// NO.26-033 | 08/2026

DISCUSSION PAPER

// RAPHAELA ANDRES, AMBRE ELSAS-NICOLLE, MICHELANGELO ROSSI, AND MARK J. TREMBLAY

Staying for Discovery, Monetizing Elsewhere: Platform Governance and Complementor Bypassing

Staying for Discovery, Monetizing Elsewhere: Platform Governance and Complementor Bypassing*

Raphaela Andres[†] Ambre Elsas-Nicolle[‡] Michelangelo Rossi[§] Mark J. Tremblay[¶]

This version: July 2026

Abstract

Digital platforms govern not only how ecosystem value is created, but also how complementors capture it. In this paper, we study how content creators respond when a dominant platform restricts monetization while remaining essential for audience discovery. Our setting is YouTube's revision of its monetization regime, implemented in the aftermath of the 2017 advertiser boycott known as the "Adpocalypse". Under the new regime, YouTube tightened content moderation and monetization rules, reducing the level and predictability of creators' advertising revenues. Creators could partially bypass YouTube's monetization layer through Patreon, a membership platform enabling recurring direct payments from users. Using monthly panel data on approximately 8,400 video creators active on Patreon between August 2017 and August 2018, we compare creators with a pre-shock YouTube link to creators without an observed link in a matched difference-in-differences design. Following the policy change, exposed creators increased their production of member-only Patreon content without reducing freely accessible posts. They also accumulated paying patrons and recurring membership revenue at a significantly faster monthly rate, with a differential increase of approximately 9%. The increase in paid content was strongest among creators whose content faced greater pre-existing monetization risk. The findings suggest that platform governance can reallocate complementor effort and value capture across platforms: creators can shift monetization toward an auxiliary channel while continuing to rely on the dominant platform for audience discovery.

Keywords: Platform governance; Complementors; Platform bypassing; Creator economy; Monetization

JEL Codes: L10, L22, L82, L86.

---

*This paper supersedes a previous version entitled "YouTube Adpocalypse: The YouTubers' Journey from Ad-Based to Patron-Based Revenues".

We would like to thank Irene Bertschek, Donato Cutolo, Ralf Elsas-Nicolle, Chris Forman, Anna Kerkhof, Reinhold Kesler, Ulrich Laitenberger, Catherine Ma, Leonardo Madio, Poonacha Medappa, Mary O'Mahony, Maria Rasskazova, Imke Reimers, Shiva Shekhar, Murat Tunc, Joel Waldfogel, Mike Ward, Xi Wu, and Michael Zhang for helpful comments. All errors are our own.

[†]ZEW Mannheim. E-mail: Raphaela.Andres@zew.de

[‡]Mines Paris, PSL University, CEDP, i3 UMR9217 CNRS. E-mail: ambre.elsas-nicolle@minesparis.psl.eu

[§]*Corresponding author*. HEC Paris and CESifo. E-mail: rossim@hec.fr

[¶]Lee Business School, University of Nevada, Las Vegas, and CESifo. E-mail: mark.tremblay@unlv.edu

1 Introduction

Digital platforms play a central role in enabling value creation by reducing transaction costs, aggregating attention, facilitating discovery, and coordinating interactions among distinct groups of users. In multi-sided markets, this value creation depends on cross-group externalities: the participation of one group affects the benefits obtained by others (Rochet and Tirole, 2003; Parker and van Alstyne, 2005). Yet, platforms do not only facilitate these interactions. Through their rules, incentives, and control over access, they also determine how jointly created value is monetized and distributed. Tension between value creation and value capture is particularly salient in ecosystems built around complementors. Complementors such as app developers, sellers, and content creators provide the variety and innovation that attract users to a platform[1] (Boudreau, 2012; Jacobides et al., 2018), but they are often small and economically dependent relative to the platform owner (Cutolo and Kenney, 2021). Platform governance choices (including pricing structures, revenue sharing arrangements, eligibility requirements, and content policies) can therefore affect complementors' incentives and allocation of effort (Kretschmer et al., 2022).

This paper studies how complementors respond when a platform's governance decision unexpectedly reduces their ability to capture value. We examine content creators operating on YouTube, a dominant advertising-funded platform that acts as an "audience maker" in the sense of Evans (2003). YouTube aggregates viewer attention and connects advertisers with audiences through content from independent creators. At the time of our study, advertising was a central source of creator revenue, making creators' returns dependent not only on audience demand but also on advertisers' preferences and brand safety concerns.[2] This dependence became particularly apparent during the 2017 YouTube "Adpocalypse". Following reports that advertisements were appearing alongside extremist or otherwise objectionable material, major advertisers suspended their spending on the platform. YouTube responded by tightening its content moderation and monetization policies, culminating in stricter eligibility criteria for the YouTube Partner Program in early 2018 (Loh et al., 2026). These changes were intended to restore advertiser trust and protect the platform's overall value proposition. For creators, however, they reduced both the expected level and predictability of advertising-based revenues.

Although creators could respond by reducing their activity or leaving YouTube, exiting was costly. YouTube's accumulated audience, recommendation infrastructure, and central position in online video made

[1] Related work cautions that attracting large numbers of complementors does not mechanically generate sustained network effects, particularly when complementors are unpaid and heterogeneous (see, for example, Boudreau and Jeppesen (2015)).

[2] YouTube is not exclusively advertising-funded. Creators may also receive a share of YouTube Premium subscription revenue and earn through channel memberships, fan funding features, and commerce tools. Our focus is on the monetization channel central to the 2017–2018 setting: the Adpocalypse and subsequent Partner Program changes primarily affected access to advertising revenue and were motivated by advertisers' brand safety concerns.

it difficult to replace as a source of visibility and discovery. Dependence on YouTube was nevertheless uneven across dimensions: creators had virtually no close substitute for its discovery role, but retained some flexibility in how they monetized the audiences built there.

In particular, creators could affiliate with Patreon, a membership platform that enables recurring direct payments from viewers to creators. While Patreon does not offer a substitute for YouTube's video hosting, recommendations, or audience aggregation, it provides a complementary monetization infrastructure that enables creators to convert established audience relationships into recurring revenue. Consistent with the concept of a platform auxiliary developed by Cutolo et al. (2026), Patreon expands creators' capacity to capture value *outside the focal platform's architecture*, while remaining dependent on that platform for discovery. We characterize this response as *monetization layer bypassing*: creators continue to rely on YouTube for visibility and audience formation while redirecting part of the resulting value capture toward a separately governed monetization channel.

Yet, the availability of an auxiliary monetization channel does not establish whether, or under what conditions, creators will actually use it to reallocate effort and value capture following a governance change on the dominant platform. This motivates our central research question: *How do complementors reallocate effort and value capture when governance on a dominant platform makes its monetization regime less attractive, but the platform remains important for discovery?* We examine whether YouTube's monetization shock induced creators to expand their activity on Patreon, whether this expansion represented incremental investment rather than simply placing existing output behind a paywall, and whether it translated into greater community-funded support. We also examine whether the response varied with creators' exposure to monetization risk and with the engagement of their pre-existing communities. The former captures the "severity" of the shock to the creator's existing monetization model; the latter captures the creator's capacity to convert pre-existing audience engagement into recurring direct financial support.

To answer these questions, we combine historical records collected by Graphtreon with original data from Patreon pages. The resulting monthly panel covers roughly 8,400 video creators active on Patreon between August 2017 and August 2018. We compare Patreon creators with an observed YouTube link prior to the policy change with those without a link, using a matched difference-in-differences design. The comparison allows us to estimate how exposure to YouTube's revised monetization regime affected creators' activity and economic outcomes on Patreon, while accounting for time-invariant creator characteristics and shocks common to all Patreon creators. Our main outcomes distinguish posts restricted to paying members from freely accessible posts. This distinction allows us to determine whether creators expanded directly monetized activity or simply reallocated their stock of content from free to paid access.

We find that creators exposed to YouTube's governance change increased their production of member-only content on Patreon, without a corresponding reduction in freely accessible posts. The response, therefore, reflects an expansion of directly monetized activity rather than a change in the composition of existing Patreon output. Exposed creators also accumulated paying patrons and displayed recurring earnings more rapidly, with complementary specifications implying differential increases of approximately 9%. These effects were substantially stronger among creators whose content faced greater pre-existing monetization risk, that is, among those for whom stricter advertiser suitability standards posed the largest threat to monetization. Creators with more engaged pre-shock communities also tend to exhibit a stronger paid content response, although this evidence is more sensitive to specification and weighting choices. Overall, the findings indicate that a governance intervention on one platform can reallocate both complementor effort and value capture toward another platform, while allowing complementors to preserve access to the focal platform's discovery infrastructure.

Our results contribute to the literature in three main ways. First, we extend research on platform governance and content moderation. Existing studies show that changes in advertising and monetization policies affect creator output, content diversity, and strategic positioning on the platform implementing the policy (Loh et al., 2026; Kerkhof, 2024; Rasskazova, 2026). Our study documents a different adjustment: governance can generate spillovers beyond the focal platform by redirecting complementor activity toward an auxiliary platform with a different monetization logic. Second, we contribute to research on complementor strategy and platform bypassing. Prior work shows that complementors reallocate investment in response to platform-level shocks (Wen and Zhu, 2019; Loh and Elsas-Nicolle, 2026; Elsas-Nicolle and Loh, 2025) and that platform participants may circumvent an intermediary by moving transactions off-platform (Gu, 2024; Hagiu and Wright, 2024). Conceptually, this form of bypassing differs from both platform switching and conventional disintermediation: it involves shifting part of monetization to Patreon without necessarily replacing YouTube's discovery function, or privately "completing the same transaction". Third, we provide empirical evidence on the role of platform auxiliaries in reshaping value architecture (Cutolo et al., 2026). By allowing creators to monetize relationships built on YouTube, Patreon shifts where value is captured across interconnected platforms. Our findings highlight the double-edged nature of platform governance: policies designed to protect advertiser trust and ecosystem quality can also encourage productive complementors to reduce their dependence on the platform controlled monetization layer.

2 Prior work

2.1 Platform governance and complementor responses

Digital platforms do not only facilitate interactions among market participants; they actively govern the ecosystems that form around them. Through architectural choices, participation requirements, incentive systems, and control mechanisms, platform owners coordinate the activities of autonomous complementors while retaining authority over the conditions under which they participate and capture value (Tiwana, 2013). Kretschmer et al. (2022) similarly conceptualize platform ecosystems as "meta organizations" in which governance structures coordinate otherwise independent participants and influence which assets and capabilities they deploy through the platform. This suggests that governance choices affect not only the creation of ecosystem value but also its distribution between platform owners and complementors. Demand-side scale and user preferences also shape the strategic interaction between platform owners and complementors, generating a tension between ecosystem growth, complementor investment, and value appropriation (Panico and Cennamo, 2022).[3]

Governance decisions are particularly consequential when complementors depend on a dominant platform for market access and visibility. Rietveld et al. (2020) show that platform governance strategies co-evolve with platform dominance and have important consequences for complementor performance. Cutolo and Kenney (2021) characterize platform-dependent entrepreneurs as legally autonomous actors whose access to audiences and commercial opportunities nevertheless depends on decisions made by the platform owner. Because complementors generally cannot contract over future changes in access, eligibility, or revenue-sharing rules, their investments remain exposed to governance decisions that they do not control.[4]

Complementors, however, can adapt their strategies in response to these decisions. They may adjust their level of participation, redirect investment, reposition their products, or expand their activities beyond the focal platform. For example, Wen and Zhu (2019) show that app developers facing a greater threat of entry by the platform owner reduce innovation in affected applications, while redirecting innovative effort toward unaffected and newly developed applications. This illustrates a broader strategic response to platform-level changes: complementors may reallocate activity rather than simply withdraw from the ecosystem.

Our study examines a distinct form of such reallocation: the creators in our setting do not primarily

[3] Our setting identifies an additional "adjustment" angle: complementors can preserve access to the focal platform's demand-side scale while redirecting part of value capture toward another platform architecture.

[4] More generally, interdependent actors may cooperate to create joint value while competing to capture it (Panico, 2017). The allocation of decision and property rights can therefore affect the weaker party's incentives to invest in collaborative value creation (Gambardella and Panico, 2014). In platform ecosystems, control over audience access, technical infrastructure, and monetization rules creates an analogous asymmetry between the platform owner and its complementors.

respond by switching between competing discovery platforms. Instead, they continue to rely on YouTube for visibility and audience formation while expanding their use of a separately governed channel for value capture. We refer to this response as *monetization layer bypassing*. It differs from conventional multihoming because the two platforms do not perform equivalent functions: YouTube remains the principal discovery and distribution infrastructure, whereas Patreon provides an alternative means of monetizing the audience relationships developed there. It is also distinct from platform envelopment, in which a platform leverages overlapping user relationships by bundling its functionality with that of another platform (Eisenmann et al., 2011). Creators instead combine two functionally differentiated architectures,[5] without Patreon absorbing or reproducing YouTube's discovery function.

2.2 Monetization governance and creator behavior

Content moderation and monetization eligibility are central instruments of governance in advertising-funded platforms. Such platforms must preserve advertiser trust while maintaining content suppliers' participation and investment incentives. These objectives need not be aligned. Rules intended to protect brand safety may reduce the set of content eligible for advertising or make monetization less predictable, thereby changing the expected returns to creator activity. Liu et al. (2022) formalize the relationship between monetization and moderation incentives, showing that platforms' preferred moderation strategies depend on their revenue models. In particular, advertising-funded platforms have stronger incentives to moderate extreme content than platforms financed through membership payments.

Recent empirical work documents how such changes affect creator behavior on YouTube: exposure to the March 2017 advertiser boycott reduced weekly uploads by roughly 10 percent (Rasskazova, 2026), and the subsequent tightening of Partner Program eligibility criteria reduced both upload frequency and content diversity (Loh et al., 2026), whereas an expansion in advertising opportunities induced creators to differentiate their content (Kerkhof, 2024). These findings highlight that advertising-related governance shapes the quantity, diversity, and strategic positioning of content supplied on the platform.

Related evidence indicates that complementors can also respond by changing their monetization model. For example, Kesler (2022) studies Apple's App Tracking Transparency policy, which restricted third party applications' ability to track users, and finds that developers shifted away from advertising-based monetization toward paid and in-app purchase models. This result is closely related to our mechanism because it shows that restrictions affecting one revenue model can induce complementors to increase their reliance on

[5] This differentiation is related to work showing that distinctive platform positioning can improve performance when a platform system is sufficiently differentiated from its rivals (Cennamo and Santalo, 2013). We do not, however, conceptualize Patreon as a direct competitor seeking to displace YouTube; it supplies a narrower auxiliary monetization function.

alternatives.

The existing evidence, however, has largely focused on adjustments made *within* the platform or product ecosystem in which the policy change occurred. We instead study whether a governance shock on a dominant discovery platform redirects creator effort and value capture toward a distinct platform operating under a different monetization logic. The response we are interested in, therefore, is not only a change in the content supplied on YouTube or in the revenue model associated with the same product: it is a cross-platform reallocation in which creators preserve access to YouTube's audience while intensifying their activity on Patreon.

2.3 Bypassing and auxiliary monetization

Our study also relates to research on platform bypassing and disintermediation. In the literature on *platform leakage*, groups of users initially connect through a platform but subsequently transact outside it, thereby limiting the platform's ability to appropriate value from the interaction it facilitated (Wang and Wright, 2020; Hagiu and Wright, 2024). The possibility of bypassing arises because, once users have found one another, gathered enough information, and developed sufficient trust, they may no longer require all of the platform's matching, communication, or transactional services. Gu (2024) provides empirical evidence on the role of external technologies in facilitating such behavior. Exploiting the Skype blockade in mainland China, she finds that restricting access to an external communication channel reduces off-platform transactions, particularly when search, communication, or switching costs are high. Her analysis shows how users can circumvent the particular functions through which an intermediary supports and captures value from a transaction.

The mechanism in our setting differs from "traditional" platform leakage in an important respect. Patreon does not simply allow a creator and a viewer to complete the same transaction privately that would otherwise occur on YouTube. Nor does Patreon replace YouTube's principal role in hosting, recommending, and distributing video content. Creators may continue to publish on YouTube, and viewers may continue to discover and consume their content there. What changes is the channel through which part of the creator's compensation is generated. Patreon allows creators to establish recurring direct payment relationships with audience members while remaining part of YouTube's discovery ecosystem. Monetization layer bypassing is therefore partial and functional: creators circumvent the focal platform's monetization architecture without circumventing its discovery infrastructure.

Research on subscription-based crowdfunding helps explain how this alternative infrastructure operates. Lin et al. (2022) characterize platforms such as Patreon as creator centered freemium systems in which creators combine publicly accessible content with private content reserved for paying supporters. They find that information control strategies, particularly the use of private posts, are associated with growth in creators' backer bases and "fan" engagement. Their analysis concerns creators' strategic choices within the subscription platform. We examine a different antecedent of those choices: whether an external governance shock on a dominant discovery platform induces creators to intensify their use of the subscription-based monetization channel.

The concept of *platform auxiliaries* provides a broader framework for understanding this relationship. Cutolo et al. (2026) define platform auxiliaries as third party actors that help complementors create and capture value within and beyond a focal platform, including by augmenting complementor capabilities, providing platform-specific information, or supporting activities outside the focal ecosystem. Patreon performs the latter function in our setting. Indeed, it allows creators to monetize audience relationships beyond YouTube's advertising-driven model without replicating its scale or discovery capabilities.

We extend this perspective by treating the use of an auxiliary as a strategic response to focal platform governance. Whereas the existence of an auxiliary expands the set of activities available to complementors, a deterioration in the focal platform's monetization regime can change the intensity with which complementors use it. This connects the literature on platform governance to research on bypassing without treating Patreon as either a direct competitor to YouTube or a conventional means of completing the same transaction off-platform. Our study asks whether governance on a dominant discovery platform redirects complementor effort and community-funded value capture toward an auxiliary monetization architecture. In doing so, it identifies monetization layer bypassing as a cross-platform response through which complementors can reduce their dependence on a platform controlled revenue model while continuing to benefit from the focal platform's discovery infrastructure.

3 Theoretical framework

We develop a framework in which complementors depend on a dominant platform for two distinct functions: reaching an audience and converting that attention into revenue. A governance change can reduce the attractiveness of the platform's monetization regime without eliminating its value for discovery. Under these conditions, complementors need not choose between remaining fully dependent on the platform and exiting it. They may continue to use the focal platform for discovery while expanding activity in a separately governed monetization channel. We characterize this response as monetization layer bypassing and derive predictions about alternative-channel investment and its heterogeneity across complementors.

