// NO.23-059 | 12/2023

DISCUSSION PAPER

// RAPHAELA ANDRES, MICHELANGELO ROSSI, AND MARK TREMBLAY

YouTube “Adpocalypse”: The Youtubers’ Journey From Ad-Based to Patron-Based Revenues

YouTube “Adpocalypse”: The YouTubers’ Journey from Ad-Based to Patron-Based Revenues*

Raphaela Andres, Michelangelo Rossi, and Mark Tremblay†

December 3, 2023

Abstract

In the past decade, the Creator Economy has witnessed unprecedented growth. This dynamic ecosystem thrives on a multi-sided business model, connecting content creators, users, and advertisers. However, matching the needs of different stakeholders is a complex challenge, as evidenced by the impact of the YouTube “Adpocalypse” in 2017, when major advertisers fled Youtube due to concerns about their ads appearing alongside objectionable content. This paper explores the response by content creators that use both Youtube and Patreon to YouTube’s content moderation policies following the “Adpocalypse”. We find that these content creators shift their efforts toward Patreon which uses a subscription fee model instead of an ad-based model; as a result, consumers subsequently increase their use of Patreon through memberships, comments, and likes. However, we also find that Youtube’s content moderation, and the shift by content creators and consumers that follows, results in an increase in toxicity on Patreon.

JEL Codes: L10, L20, L82.

Keywords: Patreon, Platform Competition, Multi-homing, Content Creators.

---

*We would like to thank Anna Kerkhof for providing helpful feedback and Irene Bertschek, Reinhold Kesler, Ulrich Laitenberger, Leonardo Madio, Mary O’Mahony, Joel Waldfogel, Mike Ward, and Michael Zhang for helpful comments. All errors are our own.

†R. Andres: Télécom Paris, Institut Polytechnique de Paris, France; and ZEW Mannheim, Germany. E-mail: Raphaela.Andres@zew.de. M.Rossi: Télécom Paris, Institut Polytechnique de Paris, France. E-mail: michelangelo.rossi@telecom-paris.fr. M.Tremblay: Lee Business School, University of Nevada, Las Vegas, USA. E-mail: mark.tremblay@unlv.edu.

1 Introduction

The Creator Economy has experienced an unparalleled increase in growth, fundamentally altering how individuals, ranging from artists to influencers, build careers and achieve financial autonomy. Over the past few decades, this dynamic ecosystem has flourished, and forecasts suggest it will continue to drive a multi-billion-dollar industry (Florida, 2022; El Sanyoura and Anderson, 2022). At the core of this transformation are content creators, whose contributions are fundamental to the platform business model. Unlike media traditional industries where content creation was centralized, modern digital platforms entrust this task to an array of independent content creators (Bhargava, 2022). One of the first platforms pioneering content creation at scale is YouTube, which currently boasts more than 100 million channels and over 500 hours of content uploaded each minute.[1]

Platforms like YouTube, TikTok, or Twitch have thrived on a multi-sided business model (Cusumano, Gawer and Yoffie, 2019), capitalizing on the cross-network effects of three primary stakeholders: content creators, users, and advertisers. Content creators craft content that captivates users who not only consume this content but also engage with the advertising embedded in it. Advertisers, in turn, benefit from the connections with users and compensate the platform which then rewards the creators. However, the viability of this model depends on the satisfaction of multiple actors and the proper matching of users, content, and advertising.

To meet the challenges of this type of business model, some online platforms have recently explored ways to encourage consumers to pay directly for content (thus reducing dependence on the advertising side) or to have greater control over the content produced by third-parties (in order to properly match their content with advertising and user needs). In the latter case, the platform assumes the role of overseeing and curating the content produced by creators. This intricate interplay between enabling content creators and controlling the ecosystem underscores a critical trade-off within the Creator Economy as it has been described by Boudreau and Hagiu (2009), Hagiu and Wright (2015), and Hagiu and Wright (2019).

These important connections across users suggests that platform strategies for innovating and adjusting multi-sided business models cannot operate in isolation. Content creators face low barri-

[1] See: https://blog.YouTube/press

ers to entry on these platforms allowing them to switch between, or establish a presence on, multiple platforms simultaneously. Hence, when a platform aims to explore changes in its business model, such as increasing or decreasing content moderation or curation, it is critical for platforms to analyze the response of content creators. This includes their potential migration to another platform or increased efforts in creating content on competing platforms.

In this paper, we offer empirical evidence of how content creators and users react to changes in a platform's moderation policies. Our focus centers on two of the most successful platforms for content creators: YouTube and Patreon. YouTube primarily monetizes content through advertising revenue. Creators can enable ads on their videos and they earn a share of the revenue generated from ad views and interactions. Patreon is a platform designed to help creators generate income directly from their supporters or "patrons" through monthly subscriptions. Creators on Patreon offer exclusive content to their subscribers in exchange for recurring monthly payments. Patreon primarily serves as a "membership portal" and does not actively direct users to discover other creators.[2] As a result, content creators are often required to establish a presence on other platforms, such as YouTube or Facebook, in order to attract users who might otherwise remain unaware of their content. As depicted in Figure 1, many content creators on YouTube actively guide their audience to their Patreon page by explicitly mentioning it within their videos or in the video description. In addition, Patreon's content moderation policies are much more permissive and do not have the imperative to accommodate the needs of advertisers by embedding content that aligns with their values, allowing for a more creator-centric environment.

To determine how content moderation impacts content creator efforts across platforms that offer different revenue models, we take advantage of one of the most notable advertising boycotts in YouTube's history: the YouTube "Adpocalypse". In 2017, several institutional advertisers discovered that their YouTube ads were appearing alongside content that was inappropriate or extremist. This realization led to a mass exodus of major advertisers, resulting in a significant drop in advertising revenue. In response to this advertising boycott, YouTube introduced stricter content moderation policies in January 2018. These changes aimed to improve the platform's ad-friendly environment by implementing stricter review and demonetization procedures for content creators

[2] To better understand how the platform relies on users knowing the exact name of the creator they want to support, please see: https://support.patreon.com/hc/en-us/articles/360028768352-Find-creators-on-Patreon.

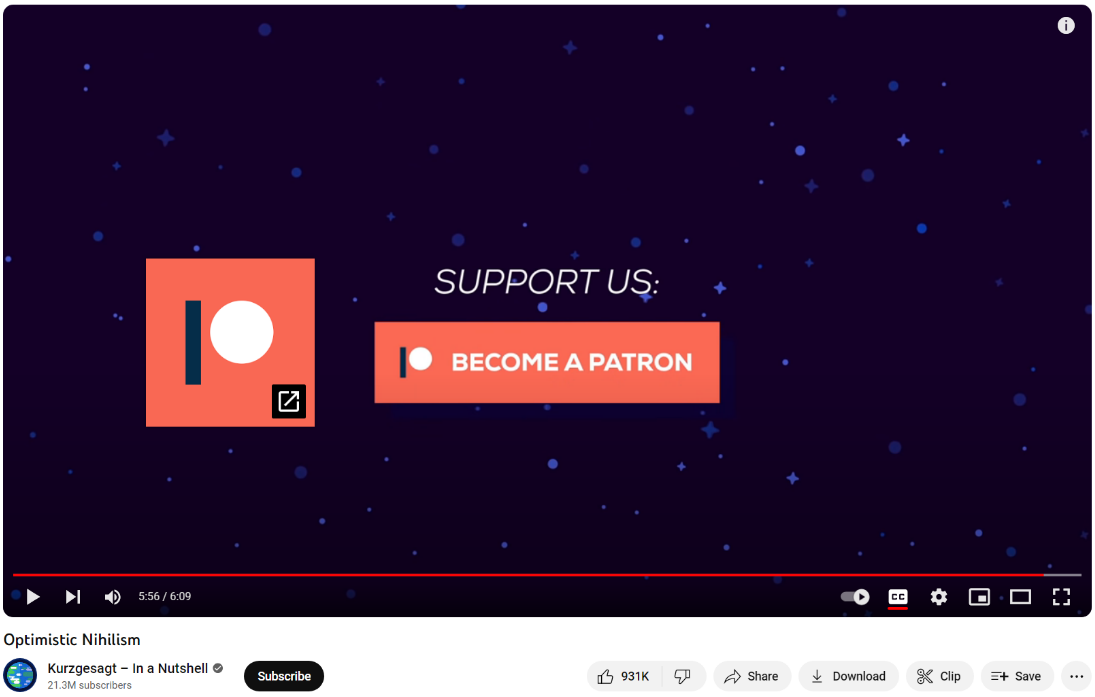

(a) Reference to Patreon during a Video

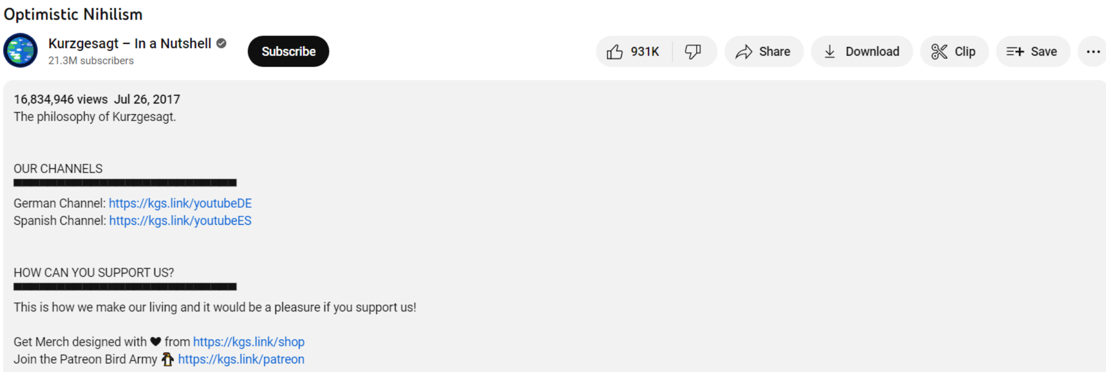

(b) Reference to Patreon in the Video Description

Figure 1: References by a YouTube Content Creator (Kurzgesagt) to their Patreon webpage

who violated the new guidelines. Our analysis focuses on content creators who are active on Patreon, and we examine changes in content provision before and after the YouTube "Adpocalypse" between creators who simultaneously use both YouTube and Patreon and those who operate only on Patreon (not YouTube).

