK channels after the settlement, whereas the non-MFK scores are more stable. This provides model-free evidence that MFK channels reduced content originality after the YouTube settlement. For ease of interpretation, our summary statistics (Table 1) and difference-in-differences analysis (Table 3) dichotomize content originality: videos are categorized as being original if their originality score exceeds 0.5.

Figure 5 shows the evolution of views and subscriptions by channel MFK type over the shorter period for which we observe demand-side outcomes: October 2019 to December 2020. Average log of weekly views for non-MFK channels is quite flat over this period except for an increase beginning in March 2020 that coincides with COVID lockdown restrictions in many countries. By contrast, the log of views falls for both MFK and mixed channel early in 2020 and—with the exception of an apparent, transient lockdown-related boost to views and subscriptions—remains below 2019 levels. Average log of subscriptions trends downward for all channel types in 2020. However, MFK and mixed channels exhibit a substantial downward level shift in 2020 relative to

²⁹See: <https://www.theatlantic.com/technology/archive/2017/07/what-youtube-reveals-about-the-toddler-minid/534765/>

Figure 3: Channel Category Evolution: Pre-announcement to Post-Settlement



Note: Sankey plots with channel MFK categorization. Within each plot, the left-axis labels classify channel MFK categories based on the pre-announcement period, whereas the right-axis labels reclassifies channels based on their 2020 (post-settlement) video releases.

non-MFK channels.


Overall, our descriptive analysis suggests that YouTube channels reduce their production of MFK videos and reduce their content originality after the YouTube settlement. Moreover, Figure 5 suggests that views and subscriptions for MFK content fall. In the next sections, we explore these relationships with greater rigor.

4 Empirical Approach

We use a difference-in-differences specification to examine the impact of the YouTube settlement. Our model separately considers the impact on two treatment groups—MFK and mixed channels—using non-MFK channels as the control group. Our model also separates the post-settlement period and the interim period between the announcement and the implementation of the YouTube settlement. We do so because channels may react to the announcement by adjusting their video production before the implementation date, which would otherwise bias our estimates of the YouTube settlement's impact.

Our preferred specification is a two-way fixed effect model as follows:

Figure 4: Before-After Comparisons of Content Originality Score Distributions




Note: The originality score here is given at the channel-week level and averages across videos if multiple are released in a week.

$$y_{it} = \lambda_1 MFK_i \cdot Announcement_t + \lambda_2 Mixed_i \cdot Announcement_t + \beta_1 MFK_i \cdot Settlement_t + \beta_2 Mixed_i \cdot Settlement_t + \theta_i + \delta_t + \epsilon_{it}, \quad (1)$$

where i denotes the channel, t denotes the week and y_{it} is the outcome of interest. The outcome variables include the log of video releases, the MFK share of video releases (conditional on releasing at least one video), and the log of new views and subscriptions.³⁰ The variables MFK_i and $Mixed_i$ are indicators equal to one if the channel i is MFK or mixed, respectively. $Announcement_t$ indicates the transition period between the settlement announcement on September 4, 2019 and

³⁰Due to observations with values of zero, our dependent variable is transformed as $\ln(y_{it} + 1)$ for video releases, views, and subscriptions.

Figure 5: Average Views and Subscriptions by Made-for-Kids Type



Note: Outcome variables are flow (not stock) variables and are transformed using $\ln(1 + y)$.

December 31, 2019. $Settlement_t$ indicates the post-settlement period, which is January 1, 2020 onward. The parameters of primary interest are β_1 and β_2 , which estimate the difference in outcomes post-settlement relative to pre-announcement. Under the parallel trends assumption (see e.g., Imbens & Rubin, 2015), β_1 and β_2 estimate the average treatment effect of the YouTube settlement's implementation on the MFK and mixed channels respectively. Furthermore, λ_1 and λ_2 estimate the impact of the settlement's announcement period. We include channel (θ_i) and week (δ_t) fixed effects to condition on channel attributes that do not vary over time and time trends common to all channels.³¹

Appendix A presents robustness checks that exploit alternative identification assumptions. A potential concern with the difference-in-differences approach is that the non-MFK channels may be indirectly treated by the settlement. In particular, some advertiser spending on personalized ads may shift from MFK content to non-MFK channels. If this phenomenon is economically important, it would inflate non-MFK productivity and lead us to overestimate the impact of the YouTube settlement. Though we can not observe counterfactual non-MFK output without this revenue spillover, we note that MFK video output in 2020 appears to continue its prior trend and does not exhibit a level shift in 2020 (see Figures 1 and 2). Still, our panel difference approach addresses the contamination concern by ignoring non-MFK channels, and instead using MFK/mixed channels from the prior year (2018-2019) as the control group. Appendix A also considers a related, triple-differences approach that further differences the non-MFK channel outcome trend from our panel

³¹The model omits variables *Announcement*, *Settlement*, *MFK*, and *Mixed* as they are subsumed by the week- and channel-fixed effects.

differences estimates. Nevertheless, we lack demand-side data from 2018 to provide comparable estimates for user views and subscriptions.

5 Results

We apply our difference-in-differences approach to estimate the impact of the YouTube settlement on our supply-side (Section 5.1), content quality (Section 5.2), and demand-side (Section 5.3) outcomes.

5.1 Supply-Side Results

The YouTube settlement prevents MFK content from showing personalized ads, which reduces ad prices: as much as 73% in the case of one MFK creator who shared data with us. We expect the reduction in ad price and revenue will reduce the supply of MFK content through two mechanisms: reducing overall video production and channels pivoting towards non-MFK content. Lowering MFK ad prices reduces the revenue gains from releasing MFK content relative to non-MFK content. As such, channels would want to pivot away from creating MFK content towards more lucrative non-MFK content. However, channels are constrained from pivoting into different content types by the channel’s expertise as well as their audience’s demands, particularly for those channels that specialize in MFK content. Channels that cannot pivot effectively will decrease their overall video production as a result.

Table 2 shows the average treatment effect on content supply from our difference-in-differences specification (equation 1). The first column shows significant reductions in the number of video releases for both MFK (-0.129) and mixed (-0.136) channels. The corresponding marginal effect estimates (see Appendix F for details) are -18.4% for MFK and -15.6% for mixed channels. The second column shows that the settlement also reduced the share of MFK content among new video releases. The coefficients show statistically significant reductions in MFK share for both MFK (-2.7 p.p.) and mixed channels (-9.5 p.p.). The mixed channel point estimate is almost four times as large as the MFK estimate, and represents a 35.6% drop relative to mixed channels’ pre-announcement average (26.7%). For both outcomes, the announcement period estimates indicate that channels started to make these changes prior to the settlement’s implementation. In par-

ticular, our anticipation period estimates are statistically significant ($p < 0.05$) and range from one-third and two-thirds as large as their corresponding post-settlement estimates. In sum, the results in Table 2 comport with our model-free evidence for channel output (Figure 2) and MFK video mix types (Figure 3).

Table 2: Difference-in-Differences Estimates

	log (Videos + 1)	Share of MFK Videos
MFK $\times$ Post-Settlement	-0.129*** (0.022)	-0.027*** (0.005)
Mixed $\times$ Post-Settlement	-0.136*** (0.024)	-0.095*** (0.010)
MFK $\times$ Anticipation Period	-0.042** (0.017)	-0.015*** (0.003)
Mixed $\times$ Anticipation Period	-0.075*** (0.022)	-0.040*** (0.007)
Week, Channel Fixed Effects	Y	Y
Marginal Effect		
MFK	-18.40%	
Mixed	-15.63%	
Adj. R^2	0.663	0.932
N	650,830	380,270

Notes: *Post-Settlement* is defined as weeks after Jan. 1, 2020. *Anticipation Period* is defined as weeks between Sep.4, 2019 to Jan.1, 2020. The sample size for share of MFK videos is smaller because it is conditional on releasing content in a given week. All specifications include week and channel fixed effects. Computation of marginal effects detailed in Appendix F. Robust standard errors clustered at the channel level are in parentheses. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

Appendix A considers the robustness of our findings to alternative identification strategies. These have the same sign pattern and our post-settlement estimates remain highly significant ($p < 0.01$) with the exception of the video release findings for mixed channels. Our panel difference estimates for the impact on MFK channel output are almost identical (-0.131, $p < 0.01$), which allays concerns that a SUTVA violation inflates our results. Our panel difference estimates for the impact on mixed channel output are more conservative (-0.052, $p < 0.1$). This addresses the potential concern that mixed channel releases violate the parallel trends assumption: Figure 2 shows that mixed channel output generally falls while non-MFK channel output grows during our sample. Still, our time-varying treatment effects analysis below suggests this differential pre-trend issue is specific to this single outcome for mixed channels alone. Our panel difference estimates for the impact on MFK share are similar to Table 2, but more conservative. Our triple difference videos release estimates are more conservative than the panel difference estimates, whereas the

triple difference MFK share estimates show somewhat larger reductions.

Appendix B considers the generalizability of our findings to YouTube more broadly. We consider the full sample of 73,354 top channels (see also Section 3 and Appendix H) that release any content during our pre-announcement period. This wider sample includes channels for all countries and all content categories. The treatment effect coefficient estimates from our global sample are close to our main sample estimates above. In particular, our estimated coefficients for MFK channels are -0.143 for video releases and -0.030 for MFK share, while those for mixed channels are -0.132 and -0.077 respectively. This suggests that our findings generalize well to the settlement's broader impact on YouTube's global platform. In related research, Kircher & Foerderer (2023b) also find that video uploads fall by 9.6% and that MFK share falls by 12.8 p.p. in their sample of educational channels from certain English-speaking countries.

These results show that channels reduce their overall content supply in response to lost ad revenue. If indeed ad prices fell by more than 70% as our anecdotal evidence suggests, then our estimates imply that the supply of MFK content is rather inelastic with respect to the price of advertising. Several theories may reconcile this finding. First, some MFK content creators may receive ad revenue from YouTube Kids, which never used personalized ads and was thus unaffected by the settlement. Second, some creators have alternative revenue streams beyond YouTube ads. For instance, Panjwani & Xiong (2023) show that some channels employ subscription-like services (e.g., via Patreon) as well as own-channel monetization strategies like affiliate marketing and sponsored content. In addition, a few big channels (e.g., Disney) earn revenue from distributing content elsewhere or selling licensed merchandise. Third, Section 5.2 shows that, because MFK channels increase their reliance on duplicate content after the settlement, our supply-side results mask a greater decline in *original* content releases.

The greater impact on MFK shares for mixed channels suggests that mixed owners could substitute into non-MFK content production more easily. This greater ease could be a function of both the channel's content expertise and the receptiveness of the channel's audience to non-MFK content. A third possibility is that mixed producers produce more content on the boundary between MFK and non-MFK content, so that the channel converts their production to non-MFK content with smaller adjustments. As an extreme example, consider a channel that moves from content directed at 12-year-olds to content directed at 13-year-olds (COPPA covers only children younger

than 13). As a fourth possibility, mixed channels that produce content on the MFK boundary may have greater discretion to label their content as non-MFK content.³² The third and fourth possibilities suggest that the mixed channel MFK-share results should be interpreted with caution: the effective reduction in content for MFK viewers may be overstated.