Figure 1. Intermediaries in a multi-sided market

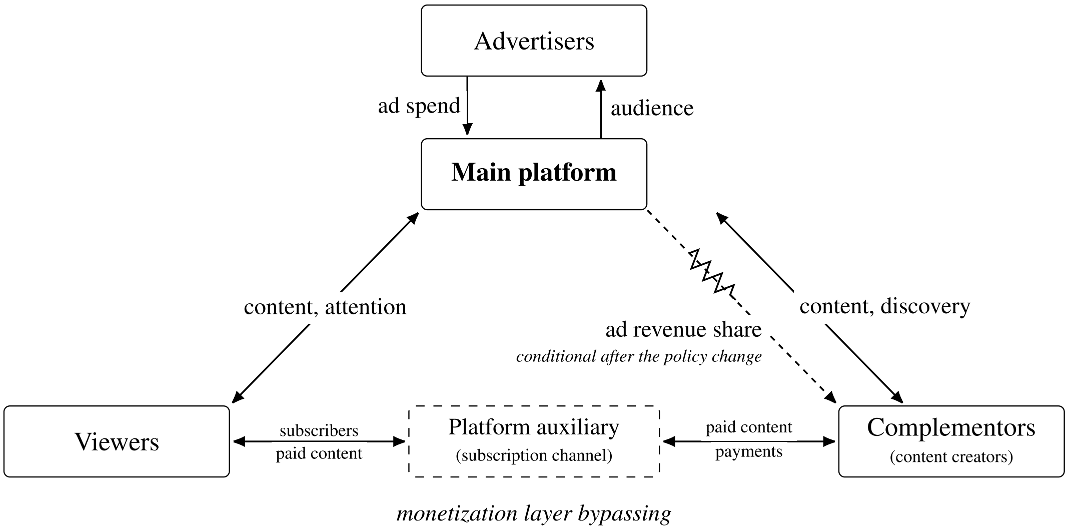

3.1 Decoupling discovery and monetization

Consider a multi-sided market in which a dominant digital platform coordinates interactions among content consumers, complementors, and advertisers (see Figure 1). The platform operates as an audience maker in the sense of Evans (2003): it aggregates consumer attention and matches viewers with content produced by independent complementors. Complementor output attracts viewers, viewer attention attracts advertisers, and the platform returns a portion of the resulting advertising revenue to complementors through revenue sharing and monetization programs. Complementors are therefore central to value creation, but the platform governs the conditions under which they capture part of that value.

Complementor dependence in this setting has two dimensions. First, complementors rely on the platform's discovery infrastructure (its accumulated audience, recommendation technology, search functions, and distribution system) to gain visibility and attract direct attention to their content. Second, they rely on the platform's monetization architecture (the platform's eligibility criteria, content classifications, enforcement procedures, and revenue sharing rules) to convert that attention into income. These functions are usually "bundled" within the same platform ecosystem, but they are not strategically inseparable. Discovery depends primarily on the platform's audience and matching infrastructure, whereas monetization depends on governance rules that the platform owner can revise. Thus, a complementor may derive substantial value from the first function even when the terms governing the second become less favorable.

This separation creates a role for platform auxiliaries. Cutolo et al. (2026) define platform auxiliaries as independent actors that provide resources supporting complementors' value creation and value capture activities within and beyond a focal platform. An auxiliary need not reproduce the focal platform's complete set of functions. Its value may instead lie in addressing a narrower complementor need that the platform owner imperfectly serves. In our setting, Patreon does not replicate YouTube's audience scale, video distribution infrastructure, or discovery capabilities. It provides a separately governed channel for creators to monetize user relationships through recurring direct payments. Creators can thus continue to rely on YouTube for attention while using Patreon to change how part of the resulting value is captured.

The viability of this structure depends on the nature of the audience relationship. Some assets are closely tied to the focal platform: recommendations, search ranking, in-platform reputation signals, and access to the platform's installed user base derive their value from the platform owner's infrastructure. Other assets may instead be complementor-specific. Viewers may develop familiarity with a particular creator, identify with the creator's output, or participate in a community organized around that creator.[6] When such relationships are tied to the creator rather than exclusively to the platform interface, they can potentially be mobilized across platforms.

This form of audience portability does not require viewers to leave the focal platform or change where they consume content. Viewers may continue to discover and watch a creator's output through the focal platform while supporting that creator through a separate payment channel. What moves is not necessarily audience attention, but the transaction through which part of that attention is monetized. Crowdfunding and patronage systems make such a separation feasible by enabling many users to make relatively small direct contributions (Belleflamme et al., 2014; Mollick, 2014). The focal platform may therefore remain indispensable for discovery without remaining the exclusive site of value capture. Subscription-based crowdfunding typically allows creators to combine public and member-only activity. Lin et al. (2022) characterize this architecture as a creator-centered freemium model in which freely accessible content attracts and maintains an audience, while private content and exclusive benefits provide reasons to become a paying supporter. Their analysis links private posting strategies to growth and engagement among supporters on the subscription platform. In our setting, this architecture allows an audience formed or maintained through YouTube to support a creator through recurring Patreon memberships.

Creator investment and audience support may consequently reinforce one another. Increasing member-only content can make membership more attractive, while a larger patron base raises the expected return to producing additional member-only content. Changes in creator activity, patrons, and earnings do not follow a simple sequence in which either supply or demand moves first. We treat patrons and earnings as measures of the community-funded value created through the auxiliary channel, not as evidence of a

[6] Research on online communities shows that social capital, reciprocity, and repeated interaction can sustain contribution and participation in computer mediated settings (Wasko and Faraj, 2005; Faraj and Johnson, 2011).

separately identified audience-side mechanism. The resulting configuration is neither a complete platform exit nor conventional multihoming across equivalent platforms. The complementor remains embedded in the focal platform's discovery ecosystem but reduces exclusive dependence on its monetization rules.

3.2 Governance shocks and alternative channel investment

Monetization rules are a central instrument of platform governance in advertising-funded ecosystems. Eligibility thresholds, advertiser suitability classifications, enforcement procedures, and content policies jointly determine whether complementor activity can generate advertising income. A tightening of these rules can therefore constrain value capture even when the relevant content remains available for distribution and continues to attract viewers.

Such a change can affect complementors through two related channels. First, it may reduce the expected level of advertising-based returns. Higher eligibility thresholds, more restrictive definitions of advertiser suitable content, or more selective access to advertisements lower the probability that additional activity will generate income. Second, it may reduce the predictability of those returns. When enforcement is opaque, or its application is difficult to anticipate, complementors face greater uncertainty about whether particular content will qualify for monetization and whether eligibility will be maintained after publication. The two channels have somewhat different implications, in principle. A reduction in expected revenue lowers the mean return to activity dependent on the focal platform's advertising system. Greater unpredictability increases the risk associated with that activity, thereby increasing the value of diversification. We do not attempt to identify these channels separately. For the strategic response we study, they operate in the same direction: both reduce the relative attractiveness of exerting additional effort. A deterioration in monetization opportunities does not necessarily induce complementors to exit the focal platform. Exit would "protect" them from its monetization rules but would also sacrifice access to its discovery infrastructure.[7] An auxiliary offers an intermediate response. It allows complementors to retain the focal platform's discovery benefits while making part of their income less dependent on the monetization rules that have become less attractive.

We define *alternative channel investment* as the expansion of complementor activity in a separately governed value capture channel that does not require leaving the focal platform's discovery infrastructure. The construct is narrower than the existence of an account or link on an auxiliary platform. It requires activities that increase the auxiliary channel's economic value, such as creating content or offering benefits available only to paying members. It is also distinct from transferring the complementor's entire activity

[7] As discussed earlier, when the platform's accumulated audience and recommendation system are difficult to reproduce, the cost of abandoning discovery may exceed the benefit of avoiding its monetization regime.

to another platform. The complementor continues to use the focal platform for visibility while committing additional effort to a different monetization architecture.

This logic resembles other forms of strategic investment redirection, although the relevant mechanism differs.[8] In our setting, redirection occurs across monetization channels rather than product markets: creators preserve the focal platform's discovery function while increasing activity whose returns are generated through an externally governed membership channel. The distinction between investment and content "reuse" is important for the observable prediction. An increase in member-only content does not necessarily imply an expansion of activity if creators place an existing stock of public output behind a paywall. Under that scenario, paid content would rise while free content declined correspondingly. Alternative-channel investment instead implies an increase in activity tied to direct monetization without a compensating reduction in publicly accessible activity. The auxiliary becomes a destination for incremental activity rather than a place to "reuse" or "repackage" existing output.

This prediction does not require creators to face a fixed effort budget. Free and paid activities can complement one another: public content can attract users, showcase the creator's output, and promote paid offerings. Consistent with this, evidence from knowledge-sharing platforms shows that paid opportunities may stimulate rather than crowd out free contributions among users of the paid feature (Wang et al., 2022; Cong et al., 2025). Our argument therefore predicts greater activity directly tied to the alternative monetization channel without a corresponding reduction in public activity on that channel.[9]

> **Hypothesis 1** *(alternative channel investment)*. Following a governance induced deterioration in the expected level or predictability of monetization on the dominant platform, exposed complementors increase directly monetized activity in an alternative channel without a corresponding reduction in publicly accessible activity on that channel.

3.3 Heterogeneous complementor responses

The response predicted by Hypothesis 1 should not be uniform across complementors. A governance change may substantially threaten the existing monetization model of some complementors while having limited consequences for others. Moreover, complementors differ in their ability to translate an incentive to diversify into a viable direct monetization strategy. We distinguish between two conditions: *displacement pres-*

[8] Wen and Zhu (2019), for example, show that developers facing an increased threat of platform owner entry reduced innovation in affected applications while redirecting innovative activity toward unaffected and new applications.

[9] Because our dataset captures Patreon activity but not creators' complete allocation of effort across Patreon, YouTube, and other platforms, the empirical implication is necessarily narrow. Indeed, we test whether creators exposed to the focal platform governance change expand member-only Patreon content and whether this occurs without a corresponding decline in freely accessible Patreon posts.

sure, which creates an incentive to reduce dependence on the focal monetization regime, and *conversion capacity*, which determines whether the complementor can successfully create (and capture) value through the alternative channel.

3.3.1 Monetization risk exposure

Advertising-funded platforms evaluate content partly according to advertisers' preferences and brand safety concerns. Content, therefore, differs in its exposure to monetization risk even before a governance change occurs. Material involving politics, social controversy, sensitive events, or other advertiser sensitive subjects may be more likely to encounter restrictions than comparatively safe material. When the platform tightens its eligibility or advertiser suitability standards, the same formal rule change can consequently produce different economic consequences across content categories.

For creators whose output lies close to the advertiser sensitive boundary, stricter screening or enforcement increases the likelihood that content will not qualify for advertising and may make monetization decisions harder to anticipate. The expected marginal return to advertising dependent activity, therefore, declines more sharply for these creators. For complementors whose content remains within the advertiser safe region, the focal platform's monetization system remains comparatively reliable, and the incentive to expand a separate revenue channel is weaker. A direct payment channel is particularly relevant to the first group because it weakens the connection between the focal platform's advertiser suitability decisions and creator income. The auxiliary platform remains governed by its own rules, but payments made through it do not depend directly on whether the focal platform classifies a particular video as suitable for advertisers. Complementors facing higher monetization risk should therefore derive greater strategic value from expanding activity tied to membership payments.

This variation in monetization risk captures heterogeneity in the intensity of the shock rather than a capability developed in response to it. Every exposed complementor may face some deterioration in expected monetization conditions, but the magnitude of that deterioration varies with the advertiser's sensitivity to the content already produced. The pressure to redirect activity should accordingly be strongest where the governance change most severely threatens the existing advertising-based revenue model.

> **Hypothesis 2** *(monetization risk exposure).* The governance induced increase in alternative channel investment is stronger for complementors whose content has greater pre-shock exposure to advertiser sensitivity and monetization risk.

3.3.2 Community engagement and conversion capacity

Displacement pressure provides a reason to seek an alternative, but it does not guarantee that direct monetization will succeed. The creator must also be able to translate audience attention into recurring support. This conversion capacity depends not only on how many users consume the creator's output, but also on the nature of the relationship between the creator and the community.

Popularity and engagement are distinct in this respect. A large audience provides a broad pool of potential supporters, but its members may interact only weakly with the creator and have little willingness to follow links or pay for additional access. A more engaged community is characterized by active interaction, familiarity, reciprocity, and an expectation of continued participation. These characteristics increase the likelihood that users will respond to a direct monetization offer and that the creator will find investment in member-only content worthwhile. Pre-existing engagement is therefore an observable indicator of a "relational asset" that may be mobilized across platforms. Highly engaged users are more familiar with the creator's output, have stronger reasons to value continued production, and may have a higher willingness to pay for exclusive content or closer interaction. For the creator, these relationships increase the expected return to investing in a membership channel. By contrast, a creator with a weakly engaged audience may face substantial displacement pressure but still be unable to generate direct support.

Conversion capacity differs from monetization risk exposure in two respects. First, it relates to the expected return on the alternative channel rather than the expected cost of continued reliance on the focal platform. Second, it is a pre-existing and unevenly distributed relational endowment. Exposure to a governance change can arise immediately from the nature of the creator's content, whereas an engaged community generally must be developed through prior interaction. The engagement moderator should therefore be more contingent than the monetization risk moderator: its effect should be stronger among creators who had already developed a community capable of supporting direct monetization before the shock.[10]

> **Hypothesis 3** *(community conversion capacity).* The governance induced increase in alternative channel investment is stronger for complementors with more highly engaged pre-shock communities.

The two moderators are conceptually distinct but need not be statistically independent. Monetization risk exposure provides the "push away" from exclusive reliance on advertising-based monetization, whereas community engagement provides the capacity to generate returns from direct support. A creator may face

[10] Our empirical measure, pre-shock comments on creator content, does not directly observe audience portability, willingness to pay, or movement between platforms. It captures active interaction and is therefore treated as an indicator, rather than a complete measure, of conversion capacity. The prediction is that complementors with stronger pre-existing community engagement exhibit greater expansion in directly monetized activity after the governance shock.

strong pressure but have little capacity to act on it, or have an engaged community but little reason to alter an advertising model that remains reliable.

3.3.3 Scope conditions

The framework applies most directly under three conditions. First, the focal platform must remain valuable for discovery and audience aggregation. When governance removes access to distribution altogether, complementors are more likely to exit or migrate to a different discovery platform. Second, the governance change must primarily deteriorate the conditions of value capture rather than eliminate the complementor's ability to reach users. Monetization layer bypassing is possible precisely because creators can preserve the discovery function while "escaping", to some degree, the monetization model. Third, an auxiliary channel must allow complementors to transact with, or obtain support from, an audience created elsewhere. This requires both a means of communicating the alternative offer and an audience relationship that is at least partly creator-specific and "portable" across platforms, rather than platform-specific. Alternative channel investment is therefore a conditional response to platform governance, not a universal consequence of stricter rules. It becomes strategically relevant when discovery remains difficult to replace, monetization becomes less attractive, and creator–audience relationships can support value capture through a separately governed channel.

4 Institutional setting

This section describes the two platforms central to the analysis and the sequence of events that produced the focal governance shock. YouTube provides creators with video distribution, audience discovery, and advertising-based monetization. Patreon provides a separately governed membership channel through which creators can obtain recurring payments from their audiences. The 2017 advertiser boycott and YouTube's subsequent policy response affected creators' access to monetization while leaving their formal access to YouTube's distribution infrastructure largely intact.

YouTube and Patreon Digital content platforms rely on independent creators to produce the material that attracts users and advertiser demand. On YouTube, creators upload and distribute videos, while its search and recommendation systems help viewers discover content. As of July 2026, YouTube reports billions of monthly logged-in users and more than 20 million videos uploaded to the platform each day; YouTube

Shorts alone averages over 200 billion daily views.[11] Its large installed audience and discovery infrastructure make it difficult for individual creators to reproduce the visibility available through the platform.

During our study period, participation in the YouTube Partner Program (YPP) was the primary way creators earned advertising revenue. Eligible creators could allow advertisements to appear alongside their videos and receive a share of the resulting revenue. Access to this revenue depended on YouTube's eligibility requirements, content policies, advertiser suitability classifications, and enforcement decisions. Creators, therefore, depended on YouTube in two related but distinct ways: the platform helped them reach viewers, and it governed whether and how the attention generated by those viewers could be converted into advertising income. Creators could also pursue revenue outside YouTube's own monetization architecture. One such option was Patreon, a membership platform that allows users to make recurring payments directly to creators. Creators typically use Patreon to offer paying supporters benefits such as member-only posts, early access to content, supplementary material, or closer interaction. A YouTube subscriber can follow a creator and access their publicly available videos "for free", whereas a Patreon patron enters into a recurring financial relationship with the creator.[12]

Importantly, Patreon does not provide a substitute for YouTube's core functions. It does not offer a comparable video distribution system, recommendation infrastructure, or general audience discovery environment. Its role is complementary: it provides the payment and access control infrastructure through which creators can monetize relationships with users who may have discovered them elsewhere. Creators can place links to their Patreon pages in YouTube videos, channel pages, or video descriptions, as illustrated in Figure A.1. These links allow viewers reached through YouTube to support the creator via a separate platform. The two platforms, therefore, serve different functions in creators' activities: YouTube remains important for visibility, audience building, and content distribution, while Patreon enables creators to receive recurring payments without relying on YouTube's advertising eligibility decisions. A creator can consequently continue publishing and attracting viewers on YouTube while increasing the share of income generated through Patreon. These institutional features make Patreon an auxiliary monetization channel in the sense developed in Section 3: it allows creators to redirect part of value capture without replacing YouTube's discovery infrastructure.

The advertiser boycott and YouTube's revised monetization regime In March 2017, news reports revealed that advertisements from prominent brands and public organizations had appeared alongside extrem-

[11] See https://blog.youtube/press/ (Accessed July 21, 2026).