To inform our empirical strategy, we start by developing a simple theoretical model that studies how content creators allocate their efforts across two potential platforms that offer different monetization features: YouTube, with an ad-based monetization, and Patreon, which relies on users' monthly subscriptions. On both platforms, content creators cater to audiences with different preferences for content toxicity. At the same time, YouTube's content moderation policy can change. It can be lenient, where content toxicity has no direct impact on a creator's revenue, or it can be more controlling, where a creator's content toxicity can negatively impact their revenue. In contrast, the level of toxicity on Patreon does not have a direct impact on revenue since contents are funded and displayed privately. Our theoretical framework shows that in response to YouTube's shift from a lenient to a more controlling moderation policy, content creators redirect their efforts and engagement to Patreon. As a direct result, the level of toxicity decreases on YouTube and increases on Patreon.

Following the same approach as the model, we examine Patreon content creators in the "Video" category from August 2017 to August 2018. We focus on content creators who multi-home (participate on both Youtube and Patreon) and thus are affected by YouTube's content moderation. Using a difference-in-differences (DiD) design, we estimate the causal effect of YouTube's content moderation shock by comparing these multi-homing creators to those who are never active on YouTube (and thus form a control group for our analysis). To do this, we use data from Graphtreon, a site that regularly scrapes Patreon, and we supplement this data with our own scraping of Patreon.

With this empirical design, we validate our model's predictions and uncover additional insights relevant to platforms, creators, and policymakers. First, we find that in the aftermath of the YouTube "Adpocalypse", content creators who engage in multi-homing increase the amount of subscriber-only content they offer on Patreon. In addition, the number of Patreon subscriptions increases. These findings are consistent with the notion that following YouTube's moderation policy change, content creators are motivated to shift their efforts and user base to Patreon as its relative profitability increases. This observation is consistent with the study by El-Komboz, Kerkhof and

Loh (2023), which examines how the same shock affects YouTube content creators. In addition, our research shows that the number of likes on subscriber-only content also increases, suggesting that multi-homing content creators not only produce a greater quantity of content, but also improve its quality.

For platforms, these results suggest that the different monetization strategies for content creators present a form a platform differentiation (similar to a result found in Casner and Teh (2022)). For an incumbent platform, implementing restrictions to the existing monetization strategy will damage the content creator side of the market and this can reduce consumer engagement through the indirect network effects that connect the two groups. However, these losses could be mitigated by simultaneously implementing alternative monetization strategies; for example, if YouTube had also implemented a subscription model at the time of the policy change that does not follow the same restrictions required to advertise.[3] For content creators, our results highlight how multi-homing increases flexibility in monetization which can insulate profits from negative platform level shocks.

We also measure content toxicity by analyzing the titles of the Patreon content employing pre-trained algorithms provided by Jigsaw and Google's Counter Abuse Technology team (Mondal, Silva and Benevenuto, 2017; ElSherief, Kulkarni, Nguyen, Wang and Belding, 2018; Han and Tsvetkov, 2020). Doing so, we can measure toxicity of free contents (available also on other platforms such as YouTube) and subscriber-only contents and this helps us understand how the toxicity of different content evolves before and after the YouTube "Adpocalypse". Our results suggest that the overall toxicity of all subscriber-only contents (only accessible via Patreon) increases after the shock. Accordingly, and confirming our theory, toxicity migrates to Patreon in response to YouTube's content restrictions. This offers entrant platforms another avenue to increase growth as smaller platforms are less likely to come under scrutiny for such content (until they become sufficiently large). From a policy perspective, this suggests that targeting large online companies may not actually reduce internet wide toxicity as the creators and viewers can simply migrate to another platform. In fact, this migration of controversial content to another platform partially mirrors the emergence of the platform "Truth Social" following the Donald Trump Twitter ban.

[3] YouTube has a "membership program" that is similar to a subscription service, but its restrictions are consistent with the requirements for earning revenue from advertising.

2 Literature Review and Contribution

Broadly speaking, our work contributes to the literature on platform competition in two-sided markets, where content creators generate income through consumer involvement and thus create the indirect network effects that connect the two groups. However, this literature, founded by Rochet and Tirole (2003), Parker and Van Alstyne (2005), and Armstrong (2006), often assumes that platforms are symmetric in size and monetization strategies which is not the case in our context. Furthermore, this literature often concludes that multi-homing users are worse-off than single-homing users, in contrast with our findings.

Only a few empirical studies examine cases of competition between asymmetric platforms. For example, Farronato and Fradkin (2022) reveal that the differentiation between Airbnb listings and hotel rooms can result in substantial welfare gains from home-sharing when cities host large events that generate an influx of tourism beyond the local hotel capacity. This highlights how a platform entrant that differentiates itself from the industry incumbent is better positioned to increase its market share. Li and Zhu (2021) consider how daily deal platforms might induce high performance users on a rival to multi-home and how platforms try to prevent other platforms from inducing their users to multi-home. While no such actions were taken by YouTube or Patreon in our setting, our work complements theirs by revealing how alternative business models can also play an important role in benefiting the users that multi-home.

One paper that closely relates to ours is El-Komboz et al. (2023) who consider the direct impact of the YouTube "Adpocalypse‘" on YouTube content creators. Using a regression discontinuity design, they found that content creators who no longer qualify for platform monetization channels produce both fewer pieces of content and lower quality content compared to creators who remained eligible. Although they do not consider platform competition or creator multi-homing, their results are in line with our own findings, indicating a shift in content creator efforts from YouTube to Patreon in response to YouTube's monetization restrictions. Another paper that is similar to ours is Kesler (2022). He shows that Apple's change in iOS policy which reduced third-party apps' ability to track consumer locations results in third-party apps pivoting their monetization strategies away from within app advertising and shift it toward payment for or within apps. Combined with our findings, these results suggest that flexibility in monetization strategies by third-party contributors can help improve third-party profitability in the face of policy restrictions by the platform.

Naturally, our work on content creators contributes to the literature on platforms that manage user-generated content.[4] For example, Sen, Grad, Ferreira and Claussen (2023) determine how different moderation strategies impact platform content when content production occurs through professionals and through user-generated content. Similarly, Paridar, Ameri and Honka (2023) consider how monetary and non-monetary rewards can incentivize specific types of content on a user-generated content platform for board games. Unlike the two previous studies (and similar to our setting), Wlomert, Papies, Clement and Spann (2023) take an across platform approach to considering how previously protected music that becomes available on a platform that hosts user-generated content can reduce artist and music producer revenues through demand cannibalization. We contribute to this literature by revealing how a content moderation policy on one platform will impact the production and consumption of user-generated content on another platform that offers content creators with an alternative way to monetize their content.

Lastly, our paper contributes to the strand of literature analyzing content moderation and the development of harmful online content. Content moderation has been shown to be effective in reducing harmful content, being through voluntary platform measures (Chandrasekharan, Pavalanathan, Srinivasan, Glynn, Eisenstein and Gilbert, 2017, Srinivasan, Danescu-Niculescu-Mizil, Lee and Tan, 2019) as well as stipulated by law (Andres and Slivko, 2021). However, most studies stay on the same platform when analyzing content moderation efforts. Remarkable examples include Ali, Saeed, Aldreabi, Blackburn, De Cristofaro, Zannettou and Stringhini (2021) and Rauchfleisch and Kaiser (2021), who analyze the phenomenon of deplatforming, i.e. users migrating to another platform with laxer moderation efforts. Both studies find that toxic users tend to migrate or multi-home to other platforms, and even increase their toxicity on these platforms. Alternatively, Bhargava (2023) considers how taxing advertising-based platforms to subsidize subscription-based platforms impacts platform content moderation practices. In line with these studies, our paper documents an increase in toxicity on the platform Patreon triggered by content moderation efforts on YouTube.[5]

[4]See Luca (2015) for an overview of the literature on user-generated content.

[5]Lefouili and Madio (2022) and Teh (2022) consider the platform's perspective of moderating the creator side and how such restrictions might impact the overall quality of the platform, creator competition within the platform, and platform liability from content that harms society.

3 Theoretical Framework

Consider a simple model for how content creators (CCs) allocate their content efforts across two potential platforms (YouTube and Patreon) that have different monetization features. We assume that the cost of content production is zero but that every CC is limited in the amount of effort they can exert to develop content across platforms. More specifically, let E denote the total amount of effort that a CC can offer, let e_{YT} denote the amount of effort toward YouTube, and let e_P denote the amount of effort toward Patreon so that $E = e_{YT} + e_P$.[6]

In terms of revenues, YouTube and Patreon have different business models and pay their CCs differently. YouTube pays for content through advertising revenues that are split between YouTube and its CCs. Let $r_{YT}(t_{YT}, \mathbb{B})$ capture the marginal advertising revenue from one unit of effort on YouTube, where t_{YT} denotes the level of toxicity of the CC on YouTube and $\mathbb{B} \in \{0,1\}$ denotes the leniency of the YouTube moderation policy. If $\mathbb{B} = 0$, then the YouTube moderation policy is lenient, and t_{YT} has no impact on r_{YT}. If $\mathbb{B} = 1$, YouTube is more likely to control and demonetize toxic content. Thus, a greater t_{YT} reduces r_{YT}. In practice, CCs have control over the content that they post; at the same time, CCs differ across audience toxicity preferences. To model these features, suppose that each CC receives a toxicity draw given by $\tau \sim F(\cdot)$ so that deviations away from this toxicity level are costly to the CC in the form of $(\tau - t_{YT})^2$ for content on YouTube.