5.1.1 Time-Varying Treatment Effects

Figure 6 shows our time-varying treatment estimates (see Appendix C for details) for the supply-side outcomes separately for MFK and mixed channels. If the parallel trends assumption holds, then treatment and control groups should follow the same trends prior to treatment, and the pre-announcement coefficient estimates should all be close to zero and without a trend. Virtually none of the estimated pre-announcement coefficients is statistically different from zero. However, the coefficient estimates for mixed channel video production in Figure 6b exhibit a downward pre-trend, which is consistent with Figure 2. As we discuss above, our panel differences estimates (Appendix A) address this concern: yielding more conservative, but still negative estimates. For the MFK share, the pre-announcement coefficient estimates do not have distinguishable trends, but start to drop below zero after the settlement announcement.


Consistent with our Table 2 estimates, Figure 6 shows a sharp and negative drop in both the number of video releases and the share of MFK videos around the settlement enforcement date. The drops in MFK channel video releases and mixed channel MFK-share are especially stark. The post-settlement time-varying treatment effects are relatively stable around the average treatment effect for MFK channels. For mixed channels, both the number of video releases and the MFK share appear to continue a slight downward trend throughout the post-treatment period. Figure 6 shows evidence of anticipatory adjustments during the announcement period.

5.1.2 Heterogenous Treatment Effects

Next, we examine whether the effect of the YouTube settlement varies by channel size. To this end, we build on our main model by adding interactions to estimate size-level heterogeneous treatment effects. Appendix D considers heterogeneity by content category, and notably shows

³²Because YouTube’s automated classification algorithm can override the channel’s MFK classification, the extent to which this discretion exists depends on the sensitivity of the algorithm. Channels can appeal the algorithm’s MFK classification, but this action involves hassle and delay.

Figure 6: Time-Varying Treatment Effects on Supply-Side Outcomes



Note: The solid lines indicate the point estimates, while the bands show the 95% confidence intervals.


that MFK channels in the education category exhibit a similar reduction (coefficient estimate of -0.107) in content output.

We consider whether the YouTube settlement had a disproportionate impact on smaller channels. Existing research finds that the GDPR hurt smaller firms more (Johnson et al., 2023; Peukert et al., 2022). In our setting, removing platform personalization should reduce the platform's ability to match users to relevant content. Based on past literature on personalized search and recommendations (e.g., Korganbekova & Zuber 2023; Sun et al. 2022), we expect that eliminating such personalization elements to most reduce views for smaller channels. Since ad revenue is proportional to views, reducing content views should reduce equilibrium content supply. Thus, smaller channels may reduce their content supply more, all else equal.

Figure 7 reports results by channel size quartile, where we split channels by their number of subscribers as of October 2019. Table I.2 of Appendix I reports the full estimation results. Figure 7a reports marginal effects for video uploads. For both MFK and mixed channels, smaller channels reduce their video releases by more after the YouTube settlement. The smallest quartile MFK and mixed channels exhibit -24.2% and -22.2% reductions respectively in video releases, both of which are statistically significant. On the other extreme, the largest quartile of both MFK and

mixed channels have a -11.1% and -9.5% reduction in video releases ($p < 0.1$ and $p < 0.05$, respectively). Turning to the MFK share outcome, the impact on mixed channels appears to increase with channel size, with the smallest and largest quartiles experiencing -4.8 and -11.7 percentage point reductions, respectively. In contrast, the MFK share estimates for pure MFK channels are smaller than those for mixed channels and do not exhibit a clear pattern by channel size.

Figure 7: Treatment Effects on Content Supply by Channel Subscription Size



Note: In the figures above, the dots indicate the marginal effect estimates (Figure 7a) and coefficient estimates (Figure 7b), while the vertical lines indicate the 95% confidence intervals. Computation of marginal effects detailed in Appendix F. <25%: bottom quartile; >75%: top quartile.

5.2 Content Quality Results

YouTube channels may respond to reduced ad revenue by cutting both MFK content quality and quantity, since quality is costly to provide. Content quality or vertical differentiation is difficult to measure: online platforms are known for cultivating horizontally differentiated content (Waldfo-

gel, 2017). We investigate the settlement’s impact on two content attributes, original content and manual captioning, which we believe are objective quality metrics. We also analyze user ratings as a subjective quality metric. Below, we describe our outcome variables in more detail before presenting our impact estimates.

5.2.1 Quality metrics

First, we consider content originality as one dimension of content quality. Channels may re-release older content to reduce production costs. For instance, a channel that produces songs for kids could create a compilation video containing songs that the channel previously released. Moreover, channel creators have an incentive to maintain their content release cadence: A regular release schedule keeps users engaged with the channel and thereby maintains the content’s ranking on YouTube. Some channels may therefore use duplicate content to maintain their release schedule.

We develop an algorithm to quantify content originality using YouTube video transcripts. The algorithm compares each video to all prior videos released by the channel to identify repeated text passages in the video. The algorithm then calculates the video’s originality score as one minus the ratio of repeated content to the total content. Due to computational constraints, we exclude the top 5% of channels by their total video releases. We also restrict our attention to videos with English transcripts, which yields originality scores for 649,035 videos in our sample. We describe details of the algorithm and the sample inclusion criteria in Appendix E.

We define an originality content indicator as videos whose originality score exceeds 50%. This is meant to capture high levels of duplicate content that are worse for viewers, all else equal. We emphasize that children may value some content repetition: repetition may even be valuable pedagogically. Nonetheless, our originality metric identifies duplication between videos and ignores repetition within a video. Moreover, we look for *changes* in content duplication, which excludes repeated content elements across videos like opening scenes. Parents and children may value compilations as they offer lengthy videos from a preferred creator. The demand for compilation videos plausibly increased when YouTube deprecated playlists for MFK content as part of the settlement. Nevertheless, compilations offer less flexibility to select and order content than playlists. Moreover, compilations restrict viewers to a single creator, which further limits competition.

Our second quality metric looks at video captions. Specifically, we use YouTube’s Data API to examine whether channels provide manually-generated video captions or instead rely on YouTube’s automated caption algorithm. We have this metric for the subsample of about 1 million videos in our sample for which we have transcripts in English. YouTube began offering AI-generated captions in 2009 as a feature to improve accessibility. However, a previous report suggests that YouTube’s automatic caption accuracy rate can be as low as 60%-70%.³³ Manual captioning is a vertical attribute because it can improve transcription accuracy, which improves accessibility for users with hearing loss and users who watch without sound (e.g., while on public transit). More generally, accurate captions can help children learn to read and improve their vocabulary (Bird & Williams, 2002; Kothari & Bandyopadhyay, 2014; Linebarger, 2001). However, manual transcripts are costly for channels to provide,³⁴ so channels may increase their reliance on automated captioning when they lose ad revenue. Prior to the FTC’s announcement, manual caption rates were 14.3% for MFK channels and 15.9% for non-MFK channels (Table 1).

Lastly, we consider user content ratings as a subjective quality metric. If creators reduce content quality, user ratings may fall in response. Further, user ratings may also fall because deactivating YouTube’s personalization elements reduce the match quality between users and content. That is, holding video content quality constant, the lack of user data available to YouTube means that users will be less likely to find content that matches their horizontal preferences. We use video like counts to create a normalized user-rating metric as follows:³⁵

$$Quality_{itm} = \frac{likes_{itm}/views_{itm}}{Avg_{m,Pre-announcement}(likes_{itm}/views_{itm})}$$

for channel i ’s videos belonging to MFK type m and released in week t . Note that our like and view variables are cumulative counts as of July 2022. The numerator is the like/view ratio at the weekly level, separately constructed for the channel’s MFK and non-MFK videos. The denominator normalizes the like/view ratio using its average level prior to the settlement announcement,

³³<https://itss.d.umn.edu/centers-locations/media-hub/media-accessibility-services/captioning-and-captioning-services/correct>

³⁴On <https://www.rev.com/pricing>, one provider charges \$1.50 per minute for human transcription and \$0.25 per minute for automated transcription. The provider advertises 99% accuracy rates for human transcription and over 90% accuracy for automated transcription.

³⁵YouTube no longer provides data on the number of dislikes, so we cannot use this metric to construct a quality measure following past work (e.g., Kerkhof, 2020).

separately by MFK video type, m . We normalize by type m because the like-to-view ratio is eight times lower for MFK channels than non-MFK channels (see Table 1), which could reflect children’s greater appetite for repeat content viewing.

5.2.2 Estimates

We modify our baseline difference-in-differences specification (equation 1) for our three quality outcomes: i.e., original content share, manual captioning share, and the like/view ratio. Our unit of analysis shifts to the channel-week-MFK type level, so that we can separately examine the quality of MFK and non-MFK content for mixed channels.³⁶ We do so as our baseline model would otherwise conflate mixed channels’ shift away from MFK content with genuine quality changes. To see this, consider the like-view ratio, which is eight times lower for MFK channels than non-MFK channels (see Table 1). Without this modification, the mixed channels’ quality metric would artificially increase post-settlement, as those that pivot toward non-MFK content would obtain higher like-to-view ratios.

Table 3 presents the difference-in-differences estimates for our three quality metrics. We find statistically significant evidence that all three quality measures for MFK channels drop after the YouTube settlement ($p < 0.01$). Both objective quality metrics demonstrate negative estimates for the anticipation period as well as for MFK content on mixed channels, though only the former are statistically significant ($p < 0.05$). Table 3 shows that the share of original content falls by 7.5 p.p. for MFK channels post-settlement. This is a 11.2% decrease relative to the pre-announcement mean of 0.672. Appendix E.2 shows that our findings are robust to alternate outcome variables: our continuous originality score and varying the score threshold (0.5) for defining our original content indicator. We find that the manual captioning share falls 3.8 p.p., which represents a 26.6% drop relative to the pre-announcement period mean (0.143). Table 3 shows that the normalized like-view ratio falls 0.103 points for MFK content by MFK channels after the settlement, implying a 10.3% decline. Though only marginally significant, the like-view ratio increases 0.063 points for MFK content by mixed channels after the settlement. Appendix B, however, estimates this model for our global YouTube channel sample and obtains negative coefficient estimates for both MFK channels ($p < 0.01$) and MFK content on mixed channels ($p < 0.05$). Appendix C provides

³⁶We omit the few observations with non-MFK content on MFK channels and MFK content on non-MFK.

corresponding time-varying treatment effect estimates showing evidence of parallel pre-trends.

Table 3: Impact on Video Quality: Difference-in-Differences Estimates

	Original Content Share ^a	Manual Captioning Share ^b	Normalized Like/View ^c
Post-Settlement			
× MFK	-0.077*** (0.010)	-0.038*** (0.013)	-0.103*** (0.029)
× Mixed (MFK videos)	-0.021 (0.014)	-0.024 (0.020)	0.063* (0.037)
× Mixed (non-MFK videos)	-0.003 (0.005)	-0.012 (0.012)	0.027 (0.021)
Anticipation Period			
× MFK	-0.030*** (0.009)	-0.021** (0.010)	0.021 (0.026)
× Mixed (MFK videos)	-0.001 (0.009)	-0.019 (0.016)	0.059 (0.037)
× Mixed (non-MFK videos)	0.003 (0.005)	-0.002 (0.010)	0.019 (0.016)
Week, channel-MFK type FE	Y	Y	Y
Adj. R^2	0.630	0.712	0.731
N	254,700	278,166	387,483

Notes: Post-Settlement is defined as weeks after Jan. 1, 2020. Anticipation Period is defined as weeks between Sep.4, 2019 to Jan.1, 2020. All specifications include week and channel-MFK type fixed effects and omit the few observations with MFK (non-MFK) videos on non-MFK (MFK) channels. Robust standard errors clustered at the channel level are in parentheses. ***significant at 1% level; **significant at 5% level; *significant at 10% level. ^aThe original content share takes the channel-weekly average of a video indicator for content with originality score ≥ 0.5 . The originality results restrict the sample notably videos with transcripts in English and to exclude the top 5% of channels by cumulative video releases. See Appendix E for details. ^bThe manual caption share is given for the subset of about 1 million videos with transcripts in English. ^cThe normalized like/view outcome variable uses YouTube's video-level like and view data, which are cumulative counts as of July 2022. The like/view sample alone omits MFK content on non-MFK channels and vice-versa due to the large differences in their respective conditional means.

