

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

El-Komboz, Lena Abou; Kerkhof, Anna; Loh, Johannes

Working Paper Platform Partnership Programs and Content Supply: Evidence from the YouTube "Adpocalypse"

CESifo Working Paper, No. 10363

Provided in Cooperation with: Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: El-Komboz, Lena Abou; Kerkhof, Anna; Loh, Johannes (2023) : Platform Partnership Programs and Content Supply: Evidence from the YouTube "Adpocalypse", CESifo Working Paper, No. 10363, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at: https://hdl.handle.net/10419/272007

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

CESIFO WORKING PAPERS

Platform Partnership Programs and Content Supply: Evidence from the YouTube "Adpocalypse"

Lena Abou El-Komboz, Anna Kerkhof, Johannes Loh



Impressum:

CESifo Working Papers ISSN 2364-1428 (electronic version) Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute Poschingerstr. 5, 81679 Munich, Germany Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de Editor: Clemens Fuest https://www.cesifo.org/en/wp An electronic version of the paper may be downloaded • from the SSRN website: www.SSRN.com

- from the RePEc website: <u>www.RePEc.org</u>
- from the CESifo website: <u>https://www.cesifo.org/en/wp</u>

Platform Partnership Programs and Content Supply: Evidence from the YouTube "Adpocalypse"

Abstract

Many digital platform host content produced by independent creators and rely on advertising as their primary source of revenues. They commonly use partnership programs to put into place incentives for creators to produce high-quality content by sharing part of the advertising revenue with them. In addition, these programs let them exercise control over their participation to prevent the presence of "bad-faith" actors who can otherwise harm the integrity of the platform. However, the rules governing access to such programs may have to be adjusted over time, which in turn may disrupt creators' motivation to produce content. This paper studies a rule change on YouTube that made access to its partner program more restrictive. This also removed all former participants who did not meet the new requirements and made it impossible for them to continue commercializing their content on the platform. Using a regression discontinuity design, we provide causal evidence that affected creators subsequently reduced the frequency of their uploads and provided content of lower quality and diversity. We also investigate and discuss effect heterogeneity between mainstream and niche as well as more and less experienced creators to learn about the underlying financial and non-pecuniary motivations. Our findings provide novel insights about the effective governance of ad-based platforms using partnership programs.

Keywords: platform governance, partnership programs, content supply, ad-based business models, access restrictions.

Lena Abou El-Komboz ifo Institute – Leibniz Institute for Economic Research at the University of Munich / Germany abou-el-komboz@ifo.de

Anna Kerkhof ifo Institute – Leibniz Institute for Economic Research at the University of Munich / Germany kerkhof@ifo.de Johannes Loh* BI Norwegian Business School Oslo / Norway johannes.loh@bi.no

*corresponding author

April 1, 2023

We thank Sam Cao, Verena Dorner, Luise Eisfeld, Oliver Falck, Philipp Hukal, Tobias Kretschmer, Tim Meyer, Christian Peukert, and participants of the 6th Digital Economics Workshop in Lausanne, for useful comments and suggestions. The paper has benefitted from feedback from conference participants at the 6th Digital Economics Workshop in Lausanne and the SMS annual meeting in London, as well as seminar participants at Ludwig-Maximilians-University, the ifo institute, and BI Norwegian Business School. Anna Kerkhof gratefully acknowledges financial support by the Joachim-Herz-Stiftung. This project is funded by the Bavarian State Ministry of Science and the Arts in the framework of the bidt Graduate Center for Postdocs. Lena Abou El-Komboz acknowledges support by the Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119).

1 Introduction

Recent years saw the emergence of large digital platforms that dominate the delivery of various types of online media, such as software, music, news, and videos (Wu and Zhu, 2022; Waldfogel, 2017). These platforms commonly rely on independent creators who produce the content, which then attracts an audience of end consumers (Jain and Qian, 2021). Popular examples include YouTube, which has more than 2 billion monthly users, with creators uploading more than 500 hours of video content per minute.¹ Likewise, on the social media platform Twitter, more than 350 million monthly active users themselves generate content to be consumed by their peers.² In addition, many consumers have come to expect the offered content to be free, which is why ad-based business models are a prevalent source of revenues (Sun and Zhu, 2013). Hence, platforms in this "creator economy" reside over multi-sided ecosystems that connect content creators, consumers, and advertisers (Bhargava, 2022).

These platforms have to attract and engage consumers to ensure a steady inflow of ad revenue. This means that they need to put into place incentives for creators to produce a steady stream of new content of high quality (Huang et al., 2022). Commonly they therefore share part of the generated ad revenue with creators (Tang et al., 2012). In addition, platforms can face an inflow of low-quality or "bad faith" actors, which can threaten the health and integrity of their ecosystems (Geva et al., 2019). For instance, the crash of the video game industry in the early 1980s has been attributed to the high number of games with low quality or obscene content (Pursey, 2022). Further, the social media platform Twitter has rolled out its crowdsourced fact-checking project "Community Notes" to deal with false information (Coleman, 2021), and YouTube has faced intense public backlash and advertiser boycotts in 2017 due to the presence of hate-speech and other problematic content (Statt, 2017). To address both challenges, some content-based platforms have created so-called partnership programs as a means to regulate which creators are able to monetize their content. For example, the streaming platform Twitch requires a minimum amount of content supply, consumer engagement, and conformity with guidelines for creators to become eligible for their basic "Affiliate"³ and more prestigious "Partner"⁴ programs. YouTube applies similar criteria, but paired with a degree

¹See https://blog.youtube/press, accessed on 13 February 2023

²See https://www.statista.com/statistics/303681/twitter-users-worldwide/, accessed on 13 February 2023

³See https://help.twitch.tv/s/article/joining-the-affiliate-program, accessed on 13 February 2023 ⁴See https://help.twitch.tv/s/article/partner-program-overview, accessed on 13 February 2023

of manual curation to govern participation in their partner program, which lets creators earn part of the ad revenue.⁵ These programs are useful for platforms: They incentivize the creation of high quality content while regulating access to them, which protects advertisers from being associated with "bad faith" actors.

However, regulating access to partnership programs is a non-trivial governance challenge for platform owners: On the one hand, if the eligibility criteria for creators are too restrictive, their effectiveness as an incentive device will be limited. On the other hand, if they are too open, then their control function will be compromised. Therefore, the decision about how selective access should be is a balancing act (Boudreau, 2010). In addition, as a platform evolves over time, the criteria for access may have to be adjusted as well (Rietveld et al., 2020). This, however, can be a disruptive event (Jacobides et al., 2018), which can create confusion and uncertainty (Jhaver et al., 2018) and ultimately lead to unanticipated and undesirable reactions by creators (Gawer and Henderson, 2007; Tiwana, 2015). This issue is exacerbated in the context of the creator economy, in which the supply of content is determined by a heterogeneous mix of financial and non-pecuniary (such as status- and identity-based) sources of motivation (Ma and Agarwal, 2007), which makes reactions to governance attempts hard to predict (Boudreau and Hagiu, 2009).

In this paper, we study how creators on YouTube reacted to a change to the access requirements to its partnership program. After facing intensive public backlash and advertiser boycotts as part of the so-called YouTube "Adpocalypse" (Alexander, 2019), the platform significantly increased the eligibility criteria in an effort to exert more control over who participates. This made it not only harder for new creators to become "YouTube Partners", but it also removed all former program participants who did not meet the new criteria at the time of the rule change. Hence, these creators where not completely shut out of the platform, but they lost their partner status as well as the possibility to monetize their content. To study how this had an effect on their subsequent motivations to create content on the platform, we ask the following research questions: *How did YouTube creators react to losing access to the partnership program in terms of content supply? And how did this reaction vary across different creator types?* In particular, we investigate the impact of the rule change on the amount, quality, and diversity of creator's subsequent content production.

Prior research has studied how regulating access to a platform as such is related to the supply of complementary products (Boudreau, 2010), highlighting an important trade-off in

⁵See https://support.google.com/youtube/answer/72851, accessed on 13 February 2023

more open systems: They benefit from increased network effects due to a larger number of participants (Eisenmann et al., 2006), which however comes at the expense of the quality and innovativeness of the products on offer (Boudreau, 2012; Parker and Van Alstyne, 2018). However, we know less about how regulating access to partnership programs affects the supply of content. This is distinct, because the rules governing access to such programs do not prevent creators' participation in the platform ecosystem altogether. In addition, the questions of how existing creators react to a change in the degree of access control, and what the role of complementor heterogeneity is, are hitherto unanswered. However, these aspects are important determinants of partnership programs' effectiveness in incentivizing the creation of content while protecting a platform's integrity.

Empirically, we leverage a unique dataset about Youtube creators' content supply and estimate the causal effect of losing access to the partnership program within a regression discontinuity design. The new eligibility criteria for program participation constitute a clear threshold in creators' subscriber count at the time of the rule change. We therefore compare the subsequent creation of videos between those who are just below (lost access) and just above (remained in the program) that threshold. In our analysis of German creators, we find that those who lost access to the program subsequently decreased their frequency of video uploads, and their content was both of lower quality and diversity. We also provide evidence for and discuss heterogeneity in these effects between mainstream and niche creators, as well as along their prechange experience to learn about how the rule change had an impact on different financial and non-pecuniary motivational sources. In particular, our results suggest that the loss of financial incentives is not sufficient in explaining heterogeneity in reactions across these creator types. Instead, we attribute this to non-pecuniary motivations arising from the loss of status (Toubia and Stephen, 2013) or their identity-based attachment (Ren et al., 2007) to the platform.

Our findings have important implications for the successful governance of platform ecosystems when using partnership programs. While primarily aimed at creating financial incentives for creators, our results imply that the acceptance to such programs also comes with additional, non-pecuniary incentives. This, in turn, drives heterogeneity in adverse reactions when program access is denied. Our contribution to this stream is therefore two-fold: First, we highlight partnership programs as a control device for platform owners, as well as some of the factors that can determine their (un-)successful implementation. And second, our results show the risks associated with "one-size-fits-all" governance attempts. Instead, platform owners need to heed the limitations of governance practices that are applied indiscriminately across the entire ecosystem.

