Advertising and Content Differentiation: Evidence from YouTube*

Anna Kerkhof†

May 6, 2022

Abstract

Many media outlets depend on advertising revenue to finance their operations, but the effect of advertising on media outlets' content choice is largely unexplored. This paper exploits two institutional features from YouTube to show that an exogenous increase in advertising quantity induces YouTubers to differentiate their video content from their competitors. A plausible mechanism is that YouTubers avoid competition: Viewers typically perceive advertising as a nuisance and therefore as an implicit price they have to pay; thus, they could switch to a competitor if a YouTuber increased her advertising quantity. This is less likely, however, if the YouTuber differentiates her content from the mainstream and moves to a niche.

JEL Codes: D22, L15, L82, L86

Keywords: advertising, content differentiation, economics of digitization, horizontal product differentiation, long tail, media diversity, user-generated content, YouTube

*I thank Johannes Münster for helpful comments, discussion, and guidance. Moreover, I thank the Editor Steffen Huck, two anonymous referees, Maja Adena, Luis Aguiar, Ruediger Bachmann, Kai Barron, Felix Bierbrauer, Joerg Claussen, Gregory Crawford, Milena Djourelova, Dante Donati, Ruben Durante, Tomaso Duso, Matthew Ellman, Ruben Enikolopov, Thomas Fackler, Oliver Falck, Marcel Garz, Fabrizio Germano, Sampreet Goraya, Paul Heidhues, Matthias Heinz, Ulrich Kaiser, Reinhold Kesler, Tobias Kretschmer, Ulrich Laitenberger, Jonas Loebbing, Johannes Loh, Leonardo Madio, Simon Martin, Sebastian Mertesacker, Axel Ockenfels, Martin Peitz, Christian Peukert, Susanne Prantl, Maximilian Schäfer, Paul Schempp, Daniel Sturm, Frederic Thenee, Kevin Tran, Dainis Zegners, and participants of numerous workshops and conferences for helpful comments. Lolakshi Rajlakshmi, Sarah Recht, and especially Aleksandar Vasic provided excellent research assistance. All errors are my own.

†ifo Institute for Economic Research, LMU Munich, and CESifo. I gratefully acknowledge financial support from the Joachim-Herz-Stiftung and the Bavarian Research Institute for Digital Transformation (bidt). E-mail: kerkhof[at]ifo.de.

1. Introduction

Media consumption is an integral part of our everyday lives. American adults spend several hours per day reading, watching, and listening to media content[1]; similar numbers hold for citizens from the EU.[2] The value of media consumption crucially depends on the match between consumers' preferences and the media content that is being provided. Preferences over media content, however, differ substantially between different groups of consumers: men and women, for instance, prefer very different types of media content (Anderson and Waldfogel, 2015). Thus, the more horizontally differentiated the content in media markets, the more likely it is that all consumers' preferences are served and that their value of media consumption is maximized (Waldfogel, 2007). Duplication of media content, in contrast, may lead to foregone consumer surplus, business stealing, and the inefficient duplication of fixed costs.

What drives content differentiation in media markets? Seminal analyses of horizontal product differentiation have singled out two general effects: A direct effect induces firms to move closer to their competitors to increase their market share, leading to minimum differentiation (Hotelling, 1929); a strategic effect prompts firms to move away from their competitors to soften price competition, leading to maximum differentiation (d'Aspremont et al., 1979; Economides, 1984, 1986) – when prices are fixed, the direct effect prevails. Whether and to which extent this logic also applies to media markets is unclear, though. First, media outlets typically charge their consumers a low or even no price at all and generate revenue through advertising instead. While pioneering models on media outlets' content choice argue that, in the absence of price competition, advertising leads to minimum content differentiation (Steiner, 1952; Beebe, 1977), more recent papers acknowledge that many consumers perceive advertising as nuisance and therefore as a "price" they have to pay (Wilbur, 2008; Huang et al., 2018; Anderson and Jullien, 2016). Hence, when incentivized by ad revenue, media outlets could also prefer to differentiate their content from each other to soften competition in the ad "price". Second, modern media markets are far from the duopoly setting that is typically considered in the theoretical literature. Although some baseline intuitions also hold in more general settings of spatial competition[3], it is an open empirical question whether the predictions from a Hotelling-style environment apply to markets where many (sometimes even thousands of) media outlets compete.