[12] A "patron" on Patreon pays a monthly subscription of at least \$1. Patreon retains between 5 and 12% of the monthly subscription revenue generated by the creator's patrons.

ist or otherwise objectionable videos on YouTube. Major advertisers responded by suspending or reducing their spending on the platform (Mostrous and Dean, 2017; Nicas, 2017). The episode and the broader period of advertiser pressure and platform policy revision that followed became known as the YouTube "Adpocalypse".[13] The controversy exposed a central tension in YouTube's advertising-funded ecosystem. Advertisers valued the platform's audience reach but were concerned about the content their brands might be associated with. Creators, by contrast, depended on advertising revenue but did not fully control how YouTube classified their videos or how advertisers responded to the surrounding content environment. YouTube, therefore, faced pressure to strengthen advertiser trust while maintaining creators' participation and investment incentives.

Prior to the boycott, entry into YouTube's monetization program was relatively permissive. Content enforcement combined automated detection, user reporting, and human review, but contemporary reporting documented difficulties in identifying and consistently handling advertiser sensitive material. Monetization decisions were also less restrictive than they became following the boycott. The subsequent policy response involved both expanding enforcement and tightening access to advertising revenue.

The March 2017 boycott preceded our observation window. The empirical shock examined in this paper is the later rollout of YouTube's revised enforcement and monetization regime, which occurred between December 2017 and February 2018. In December 2017, YouTube announced an expansion of the resources devoted to identifying inappropriate content and enforcing its policies. The announcement explicitly linked these measures to advertiser trust. YouTube's then-CEO stated that the platform sought to give advertisers greater assurance about the content accompanying their advertisements while also protecting legitimate creators from the consequences of harmful activity elsewhere on the platform.[14]

In January 2018, YouTube announced a substantial tightening of eligibility for the YPP. Beginning on February 20, channels generally needed at least 1,000 subscribers and 4,000 hours of watch time during the previous twelve months to qualify for monetization.[15] Channels that did not meet these thresholds could continue to upload videos, maintain their audiences, and obtain exposure through YouTube, but they could no longer earn advertising revenue through the program. Channels that remained eligible were subject to closer review for compliance with YouTube's Community Guidelines and advertiser suitability policies. Repeated or serious violations could result in limited monetization, demonetization, or removal from the program.

[13] See also `https://www.theverge.com/2019/4/5/18287318/youtube-logan-paul-pewdiepie-demonetization-adpocalypse-premium-influencers-creators`.

[14] See YouTube, *Expanding our work against abuse of our platform*, December 2017. `https://blog.youtube/news-and-events/expanding-our-work-against-abuse-of-our/`

[15] See YouTube, *Additional Changes to the YouTube Partner Program (YPP) to Better Protect Creators*, January 2018. `https://blog.youtube/news-and-events/additional-changes-to-youtube-partner/`

The revised regime, therefore, operated primarily on the monetization front. It did not formally prevent affected creators from publishing content or maintaining a presence on YouTube. Instead, it changed the conditions under which their activity could generate advertising income. The thresholds directly removed smaller channels from the monetization program, while the expanded review and advertiser suitability standards increased the relevance of content classification for creators that remained eligible. The policy change could consequently reduce both the expected level and the predictability of advertising-based returns.

The consequences were unlikely to be uniform across creators. YouTube's revised approach was motivated by advertiser concerns about brand safety, and content categories differ in the extent to which they raise such concerns. Creators producing political commentary, controversial humor, sensitive news, or other advertiser risky material were more likely to have their videos receive limited or no advertising. Creators producing comparatively safe content faced less direct exposure. The same governance change could therefore impose substantially different monetization pressure depending on the type of content a creator produced, consistent with evidence that the Adpocalypse and subsequent YPP changes affected creator behavior heterogeneously (Loh et al., 2026; Rasskazova, 2026).

For creators already active on Patreon, the policy rollout created an opportunity to reduce reliance on YouTube's advertising system without giving up the platform's discovery benefits. They could continue using YouTube to publish content and reach viewers while increasing member-only production or encouraging recurring support through Patreon. The institutional setting thus combines the three conditions required by our framework. First, YouTube remained valuable for audience discovery and distribution. Second, the December 2017–February 2018 policy rollout deteriorated the terms of the platform's monetization without eliminating access to discovery. Third, Patreon offered a separate channel through which creators could monetize relationships with audiences reached via YouTube.

5 Data

5.1 Data sources and panel construction

We construct a monthly panel of Patreon creators by combining historical records collected by Graphtreon[16] with information scraped from Patreon pages. Graphtreon tracks creators with at least one paying patron and records monthly measures of creator activity, including patron counts, publicly displayed earnings, creator categories, and links to external platforms. The observation window extends from August 2017 through August 2018. It therefore covers the months preceding YouTube's revised monetization regime, the

[16] See https://graphtreon.com/.

Figure 2. Examples of posts on Patreon

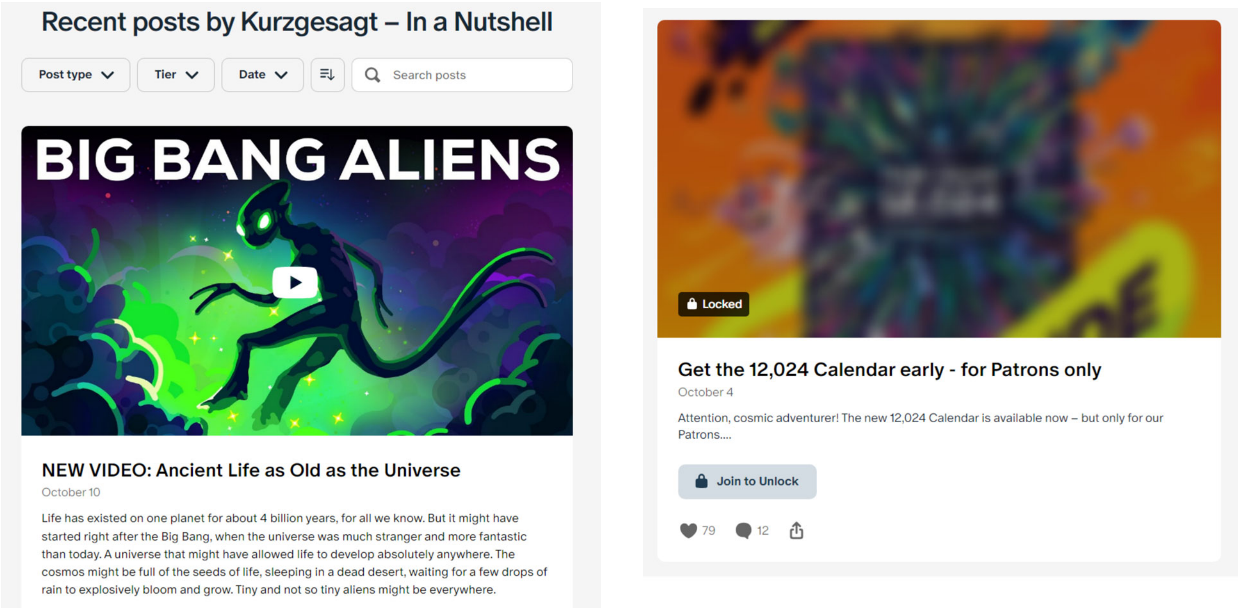

(a) Public post **(b)** Patrons-only post

Notes: The figure illustrates the information visible on a creator's Patreon page for a publicly accessible post and a post restricted to paying members. The example is drawn from the Patreon page of Kurzgesagt.

December 2017–February 2018 rollout period, and the subsequent months in which creators could adjust their Patreon activity.

We supplement these records with a spring 2022 scrape of Patreon pages, collecting the historical posts that were accessible at that time. At the post level, we observe the publication date, title, displayed likes and comments, media type, and access status, which identifies whether a post is publicly available or restricted to paying members. Figure 2 illustrates the two forms of content.

We aggregate the post-level records to the creator–month level using each post's publication date. *Free content* consists of posts accessible without a paid membership. We divide member-only posts into base content, available at the lowest observed membership tier, and premium content, which requires a higher membership tier. The empirical analysis combines these two categories into *paid content*. *Total content* is the sum of free and paid posts. We also construct monthly counts of images and videos and aggregate the likes and comments displayed on posts published by a creator during the month.

The ability to distinguish between paid and freely accessible content makes these data particularly well suited to testing our predictions. Because a paid post is one whose accessibility is conditioned on a recurring membership payment, the count of paid posts is a direct behavioral measure of *alternative channel investment* as defined in Section 3: activity that increases the economic value of the auxiliary channel. Free

Table 1. Number of observations in our samples

Sample restriction	Creators	YouTube-linked	Comparison	Creator–months
Full scraped universe, Aug. 2017–Aug. 2018	11,371	8,194	3,177	131,006
Video category, excluding adult content	10,248	7,553	2,695	116,773
Main stayer sample	8,436	7,336	1,100	108,025
CEM common-support sample	7,065	6,030	1,035	90,396

Notes: Counts are applied sequentially. YouTube-linked and comparison counts use the treatment definition used in estimation: a YouTube link observed through December 2017, strictly before the January 2018 post period. The stayer restriction requires observation in the window at least once before December 2017 and at least once after February 2018. The CEM row retains creators in matched strata containing both YouTube-linked and comparison creators.

posts, by contrast, support audience acquisition and community maintenance but are not gated on payment. Observing the two separately allows us to distinguish an expansion of directly monetized activity from a within-channel reallocation of content behind a paywall.[17]

5.2 Sample construction

Table 1 reports the sequential construction of the analysis sample. The initial August 2017–August 2018 panel contains 11,371 creators and 131,006 creator–month observations. We restrict this universe to creators assigned to Patreon's *Video* category and exclude creators classified as producing adult content. The table reports creator and creator–month counts at each stage, split according to the treatment definition used in estimation (Section 6). Our main analysis focuses on *stayers*: creators observed within the window at least once before December 2017 and at least once after February 2018. It isolates within creator adjustment from compositional change generated by entry and exit, and it aligns with the theory, which focuses on how complementors that had *already* established an auxiliary presence reconfigure value capture rather than the decision to affiliate with an auxiliary for the first time. Robustness analyses reintroduce entrants and exiters and impose a strictly balanced panel.

The final row of Table 1 applies the common support indicator from the coarsened exact matching procedure (see Section 6), keeping creators assigned to strata containing both creators with and without a YouTube link. Recorded entry into the Video category varies substantially over the window and rises sharply during the rollout. This motivates the need to separate intensive-margin adjustment among stayers from compositional change. Appendix Figure A.2 reports the monthly entry and exit of content creators.

[17] The combined data enable us to observe both investment in Patreon's membership channel and the economic value it creates, but they impose two coverage restrictions. First, earnings are observed only when a creator chooses to display them publicly. Second, the content data include only pages and historical posts that remained accessible when we conducted the scrape. Deleted pages and posts that were no longer available in 2022 cannot be recovered.

Table 2. Patreon creator characteristics

Variable	*N*	Mean	Median	SD	Min	Max
Patrons	11,371	83.75	10.31	369.15	1.00	12,206.00
Earnings (USD)	9,465	320.88	53.99	1,175.65	0.01	38,877.80
YouTube link (pre-shock)	11,371	0.71	1.00	0.45	0.00	1.00
Contents per month	11,371	4.08	1.46	7.46	0.00	169.92
Comments per month	11,371	18.06	0.77	79.70	0.00	2,076.00
Likes per month	11,371	40.76	1.85	189.87	0.00	6,154.50
Facebook likes	5,917	31,865.07	1,131.60	258,308.59	0.00	8,742,815.00
Twitter followers	8,580	22,887.83	856.37	846,570.16	0.00	55,263,868.00
Monetization risk (1–4)	11,371	3.22	3.00	0.97	1.00	4.00

Notes: One observation per Patreon creator observed from August 2017 through August 2018. Monthly variables are first averaged within creator. Earnings are included only for creators who publicly display them, so the number of observations is smaller for this variable. YouTube link equals one when a YouTube link is observed strictly before December 2017. This descriptive variable uses the stricter pre-rollout definition and differs from the treatment indicator used in estimation, which incorporates link information through December 2017. Monetization risk is coded 1 = High, 2 = Medium, 3 = Low, and 4 = Unclear/unclassified.

5.3 Creator characteristics and outcome variables

Table 2 reports creator-level characteristics for the full sample (i.e., scraped universe of content creators on Patreon). Monthly measures are first averaged within-creator, leaving one observation per individual. The table describes Patreon activity, community engagement, external social media reach, and monetization risk classification. The distributions are strongly skewed to the right: for patrons, earnings, comments, likes, and social media reach, the means are substantially greater than the corresponding medians. The smaller number of observations for earnings reflects creators' decisions not to disclose their monthly earnings publicly.

Two of these characteristics are of particular interest. Community engagement, captured here by the comments a creator's posts attract, is our proxy for *conversion capacity*: it captures repeated interaction and accumulated relational capital, thereby increasing, as discussed in Section 3, a creator's capacity to mobilize its audience through a direct monetization channel.[18] External social media reach (Facebook and Twitter), by contrast, measures potential exposure rather than active engagement, and serves both as a balance check and as a reminder that reach and engagement are distinct elements.

Table 3 describes the creator–month content records for the non-adult Video sample. It separates total, free, base-tier, and premium content and reports the number of images and videos posted. The zero medians for several content categories indicate that the distributions contain a large mass at zero, alongside a comparatively small number of highly active creator–months. The primary content outcomes are *paid content*, *free content*, and *total content*. Paid content combines base-tier and premium posts and therefore

[18] We see it as an indicator of that capacity rather than a direct measure of it, and we note the caveat that pre-shock engagement may partly reflect a creator's pre-existing capability on Patreon itself, a point we return to when interpreting the heterogeneity results.

Table 3. Patreon creator–month content characteristics

Variable	*N*	Mean	Median	SD	Min	Max
Number of contents	116,773	3.64	1.00	7.34	0.00	271.00
Number of free contents	116,773	1.31	0.00	3.81	0.00	162.00
Number of base-tier contents	116,773	1.38	0.00	3.75	0.00	113.00
Number of premium contents	116,773	0.95	0.00	3.75	0.00	252.00
Number of images	116,773	2.86	0.00	6.38	0.00	230.00
Number of videos	116,773	1.98	0.00	4.97	0.00	216.00
Sensitive/controversial creator	99,164	0.15	0.00	0.36	0.00	1.00
Monetization risk (1–4)	116,773	3.01	3.00	0.99	1.00	4.00

Notes: Creator–month observations for non-adult creators in Patreon's Video category, August 2017 through August 2018. Base-tier content is available at the lowest observed membership tier; premium content requires a higher membership tier. Monetization risk is coded 1 = High, 2 = Medium, 3 = Low, and 4 = Unclear/unclassified.

captures activity whose accessibility depends directly on a recurring membership payment. This variable is our proxy for alternative channel investment. Free content captures publicly accessible posts, which may support audience acquisition, communication, and community maintenance without requiring membership. Observing the two types separately allows us to test the theory's prediction that paid activity rises *without* a compensating decline in free output. The paid-minus-free difference and the paid share of posts serve as supplementary measures of content composition.

Two additional outcomes measure value creation on Patreon. *Patrons* is the number of paying supporters on a creator's page; *earnings* is the monthly amount displayed in U.S. dollars. Because both variables are cumulative stocks, differences in their levels may grow over time simply because one group was already accumulating patrons or earnings faster before the governance change. Such divergence would not necessarily indicate a response to the policy rollout. Our main monetization analyses, therefore, focus on net monthly changes in patrons and earnings, calculated only between consecutive observed months. Because our framework treats value creation as jointly determined by creators and audiences, we interpret these outcomes as measures of community-funded value capture rather than as evidence of a separately identified demand response.

5.4 Presence on YouTube and advertiser sensitivity classifications

Presence on YouTube We identify creators' connections to YouTube from the information contained in the Patreon and Graphtreon records. The YouTube Partner Program imposed almost no eligibility restrictions before February 2018 (Section 4). Thus, a displayed link is a close indicator of participation in, and reliance on, YouTube.[19]

Monetization risks based on topics We measure exposure to advertiser sensitivity through a creator-level classification of the content posted on Patreon. This classification is the basis for our proxy for *displacement pressure*: it captures how close a creator's content is to the advertiser sensitive boundary affected by the governance change, and therefore, how strongly the shock is expected to reduce that creator's access to advertising revenue. Each creator is assigned to one of four mutually exclusive groups. The high sensitivity group contains creators classified as "Sensitive/Controversial". The medium sensitivity group contains "Culture/History/Religion" and "Commentary/Reactions". All remaining categories form the low sensitivity group. Creators classified as "Other/Unclear" remain in a separate unclassified category.[20]

We refer to this variable throughout as *monetization risk*, coded 1 for the high sensitivity group through 4 for unclassified creators. Table A.2 presents some descriptive statistics. In the stayer sample, 1,173 creators fall in the high sensitivity group, 815 in the medium group, and 3,702 in the low group, with 2,746 unclassified.

Monetization risks based on keywords For our robustness analyses presented in Section 7, we consider a second classification of monetization risk. Rather than inferring a creator's subject matter and grouping categories by their likely appeal to advertisers, this measure evaluates the text of a creator's posts against lists of advertiser "restricted terms" spanning violence, harm, profanity, adult content, controversy, sensitive events, shocking material and hate, and assigns each creator to a high, medium or low sensitivity level, with creators whose text provides insufficient signal left unclassified.[21] Thus, unlike the topic-based classification, it measures advertiser sensitivity from the occurrence and contextual use of restricted terms in creators' post text. Because those terms concern both sensitive subject matter and linguistic expression, the measure does not provide a pure distinction between what creators discuss and how they express it. In the stayer sample, the measure classifies 6,216 creators (2,616 high, 802 medium, and 2,798 low), leaving 1,501 with insufficient information and 719 without a classification. We discuss the overlap between the two classifications in Appendix Section A.3.4.

[19] Naturally, it does not record the intensity of that reliance or whether a given creator was individually demonetized, and creators without an observed link may still have been active on YouTube or on other platforms. Table A.1, presented in Appendix A, reports the prevalence of YouTube connections under a stricter descriptive definition, requiring the link to be observed before December 2017, for the full creator universe, the non-adult Video sample, the main stayer sample, and the CEM common support sample. The treatment coding used in the regressions, which uses link information through December 2017, is specified in Section 6.

[20] We provide full details on the categorization procedure in Appendix B.

[21] Appendix B provides full details, including the sources of the words and the procedure used to assign them to categories.

6 Empirical strategy

Sample and outcomes. The unit of observation is a creator–month. The main analysis uses creators observed on Patreon both before and after YouTube's December 2017–February 2018 policy rollout (the *stayers*). This restriction isolates within-creator adjustment among complementors that had already established an alternative monetization channel.