We perform a back-of-the envelope calculation to compute the combined impact of the YouTube settlement using our estimates of the impact on the supply of video releases (-18.4%), the MFK content share (-2.7 p.p.), as well as original content (-7.7 p.p.) for MFK channels. We calculate that original, MFK video releases decline about 24.9% for MFK channels: i.e., $18.4\% + (1 - 18.4\%)(7.7\% + 2.7\%)$. In particular, the relative reduction in the share of original content (11.2%) for MFK channels is about two thirds as large as the reduction in their supply of video releases (18.4%), though we acknowledge that the former results apply to only a subset of our video data. While other scholars also study the impact of the YouTube settlement (Kircher & Foerderer, 2023b), we are the first to note this important dimension. The reduction in original content suggests a partial answer to the apparent inelasticity of content supply: the total decline in *original* content (18.4%

+ (1-18.4%)*7.7%) is about a third greater than our supply estimate alone (18.4%). Our study also provides novel evidence that the share of videos with manual captions falls by over a quarter.

5.3 Demand-Side Results

We now turn to the impact of the YouTube settlement on viewer demand for MFK content. The settlement can affect content demand in two ways. First, the reduction in content supply (due to lower ad prices) should lead to a reduction in content viewing. Since YouTube content is free, the equilibrium demand is not moderated by content price. Instead, viewers allocate their attention to either other YouTube content or alternatives. Second, the loss of platform personalization will hinder content discovery, especially for small channels. Past literature suggests that personalized recommendations and search results benefit niche firms and products more. When YouTube is less able to help viewers find content that matches their taste, the demand for smaller channels should fall more. Moreover, this demand reduction can be self-reinforcing: reduced viewer demand leads to less revenue for small channels and in turn less content creation, which further reduces viewer demand. Shiller et al. (2018) describe a similar self-reinforcing cycle regarding the impact of ad-blocking on user demand for website content.

We examine the impact of the YouTube settlement on two user demand metrics: the number of new views and the number of new subscriptions at the channel-week level. We follow a similar identification approach as in Section 5.1. However, our demand metric panel starts at the beginning of October due to the data limitation discussed in Section 3.1. Consequently, our estimation omits the $Announcement_t$ terms in equation (1) and compares the post-settlement period to the last three months of the announcement period. Though anticipatory changes are possible on the supply side, viewers should only change their viewing behavior in response to changes in content releases, which largely occur after the settlement's implementation. To the extent that the content supply began to falter during the announcement period, our demand-side estimates are conservative. Recall from Section 3 that our new subscription metric is more volatile because YouTube provides rounded subscription counts.

Table 4 shows our demand-side average treatment effect estimates for MFK and mixed channels over the post-settlement period. The parameter estimates show a negative and statistically

significant reduction in all demand metrics for both MFK and mixed channels relative to the control. For views, we estimate a -19.9% marginal effect for MFK channels and a -12.9% marginal effect for mixed channels. The impact on new subscriptions is larger with MFK channels falling 25.0% and mixed channels falling 23.9%.

The larger decrease in new subscriptions than views could reflect a greater impact on channel discovery. In particular, eliminating platform personalization yields worse matches between viewers and content, which may limit new matches between viewers and channels. However, the result could also reflect the fact that the YouTube settlement eroded one key benefit of subscribing to a channel: receiving new content notifications.

Table 4: Settlement Impact on Views & Subscriptions: Difference-in-Differences Estimates

	log (Views+1)	log (Subscriptions+1)
MFK $\times$ Post-Settlement	-0.222*** (0.038)	-0.288*** (0.084)
Mixed $\times$ Post-Settlement	-0.139*** (0.041)	-0.273*** (0.085)
Week, channel fixed effects	Y	Y
Marginal Effects		
MFK	-19.89%	-25.00%
Mixed	-12.94%	-23.91%
Adj. R^2	0.867	0.574
N	325,258	325,258


Notes: *Post-Settlement* is defined as weeks after Jan. 1, 2020. Computation of marginal effects detailed in Appendix F. Robust standard errors clustered at the channel level are in parentheses. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

5.3.1 Time-Varying Treatment Effects

Figure 8 shows our time-varying treatment effect estimates for our demand-side outcomes. If non-MFK channels are a valid control group, the pre-treatment coefficient estimates should all be close to zero. Figure 8 confirms that none of the estimated pre-treatment estimates are statistically significantly different from zero. Unlike the announcement period changes on the supply side (Figure 6), we see little evidence of an anticipatory pullback in demand during the announcement period. This pattern supports our conjecture that the announcement-to-implementation period comparison provides a valid measure of how the YouTube settlement affects viewer demand.

Figure 8 shows a sharp drop in subscriptions and views in January 2020 for both mixed and

Figure 8: Time-Varying Treatment Effects on Demand-Side Outcomes



Note: The solid lines indicate the point estimates, while the bands show the 95% confidence intervals.

MFK channels relative to the non-MFK control group. We note a modest bounce back in demand around March 2020, which marks the start of the COVID pandemic. During this time, schools in many countries closed down, which could increase the demand for MFK content on YouTube. Our demand estimates soon return to their January 2020 levels. For unknown reasons, the demand effect of the YouTube settlement attenuates starting Fall 2020, particularly for subscriptions to MFK channels. This appears to arise despite the persistent reduction in new content creation observed in Figure 6.

The demand boost to COVID may lead us to underestimate the post-settlement reduction in viewer demand that would have arose absent the pandemic. However, we see little evidence that the supply of MFK content rose to meet the COVID boost in demand (see Figure 6), though the COVID lockdowns created production challenges for YouTube channels as well.

5.3.2 Heterogeneous Treatment Effects

Next, we examine whether the demand-side impact of the YouTube settlement varies by channel size. As in Section 5.1.2, we split channels by their quartile of subscriber count as of October 2019. Appendix D considers differences by content category: in particular, education-category MFK

channel demand falls for both views (coefficient estimate of -0.204) and subscriptions (-0.389).


Figure 9 plots marginal effect estimates showing that the magnitude of the settlement impact falls with channel size (see Table I.3 for the full estimation results). This pattern is much more pronounced than the heterogeneity pattern observed on the supply side (Figure 7): views fall by 35.6% for the bottom-quartile MFK channels and 21.2% for the bottom-quartile mixed channels. For the largest quartile channels, MFK channel views are actually increasing (1.5%) but statistically insignificant, and mixed channel views too are increasing (8.4%) and marginally significant. The new subscription results are even more stark with the lowest quartile MFK channels falling -46.3% and mixed channels falling -38.4%. At the other end, we estimate essentially no impact on the number of new subscriptions for the largest quartile channels: 2.1% for top MFK channels and 0.3% for mixed channels. Kircher & Foerderer (2023b) find the opposite pattern for the educational channels that they consider: subscriptions rise for bottom-quartile MFK channels and fall for the rest. Moreover, the former effect dominates as Kircher & Foerderer (2023b) find an aggregate increase in MFK subscribers. Nevertheless, our results are not directly comparable, as they omit channels that exit (see our Figure 3) whereas we consider all channels.

The fact that the incidence of the demand reduction is concentrated on smaller channels is consistent with the greater reduction in supply observed for these channels in Section 5.1.2. However, the reduction in demand exceeds the reduction in supply: for example, a 24.2% reduction in the supply of smallest quartile MFK channel content versus a 35.6% reduction in views. The greater effect on the demand side is consistent with a reduction in platform personalization so that the platform pushes users to consume popular channel content rather than lower-ranked content that may be better matched to the viewer's individual preferences. This also means that the absolute decrease in demand post-settlement is more modest than our estimates indicate, which focuses on the percentage decrease.

6 Welfare Implications

We discuss the welfare implications of our findings below. Though we do not attempt a formal welfare analysis, we discuss the implications of a traditional welfare analysis as well as its inherent limitations.

Figure 9: Treatment Effects on Content Demand by Channel Subscription Size



Note: In the figures above, the dots indicate the marginal effect estimates, while the vertical lines indicate the 95% confidence intervals. Computation of marginal effects detailed in Appendix F. <25%: bottom quartile; >75%: top quartile.

A full welfare calculation for the settlement's impact on MFK creators and viewers requires analyses of consumer surplus from MFK content and from enhanced privacy, content creator's profits, and the impact on advertisers. We do not attempt to quantify the privacy benefit to MFK viewers. Although the social value of children's privacy in general is likely to be high, the marginal value of enhanced privacy in this particular context—banning persistent identifiers associated with a particular device that children use—is uncertain. We lack creator cost and revenue data to quantify the impact on creator profits. However, we provide strong evidence that creator profits fell, consistent with a reduction in MFK content supply. Moreover, creator revenue is a function of content views and ad prices: we find that the MFK channel views fell by 20% and we provide anecdotal evidence that ad prices fell 73%. We lack advertiser data to speak to advertiser impact, though we discuss related concerns in Appendix G.

A traditional welfare analysis would conclude from our results that the settlement reduced consumer surplus from MFK content. The decrease in viewer demand suggests that MFK content viewers did not find adequate substitute YouTube content to offset the reduced MFK content supply. Holding fixed the outside option, these results imply that MFK content viewers obtained less utility from watching YouTube after the settlement. Moreover, our results show the settlement also degraded content quality, which implies a further utility loss. If YouTube content is a normal good for MFK content viewers, our results imply a reduction in consumer surplus for MFK content viewers. For context, Brynjolfsson et al. (2019) use online choice experiments to evaluate the surplus associated with online product categories for adult consumers. For video streaming services like YouTube and Netflix, Brynjolfsson et al. (2019) find a median willingness-to-accept of \$1,173 per year. This suggests that the surplus generated by YouTube as a free service may be high.

However, a traditional welfare analysis may have limitations in this setting. First, viewing YouTube content could be thought of as good or bad for children. For instance, YouTube may have ill effects on mental health, though this has been understudied. Instead, economists have studied social media more generally and many have shown a negative impact on mental health in adults (e.g., Allcott et al., 2020; Braghieri et al., 2022; Mosquera et al., 2020). In the psychological literature (J. Haidt et al., *oing*), many studies find a quadratic relationship between screen time and self-reported well-being: moderate levels of daily use (i.e., around two hours) are associated with higher levels of self-reported well-being, but higher levels of use are associated with lower levels of self-reported well-being. However, few of the studies cited in J. Haidt et al. (*oing*) examine YouTube: only one examines children and none provide causal evidence linking YouTube with mental well-being.³⁷ Educational content represents 27% of MFK channels in our sample prior to the settlement and this category may best exemplify content that benefits children. However, we see similar-sized reductions in educational content production and viewing after the settlement (see Appendix D). Given the link between captions and children’s language learning, the reduction in manual captioning also suggests some educational harm to children.