On Patreon, CCs offer a menu of services for different membership fees that are paid directly by consumers. The ability to offer a menu of services allows CCs to price discriminate so that the first few units of effort on Patreon go toward the most profitable forms of content. To model this environment simply, suppose that the marginal revenue from effort diminishes on Patreon so that the marginal revenue at effort e is given by $r_P(e) = \rho - e$, where $\rho \sim G(\cdot)$ is a random draw that captures the maximum revenue generated from initial effort on Patreon which, like toxicity τ, can differ across CCs.[7] This implies that a CC's total revenue from Patreon when exerting effort level e_P is given by $\int_0^{e_P}(\rho - e)de$. Note that toxicity does not impact revenues on Patreon because Patreon is a pure "membership portal" that, to this point, has not faced scrutiny regarding toxicity on its platform. Much like on YouTube, we allow for CCs to choose the toxicity level on Patreon

[6] Note that the main predictions hold if we allow CCs to differ in the total amount of effort they exert.

[7] Naturally, CCs are also heterogenous in the amount of views they earn on YouTube. We avoid heterogeneity on YouTube so that we can make comparative statics with respect to r_{YT}. However, additionally allowing for YouTube heterogeneity does not change the predictions of our model.

so that the cost of setting toxicity t_P on Patreon is given by $(\tau - t_P)^2$ for content on Patreon.[8]

Combining the revenue streams from YouTube and Patreon, we see that the CC maximization problem is given by

$$\max_{e_{YT}, e_P, t_{YT}, t_P} \quad r_{YT}(t_{YT}, \mathbb{B}) \cdot e_{YT} - (\tau - t_{YT})^2 + \int_0^{e_P} (\rho - e) de - (\tau - t_P)^2$$
$$\text{s.t. } E = e_{YT} + e_P.$$

Solving allows us to determine equilibrium effort levels across platforms as well as CC level toxicity across platforms. In terms of effort, we have the following result:

Proposition 1. *In equilibrium, CC effort levels are given by* $e_P^* = \max\{0, \min\{E, \rho - r_{YT}\}\}$ *and* $e_{YT}^* = \max\{0, \min\{E, E + r_{YT} - \rho\}\}$.

In other words, an interior solution occurs with $e_P^* = \rho - r_{YT}$ whenever $(\rho - r_{YT}) \in (0, E)$. In this case, we have that the equilibrium marginal revenue from effort toward Patreon, given by $\rho - e_P^*$, equals the marginal revenue from YouTube effort, r_{YT}, so that the CC produces content on both platforms. Instead, if the Patreon marginal revenues are too low ($\rho < r_{YT}$), then the CC only produces on YouTube; similarly, if the Patreon marginal revenues are too high ($\rho - E > r_{YT}$), then the CC only produces on Patreon.

Our model provides some natural predictions for how different types of CCs might respond to the YouTube ad boycott. To highlight these predictions in a manner that aligns with our empirical approach, we compare the equilibrium in the pre- and post-boycott periods. Regarding effort responses to the YouTube ad boycott, we see that efforts migrate from YouTube to Patreon:

Corollary 1. *Content creator efforts migrate from YouTube toward Patreon following a YouTube ad boycott:* $\frac{de_P^*}{dr_{YT}} \in \{0, -1\}$ *and* $\frac{de_{YT}^*}{dr_{YT}} \in \{0, 1\}$.

To determine how toxicity changes across platforms, we must measure toxicity at the CC-

[8] Alternatively, each content could vary in its toxicity draw or costs from altering the toxicity draw could depend on substitution in toxicity across platforms. Incorporating these features into the model does not qualitatively change our main results.

platform level. In other words, toxicity on YouTube and Patreon are given by

$$\begin{aligned} T_{YT} &= e_{YT} \cdot t_{YT}, \\ T_P &= e_P \cdot t_P. \end{aligned}$$

Thus, in terms of toxicity at the CC-platform level, we see that

Proposition 2. *In equilibrium,* $T_{YT}(\mathbb{B}=0) \geq T_{YT}(\mathbb{B}=1)$ *and* $T_P(0) \leq T_P(1)$.

Not surprisingly, when toxic behavior becomes more costly on YouTube, toxicity migrates from YouTube to Patreon.

4 The Empirical Setting

Our empirical analysis follows a methodology similar to the approach outlined in the model of Section 3. Specifically, we direct our focus towards two distinct groups of Patreon content providers: those who multi-home on Patreon and YouTube and those who operate on Patreon without multi-homing on YouTube. To accomplish this, we use a dataset sourced from Graphtreon, a data provider that routinely collects information from Patreon. This dataset provides us with a comprehensive overview of the number of content creators on the platform over time. It includes data on whether they engage in multi-homing on YouTube and the total count of the paying subscribers for each creator active on the platform over time. However, this dataset does not include specific details about the content produced by each creator. To address this limitation, we enrich this dataset by conducting our own data scraping directly from Patreon to gather content-related information.

In the following part of this Section, we start presenting contextual information relevant to our analysis, including details about YouTube and Patreon, the YouTube Partner Program, and the chronological sequence of events related to the so-called “Adpocalypse” shock. Following this contextual overview, we provide more comprehensive descriptive statistics derived from the dataset employed in our analysis.

4.1 YouTube and Patreon

YouTube is one of the world's leading video-sharing platforms, with a vast array of user-generated content spanning a wide range of genres, from entertainment and education to gaming and vlogging. Founded in 2005, YouTube has since become one of the major platforms in the digital world. The platform enables users to upload, view, and share videos. YouTube currently boasts over 2 billion logged-in monthly users, with millions of videos uploaded daily.[9]

YouTube offers a range of monetization options for content creators, allowing them to earn revenue from their videos. The most common method of monetization on YouTube is through ad revenue. Content creators can enable ads on their videos, and they earn a share of the revenue generated from ads shown to viewers. Moreover, content creators can offer channel memberships to their audience. Viewers who become channel members pay a monthly fee and, in return, receive perks like custom badges, emojis, and exclusive content.

To unlock the various monetization features offered by the platform, content creators must officially partner with YouTube and enroll in the YouTube Partner Program (YTPP). In the following subsection, we will explore the specific criteria that creators need to satisfy to qualify for YTPP, and how these requirements have evolved over time. Apart from the monetization features designed for content creators on the YouTube platform, creators can also sponsor content or use crowdfunding platforms like Patreon or external donation links to receive direct financial support from their fans.

Patreon is a platform that enables content creators to receive direct financial support from their dedicated fans and followers. It empowers creators to establish a direct connection with their audience by offering exclusive, often more private, content in exchange for monthly subscriptions. That is, unlike a Youtube "subscriber" who does not necessarily pay anything to Youtube or the CC, a "patron" on Patreon is paying a monthly subscription for which the majority of the at least \$1 membership goes directly to the CC.[10] Patreon is one of the most renowned crowdfunding platforms, especially within the realm of online content creation, including YouTubers, podcasters, and other digital creators. In the context of YouTube, content creators can reference their Patreon webpage directly in their videos or within the video description, as illustrated in Figure 1. This

[9]For more information about updated YouTube statistics: https://www.statista.com/topics/2019/YouTube.

[10]Patreon only takes between 5-12% of the monthly subscription revenues generated from the CC's patrons.

approach allows content creators to diversify their income streams and reduce their reliance on YouTube's advertising revenue, mitigating the impact of changes in YouTube's policies, such as ad monetization eligibility.

4.2 The YouTube "Adpocalypse"

In early 2017, major advertisers and prominent brands, including well-known companies, became concerned about their advertisements appearing alongside content they deemed inappropriate or extremist.[11] In response, many of these brands temporarily suspended their advertising on YouTube. This advertising boycott, and more generally, this period in YouTube's history, is known as the YouTube "Adpocalypse".

Before the "Adpocalypse" of 2017, YouTube's approach to content moderation was considerably more relaxed compared to what it would become in the aftermath of that event. YouTube primarily relied on algorithms to monitor and moderate content. These algorithms were designed to flag and remove content that violated YouTube's Community Guidelines, including policies against hate speech, violence, explicit content, and copyright infringement. However, the accuracy of these algorithms was not flawless, and some inappropriate or borderline content often evaded detection.[12] In addition to algorithmic moderation, YouTube heavily depended on its user reporting system. Users were encouraged to report content they found offensive or in violation of guidelines. This crowd-sourced approach to moderation meant that content could be flagged based on user reports, and YouTube would then review the reported content. While this system helped identify problematic content, it also resulted in inconsistencies as flagged content was reviewed by human moderators and subject to their interpretation.

Before the "Adpocalypse", YouTube's monetization model was relatively lenient. Content creators could monetize their videos with ads, and there were fewer restrictions on what content could be monetized. In response to the events in 2017, YouTube implemented a series of demonetization policies aimed at addressing concerns raised by advertisers and improving brand safety on the

[11] For specific cases, refer to these Times and Wall Street Journal articles. For a general overview, see: https://www.theverge.com/2019/4/5/18287318/YouTube-logan-paul-pewdiepie-demonetization-adpocalypse-premium-influencers-creators.