2 Related Literature

2.1 Platform governance

Our study relates to the stream of literature on platform governance. This literature investigates the set of rules and design features put into place by platform owners to coordinate and facilitate complementors' (in our case: creators') value creation processes (Boudreau and Hagiu, 2009; Ceccagnoli et al., 2012; Wareham et al., 2014). Specifically, researchers have studied different types of platform governance, such as boundary resources or interface features (such as APIs for app developers) (Ghazawneh and Henfridsson, 2013; Tae et al., 2020), selective promotion of complements (Rietveld et al., 2019), and algorithmic or ranking-based recommendation systems to guide and facilitate interactions between consumers and complementors (e.g. Dinerstein et al., 2018; Kapoor and Agarwal, 2017; Oestreicher-Singer and Sundararajan, 2012). Most relevant for our study is the discussion of platform control, that is, the extent to which owners attempt to influence complementors' activities by restricting the types of activities they are allowed to carry out. Prior studies have investigated trade-offs related to how restrictive access to the platform should be (e.g. Boudreau, 2010, 2012; Parker and Van Alstyne, 2018) and have shown that too few restrictions can compromise the quality and integrity of the products and services on offer (Geva et al., 2019; Eaton et al., 2015). Others have investigated how "softer" control measures affect complementor activity, such as certification systems or awards (Huang et al., 2013; Rietveld et al., 2021; Foerderer et al., 2021), sending signals about desirable content (Hukal et al., 2020), or regulating access to boundary resources (Constantinides et al., 2018) or users (Claussen et al., 2013). In addition, some studies note the potentially disruptive effects of changing governance practices (Jacobides et al., 2018; Jhaver et al., 2018; Koo and Eesley, 2020), which can entail unintended and detrimental complementor reactions such as reduced performance and exit (Gawer and Henderson, 2007; Tiwana, 2015).

While these studies provide diverse insights about how platform owners govern their ecosystems, partnership programs have been overlooked as a potentially useful governance tool. In particular, they allow owners to exert control over complementor activities by regulating access to programs. This, however, does not shut them out of the ecosystem completely, but rather exposes them to incentives to supply content. In addition, while we have a good understanding about how varying degrees of access control can drive heterogeneity in the size and quality across different platforms, we know less about how complementors react to a *change* in control over their activities. This is particularly relevant in the context of access to partnership programs because it regulates the extent to which they are exposed to supply incentives, and therefore carries implications for the performance of the platform as a whole.

2.2 Ad-based business models

In addition, scholars have investigated ad-based platform business models, and in particular the dynamics of revenue sharing with creators. Two recent theoretical studies analyze platforms' optimal revenue sharing policies as a function of their size and between-creator competition (Jain and Qian, 2021), and the determinants of content creation under ad-revenue sharing, as well as several ways for the platform to optimize these processes (Bhargava, 2022). In addition, a handful of empirical studies investigate how earning a share of the generated ad revenue affects creators' subsequent content production. First, Tang et al. (2012) show that a share of the ad-revenue works as an incentive for content creators on YouTube, in addition to reputational concerns they have. Second, Sun and Zhu (2013) find that bloggers increased the quality of their content and shifted production towards more popular topics after they adopted such a model. This is consistent with early studies noting that advertising creates an incentive to shift production towards popular content in an effort to maximize "eyeballs" (Steiner, 1952; Wilbur, 2008), which leads to a duplication of mainstream at the expense of niche content (Anderson and Gabszewicz, 2006). Second and in contrast, Kerkhof (2022) finds that an increase in the amount of advertising that creators on YouTube include in their videos *decreases* their tendency to duplicate mainstream content. This is attributed to increased competitive pressure in this market segment. Third, Wu and Zhu (2022) show that authors on a creative writing platform react more strongly – in terms of effort provision and novelty – to increased competition when they earn a share of advertising revenue.

These findings provide evidence that the creation of financial incentives affects creators' supply of content, both in terms of quality and type (e.g. popular vs. niche). However, especially regarding the latter, the evidence about the direction of the effect is mixed. In addition, prior studies mainly point to *aggregate* effects of the creation of ad-based incentives, but have paid little attention to the potential role of creator *heterogeneity*. This, however, is an important factor in determining the effectiveness of incentive devices in platform-based ecosystems. For one, different creators can vary in their needs and characteristics (Boudreau and Hagiu, 2009), which makes uniform reactions to incentive mechanisms unlikely. Moreover,

this should be particularly pronounced in the context of user-generated and creative content, where creators are driven by a diverse mix of financial and non-pecuniary motivational sources (Ma and Agarwal, 2007). Hence, the questions of which type of incentives are actually put into place via partnership programs, and how this may vary across different creator types, are important determinants of their effectiveness. And finally, these studies provide evidence about the creation of incentives, but not what happens when creators lose access to them. This, however, is an important distinction as economic actors may react differently and more strongly to losses than gains (Tversky and Kahneman, 1991).

2.3 User-generated content

Finally, our study relates to the literature studying potential nudges to stimulate user-generated content (UGC). The effectiveness of financial incentives has been studied in the context of online reviews, providing nuanced insights. Cabral and Li (2015) find only a small positive influence on creator activity. Similarly, Burtch et al. (2018) find that they increase activity, but are most effective in combination with other nudges, such as social norms. Others find that the effectiveness of financial rewards is subject to heterogeneity across different creator types: Sun et al. (2017) document a stifling effect on activity among well-connected creators, but a stimulating effect for others. Similarly, Khern-am nuai et al. (2018) find that the platform experienced an inflow of new creators after financial incentives had been put into place, but that existing ones decreased their activity. These mixed findings illustrate the complex motivational sources underlying user-generated content, and that financial incentives may not be the most important stimulant. Indeed, a range of studies show the importance of intrinsic or reputationbased motivations (e.g. Ma and Agarwal, 2007; Toubia and Stephen, 2013; Tang et al., 2012). Many UGC-based platforms therefore use awards or community badges to motivate activity. Prior research largely found positive effects, with reputation-based nudges stimulating activity on StackOverflow (Anderson et al., 2013) and increasing newcomer retention on Wikipedia (Gallus, 2017). At the same time, Burtch et al. (2022) – while finding a positive effect on activity – also document that (peer) awards tend to decrease content novelty, thus demonstrating a potential downside in the context of the generation of creative content. Goes et al. (2016) highlight an additional limitation in the context of milestone-based incentive hierarchies: Consistent with motivation stemming from the pursuit of a goal (Locke and Latham, 2002), they find that milestones initially stimulate activity, but that motivations decrease immediately after the successful accomplishment.

These studies further highlight the multi-layered motivational sources on platforms that (partly) rely on user-generated content. This suggests that governing access to partnership programs may not be purely related to financial incentives, but also non-pecuniary concerns such as creator reputation and a sense of accomplishment. Therefore, removing creators from such a program is likely not the same as merely taking away financial incentives. In our study, we use these insights about complex motivational sources to explain heterogeneous creator reactions to losing access to a partner program, and argue that the financial loss is not a sufficient explanation.

3 Institutional Background

YouTube is the world's largest video sharing platform, and – as of February 2023 – the secondlargest website in terms of overall traffic behind Google.⁶ YouTube has more than 2 billion monthly users, and more than 500 hours of user-generated videos are uploaded every minute.⁷ These videos constitute YouTube's supply of complements; they are created by registered users which are commonly referred to and self-identify as "YouTubers" (Kerkhof, 2022). Creators upload videos to their own "channels", which can be subscribed to by viewers, who will subsequently become informed about the release of new content. The videos as such, however, can be viewed by anyone for free. While also offering paid premium memberships for viewers⁸, YouTube's main source of revenue is advertisements that are played before and during videos.

3.1 YouTube partner program

Creators have the option to earn money with their content by participating in the YouTube partner program (YPP).⁹ This program mainly serves two purposes. For the platform, it provides a means of quality control. Its ads only run with videos uploaded by creators who are part of this program¹⁰ to ensure that they are not shown alongside inappropriate content. In addition, creators cannot freely join the program, but they have to fulfill certain criteria – which are the subject of this study – before they are eligible to apply for membership. As part of

⁶See https://www.statista.com/statistics/1201880/most-visited-websites-worldwide/, accessed on 13 February 2023

⁷See https://blog.youtube/press, accessed on 13 February 2023

⁸See https://www.youtube.com/premium, accessed on 13 February 2023

⁹See https://creatoracademy.youtube.com/page/lesson/ypp_what-is-ypp_video, accessed on 13 February 2023

¹⁰Following a change in the platform's policy in late 2020, YouTube now also holds the option to run ads with videos outside the program as well (Koetsier, 2020). This, however, has no implications for the present study as it falls outside our sample period.

the application process, they then undergo a (partly automated, partly manual) review process to ensure that their videos follow the platform's guidelines. In turn, for the creator, it provides a means to monetize their videos as YouTube shares part of the generated revenue. However, while creators have agency about *whether or not* and *how many* ads may be shown during a video, the platform determines *which* ads are actually shown via an algorithm (Kerkhof, 2022). Accordingly, advertisers and creators have no way of directly interacting with one another. In terms of the attractiveness of the YPP, anecdotal evidence – official statistics do not exist – suggests that creators can earn about three to five USD per 1,000 video views.¹¹ As a result, YouTube relies on two tiers of creators: First, creators in the YPP are vetted and they can earn money by allowing the platform to run ads before and during their videos. Second, creators outside the program cannot earn money, and no advertisements are shown with their videos. In addition, through the eligibility criteria and application process the platform limits the access to the program and actively controls who can move from the latter to the former tier.

The eligibility criteria for the YPP changed several times since its launch in 2007. While having been quite selective in the beginning, YouTube opened it up in 2012 by removing virtually all access barriers (YouTube Official Blog, 2012). However, this entailed an inflow of "bad actors", threatening the platform's integrity. In response, in early 2017, it put into place the restriction that creators have to have accumulated a minimum of 10,000 lifetime views before being able to apply (Popper, 2017). In this study, we analyze the subsequent rule change put into place in early 2018, which instituted an additional and much more significant increase in the eligibility criteria.

3.2 YouTube "Adpocalypse" and the Rule Change in 2018

We study a change to the eligibility criteria for the YPP that occurred in February 2018. This change was preceded by YouTube facing considerable backlash and, ultimately, a large-scale boycott by its advertisers (Nicas, 2017) – commonly referred to as the YouTube "Adpocalypse" (Alexander, 2018). Despite existing access requirements, advertisements routinely appeared alongside hate speech as well as racist and anti-Semitic content. In addition, YouTube faced criticism more broadly for its lack of restrictions as to what type of content is allowed on the platform (Statt, 2017). The situation was then further exacerbated by scandals surrounding two of the platform's most prominent creators (Gillespie, 2018), entailing further scrutiny. As part of this "Adpocalypse" and in reaction, the platform announced more manual curation of

¹¹See https://influencermarketinghub.com/how-much-do-Youtubers-make/, accessed on 13 February 2023

creators allowed to the YPP in December 2017 as an effort to ensure that advertisements are only shown alongside unproblematic content (Wojcicki, 2017). In January 2018, the platform went on to reveal new criteria determining the eligibility to apply to the program, which would take effect one month later, in February 2018 (Mohan and Kyncl, 2018).

This rule change contained two important elements: First, it updated the eligibility criteria. Now, to be able to apply for the YPP, creators had to have accumulated a minimum of 1,000 subscribers and 4,000 hours of "watchtime" over the preceding twelve months. The latter is calculated by multiplying the times videos are viewed with the amount of time each viewer actually spends with the videos. Together, this made access considerably more restrictive compared to the previous requirement of 10,000 lifetime views.¹² Second, those creators who had been part of the YPP, but did not meet the new criteria would be excluded from the program, effectively making it impossible for them to earn money on the platform. In the blog post announcing these changes, it was also noted that – while affecting a "significant number of channels" – the financial ramifications of this "demonetization" would be mild as "99 % of those affected were making less than 100 [USD] per year in the last year, with 90 % earning less than 2.50 [USD] in the last month" (Mohan and Kyncl, 2018). In addition, it was announced that those creators who lost access would get the possibility to reapply once they met the new criteria.