Second, if YouTube is addictive, traditional welfare calculations would overstate the decline in surplus from viewing less YouTube content. For instance, economists have found evidence

³⁷In a correlational study, Fardouly et al. (2020) find that YouTube and Instagram users reported more body image issues, but no higher levels of anxiety or depressive symptoms than non-users.

suggesting that social media usage can be addictive for adults (Allcott et al., 2020, 2022; Aridor, 2023; Mosquera et al., 2020). The evidence for YouTube, however, is mixed. Aridor (2023) finds no such evidence of addiction for YouTube. Allcott et al. (2022) show that demand for a bundle of apps—consisting of social media, web browsing, and YouTube—collectively exhibit addiction, and find evidence to suggest that adult users have self-control problems with YouTube viewing.

Third, children may substitute towards problematic alternatives in place of YouTube. Although children can substitute towards activities that may be beneficial like playing outside, studying, or practicing piano, they also may substitute towards social media use or video content viewing (i.e., online streaming or television) that share the above concerns of addiction and harm to children. Extant literature on substitution patterns shows that adults substitute an important share of their time away from YouTube into other apps. Aridor (2023) finds that users restricted from viewing YouTube substitute into other social media apps, rather than entertainment apps. Further, he finds increased time spent on newly downloaded apps and reduced overall phone time, although how this time was spent is unclear. Allcott et al. (2022) show that restricting the use of a bundle of apps that included YouTube led users to substitute about half of their economized time into other smartphone apps.

In sum, traditional welfare analysis of content results suggests that the YouTube settlement harmed both creators and viewers. However, this analysis would overstate the harm to children if YouTube content is harmful or addictive for children. More research is therefore needed to understand the benefits, harms, and addictiveness of YouTube viewing among children as well as children’s substitutes for YouTube.

7 Conclusion

We empirically examine the privacy-for-content tradeoff by studying YouTube’s 2020 settlement with the FTC. We show that the settlement, which eliminated both personalized advertising and platform personalization for make-for-kids content, impacted both the supply and the demand for content. Made-for-kids (MFK) creators reduced content creation by about 18%. Moreover, channel creators pivot away from MFK content, with mixed channels reducing their MFK content share by over a quarter. Smaller channels experienced larger reductions in content supply. MFK

creators also reduce quality—relying more on duplicate content and cutting manual captioning—and user-quality ratings fall. Our results show that the demand for content falls for both MFK and mixed channels. The loss of platform personalization likely reduces the platform’s ability to match users to more niche content. Consistent with this conjecture, we observe large reductions in viewer demand for smaller channels, but little to no change in demand for top channels, which exacerbates the settlement’s negative effect on smaller channels’ ad revenue.

We acknowledge several limitations of our research. First, we rely on data collected after the settlement. As such, our top channel sample is favorably selected and may understate the impact on smaller channels that dropped out of the top channel list. In addition, our analysis is not designed to capture YouTube settlement’s impact on channel entry. Second, we cannot separately identify the impact of YouTube’s concurrent changes to its platform and ad personalization. Third, the COVID pandemic occurred in the middle of our event study, which may violate our identifying assumptions. Fourth, spillovers from the policy to non-MFK channel ad revenue may lead us to overestimate settlement’s impact.

A traditional welfare analysis suggests that the reduction in MFK videos is welfare reducing. That said, traditional welfare logic may overstate consumer harm if MFK content on YouTube is addictive or otherwise harmful. Even in this case, however, the settlement’s impact on children’s welfare depend on the alternatives they choose. We may still consider reduced YouTube consumption as welfare reducing, if these outside options (e.g., social media and video streaming services) are worse than YouTube content.

With these caveats in mind, our results nonetheless have implications for privacy regulation that limit the use of persistent online identifiers for content providers. For example, Congress is considering a comprehensive privacy bill that will restrict personalized advertising as well as an amendment to strengthen COPPA.³⁸ The FTC recently proposed modifications to its COPPA rule, which among other things, would heighten consent requirements for personalized advertising and further limit the use of persistent identifiers for website personalization.³⁹ Further, it has

³⁸See *American Data Privacy and Protection Act*, H.R. 8152: <https://www.congress.gov/bill/117th-congress/house-bill/8152/text>; and Senators Markey and Cassidy’s proposal “COPPA 2.0”: <https://www.markey.senate.gov/news/press-releases/senators-markey-and-cassidy-reintroduce-coppa-20-bipartisan-legislation-to-protect-online-privacy-of-children-and-teens>.

³⁹<https://www.ftc.gov/news-events/news/press-releases/2023/12/ftc-proposes-strengthening-childrens-privacy-rule-further-limit-companies-ability-monetize-childrens>.

started a rulemaking process to curtail what it refers to as “commercial surveillance.”⁴⁰ While the FTC has yet to make concrete proposals, its request for comment suggests a skepticism toward personalized advertising and, to a lesser extent, personalization features.⁴¹ As another example, the European Union’s recent Digital Services Act prohibits personalized advertising towards children under 18 on large online platforms like YouTube. Our results highlight a tradeoff between online content and privacy that policymakers should consider when drafting privacy regulation.

References

- Abou El-Komboz, L., Kerkhof, A., & Loh, J. (2022). Complementor reactions to platform control – evidence from the YouTube “Adpocalypse”.
- Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–92.
- Alcobendas, M., Kobayashi, S., Shi, K., & Shum, M. (2023). The impact of privacy measures on online advertising markets. *Available at SSRN 3782889*.
- Allcott, H., Braghieri, L., Eichmeyer, S., & Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review*, 110(3), 629–76.
- Allcott, H., Gentzkow, M., & Song, L. (2022). Digital addiction. *American Economic Review*, 112(7), 2424–63.
- Aridor, G. (2023). Market definition in the attention economy: An experimental approach.
- Beemsterboer, S. (2020). COPPA killed the video star: How the YouTube settlement shows that COPPA does more harm than good. *Illinois Business Law Journal*, 25.
- Bhargava, H. K. (2022). The creator economy: Managing ecosystem supply, revenue sharing, and platform design. *Management Science*, 68(7), 5233–5251.
- Bird, S. A. & Williams, J. N. (2002). The effect of bimodal input on implicit and explicit memory: An investigation into the benefits of within-language subtitling. *Applied Psycholinguistics*, 23(4), 509–533.
- Braghieri, L., Levy, R., & Makarin, A. (2022). Social media and mental health. *American Economic Review*, 112(11), 3660–3693.
- Brynjolfsson, E., Collis, A., & Eggers, F. (2019). Using massive online choice experiments to measure changes in well-being. *Proceedings of the National Academy of Sciences*, 116(15), 7250–7255.

⁴⁰<https://www.ftc.gov/legal-library/browse/federal-register-notices/commercial-surveillance-data-security-rulemaking>.

⁴¹The FTC also recently moved to modify its order against Meta to prevent it from using data to target ads on its platform to any user under 18. See: <https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-propose-s-blanket-prohibition-preventing-facebook-monetizing-youth-data>.

- Cecere, G. & Lemaire, S. (2023). Have i seen you before? measuring the value of tracking for digital advertising.
- Chen, J. & Roth, J. (forthcoming). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*.
- Cheyre, C., Leyden, B. T., Baviskar, S., & Acquisti, A. (2023). The impact of apple's app tracking transparency framework on the app ecosystem. *Available at SSRN 4453463*.
- Donnelly, R., Kanodia, A., & Morozov, I. (2022). Welfare effects of personalized rankings. *Available at SSRN 3649342*.
- Fardouly, J., Magson, N. R., Rapee, R. M., Johnco, C. J., & Oar, E. L. (2020). The use of social media by australian preadolescents and its links with mental health. *Journal of Clinical Psychology*, 76(7), 1304–1326.
- Goldberg, S., Johnson, G. A., & Shriver, S. K. (2024). Regulating privacy online: An economic evaluation of the GDPR. *American Economic Journal: Economic Policy*, 16(1), 325–58.
- Goldfarb, A. & Que, V. F. (2023). *The Economics of Digital Privacy*. Working Paper 30943, National Bureau of Economic Research.
- Hui, K.-L., Kircher, T., & Foerderer, J. (2023). Social disadvantages of privacy? its convergent consequences for the access to app-based education. *Available at SSRN*.
- Imbens, G. W. & Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- J.Haidt, Z.Rausch, & Twenge, J. (ongoing). Social media and mental health: A collaborative review. NYU at <https://docs.google.com/document/d/1w-HOfseF2wF9YIpXwUUtP65-olnkPyWcgF5BiAtBEy0/edit>.
- Johnson, G. A. (2022). *Economic Research on Privacy Regulation: Lessons from the GDPR and Beyond*. Working Paper 30705, National Bureau of Economic Research.
- Johnson, G. A., Shriver, S., & Goldberg, S. (2023). Privacy & market concentration: Intended & unintended consequences of the GDPR. *Management Science*.
- Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1), 33–51.
- Katz, J. & Fener, V. (2019). Is a YouTube COPPAcalypse coming? FTC rules could start demonetizing creators in 2020.
- Kerkhof, A. (2020). Advertising and content differentiation: Evidence from YouTube.
- Kesler, R. (2023). The impact of Apple's App Tracking Transparency on app monetization. *Available at SSRN 4090786*.
- Kircher, T. & Foerderer, J. (2023a). Ban targeted advertising in apps? an empirical investigation of the consequences for app development. *Management Science*, (pp. in press).
- Kircher, T. & Foerderer, J. (2023b). Does privacy undermine content provision and consumption? evidence from educational YouTube channels. *SSRN working paper*.

- Korganbekova, M. & Zuber, C. (2023). Balancing user privacy and personalization. Work in progress.
- Kothari, B. & Bandyopadhyay, T. (2014). Same language subtitling of bollywood film songs on tv: Effects on literacy. *Information Technologies & International Development*, 10(4), pp–31.
- Laub, R., Miller, K., & Skiera, B. (2022). The economic value of user tracking for publishers. SSRN working paper.
- Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (2022). Does privacy regulation harm content providers? a longitudinal analysis of the impact of the gdpr. Available at SSRN.
- Li, D. & Tsai, H.-T. (2023). Mobile apps and targeted advertising: Competitive effects of data sharing. Available at SSRN 4088166.
- Li, N., Haviv, A., & Lovett, M. (2023). Let's play fair - purchase and usage effects of influencer marketing on youtube. Available at SSRN 3884038.
- Linebarger, D. L. (2001). Learning to read from television: The effects of using captions and narration. *Journal of educational psychology*, 93(2), 288.
- Mathur, A., Narayanan, A., & Chetty, M. (2018). Endorsements on social media: An empirical study of affiliate marketing disclosures on YouTube and Pinterest. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), 1–26.
- Miller, A. R. (2022). Privacy of digital health information. In A. Goldfarb & C. Tucker (Eds.), *Economics of Privacy* chapter 3. University of Chicago Press.
- Mosquera, R., Odunowo, M., McNamara, T., Guo, X., & Petrie, R. (2020). The economic effects of facebook. *Experimental Economics*, 23, 575–602.
- Panjwani, A. & Xiong, H. (2023). Subscription revenue and media content: Evidence from YouTube. Working Paper.
- Peukert, C., Bechtold, S., Batikas, M., & Kretschmer, T. (2022). Regulatory spillovers and data governance: Evidence from the GDPR. *Marketing Science*, 41(4), 318–340.
- Phillips, N. (2019). Separate Statement of Commissioner Noah Joshua Phillips: United States of America and People of the State of New York v. Google LLC and YouTube, LLC. Federal Trade Commission Public Statement.
- PricewaterhouseCoopers (2019). *Kids digital media report 2019*. Technical report, PricewaterhouseCoopers.
- Ravichandran, D. & Korula, N. (2019). *Effect of disabling third-party cookies on publisher revenue*. Technical report, Google Inc.
- Rideout, V. & Robb, M. B. (2019). *The Common Sense Census: Media Use by Tweens and Teens, 2019*. Technical report, Common Sense.
- Shiller, B., Waldfogel, J., & Ryan, J. (2018). The effect of ad blocking on website traffic and quality. *The RAND Journal of Economics*, 49(1), 43–63.