The primary outcomes are paid, free, and total Patreon posts. Paid posts are accessible only to members, whereas free posts remain publicly available. As discussed before, an increase in paid posts without a corresponding decline in free posts therefore distinguishes an expansion of directly monetized activity from the "reuse" of existing production behind a paywall. We also examine patrons and publicly displayed earnings as measures of realized community-funded value capture. Because these variables are accumulated stocks, our preferred specifications use net monthly changes computed across consecutive observations.[22]

Treatment. Treatment status is defined based on whether a YouTube link appears on a creator's Patreon page by December 2017, strictly before the post-treatment period. Therefore, a creator is *treated* when we identify a linked YouTube channel by that date; control creators are those for whom no linked channel is observed.[23] The governance change (the *shock*) unfolded through a sequence of announcements and implementation decisions rather than a single instantaneous cutoff. Indeed, YouTube announced an expansion of its moderation-enforcement guidelines in December 2017, announced the stricter YouTube Partner Program eligibility criteria in January 2018, and enforced them from February 20, 2018.

The post-period begins in January 2018, when the monetization specific policy was announced, so December 2017 (i.e., when enforcement expansion had been announced, but monetization eligibility had not yet changed) remains in the pre-period. We treat December 2017 through February 2018 as a rollout window, shade it in the event study figures, and estimate alternative specifications that start the post period in December 2017 or February 2018.

The resulting estimates are intent-to-treat effects of exposure to the shock: they capture the differential response of creators with a visible pre-shock connection to YouTube, not the causal effect of a documented creator specific loss of advertising revenue. This interpretation is consistent with the theoretical argument because governance changes can affect expected returns and the predictability of monetization, even when individual enforcement outcomes remain unobserved.

[22] Specifications in levels, inverse-hyperbolic-sine transformations, and PPML models provide complementary evidence on magnitude and functional-form sensitivity. We present these alternative results in Section 7.3 and Appendix Section A.4.

[23] Because no information from any post-period month is included in the definition, subsequent linking decisions cannot redefine exposure after the governance change.

6.1 Difference-in-differences specification

We estimate the following difference-in-differences (DiD) specification, using the monthly panel:

$$Y_{it} = \beta\,(\text{YouTube}_i \times \text{Post}_t) + \alpha_i + \gamma_t + \varepsilon_{it} \tag{1}$$

where Y_{it} is an outcome for creator i in month t, YouTube_i identifies creators with a linked YouTube channel before the shock, and Post_t equals one from January 2018 onward. Creator fixed effects, α_i, absorb time invariant differences across creators (i.e., content category, ability, baseline platform use, persistent propensity to invest in Patreon) while month fixed effects, γ_t, absorb shocks common to Patreon creators in a given month, from platform-wide growth and seasonality to broader shifts in the creator economy. Identification thus rests on within-creator changes, not cross-sectional differences in levels: treated and comparison creators may differ substantially in size, popularity, or baseline use of Patreon, but those persistent differences are removed, and what matters is whether the outcomes of treated creators *change* differently after the shock. The coefficient of interest, β, measures the differential post-shock change among creators with a pre-shock YouTube connection relative to those without one. Hypothesis 1 is tested using the sign and statistical significance of the estimated coefficient β for paid content in Equation (1).[24] Because treatment varies at the creator level and outcomes are serially correlated across months, all standard errors are clustered at the creator level, allowing within-creator correlation over the estimation period.

Estimators. Our choice of estimator reflects the distributional and economic properties of each outcome. For content counts, linear fixed-effects models serve as the baseline: their coefficients represent changes in posts per creator-month and allow for direct comparison across paid, free, and total production. Since content outcomes are nonnegative and skewed, Poisson pseudo-maximum-likelihood (PPML) models complement the linear estimates, giving us proportional effects without requiring the conditional variance to equal the conditional mean.[25] For patrons and earnings, first-difference (flow) specifications are preferred since they measure the rate of accumulation rather than the accumulated level; the inverse-hyperbolic-sine and level specifications are reported as robustness and as evidence on economic magnitude, respectively, and PPML provides a further functional-form check.

As in standard DiD models, identification requires that, absent the YouTube monetization shock, Patreon creators with and without linked YouTube channels would have experienced similar changes in outcomes.

[24] Hypotheses 2 and 3 are tested using the additional interaction terms described below.

[25] For a discussion of why the logarithmic transformation is not appropriate for such values, see Benatia et al. (2025) and Cohn et al. (2022).

Similar outcome levels are unnecessary because creator fixed effects absorb time-invariant differences. The key requirement is that no contemporaneous factor generated a differential change among YouTube-linked creators that coincided with the rollout and independently affected Patreon activity. We assess this assumption using event-study specifications that examine pre-rollout trends and the timing of the response.

Coarsened exact matching. To ensure that the difference-in-differences analysis compares creators who are observably similar, we use coarsened exact matching (CEM), since YouTube-linked creators may have been systematically larger and more active before the shock. We coarsen and match on pre-treatment averages of content production, comments, patrons, earnings, and monetization risk; variables that are central to both treatment assignment and subsequent Patreon activity. We discuss the results from the matching procedure in Section 7.1. We also assess balance on characteristics not used in the coarsening procedure, including likes, free content production, and external social media reach (see Table A.3).

6.2 Event study specification

To examine pre-treatment dynamics and the evolution of the estimated response, we estimate the event study counterpart of Equation 1:

$$Y_{it} = \sum_{\substack{k=\text{Aug. 2017} \\ k\neq\text{Nov. 2017}}}^{\text{Aug. 2018}} \beta_k \left(\text{YouTube}_i \times \mathbf{1}\{t=k\}\right) + \alpha_i + \gamma_t + \varepsilon_{it} \tag{2}$$

November 2017, the final month before the rollout window, is the omitted category, so each β_k measures the treated-control difference in month k relative to its November 2017 value. The pre-rollout coefficients serve as a visual diagnostic: movement before the rollout would suggest that the two groups were already on different trajectories. For the monetization outcomes, the event studies are estimated on the net monthly change, consistent with the flow specifications in the corresponding tables.

Also, these event study specifications provide a useful complement to the standard difference-in-differences models because the shock is unlikely to generate an immediate, sudden, response. Creators may need time to expand member-only production, communicate with existing audiences, and build recurring support. Moreover, the policy change unfolded over several months. Our theoretical framework does not imply a discrete break at a single date; a gradual post-rollout response is consistent with adjustment to a governance change implemented progressively.

6.3 Heterogeneous complementor responses

We argued in Section 3 that alternative channel investment depends on two conditions that differ in how broadly they apply: displacement pressure bears on every exposed creator in proportion to their content's advertiser sensitivity, whereas conversion capacity is a relational endowment that many creators lack. We therefore expect the exposure to be the more general feature of the response and the capacity moderator to be more contingent. We test both with the following specification with moderators:

$$Y_{it} = \beta_1 (\text{YouTube}_i \times \text{Post}_t) + \beta_2 (\text{YouTube}_i \times \text{Post}_t \times M_i) + \delta (\text{Post}_t \times M_i) + \alpha_i + \gamma_t + \varepsilon_{it} \tag{3}$$

where M_i is a pre-treatment moderator whose time-invariant main effect, along with its interaction with YouTube_i, is absorbed by the creator fixed effects.[26] Here, β_1 captures the treatment effect for creators in the moderator's baseline category, while β_2 captures the additional treatment effect for creators in the highest category of M. The total treatment effect for high category creators is therefore $\beta_1 + \beta_2$. Our tests of Hypotheses 2 and 3 focus on β_2, which indicates whether the treatment effect differs between the high and low-category groups.

For monetization risk exposure, the moderator compares creators facing high risk against those facing low risk; a broader comparison pools medium- and low-risk creators in the reference group. For community conversion capacity, the moderator is the creators' average pre-treatment number of comments, a measure of active interaction rather than of reach alone. The approach compares the top and bottom terciles of pre-treatment engagement, contrasting highly and weakly engaged communities.[27]

7 Estimation results

We present the main estimation results in three steps. We first report the results of the matching procedure, then turn to the DiD estimates and event studies for the outcomes of interest. We then conclude with robustness tests. Throughout the main tables, we report results for: (i) the full estimation sample; (ii) the common support sample obtained with CEM; and (iii) the matched sample using CEM weights.[28]

[26] The $\text{Post}_t \times M_i$ term allows creators with different moderator values to have different post-period changes even within the comparison group; without it, a differential treatment response could not be separated from a general post-period change associated with the moderator.

[27] Alternative approaches are considered: (i) a median split keeps the full range of creators under a less selective threshold, and (ii) a specification where both moderators enter simultaneously, asking whether each is associated with a stronger response when the other is held constant. Results from these alternative specifications are presented in Appendix Section A.3.5.

[28] The full sample specifications include all creators satisfying the relevant sample restrictions. The common support specifications retain only creators assigned to CEM strata containing both YouTube-linked and comparison creators. The weighted specifications additionally account for differences in the relative numbers of treated and comparison creators within each retained

7.1 Matching quality

Unmatched creators with a linked YouTube channel are larger and more active on several dimensions, with the largest standardized differences in patrons, earnings, free content, and external audience reach, precisely the gap that motivates our matching approach. Table A.3, presented in Appendix A, shows that the CEM weights eliminate most of these imbalances: standardized differences fall to near zero for pre-treatment content production, comments, likes, patrons, earnings, and monetization risk, and remain below 5% for free content and Twitter followers. The one exception is Facebook reach, slightly above 10%. The matched estimates, therefore, compare creators occupying similar "regions" of the observed pre-treatment distribution, though matching balances what we observe rather than what we do not.[29]

7.2 Main results

7.2.1 Alternative channel investment: content produced

We begin with Hypothesis 1, which predicts that exposed complementors expand directly monetized activity on the auxiliary channel and, more specifically, do so without a decrease in freely accessible output. Table 4 presents the main evidence. Across all three specifications, YouTube-linked creators increase paid Patreon posts after the rollout: 0.191 posts per creator–month in the full sample, 0.165 in the CEM common-support sample, and 0.153 in the weighted CEM specification, each statistically significant at the 5% level. Against the mean of 1.558 paid posts in the matched estimation sample, the weighted estimate amounts to an increase of approximately 10%. The free content estimates constitute the other half of the test, and they align with our theoretical predictions: they are small and statistically insignificant across all specifications. The weighted estimate is 0.057 posts per creator–month against a matched-sample mean of 1.006: the post-shock response is concentrated in content whose accessibility depends directly on membership payments, with no comparable change in publicly accessible posts.

The additional results presented in Tables A.4 and A.5 distinguish alternative channel expansion from within-Patreon reallocation. Total Patreon content rises by 0.209 posts per creator–month in the weighted CEM specification, roughly 8% of the matched-sample mean; the paid-minus-free difference is positive but less precisely estimated in the weighted specification and is not significant. The increase in paid content, the increase in total content, and the absence of a statistically significant change in free content are more consistent with alternative channel expansion than with moving a fixed stock of posts behind a paywall. This

stratum.

[29] Propensity score matching (PSM) yields less convincing balance. Still, we discuss the results obtained using this approach in Section 7.3 and Appendix Section A.4.2.

Table 4. Alternative channel investment: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.191**	0.165**	0.153**	-0.031	0.047	0.057
	(0.083)	(0.076)	(0.076)	(0.047)	(0.043)	(0.048)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	108,025	90,396	90,396
Mean DV	2.315	1.558	1.558	1.277	1.006	1.006

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

also aligns with our theoretical predictions: Exposed ("YouTube-linked") creators increase their directly monetized activity on Patreon after the governance shock.

Figure 3 presents the event studies for the four content outcomes. The pre-rollout coefficients for total, paid-minus-free, paid, and free content show no systematic divergence relative to November 2017. Once the rollout begins, coefficients for paid content turn positive and continue to grow through the spring of 2018, while coefficients for free content remain comparatively flat.[30] Overall, the magnitude, composition, and timing of the estimated effects provide consistent support for Hypothesis 1.

7.2.2 Monetization risk exposure

Hypothesis 2 emphasizes the moderating role of monetization risk exposure. The underlying mechanism is displacement pressure: because the governance shock increases the expected cost of continued reliance on YouTube, the response should be strongest among creators whose content lies closest to the advertiser sensitive boundary affected by the policy change.

Table 5 compares creators whose content falls in the advertiser sensitive topic category with those in clearly advertiser safe categories, excluding the intermediate and unclassified groups. The estimates indicate a strong moderating role for monetization risk exposure. In the weighted CEM specification, the interaction between treatment and sensitive topic status is 0.693 for paid content, significant at the 1% level, whereas the corresponding interaction for free content is smaller and statistically insignificant. This moderating

[30] Figures A.3-A.6 reproduce the four event studies separately, at a larger scale.

Figure 3. Event studies for content outcomes

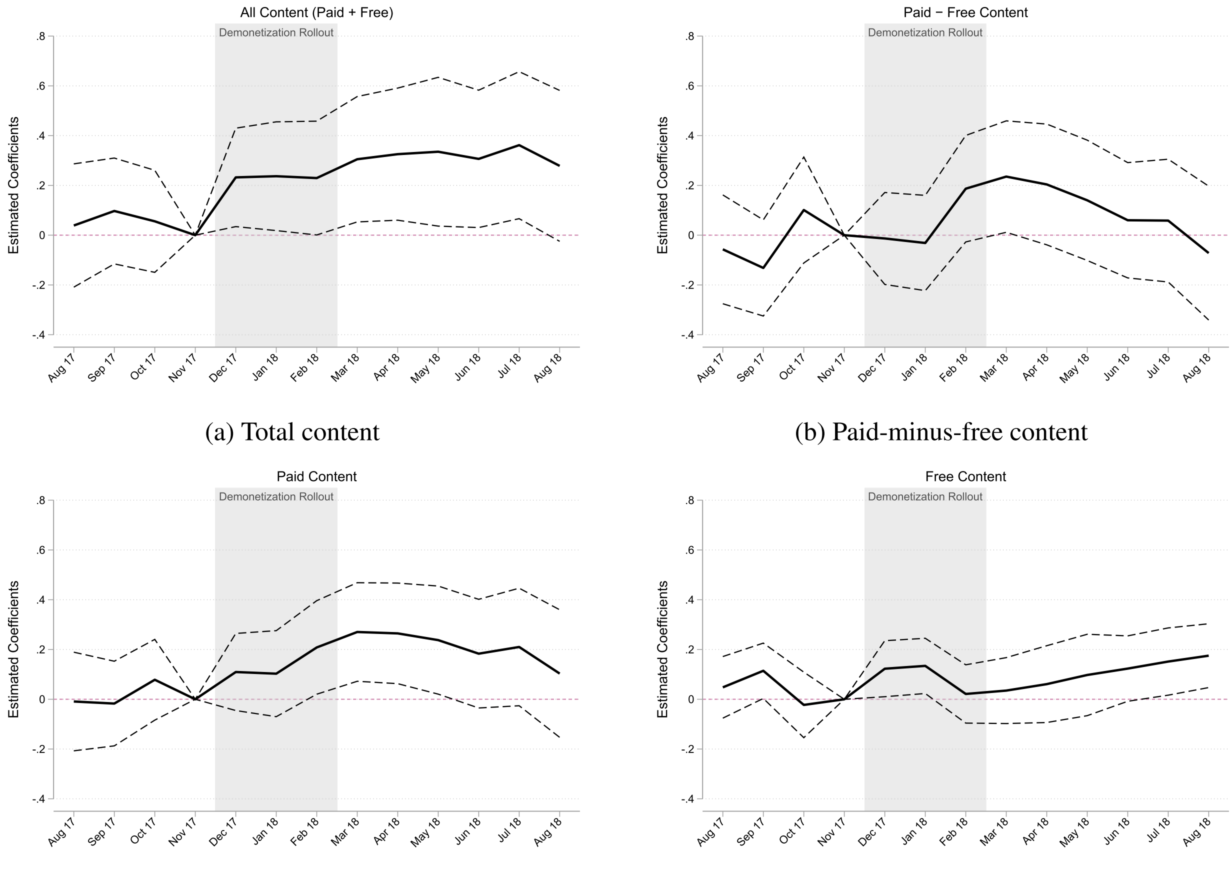

(a) Total content

(b) Paid-minus-free content

(c) Paid content

(d) Free content

Notes: Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month. Dashed lines report 95% confidence intervals based on standard errors clustered at the creator level; the shaded area denotes the December 2017–February 2018 rollout period.

effect is observed for total production as well: Table A.6 shows a weighted total-content interaction of 0.864 ($p < 0.01$), with the paid-minus-free interaction positive and significant at the 5% level. In the weighted paid content specification, the estimated effect for creators in advertiser safe categories is 0.106 posts per month and is not statistically significant; sensitive topic status is associated with an additional effect of 0.693 posts per month.[31]

Figure A.7 presents the event study estimates comparing sensitive and safe-topic creators in the unweighted CEM common-support sample. The post-rollout response appears larger and more persistent among sensitive topic creators. We nevertheless treat the figure as supplementary because the group-specific event studies do not directly estimate the dynamic counterpart of the triple-difference interaction, and the

[31] We assess the robustness of these results to the composition of the reference group. Tables A.7 and A.8 report estimates obtained by pooling intermediate- and low-sensitivity creators into the comparison group. The CEM-weighted interaction coefficients remain large (0.871 for total content and 0.682 for paid content) and are again significant at the 1% level.

Table 5. Monetization-risk exposure: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.145	0.155	0.106	-0.104	-0.008	0.013
	(0.152)	(0.142)	(0.135)	(0.076)	(0.069)	(0.083)
Treated × High Monetization Risk	0.434*	0.325	0.693***	0.244*	0.289**	0.171
	(0.244)	(0.213)	(0.240)	(0.140)	(0.136)	(0.133)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	63,375	53,508	53,508	63,375	53,508	53,508
Mean DV	2.537	1.772	1.772	1.374	1.101	1.101

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Monetization Risk is the differential effect for the high-moderator group. Post × High Monetization Risk is included as a control; its coefficient is not shown. High Monetization Risk compares creators classified High against creators classified Low; Medium and unclassified creators are excluded. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

sensitive topic series is less precisely estimated than the main event study results. The interaction estimates reported in Table 5 therefore provide the main test of Hypothesis 2. They show that the governance change did not affect all YouTube-linked creators equally: the increase in paid content is substantially larger among creators whose content is closest to the advertiser sensitivity and monetization boundaries affected by the policy change. We also explore an alternative categorization of monetization risk and discuss the results in Appendix Section A.3.4 (see, in particular, Tables A.9 and A.10).