[12] See: https://www.polygon.com/2017/11/7/16620400/YouTube-algorithm-kids-programming.

platform. Quoting the former YouTube CEO from an online post in December 2017:[13] "We want advertisers to have peace of mind that their ads are running alongside content that reflects their brand's values. Equally, we want to give creators confidence that their revenue won't be hurt by the actions of bad actors."

YouTube introduced more comprehensive and explicit content guidelines to provide content creators with a better understanding of what was considered acceptable on the platform. Moreover, YouTube modified the YouTube Partner Program (YTPP) on February 20th, 2018, which previously had almost no restrictions on content creators. Now, "(YouTube) channels with fewer than 1,000 subscribers or 4,000 watch hours will no longer be able to earn money on YouTube".[14] Moreover, channels had to comply with YouTube's Community Guidelines and Ad Policies. This meant that the content on the channel had to adhere to YouTube's rules regarding hate speech, violence, explicit content, and other policy areas. Channels with repeated violations could face demonetization or other penalties.

These changes in the YTPP not only set explicit eligibility standards but also revealed a nuanced shift in YouTube's approach to content moderation. They implied that content creators with smaller audiences might be too numerous to moderate effectively and, therefore, would not be monetized. On the other hand, creators with larger followings would have the opportunity for monetization but would need to navigate the scrutiny of the platform's content moderation policies.

4.3 Graphtreon and Patreon Data

We trace CCs' content on Patreon with the help of data from Graphtreon, a website that tracks the overall number and characteristics over time of CCs on Patreon with at least one patron (subscriber). In our analysis, the dataset is composed of Patreon snapshots from August 2017 to August 2018 for the "Video" category. We combine all these snapshots to create an unbalanced panel dataset. In each snapshot, we observe CCs if they are present on the Patreon website on the snapshot date and have at least one patron. Consequently, for each CC in the dataset, we identify entry and exit events, as well as several CC characteristics. Some of these characteristics remain constant over time, such as the CCs' category and the Patreon web link, while others are updated with

[13] For the entire post, see: https://blog.YouTube/news-and-events/expanding-our-work-against-abuse-of-our/.
[14] See: https://blog.YouTube/news-and-events/additional-changes-to-YouTube-partner/.

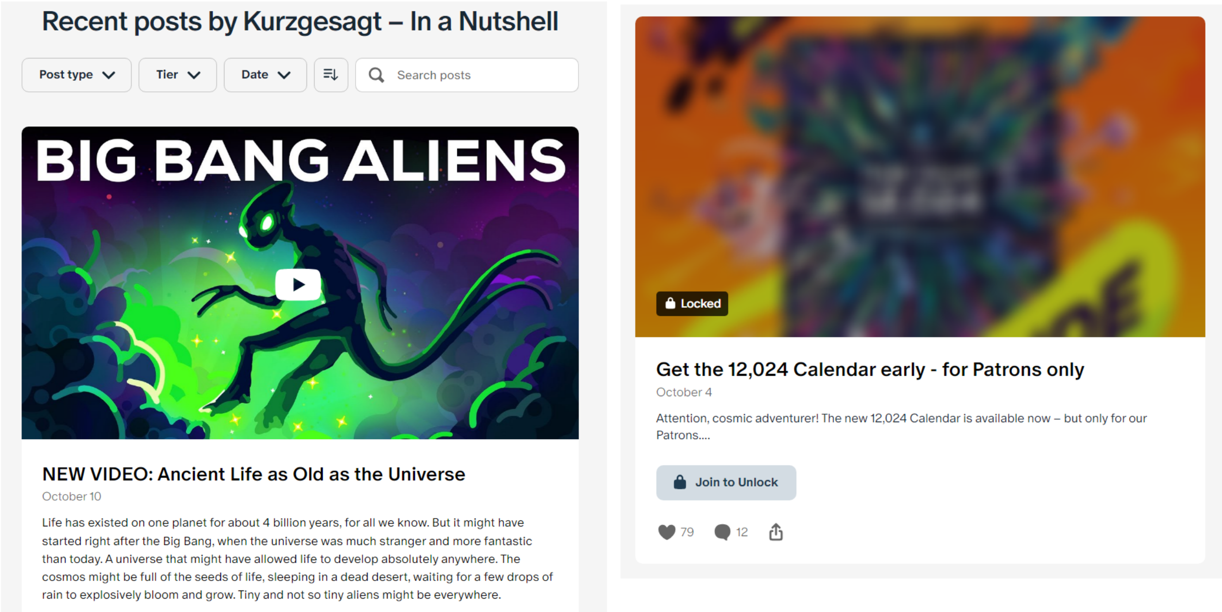

(a) Public Post

(b) Patrons-only Post

Figure 2: Observable Posts on Patreon (Kurzgesagt)

each snapshot. These dynamic characteristics include the number of patrons subscribed to each CC, monthly earnings, and whether the CC has a direct web link to other platforms like YouTube.

We enhance the Graphtreon dataset by extracting information directly from Patreon. Specifically, we leverage the direct links to Patreon provided by each CC. Not all the Patreon CCs listed in the Graphtreon dataset were still active on the platform when we began our scraping in 2022. Therefore, our analysis is limited to those CCs that were active throughout the year 2017 and whose webpages had not been deleted as of Spring 2022.

Since Patreon is a subscription-only platform, we do not have direct access to all the content posted by each CC, such as their videos (see Figure 2). However, we are able to observe important details, including the posting date, title, the number of comments and likes for each post, and whether the content is freely accessible or requires a subscription. Consequently, for each month in our tracking period, we collect data on the number of free or subscribers-only content items produced by CCs, as well as the number of likes and comments received by each piece of content.

Furthermore, through analysis of the content titles, we employ pre-trained algorithms provided by Jigsaw and Google's Counter Abuse Technology team (Mondal et al., 2017; ElSherief et al., 2018; Han and Tsvetkov, 2020) to identify and measure the level of toxicity in the content produced

Table 1: Patreon CCs Characterstics

	N	Mean	Median	SD	Min	Max
Patrons	11371	84	10	369	1	12,206
Earnings	9465	321	54	1,176	.0092	38,878
YT Presence	11371	.85	1	.36	0	1
Contents/Month	11371	4.1	1.5	7.5	0	170
Comments/Month	11371	18	.77	80	0	2,076
Likes/Month	11371	41	1.8	190	0	6,155

Notes: The dataset is sourced from Graphtreon and our independent web scraping of Patreon webpages. We narrow our focus to CCs who are active on Patreon (category "Video") for at least one month within the period spanning from August 2017 to August 2018.

by each CC over time. Before delving into the details of our identification design, we provide an overview of the dataset we have gathered regarding Patreon Content Creators (CCs). In Tables 1 and 2, we present information regarding the characteristics of CCs on Patreon. Our dataset is limited to the period from August 2017 to August 2018, focusing on CCs for whom we collected information when we scraped their Patreon webpages in Spring 2022.

Table 1 displays key statistics for the CCs within our sample. We have data for over 11,000 CCs, with an average patron count exceeding 80. Not all CCs disclose their earnings and our sample reveals that roughly 80% of the CCs choose to reveal. For those who do disclose their earnings, the average monthly income is more than 320 US$, which suggests that, on average, each patron contributes around 4 US$ per month in revenue. Many CCs on Patreon also maintain a presence on YouTube. Specifically, 85% of CCs have, at least once during the period of interest, provided a link to YouTube. Regarding the volume of content, comments, and likes they receive each month, CCs produce an average of four pieces of content per month. These pieces of content receive an average of 18 comments and 41 likes. This indicates that CCs on Patreon enjoy an actively engaged audience of patrons who interact with their content.

In Table 2, we present the same dataset, but this time we no longer average the data at the individual CC level. Instead, we consider each individual CC-month observation, which is the dimension we use in our identification design. As previously mentioned, our dataset comprises more than 11,000 unique CCs and a total of 131,000 CC-month observations. To reiterate, on average, content creators produce four pieces of content per month. These content pieces can be categorized into three types: free content, base content, which requires a minimum subscription (one dollar) to access, and premium content, which requires a subscription of more than one dollar for patrons to view.

Table 2: Patreon CCs-Month observations: Types of Contents

	N	Mean	Median	SD	Min	Max
Contents	131006	4.01	1	8.24	0	271
Free Contents	131006	1.34	0	3.90	0	162
Base Contents	131006	1.55	0	4.32	0	201
Premium Contents	131006	1.11	0	4.18	0	252
Images	131006	3.18	0	7.23	0	230
Videos	131006	1.89	0	5.13	0	216

Notes: The dataset is sourced from Graphtreon and our independent web scraping of Patreon webpages. We narrow our focus to CCs who are active on Patreon (category "Video") for at least one month within the period spanning from August 2017 to August 2018.

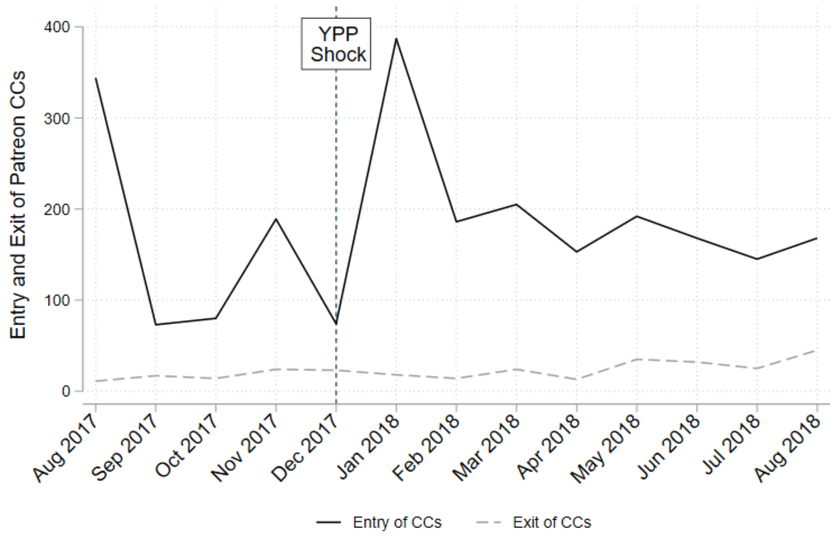

Figure 3: Entry and Exit on Patreon before and after the youTube "Adpocalypse"

Notes: The dataset is sourced from Graphtreon and our independent web scraping of Patreon webpages. We narrow our focus to CCs who are active on Patreon (category "Video") for at least one month within the period spanning from August 2017 to August 2018.