- Sun, M. & Zhu, F. (2013). Ad revenue and content commercialization: Evidence from blogs. *Management Science*, 59(10), 2314–2331.
- Sun, T., Yuan, Z., Li, C., Zhang, K., & Xu, J. (2022). The value of personal data in internet commerce: A high-stake field experiment on data regulation policy. *Available at SSRN 3962157*.
- Tang, Q., Gu, B., & Whinston, A. (2011). Content contribution under revenue sharing and reputation concern in social media: The case of YouTube. In *32nd International Conference on Information System 2011, ICIS 2011* (pp. 799–814).
- Tang, Q., Gu, B., & Whinston, A. B. (2012). Content contribution for revenue sharing and reputation in social media: A dynamic structural model. *Journal of Management Information Systems*, 29(2), 41–76.
- Urgo, J. (2020). Made For Kids & COPPA – initial look at the YouTube data.
- Waldfogel, J. (2017). How digitization has created a golden age of music, movies, books, and television. *Journal of economic perspectives*, 31(3), 195–214.
- Wooldridge, J. M. (2023). Simple approaches to nonlinear difference-in-differences with panel data. *The Econometrics Journal*, 26(3), C31–C66.
- Yeo, S. L., Schaller, A., Robb, M. B., & Radesky, J. S. (2021). Frequency and duration of advertising on popular child-directed channels on a video-sharing platform. *JAMA Network Open*, 4(5), e219890–e219890.
- Yoganarasimhan, H. (2012). Impact of social network structure on content propagation: A study using YouTube data. *Quantitative Marketing and Economics*, 10, 111–150.
- YouTube Official Blog (2023). YouTube for press.

Appendix

A Supply-side Impact Results: Robustness

We consider the robustness of our supply-side findings to alternative identification approaches. We apply two approaches, panel differences and triple differences, that use the treated channel’s outcomes from the previous year as an alternative control group. Recall that our demand-side data begins in October 2019 and therefore lacks the requisite pre-treatment data for these robustness checks.

These robustness checks address two concerns. First, the YouTube settlement may generate equilibrium spillovers to the non-MFK control group via ad prices, as the advertiser’s demand for ad personalization can only be filled by non-MFK content. The settlement would thus create a relative increase in ad prices for non-MFK content. Panel differences avoids the potential concern that using non-MFK channels as the control group violates the Stable Unit Treatment Values Assumption (SUTVA). Second, we examine the robustness of our evidence to alternative identifying assumptions. In particular, we observe some evidence of differential pre-trends between mixed and non-MFK channels (see Figures 2 & 6). Our panel difference and triple difference models will parse out the weak divergence in trends across the two groups from the causal effect estimates.

Our panel differences approach (see e.g., Goldberg et al., 2024) applies a modified two-way fixed effect estimator (see equation 1) where the treatment group is instead MFK channels during July 1, 2019 to June 30, 2020 and the control group is those same MFK channels from July 1, 2018 to June 30, 2019. Instead of time fixed effects, panel differences include week-of-year fixed effects, which capture seasonality in these channels’ output. Mixed channel panel differences estimates are analogously derived. To ensure the treatment and control groups have the same number of weeks, the post-settlement period is six months shorter than our preferred approach in Table 2—so that these estimates will diverge somewhat. Our panel difference estimates also include fewer observations, because they omit non-MFK channels. Our triple-differences approach differences the panel difference estimates with an analogous panel difference for non-MFK channels. The triple-differences approach accounts for any lingering year-over-year differences in channel outcomes

that are common to all channel types.⁴²

The second and third columns of Table A.1 present our panel difference estimates. The MFK channel estimates are close to our preferred specification estimates in Table 2: the confidence intervals overlap considerably. The video release estimate (-0.131) is essentially unchanged, which may allay the concern that the YouTube settlement contaminated the non-MFK control group by increasing its ad prices and content creation. The MFK share estimates are somewhat weaker at -2.0 rather than -2.7 percentage points. By definition, the MFK video share can only go down for MFK channels and can only go up for non-MFK channels, so difference-in-differences will yield larger (absolute) estimates. As expected, the mixed channel estimates are more modest under the panel differences specification. Our coefficient estimate for mixed channel video releases is smaller (-0.052 versus -0.136 in Table 2), though this decrease remains marginally statistically significant. The mixed channel MFK share estimate remains negative and statistically significant though is smaller in magnitude: -5.0 versus -9.5 percentage points in Table 2. As above, the increase in MFK share among non-MFK channels contributes to this. However, the modest decline in mixed channels' MFK share during 2019 also attenuates the result as does omitting the additional declines in the second half of 2020 from the sample (see Figure 6). The anticipation period estimates all remain negative though the MFK video release estimate is no longer statistically significant.

The last two columns of Table A.1 present our triple-difference estimates, which are attenuated relative to almost all of our preferred specification estimates in Table 2. The core results—the reduction in MFK channel video releases and the pivot away from MFK content for both MFK and mixed channel—remain highly statistically significant. The mixed channel video release estimate remains negative, but is small in magnitude (-0.012) and is no longer statistically significant. In sum, our robustness checks confirm our skepticism about this particular result, but increase our confidence in the rest of our supply-side results.

B Generalizability: YouTube Channels Globally

We explore the generalizability of our findings to a broader sample of YouTube channels. Recall that our sample consists of 5,066 US channels in the top 3 MFK categories (education, entertain-

⁴²However, triple differences thereby reintroduces non-MFK channels at the risk of a potential SUTVA violation.

Table A.1: Supply-side Effect Robustness Estimates

	Panel Differences ^a		Triple Differences ^b	
	log (Videos + 1)	Share of MFK Videos	log (Videos + 1)	Share of MFK Videos
MFK $\times$ Post-Settlement	-0.131*** (0.029)	-0.020*** (0.004)	-0.091*** (0.031)	-0.022*** (0.004)
Mixed $\times$ Post-Settlement	-0.052* (0.029)	-0.050*** (0.013)	-0.012 (0.031)	-0.051*** (0.013)
MFK $\times$ Anticipation Period	-0.019 (0.019)	-0.010*** (0.003)	-0.008 (0.021)	-0.013*** (0.003)
Mixed $\times$ Anticipation Period	-0.040* (0.021)	-0.023*** (0.009)	-0.029 (0.022)	-0.026*** (0.009)
Week of year, channel fixed effects	Y	Y	Y	Y
Adj. R^2	0.708	0.877	0.732	0.944
N	131,616	80,376	519,114	305,059

Notes: Post-Settlement is defined as weeks after Jan. 1, 2020. Anticipation Period is defined as weeks between Sep.4, 2019 to Jan.1, 2020. All specifications include week and channel fixed effects. Robust standard errors clustered at the channel level. ^aPanel differences compare MFK (Mixed) channels between July 1, 2019 and June 30, 2020 to their past outcomes between July 1, 2018 and June 30, 2019. ^bTriple differences augment the panel differences (i.e., comparing MFK (Mixed) channels between July 1, 2019 and June 30, 2020 to their past outcomes between July 1, 2018 and June 30, 2019) by further differencing the non-MFK panel differences estimates from these MFK (Mixed) panel difference estimates. ***significant at 1% level, **significant at 5% level; *significant at 10% level.

ment, and film & animation). Here, we consider YouTube channels globally across all content content categories. Our wider sample consists of 73,354 channels—drawn from our original list of top 100,000 channels—again that release a video in the pre-announcement period (July 1, 2018 and September 4, 2019) and remain on YouTube in 2022 (see Appendix H). We only collect the supply-side outcomes since the transcript and demand-side data are much costlier to acquire.

Table B.1 presents our resulting supply-side estimates. Our global sample estimates closely resemble our main sample estimates in Table 2. For the video release outcome, we estimate a somewhat larger (absolute) coefficient for MFK channels (-0.143 versus -0.129) and similar coefficient for mixed channels (-0.132 versus -0.136). For the MFK share outcome, we estimate a similar coefficient for MFK channels (-0.030 versus -0.027) and a somewhat smaller (absolute) coefficient for mixed channels (-0.077 versus -0.095). Though we omit our corresponding panel difference estimates for space, these estimates also resemble those in Table A.1 of Appendix A.⁴³ We also omit our heterogeneity estimates by channel size for space, however these estimates broadly resemble those in Figure 7.⁴⁴

Table B.1: Generalizability to Global Channels: Supply-Side Estimates

	log (Videos + 1)	Share of MFK Videos
MFK × Post-Settlement	-0.143*** (0.009)	-0.030*** (0.002)
Mixed × Post-Settlement	-0.132*** (0.007)	-0.077*** (0.003)
MFK × Anticipation Period	-0.043*** (0.008)	-0.011*** (0.001)
Mixed × Anticipation Period	-0.055*** (0.006)	-0.026*** (0.002)
Week, Channel Fixed Effects	Y	Y
Adj. R^2	0.672	0.896
N	9,384,571	5,662,269

Notes: Full global YouTube channel sample. Post-Settlement is defined as weeks after Jan. 1, 2020. Anticipation Period is defined as weeks between Sep.4, 2019 to Jan.1, 2020. The sample size for share of MFK videos is smaller because it is conditional on releasing content in a given week. All specifications include week and channel fixed effects. Robust standard errors clustered at the channel level are in parentheses. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

⁴³For the video release outcome, we estimate a somewhat smaller (absolute) coefficient for MFK channels (-0.111 versus -0.131), but a notably larger (absolute) coefficient for mixed channels (-0.129 versus -0.052). For the MFK share outcome, we estimate a similar coefficients for both MFK channels (-0.023 versus -0.020) and mixed channels (-0.052 versus -0.050).

⁴⁴Again, we see greater impact on video releases for smaller channels and for both MFK and mixed channels. For the MFK share outcome, MFK channels here show a greater impact for smaller channels (Figure 7 exhibited no clear trend) and mixed channels again exhibit greater impact for larger channels.