Overall, the empirical evidence supports the theoretical prediction outlined in Hypothesis 2. Creators in advertiser sensitive categories exhibit a substantially larger paid content response while their free content interaction remains statistically insignificant, consistent with creators under greater displacement pressure expanding the activity least dependent on advertiser eligibility.

7.2.3 Community conversion capacity

Hypothesis 3 emphasizes conversion capacity as a moderator. In Section 3, we argued that creators with more engaged pre-shock communities should be better positioned to convert governance pressure into direct monetization. Our empirical results provide partial support for this prediction: the effect is directionally consistent, even though less stable than for monetization risk. Comparing creators in the top and bottom terciles of pre-treatment comment volume (Tables A.11 and A.12), the high-engagement interaction for paid content is 0.562 in the full sample and 0.528 in the unweighted CEM common-support sample. Both estimates are significant at the 5% level, and the corresponding paid-minus-free interactions are likewise

Table 6. Monetization-risk exposure and community conversion capacity

	Total Content (Paid + Free)			Paid Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.068	0.106	0.066	-0.041	-0.018	-0.047
	(0.214)	(0.208)	(0.211)	(0.160)	(0.150)	(0.148)
Treated × High Monetization Risk	0.703*	0.513	0.862**	0.562*	0.373	0.752***
	(0.377)	(0.345)	(0.340)	(0.332)	(0.284)	(0.287)
Treated × High Engagement	0.415	0.710	0.501	0.675*	0.728*	0.549*
	(0.439)	(0.435)	(0.392)	(0.380)	(0.376)	(0.334)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	42,120	33,462	33,462	42,120	33,462	33,462
Mean DV	4.470	3.083	3.083	3.058	1.976	1.976

Notes: Difference-in-differences estimates with both moderator interactions entered jointly. High Monetization Risk (High vs Low classification) and High Engagement (top vs bottom tercile of pre-period comments) are defined as in the single-moderator tables; Post × moderator controls for both moderators are included (coefficients not shown). Sample: classified High/Low creators in the top or bottom engagement tercile; otherwise treatment, sample, and CEM matching as in the baseline tables. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

positive and statistically significant. Applying the CEM weights changes the picture, though: the paid content interaction falls to 0.224 and loses statistical significance, and the total content interaction, while positive, is imprecisely estimated. The median-split results point in the same direction (Table A.13). The high-engagement interaction is 0.433 in the full sample and 0.382 in the matched sample, both significant at the 5% level, but falls to an insignificant 0.216 in the CEM-weighted specification.

Table 6 enters monetization risk and prior engagement simultaneously. In the weighted paid content model, the monetization risk interaction remains large and significant at the 1% level; the engagement interaction is positive and significant at the 10% level. Also, we show that the interactions are not a byproduct of the baseline effect vanishing in restricted samples: Table A.14 shows that the average paid content effect remains positive when re-estimated in each heterogeneity subsample.[32]

7.2.4 Community-funded value capture: patrons and earnings

We next examine whether the content response documented above translates into greater value capture. Because the number of patrons and displayed earnings are stocks that accumulate over time, a persistent

[32]Two limitations are worth noting. Monetization risk and pre-treatment engagement are creator characteristics, not experimentally assigned variation, so the results presented above simply highlight that each moderator retains predictive content within the joint sample without establishing statistical complementarity or a “causal ordering” between displacement pressure and conversion capacity. Also, because the preferred weighted single moderator specification does not yield a significant interaction, we read the results as partial support for Hypothesis 3: creators with more engaged communities tend toward stronger paid content expansion, and the engagement interaction survives the inclusion of monetization risk, but the evidence remains suggestive rather than definitive.

Table 7. Net monthly change in patrons and earnings

	Net Change in Patrons			Net Change in Earnings		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.857**	0.705**	0.526*	2.184	3.333**	3.690***
	(0.428)	(0.310)	(0.269)	(1.594)	(1.398)	(1.392)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	107,138	89,620	89,620	84,888	73,867	73,867
Mean DV	2.999	1.705	1.705	8.021	5.827	5.827

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Outcomes are one-month first differences (net monthly change), defined only across consecutive observed months; the coefficient is the additional net monthly gain for treated creators. Matching: coarsened exact matching (CEM) on pre-period covariates as in the baseline. All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

pre-existing difference in a creator's growth rate would mechanically increase the level gap at any cutoff we might choose. We therefore use the net monthly changes, i.e., the first differences. The dependent variables then measure the rate at which creators accumulate patrons and earnings, thereby isolating a change in that rate at the rollout from any pre-existing divergence in levels. Table 7 reports estimates from these models.

We observe that, after the rollout, YouTube-linked creators obtain patrons at a faster monthly rate: 0.52 to 0.86 additional net patrons per creator–month, significant across all three specifications. The earnings flow is 3.3 to 3.7 additional dollars per creator–month in the two CEM specifications, significant at the 5% and 1% levels; the full-sample earnings estimate is positive but imprecise, so we read the earnings flow from the common-support specifications where treated and comparison creators are matched. Relative to matched-sample mean monthly increases of approximately 1.7 patrons and $5.8 in displayed earnings, these effects represent roughly one-third to one-half of baseline accumulation rates.

Figure 4 provides the clearest evidence for these outcomes. First differences remain near zero before the rollout, increase as it begins, and remain high, indicating a discrete rise in accumulation rates rather than a gradually widening level gap.

Additional analyses using transformed outcomes provide further support for the flow-based results. The inverse-hyperbolic-sine specifications in Table A.15 reduce the influence of the highly skewed stock distributions and imply differential increases of approximately 9% in patrons and displayed earnings. In levels (see Table A.16), the weighted CEM estimates correspond to approximately 3.5 additional patrons and $14.1

Figure 4. Event studies for the net monthly change in patrons and earnings

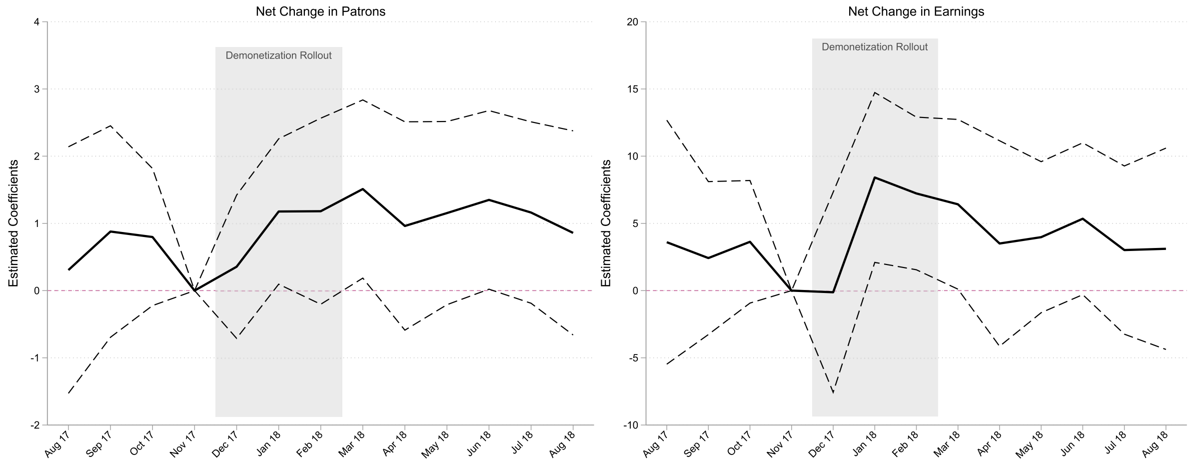

Notes. Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month. Dashed lines report 95% confidence intervals based on standard errors clustered at the creator level; the shaded area denotes the December 2017–February 2018 rollout period.

in additional displayed monthly earnings; we report levels only to measure economic magnitude, since the level of a stock outcome mechanically embeds any pre-existing growth difference and is correspondingly less stable across samples.

Overall, these analyses suggest that complementors' increased activity on the auxiliary platform is associated with greater value creation, as reflected in higher patron counts and recurring revenue. As mentioned earlier, faster patron growth may reflect greater willingness to support creators following the governance shock, a response to the additional paid content offered, or a mutually reinforcing process between creator investment and audience support. The event-study estimates cannot distinguish among these mechanisms, and earnings are observed only for creators who choose to disclose them publicly.

7.3 Robustness of the findings

The results are, overall, robust to the additional checks reported in detail in Appendix A.4. These include PPML specifications suited to skewed count outcomes (see Tables A.17-A.19); propensity score matching and more restrictive CEM procedures (see Tables A.20-A.22); alternative post-treatment dates in December 2017 and February 2018 (see Tables A.23 and A.24); samples that reintroduce entrants and exiters (see Table A.25), impose a strictly balanced panel (Table A.26), or exclude the top percentile of pre-shock patrons and earnings (Table A.27); and placebo cutoffs both immediately before the rollout and during an earlier period with no comparable governance change (Tables A.28-A.30).

Across these analyses, the most consistent finding concerns content production on Patreon. Exposed creators increase paid content, while free content remains statistically unchanged. This result is robust to alternative rollout dates and sample definitions and is substantially stronger among creators facing greater monetization risk. Other findings are less robust across specifications. The estimated effect of conversion capacity is positive but sensitive to the use of CEM weights. Patron and earnings effects are significant in the first-difference and inverse-hyperbolic-sine specifications. The patron result is less robust under PPML, whereas earnings remain positive and statistically significant in levels and PPML specifications.

8 Conclusion

This paper examines how complementors respond when governance on a dominant platform reduces the attractiveness of its monetization regime while leaving its discovery infrastructure largely intact. We study YouTube's revised moderation and monetization policies following the 2017 Adpocalypse and focus on creators' responses on Patreon, a separately governed membership platform. This setting allows us to examine a strategic response that lies between continued dependence and platform exit: creators can remain reliant on YouTube for visibility and audience formation while shifting part of value capture toward Patreon.

The most robust finding relates to creators' activity on the auxiliary platform. Following the policy rollout, Patreon creators with a pre-shock YouTube connection increased their production of member-only content without a corresponding decline in freely accessible posts. Thus, the response appears to represent an expansion of directly monetized activity, rather than the "reuse" of existing output behind a paywall. The response is also substantially stronger among creators whose content falls within advertiser-sensitive categories. This exposure moderator supports the proposed displacement pressure mechanism: complementors facing greater deterioration in the expected attractiveness of advertising-based monetization have stronger incentives to invest in a channel less dependent on advertiser eligibility. Pre-shock community engagement is positively associated with the paid content response, consistent with conversion capacity, although this result is more sensitive to weighting and should be interpreted as suggestive. Finally, our findings indicate that exposed creators experienced faster growth in patrons and recurring earnings.

Our findings suggest that platform governance can redistribute value capture across platforms. The observed Patreon response is consistent with complementors reducing their dependence on a dominant platform's monetization rules without necessarily leaving its discovery ecosystem. Understanding platform power, therefore, requires examining not only whether complementors remain on or exit a platform, but also how value capture is reallocated between focal and auxiliary platforms.

Theoretical contributions. The findings contribute to research on platform governance by showing that governance interventions can reshape complementor behavior beyond the boundaries of the platform implementing them. Existing work has primarily examined how complementors adjust their participation, output, or positioning within the focal platform. We identify a cross-platform response: a change in governance on a dominant discovery platform redirects complementor activity toward a platform performing a different function. The relevant strategic consequence of governance is therefore not limited to exit, reduced participation, or changes in on-platform production: governance can also affect *where* complementors invest and capture value in interconnected ecosystems.

We characterize this response as *monetization layer bypassing*. Unlike conventional multihoming, it does not involve allocating activity across platforms offering comparable discovery or distribution functions. Also, unlike conventional platform leakage, it does not involve directly or privately completing the same transaction that the focal platform would otherwise intermediate. Instead, the mechanism involves greater reliance on a separately governed monetization channel without requiring leaving the dominant platform's audience-making infrastructure. This distinction highlights that platform functions commonly "bundled" within a single ecosystem (discovery, distribution, and monetization) may remain strategically separable.

The paper also provides evidence that platform auxiliaries can become more important in response to changes in focal-platform governance. An auxiliary does not need to reproduce the dominant platform's full technological or audience infrastructure to reshape its value architecture. By supplying a narrower capability (in this case, recurring direct payments and access-controlled content), it can allow complementors to reduce dependence on a particular layer of the focal platform while remaining embedded in the broader ecosystem. The strategic relevance of an auxiliary is therefore governance-contingent: its value increases when the terms governing the corresponding function on the focal platform deteriorate.

Managerial implications. For platform owners, the findings suggest that continued complementor participation should not necessarily be interpreted as evidence that a governance intervention has imposed few strategic costs. Complementors may remain active because the platform's discovery infrastructure is difficult to replace, while simultaneously shifting part of the value capture to auxiliary channels. A dominant position in audience discovery does not, therefore, guarantee equivalent control over monetization. Platforms should evaluate governance changes not only by their immediate effects on advertiser trust, policy compliance, and on-platform activity, but also by their consequences for complementors' reliance on adjacent layers of the ecosystem.

This is particularly important when governance rules are broad, difficult to anticipate, or uneven in their consequences across complementors. Greater transparency concerning monetization eligibility, clearer enforcement and appeal procedures, and credible alternative monetization options *within* the platform may reduce complementors' incentives to shift value capture elsewhere. The results also suggest that governance changes can redistribute activity unevenly: complementors whose production lies closest to the boundaries redrawn by the policy are especially likely to seek alternative channels.

For complementors, our findings illustrate the strategic value of separating audience formation from value capture. Developing direct relationships with users and maintaining access to auxiliary monetization channels can reduce exposure to a single platform controlled revenue model. Such diversification does not necessarily eliminate platform dependence. In our setting, YouTube remained an important potential source of discovery. Creators' ability to benefit from Patreon depends on mobilizing audiences through a direct payment relationship. Auxiliary channels, therefore, provide partial autonomy rather than complete independence.

Limitations and future research. Our study has several limitations. First, our treatment captures exposure to a system-wide governance change through an observable pre-shock connection to YouTube. This approach is consistent with our theoretical focus on how complementors respond to a change in the platform's governance environment, rather than on the consequences of a documented advertising revenue loss for a particular creator. The estimates should therefore be interpreted as *intent-to-treat effects* of exposure to the governance change, encompassing responses to both realized and anticipated monetization constraints. Our moderators are similarly based on information observable on creators' Patreon pages. Monetization risk captures the positioning of creators' content rather than a direct classification of their individual YouTube videos, while pre-shock comments indicate, rather than directly measure, their capacity to convert audience engagement into community support.

Second, our data capture the auxiliary side of the cross-platform response. This focus aligns with our theory, which predicts that complementors will increase their use of a separately governed monetization channel. It allows us to establish that exposed creators expanded paid activity and value capture on Patreon without reducing their free output on Patreon, though we do not observe a change in the intensity of their production on YouTube. Within Patreon, increases in patrons and earnings represent realized community-funded value capture but may reflect both greater creator investment in paid content and stronger audience willingness to provide support. Our design identifies their combined effect rather than separately decomposing these creator- and audience-side mechanisms. Finally, earnings are observed only for creators who publicly disclose them, and post data include only pages and historical content accessible when collected.

Future research could combine activity and revenue data across focal and auxiliary platforms to capture the full reallocation of complementor effort and value capture. Also, further work could examine how focal platforms respond as auxiliary channels become economically important, whether such channels create new forms of dependence and governance risk, and how cross-platform reallocation affects content quality, advertiser participation, consumer welfare, and the long-run distribution of ecosystem value.

References

Belleflamme, Paul, Thomas Lambert, and Armin Schwienbacher (2014) 'Crowdfunding: Tapping the right crowd.' *Journal of Business Venturing* 29(5), 585–609

Benatia, David, Christophe Bellégo, and Louis Pape (2025) 'Dealing with logs and zeros in regression models.' *Working Paper*

Boudreau, Kevin J. (2012) 'Let A Thousand Flowers Bloom? An Early Look at Large Numbers of Software App Developers and Patterns of Innovation.' *Organization Science* 23(5), 1409–1427

Boudreau, Kevin J., and Lars B. Jeppesen (2015) 'Unpaid Crowd Complementors: The Platform Network Effect Mirage.' *Strategic Management Journal* 36(12), 1761–1777

Cennamo, Carmelo, and Juan Santalo (2013) 'Platform Competition: Strategic Trade-Offs in Platform Markets.' *Strategic Management Journal* 34(11), 1331–1350

Cohn, Jonathan B., Zack Liu, and Malcolm I. Wardlaw (2022) 'Count (and count-like) data in finance.' *Journal of Financial Economics* 146(2), 529–551

Cong, Ziwei, Ying Zhao, and Zilei Zhang (2025) 'Understanding users' content contribution behavior when knowledge can be priced.' *Marketing Science* 44(3), 655–670

Cutolo, Donato, and Martin Kenney (2021) 'Platform-dependent entrepreneurs: Power asymmetries, risks, and strategies in the platform economy.' *Academy of Management Perspectives* 35(4), 584–605

Cutolo, Donato, Andrew Hargadon, and Martin Kenney (2026) 'Hidden in plain sight: Conceptualizing platform auxiliaries and their conflicting impacts on platforms' value architecture.' *Journal of Management Studies*

Eisenmann, Thomas, Geoffrey Parker, and Marshall Van Alstyne (2011) 'Platform Envelopment.' *Strategic Management Journal* 32(12), 1270–1285

Elsas-Nicolle, Ambre, and Johannes Loh (2025) 'Competition-as-a-service: How subscription bundles reshape platform competition and seller performance.' *Working Paper*

Evans, David S. (2003) 'Some empirical aspects of multi-sided platform industries.' *Review of Network Economics* 2(3), 191–209

Faraj, Samer, and Steven L. Johnson (2011) 'Network exchange patterns in online communities.' *Organization Science* 22(6), 1464–1480

Gambardella, Alfonso, and Claudio Panico (2014) 'On the management of open innovation.' *Research Policy* 43(5), 903–913

Gu, Grace Y. (2024) 'Technology and disintermediation in online marketplaces.' *Management Science* 70(11), 7868–7891

Hagiu, Andrei, and Julian Wright (2024) 'Marketplace leakage.' *Management Science* 70(3), 1529–1553

Jacobides, Michael G., Carmelo Cennamo, and Annabelle Gawer (2018) 'Towards a theory of ecosystems.' *Strategic Management Journal* 39(8), 2255–2276