In the upcoming sections, we will merge base and premium content and solely focus on analyzing the distinctions between free and paid content. However, it is worth noting that these three types of content (free, base, and premium) constitute a similar proportion of the total number of content pieces. Content posted on Patreon typically includes images and videos, as evidenced by the fact that, on average, CCs post more than 3 images and nearly two videos every month.

Before we delve into our identification strategy, we provide an overview of how the number of CCs evolved around the time when YouTube implemented a stricter content moderation policy. In Figure 3, we present a graphical representation of the number of CCs in the "Video" category that either joined or left Patreon between August 2017 and August 2018. As can be seen in the figure, the platform experienced growth during this period, with more new CCs joining than leaving. Notably, we can see specific spikes in the months of August 2017, November 2017, and January 2018, which coincide with the timing of the YouTube ad boycott. Additionally, throughout 2018, the influx of new CCs has consistently exceeded the levels seen in the last few months of 2017. However, despite these fluctuations in new entrants, the total number of CCs on Patreon remained relatively stable. The total number of CCs present from August 2017 to August 2018 exceeded 11,000, while the number of new CCs joining the platforms in late 2017 or early 2018 is less than 1200. This suggests that while there were fluctuations in the number of new arrivals, the overall CC population on Patreon experienced only modest changes.

5 Identification Strategy

With our empirical design, we aim to study how CCs respond to the changes in the YouTube monetization policy due to the "Adpocalypse" ad boycott. To do so, we investigate the evolution in the quantity and nature of content produced, as well as changes in the number of patrons that pay monlthy subscriptions to two sets of Patreon CCs from August 2017 to August 2018. We focus on CCs that were multi-homing Patreon and YouTube, as they constitute the "treated" group directly affected by the introduction of the YouTube moderation policy. We identify these CCs based on whether they display a YouTube link on their Patreon webpage at least once during the period of analysis. Conversely, Patreon CCs without YouTube links on their webpages serve as our control group, as they remain unaffected by the YouTube policy changes.

The change in YouTube's moderation policy, specifically the tightened eligibility criteria for the YouTube Partner Program, was formally implemented in February of 2018.[15] However, the shift in YouTube's moderation approach began several months earlier during the "Adpocalypse" of 2017 and the resulting ad boycotts. We cannot trace back the exact moment when YouTube altered its approach and content creators became aware of this change. However, in December 2017, the former YouTube CEO, Susan Wojcicki, officially announced expanded efforts to combat abuse on the platform.[16] Thus, we use December 2017 as the month when content creators became aware of the forthcoming changes. Some content creators may have anticipated the changes prior to this date, but we consider that December 2017 represents the time when most content creators started taking measures to respond to the upcoming moderation policy.

Thus, we use a Difference-in-Difference (DiD) identification approach and compare Patreon CCs in the "treated" group (multi-homing on YouTube) and "control" group (not multi-homing on YouTube) from August 2017 to August 2018 before and after December 2017. TThe estimating regression to capture the causal impact of the moderation policy on Patreon CCs' activity is:

$$y_{it} = \alpha_i + \rho_t + \beta_1 YT_i + \beta_2 After_t + \beta_3 YT_i \times After_t + \varepsilon_{it}, \tag{5.1}$$

where y_{it} denotes the variable of interest across different dimensions of CCs' activity affected by the YouTube policy (e.g., number of subscribers or contents), α_i and ρ_t are the CC and time fixed effects, YT_i is a dummy variable equal to 1 if the CC i has multi-homed on Patreon and YouTube for at least one month during the period of analysis and equal to 0 if the CC i has never multi-homed on YouTube, and $After_t$ is equal to 1 for all snapshots after December 2017 and equal to 0 otherwise. The coefficient β_3 captures the impact of the change in the YouTube monetization policy under the assumption that Patreon CCs that do not multi-home on YouTube provide a good counterfactual group for the evolution of the variables of interest of the CCs that are active on Patreon and YouTube.

In the next section, we present the results of the Difference-in-Differences (DiD) approach for all the variables of interest. Additionally, for each variable, we check for the absence of potential pre-trends between the different groups of CCs using an event-study approach. We use the fol-

[15] See: https://blog.YouTube/news-and-events/additional-changes-to-YouTube-partner/.

[16] See: https://blog.YouTube/news-and-events/expanding-our-work-against-abuse-of-our/.

lowing lead-lag model in which y_{it} is regressed on the product of the dummy variable YT_i and a comprehensive set of dummy variables for each snapshot. The model controls for CC fixed effects and time fixed effects:

$$y_{it} = \alpha_i + \rho_t + \sum_{\tau=Sep17}^{Jul18} \beta_\tau YT_i \times 1(t = \tau) + \varepsilon_{it}. \tag{5.2}$$

In all graphs, the coefficients corresponding to months before December 2017 are approximately zero and do not display a discernible trend. Consequently, the evolution of the variables of interest for the treated and control CCs appeared to be similar before the moderation policy change. This observation supports the parallel trend assumption and confirms the validity of our empirical design.

6 Results

Using the DiD identification design outlined in Section 5, we can study different dimensions related to the number of patron subscriptions and the quantity and the nature of the contents of Patreon CCs. Before focusing on each variable, here we provide an overview of the main results.

CCs multi-homing on YouTube produce more content that is exclusively available to their subscribers on Patreon. Instead, the quantity of the public content remains unchanged. Together, these results suggest that CCs affected by the moderation policy change are prioritizing the now more attractive monetization strategy offered by Patreon. Also the number of patrons subscribing to Patreon increases, possibly due to active recommendations from CCs encouraging viewers to transition from YouTube to Patreon. This rise in the number of subscribers also has a significant impact on earnings.

Moreover, we can show that treated CCs are also investing more effort on Patreon, not only in terms of quantity (more content) but also in the quality of content. Notably, we observe an increase in the number of likes on Patreon for those CCs affected by the content moderation policy compared to CCs who do not multi-home on YouTube. The increase in the number of likes is primarily noticeable for subscription-only content, and this positive effect remains robust even after controlling for changes in the number of content and patrons. It is essential to emphasize this point, as one might anticipate that more content and more patrons would mechanically increase the number of likes. A similar effect is evident for comments, with more comments on subscribers-only content posted on Patreon. However, once we account for the number of content pieces and the number of patrons, the effect becomes null, suggesting that changes in the number of comments are primarily linked to increases in subscribers and content volume. Finally, the overall toxicity of Patreon increases, mainly due to CCs posting more content on this platform. However, there is no observed increase in the average toxicity of the content itself.

In the remainder of this section, our focus will shift to each individual result, where we will describe the DiD table and the event study figures in detail.

6.1 Number of Content (All, Free, and Paid)

We begin by examining the impact of the content moderation shock on the content generated by CCs on Patreon. In this context, the outcome variable, denoted as y_{it} in Equations 5.1 and 5.2, is defined either as a binary variable indicating whether a CC produced at least one piece of content in a given month (as shown Columns 1-3 in Table 3) or as the logarithm of the number of content items generated by each CC within that month (Columns 1-3 in Table 3). These two settings correspond to the extensive and intensive margins of the effect. While the first setting looks at whether content moderation leads to CCs posting more frequently in more months, the second establishes whether CCs create a higher number of contents per month. In both settings, our analysis considers all CCs active on Patreon in the "Video" category, and we focus on those who joined the platform before the implementation of the YTPP and remained active afterwards.

The findings from both Tables reveal a significant trend: CCs who engage in multi-homing on YouTube post a greater volume of content compared to their non-multi-homing counterparts: They generate content in more months as well as more content per month. Moreover, this effect is primarily driven by the production of paid content, which is accessible exclusively through subscription, as opposed to free content that is available to all users.

To visualize these trends, Figure 4 showcases three event study graphs utilizing the logarithm of the number of total content items, free content, and paid content. All three graphs demonstrate a lack of noticeable pre-trend behavior before the YTPP implementation. However, an effect

Table 3: Difference-in-Differences: Presence and Number of Content (All, Free, and Paid)

	Presence of Content			Log No. of Contents		
	All	Free	Paid	All	Free	Paid
$post_t^{Dec2017} \times YT_i$	0.00	0.01*	0.01	0.08***	-0.02	0.07***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)	(0.03)
User FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Mean	.533	.32	.403	1.39	.929	1.24
R^2	0.61	0.53	0.64	0.72	0.66	0.70
N	108315	108315	108315	57006	33394	42886

Notes: Standard errors clustered by CC are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

emerges after January 2018, particularly concerning the production of content and paid content.

In order to tackle potential concerns about the specific timing of when the impact of YouTube's moderation changes should come through on Patreon, we rerun the above analyses using different post cutoffs. Tables 8 and 9 in Appendix B confirm the previous findings: Multi-homing CCs produce more content, and this increase is primarily driven by an increased number of paid contents.