Still, YouTube faced negative reactions following this announcement, primarily from creators who were affected by the rule change (Alexander, 2018). Points of criticism are diverse. Some creators worried about the loss of earnings, while others felt treated unfairly, that is, they were not primarily concerned about the financial ramifications, but rather felt disappointed to lose the platform's endorsement after having been part of it for a long time. Others still considered it a sign that YouTube generally shifted its focus from smaller creators to larger, more prominent ones, regarding it as an indication that "the golden age of YouTube is over" (Alexander, 2019).

The rule change presents a well-suited opportunity to study how a platform's increased use of control mechanisms affects the supply of content. First, the primary reason for the platform to increase access requirements has been to make more manual curation feasible. In light of the sheer amount of videos that are uploaded to the platform, limiting the number of participating creators has arguably been necessary. Second, the rule change occurred in reaction to advertiser boycotts, which in turn had been sparked by problematic content, such as hate-speech by many

¹²For reference, a view is counted when a video is watched for at least 30 seconds. Accordingly, before the change the requirement of 10,000 views translates to a minimum of 83,33 hours of watchtime $(\frac{10,000}{60\times60})$, without any further restriction on the time span across which a creator may attain this goal.

small channels, but also controversies surrounding popular creators. However, *all* creators who did not meet the new requirements lost access to the platform. In other words, the majority did not become demonstrated due to their content being problematic, but simply because they did not meet the new criteria. We therefore study creators who were not the *cause* of the rule change, but who were very much *affected* by it.

4 Hypotheses Development

We study how the change of the access requirements for the YouTube partnership program affected the subsequent supply of content by those who lost access compared to those who got to remain. In particular, we develop hypotheses about how affected creators adjusted the extend of their *activity*, as well as the *quality* and *diversity* of their content. In addition, we explore two sources of creator heterogeneity, namely (i) between those who predominantly created *mainstream or niche* content, and (ii) along their *experience* on the platform before the rule change.

4.1 Effects on Creator Activity, Content Quality, and Diversity

When the rule change was communicated via the YouTube blog, it was not only announced that ineligible participants will be removed from the partner program, but they were also explicitly encouraged to reapply once they meet the new requirements. This bears opposing implications for the expected reactions from affected creators. On the one hand, the encouragement to reapply could have acted as a motivator. The new access requirements present a clear goal for creators to attain, which can spark an increased provision of effort to grow their channel to the necessary size and engagement. Consistent with this idea, existing research provides evidence that the pursuit of such goals can induce activity (Locke and Latham, 2002). Similar to our empirical context, Goes et al. (2016) have shown that volunteer contributors in an knowledgeexchange platform increased their effort to reach a higher "rank" in the community's status hierarchy, which they attribute to goal attainment incentives. Hence, it is possible that former participants of the YouTube partner program were subject to similar motivations after losing access, which implies an *increase* in effort provision afterwards.

On the other hand, however, there are several reasons to expect that affected creators were less motivated after losing access to the program. First, they lost the ability to monetize their content on the platform. Even if their earnings were low before the rule change, this removed any financial incentives that may have been a driving force behind their content production. Second, many creators are not only, or even mainly, driven by financial concerns. Rather, they follow non-pecuniary motivations, such as building a reputation or attaining status among their peers and audience (Toubia and Stephen, 2013), or their own "intrinsic" enjoyment of creating content (Shah, 2006). Losing access to the partner program deteriorates these motivational sources as well: Attaining the status of "YouTube partner" benefits their reputation and lets them form an attachment with the community and the platform, both of which are subsequently lost due to the rule change. In addition, the actions of the platform owner can send a signal about desirable content (Hukal et al., 2020), and the loss of the partner status may then be perceived as the withdrawal of its endorsement (Ho and Rai, 2017). Together, the rule change therefore deteriorated both financial and these non-pecuniary incentives, which implies a *decrease* in effort provision afterwards.¹³

Together, expectations about how the rule change affected subsequent effort provision are subject to a degree of ambiguity. Still, we expect the deterioration of financial and non-pecuniary incentives to dominate a potential motivation from re-attaining the "goal" of becoming YouTube partner. While rejoining the program can reinstate the ability to monetize content, it is not clear that it would also restore non-pecuniary benefits. For one, the potential status may not carry the same reputational benefits after it had been lost before. And if the removal from the program is perceived as a signal that the platform does not endorse or value their content, then creators may lack the intrinsic motivation to further put effort into its support. In addition, because the earnings potential among affected creators is relatively low, the potential of their re-attainment is unlikely to be a sufficient motivator to outweigh the detrimental deterioration of the non-pecuniary incentives. Therefore, we expect the average creator to reduce their effort in producing content in reaction to losing access to the partner program. This can manifest in two ways: First, they can reduce the frequency at which they provide content, that is, they produce and upload fewer videos to YouTube. And second, they may put less effort into the production value of their videos, entailing reduced quality and worse audience reception. Therefore, we hypothesize:

Hypothesis 1. (Activity) After the rule change, the frequency of video uploads is <u>lower</u> for creators who lost access to the partner program, compared to those who remained.

Hypothesis 2. (Quality) After the rule change, the quality of videos is lower for creators who

 $^{^{13}}$ In principle, a third possibility exists: If creators are purely intrinsically motivated – that is not driven by, say, financial or reputational concerns – they may not react to the rule change at all. In our empirical analysis, we would then not find significant effects.

lost access to the partner program, compared to those who remained.

In addition to adjusting the amount and quality of content creators produce, losing access to the partner project may also change how they approach the type of content they produce. In particular, we are interested in content diversity, that is, how many different topics, subject areas, or genres they cover on their channel. How broad or narrow the scope of the covered topics should be is an important aspect of creators' content strategy, and we expect this decision to be shaped by the incentives put into place by the platform owner. Consistent with this notion, some prior research cautions a inefficient duplication of mainstream content when creators receive a share of the ad revenue (Sun and Zhu, 2013; Anderson and Gabszewicz, 2006; Wilbur, 2008), unless the competitive pressure is to great in this segment (Kerkhof, 2022).¹⁴ Likewise, we posit that pursuing a content strategy of greater diversity also reflects an effort to increase the potential audience: covering a greater range of topics bears the potential to match a broader range of consumer tastes.

Similar to our discussion above, this implies opposing expectations about how losing access to the partnership program affects creators' subsequent content diversity. On the one hand, deteriorating financial and non-pecuniary (e.g. reputation-based) incentives implies reduced creator effort to maximize their audience. In addition, prior research suggests that such extrinsic sources of motivation can crowd out intrinsic sources of motivation (Khern-am nuai et al., 2018; Toubia and Stephen, 2013), which implies that creators would focus on producing content they personally enjoy. Both reasons suggest *reduced* subsequent content diversity. On the other hand, if creators primarily strive to regain access as quick as possible, they will increase their content diversity to attract a broader audience. In addition, they could also perceive losing access as negative performance feedback from the platform. This, in turn, can trigger "problemistic search" (Posen et al., 2018), that is, they start experimenting with different approaches to their content strategy. Both suggest that creators *increase* their subsequent content diversity.

Again, it is ultimately an empirical question which of the two factors is dominant. Here, we follow the same logic as before and expect the effect of deteriorating incentives to weigh more heavily than efforts to regain access to the partner program, and we therefore hypothesize:

Hypothesis 3. (*Diversity*) After the rule change, the diversity of content is <u>lower</u> for creators who lost access to the partner program, compared to those who remained.

¹⁴We also investigate creators' subsequent tendency to produce mainstream content as a piece of additional analysis, but do not find an effect.

4.2 Heterogeneous Effects

4.2.1 Mainstream and Niche Creators

Prior research notes that relying on ad-revenue can create incentives to predominantly produce popular content in an effort to maximize the audience (e.g. Wilbur, 2008; Sun and Zhu, 2013). We take this as a point of departure to discuss one dimension of creator heterogeneity which may entail differences in their reaction to losing access to the partner program. In particular, we explore differences between those creators who predominantly produced mainstream content, and those that rather positioned themselves in niche segments. Mainstream creators focused on producing popular content before the rule change, which we take as an indication that they are relatively more driven by extrinsic sources of motivation, such as earnings or reputational concerns. In contrast, niche creators deliberately position themselves in segments of relatively lower demand (Zhu and Zhang, 2010). Compared to their mainstream counterpart, we therefore consider them to be relatively more driven by intrinsic motivations, that is, they produce content they personally value and enjoy. Hence, comparing the two allows us to provide an empirical test about the relative importance of extrinsic and intrinsic motivations in driving their reaction to losing access to the partner program.

As pointed out before, the rule change likely deteriorated incentives to produce content. This affected both extrinsic motivations through the loss of earnings and the partner status, as well as intrinsic motivations if the rule change is perceived as a withdrawal of the platform's appreciation (Ho and Rai, 2017) or a signal that the content is not valued anymore (Hukal et al., 2020). The main difference between the expected reaction between mainstream and niche creators then stems from their incentives to try and regain access to the program. Here, we posit that mainstream creators have a better opportunity to recover some of the motivational drivers by rejoining the program. It restores their ability to monetize their content, hence generating financial incentives. In addition, while unclear how much reputational benefit it still carries after it has been lost once, it still restores their status of YouTube partner. In contrast, it is unlikely that niche creators can similarly recover their intrinsic motivations after having experienced the loss of platform endorsement, which sends a signal that their content is not valued. Together, while we still expect an overall negative effect of losing access to the program on activity, quality, and diversity, we expect it to be stronger for niche than mainstream creators. We therefore hypothesize:

Hypothesis 4. (Mainstream/Niche) The effects of losing access to the partner program on

(a) the frequency of video uploads, (b) content quality, and (c) content diversity is stronger for niche than mainstream creators.