Kerkhof, Anna (2024) ‘Advertising and content differentiation: Evidence from YouTube.’ *The Economic Journal* 134(663), 2912–2950

Kesler, Reinhold (2022) ‘The impact of Apple’s App Tracking Transparency on app monetization.’ *Working Paper*

Kretschmer, Tobias, Aija Leiponen, Melissa Schilling, and Gurneeta Vasudeva (2022) ‘Platform ecosystems as meta-organizations: Implications for platform strategies.’ *Strategic Management Journal* 43(3), 405–424

Lin, Yu-Kai, Arun Rai, and Yukun Yang (2022) ‘Information control for creator brand management in subscription-based crowdfunding.’ *Information Systems Research* 33(3), 846–866

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

Loh, Johannes, and Ambre Elsas-Nicolle (2026) ‘Platform competition and strategic trade-offs for complementors: Heterogeneous reactions to the entry of a new platform.’ *Strategic Management Journal* 47(6), 1671–1729

Loh, Johannes, Lena Abou El-Komboz, and Anna Kerkhof (2026) ‘Value capture restrictions and complementary innovation in platform ecosystems.’ *Working Paper*

Mollick, Ethan (2014) ‘The Dynamics of Crowdfunding: An Exploratory Study.’ *Journal of Business Venturing* 29(1), 1–16

Mostrous, Alexi, and Jonathan Dean (2017) ‘Top brands pull Google adverts in protest at hate video links.’ *The Times*

Nicas, Jack (2017) ‘Google’s YouTube has continued showing brands’ ads with racist and other objectionable videos.’ *The Wall Street Journal*

Panico, Claudio (2017) ‘Strategic interaction in alliances.’ *Strategic Management Journal* 38(8), 1646–1667

Panico, Claudio, and Carmelo Cennamo (2022) ‘User preferences and strategic interactions in platform ecosystems.’ *Strategic Management Journal* 43(3), 507–529

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

Rasskazova, Maria (2026) ‘Content demonetisation and creators’ participation: Study of YouTube’s first Adpocalypse.’ *Working Paper*

Rietveld, Joost, Joe N. Ploog, and David B. Nieborg (2020) ‘Coevolution of Platform Dominance and Governance Strategies: Effects on Complementor Performance Outcomes.’ *Academy of Management Discoveries* 6(3), 488–513

Rochet, Jean-Charles, and Jean Tirole (2003) ‘Platform Competition in Two-Sided Markets.’ *Journal of the European Economic Association* 1(4), 990–1029

Tiwana, Amrit (2013) *Platform Ecosystems: Aligning Architecture, Governance, and Strategy* (Morgan Kaufmann)

Wang, Chengsi, and Julian Wright (2020) ‘Search platforms: Showrooming and price parity clauses.’ *The RAND Journal of Economics* 51(1), 32–58

Wang, Jing, Gen Li, and Kai-Lung Hui (2022) ‘Monetary incentives and knowledge spillover: Evidence from a natural experiment.’ *Management Science* 68(5), 3549–3572

Wasko, Molly McLure, and Samer Faraj (2005) ‘Why Should I Share? Examining Social Capital and Knowledge Contribution in Electronic Networks of Practice.’ *MIS Quarterly* 29, 35–57

Wen, Wen, and Feng Zhu (2019) ‘Threat of platform-owner entry and complementor responses: Evidence from the mobile app market.’ *Strategic Management Journal* 40(9), 1336–1367

Appendix A

A.1 Institutional Background

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

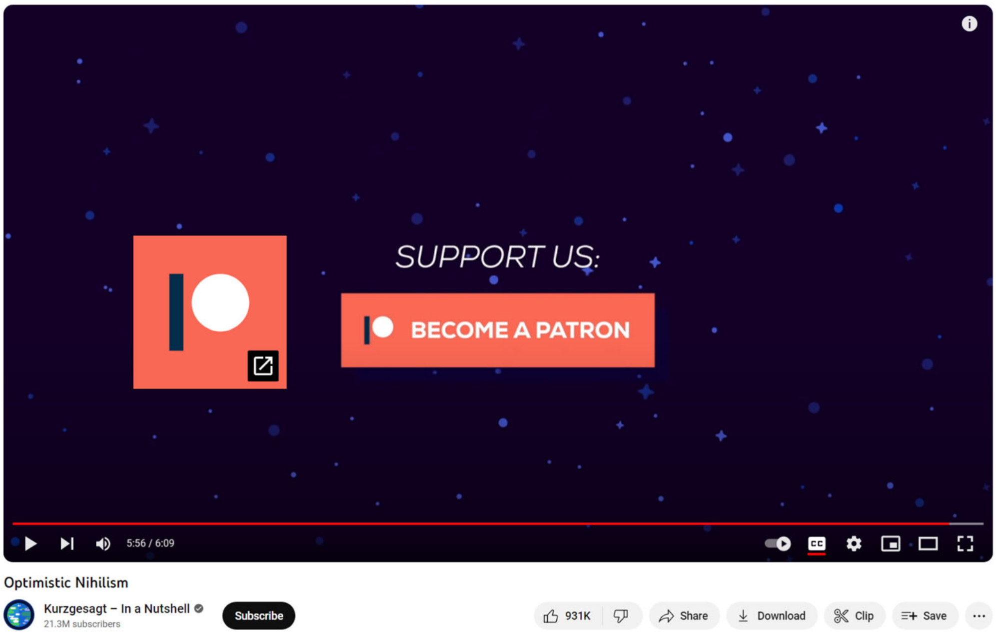

(a) Reference to Patreon during a Video

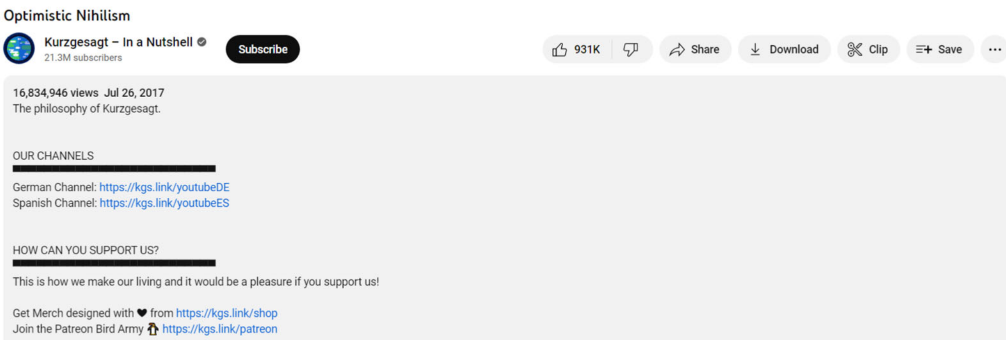

(b) Reference to Patreon in the Video Description

A.2 Additional descriptive statistics

Figure A.2. Entry and exit of Video creators on Patreon

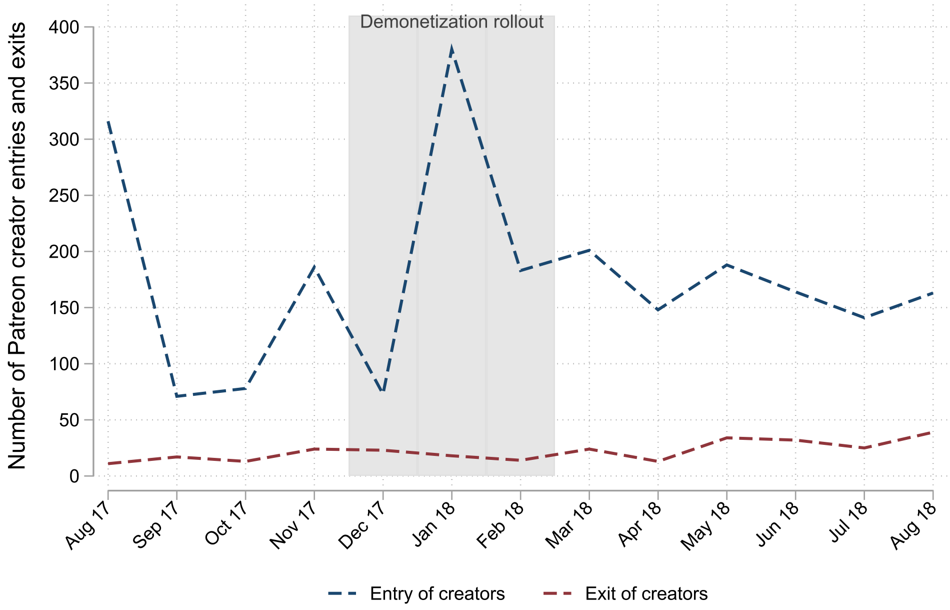

Notes: The lines report monthly entry and exit events among creators assigned to Patreon's Video category. Entry and exit are based on each creator's first and last observed month in the underlying panel. The shaded region denotes the December 2017–February 2018 demonetization rollout period.

Table A.1. Pre-shock YouTube links across samples

Sample	Creators	Pre-shock YouTube links	Share (%)
Full scraped universe	11,371	8,110	71.32
Video category, excluding adult content	10,248	7,470	72.89
Main stayer sample	8,436	7,317	86.74
CEM common-support sample	7,065	6,013	85.11

Notes: A pre-shock YouTube link equals one when YouTube subscriber information is observed strictly before December 2017, i.e., before the rollout window begins. This descriptive definition is intentionally narrower than the treatment used in estimation, which incorporates information through December 2017 (strictly before the January 2018 post period).

Table A.2. Monetization risk classifications

Monetization-risk group	Main sample	YouTube-linked	Comparison	Main-sample share (%)
High	1,173	1,019	154	13.90
Medium	815	697	118	9.66
Low	3,702	3,208	494	43.88
Unclear/unclassified	2,746	2,412	334	32.55

Notes: Creator counts. The full universe includes all Patreon creators observed from August 2017 through August 2018. The main sample contains non-adult Video creators observed in the window before December 2017 and after February 2018. YouTube-linked status follows the treatment definition used in estimation. Creators without an entry in the thematic classification files appear as Unclear/unclassified.

Table A.3. Covariate balance, main CEM specification

	Unmatched			Matched (weighted)			
	Treated	Control	Std. diff. (%)	Treated	Control	Std. diff. (%)	$p-value$
Contents per month	3.72	3.26	8.3	2.62	2.68	-1.1	0.651
Comments per month	14.25	9.92	9.0	4.49	4.20	0.6	0.577
Likes per month	25.28	21.01	4.3	8.49	9.08	-0.6	0.699
Free contents per month	1.45	1.12	11.6	1.11	0.98	4.7	0.087
Patrons	70.95	35.59	14.7	32.71	30.48	0.9	0.475
Earnings (USD)	322.88	185.50	14.4	172.71	183.73	-1.2	0.595
Facebook likes	28633.90	9181.86	11.7	25350.86	7796.83	10.6	0.000
Twitter followers	10010.71	5114.08	8.2	6414.30	3975.10	4.1	0.002
Monetization risk (1–4)	2.96	2.92	4.0	3.01	3.01	0.0	1.000
Creators	7,336	1,100		6,030	1,035		

Notes: Covariates are creator-level averages over the window pre-period (August–November 2017), computed on the main stayer sample. Coarsened exact matching coarsens contents, comments, patrons, and earnings into bins and matches on the resulting strata; monetization risk enters as exact categories rather than as a coarsened numeric variable, so the four groups (High, Medium, Low, Unclear/unclassified) are never pooled with one another and balance on this covariate is exact by construction. Its mean is reported on the 1–4 coding for completeness only. Standardized differences are in percent and use the pooled, unmatched standard deviation in both panels. Matched-sample p-values come from weighted regressions of each covariate on the treatment indicator with heteroskedasticity-robust standard errors. Matched columns use CEM weights. Facebook and Twitter reach are excluded from the coarsening and assessed ex-post.

A.3 Additional estimation results

A.3.1 Main specification

Table A.4. Total content and paid-minus-free composition

	All Content (Paid + Free)			Paid - Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.161	0.212**	0.209**	0.222**	0.118	0.096
	(0.102)	(0.095)	(0.097)	(0.089)	(0.079)	(0.083)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	108,025	90,396	90,396
Mean DV	3.593	2.564	2.564	1.038	0.552	0.552

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.5. Paid share of content: expansion versus substitution

	Paid Share of Content			Paid Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.007	0.003	0.005	0.191**	0.165**	0.153**
	(0.008)	(0.009)	(0.014)	(0.083)	(0.076)	(0.076)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	56,866	42,835	42,835	108,025	90,396	90,396
Mean DV	0.608	0.583	0.583	2.315	1.558	1.558

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Paid share is paid content divided by total content, defined only for creator-months with at least one post; treatment can affect this conditioning. Matching: CEM as in the baseline. All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.3.2 Event studies

Figure A.3. Event study: total content

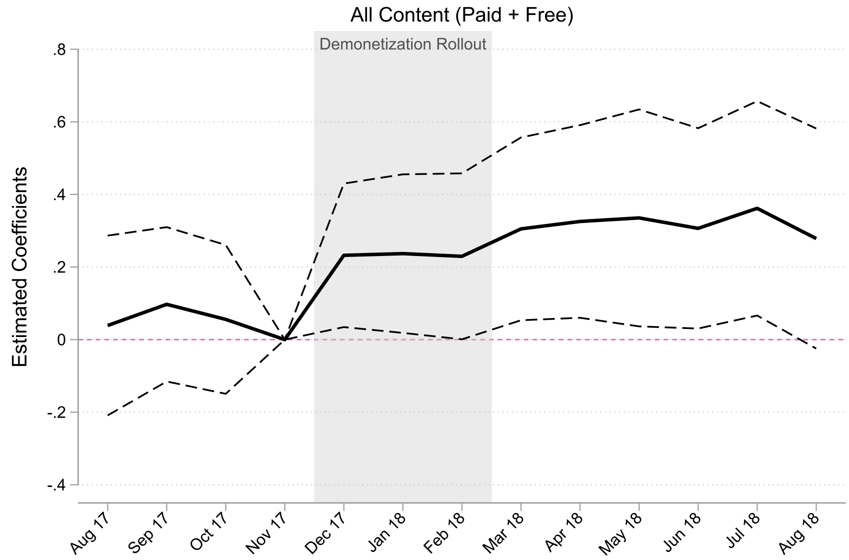

Notes: Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month; dashed lines report 95% confidence intervals based on creator-clustered standard errors, and the shaded area denotes the December 2017–February 2018 rollout period.

Figure A.4. Event study: paid-minus-free content

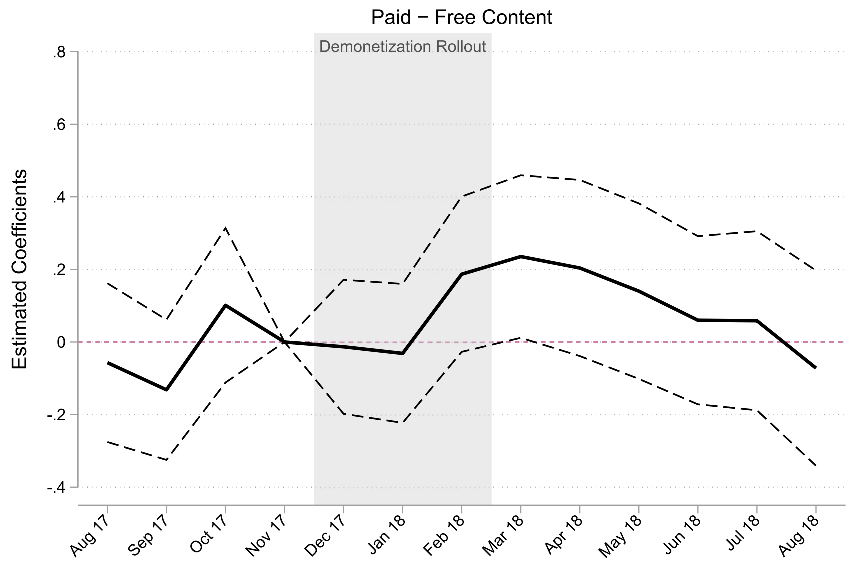

Notes: Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month; dashed lines report 95% confidence intervals based on creator-clustered standard errors, and the shaded area denotes the December 2017–February 2018 rollout period.

Figure A.5. Event study: paid content

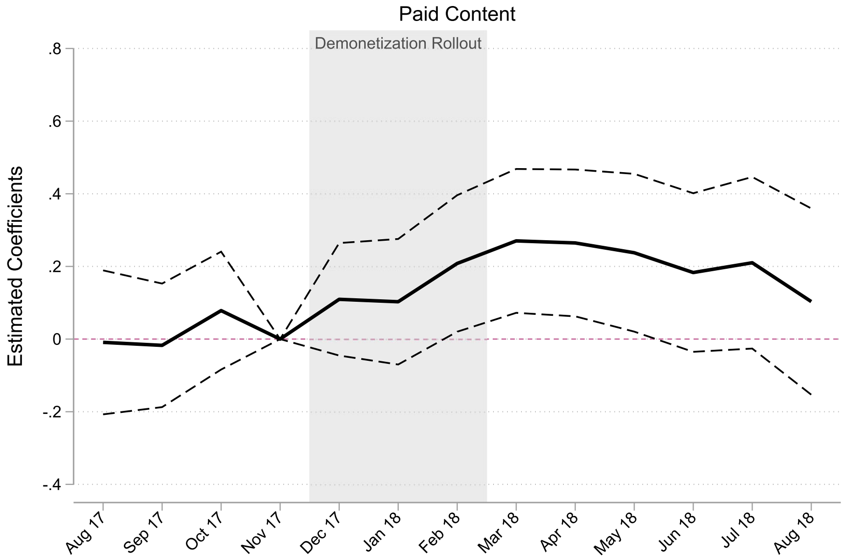

Notes: Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month; dashed lines report 95% confidence intervals based on creator-clustered standard errors, and the shaded area denotes the December 2017–February 2018 rollout period.