6.2 Number of Patrons and Earnings

The rise in the quantity of content directly reflects changes in the behavior of CCs. However, the impact of the YouTube content moderation extends beyond just CCs; it also profoundly influences the behavior of all platform users. This influence is evident in Table 4, where we observe a significant and dramatic increase in the number of patrons subscribing to CCs who engage in multi-homing on YouTube, compared to those CCs who do not multi-home. In line with these findings, CCs' earnings on Patreon experience a notable surge following the implementation of the content moderation, driven by the substantial increase in the number of patron subscriptions.

To visualize these effects, Figure 5 presents the event study for the logarithm of the number of patrons and earnings. In both cases, a clear and positive effect becomes evident in the post-YTPP period. However, it is worth noting that we may also observe some degree of anticipation among CCs, as they appear to have started attracting users to Patreon even before December 2017, possibly in anticipation of the impending tightening of rules.

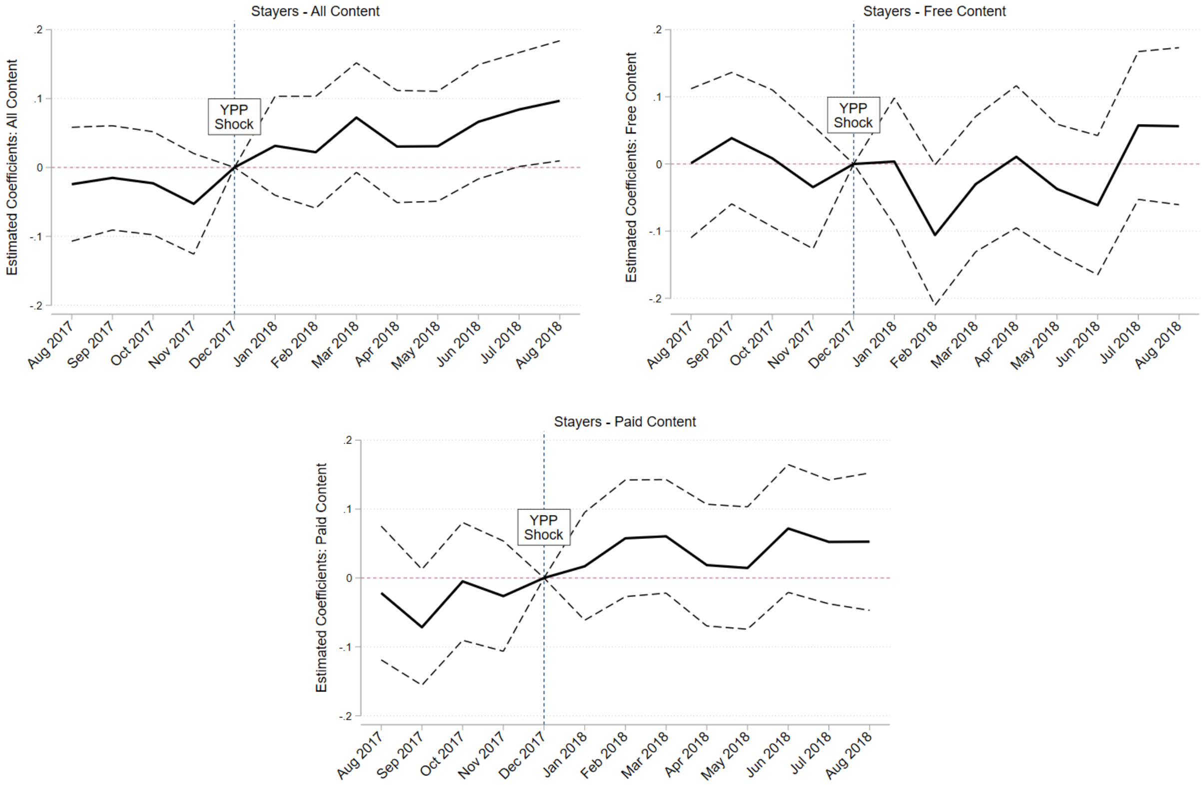

Figure 4: Event study: log of Number of Content (All, Free, and Paid) - Stayers

Notes: In line with Equation 5.2, the log of $Content_{it}$, $Content_{it}^{Free}$, $Content_{it}^{Paid}$ are regressed on CC fixed effects; time fixed effects, and on the products between YT_i and a full set of dummy variables for each month. The graphs plot the estimated coefficients on these products. The value of the coefficient corresponding to December 2017 is normalized to zero. The sample includes months between August 2017 and August 2018. Standard errors (5%) are clustered by user.

Table 4: Difference-in-Differences: log of Number of Patrons and Earnings

	Log Patrons	Log Earnings
$post_t^{Dec2017} \times YT_i$	0.12***	0.13***
	(0.01)	(0.02)
User FE	✓	✓
Time FE	✓	✓
Mean	2.72	4.03
R^2	.968	.95
N	108,315	86,057

Notes: Standard errors clustered by CC are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

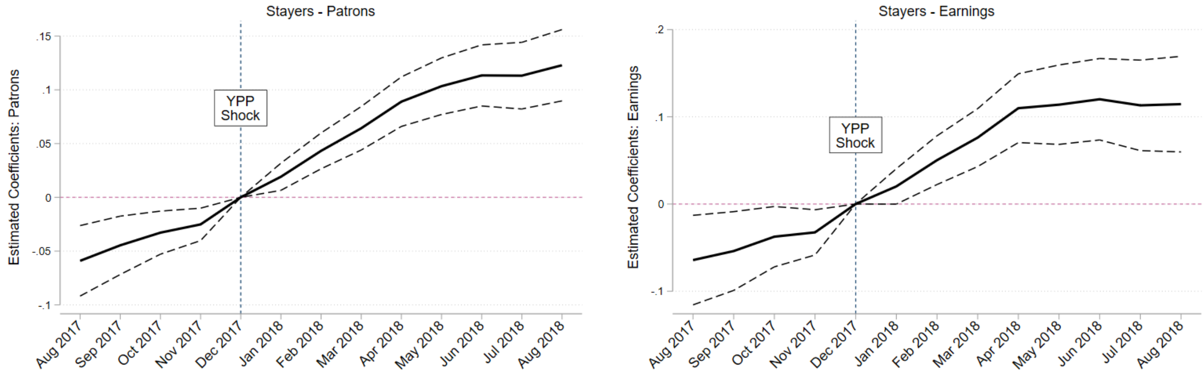

Figure 5: Event study: log of Number of Patrons and Earnings - Stayers

Notes: In line with Equation 5.2, the log of $Patrons_{it}$, $Earnings_{it}$ are regressed on CC fixed effects; time fixed effects, and on the products between YT_i and a full set of dummy variables for each month. The graphs plot the estimated coefficients on these products. The value of the coefficient corresponding to December 2017 is normalized to zero. The sample includes months between August 2017 and August 2018. Standard errors (5%) are clustered by user.

6.3 Number of Likes

Following our examination of the cross-platform impact on the quantity of content and the total number of subscribed users, we now delve into its effect on content quality, specifically measured by the number of likes received by CCs and month. The first three columns in Table 5 replicate the analysis previously presented in Table 3, focusing on the impact of the rule change on the logarithm of the number of likes for all content, free content, and paid content, employing the same specifications as outlined in Equation 5.1. Notably, the results demonstrate that the increase in the number of likes is indeed evident, with a noteworthy effect primarily driven by likes for paid content.

While these results are intriguing, it is reasonable to question whether these effects are merely a mechanical consequence of the increased volume of content and patrons. One might expect that more content would naturally lead to more likes, and an expanded patron base might also correlate with increased likes. To address this concern, the second half of Table 5 employs a classical DiD design while controlling for both the total number of patrons and the quantity of content produced in a given month. This approach departs somewhat from the classical DiD design, as we typically would not control for a variable affected by the very shock under investigation. However, we believe it is instructive to demonstrate that the rise in likes appears only partially linked to these other variables. Remarkably, Columns 4-6 of Table 5 reveals a significant increase in the number of likes for all content, driven by paid content, even after accounting for the mentioned control variables.

We acknowledge that the variation in likes may also depend on the characteristics of the patrons, who may vary in their inclination to leave a like. Nonetheless, these findings suggest that the YTPP's impact on likes extends beyond a simple mechanical effect and is indicative of genuine shifts in user engagement and content quality. As before, we visualize these effects in Figure 6 where we presents the event study for the logarithm of the number of likes for different types of content. In all cases, a clear and positive effect becomes evident in the post-shock period for the likes on all content and paid content.