This paper studies the effect of advertising on content differentiation on YouTube – the second-most visited website in the world[4] – to resolve the open question. To this end, I study unique and newly collected data on several thousand German YouTubers and more than a million videos that have not been investigated before. I demonstrate that

[1] See `https://www.nielsen.com/us/en/insights/article/2018/time-flies-us-adults-now-spend-nearly-half-a-day-interacting-with-media/` (Aug 2020).

[2] See `https://www.emarketer.com` (Aug 2020).

[3] E.g., Appendix A uses the spokes model by Chen and Riordan (2007) to motivate the empirical analysis

[4] See `www.alexa.com/topsites` (Aug 2020).

an exogenous increase in the feasible advertising quantity leads to a considerable decrease in the YouTubers' probability to duplicate mainstream content, i.e., the type of content that attracts the largest number of views. The result is likely to be driven by an intuitive economic mechanism: Mainstream content is provided by many competing YouTubers; thus, viewers could easily switch to a competitor if a YouTuber increased her advertising quantity. Switching is less likely, however, if the YouTuber differentiates her content from the mainstream, moves towards a niche, and thereby softens competition in the ad "price." The empirical results are therefore in accordance with recent considerations that acknowledge a conceptual similarity of advertising and subscription prices as determinants of content differentiation in media markets. Moreover, the results are in line with more general predictions about the interplay of price competition and horizontal product differentiation: When prices are fixed (or advertising quantities are limited), firms locate close to their competitors. If price competition is possible (if the feasible advertising quantity goes up), the strategic effect becomes stronger, whereby product differentiation increases.[5]

I exploit two institutional features of YouTube to identify causal effects. First, I use the "ten minutes trick", which is a discontinuity in YouTube's mapping from video duration to the technically feasible number of ad breaks per video: If a video is shorter than ten minutes, YouTubers can permit for exactly one ad break in it. If the video is ten minutes or longer, YouTubers face no such limitation. Second, the ten minutes trick was unknown to the majority of YouTubers until Nov 2015, when YouTube launched a new ad break tool that made its existence prominent to the community. I focus on a subsample of YouTubers who were unaware of the ten minutes trick before Nov 2015; in other words, I consider a subsample of YouTubers whose advertising quantity was limited. Then, I compare the change in the probability to duplicate mainstream content of YouTubers whose feasible advertising quantity increased after Nov 2015 (YouTubers who increased their share of videos that are ten minutes or longer) with the change in content of YouTubers whose feasible advertising quantity remained constant in a difference-in-differences framework.

The increase in the feasible advertising quantity is endogenous, though – after all, the YouTubers have perfect control over their videos' duration, and particularly money-loving YouTubers could be especially eager to produce longer videos to show more ads. To take this into account, I use the YouTubers' *median video duration before Nov 2015* – in a sense, their "closeness" to the ten minutes threshold – to instrument for the increase in the feasible advertising quantity. I argue that the instrument is relevant and exogenous. On the one hand, a YouTuber's median video duration before Nov 2015 is correlated to the potential increase in the feasible number of ad breaks per video, as extending the videos' duration to at least ten minutes is easier for YouTubers with videos close to the threshold. On the other hand, the YouTubers in my sample did not choose their videos' duration before Nov 2015 bearing the ten minutes trick in mind, because they were unaware of the feature. As a result, the median video duration before Nov 2015 is exogenous to omitted variables that might drive an increase in the feasible advertising quantity (e.g., commercial

[5] I further support the argument with a theoretical model à la Chen and Riordan (2007) in Appendix A.

interests) and has furthermore no effect on the YouTubers' probability to duplicate mainstream content except through the potential increase in advertising quantity. A broad range of validity checks supports the identification strategy.

The analysis of around one million YouTube videos shows that an increase in the feasible number of ad breaks per video leads to a twenty percentage point reduction in the YouTubers' probability to duplicate mainstream content. The effect size is considerable: it corresponds to around 40% of a standard deviation in the dependent variable and to around 50% of its baseline value. The large sample size allows me to conduct several subgroup analyses to study effect heterogeneity. I find that the positive effect of advertising on content differentiation is driven by the YouTubers who have at least $1,000$ subscribers, i.e., the YouTubers whose additional ad revenue is likely to exceed the costs from adapting their videos' content. Moreover, I find heterogeneity along video categories: some categories are more flexible in terms of their typical video duration than others, hence, exploiting the ten minutes trick is easier (e.g., a music clip is typically between three and five minutes long and cannot be easily extended). Finally, I show that YouTubers who produced a large proportion of mainstream content even before the launch of the new ad break tool – i.e., the relatively well-known trend-setters – are less likely to move to a niche after Nov 2015, as they experience relatively low competitive pressure.