4.2.2 Creator Experience

Creators' sources of motivation have been shown to vary by their experience (e.g. Loh and Kretschmer, 2023; Panciera et al., 2009), which is why we expect heterogeneity in the effect of losing access to the partner program along this dimension. In the context of volunteer or usergenerated content, prior research argues that creators can form a bond with the community of their peers and also create a common identity attached to their platform (Ren et al., 2007). Because less experienced creators have not had the opportunity to engage with the community or gain recognition through their platform activity, their degree of self-identification as "YouTubers" should be lower than for more experienced creators. Consistent with this, some research on community awards has shown that new creators react more strongly to external signals of approval and recognition, in terms of creating novel content on Reddit (Burtch et al., 2022) and in the retention of new contributors on Wikipedia (Gallus, 2017). However, these studies investigate a reputational *qain* for those creators who did not have a high status beforehand. In contrast, the rule change we study constitutes a loss of reputation because they no longer enjoy the status of YouTube partner. Therefore, we expect the opposite effect here, namely that more experienced creators will display a stronger reaction. The reason is two-fold: First, their greater cumulative effort before the rule change is indicative of their attachment to and identification with the platform. Therefore, we expect them to place higher value on the recognition from having been accepted into the partner program in the first place. And second, receiving the negative signal about the desirability of their content from the platform is likely stronger for them than for less experienced creators precisely because their relatively greater past efforts are unappreciated. Therefore, we expect that the deterioration of these non-pecuniary motivations are more severe for more experienced creators. However, like before, we may also expect that these are reasons for them to attempt to regain access to the program in an effort to restore these benefits. But, in keeping with our prior arguments, we still expect the incentive deterioration to weigh more heavily, and hypothesize:

Hypothesis 5. (*Experience*) The effects of losing access to the partner program on (a) the frequency of video uploads, (b) content quality, and (c) content diversity is <u>stronger</u> for more experienced creators.

5 Data and Methods

5.1 Data Set

To analyze how the supply of videos changed for creators who lost access to the YPP after the rule change, we combine information from two waves of data collection via the YouTube Data API. First, we use the same snapshot as Kerkhof (2022) who has obtained information about all active German YouTube channels as of December 2017 – that is just before the rule change –, including whether or not they have participated in the YPP at that point in time. This piece of information is unique and crucial to our analysis, as it is impossible to assess historic information on a creator's program participation otherwise. From that snapshot we select all creators who were part of the YPP and had between 500 and 5,000 subscribers by the end of 2017^{15} , that is whom we consider "at risk" of losing access to the program.

For the second wave, we accessed the YouTube Data API from September to November 2020 to obtain a snapshot containing updated information for the selected sample of creators, which lets us track their upload history since January 2018. Combining the two snapshots provides us with crucial information for the construction of our regression samples and key variables. Specifically, we obtained cross-sectional information at the creator level, such as their subscriber count in December 2017 (first snapshot) and November 2020 (second snapshot) and their total number of videos. In addition, we collected information at the video-level, such as the number of views, likes, dislikes, duration, date of upload, keywords, and the video category. This lets us track each creator's video uploads over time, and provides us with information about the extent of their activity (e.g. upload frequency), as well as if and how their content strategy has changed over time.

One shortcoming is that a creator's watchtime over the past twelve months is not directly provided via the API. However, this measure is crucial for our analysis as it is part of the new eligibility criteria instituted by the rule change under study. Since the measure is not publicly available, we compute it ourselves using the length of all videos that a creator has uploaded in the twelve months before the rule change and the number of views that these videos have accumulated. However, viewers often do not watch the entire video. For example, Maggi et al. (2018) find that the least popular videos are watched to about 50% on average, and the most

¹⁵The distribution of subscribers over creators is heavily skewed; in other words, there are many more creators with few than creators with many subscribers (the median number of subscribers in Kerkhof (2022) is around 1,000). Thus, to ensure that we have a sufficient number of observations in our main analysis, we decided to use a relatively large initial bandwidth regarding the upper subscriber bound.

popular ones are watched to about 75% on average. Since our sample consists of relatively unknown creators with small channels, we conservatively presume in our main analysis that 50% of each video is watched. We then simply multiply each video's duration by 0.5 and subsequently take the sum over all video views in the twelve months before the rule change in February 2018. Note that not all selected creators appear in the second snapshot. This is the case if they exited the platform between December 2017 and November 2020. We further discuss the characteristics of creators who exited the platform below.

5.2 Empirical Framework

5.2.1 Identification Strategy

We want to estimate the causal effect of losing access to the partner program on subsequent creator activity. As this is determined by their watchtime and subscriber counts at the time of the rule change, we face the challenge of separating the effect from unobserved creator characteristics that may otherwise drive our estimates. Specifically, those who are better able to produce engaging content will have more subscribers and more watchtime (hence they do not lose access to the program), and they may be less inclined to change their video supply. At the same time, the clearly defined subscriber and watchtime thresholds provide us with a quasi-experimental setting to implement a regression discontinuity design to estimate the causal effect: For creators very close to these thresholds the treatment of losing access to the program can be considered to be as good as random. In other words, those of similar quality likely still exhibit small subscriber count and watchtime differences, and we can attribute differences in behavior after the rule change to losing or retaining access to the partner program.

In principle, there are two running variables determining treatment status: subscriber count and watchtime. In our main analysis, we focus on the subscriber threshold for two reasons. First, the definition of watchtime is not always clear to creators, and the metric is not as salient as their subscriber count.¹⁶ As a consequence, we do not expect that being just below the watchtime threshold after the rule change has an equivalently large impact on creators' behavior than being just below the subscriber threshold. Second, in contrast to the subscriber count, we can only approximate watchtime and measure it with noise. Hence, we consider creators just above and just below the subscriber threshold to identify the causal effect of losing access to the YPP and – as proposed by (Papay et al., 2010) – include the second running variable, watchtime,

¹⁶See https://www.qqtube.com/article/youtube-retention-vs-watch-time, accessed on 13 February 2023).

as a control. Our key identifying assumption is that the two groups are comparable within a reasonably narrow bandwidth, and that subsequent differences in behavior can be attributed to the rule change (Lee and Lemieux, 2010; Imbens and Lemieux, 2008). Specifically, for our main analysis, we select creators who have at least 900, but no more than 1,200 subscribers (we provide robustness checks using both a narrower and wider bandwidth in section 6.2). As there are more creators just below the 1,000 subscriber threshold than above, this selection ensures that we analyze comparable creators that are sufficiently close to the threshold, while also maintaining a reasonable sample size that is balanced between treatment and control groups. We can then obtain an unbiased estimate for the effect of losing access by estimating the following equation:

$$Y_i = \alpha + f(Subscribers_i) + \beta \cdot LostAccess_i + \gamma \cdot Watchtime_i + \epsilon_i, \tag{1}$$

where Y_i is the outcome of interest, $LostAccess_i$ is a dummy variable indicating whether creator i lost access to the YPP due to an insufficient number of subscribers, and $Subscribers_i$ is our running variable, the subscriber count. $Watchtime_i$ controls for creators' watchtime in the twelve months before the rule change. The coefficient of interest in equation (1) is β , which gives us the local average treatment effect (LATE), that is, the effect of losing access to the YPP, for each regression (Imbens and Angrist, 1994).

The function $f(\cdot)$ captures the underlying relationship between the subscriber count and our outcomes of interest; in particular, we implement a local linear regression approach by letting $f(\cdot)$ be a linear function of subscribers.¹⁷ In addition, we let the slopes of our fitted lines differ on each side of the subscriber threshold by interacting $f(\cdot)$ with *LostAccess*_i to control for differential trends in *Subscribers*_i. Following Calonico et al. (2020) we use a triangular Kernel function which assigns zero weight to all observations outside of our specified bandwidth, and positive weights to all observations within our bandwidth. The weight is maximized at the threshold, and declines symmetrically and linearly going away from the threshold.

Finally, a potential threat to identification is that creators very close to the threshold may manipulate their subscriber count to remain in the YPP, known as bunching. Therefore, we implement a "donut" by excluding creators with a subscriber count between 990 and 1010 or watchtime between 3950 and 4050 hours. In addition, we perform a test for continuity in the subscriber count distribution around the threshold in section 5.4.

We use the specification described in equation (1) for the entirety of our analysis, but use different outcomes of interests to investigate different aspects of creator behavior. For each, we

 $^{^{17}\}mathrm{We}$ provide robustness checks using quadratic and cubic model fits in section 6.2

perform the analysis both before and after the rule change. Differences in behavior between treated and untreated should only exist after the rule change, but not before. Hence, if the coefficient of interest β is insignificant in the before period, but significant after, we can attribute it to the rule change.

5.2.2 Outcome variables

We consider three main outcome variables: creator activity, the quality of their content, and its diversity. To test hypothesis 1, we measure creator activity in terms of their upload frequency, that is, the average number of monthly video uploads in the six months before and after the rule change.¹⁸

In hypothesis 2, we make predictions about creators' subsequent content quality. To measure this, we leverage information about viewer engagement in terms of "likes" and "dislikes" videos receive. Specifically, we compute the share of likes over all viewer reactions, $\frac{\text{Likes}}{\text{Likes+Dislikes}}$, for each video. Based on that, we derive a creator's average quality of content in the six months before and after the rule change.

Third, hypothesis 3 is about the extent creators who lost access to the YPP change the diversity of their videos compared to those who did not. To this end, we consider the videos' keywords, which are illustrative terms that creators assign to their videos to let YouTube know what the video is about. For example, a funny cat video might be given the keywords "funny", "cat", and "pet". To measure content diversity, we compute a creator's average monthly number of unique keywords in the six months before and after the rule change. A larger number of unique keywords then indicates more diversity in a creators' coverage of different topic, which in turn shows increased experimentation with different types of content, or attempts to appeal to a broader audience that exhibits a wider range of tastes for horizontal content attributes.

5.2.3 Creator heterogeneity

We study creator heterogeneity along two dimensions: First, in hypothesis 4 we formulate expectations about heterogenous reactions between mainstream and niche creators. To distinguish between the two types, we adopt the mainstream measure from Kerkhof (2022). Specifically, for each month and video category, keywords are ranked by how many video views they attract, that is by their popularity. The upper one percent of keywords in this distribution is then classified as "mainstream". Based on that, we classify videos who exhibit at least one such

 $^{^{18}}$ Section 6.2 shows that our results are robust to using alternative time windows for all three outcome variables.

keyword as "mainstream content". Finally, we compute a creator's proportion of "mainstream content" in the six months before and after the rule change. Since we find that the vast majority of creators exclusively uploads mainstream content, we classify a creator as "niche" if her share of mainstream videos is smaller than one, and as "mainstream" otherwise.

Second, hypothesis 5 is about creator experience. To measure this, we consider the date of her first video upload and compute the age of her channel in terms of how many months she has been active on the platform until the rule change. As a robustness check, we also measure creator experience in terms of the number of video uploads before the rule change.

5.3 Summary Statistics

Table 1 shows summary statistics for our main estimation sample of 484 creators who fall within the subscriber count bandwidth that we use (900 to 1200) and who remained active on the platform after the rule change. The majority (80 %) are below the threshold and lost access to the program. This is the case despite our asynchronous bandwidth selection and shows the "long tail" distribution in subscriber count - as it is the case in most media markets, most content creators are relatively unsuccessful (Anderson, 2004). The average creator had 1026.43 subscribers at the time of the rule change and accumulated about 5307.68 hours of watchtime in the twelve months before February 2018. Within the subsequent six months, the average creator uploaded an average of 3.248 videos per month, using a monthly average of 41.62 unique keywords, and received an average monthly like share of 91 %. The high share of likes is likely due to rating inflation, which is common on digital platforms (Zervas et al., 2021). Moreover, the average creator has uploaded a total of 128.88 videos to her channel before the rule change was implemented. With 60 % the majority of creators in our main sample are classified as mainstream. This is not surprising, as - by definition - this is the most popular and prevalent type of content on the platform. Finally, the average creator has been active for about 35.57 months – that is, for nearly three years – before the rule change.