Figure A.6. Event study: free content

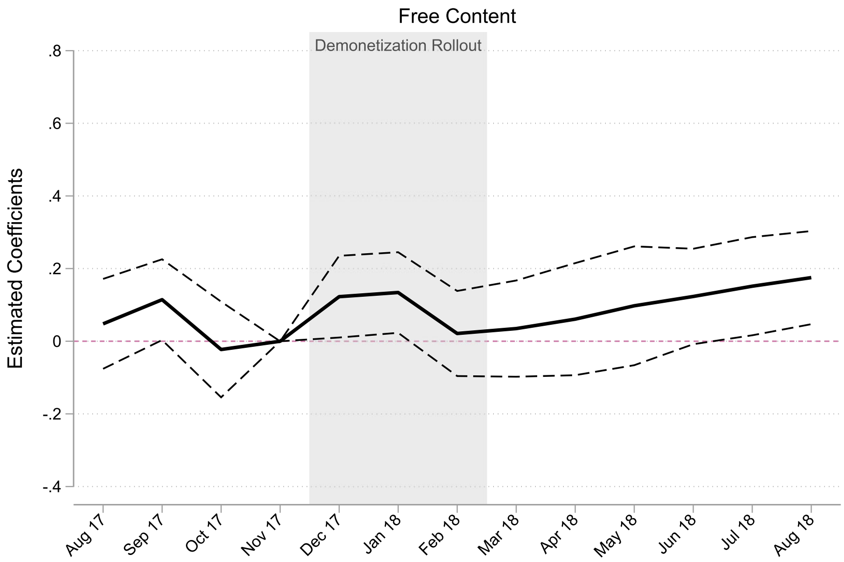

Notes: Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month; dashed lines report 95% confidence intervals based on creator-clustered standard errors, and the shaded area denotes the December 2017–February 2018 rollout period.

A.3.3 Monetization risk

Table A.6. Monetization risk (High versus Low): total content and composition

	Total Content (Paid + Free)			Paid - Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.041	0.147	0.119	0.250	0.163	0.093
	(0.184)	(0.173)	(0.171)	(0.155)	(0.141)	(0.144)
Treated × High Monetization Risk	0.678**	0.614**	0.864***	0.189	0.036	0.521**
	(0.288)	(0.266)	(0.282)	(0.274)	(0.237)	(0.266)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	63,375	53,508	53,508	63,375	53,508	53,508
Mean DV	3.911	2.873	2.873	1.163	0.671	0.671

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Monetization Risk is the differential effect for the high-moderator group. Post × High Monetization Risk is included as a control; its coefficient is not shown. High Monetization Risk compares creators classified High against creators classified Low; Medium and unclassified creators are excluded. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.7. Monetization risk (High versus Medium/Low): total content and composition

	Total Content (Paid + Free)			Paid - Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.053	0.117	0.111	0.224*	0.170	0.122
	(0.157)	(0.148)	(0.149)	(0.132)	(0.120)	(0.126)
Treated × High Monetization Risk	0.666**	0.644**	0.871***	0.215	0.029	0.492*
	(0.271)	(0.251)	(0.270)	(0.262)	(0.225)	(0.256)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	73,970	61,230	61,230	73,970	61,230	61,230
Mean DV	3.857	2.808	2.808	1.156	0.644	0.644

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Monetization Risk is the differential effect for the high-moderator group. Post × High Monetization Risk is included as a control; its coefficient is not shown. High Monetization Risk equals one for creators classified High; Medium and Low form the comparison group; unclassified creators are excluded. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.8. Monetization risk (High versus Medium/Low): paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.139	0.143	0.117	-0.085	-0.027	-0.006
	(0.128)	(0.119)	(0.117)	(0.068)	(0.062)	(0.073)
Treated × High Monetization Risk	0.441*	0.336*	0.682***	0.226*	0.307**	0.189
	(0.229)	(0.198)	(0.230)	(0.136)	(0.132)	(0.128)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	73,970	61,230	61,230	73,970	61,230	61,230
Mean DV	2.506	1.726	1.726	1.351	1.082	1.082

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Monetization Risk is the differential effect for the high-moderator group. Post × High Monetization Risk is included as a control; its coefficient is not shown. High Monetization Risk equals one for creators classified High; Medium and Low form the comparison group; unclassified creators are excluded. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Figure A.7. Event studies for paid content by monetization risk group

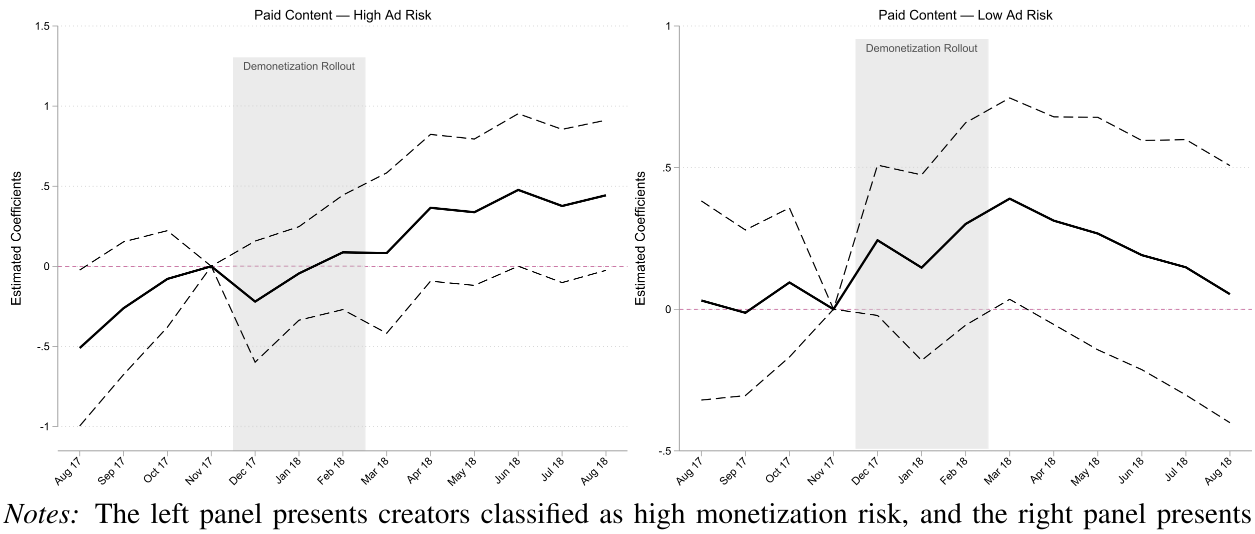

Notes: The left panel presents creators classified as high monetization risk, and the right panel presents creators classified as low monetization risk. Estimates use the unweighted CEM common-support sample. November 2017 is the omitted month; dashed lines report 95% confidence intervals based on creator-clustered standard errors, and the shaded area denotes the December 2017–February 2018 rollout period.

A.3.4 Alternative monetization categories

In our main specification, we classify monetization risk based on the topics covered by creators (see Section 5.4). A natural question is whether the moderating effect persists under an alternative operationalization that captures different textual features. We therefore replicate the analysis using a second, independently constructed classification that assigns creators an advertiser-sensitivity level based on lists of advertiser-restricted keywords rather than inferred content topics. Appendix B, Section B.1, describes this approach in detail.

First, we examine how the two classifications relate to one another. The main measure assigns monetization risk based on creators' broad substantive topic categories. The alternative measure instead evaluates the occurrence and contextual use of terms drawn from advertiser exclusion lists. Because those lists cover both sensitive subject matter and potentially objectionable language, the keyword-based measure may capture aspects of both what creators discuss and how they discuss it. Descriptive statistics show that the two classifications are nearly orthogonal at the creator level. We therefore interpret them as distinct operationalizations of advertiser sensitivity rather than as pure measures of topic and linguistic register, respectively. We then use the alternative classification as a moderator in our main regressions. The results, reported in Table A.9, show that the interaction coefficient is not statistically significant. In the weighted CEM specification, the interaction between treatment and high keyword based sensitivity is -0.340 for paid content, statistically insignificant.

Table A.9. Alternative exposure measure: keyword based advertiser sensitivity

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.368***	0.342***	0.339***	-0.052	0.034	0.072
	(0.129)	(0.120)	(0.118)	(0.068)	(0.064)	(0.076)
Treated × High Keyword Sensitivity	-0.324	-0.337	-0.340	0.108	0.110	-0.015
	(0.247)	(0.229)	(0.229)	(0.142)	(0.134)	(0.135)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	70,382	57,486	57,486	70,382	57,486	57,486
Mean DV	2.780	1.864	1.864	1.513	1.217	1.217

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Keyword Sensitivity is the differential effect for the high-moderator group. Post × High Keyword Sensitivity is included as a control; its coefficient is not shown. High Keyword Sensitivity equals one for creators classified as High and zero for creators classified as Low under the keyword-based advertiser-sensitivity measure. Medium and unclassified creators are excluded. This classification is statistically unrelated to the topic-based measure used in the preceding tables. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

We argue that this difference reflects sample composition rather than a substantive contradiction. Because the two classifications are nearly orthogonal, the group identified as high sensitivity by the keyword based measure includes creators with widely varying levels of topical exposure. This motivates estimating an additional model that combines the two classifications, yielding interactions with four mutually exclusive indicator variables.

Table A.10 presents the results from this approach. Within each keyword based group, the estimated effect is larger for creators in sensitive topic categories: 0.585 versus 0.363 among creators classified as keyword-low, and 1.207 versus −0.269 among those classified as keyword-high. When the topic dimension is ignored, the keyword-high group therefore combines the largest and smallest effects across the four categories, generating the negative pooled interaction.[33]

Table A.10. Treatment effect by topic and keyword exposure

	Paid Content		
	(1)	(2)	(3)
Low topic risk × Low keyword sensitivity	0.531**	0.415**	0.363*
	(0.233)	(0.208)	(0.202)
High topic risk × Low keyword sensitivity	0.369*	0.497**	0.585***
	(0.189)	(0.204)	(0.218)
Low topic risk × High keyword sensitivity	-0.371	-0.202	-0.269
	(0.342)	(0.340)	(0.304)
High topic risk × High keyword sensitivity	1.163**	0.716**	1.207**
	(0.477)	(0.360)	(0.478)
CEM matched sample		✓	✓
CEM weights			✓
p (cells equal)	0.048	0.246	0.038
Creators	3485	2880	2880
N	45,305	37,440	37,440
Mean DV	2.953	2.032	2.032

Notes: Difference-in-differences estimates of the treatment effect within each category formed by the topic-based and keyword-based advertiser-sensitivity classifications. The specification is fully saturated, interacting the four categories with both treatment and the post-period indicator; category main effects are absorbed by creator fixed effects. The sample includes only creators classified as High or Low under both measures; Medium and unclassified creators are excluded. Column (1) reports the full sample, Column (2) the CEM-matched sample, and Column (3) the CEM-matched sample using matching weights. All specifications include creator and month fixed effects. The reported p-value tests the joint null that the treatment effects are equal across the four categories. Standard errors clustered at the creator level are reported in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

[33] This approach has an important limitation. Both classifications are derived from the text of Patreon posts rather than from the YouTube content directly affected by the policy change. Subject matter is likely to carry across the two platforms: a creator who writes about history on Patreon is also likely to produce history related videos on YouTube, whereas linguistic register may not. For example, a creator may speak informally in videos while using more neutral language in written posts.

A.3.5 Pre-treatment engagement

Table A.11. Pre-treatment engagement (top versus bottom tercile): total content and composition

	Total Content (Paid + Free)			Paid - Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.153	0.138	0.223**	0.032	0.004	0.055
	(0.115)	(0.114)	(0.113)	(0.083)	(0.082)	(0.081)
Treated × High Engagement	0.450	0.559*	0.229	0.673**	0.497*	0.218
	(0.314)	(0.312)	(0.293)	(0.278)	(0.275)	(0.260)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	72,206	56,650	56,650	72,206	56,650	56,650
Mean DV	4.115	2.699	2.699	1.478	0.731	0.731

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Engagement is the differential effect for the high-moderator group. Post × High Engagement is included as a control; its coefficient is not shown. High Engagement compares the top and bottom terciles of pre-period average monthly comments (terciles across creators); the middle tercile is excluded. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.12. Pre-treatment engagement (top versus bottom tercile): paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.092	0.071	0.139*	0.061	0.067	0.084
	(0.083)	(0.081)	(0.080)	(0.057)	(0.057)	(0.057)
Treated × High Engagement	0.562**	0.528**	0.224	-0.111	0.031	0.006
	(0.271)	(0.269)	(0.248)	(0.120)	(0.120)	(0.124)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	72,206	56,650	56,650	72,206	56,650	56,650
Mean DV	2.796	1.715	1.715	1.319	0.984	0.984

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Engagement is the differential effect for the high-moderator group. Post × High Engagement is included as a control; its coefficient is not shown. High Engagement compares the top and bottom terciles of pre-period average monthly comments (terciles across creators); the middle tercile is excluded. Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.13. Pre-treatment engagement (median split): paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.078	0.061	0.104	0.048	0.064	0.078
	(0.068)	(0.068)	(0.072)	(0.051)	(0.051)	(0.059)
Treated × High Engagement	0.433**	0.382**	0.216	-0.077	0.023	0.010
	(0.195)	(0.187)	(0.185)	(0.102)	(0.093)	(0.100)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	108,025	90,396	90,396
Mean DV	2.315	1.558	1.558	1.277	1.006	1.006

Notes: Difference-in-differences estimates with moderator interactions. Treated is the estimated effect for the low-moderator group; Treated × High Engagement is the differential effect for the high-moderator group. Post × High Engagement is included as a control; its coefficient is not shown. High Engagement equals one for creators with above-median pre-period average monthly comments (median across creators). Treatment (January 2018 cutoff), sample (stayers, Aug 2017–Aug 2018, no adult content), and CEM matching as in the baseline tables. Creator and month fixed effects; the sample is limited to creators with a non-missing moderator. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.3.6 Other additional results

Table A.14. Baseline paid content effect in the heterogeneity estimation samples

	(1)	(2)	(3)	(4)
Treated	0.191**	0.249**	0.219**	0.299*
	(0.083)	(0.124)	(0.112)	(0.170)
Sample	Baseline	High/low monetization risk	Top/bottom engagement terciles	Joint-moderator sample
N	108,025	63,375	72,206	42,120

Notes: Baseline difference-in-differences estimate for paid content, re-estimated without moderator interactions on: (1) the main stayer sample; (2) creators classified as having high or low monetization risk, excluding medium-risk and unclassified creators; (3) creators in the top or bottom tercile of pre-treatment average monthly comments, excluding the middle tercile; and (4) creators satisfying both the monetization-risk and engagement restrictions. All specifications include creator and month fixed effects. Standard errors are clustered at the creator level. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.15. Inverse-hyperbolic-sine patrons and earnings

	asinh(Patrons)			asinh(Earnings)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.089***	0.088***	0.086***	0.098***	0.098***	0.091***
	(0.014)	(0.015)	(0.017)	(0.022)	(0.022)	(0.026)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	85,818	74,671	74,671
Mean DV	3.442	3.191	3.191	4.741	4.546	4.546

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Robustness for the monetization result. Outcomes are inverse-hyperbolic-sine transformed; coefficients approximate proportional effects. Matching: CEM as in the baseline. All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.16. Patrons and earnings in levels

	Patrons			Earnings		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	9.014***	2.262	3.473**	29.076***	14.763***	14.069***
	(2.711)	(2.317)	(1.725)	(6.758)	(5.380)	(5.227)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	85,818	74,671	74,671
Mean DV	98.239	48.988	48.988	346.292	208.497	208.497

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Levels are reported for magnitude context only; the headline monetization specification is the flow (first-difference) table. Matching: CEM as in the baseline. All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.4 Robustness checks

In this section, we assess the sensitivity of the results to alternative functional forms, treated-control comparability, treatment timing, sample composition, and short-run pre-treatment changes.

A.4.1 Alternative functional form

The PPML estimates, suited to the nonnegative and skewed outcome distributions, closely mirror the linear results for content: Tables A.17 and A.18 put the weighted CEM semi-elasticity at approximately 8% for total content and 10% for paid content, with the free content estimate small and statistically insignificant. For monetization, Table A.19 yields a positive but imprecise patrons estimate alongside an earnings effect of approximately 6%, significant at the 5% level. The functional form evidence is accordingly strongest for the content outcomes and earnings; it does not deliver a statistically precise patron effect.

Table A.17. PPML estimates: total content

	Total Content (Paid + Free)		
	(1)	(2)	(3)
Treated	0.055*	0.083**	0.083**
	(0.032)	(0.038)	(0.039)
CEM matched sample		✓	✓
CEM weights			✓
N	91,549	74,737	74,737
Mean DV	4.240	3.101	3.101

Notes: Poisson pseudo-maximum-likelihood (PPML) difference-in-differences estimates; coefficients are semi-elasticities. Treated equals one for creators with a YouTube presence before January 2018, using link information through December 2017. Sample: stayers, Aug 2017–Aug 2018, excluding adult content; creator–month observations. Column (1): full sample; (2): CEM-matched sample; (3): CEM-matched sample with matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Observations separated by a fixed effect or from creators with all-zero outcomes are dropped, so N differs from the OLS tables. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.18. PPML estimates: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.090**	0.104**	0.097**	-0.008	0.056	0.067
	(0.039)	(0.047)	(0.048)	(0.044)	(0.048)	(0.054)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	75,415	59,801	59,801	78,259	63,375	63,375
Mean DV	3.317	2.355	2.355	1.763	1.435	1.435

Notes: Poisson pseudo-maximum-likelihood (PPML) difference-in-differences estimates; coefficients are semi-elasticities. Treated equals one for creators with a YouTube presence before January 2018, using link information through December 2017. Sample: stayers, Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Observations separated by a fixed effect or from creators with all-zero outcomes are dropped, so N differs from the OLS tables. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.19. PPML estimates: patrons and earnings

	Patrons			Earnings		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	-0.022	0.021	0.046	0.039	0.059**	0.057**
	(0.039)	(0.046)	(0.033)	(0.029)	(0.026)	(0.024)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	85,818	74,671	74,671
Mean DV	98.239	48.988	48.988	346.292	208.497	208.497

Notes: Poisson pseudo-maximum-likelihood (PPML) difference-in-differences estimates; coefficients are semi-elasticities. Treated equals one for creators with a YouTube presence before January 2018, using link information through December 2017. Sample: stayers, Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Observations separated by a fixed effect or from creators with all-zero outcomes are dropped, so N differs from the OLS tables. Mean DV is the estimation-sample mean. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.4.2 Alternative matching

Table A.20 presents the covariate balance obtained with the PSM approach. Table A.21 presents the results obtained with this alternative matching. The coefficient of interest remains positive. As shown in Table A.22, we obtain comparable results with a tighter CEM specification, although the weighted estimates are less precise. We place less focus on the PSM estimates because post-matching balance is weaker for several activity and audience characteristics than under the main CEM design.