The empirical findings support competition among YouTubers as a major economic mechanism behind my main results. In particular, I show that mainstream content – i.e., content in high demand by the audience – is also *supplied* by many competing YouTubers. Thus, YouTubers who duplicate mainstream content run the risk of losing their audience to a competitor if they increase their advertising quantity. This is less likely, however, if YouTubers differentiate their content from the mainstream, move towards a niche, and thereby soften competition in the ad "price." Matching this idea, I demonstrate that the number of competing options to each video decreases both on the extensive and on the intensive margin if YouTubers increase their feasible advertising quantity.

To better interpret the effect and relate it to the bigger picture, note that my empirical strategy isolates a strategy change for a relatively small proportion of YouTubers, while the majority of YouTubers in my sample are unaffected. Hence, the evidence that I present should ultimately be understood as a clean documentation of the effect of advertising on content choice for YouTubers who are willing to adapt their content, and not as an effect that involves the entire market.

My paper makes two main contributions to the literature. First, I advance the knowledge on the effect of advertising on content differentiation in media markets. To my knowledge, this is the first paper that provides evidence of a causal *positive* effect of advertising on content differentiation, whereby it challenges the widespread – public and academic – opinion that media outlets duplicate mainstream content when incentivized by ad revenue (e.g., Herman and McChesney, 1997; Hamilton, 2004; McChesney, 2004). This is a major insight, especially because the media's options to generate ad revenue are often subject to external regulation.[6]

Second, my results contribute to recent discussions about the effect of digitization on content differentiation and diversity in media markets (Waldfogel, 2017, 2018). The cost structure of traditional media markets – high fixed and low marginal costs – impedes media diversity, as the number of outlets that can co-exist is limited. Goldfarb and Tucker (2019), however, point out that digital technology has "reduced the cost of storage, computation, and transmission of data" (p.3). As a result, online media outlets can afford to provide niche content, and enhanced search technologies simultaneously enable consumers to find it – a phenomenon that Anderson (2006) summarizes as "the long tail."[7] YouTube serves as a point in case to study the determinants of content differentiation in modern media markets (e.g., online news markets or alternative user-generated content platforms). In particular, my paper shows that the prospect of advertising revenue provides additional incentives for media outlets to differentiate their content that – in combination with enhanced technology – can help to increase media diversity.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. In Section 3, I provide background information on YouTube, its monetization policy, and the institutional features that the empirical strategy builds on. A stylized example introduces the central ideas of identification in Section 4, before I illustrate the data collection process and how I construct a dataset that is suitable for the analysis in Section 5. Section 6 discusses the details of the empirical strategy; the results are presented in Section 7. Next, in Section 8, I explore competition as a plausible economic mechanism behind these results. Section 9 investigates changes in video quality; Section 10 concludes.

2. Related literature

The paper is related to four strands of literature that partially overlap. First, it is linked to the classic literature on spatial competition, including the seminal theoretical papers by Hotelling (1929) and Salop (1979), as well as more recent contributions (e.g., Gabszewicz and Thisse, 1986; Anderson et al., 1992; Chen and Riordan, 2007), where it especially adds to empirical evidence on the relationship between price competition and horizontal product differentiation (e.g., Numan and Willekens, 2012; Netz and Taylor, 2002; Davcik and Sharma, 2015).

Second, the paper contributes to the literature on content choice in media markets. Most closely related are the papers by Seamans and Zhu (2014) and Sun and Zhu (2013). While Seamans and Zhu (2014) demonstrate that newspaper subscription prices are positively correlated to content differentiation, this paper demonstrates that an increase in advertising quantity can have an analogous effect. Sun and Zhu (2013) study the introduction of an ad-revenue-sharing program on a major Chinese online platform and find that ad-

[6]The Audiovisual Media Services Directive, for instance, requires that the proportion of television advertising and teleshopping spots within a given clock hour shall not exceed 20% (Article 23 §1).

[7]See also Brynjolfsson et al. (2003, 2011) for a discussion on the long tail and how consumer surplus benefits from increased product variety.

vertising leads to more duplication of mainstream content. Our results complement each other: While ad breaks before or during YouTube videos are likely to be a true nuisance to viewers, Sun and Zhu (2013) explicitly state that the ads appearing on the bloggers' posts are not intrusive (p. 2317), which means that only a direct, but no strategic effect operates in their setting.