Crucially, out of the total number of creators that appeared in the first wave of data collection, roughly 46 % had exited the platform by the time of the second wave. Although it is interesting to compare the characteristics of creators who have and have not exited the platform between the two snapshots in 2017 and 2020, we must interpret any differences with care. In particular, it is not clear that these creators exited as a result of the rule change or because they were less successful than the creators who remained active. We show differences between exiting and non-exiting creators in table 2. In fact, we observe that creators who exited had even more subscribers and watchtime on average than creators who stayed; similarly, they had on average a higher like share and a smaller proportion lost access to the YPP. However, given that the creators whom we consider are most likely to operate their channel as a hobby in their free time, it is more plausible to assume that many of them simply lost time or interest in participating on the platform. As our main objective is to study creator behavior *after* the rule change, we focus on creators who stayed. In addition, we find differences between the groups in their upload frequency and the use of unique keywords. To account for potential selection into remaining on the platform, we estimate all regression with a Heckman two-stage procedure as a robustness check. In each case, the resulting inverse Mills Ratio is statistically insignificant and our main estimates remain unchanged, indicating that selection issues are no concern.

=== Table 1 here ===

=== Table 2 here ===

5.4 Test for Quasi-Random Assignment

The main identifying assumption of our empirical approach is that losing access to the partnership program is as good as random within the specified subscriber count bandwidth (see e.g. Flammer, 2015). In other words, we assume that creators just above and just below the 1,000 subscribers threshold are similar in all observed and unobserved characteristics except continued access to the YPP after February 2018. We perform two tests for the validity of this assumption.

First, we show that the distribution of subscriber counts is continuous around the threshold. If we would detect a discontinuity at the threshold, this would indicate that assignment to the treatment is in fact not as good as random. For instance, it may be that creators who had been below the threshold before the rule change show efforts to increase their subscriber count to not lose access to the partner program by the time of the rule change. Following Cattaneo et al. (2017), we conduct an automatic manipulation test which does not reject the null of continuity around the threshold (p = 0.99; Figure 1 visualizes the test). Thus, we find no evidence for a violation of the continuity assumption in our sample.

$$===$$
 Figure 1 here $===$

Second, we check if creators just above and just below the threshold show differences in their behavior even before the rule change. In this case, any observed differences in our main outcomes after the rule change may not be a result of lost access to the YPP, but rather of underlying differences between the two groups of creators. To test this, we estimate equation (1) exclusively based on observations from the six months before the rule change. Table 3 shows the results. We do not find any statistically significant differences in behavior between creators just above and just below the subscriber threshold before the rule change. Hence, it is plausible to assume that any differences in our outcome variable that we observe afterwards can be attributed to losing access to the YPP.

=== Table 3 here ===

6 Results

6.1 Main Analysis

In hypothesis 1 we argue that losing access to the partner program causes a decrease in creator activity. We test our expectation in model 1 of table 4 and find evidence to support the hypothesis. We find that, each month, creators who lost access to the partner program uploaded significantly fewer videos to the platform compared to those that did not lose access ($\beta = -2.816$, p < 0.05). Specifically, they uploaded 2.816 fewer videos, which – considering the sample mean of 3.248 – is a sizable effect. We show the accompanying RDD plot in figure 2a. These provide evidence that the rule change indeed had a deteriorating effect on affected creators' incentives to provide effort in producing content for the platform. Further, it shows that this negative effect, on average, outweighs a potential motivational nudge that arises from creators attempting to regain access to the program.

$$===$$
 Table 4 here $===$

Next, we formulated similar expectations in hypothesis 2, namely that creators who lost access would also decrease the quality of the content – an additional symptom of reduced effort provision. We test this in model 1 of table 5, using the average like share as dependent variable. Consistent with our hypothesis, we find significantly lower content quality among affected creators than those who remained in the partner program ($\beta = -0.048$, p < 0.1). Again, this is visualized in figure 2b. Specifically, the average like share is 4.8 percentage points lower for the affected creators. While this may not seem much considering the relatively high sample mean of 91 %, we do have to consider that ratings on YouTube are concentrated in the higher value ranges. In light of this, our estimate suggests a meaningful difference in content quality between treated and untreated creators after the rule change. In addition, this result adds further evidence that the change deteriorated incentives to create content on the platform, which therefore manifests in both subsequent lower activity and quality.

$$===$$
 Table 5 here $===$

We are also interested in how the rule change affected creators' content strategy, and in particular its diversity. In hypothesis 3, we argue that it would cause a decrease in content diversity, and we test this expectation in model 1 of table 2c. We find that creators who lost access to the partner program exhibit a lower subsequent degree of content diversity compared to those who remained ($\beta = -25.38$, p < 0.05). We also show the RDD plot in figure 2c. On average, affected creators use 25.38 fewer unique keywords per month to describe their videos. Again, considering the sample mean of 41.62, this implies a sizeable reduction in the range of topics creators cover on their channels. In addition, this provides support for our hypothesis and indicates that deteriorating incentives reduce creators' attempts to satisfy a broad range of viewer tastes to maximize their audience. Instead, they adjust their strategy towards a more focused approach, perhaps because they now follow more intrinsic sources of motivation. In other words, being "free" of financial or reputational concerns, the rule change enabled them to produce more content they themselves enjoy.

=== Table 6 here ===

=== Figure 2 here ===

In a next step, we investigate potential sources of heterogeneity in the effect of losing access to the program. First, in hypothesis 4, we expect the reaction to the rule change to differ between mainstream and niche creators. In particular, we hypothesize the negative effects on (a) activity, (b) content quality, and (c) content diversity to be stronger for niche than mainstream creators. Results for activity are reported in models 2 and 3 of table 4, in which we split the sample between the two types. We only find a statistically significant effect for mainstream creators. In addition, the estimated coefficient is considerably larger than for niche creators. This goes against our expectation that mainstream creators' greater motivation to regain access to the partner program partly offsets the overall negative effect from deteriorating incentives. Instead, these results suggest that taking away extrinsic benefits has a stronger effect on those creators who originally positioned themselves in the most popular segments. In addition, niche creators intrinsic motivation had been the primary driver of their activity all along, and that the rule change did not affect this. Hence, we do not find support for hypothesis 4a.

We test hypothesis 4b in models 2 and 3 of table 5, again splitting the sample between mainstream and niche, but this time using average like shares as the dependent variable. We again only find a statistically significant effect for mainstream (model 2, $\beta = -0.061$, p < 0.05), but not niche creators (model 3, $\beta = -0.034$, p > 0.1). Hence, we do not find support for hypothesis 4b either. Still, this finding produces interesting insights about underlying motivations: Even after losing access to the partner program, niche creators are unwilling to compromise the quality of their content. A likely explanation is that they are indeed driven by intrinsic motivation and seek to produce content they themselves enjoy, even after the platform withdrew its endorsement.

Next, hypothesis 4c is tested in models 2 and 3 of table 6, using a similar sample split and investigating creators' use of unique keywords to describe their content. As expected, we find that niche creators exhibit a greater reduction in diversity than mainstream creators after the rule change. The former uses an average of 33.294, and the latter of 19.803 fewer keywords to describe their content after the rule change. At the same time, the coefficient for niche creators is estimated at reduced precision and therefore just statistically insignificant (p = 0.12). This is likely the result of the smaller sample size in model 3. Together, we still take this as support for hypothesis 4c. While both types reduce the diversity of their content, niche creators do so to a larger extent.

Together, hypothesis 4 only receives limited support. However, these findings still paint a nuanced pattern in the heterogeneous reactions between mainstream and niche creators, which we believe to be consistent with intrinsic motivations. Compared to mainstream, niche creators are unwilling to compromise the quality of their content. But at the same time, they greatly reduce the scope of topics they cover. Together, this suggests that they rather hone into the areas they are personally excited about, which motivates them to keep quality up. In contrast, we find patterns that are consistent with deteriorating extrinsic motivations for mainstream creators, who reduce their activity, content quality, and – albeit to a lower extent – content diversity.

In a final string of analyses, we are interested in potential heterogeneous effects along the dimension of pre-change experience. In hypotheses 5, we formulated the expectations that the negative effect of losing access should be stronger for the more experienced than less experienced creators. Similar to before, we therefore perform a series of sample splits at the median of the experience distribution. We test hypothesis 5a, which is about heterogeneity in the activity effect, in models 4 and 5 of table 4, and find strong support. Only more experienced creators (model 4, $\beta = -3.82$, p < 0.1) creators reduce the frequency at which they upload videos to the platform after the rule change. In contrast, we find no evidence for an effect for less experienced (model 5, $\beta = -0.841$, p > 0.1). This suggests that the activity effect is completely driven by more experienced creators, and provides evidence for the relevance of non-pecuniary sources of motivation. We argued that, while financial motivations are unlikely to be a function of prior efforts, more experienced creators had formed a deeper identity-based attachment to the platform. In addition, they had spent a longer time providing effort in the past, which should make a signal of lost valuation and endorsement from the platform all the more impactful.

We similarly test hypothesis 5b in models 4 and 5 of table 5, and find no support. We only find a significant negative effect on content quality for less experienced creators (model 5, $\beta = -0.045$, p < 0.1). However, the size of the estimated coefficient is similar to the estimate for more experienced creators (model 4, $\beta = -0.048$, p > 0.1). While this suggests that both types tended to decrease the quality of their content, we do not find evidence for differences in effect sizes between them.

And finally, we study heterogeneity in the effect on content diversity, finding strong support for hypothesis 5c. While more experienced creators show a sizeable and statistically significant reduction in the diversity of their content (model 4, $\beta = -34.082$, p < 0.05), we do not find an effect for less experienced creators (model 5, $\beta = -8.393$, p > 0.1). Again, this demonstrates the importance of non-pecuniary motivations, which in this case drive more experienced creators to reduce the scope of the topics they cover.

In all, we find support for hypothesis 5, at least in terms of creator activity and content diversity. Our results suggest that both effects are completely driven by more experienced creators, who likely formed a stronger identity as "YouTuber" and who had put a lot of effort into producing content before the rule change. Therefore, withdrawn platform support and the potential reputation decline weighs more heavily for them than for less experienced creators.

6.2 Robustness Checks

We run a series of regressions to test the robustness of our results to alternative modeling and variable choices. First, we use a bandwidth of 900 to 1200 subscribers in our running variable. There is a trade-off in choosing this range: The wider the bandwidth, the less comparable become observations in the treated and untreated groups. But the narrower, the smaller the sample becomes, which may entail reduced statistical power. We test the sensitivity of our results to different ranges in table 7. We use a wider range in models 1, 3, and 5, and a narrower range in models 2, 4, and 6. For the latter, as expected, standard errors increase due to the smaller sample, yielding statistically insignificant estimates for all outcomes. However, coefficient sizes are very similar to our main specification, which shows that the sample composition does not seem to be affected much. For the former, the coefficient sizes become considerable smaller compared to our main specification, and – with the exception of model 5 – they become statistically insignificant. Still, they show the same sign as in our main specification. Hence, while we do not believe that these tests are a serious cause for concern about the validity of our modeling choice – especially given our small sample sizes –, we do determine that the precision of our estimates declines with different bandwidth choices.