Table 5: Difference-in-Differences: log of Number of Likes (All, Free, and Paid)

	Unconditioned			Conditioned		
	All Contents	Free Contents	Paid Contents	All Contents	Free Contents	Paid Contents
$post_t^{Dec2017} \times YT_i$	0.14***	0.05	0.15***	0.08***	0.01	0.10***
	(0.03)	(0.04)	(0.04)	(0.03)	(0.04)	(0.03)
User FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
No. Patrons				✓	✓	✓
No. Contents				✓	✓	✓
Mean	2.73	2.28	2.7	2.73	2.28	2.7
R^2	.854	.827	.843	.888	.846	.872
N	47,642	26,875	35,686	47,642	26,875	35,686

Notes: The first three columns replicate the previous analysis in a simple DiD framework. Columns 4-6 additionally control for the number of patrons and contents. Standard errors clustered by CC are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.4 Number of Comments

The quality of content could also be measured by considering the number of comments generated by a content or CC. Table 6 provides insights into this aspect, indicating an increase in the number of comments for both all content and paid content when employing the same specification as outlined in Equation 5.1. However, when we reevaluate this analysis while accounting for the number of content items and subscribers in the second part of Table 6, we fail to detect a significant effect. This observation suggests that the number of comments per content and per subscriber remains relatively stable. It implies that comments are more likely a metric of audience engagement, closely

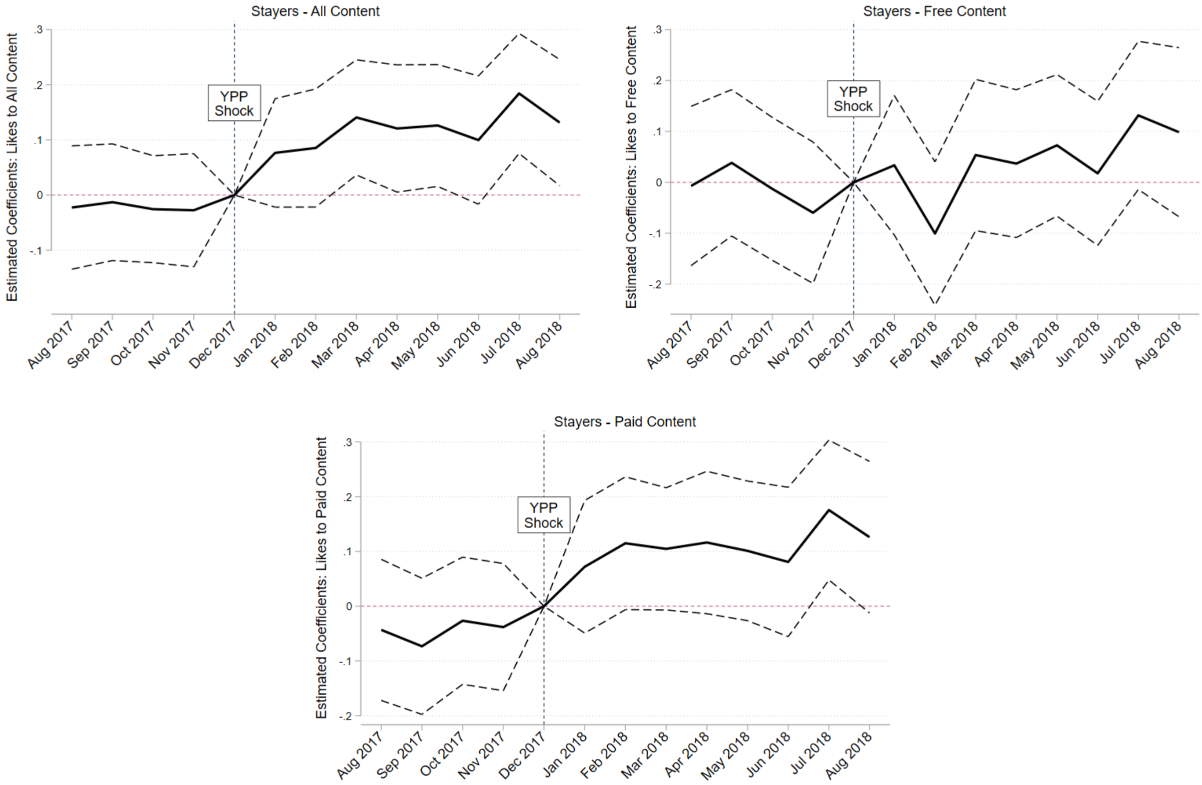

Figure 6: Event study: log of Number of Likes on Content (All, Free, and Paid) - Stayers

Notes: In line with Equation 5.2, the log of $Likes_{it}$, $Likes_{it}^{Free}$, $Likes_{it}^{Paid}$ are regressed on CC fixed effects; time fixed effects, and on the products between YT_i and a full set of dummy variables for each month. The graphs plot the estimated coefficients on these products. The value of the coefficient corresponding to December 2017 is normalized to zero. The sample includes months between August 2017 and August 2018. Standard errors (5%) are clustered by user.

aligned with the overall size of the audience, rather than a reliable indicator of content quality, as we previously interpreted the increase in the number of likes to be.

6.5 Content Toxicity

We conclude our series of findings by examining the measure of toxicity within the content generated by CCs. Utilizing our content titles and employing the pre-trained algorithms provided by Jigsaw and Google's Counter Abuse Technology team (Perspective), we measure the toxicity of each piece of content as a continous score ranging from 0 to 1. In the first part of Table 7, we aggregate the toxicity levels of all content produced by each CC in a given month. Conversely, the second part of Table 7 employs the average toxicity level of the content. These analyses serve dis-

Table 6: Difference-in-Differences: log of Number of Comments (All, Free, and Paid)

	Unconditioned			Conditioned		
	All Contents	Free Contents	Paid Contents	All Contents	Free Contents	Paid Contents
$post_t^{Dec2017} \times YT_i$	0.11***	0.06	0.10**	0.06	0.02	0.04
	(0.04)	(0.05)	(0.05)	(0.04)	(0.05)	(0.04)
User FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
No. Patrons				✓	✓	✓
No. Contents				✓	✓	✓
Mean	2.43	1.93	2.46	2.43	1.93	2.46
R^2	.778	.707	.758	.809	.725	.786
N	39,020	19,400	30,064	39,020	19,400	30,064

Notes: The first three columns replicate the previous analysis in a simple DiD framework. Columns 4-6 additionally control for the number of patrons and contents. Standard errors clustered by CC are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Difference-in-Differences: Content Toxicity (All, Free, and Paid)

	Log Total Toxicity			Log Average Toxicity		
	All Contents	Free Contents	Paid Contents	All Contents	Free Contents	Paid Contents
$post_t^{Dec2017} \times YT_i$	0.09***	-0.01	0.08**	0.01	0.01	0.01
	(0.03)	(0.04)	(0.04)	(0.02)	(0.03)	(0.02)
User FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Mean	-4.61	-5.12	-4.81	-6.01	-6.05	-6.05
R^2	.636	.578	.622	.421	.396	.415
N	56,832	33,315	42,794	56,832	33,315	42,794

Notes: Standard errors clustered by CC are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

tinct purposes: While the total toxicity enables us to assess whether Patreon becomes more toxic, in total, following the content moderation shock on YouTube; the average toxicity measures allow us to determine whether the individual contents being produced are more toxic.

Our analysis of these two toxicity measures reveals a noteworthy outcome: CCs who engage in multi-homing on YouTube do not appear to become more toxic themselves. However, as they increase their content production, especially in the realm of paid content, the overall toxicity level on the platform experiences an upsurge.

7 Managerial Implications

Our findings speak directly to several players in the content creator space: incumbent and entrant platforms considering content moderation and monetization policies, content creators evaluating such policies, and legislators interested in reducing toxicity. For platforms implementing content moderation policies, our work shows that such policies may leave the platform open to competition from platforms offering alternative monetization strategies. This suggests that the losses associated with introducing content moderation could be limited if the platform also provides alternative monetization opportunities to its content creators. More generally, we reveal an important relationship between content moderation and platform monetization that potential platform entrants should be aware of: an advertising model where content is open to the public may result in future content moderation issues that could be avoided under a membership model where content is funded by consumers and is therefore able to remain private. Naturally, a platforms' choices in monetization and moderation policies impact content creators. In particular, content creators that are able to adapt their content to multiple platform business models are better able to monetize their content, emphasizing the importance of developing a variety of content. Lastly, our work suggests that a legislator interested in reducing toxicity may only be able to reduce publicly observable toxicity when targeting content moderation on a dominant platform as entrant platforms may offer an alternative option for such content to migrate.

This connection between content moderation and platform monetization appears in other industries as well. For example, Elon Musk has speculated that changing the monetization strategy of X, formerly Twitter, so that users generating content must pay an annual fee will also serve as a content moderation policy by blocking bot users.[17] Alternatively, Instagram competing with Onlyfans relates to our work on platform competition when platforms have different monetization strategies that allow for public versus private content. More specifically, in response to Onlyfans' successful use of the membership model with private content (unlike YouTube's null response to Patreon's success), Instagram launched a similar subscription model for private content to better compete with Onlyfans. These examples, combined with our own findings, highlight how important it is for platforms to account for the interplay between content moderation and monetization.

[17] Specifically, in an X post on October 17, 2023: "read for free, but $1/year to write. It's the only way to fight bots without blocking real users."

8 Conclusion

Our study delves into the complex dynamics of the content creator industry, with a particular focus on the aftermath of the YouTube "Adpocalypse" and the subsequent restrictions in monetization. As we examine the shifts in content creators' strategies and consumer behavior across platforms, several key findings emerge that have implications for various stakeholders in this ecosystem.

First, our theoretical model and empirical evidence confirm that content creators respond strategically to changes in moderation policies that could potentially lead to demonetization. They diversify their monetization strategies, with an increased focus on platforms like Patreon, which offers alternative revenue streams such as monthly subscriptions. This diversification not only allows content creators to adapt to changing market conditions but also results in an improvement in the quality and quantity of content available on Patreon. This finding underscores the importance of flexibility in monetization for content creators, which can insulate their profits from platform-level shocks and ultimately enhance their ability to thrive in a rapidly evolving industry. For platforms, our research highlights the significance of offering diverse monetization options for content creators. Differentiation in monetization strategies can attract and retain a diverse range of content creators, fostering a more robust ecosystem.

Regarding the issue of toxic content, our study reveals a migration of toxicity to Patreon following the YouTube restrictions. This migration pattern suggests that efforts to target large online companies alone may not effectively reduce internet-wide toxicity, as creators and viewers can simply shift to other platforms with fewer content moderation measures. This phenomenon underscores the challenges faced by policymakers in addressing online toxicity and highlights the need for a more comprehensive, industry-wide approach to tackling this issue.

Thus, our research provides valuable insights into the ever-evolving content creator industry, shedding light on how content creators, platforms, and policymakers respond to significant disruptions and challenges. As this industry continues to evolve, it will be essential for all stakeholders to adapt, innovate, and collaborate to create a safer and dynamic online content ecosystem.

Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

References

Ali, Shiza, Mohammad Hammas Saeed, Esraa Aldreabi, Jeremy Blackburn, Emiliano De Cristofaro, Savvas Zannettou, and Gianluca Stringhini (2021) 'Understanding the effect of deplatforming on social networks.' In '13th ACM Web Science Conference 2021' pp. 187–195

Andres, Raphaela, and Olga Slivko (2021) 'Combating online hate speech: The impact of legislation on Twitter.' *ZEW-Centre for European Economic Research Discussion Paper*

Armstrong, Mark (2006) 'Competition in two-sided markets.' *The RAND journal of economics* 37(3), 668–691

Bhargava, Hemant K (2022) 'The creator economy: Managing ecosystem supply, revenue sharing, and platform design.' *Management Science* 68(7), 5233–5251

_ (2023) '"if it's enraging, it's engaging" - exploitive design in information platforms and the attention economy'

Boudreau, Kevin J, and Andrei Hagiu (2009) 'Platform rules: Multi-sided platforms as regulators.' *Platforms, markets and innovation* 1, 163–191

Casner, Ben, and Tat-How Teh (2022) 'Content-hosting platforms: discovery, membership, or both?' *SSRN Working Paper*

Chandrasekharan, Eshwar, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert (2017) 'You can't stay here: The efficacy of reddit's 2015 ban examined through hate speech.' *Proceedings of the ACM on Human-Computer Interaction* (1. CSCW), 1–22

Cusumano, Michael A, Annabelle Gawer, and David B Yoffie (2019) *The business of platforms: Strategy in the age of digital competition, innovation, and power,* vol. 320 (Harper Business New York)

El-Komboz, Lena Abou, Anna Kerkhof, and Johannes Loh (2023) 'Platform partnership programs and content supply: Evidence from the Youtube "Adpocalypse"'

El Sanyoura, Lana, and Ashton Anderson (2022) 'Quantifying the creator economy: A large-scale analysis of patreon.' In 'Proceedings of the International AAAI Conference on Web and Social Media,' vol. 16 pp. 829–840

ElSherief, Mai, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding (2018) 'Hate lingo: A target-based linguistic analysis of hate speech in social media.' In 'Proceedings of the international AAAI conference on web and social media,' vol. 12

Evans, David S, and Richard Schmalensee (2013) 'The antitrust analysis of multi-sided platform businesses.' Technical Report, National Bureau of Economic Research

Farronato, Chiara, and Andrey Fradkin (2022) 'The welfare effects of peer entry: the case of Airbnb and the accommodation industry.' *American Economic Review* 112(6), 1782–1817

Florida, Richard (2022) 'The rise of the creator economy'

Hagiu, Andrei, and Julian Wright (2015) 'Marketplace or reseller?' *Management Science* 61(1), 184–203

_ (2019) 'Controlling vs. enabling.' *Management Science* 65(2), 577–595

Han, Xiaochuang, and Yulia Tsvetkov (2020) 'Fortifying toxic speech detectors against veiled toxicity.' *arXiv preprint arXiv:2010.03154*

Kerkhof, Anna (2020) 'Advertising and content differentiation: Evidence from Youtube'

Kesler, Reinhold (2022) 'The impact of apple's app tracking transparency on app monetization.' *Available at SSRN 4090786*

Lefouili, Yassine, and Leonardo Madio (2022) 'The economics of platform liability.' *European Journal of Law and Economics* 53(3), 319–351

Li, Hui, and Feng Zhu (2021) 'Information transparency, multihoming, and platform competition: A natural experiment in the daily deals market.' *Management Science* 67(7), 4384–4407

Liu, Yi, Pinar Yildirim, and Z John Zhang (2022) 'Implications of revenue models and technology for content moderation strategies.' *Marketing Science* 41(4), 831–847

Luca, Michael (2015) 'User-generated content and social media.' In 'Handbook of Media Economics,' vol. 1 pp. 563–592

Mondal, Mainack, Leandro Araújo Silva, and Fabrício Benevenuto (2017) 'A measurement study of hate speech in social media.' In 'Proceedings of the 28th ACM conference on hypertext and social media' pp. 85–94

Paridar, Mahsa, Mina Ameri, and Elisabeth Honka (2023) 'More, faster, and better? effects of rewards on incentivizing the creation of user-generated content.' *Working Paper*

Parker, Geoffrey G, and Marshall W Van Alstyne (2005) 'Two-sided network effects: A theory of information product design.' *Management science* 51(10), 1494–1504

Rauchfleisch, Adrian, and Jonas Kaiser (2021) 'Deplatforming the far-right: An analysis of youtube and bitchute.' *Available at SSRN: 3867818*

Rochet, Jean-Charles, and Jean Tirole (2003) 'Platform competition in two-sided markets.' *Journal of the European Economic Association* 1(4), 990–1029

Sen, Ananya, Tom Grad, Pedro Ferreira, and Jörg Claussen (2023) '(how) does user-generated content impact content generated by professionals? evidence from local news.' *Management Science*

Srinivasan, Kumar Bhargav, Cristian Danescu-Niculescu-Mizil, Lillian Lee, and Chenhao Tan (2019) 'Content removal as a moderation strategy: Compliance and other outcomes in the changemyview community.' *Proceedings of the ACM on Human-Computer Interaction* (3. CSCW), 1–21

Teh, Tat-How (2022) 'Platform governance.' *American Economic Journal: Microeconomics* 14(3), 213–254

Wlomert, Nils, Dominik Papies, Michel Clement, and Martin Spann (2023) 'The interplay of user-generated content, content industry revenues, and platform regulation: Quasi-experimental evidence from youtube.' *Marketing Science*

A Appendix of Proofs

Proof of Proposition 1: Considering equilibrium effort levels reduces the CC's four variable constrained maximization problem to a simple two variable Lagrange maximization problem:

$$\max_{e_{YT},e_P} \quad r_{YT}(t_{YT},\mathbb{B})\cdot e_{YT} + \int_0^{e_P}(\rho - e)de,$$
$$\text{s.t. } E = e_{YT}+e_P.$$

In this case, an interior solution occurs with $e_P^* = \rho - r_{YT}$ and $e_{YT}^* = E - \rho - r_{YT}$ whenever $(\rho - r_{YT}) \in (0,E)$. Instead, if $\rho < r_{YT}$, then $e_P^* = 0$ and $e_{YT}^* = E$, and if $\rho - E > r_{YT}$), then $e_P^* = E$ and $e_{YT}^* = 0$. □

Proof of Proposition 2: If no ad boycott exists ($\mathbb{B} = 0$), then $r_{YT}(t_{YT},\mathbb{B})$ is not effected by t_{YT} so that content creators implement their default level of toxicity: $t_{YT}^* = t_P^* = \tau$. Instead, if an ad boycott exists so that $r_{YT}(t_{YT},\mathbb{B})$ is decreasing in t_{YT}, then t_{YT}^* is given by

$$0 = \frac{dr_{YT}(t_{YT},1)}{dt_{YT}}\cdot e_{YT} + 2(\tau - t_{YT}),$$

where $\frac{dr_{YT}(t_{YT},1)}{dt_{YT}} < 0$. Solving implies that

$$t_{YT}^* = \frac{dr_{YT}(t_{YT},1)}{dt_{YT}}\cdot\frac{e_{YT}}{2} + \tau < \tau.$$

Combining this with Corollary 1 implies that $T_{YT}(\mathbb{B}=0) \geq T_{YT}(\mathbb{B}=1)$ and $T_P(0) \leq T_P(1)$. □

B Further Tables

Table 8: Robustness Check: Varying the Post Cutoff, Extensive Margin

	All Contents			Free Contents			Paid Contents		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$post_t^{Dec2017} \times YT_i$	0.00			0.01*			0.01		
	(0.01)			(0.01)			(0.01)		
$post_t^{Jan2018} \times YT_i$		0.01			0.01			0.01**	
		(0.01)			(0.01)			(0.01)	
$post_t^{Feb2018} \times YT_i$			0.01			0.01*			0.02**
			(0.01)			(0.01)			(0.01)
User FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Post Cutoff	Dec 17	Jan 18	Feb 18	Dec 17	Jan 18	Feb 18	Dec 17	Jan 18	Feb 18
Mean	.533	.533	.533	.32	.32	.32	.403	.403	.403
R^2	.613	.613	.613	.529	.529	.529	.64	.64	.64
N	108,315	108,315	108,315	108,315	108,315	108,315	108,315	108,315	108,315

Notes: Standard errors clustered by CC are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Robustness Check: Varying the Post Cutoff, Intensive Margin

	All Contents			Free Contents			Paid Contents		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$post_t^{Dec2017} \times YT_i$	0.08***			-0.02			0.07***		
	(0.02)			(0.03)			(0.03)		
$post_t^{Jan2018} \times YT_i$		0.07***			-0.02			0.06**	
		(0.02)			(0.03)			(0.03)	
$post_t^{Feb2018} \times YT_i$			0.07***			0.01			0.05**
			(0.02)			(0.03)			(0.02)
User FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Post Cutoff	Dec 17	Jan 18	Feb 18	Dec 17	Jan 18	Feb 18	Dec 17	Jan 18	Feb 18
Mean	1.39	1.39	1.39	.929	.929	.929	1.24	1.24	1.24
R^2	.716	.716	.716	.661	.661	.661	.704	.704	.704
N	57,006	57,006	57,006	33,394	33,394	33,394	42,886	42,886	42,886

Notes: Standard errors clustered by CC are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Download ZEW Discussion Papers:

https://www.zew.de/en/publications/zew-discussion-papers

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

//
IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de