Third, my work relates to analyses on further determinants of content differentiation in media markets. E.g., Berry and Waldfogel (2001) and Sweeting (2010) show that ownership concentration in the US radio market increased the number of formats relative to the number of stations, most likely to avoid business stealing effects stations with common owners. Similarly, George (2007) finds that ownership concentration in US newspaper markets increases differentiation and the variety of topics.

Finally, the paper contributes to the large and growing literature on user-generated content (UGC) (see Luca, 2016b, for a survey). First, I present a novel empirical strategy to identify causal effects on a UGC platform. While existing approaches use variation in institutional features *across* platforms (e.g., Chevalier and Mayzlin, 2006; Mayzlin et al., 2014), *within* platforms (Anderson and Magruder, 2012; Luca, 2016a), or conduct randomized experiments (Bond et al., 2012; Aral and Walker, 2012), I exploit two distinctive features of YouTube's monetization policy to identify the causal effect of advertising on the YouTubers' content choice. Second, my paper explores how monetization affects UGC. Since many other UGC platforms such as Wikipedia, TripAdvisor or Twitter do not allow their contributors to earn money, YouTube offers a unique environment to study this question. While previous analyses show that users mainly contribute UGC for reputational reasons (Wang, 2010; Anderson et al., 2013; Easley and Ghosh, 2013), the results of this paper demonstrate that economic considerations matter, too.

3. YouTube: Background

3.1. Platform, audience, and contributors

YouTube is a video sharing platform founded in 2005 and acquired by Google in 2006. Its reach is tremendous: with 800 million unique users and 15 billion visits per month, it is the second-most popular website in the world (after google.com).[8] As of March 2022, more than a billion hours of video content from YouTube are watched every day.[9]

YouTube is based on user-generated content. While unregistered users are limited to watching, registered users can upload, share, and comment on videos. Registered users who upload videos on a regular basis are called *YouTubers*; YouTubers, in turn, operate a YouTube *channel* under their user name to distribute their videos.[10] Appendix A presents a stylized theoretical framework that illustrates the economic incentives of

[8] See `www.alexa.com/topsites` (March 2022).
[9] See `https://blog.youtube/press/` (March 2022).
[10] I use the terms "YouTuber" and "channel" synonymously; cases where one YouTuber operates several channels are rare.

viewers, YouTubers, and YouTube itself in detail.

3.2. Monetization

YouTubers have the option to monetize their content; in particular, they can generate advertising revenue by permitting YouTube to show ads to viewers before or during their videos. However, while YouTubers can permit that ads *may* be shown, YouTube's algorithm determines *if* and *which* ad is displayed to a particular viewer. Thus, there is no direct relationship between YouTubers and advertisers.[11] According to anecdotal evidence – official statistics do not exist – YouTubers earn about three to five USD per 1,000 views per ad per video.[12]

Monetization via ad breaks is not open to all YouTubers, though. First, a YouTuber's content must be advertiser-friendly, i.e., free of violence, sex, and crime.[13] In early 2017, YouTube introduced a new policy of automated demonetization of non-advertiser-friendly content (also known as "adpocalypse") that aims at videos on sensitive social issues, tragedy, or conflict; many YouTubers reported losing more than half of their income as a result.[14] Second, while not bounded to a subscriber threshold before, YouTube disabled the monetization option for YouTubers with fewer than 1,000 subscribers in Feb 2018 (Abou El-Komboz et al., 2022). This policy, too, is a reaction to advertisers' complaints about their products appearing next to dubious video content.[15] The subscriber threshold, YouTube argues, gives them enough information to determine the validity of a YouTuber's channel and to confirm that it is following the YouTube community guidelines and advertiser policies.[16]

3.3. The ten minutes trick

YouTube's monetization policy exhibits one distinctive feature, which is known as the "ten minutes trick." The ten minutes trick refers to a discontinuity in YouTube's mapping from a video's duration to the technically feasible number of ad breaks that the YouTuber can permit. If a video is shorter than ten minutes, YouTubers can permit for exactly one ad break in it. If, on the other hand, the video is ten minutes or longer, YouTubers face no technical restriction on the number of ad breaks.[17] Hence, the ten minutes trick can be summarized as

$$feasible\ number\ of\ ad\ breaks = \begin{cases} 1 & if\ video\ duration < 10\ min \\ \infty & if\ video\ duration \geq 10\ min. \end{cases} \tag{1}$$

[11] In addition to permitting for ad breaks in their videos, YouTubers might also earn money through product placement and affiliate links. In this case, there exists a contractual basis with the advertiser.