Table B.2: Generalizability to Global Channels: Video Quality Estimates

Treatment Group: ^b	Normalized Like/View ^a		
	MFK	Mixed (MFK videos)	Mixed (non-MFK videos)
Post-Settlement interactions:	-0.131*** (0.011)	-0.044*** (0.014)	-0.017** (0.007)
Anticipation Period interactions:	-0.008 (0.009)	0.009 (0.011)	-0.004 (0.005)
Week, channel-MFK type FE	Y	Y	Y
Adj. R^2	0.824	0.824	0.826
N	5,048,679	4,992,315	5,291,932

Notes: Full global YouTube channel sample. Post-Settlement is defined as weeks after Jan. 1, 2020. Anticipation Period is defined as weeks between Sep.4, 2019 to Jan.1, 2020. All specifications include week and channel-MFK type fixed effects and omit the few observations with MFK (non-MFK) videos on non-MFK (MFK) channels. Robust standard errors clustered at the channel level are in parentheses. ***significant at 1% level; **significant at 5% level; *significant at 10% level. ^aThe normalized like/view outcome variable uses YouTube's video-level like and view data, which are cumulative counts as of July 2022. The like/view sample alone omits MFK content on non-MFK channels and vice-versa due to the large differences in their respective conditional means. ^bWe separately report difference-in-difference specifications that compare the MFK, mixed (MFK videos), and mixed (non-MFK videos) channels respectively to the non-MFK channels.

Table B.2 presents our global sample video quality results. However, we only present the user rating results due to the computational cost of computing originality scores as well as the greater cost of using the YouTube caption API. For MFK channels, we estimate a somewhat larger reduction in the like/view ratio (-0.131 versus -0.103), though the 95% confidence intervals of both estimates overlap. However, we now estimate a negative impact for mixed channels for both MFK content (-0.044, $p < 0.05$) and non-MFK content (-0.017, $p < 0.1$).

C Time-Varying Treatment Effect Estimates

We estimate a flexible event-study model that allows treatment effects to vary by week in the pre- and post-settlement period:

$$y_{it} = \sum_{t=1}^T \beta_t \cdot Treated_i + \theta_i + \delta_t + \epsilon_{it} \quad (2)$$

where y_{it} is the outcome variable, $Treated_i \in \{MFK_i, Mixed_i\}$, and θ_i and δ_t are channel- and week-level fixed effects. The β_t parameters estimate the mean difference between the treated (MFK and mixed) and control (non-MFK) channels for week t separately, conditional on channel fixed effects. This model allows us to examine the persistence of the treatment effect and to test whether the treatment and control groups follow parallel trends prior to treatment. We normalize


the β_t estimates such that the pre-announcement period estimates are centered around zero.⁴⁵ We include our resulting supply-side (Figure 6) and demand-side (Figure 8) time-varying effect estimates in the main text.

Figures C.1 and C.2 presents our time-varying treatment effect estimates for our objective quality measures: content originality share and manual captioning respectively. Consistent with our Table 3 results, we see no statistically significant changes for MFK or non-MFK videos by mixed channels for either metric. In all cases, we see no estimates that are significantly different from zero in the pre-announcement period. For both outcomes, we see a downward trend in the quality outcomes after the announcement period and a fairly constant, negative treatment effect throughout 2020.

Figure C.3 presents our time-varying treatment effect estimates for the like-view ratio. The pre-trends appear to be parallel. The negative impact on MFK channels appears somewhat delayed relative to other outcomes. Also, the like-to-view ratio falls in a relative but not an absolute sense: the raw trends are flat for MFK channels but increasing post-settlement for non-MFK channels. Our Section 5.2.2 results provide an explanation: MFK channels reduce their investment in content quality as evidenced by our two objective quality measures. Nevertheless, we acknowledge the possibility that non-MFK content may be an inadequate control group for our like-view ratio metric—given the different liking behavior of MFK viewers. Recall from Table 3 that the estimated quality impact for MFK content by mixed channels is positive (0.063) but marginally significant ($p < 0.1$). However, Figure C.3b shows this result attenuates to essentially zero in the second half of 2020, and the equivalent result for our global sample (Table B.2) instead shows a negative and highly significant impact. If the positive impact result for our main sample is indeed real, we speculate that—in pivoting away from MFK content—mixed creators discontinue their inferior MFK content. In other words, the mixed channels’ MFK quality result could reflect a selection effect rather than an increase in their MFK content quality.


⁴⁵We choose this normalization rather than choosing the last week pre-treatment as the baseline, because the weekly level averages are noisy and for consistency with our base model (equation 1).

Figure C.1: Time-Varying Treatment Effect on Content Originality Share




Note: In the figures above, the solid blue lines indicate the point estimates, while the bands show the 95% confidence intervals.

Figure C.2: Time-Varying Treatment Effect on Manual Captioning



Note: In the figures above, the solid blue lines indicate the point estimates, while the bands show the 95% confidence intervals.

Figure C.3: Time-Varying Treatment Effect on Like/View Ratio



Note: In the figures above, the solid blue lines indicate the point estimates, while the bands show the 95% confidence intervals.

D Impact Heterogeneity by Content Category

We consider impact heterogeneity by content category for our key demand- and supply-side outcomes. Table D.1 shows our supply-side heterogeneity estimates by the channel's content category. The entertainment category sees the largest reductions across the board. After the settlement, MFK and mixed channel treatment effect coefficients are -0.185 and -0.168, and the MFK share falls -3.8 and -10.6 percentage points, respectively. In the education category, the treatment effect estimates for MFK channels are -0.107 for video releases and a -2.5 percentage points reduction in the MFK share. Mixed channels in the educational category, however, show no statistically significant effect on content creation (0.009) or MFK share (-3.9 p.p.). In the animation & film category, MFK channels show a statistically insignificant drop in video releases (-0.054) whereas mixed channel production show a statistically significant drop (-0.104). For this category, the MFK channels' MFK share falls -0.5 percentage points and the mixed channels fall -7.8 percentage points. The effects for channels that produce "other" content are imprecisely estimated. However, they exhibit large and statistically significant reductions in the MFK share.

The education category results are particularly interesting, as this category may be seen as the most beneficial to children (see also Kircher & Foerderer, 2023b). The greater impact on ed-

educational and entertainment content production over the film & animation category may reflect different channel economics in these categories. Film and animation content may be expensive to produce, but this category includes film clips, previews, and reviews. So creators in this category may be reusing material from other creators. What is more, some channels in this category are associated with large content producers (e.g., Disney) that use YouTube to promote their movie and TV content. Such large producers may have additional revenue streams so are less vulnerable to a reduction in YouTube revenue. MFK channels in the animation & film category have the smallest reduction in MFK share, which may arise because they are particularly constrained by their audience and content expertise.

Table D.1: Treatment Effects by Channel Content Type

Content Category	log (Videos + 1)		Share of MFK Videos	
	MFK	Mixed	MFK	Mixed
Education	-0.107***	-0.009	-0.025***	-0.039
Entertainment	-0.185***	-0.168***	-0.038***	-0.106***
Animation & Film	-0.054	-0.104**	-0.005**	-0.078***
Other	-0.016	-0.071	-0.074**	-0.093***

Notes: This table presents primary coefficient estimates alone. Difference-in-differences estimates compare outcomes pre-announcement and post-settlement. All specifications include week and channel fixed effects. Robust standard errors clustered at the channel level, and are omitted for space and readability. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

Table D.2 considers heterogeneous demand-side treatment effects by content category. Unlike our supply-side estimates in Table D.1, the entertainment category no longer exhibits the largest reduction among the three main categories in all instances. However, the entertainment category does see the largest reduction in views for the MFK channels (-0.286). Note that the estimated impact coefficients for education MFK channels is -0.204 for views and -0.389 for subscriptions ($p < 0.01$, respectively).

E Content Originality

To evaluate video quality, we quantify the share of original—as opposed to duplicate—content. In Section E.1, we describe our scoring algorithm, which uses video transcripts to compute the share of duplicate content between each video and its channel’s prior released content. We present the algorithm’s performance on two example channels for which we know the ground truth. We

Table D.2: Demand-Side Impact Heterogeneity by Channel Content Type

Content Category	Log (Views + 1)		Log (Subscriptions + 1)	
	MFK	Mixed	MFK	Mixed
Education	-0.204***	-0.307	-0.389***	-0.070
Entertainment	-0.286***	-0.122**	-0.157	-0.312***
Animation & Film	-0.169**	-0.088	-0.629***	-0.369**
Other	-0.378	-0.115	-0.651	-0.059

Notes: This table presents primary coefficient estimates alone. Difference-in-differences estimates compare outcomes pre- and post-January 2020. All specifications include week and channel fixed effects. Robust standard errors clustered at the channel level, and are omitted for space and readability. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

also describe the sample inclusion criteria for this analysis, which is driven by data availability and computational constraints. In Section E.2, we consider robustness of our Table 3 results to alternate definitions of our originality score outcome.

E.1 Originality Score Construction

Our originality scoring algorithm looks for matching text passages between a target video and its channel’s preceding videos. The goal is to detect duplicate content like recompilation videos that reuse a channel’s previously released content. Though we expect some natural content duplication for instance due to opening and closing credits in videos, we expect that such repeated elements represent a small share of total content that is stable over time. Our originality metric ignores within-video repetition, which MFK channels may employ to promote learning. A channel could evade detection by deleting an earlier video, leading us to undercount duplicate content. However, the channel would thereby lose the associated engagement data (e.g., likes and views) and the associated YouTube ranking benefit, which makes this practice unlikely.

We begin by pre-processing the transcript data using the NLTK library.⁴⁶ We remove punctuation and stop words like “the” and “a”. We then use stemming to group words with the same root: e.g., “opens” and “opening” become “open.” These steps are standard data cleaning procedures for natural language processing, which standardize the text and improve model accuracy. Text standardization is particularly important in our setting, as the same content may correspond to slight captioning difference due to captioning inaccuracies (either manual or automated).

Our algorithm’s core step identifies identical passages between a target and a prior bench-

⁴⁶<https://www.nltk.org/>

mark video. We tokenize the transcript into trigrams: e.g., “Mary had a little lamb” becomes the trigrams “Mary had a”, “had a little”, and “a little lamb.” We then identify consecutive matching trigrams between the transcripts. To limit false positives (e.g., commonly occurring phrases like “how are you?”), we only classify a passage as duplicated if it has 8 or more consecutive matching trigrams (i.e., 10 or more words) in the target document. We found that this tuning parameter had a weak, positive relationship with the resulting originality score. We believe that 10 duplicate words in sequence is a conservative threshold that rules out short, common phrases and instead captures longer, more complex phrases. That said, we have tested the model’s performance with different thresholds (from 8 to 16 consecutive duplicated words) and find the model is similarly capable of separating original and duplicate content in the below examples.


The algorithm’s inner loop iterates over a channel’s prior content in release order while applying its core step of identifying identical passages in the target video’s transcript. To improve efficiency, we flag and remove duplicate passages from the target transcript for subsequent iterations. This approach reduces the size of the target document in subsequent iterations at some risk of breaking up longer duplicate passages that may be detected in subsequent iterations. Once complete, the originality score s_v for target video v is given by

$$s_v = 1 - \frac{N_v^{duplicate}}{N_v^{total}}$$

where $N_v^{duplicate}$ denotes v ’s total words arising in duplicate passages and N_v^{total} denotes v ’s total number of words. Finally, the algorithm’s outer loops iterate over all (target) videos that a channel released during our sample period and across all channels.