=== Table 7 here ===

Second, we use local linear regressions in our main analysis, following recommendations from the literature (Gelman and Imbens, 2019). Still, we test the robustness of our results to using higher order polynomials, which we report in table 8. Specifically, we fit $f(Subscribers_i)$ in equation 1 with quadratic terms in models 1, 3, and 5, and we use cubic terms in models 2, 4, and 6. Across the board, standard errors are increased, rendering the results statistically insignificant. At the same time, however, coefficient sizes tend to be larger than in our main specification, and they continue showing the same sign. Therefore, we consider results qualitatively robust, which is also supported visually by an RDD plot of the quadratic model fit, which we show in figure 3 (the cubic model fit is also consistent). In addition, the reduced precision of the estimates can likely be attributed to the increased statistical power required to estimate more complex regression models with reduced sample sizes (see e.g. Gelman, 2018). In all, we are therefore not concerned about the validity of our choice of model fit, but we do determine that the small sample size imposes certain limitations in terms of statistical power and the precision of our estimates.

=== Table 8 here ===

=== Figure 3 here ===

Third, our outcomes of interest are calculated based on creator behavior in the six months after the rule change in our main specification. Here, we test the robustness to using alternative time windows of three and twelve months. Results are reported in table 9. They are largely consistent with our main specification. For creator activity (models 1 and 2) and content diversity (models 5 and 6), effect sizes are slightly larger when using a shorter window, and slightly smaller when using a longer window. This may suggest that reactions in these dimensions manifest in the relative short run after the rule change. However, the pattern is reversed for content quality. Here, the coefficient is slightly larger than in our main specification when using a larger time window, and smaller and statistically insignificant when using a shorter one. This may suggest that adjustments to quality only unfold over time. Together, however, these patterns do not cause concerns about our choices in measuring our outcomes.

=== Table 9 here ===

And finally, we measure creator experience in terms of their time spent on the platform prior to the rule change. Here, we use an alternative measure in their number of videos they had uploaded at that time. The results are reported in table 10. They are qualitatively consistent with our main specification. Still, two difference are of note: First, the difference between more and less experienced creators becomes more pronounced in terms of how their content quality changes (models 3 and 4). And second, in contrast to our main specification, we do find a negative effect on content diversity for less experienced creators here. However, this effect is still considerably smaller than for the more experienced, which is consistent with hypothesis 5c. In all, our main results are therefore robust to this alternative measure.

=== Table 10 here ===

7 Discussion and Conclusion

We study how an increase in the eligibility criteria to the YouTube partner program affected the supply of videos from those creators who lost access to it as a result. Such partner programs are a useful governance tool for platforms: They create incentives to supply high-quality, while also letting them exert some control over creators' activities to prevent "bad faith" and low-quality content. In our setting, following widespread criticism and advertiser boycotts, YouTube significantly increased the criteria to be eligible for its partner program, which made it harder for new creators to participate. In addition, it removed all former participants who did not meet the new requirements. We empirically analyze their reaction to this rule change and specifically investigate how this had an effect on their activity as well as the quality and diversity of the content they create on the platform.

We use a regression discontinuity design to estimate the causal effect of losing access to the program on subsequent creator behavior. The new criteria provide us with a clear threshold in their subscriber count at the time of the rule change. Hence, we compare the supply of videos between those who were just below that threshold to those just above. In our empirical analysis of German creators we find that those who lost access to the partner program significantly reduced the frequency at which they uploaded new videos to the platform. In addition, both the quality and diversity of their content decreased. Together, these effects speak to a deterioration of incentives due to the rule change. Further, we find that mainstream creators showed a stronger negative reaction than niche creators in terms of their activity and content quality, but not diversity. We attribute this to the relative importance of extrinsic and intrinsic motivations: Because niche creators enjoy producing content they themselves like, they are unwilling to compromise its amount and quality. At the same time, they focus on a more narrow range of topics after the change. In contrast, because mainstream creators are likely to be relatively more driven by financial and reputational concerns (given their ex-ante positioning in popular segments), the deterioration of these extrinsic motivators causes their adverse reaction to the rule change. And finally, we find that only more, and not less experienced creators show a reaction to losing access to the partner program. Because financial considerations are unlikely to be a function of experience, we attribute this to identity- and attachment-related motivations. More experienced creators had spent more effort on the platform in the past, and they had the opportunity to form a connection to the platform. Hence, they are more likely to identify as "YouTubers". Then, losing the partner status and receiving a signal that the content is not valued by the platform explains the negative reaction only by more, but not less experienced creators.

Our results generate useful insights for managers of platform ecosystems. While revenue sharing programs can be a useful tool to motivate high-quality content creation, changing the conditions under which creators can participate should be done with caution. Because creators are not only (or even mainly) motivated by financial considerations, such changes can affect their platform support in ways that are hard to predict and potentially adverse. In our setting, YouTube clearly tried to exercise caution by attempting to explain the rationale for their rule change (Wojcicki, 2017) and by highlighting the possibility to reapply (Mohan and Kyncl, 2018). Nevertheless, the reactions we uncover clearly show decreased motivations by affected creators. Further, our study also highlights a challenge inherent in large ecosystems: Because it is infeasible to tailor rules to individual creators, most changes will be applied indiscriminately among the whole population. In our setting, the issues of problematic content were caused by a small subset of "bad faith" actors. However, this still necessitated wholesale changes to access requirements for the partner program. As a result, all creators were affected, which then entailed adverse reactions by otherwise (potentially) valuable ecosystem participants.

Our findings contain contributions to two streams of literature. First, we provide novel insights about changes to a platform's governance strategy, and shifts towards more control in particular. Prior studies show that exerting control is necessary to secure a set of highquality and innovative complements (Boudreau, 2012; Parker and Van Alstyne, 2018) or to avoid uncontrolled creativity (Geva et al., 2019) and threats to a platform's integrity (Eaton et al., 2015). We add to this discussion by highlighting partnership programs as one governance tool platforms can use to exert more control over complementors' activities, without shutting them out of the platform completely. In particular, we show how an increase in control (by making access to the program more restrictive) can have detrimental effects on the motivation of creators to subsequently exert effort in supporting the platform. Further, recent studies note that it is hard to predict complementor reactions to rule changes more generally (e.g. Jhaver et al., 2018; Koo and Eesley, 2020), due to heterogeneous needs and characteristics (Boudreau and Hagiu, 2009). This is confirmed by our results, which show that different complementor types draw on different motivational sources, which explains heterogeneity in their reaction to losing access to the partner program. Together, our contribution to this stream is therefore twofold: We provide insights about how platform partnership programs affect multifaceted content supply motivations, which are important determinants of the effectiveness of such governance attempts. Further, while previous studies have considered ecosystem-level compositional implications of restricting access to the platform, we provide evidence about complementor-level effects of *changing* levels of platform control. Hence, we add novel insights to the discussion around unintended consequences of governance attempts (Tiwana, 2015; Gawer and Henderson, 2007). Finally, we provide evidence on how complementor-level heterogeneity in their content strategy and experience determines their reaction to governance attempts, which drives effective management of platform ecosystems.

Second, we contribute to the literature on ad-based platform business models that rely on external contributions from independent creators for value creation. Prior studies have shown how sharing part of the revenues with creators affects their supply of content (Sun and Zhu, 2013; Tang et al., 2012). We contribute to this discussion in two ways. While prior studies have focused on the aggregate effects of financial incentives, we discuss and provide evidence for the role of creator heterogeneity in determining the effectiveness of revenue-sharing schemes. In so doing, we shed light on how complex financial and non-pecuniary (such as status- or identitybased) sources of motivation drive the effectiveness of these schemes. In addition, we study the effects of losing, rather than gaining, access to these sources of motivation. In particular under unclear and diverse sources of motivations, this is an important distinction. For example, while financial incentives may be regained after they have been once lost, it is not clear that creators are able to recover status benefits in a similar way. Hence, the heterogeneous reactions we uncover imply that the financial loss is not a sufficient explanation. And finally, prior work has studied creators' positioning within popular or niche segments as an outcome of interest. Instead, we study their ex-ante positioning as a determinant of their reactions to the rule change, which we identify as an important indication of the motivational sources they draw from.

Finally, our study contains several limitations and point towards future research opportunities. First, we study a small subset of creators in a large platform ecosystem. To identify causal effects of losing access to the YouTube partner program we focus on those who fall within a narrow range of subscriber counts. In addition, these creators are relatively small and unlikely to (individually) contribute a lot of value to the platform. Therefore, they may not be representative of the entire ecosystem. For instance, larger, highly successful creators may enjoy a stronger bargaining position vis-à-vis the platform, eliciting a different reaction to a rule change similar to the one we study. Therefore, future research could look at heterogeneity along the dimension of creator size or success. Second, we are unable to clearly separate financial and non-pecuniary incentives in our study. Instead, we approximate this distinction by investigating heterogeneous reactions between different creator types. Hence, future research could make this distinction explicitly, perhaps in the form of an experiment. Lastly, lessons learned from studying YouTube may not perfectly translate to other empirical contexts. For one, the ecosystem offers both professional and user-generated content. Hence, reactions to increased control may differ in other platform settings, such as app stores, video game consoles, or e-commerce. Moreover, YouTube enjoys a near-monopolistic market position, which may allow it to implement more restrictive rules without many repercussions. Creator reactions could therefore be even stronger in the case of platforms that face rivals, which could provide attractive outside options for them. And finally, the creator economy offers alternative ways to monetize content, most prominently via crowdfunding platforms such as Patreon. This may further limit the effectiveness of financial incentives as a governance tool. Hence, future research should investigate changes to platform control mechanisms in more diverse settings.