Table A.20. Covariate balance, PSM specification (5-NN)

	Unmatched			Matched (weighted)			
	Treated	Control	Std. diff. (%)	Treated	Control	Std. diff. (%)	$p-value$
Contents per month	3.72	3.26	8.3	3.63	3.81	-3.3	0.408
Comments per month	14.25	9.92	9.0	11.50	15.13	-7.5	0.069
Likes per month	25.28	21.01	4.3	20.81	30.94	-10.2	0.021
Free contents per month	1.45	1.12	11.6	1.44	1.27	6.1	0.081
Patrons	70.95	35.59	14.7	54.89	52.88	0.8	0.780
Earnings (USD)	322.88	185.50	14.4	261.22	273.09	-1.2	0.722
Facebook likes	28633.90	9181.86	11.7	27275.74	8625.68	11.2	0.000
Twitter followers	10010.71	5114.08	8.2	8289.75	6444.82	3.1	0.149
Monetization risk (1–4)	2.96	2.92	4.0	2.95	2.99	-3.9	0.329
Creators	7,336	1,100		7,209	1,100		

Notes: Covariates are creator-level averages over the window pre-period (August–November 2017), computed on the main stayer sample. Standardized differences are in percent and use the pooled unmatched standard deviation in both panels. Matched-sample $p-values$ come from weighted regressions of each covariate on the treatment indicator with heteroskedasticity-robust standard errors. Matched columns use PSM weights: five-nearest-neighbour propensity-score matching with replacement, ties, and common support.

Table A.21. Propensity-score matching: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.191**	0.168**	0.183*	-0.031	-0.031	0.004
	(0.083)	(0.082)	(0.097)	(0.047)	(0.047)	(0.058)
PSM matched sample		✓	✓		✓	✓
PSM weights			✓			✓
N	108,025	106,400	106,400	108,025	106,400	106,400
Mean DV	2.315	2.219	2.219	1.277	1.266	1.266

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): PSM-matched sample; (3)/(6): PSM-matched, matching weights. Matching: propensity-score matching (logit on the CEM covariates; five nearest neighbours with replacement, ties, common support). All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.22. Tighter CEM specification: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.191**	0.210***	0.136	-0.031	0.034	0.035
	(0.083)	(0.072)	(0.096)	(0.047)	(0.046)	(0.050)
Tight CEM matched sample		✓	✓		✓	✓
Tight CEM weights			✓			✓
N	108,025	69,470	69,470	108,025	69,470	69,470
Mean DV	2.315	1.300	1.300	1.277	0.879	0.879

Notes: Difference-in-differences estimates. Treated equals one for creators with a YouTube presence before January 2018 (link information through December 2017), in all months from January 2018 onward (announcement of the stricter YouTube Partner Program criteria). Sample: Patreon video creators active before and after the shock (stayers), Aug 2017–Aug 2018, excluding adult content; creator–month observations. Columns (1)/(4): full sample; (2)/(5): Tight CEM-matched sample; (3)/(6): Tight CEM-matched, matching weights. Matching: CEM additionally coarsening pre-period Facebook and Twitter reach (tighter balance, lower retention of treated creators). All specifications include creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.4.3 Alternative treatment dates

Our results also remain stable when using alternative treatment dates. Starting the post-treatment period in December 2017, at the onset of the enforcement expansion rollout, yields a weighted CEM estimate of 0.209 for paid content (see Table A.23). Starting it in February 2018, when the stricter eligibility criteria were enforced, yields an estimate of 0.217 (see Table A.24). Both estimates are significant at the 1% level, while the corresponding estimates for free content are small under both definitions.

Table A.23. Alternative cutoff in December 2017: paid and free content.

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.206**	0.175**	0.209***	-0.018	0.068	0.078*
	(0.084)	(0.078)	(0.080)	(0.048)	(0.045)	(0.047)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	108,025	90,396	90,396
Mean DV	2.315	1.558	1.558	1.277	1.006	1.006

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Treatment switches on in December 2017 (start of the enforcement-expansion rollout); the baseline dates the post period from the January 2018 announcement of stricter YPP eligibility. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.24. Alternative cutoff in February 2018: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.168**	0.167**	0.217***	-0.044	0.028	0.021
	(0.082)	(0.074)	(0.076)	(0.046)	(0.043)	(0.048)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	108,025	90,396	90,396	108,025	90,396	90,396
Mean DV	2.315	1.558	1.558	1.277	1.006	1.006

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Treatment switches on in February 2018 (enforcement of stricter YPP eligibility); December 2017–January 2018 are coded as pre-period. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.4.4 Alternative samples

Our results also hold across alternative sample definitions. Including entrants and exiters (see Table A.25), restricting to a strictly balanced panel (Table A.26), and excluding creators above the 99th percentile of pre-shock patrons or earnings (Table A.27) all yield positive and statistically significant weighted estimates for paid content, with free content statistically unchanged. Therefore, we see that the paid content response is neither confined to the baseline stayer definition nor driven by the upper tail of the pre-treatment patron and earnings distributions.

Table A.25. All creators, including entrants and exiters: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.193**	0.165**	0.153**	-0.032	0.047	0.057
	(0.083)	(0.076)	(0.076)	(0.047)	(0.043)	(0.048)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	116,552	90,396	90,396	116,552	90,396	90,396
Mean DV	2.331	1.558	1.558	1.309	1.006	1.006

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Sample includes all video creators observed in the window; entrants and exiters enter whenever outcomes are observed. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.26. Strictly balanced panel: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.203**	0.183**	0.183**	-0.036	0.042	0.054
	(0.087)	(0.080)	(0.080)	(0.049)	(0.045)	(0.051)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	100,321	83,603	83,603	100,321	83,603	83,603
Mean DV	2.358	1.599	1.599	1.288	1.015	1.015

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Sample restricted to creators observed in every month from August 2017 through August 2018 (balanced panel). Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.27. Excluding the top 1% of creators: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.191**	0.156**	0.150**	-0.034	0.047	0.057
	(0.083)	(0.076)	(0.076)	(0.046)	(0.043)	(0.048)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	106,595	90,266	90,266	106,595	90,266	90,266
Mean DV	2.268	1.555	1.555	1.264	1.004	1.004

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Sample excludes creators above the 99th percentile of pre-period average patrons or pre-period average earnings. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

A.4.5 Placebo treatment timing

We also probe the identification of the monetization result with placebo timing.

First, we consider a short window placebo and assign treatment from November 2017, a false cutoff immediately preceding the rollout. Results are presented in Tables A.28 and A.29. The placebo content coefficients are statistically insignificant, as are the monetization coefficients in the matched specifications.[34]

A second placebo shifts the design to an earlier period without a comparable governance change, estimating the net monthly change in patrons and earnings within a window ending before the March 2017 advertiser boycott.[35] The matched placebo coefficients are small and statistically insignificant (Table A.30). In the full, unmatched sample, the clean placebo is negative, indicating that the unmatched treated and comparison creators were not on parallel "flow" paths in earlier periods either, reinforcing the matched comparisons on which the design relies. The absence of a flow effect during a genuine no shock window, combined with a significant effect at the actual rollout, suggests that the monetization result captures a change in the rate of community-funded support at the time of the policy change rather than pre-existing differential growth.

Table A.28. Short window placebo: paid and free content

	Paid Content			Free Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.092	0.030	0.140	-0.021	0.020	0.023
	(0.082)	(0.069)	(0.096)	(0.054)	(0.048)	(0.054)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	41,287	34,561	34,561	41,287	34,561	34,561
Mean DV	2.299	1.516	1.516	1.325	1.011	1.011

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Placebo: treatment assigned from November 2017 onward, estimated on the August–December 2017 pre-period only. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

[34] The positive and statistically significant earnings coefficient in the unmatched full sample, however, highlights the importance of relying on the matched comparisons.

[35] Here, we use additional historical Graphtreon records covering August 2016-February 2017.

Table A.29. Short window placebo: patrons and earnings

	Patrons			Earnings		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	1.537	-0.527	0.384	9.268**	-0.443	-3.083
	(1.712)	(1.338)	(0.939)	(4.646)	(4.087)	(5.912)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	41,287	34,561	34,561	34,238	29,689	29,689
Mean DV	86.709	42.403	42.403	336.677	195.162	195.162

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Placebo: treatment assigned from November 2017 onward, estimated on the August–December 2017 pre-period only. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A.30. Clean window placebo (pre-March 2017): net monthly change in patrons and earnings

	Net Change in Patrons			Net Change in Earnings		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	-2.914***	-0.516	-0.275	-12.051**	-4.252	-3.140
	(0.780)	(0.363)	(0.489)	(4.713)	(3.213)	(3.170)
CEM matched sample		✓	✓		✓	✓
CEM weights			✓			✓
N	28,321	23,700	23,700	27,996	23,479	23,479
Mean DV	3.498	1.640	1.640	13.077	6.702	6.702

Notes: Difference-in-differences estimates; design as in the baseline tables except as noted. Clean placebo (Aug 2016–Feb 2017; Dec 2016 cutoff; pre-March 2017); outcomes in one-month first differences. Columns (1)/(4): full sample; (2)/(5): CEM-matched sample; (3)/(6): CEM-matched, matching weights. Matching: coarsened exact matching (CEM) on pre-period content, comments, patrons, earnings, and monetization risk. Creator and month fixed effects. Mean DV is the estimation-sample mean of the dependent variable. Standard errors clustered at the creator level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Appendix B

B.1 Monetization risk categories

To classify creators into our monetization risk categories, we use two approaches based on textual information from their Patreon pages. The main analysis relies on categories defined by content topic. In Section A.3.4, we also report results based on an alternative classification using advertiser sensitive keywords.

B.1.1 Monetization risk based on topics

In this approach, we classify content creators into topic categories using textual information extracted from their Patreon posts. The classification draws on post titles and accompanying text, allowing us to capture a richer, more detailed description of the content that creators offer their communities. Using this procedure, we assign at least one topic category to 97% of creators. The remaining cases lack sufficient meaningful textual information to support a reliable classification. The categorization procedure proceeds in four stages: (i) topic identification using topic modeling and a large language model, (ii) cleaning and filtering of non-substantive topics, (iii) consolidation into broad analytical categories, and (iv) classification into monetization risk groups.

Although we report statistics across four categories, i.e., unclear, low, medium, and high (See Table 3), our regressions use a binary indicator, "High" equal to one for content creators in the high category (see, for example, Tables 5 and 6).

Stage 1: Topic identification using BERTopic and LLM labeling We first assign topic labels at the individual-post level using a BERTopic-based topic-modeling framework. Text embeddings are generated with the `snowflake-arctic-embed2:568m` model, which captures semantic similarities across titles and accompanying notes. Topic labels are then generated using the `gemini-3-flash-preview` large language model. This procedure produces a primary topic and a ranked list of additional topic descriptors for each post. We then aggregate these post-level classifications to the creator level. For each creator, we identify one primary topic and up to two secondary topics based on the frequency and prominence of the post-level topic assignments.

Stage 2: Cleaning and filtering of non-substantive topics The raw topic output includes both substantive content topics and labels that describe administrative communication, content formats, or community management activities rather than the subject matter of the creator's content. To ensure that the resulting categories capture substantive themes, we systematically clean the topic labels.[36] We then normalize the remaining labels by converting them to lowercase, trimming whitespace, and removing extraneous characters in preparation for the construction of broader categories.

Stage 3: Construction of broad content categories We consolidate the cleaned topic labels into broader thematic categories using rule-based keyword matching implemented through regular expressions. Each creator is assigned to one mutually exclusive broad category according to the first applicable matching rule. Figure B.1 presents the resulting categories and their shares in the sample.

Figure B.1. Text-based categories in the Patreon sample

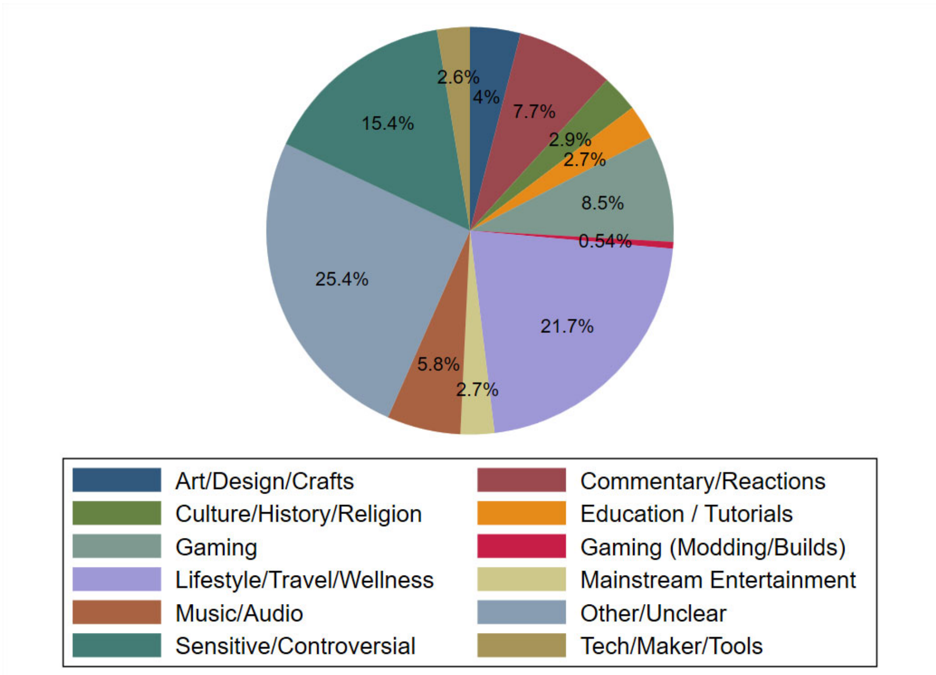

Notes: The figure is based on 10,248 unique content creators.

[36] In particular, we remove or recode as "N/A" labels corresponding to: (i) platform logistics and rewards, such as patron appreciation posts, milestone updates, reward fulfillment, early-access announcements, merchandise, and Discord access; (ii) scheduling and project updates, such as weekly schedules, video-release calendars, general updates, and works in progress; (iii) engagement formats, such as polls, Q&A threads, giveaways, and introductions; (iv) generic media formats, such as "Podcast Episodes," "Behind the Scenes," "Sneak Peeks," and "Video Ideas"; (v) celebrations and time-specific announcements, such as holiday greetings and anniversaries; and (vi) archives and one-off collections, such as digital wallpaper archives or named episode collections. When a creator's primary topic is non-substantive and therefore coded as "N/A," we replace it sequentially with the first available substantive secondary topic. This procedure ensures that each creator is assigned the most informative content descriptor available.

Stage 4: Monetization risk classification To operationalize monetization risk, we map the broad thematic categories into three levels of advertising sensitivity, together with a separate unclassified group. We define: (i) high risk, comprising creators in the Sensitive/Controversial category; (ii) medium risk, comprising creators in the Culture/History/Religion and Commentary/Reactions categories; (iii) low risk, comprising creators in all remaining substantive categories; and (iv) unclear, comprising creators classified as Other/Unclear. This classification is intended to capture differences in creators' exposure to advertiser concerns and, consequently, to the risk of limited or unavailable advertising monetization. It therefore provides a topic-based measure of exposure to the governance shock examined in the paper.

B.1.2 Monetization risk based on advertiser sensitive keywords

Our second approach classifies creators using advertiser-sensitive terms appearing in their Patreon content. Whereas the main classification assigns risk on the basis of creators' broad substantive topic categories, this alternative measure evaluates the occurrence and contextual use of terms drawn from advertiser exclusion lists. We begin with a list of keywords compiled by Adalytics[37] from exclusion lists associated with approximately 350 major U.S. advertisers. The approach is similar in spirit to Rasskazova (2026), although we apply the keyword information to the titles and accompanying text of Patreon posts. The categorization procedure proceeds in three stages: (i) preparation of the textual evidence and identification of advertiser-sensitive keywords; (ii) contextual assessment using a large language model; and (iii) aggregation of the resulting classifications to the creator level.

For the regression analyses, we construct a binary moderator comparing creators classified as High with those classified as Low under this keyword-based advertiser-sensitivity measure. Creators classified as Medium or Unknown are excluded from these specifications (See Tables A.9 and A.10).

Stage 1: Text preparation and keyword screening We first clean and normalize the titles and accompanying text of creators' Patreon posts. Because some creators have a large number of posts, the text is divided into bounded evidence chunks that can be assessed separately. We then identify occurrences of the advertiser sensitive *terms* collected by Adalytics using normalized phrase-boundary matching. The matched terms are organized into eight broader issue families representing common sources of advertiser sensitivity. These include (1) inappropriate language, (2) violence, (3) adult content, (4) shocking content, (5) harmful

[37] See `https://adalytics.io/blog/major-brands-appear-to-be-blocking-ad-targeting-keywords`.

or dangerous acts, (6) hateful content, (7) controversial issues, and (8) sensitive events. A post or creator may be associated with more than one issue family.

Stage 2: Contextual classification using an LLM We next use the `gemini-2.5-flash-lite` large language model to assess the use of the sensitive keywords in context. For each evidence chunk, the model receives the cleaned Patreon text, the identified keyword matches and issue families, and instructions to distinguish the simple mention of a sensitive term from its endorsement, promotion, or graphic treatment.

We obtain a *general advertiser sensitivity* assessment. This assessment captures whether the content is likely to raise broad brand safety concerns (based on the advertiser's keywords), without implying that the content violates YouTube's policies or that the creator is ineligible for monetization. Each evidence chunk is classified as High, Medium, Low, or Unknown. The Unknown category is used when the available Patreon text does not provide sufficient information for a reliable assessment.

Stage 3: Creator-level aggregation and risk classification Because a creator may have several evidence chunks, we aggregate the chunk-level classifications to the creator level using the most sensitive classification observed. A creator is classified as High if at least one evidence chunk is rated High. In the absence of a High classification, the creator is classified as Medium if at least one chunk is rated Medium, and as Low if at least one chunk is rated Low. Creators for whom none of the available evidence supports a classification remain in the Unknown category.

The resulting measure provides an alternative operationalization of monetization risk based on the occurrence and contextual use of advertiser-restricted terms in Patreon post text. Unlike the topic-based measure, it does not assign risk on the basis of a creator's broad substantive category. It may nevertheless capture both subject matter and linguistic expression, because the underlying exclusion lists contain terms associated with sensitive topics as well as inappropriate language. Because the measure is derived from Patreon posts rather than directly from YouTube videos, we interpret it as a proxy for exposure to advertiser sensitivity, not as a direct measure of demonetization or policy enforcement.

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