To assess model performance, we tested the algorithm on all videos uploaded by two example channels. In particular, we focused on a MFK channel that creates songs for children and had released compilation videos containing previously released songs. This channel usually distinguishes its compilation videos by including the phrase “+ more” at the end of the title: e.g., “Mary Had a Little Lamb + more.” This rule imperfectly labels duplicate content in both directions: the channel’s compilations can include new songs and the channel’s single-song videos can present the same song in different ways (e.g., live action versus animation). Nevertheless, this title label provides some ground truth to help assess our scoring algorithm’s performance. Figure E.1a)

Figure E.1: Originality Score: Two YouTube Channel Examples



demonstrates the resulting originality score's by title label. As expected, videos with “+ more” in the title have lower mean originality score (0.31) than the others (0.76). . We manually inspected the two outlier observations with “+ more” in the title and $s_v = 1$: these were not compilations videos, but product-release ads for the channel. In comparison, we also examine a randomly selected non-MFK channel. Casual inspection suggested that this channel did not release compilations. Figure E.1b) confirms that our algorithm's originality scores are consistently high ($s \geq 0.9$) throughout.⁴⁷

We restrict the sample of videos for which we calculate the originality score in three key ways. First, we only examine English-language videos as many of our text analysis tools apply to the English language. Second, we only examine videos that possess caption data. Coverage here is high because YouTube provides automated captions by default. However, some videos such as classical music recordings lack caption data. In many cases, the captions are disabled (presumably

⁴⁷In addition to the channel in Figure E.1a, we examined six total channels to consider exemplars of MFK and non-MFK channels in each of our three content categories. Collectively, the three non-MFK channels possess originality scores above 0.9 for 99.6% of all videos. We found another MFK channel that often releases duplicate content (again, using song compilations) with mean $s = 0.30$, a channel with consistently high originality scores (mean $s = 0.94$), and a third channel with minimal talking that our algorithm scored as having consistently low originality score (mean $s = 0.09$).

by the channel owner) and are unavailable.⁴⁸ Together, these two criteria restrict our sample of videos from about 1.8 million videos down to 1.0 million videos during our sample. This is the video sample used to construct our manual captioning indicator variable (see Section 5.2).

Third, we remove channels with a large number of video uploads due to their computational burden and the cost of data acquisition. We drop the top 5% of channels by total videos uploaded according to Social Blade. These outliers have more than 3,579 video uploads each and collectively account for 46.2% of all videos uploaded in our sample. We remove these outliers for two reasons. YouTube limits the quantity of data that can be pulled per day using its Data API and caption data is particularly costly: each video’s caption data costs 100 credits whereas the remaining video-level data that we collected costs at most 1 credit per video. Also, our algorithm’s time requirement grows essentially quadratically in the number of videos it must evaluate per channel.⁴⁹ This is sufficiently computationally burdensome that we dropped another 25 of the largest channels in our data: our scoring algorithm timed out on our research computing cluster despite repeated attempts. Finally, we obtain our originality score variable for 649,035 videos in our sample.⁵⁰

E.2 Robustness: Content Originality Impact

We consider the robustness of our content originality results to alternate outcome variables in Table E.1. We find that our estimates are qualitatively similar: in particular, all estimates show a statistically significant reduction in original content for MFK channels. First, we use the continuous originality scores defined above. The average originality score falls -0.059 for MFK channels. This is similar to the corresponding estimate for the dichotomized original content share (-0.077). This represents a 8.5% drop relative to the pre-announcement, MFK channel mean of 0.691. Second, we consider alternate thresholds s' to define our original content indicator: i.e., $s' \in \{0.25, 0.5, 0.75, 0.9\}$. The 95% confidence intervals for the resulting MFK channel impact esti-

⁴⁸When pinging the API, we found that 36% of videos did not return a transcript. Of these, the API indicated that no transcript was found in English for one quarter of the cases and that captions were disabled for the remaining three quarters.

⁴⁹To generate originality scores for N videos, this requires $1 + 2 + \dots + (N - 1) = \frac{N-1}{2}N$ video-pair comparisons. For our censoring threshold of 3,579 videos, this implies 6,402,831 comparisons. More generally, we only need to compute the originality scores for the videos that were released during our sample period. If our sample runs from the channel’s N th to M th videos, this requires $(N - 1) + N + (N + 1) + \dots + (M - 1) = \frac{M-N+1}{2}(N + M - 2)$ comparisons.

⁵⁰Compared to our sample of 1.0 million videos with English transcripts, more than 99% of the missing scores result from excluding the top 5% of channels by video uploads. 0.6% of missing scores result from videos that contained only stop words and were therefore empty after our data pre-processing. Only 0.2% of missing scores results from dropping the 25 channels that timed out on our research computing cluster.

Table E.1: Impact on Video Quality Robustness: Difference-in-Differences Estimates

Score thresholds	Originality score ^a	Original content share with different thresholds ^b			
	NA	$s \geq 0.25$	$s \geq 0.5$	$s \geq 0.75$	$s \geq 0.9$
MFK	-0.059***	-0.054***	-0.077***	-0.076***	-0.073***
× Post-Settlement	(0.008)	(0.008)	(0.010)	(0.011)	(0.012)
Mixed (MFK videos)	-0.020*	-0.005	-0.021	-0.040**	-0.034**
× Post-Settlement	(0.011)	(0.010)	(0.014)	(0.017)	(0.016)
Mixed (non-MFK videos)	0.000	0.001	-0.003	-0.003	0.002
× Post-Settlement	(0.004)	(0.004)	(0.005)	(0.006)	(0.007)
MFK	-0.026***	-0.022***	-0.030***	-0.032***	-0.034***
× Anticipation Period	(0.007)	(0.007)	(0.009)	(0.010)	(0.011)
Mixed (MFK videos)	-0.002	0.004	-0.001	-0.016	-0.012
× Anticipation Period	(0.008)	(0.007)	(0.009)	(0.012)	(0.013)
Mixed (non-MFK videos)	0.004	0.006	0.003	0.001	0.007
× Anticipation Period	(0.004)	(0.004)	(0.005)	(0.006)	(0.007)
Fixed effects: week, channel-video type	Y	Y	Y	Y	Y
Adj. R^2	0.679	0.554	0.630	0.638	0.612
N	254,700	254,700	254,700	254,700	254,700

Notes: The originality results restrict the sample notably to exclude the top 5% of channels by cumulative video releases. See Appendix E for details. Post-Settlement is defined as weeks after Jan. 1, 2020. Anticipation Period is defined as weeks between Sep.4, 2019 to Jan.1, 2020. All specifications include week and channel fixed effects and omit the few observations with MFK (non-MFK) videos on non-MFK (MFK) channels. Robust standard errors clustered at the channel level are in parentheses. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

^aUses the raw originality score as described in Appendix E.1. ^bThe original content share takes the channel-weekly average of a video indicator for content with originality score $\geq s'$ for different thresholds s' .

mates overlap substantially across all thresholds. The $s' = 0.25$ threshold yields a somewhat lower estimate (-0.053), but the rest closely resemble the estimate for our preferred ($s' = 0.5$) threshold. Though our preferred threshold ($s' = 0.5$) yields the largest estimate among the four considered thresholds, this was incidental: we selected this threshold beforehand for interpretability (i.e., majority original content). Third, note that mixed channels' MFK content now shows a marginally significant drop ($p < 0.1$) in the originality score and a significant drop ($p < 0.05$) in original content share for higher thresholds ($s' \in \{0.75, 0.9\}$). This suggests that mixed channels increased their reliance on duplicate MFK content to a lesser extent than MFK channels and did so by including only some duplicate content in their videos rather than relying on compilation videos with little original content.

F Marginal effect calculation

Below, we describe how we calculate the marginal effects for our $\log(y + 1)$ transformed outcomes. Our marginal effects calculation computes the proportional treatment effect. Adapting the potential outcome notation from Chen & Roth (forthcoming, section 5.2), the average proportional treatment effect on the treated for MFK channels is given by

$$\theta_{ATT\%} = \frac{E[y_{it}(1) | MFK_i = 1, Settlement_t = 1] - E[y_{it}(0) | MFK_i = 1, Settlement_t = 1]}{E[y_{it}(0) | MFK_i = 1, Settlement_t = 1]}$$

where $E[y_{it}(1) | MFK_i = 1, Settlement_t = 1]$ represents the average outcomes among treated units (i.e., $MFK_i = 1$) in the post-period (i.e., $Settlement_t = 1$) and $E[y_{it}(0) | MFK_i = 1, Settlement_t = 1]$ represents the average counterfactual outcomes (in the absence of treatment) among the treated units in the post-period. The expression for mixed channels is analogous.

We observe $E[y_{it}(1) | MFK_i = 1, Settlement_t = 1]$ in the data: this is the average outcome variable among MFK channels in 2020. However, we do not observe the counterfactual average $E[y_{it}(0) | MFK_i = 1, Settlement_t = 1]$. We instead use our model estimates to infer this quantity, by essentially subtracting our (constant) treatment effect estimate from the observed outcomes.⁵¹ For MFK channels, this is given by

$$\begin{aligned} E[y_{it}(0) | \widehat{MFK_i = 1}, Settlement_t = 1] &= E\left[\frac{1 + y_{it}}{\exp(\beta_1)} - 1 | MFK_i = 1, Settlement_t = 1\right] \\ &= \frac{Avg[1 + y_{it} | MFK_i = 1, Settlement_t = 1]}{\exp(\beta_1)} - 1 \end{aligned} \quad (3)$$

where β_1 is the difference-in-differences coefficient estimate from equation (1). The calculation for mixed channels is analogous and instead uses the set of mixed channels, post-settlement and the corresponding β_2 coefficient. To construct 95% confidence intervals—e.g., for Figures 7a and 9—we replace the relevant coefficient estimate in equation (3) with the upper (lower) limit of the interval. Due to this transformation, the resulting confidence intervals are asymmetric about the mean.

⁵¹Chen & Roth (forthcoming) instead estimate a related model using the ratio version from Wooldridge (2023).

G Advertising on YouTube MFK content

Though this is out of scope for our study, the YouTube settlement also impacted the ads that MFK viewers see. We lack pre-settlement data to study this, but collected our own post-settlement survey. We discuss this survey and the advertising dimension of the YouTube settlement below. This provides additional institutional context for the reader.

G.1 Post-Settlement Ad Survey

We conducted our own informal survey of ads on MFK content in 2023 and found that a minority of these featured products that appeared to be directed at children (e.g., toys, children’s movies). The remaining ads featured products that are better suited for adults including ads for vehicles, telecommunications, and tax preparation tools.

We selected 300 MFK videos⁵² at random from our sample. Using an incognito Chrome browser in the United States in early 2023, we visited each video once, returned to half of the videos a second time, and recorded data on all the pre-roll advertisements in both cases. YouTube showed ads on only 30 of these videos but showed two pre-roll ads on 22 of these. The videos with ads changed between the first and second visits, suggesting a probabilistic component to ad allocation. Most of the ads were skippable (after 5 seconds): more than half (15/24) of the 15-second ads and all of the 30-second or longer ads.