References

- Alexander, J. (2018). YouTube's lesser-known creators worry for the future after major monetization changes (update). Polygon (January 17). https://www.polygon.com/2018/ 1/17/16900474/youtube-monetization-small-creators-adsense.
- Alexander, J. (2019). The golden age of YouTube is over. The Verge (April 5). https://www.theverge.com/2019/4/5/18287318/youtube-logan-paul-pewdiepie-demonetization-adpocalypse-premium-influencers-creators.
- Anderson, A., Huttenlocher, D., Kleinberg, J., and Leskovec, J. (2013). Steering User Behavior with Badges. In Proceedings of the 22nd International Conference on World Wide Web, pages 95–106.
- Anderson, C. (2004). The Long Tail. Wired Magazine (October), pages 170–177.
- Anderson, S. P. and Gabszewicz, J. J. (2006). The media and advertising: a tale of two-sided markets. Handbook of the Economics of Art and Culture, 1:567–614.
- Bhargava, H. K. (2022). The creator economy: Managing ecosystem supply, revenue sharing, and platform design. *Management Science*, 68(7):5233–5251.
- Boudreau, K. (2010). Open Platform Strategies and Innovation: Granting Access vs. Devolving Control. Management Science, 56(10):1849–1872.
- Boudreau, K. 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, K. J. and Hagiu, A. (2009). Platform Rules: Multi-sided Platforms as Regulators. *Platforms, Markets and Innovation*, 1:163–191.
- Burtch, G., He, Q., Hong, Y., and Lee, D. (2022). How do peer awards motivate creative content? experimental evidence from reddit. *Management Science*, 68(5):3488–3506.
- Burtch, G., Hong, Y., Bapna, R., and Griskevicius, V. (2018). Stimulating Online Reviews by Combining Financial Incentives and Social Norms. *Management Science*, 64(5):2065–2082.
- Cabral, L. and Li, L. (2015). A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay. *Management Science*, 61(9):2052–2063.
- Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2020). Optimal Bandwidth Choice for Robust Bias-corrected Inference in Regression Discontinuity Designs. *The Econometrics Journal*, 23(2):192–210.
- Cattaneo, M. D., Jansson, M., and Ma, X. (2017). Simple local regression distribution estimators. Working Paper.
- Ceccagnoli, M., Forman, C., Huang, P., and Wu, D. (2012). Cocreation of Value in a Platform Ecosystem! The Case of Enterprise Software. *MIS quarterly*, pages 263–290.
- Claussen, J., Kretschmer, T., and Mayrhofer, P. (2013). The Effects of Rewarding User Engagement: The Case of Facebook Apps. *Information Systems Research*, 24(1):186–200.
- Coleman, K. (2021). Introducing Birdwatch, a community-based approach to misinformation. *Twitter Blog* (January 25). https://blog.twitter.com/en_us/topics/product/2021/ introducing-birdwatch-a-community-based-approach-to-misinformation/.
- Constantinides, P., Henfridsson, O., and Parker, G. G. (2018). Introduction—Platforms and Infrastructures in the Digital Age. 29(2):381–400.
- Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer Price Search and Platform Design in Internet Commerce. American Economic Review, 108(7):1820–1859.
- Eaton, B., Elaluf-Calderwood, S., Sørensen, C., and Yoo, Y. (2015). Distributed Tuning of Boundary Resources: The Case of Apple's iOS Service System. *MIS Quarterly*, 39(1):217– 244.
- Eisenmann, T., Parker, G., and Van Alstyne, M. W. (2006). Strategies for Two-Sided Markets. *Harvard Business Review*, 84:92–101.
- Flammer, C. (2015). Does corporate social responsibility lead to superior financial performance?

a regression discontinuity approach. Management Science, 61(11):2549–2568.

- Foerderer, J., Lueker, N., and Heinzl, A. (2021). And the winner is...? the desirable and undesirable effects of platform awards. *Information Systems Research*, 32(4):1155–1172.
- Gallus, J. (2017). Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia. *Management Science*, 63(12):3999–4015.
- Gawer, A. and Henderson, R. (2007). Platform Owner Entry and Innovation in Complementary Markets: Evidence from Intel. Journal of Economics & Management Strategy, 16(1):1–34.
- Gelman, A. (2018). You need 16 times the sample size to estimate an interaction than to estimate a main effect. Statistical Modeling, Causal Inference, and Social Science Blog (March 15). https://statmodeling.stat.columbia.edu/2018/03/15/need-16-timessample-size-estimate-interaction-estimate-main-effect/.
- Gelman, A. and Imbens, G. (2019). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 37(3):447–456.
- Geva, H., Barzilay, O., and Oestreicher-Singer, G. (2019). A Potato Salad with a Lemon Twist: Using a Supply-Side Shock to Study the Impact of Opportunistic Behavior on Crowdfunding Platforms. *MIS Quarterly*, 43(4):1227–1248.
- Ghazawneh, A. and Henfridsson, O. (2013). Balancing Platform Control and External Contribution in Third-Party Development: The Boundary Resources Model. *Information* Systems Journal, 23(2):173–192.
- Gillespie, T. (2018). The Logan Paul YouTube controversy and what we should expect from internet platforms. *Vox* (January 16). https://www.vox.com/the-big-idea/2018/1/12/16881046/logan-paul-youtube-controversy-internet-companies.
- Goes, P. B., Guo, C., and Lin, M. (2016). Do Incentive Hierarchies Induce User Effort? Evidence from an Online Knowledge Exchange. *Information Systems Research*, 27(3):497–516.
- Ho, S. Y. and Rai, A. (2017). Continued voluntary participation intention in firm-participating open source software projects. *Information Systems Research*, 28(3):603–625.
- Huang, P., Ceccagnoli, M., Forman, C., and Wu, D. (2013). Appropriability mechanisms and the platform partnership decision: Evidence from enterprise software. *Management Science*, 59(1):102–121.
- Huang, P., Lyu, G., and Xu, Y. (2022). Quality regulation on two-sided platforms: Exclusion, subsidization, and first-party applications. *Management Science*, 68(6):4415–4434.
- Hukal, P., Henfridsson, O., Shaikh, M., and Parker, G. (2020). Platform Signaling for Generating Platform Content. MIS Quarterly, 44(3):1177–1205.
- Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, pages 467–475.
- Imbens, G. W. and Lemieux, T. (2008). Regression Discontinuity Designs: A Guide to Practice. Journal of Econometrics, 142(2):615–635.
- Jacobides, M. G., Cennamo, C., and Gawer, A. (2018). Towards a Theory of Ecosystems. Strategic Management Journal, 39(8):2255–2276.
- Jain, S. and Qian, K. (2021). Compensating online content producers: A theoretical analysis. Management Science, 67(11):7075–7090.
- Jhaver, S., Karpfen, Y., and Antin, J. (2018). Algorithmic Anxiety and Coping Strategies of Airbnb Hosts. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, pages 1–12.
- Kapoor, R. and Agarwal, S. (2017). Sustaining Superior Performance in Business Ecosystems: Evidence from Application Software Developers in the iOS and Android Smartphone Ecosystems. Organization Science, 28(3):531–551.
- Kerkhof, A. (2022). Advertising and Content Differentiation: Evidence from YouTube. *Working Paper*.
- Khern-am nuai, W., Kannan, K., and Ghasemkhani, H. (2018). Extrinsic Versus Intrinsic

Rewards for Contributing Reviews in an Online Platform. *Information Systems Research*, 29(4):871–892.

- Koetsier, J. (2020). YouTube Will Now Show Ads On All Videos Even If Creators Don't Want Them. Forbes (November 18). https://www.forbes.com/sites/johnkoetsier/2020/ 11/18/youtube-will-now-show-ads-on-all-videos-even-if-creators-dont-wantthem/?sh=155ae06e4913.
- Koo, W. W. and Eesley, C. E. (2020). Platform Governance and the Rural–Urban Divide: Sellers' Responses to Design Change. *Strategic Management Journal*.
- Lee, D. S. and Lemieux, T. (2010). Regression Discontinuity Designs in Economics. Journal of Economic Literature, 48(2):281–355.
- Locke, E. A. and Latham, G. P. (2002). Building a Practically Useful Theory of Goal Setting and Task Motivation: A 35-year Odyssey. *American Psychologist*, 57(9):705.
- Loh, J. and Kretschmer, T. (2023). Online communities on competing platforms: Evidence from game wikis. *Strategic Management Journal*, 44(2):441–476.
- Ma, M. and Agarwal, R. (2007). Through a Glass Darkly: Information Technology Design, Identity Verification, and Knowledge Contribution in Online Communities. *Information Systems Research*, 18(1):42–67.
- Maggi, L., Gkatzikis, L., Paschos, G., and Leguay, J. (2018). Adapting caching to audience retention rate. *Computer Communications*, 116:159–171.
- Mohan, N. and Kyncl, R. (2018). Additional Changes to the YouTube Partner Program (YPP) to Better Protect Creators. *YouTube Official Blog* (January 16). https://blog.youtube/newsand-events/additional-changes-to-youtube-partner.
- Nicas, J. (2017). Google's YouTube Has Continued Showing Brands' Ads With Racist and Other Objectionable Videos. *The Wall Street Journal* (March 24). https://www.wsj.com/articles/googles-youtube-has-continued-showing-brandsads-with-racist-and-other-objectionable-videos-1490380551.
- Oestreicher-Singer, G. and Sundararajan, A. (2012). Recommendation Networks and the Long Tail of Electronic Commerce. *MIS Quarterly*, 36(1):65–83.
- Panciera, K., Halfaker, A., and Terveen, L. (2009). Wikipedians are born, not made: a study of power editors on Wikipedia. In *Proceedings of the ACM 2009 international conference on* Supporting group work, pages 51–60.
- Papay, J. P., Murnane, R. J., and Willett, J. B. (2010). The Consequences of High School Exit Examinations for Low-performing Urban Students: Evidence from Massachusetts. *Educational Evaluation and Policy Analysis*, 32(1):5–23.
- Parker, G. and Van Alstyne, M. (2018). Innovation, Openness, and Platform Control. Management Science, 64(7):3015–3032.
- Popper, B. (2017). YouTube will no longer allow creators to make money until they reach 10,000 views. *The Verge* (April 6). https://www.theverge.com/2019/4/5/18287318/youtube-logan-paul-pewdiepie-demonetization-adpocalypse-premium-influencers-creators.
- Posen, H. E., Keil, T., Kim, S., and Meissner, F. D. (2018). Renewing research on problemistic search—a review and research agenda. *Academy of Management Annals*, 12(1):208–251.
- Pursey, J. (2022). 7 Facts About The 1983 Video Game Crash. *GameRant* (April 17). https://gamerant.com/1983-video-game-crash-facts-details-trivia/.
- Ren, Y., Kraut, R., and Kiesler, S. (2007). Applying Common Identity and Bond Theory to Design of Online Communities. Organization Studies, 28(3):377–408.
- Rietveld, J., Ploog, J. N., and Nieborg, D. B. (2020). Coevolution of Platform Dominance and Governance Strategies: Effects on Complementor Performance Outcomes. Academy of Management Discoveries, 6(3):488–513.
- Rietveld, J., Schilling, M. A., and Bellavitis, C. (2019). Platform Strategy: Managing Ecosystem Value through Selective Promotion of Complements. *Organization Science*, 30(6):1232–1251.