[12] See `influencermarketinghub.com/how-much-do-youtubers-make/` (Dec 2018).

[13] See `support.google.com/youtube/answer/6162278?hl=en` (Dec 2018).

[14] See `nymag.com/intelligencer/2017/12/can-youtube-survive-the-adpocalypse.html` (Dec 2018).

[15] See `turbofuture.com/internet/YouTube-Screwed-Small-YouTube-Channels-With-Their-New-Memorization-Policy` (Dec 2018).

[16] `support.google.com/youtube/answer/72857?hl=en\&ref_topic=6029709` (Dec 2018).

[17] `support.google.com/youtube/answer/6175006?hl=en` (Oct 2018).

While the ten minutes trick had long been a hidden feature, it gained sudden prominence in Nov 2015, when YouTube launched a new ad break tool for YouTubers.[18] The tool had two effects. First and foremost, it made the ten minutes trick apparent. In its old version, only a small additional input box would appear for videos exhibiting the ten minutes threshold (A in Figure A.18). In contrast to that, the option to embed additional ad breaks is now permanently visible and points YouTubers to its existence (B in Figure A.19). Second, editing additional ad breaks became less cumbersome. The new tool allows YouTubers to drag ad breaks back and forth on their video time line and it also offers a preview option to check whether an ad appears at an appropriate point in time during the video (C and D in Figure A.19). The old version, in contrast, required typing and re-typing the point in time where the ad breaks were supposed to appear (A in Figure A.18).

4. Identification: Stylized example

An ideal experiment would randomly assign some YouTubers to the option of showing just one, and others to the option of showing several ads per video to their viewers, and then compare the groups' probabilities to duplicate mainstream content. Given that the YouTubers' real life monetization settings are endogenous, however, the identification of a causal link from advertising quantity to individual content choice requires a thoughtful empirical strategy. Though highly stylized, this section illustrates how combining the ten minutes trick with the launch of the new ad break tool yields variation in the YouTubers' feasible number of ad breaks per video that I exploit to identify the causal effect of interest.

Figure 1 illustrates YouTube's mapping from video duration to the technically feasible number of ad breaks per video as illustrated above. Consider three hypothetical YouTubers A, B, and C *before* Nov 2015, where A's videos are very short, B's videos are close to but still below the ten minutes threshold, and C's videos are longer than that. Hence, while A and B may only permit for one ad break per video, C faces no such limitation. Note that this is no regression discontinuity setting, because the YouTubers have perfect control over their videos' duration. In particular, C could have chosen her videos' duration strategically to benefit from the jump in the feasible number of ad breaks per video.

Next, consider the launch of the new ad break tool in Nov 2015. While C is unaffected, A and B realize that they can increase the feasible number of ad breaks per video by uploading videos that are ten minutes or longer. Pushing her video duration beyond the threshold, however, is easier to accomplish for B than for A. The key identifying assumption is that although a YouTuber has perfect control over her videos' duration, A and B, who were initially ignorant of the threshold's existence, did not choose their videos' *distance* to the ten minutes threshold having the discontinuity in mind. As a consequence, the costs of moving beyond the threshold *after* it became prominent – and thereby also the probability to actually do so – is exogenous to unobserved characteristics such as, for

[18]See `www.youtube.com/watch?v=z58Ed6q6xQg` (Oct 2018).

instance, commercial incentives that may also drive a YouTuber's decision to increase her feasible advertising quantity.[19]

I exploit the variation in the YouTubers' costs to move beyond the threshold as follows: First, I consider only YouTubers like A and B, i.e., YouTubers below the ten minutes threshold before Nov 2015, whose feasible advertising quantity was therefore restricted. Then, I compare the change in the probability to upload mainstream content before and after Nov 2015 of YouTubers who could increase the feasible number of ad breaks per video by uploading videos that are ten minutes or longer to YouTubers whose feasible advertising quantity remained constant in a difference-in-differences framework. Finally, I account for endogeneity in the increase of feasible advertising quantity by using a YouTuber's "closeness" to the ten minutes threshold before Nov 2015 as an instrument. Thus, my empirical strategy boils down to exploiting exogenous variation between YouTubers who were close to the threshold before Nov 2015 to YouTubers who were further away from it (in contrast to comparing YouTubers just left to the threshold to YouTubers just right to it, as one would do in a regression discontinuity design). A detailed discussion of the empirical strategy follows in Section 6.