Our survey identified 52 total ads inclusive of repeats. We classified 12 (23%) of these ads as well suited to children. The remaining 40 ads were better suited to adults or a general audience. We list the advertised products by our categorization below:

- *Ads targeting kids:* Super Mario Bros. Movie (kids movie) (2 times), Disney Plus (streaming subscription), Zuru X Shot (toy gun), Hello Thinkster (math tutoring for kids), Elmer’s Squishies (slime toy), Disney World (theme parks) (4 times), Thomas and Friends (show for kids, full episode), NFL Flag (flag football for kids through the NFL), and X-Shot Skins (toy gun).
- *Ads targeting adults or a general audience:* Morgan and Morgan (injury law firm) (8 times), ADP

⁵²23 of these were unavailable: 3 were labeled as “private” and 20 were “unavailable.”

(HR payroll software) (2 times), Turbo Tax (tax software), Wix (small business sales platform) (3 times), Xfinity (mobile service), AT&T (cell phone), Samsung (tablets and phones) (2 times), Enterprise (car rental), Grammarly (spelling and grammar checking software) (3 times), Veterinary Emergency Medicine Group (veterinary emergency medicine group), Hefty (trash bags), Lowe’s (hardware and appliance store), Gator Album by Pouya (rap music album), Square Space (website builder), LL Flooring (hardwood floors), Old Spice (body wash and hygiene products), Instrumart Flow Meter (a meter for industrial purposes), Honda CR-V Hybrid (SUV car) (3 times), Acura MDX (SUV car), Interstate Batteries (car battery), Canva (graphic design software), and Roofing Leads (roofing).

G.2 Discussion

In addition to our survey, more evidence suggests that MFK viewers see ads on YouTube that target a general or older audience. Surveys of YouTube’s MFK content found age-inappropriate ads on 20% of videos (?) and 6% of ads (Yeo et al., 2021), including ads for politics, lingerie, alcohol, or containing violence. YouTube has since banned age-inappropriate ad topics from being displayed on MFK content.⁵³

Advertisers may show general audience ads in MFK content for several reasons. First, advertisers may view MFK content as an arbitrage opportunity: their target audience may be indirectly exposed often enough at MFK’s lower ad price to be worthwhile. Second, advertisers may be inattentive or confused about excluding MFK content from their campaigns. YouTube is a closed platform where advertisers specify their campaign parameters to Google to place ads on their behalf. Advertisers (or their agents) can opt to exclude MFK content, but MFK content is included by default. Moreover, Google’s related setting is labeled “content suitable for families”,⁵⁴ which may be confusing. Google recently confirmed that this setting includes MFK content,⁵⁵ but the extent of additional family-friendly content is unclear.

While all the above advertisement evidence is taken after the settlement, we expect that children often saw ads targeted at older users prior to the settlement as well. Children under the age

⁵³<https://support.google.com/adspolicy/answer/9683742>; <https://support.google.com/youtube/answer/9713557>.

⁵⁴<https://support.google.com/google-ads/answer/12764663>.

⁵⁵<https://blog.google/products/ads-commerce/our-strict-privacy-standards-around-made-for-kids-content/>.

of 13 may use their parents’ devices to watch YouTube, and we expect that children spend relatively little time on commercially-relevant websites. Moreover, we conjecture that child-directed content was harder for advertisers to avoid before YouTube developed MFK-detection algorithms and publicly identified MFK content. For these reasons, we suspect that personalized advertising on MFK content prior to 2020 targeted the parents—rather than their children—to a significant extent.

We suspect that ad prices on MFK content would have been lower if advertisers made greater effort to avoid MFK content. By showing adult and general audience ads on MFK content, advertisers are effectively subsidizing MFK content. The consequence of this cross-subsidization for the decline in ad prices post-settlement is unclear, however. We conjecture that ad prices fell by more than they would on a platform whose audience is kids alone: parents have greater spending power and spend more time on commercially-relevant websites, so personalized advertising is particularly valuable when children receive ads directed at their parents.

In August 2023, YouTube faced renewed criticism for its advertising on MFK content.⁵⁶ Adalytics, an ad technology transparency firm, criticized YouTube for the adult and general audience ads appearing on MFK content. Moreover, Adalytics alleged that behaviorally targeted ads were appearing on MFK content. Google refuted the latter claim and explained that behaviorally-targeted ads can appear on the non-MFK videos of majority-MFK channels.⁵⁷ Google also explained that its “affinity audience segments” also included contextual targeting, and could therefore reach users on relevant MFK videos.

Finally, we may expect that lower ad prices on YouTube’s MFK content may have induced some entry by child-relevant advertisers. However, YouTube now restricts a number of industries from showing ads on YouTube Kids and has indicated that these policies extend to MFK content on YouTube.⁵⁸ These industries include age-sensitive media content, beauty and fitness, dating or relationship, food and beverages, illegal or regulated products, online or virtual communities, political ads, religious ads, and video games.⁵⁹

⁵⁶<https://www.nytimes.com/2023/08/17/technology/youtube-google-children-privacy.html>.

⁵⁷<https://blog.google/products/ads-commerce/our-strict-privacy-standards-around-made-for-kids-content/>.

⁵⁸[https://www.markey.senate.gov/imo/media/doc/Response to Sen Markey_Dec13 2019.pdf](https://www.markey.senate.gov/imo/media/doc/Response%20to%20Sen%20Markey_Dec13%202019.pdf)

⁵⁹<https://support.google.com/youtube/answer/6168681>.

H YouTube API Video Data

We obtain our video-level data from the YouTube Data API. YouTube’s Data API provides rich data on YouTube’s platform.⁶⁰ However, the data are current as of the time YouTube provides the data: YouTube does not offer historical data. We discuss some resulting implications for our data below.

We obtained our main video data sample as of July 2022. This first data pull contains 5,541 channels that Social Blade categorizes as being both produced by channels in the United States and belonging to either “film & animation,” “education,” or “entertainment” categories. Our main sample represents the 91.4% (5,066) of these channels that release a video during our pre-announcement period. In December 2022, we again used Social Blade’s list of top 100,000 YouTube channels (by subscribers as of June 2021) to pull the YouTube video data this time for all channels. In both cases, we downloaded the data associated with videos released during our sample period of July 1, 2018 to December 31, 2020.

First, YouTube does not provide data for channels that are no longer on its platform. This can arise because YouTube has banned the channels or because the creator deleted the channel. From our global channel pull, we noted that 17,730 (17.7%) of channels were missing. Though this affects the interpretation of the results, removing banned channels has some advantages. YouTube may ban a channel for copyrights violation. However, we are focused on original content creation rather than content pirating. YouTube may also ban content that violates its terms of service in other ways. This can have the advantage of removing violent and objectionable content. For instance, YouTube faced a wave of criticism around 2017 (i.e., “Elsagate”) for hosting videos with child-friendly characters engaging in child-inappropriate behaviors.⁶¹ YouTube removed (or demonetized) many such channels.⁶² Between pulls, we found that 103 channels in our main sample were missing in the second pull. Nevertheless, our main sample employs the data as of July 2022—attempting to strike a balance between these advantages of data deletion and the representativeness of channels during our time period of study.

⁶⁰YouTube offers a program for researchers (<https://research.youtube/>), which is particularly helpful for researchers seeking to obtain large quantities of API data.

⁶¹<https://web.archive.org/web/20181017005522/https://medium.com/@jamesbridle/something-is-wrong-on-the-internet-c39c471271d2>.

⁶²<https://web.archive.org/web/20171128192652/https://www.buzzfeed.com/blakemontgomery/youtube-has-deleted-hundreds-of-thousands-of-disturbing>.

Second, the video-level data can also exhibit deletion and modification. The half-year gap between our data pulls helps us quantify these issues. We observe some “melt” in the data—i.e., videos disappearing from the Youtube platform—between data pulls. However, 94.5% of the videos observed in our main sample (from July 2022) are still observed in December 2022.⁶³ The data melt appears unrelated to whether the video is MFK: 95.6% of MFK videos and 94.3% of non-MFK videos remain for the later data pull. Finally, the data melt is concentrated in a few channels: 62% of channels have the same number of videos, while only 11.5% channels have a difference larger than 5 videos with respect to the original pull.

We also observe some MFK relabeling between data pulls at the video level, though this is rare. In particular, 4,667 videos were marked as non-MFK in the original pull but relabeled as MFK in the second pull and 1,855 videos were marked MFK in the original pull but relabeled as non-MFK in the second pull. In total, 6,522 videos were relabeled, representing less than 0.4% of the sample. Nonetheless, deletions and MFK relabeling have a negligible impact on the share of MFK videos in our main sample between the first (11.04%) and second (11.32%) data pulls.

I Additional Exhibits

Below, we present additional exhibits related to summary statistics as well as our demand- and supply-side heterogeneity analyses.

Table I.1: Summary Statistics: Pre-Announcement versus Post-Settlement Means

	Pre-Announcement Means			Post-Settlement Means		
	MFK	Mixed	NFK	MFK	Mixed	NFK
Weekly Video Releases	1.861	3.842	2.629	1.560	3.696	2.840
MFK Share of Releases	1	0.318	0	0.969	0.199	0.003
Original Content Share	0.672	0.904	0.966	0.572	0.896	0.955
Manual Caption Share	0.143	0.142	0.159	0.086	0.116	0.129
Like/View Ratio	0.004	0.016	0.032	0.005	0.021	0.038

Notes: Conditional mean values for the main sample of 5,066 American top YouTube channels where the unit of observation is a channel-week. All variables related to video content (e.g., MFK share and like/view ratio) omits weeks in which the channel releases no videos and represents the average across videos when the channel releases multiple videos in a week. Original content is defined as videos with an originality score over 0.5 (see Appendix E for details).

⁶³Some video additions also appear, though these are rare. 4,825 videos appear in the second pull, but not in the original pull.

Figure I.1: MFK Video Share Histogram among Mixed Channels




Table I.2: Treatment Effects by Channel Subscription Size

Channel Subscription Quartile	log (Videos + 1)		Share of MFK Videos	
	MFK	Mixed	MFK	Mixed
0th-25th percentile	-0.186***	-0.212***	-0.014***	-0.048***
25th-50th percentile	-0.098**	-0.160***	-0.020***	-0.107***
50th-75th percentile	-0.125***	-0.084*	-0.055***	-0.100***
75th-100th percentile	-0.081*	-0.081**	-0.018***	-0.117***

Notes: This table presents primary coefficient estimates alone. Difference-in-differences estimates compare outcomes pre-announcement and post-settlement. All specifications include week and channel fixed effects. Robust standard errors clustered at the channel level, and are omitted for space and readability. ***significant at 1% level; **significant at 5% level; *significant at 10% level.

Table I.3: Demand-Side Impact Heterogeneity by Channel Subscription Size

Channel Subscription Quartile	Log (Views + 1)		Log (Subscriptions + 1)	
	MFK	Mixed	MFK	Mixed
0th-25th percentile	-0.441***	-0.239**	-0.622***	-0.486***
25th-50th percentile	-0.216***	-0.215**	-0.312**	-0.421***
50th-75th percentile	-0.246***	-0.115*	-0.302*	-0.144
75th-100th percentile	0.015	0.081*	0.021	0.003

Notes: This table presents primary coefficient estimates alone. Difference-in-differences estimates compare outcomes pre- and post-January 2020. All specifications include week and channel fixed effects. Robust standard errors clustered at the channel level, and are omitted for space and readability. ***significant at 1% level; **significant at 5% level; *significant at 10% level.