- Rietveld, J., Seamans, R., and Meggiorin, K. (2021). Market orchestrators: The effects of certification on platforms and their complementors. *Strategy Science*, 6(3):244–264.
- Shah, S. K. (2006). Motivation, Governance, and the Viability of Hybrid Forms in Open Source Software Development. *Management Science*, 52(7):1000–1014.
- Statt, N. (2017). YouTube is facing a full-scale advertising boycott over hate speech. The Verge (March 24). https://www.theverge.com/2017/3/24/15053990/google-youtubeadvertising-boycott-hate-speech.
- Steiner, P. O. (1952). Program patterns and preferences, and the workability of competition in radio broadcasting. *The Quarterly Journal of Economics*, 66(2):194–223.
- Sun, M. and Zhu, F. (2013). Ad revenue and content commercialization: Evidence from blogs. Management Science, 59(10):2314–2331.
- Sun, Y., Dong, X., and McIntyre, S. (2017). Motivation of User-Generated Content: Social Connectedness Moderates the Effects of Monetary Rewards. *Marketing Science*, 36(3):329– 337.
- Tae, C. J., Luo, X., and Lin, Z. (2020). Capacity-Constrained Entrepreneurs and Their Product Portfolio Size: The Response to a Platform Design Change on a Chinese Sharing Economy Platform. *Strategic Entrepreneurship Journal*, 14(3):302–328.
- Tang, Q., Gu, B., and Whinston, A. B. (2012). Content contribution for revenue sharing and reputation in social media: A dynamic structural model. *Journal of Management Information* Systems, 29(2):41–76.
- Tiwana, A. (2015). Platform Desertion by App Developers. Journal of Management Information Systems, 32(4):40–77.
- Toubia, O. and Stephen, A. T. (2013). Intrinsic vs. Image-Related Utility in Social Media: Why do People Contribute Content to Twitter? *Marketing Science*, 32(3):368–392.
- Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. The Quarterly Journal of Economics, 106(4):1039–1061.
- Waldfogel, J. (2017). How digitization has created a golden age of music, movies, books, and television. *Journal of Economic Perspectives*, 31(3):195–214.
- Wareham, J., Fox, P. B., and Cano Giner, J. L. (2014). Technology Ecosystem Governance. Organization Science, 25(4):1195–1215.
- Wilbur, K. C. (2008). A two-sided, empirical model of television advertising and viewing markets. Marketing science, 27(3):356–378.
- Wojcicki, S. (2017). Expanding our work against abuse of our platform. YouTube Official Blog (December 5). https://blog.youtube/news-and-events/expanding-our-work-against-abuse-of-our.
- Wu, Y. and Zhu, F. (2022). Competition, contracts, and creativity: Evidence from novel writing in a platform market. *Management Science*, 68(12):8613–8634.
- YouTube Official Blog (2012). Being a YouTube Creator Just Got Even More Rewarding. YouTube Official Blog (April 12). https://blog.youtube/news-and-events/beingyoutube-creator-just-got-even.
- Zervas, G., Proserpio, D., and Byers, J. W. (2021). A first look at online reputation on airbnb, where every stay is above average. *Marketing Letters*, 32(1):1–16.
- Zhu, F. and Zhang, X. (2010). Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics. *Journal of Marketing*, 74(2):133–148.

Figures



Figure 1 Test for continuity at the subscriber threshold











Figure 3 RDD plots: Robustness – Quadratic model fit





Tables

 Table 1 Summary statistics

	Mean	Std. Dev.	Min.	Max.	Ν
Lost Access	0.80	0.40	0	1.00	484
Subscriber count	1026.43	75.48	901	1196.00	484
Watchtime	5307.68	11845.07	3	98600.34	484
Upload frequency	3.248	5.708	0.14	68.29	428
Like share	0.91	0.11	0	1.00	425
Unique keywords	41.62	44.81	1	480.00	428
Lifetime video uploads	128.88	131.69	1	616.00	484
Mainstream	0.60	0.49	0	1.00	484
Age	35.57	17.19	1	107.00	484

The summary statistics are based on our main sample of creators who did not exit the platform after the rule change.

	Non-Exit	\mathbf{Exit}	Diff.	
Subscriber count	1026.82 (77.33)	1097.15 (83.17)	-70.33 (118.06)	* * *
Watchtime	5308.02 (11786.76)	$\begin{array}{c} 6444.02 \\ (20154.47) \end{array}$	-1136.00 (23859.67)	
Upload frequency	$\begin{array}{c} 0.76 \\ (0.05) \end{array}$	$0.6 \\ (0.04)$	$0.16 \\ (0.06)$	* * *
Like share	$0.91 \\ (0.11)$	$0.92 \\ (0.09)$	-0.01 (0.14)	
Unique keywords	$33.63 \ (35.33)$	$29.21 \\ (23.97)$	$4.42 \\ (45.17)$	* * *
Mainstream	$0.59 \\ (0.49)$	$\begin{array}{c} 0.53 \\ (0.50) \end{array}$	$\begin{array}{c} 0.06 \\ (0.73) \end{array}$	*
Experience	$35.59 \ (17.16)$	35.09 (22.15)	0.50 (28.92)	

Table 2 Difference in means between Exit and Non-Exit

*** p < 0.01, ** p < 0.05, * p < 0.1. Standard deviations in parentheses. The tables reports the results of two-tailed t-tests on the difference in means between creators who appeared in the second wave of data collection (Non-Exit) and those that did not (Exit).

Table	3	Manipulation	test
-------	---	--------------	------

	Upload Frequency (1)	Like Share (2)	Unique Tags (3)
Lost Access	-0.173 (0.209)	-0.035 (0.024)	-4.668 (6.721)
Bandwidth Observations	[900, 1200] 428	$[900, 1200] \\ 425$	$[900, 1200] \\ 428$

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models use local linear regressions and control for watchtime. The sample consists of creators who did not exit the platform before the first and second waves of data collection. The outcome variables are calculated based on activity in the six months before the rule change.

	Upload Frequency						
			Expe	rience			
	$\begin{array}{c} \text{All} \\ (1) \end{array}$	Mainstream (2)	$\begin{array}{c} { m Niche} \ (3) \end{array}$	High (4)	$ \begin{array}{c} \text{Low} \\ (5) \end{array} $		
Lost Access	-2.816^{**} (1.387)	-3.617^{*} (2.098)	-1.278 (1.427)	-3.82^{*} (2.038)	-0.841 (1.559)		
Bandwidth Observations	[900, 1200] 428	[900, 1200] 253	[900, 1200] 175	[900, 1200] 241	[900, 1200] 187		

 Table 4 Main Results: Creator activity

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watchtime. Sample contains creators who did not exit the platform between the first and second waves of data collection.

 Table 5 Main Results: Content quality

		Like Share					
	Experience				rience		
	$\begin{array}{c} \text{All} \\ (1) \end{array}$	Mainstream (2)	$\stackrel{ m Niche}{ m (3)}$	High (4)	$ \begin{array}{c} \text{Low} \\ (5) \end{array} $		
Lost Access	-0.048^{*} (0.025)	-0.061^{**} (0.030)	-0.034 (0.042)	-0.048 (0.035)	-0.045^{*} (0.026)		
Bandwidth Observations	$[900, 1200] \\ 425$	[900, 1200] 251	[900, 1200] 174	$[900, 1200] \\ 240$	[900, 1200] 185		

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watchtime. Sample contains creators who did not exit the platform between the first and second waves of data collection.

	Unique Keywords						
				Experience			
	All	Mainstream	Niche	High	Low		
	(1)	(2)	(3)	(4)	(5)		
Lost Access	-25.380^{**} (10.485)	-19.803^{*} (10.393)	$-33.294 \\ (21.493)$	-34.082^{**} (15.873)	$-8.393 \\ (12.517)$		
Bandwidth Observations	[900, 1200] 428	[900, 1200] 253	[900, 1200] 175	[900, 1200] 241	[900, 1200] 187		

 Table 6 Main Results: Content diversity

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watch time. Sample contains creators who did not exit the platform between the first and second waves of data collection.

 Table 7 Robustness: Bandwidth

	Upload F	Upload Frequency		Like Share		Keywords
	(1) wide	(2) narrow	(3) wide	(4) narrow	(5) wide	(6) narrow
Lost Access	-1.207 (0.809)	-3.079 (2.154)	-0.010 (0.016)	-0.039 (0.045)	-11.503^{*} (6.136)	-30.236 (18.554)
Bandwidth Observations	[800,1400] 827	$[950,\!1100]\\263$	[800,1400] 819	[950,1100] 261	[800,1400] 827	$[950,\!1100]\\263$

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watchtime. Sample contains creators who did not exit the platform between the first and second waves of data collection. The table reports results using different bandwidths of the running variable (subscriber count).

	Upload Frequency		Like	Like Share		Keywords
	(1) quadratic	(2) cubic	(3) quadratic	(4) cubic	(5) quadratic	(6) cubic
Lost Access	-3.649 (2.393)	$-2.1 \\ (3.595)$	-0.048 (0.047)	-0.021 (0.091)	-32.774 (20.703)	-38.769 (35.788)
Bandwidth Polynomials Observations	$[900,1200] \\ 2 \\ 428$	$[900,1200] \\ 3 \\ 428$	$[900,1200] \\ 2 \\ 425$	$[900,1200] \\ 3 \\ 425$	$[900,1200] \\ 2 \\ 428$	$[900,1200] \\ 3 \\ 428$

 ${\bf Table \ 8 \ Robustness: \ Higher \ order \ polynomials}$

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watchtime. Sample contains creators who did not exit the platform between the first and second waves of data collection. The table reports results using different higher order polynomials to fit $f(Subscribers_i)$ in equation 1.

	Upload Frequency		Like	Share	Unique Keywords	
	(1) 3 months	(2) 12 months	(3) 3 months	$(4) \\ 12 months$	(5) 3 months	(6) 12 months
Lost Access	-4.281^{**} (2.111)	-2.182^{**} (0.981)	-0.036 (0.025)	-0.051^{*} (0.028)	-30.583^{**} (12.462)	-23.706^{**} (9.438)
Bandwidth Observations	[900, 1200] 359	[900, 1200] 465	[900, 1200] 357	[900, 1200] 463	[900, 1200] 359	[900, 1200] 465

 Table 9 Robustness: Different time windows

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watchtime. Sample contains creators who did not exit the platform between the first and second waves of data collection. The table reports results using measures calculated based on creator activity within three and twelve months after the rule change.

	Upload Frequency		Like	Like Share		Keywords
	(1) More	(2) Less	(3) More	(4) Less	(5) More	(6) Less
Lost Access	-4.897^{**} (2.44)	$-0.325 \ (0.369)$	-0.018 (0.030)	$egin{array}{c} -0.075^{**} \ (0.038) \end{array}$	-32.967^{*} (18.993)	-15.466^{**} (7.542)
Bandwidth Observations	$[900, 1200] \\ 222$	[900, 1200] 206	$[900, 1200] \\ 222$	$[900, 1200] \\ 203$	$[900, 1200] \\ 222$	$[900, 1200] \\ 206$

 ${\bf Table \ 10 \ Robustness: \ Alternative \ Experience \ Measure}$

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors in parentheses are reported. All models control for watchtime. Sample contains creators who did not exit the platform between the first and second waves of data collection. The table reports results using the number of videos uploaded before the rule change as an experience measure.