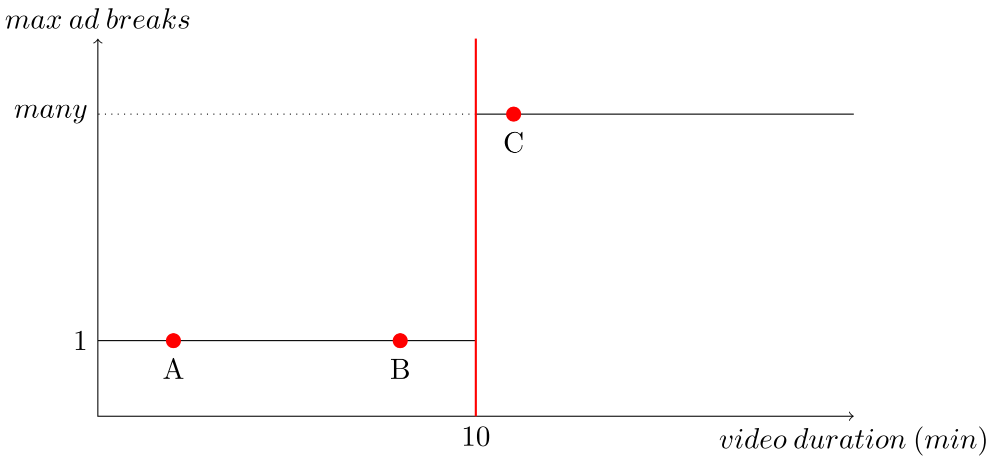

Figure 1: Stylized example of the identification strategy.

[19] To be precise, A and B could correspond to three types of YouTubers: (i) those who did not know about the threshold, as discussed above, (ii) those who knew about the threshold, but found it too cumbersome to permit for additional ad breaks, and (iii) those who knew but did not want to increase their videos' duration. The logic that applies to YouTubers in group (i) holds for YouTubers in group (ii) as well. YouTubers in group (iii) can be interpreted as "never-takers", see Section 6 for a discussion of instrument heterogeneity.

5. Data

5.1. Data collection

To carry out the above analysis, I collect data via the YouTube Data API and HTML webscraping. I start by using the website `channelcrawler.com` to compile a list of all active German YouTubers as of Oct 2017. Based on this, I collect data on the YouTuber level – including a full history of video uploads – from the YouTube Data API. Finally, I retrieve data on the video level, including the date of upload, video duration, views, likes, dislikes, video category, and video tags. Video tags are descriptive keywords that YouTubers can add to their videos to help viewers find them (see Section 5.2 and Appendix B.1 for further discussion). Note that views, likes and dislikes are accumulative measures; thus, I retrieve these numbers as they are on the day of data collection.

Data on the YouTubers' monetization settings is, unfortunately, highly limited. In particular, the YouTube Data API does not provide any information on the exact number of ad breaks per video. While this information is in principle available in a video's HTML code, YouTube prohibits any automated data collection that is "faster than a human"[20], making it impossible to crawl every single video in the dataset within a reasonable amount of time. To retrieve data on the YouTubers' monetization settings nonetheless, I make a compromise: I let a webscraper crawl twenty randomly drawn videos per YouTuber. If it detects at least one ad break in at least one video, I classify the YouTuber as "advertising YouTuber", and as "non-advertising YouTuber" otherwise.[21] This allows me to collect monetization data on the YouTuber level for all YouTubers in my dataset, but forgoes more fine-grained information on the video level.[22]

5.2. Measuring mainstream content

5.2.1. Definition

Similar to the procedure in Sun and Zhu (2013), I use the number of video views and the videos' *tags* to generate a measure for mainstream content. Video tags are descriptive keywords that let YouTube understand what a video is about and let viewers find the video via YouTube's search engine (e.g., a funny cat video might be given the tags *funny*, *cat*, and *pet*). YouTubers can enter such tags through a specific template when they upload their videos. In my main sample, the average number of tags per video is equal to 11.7, and the median number is equal to 12.[23]

[20] See `www.youtube.com/static?gl=de\&template=terms\&hl=en` (Oct 2018).

[21] The webscraper pauses for eight seconds before proceeding to the next video; crawling each video this way would take several years. Crawling twenty videos per YouTuber, in contrast, is feasible within three to four months. Appendix G.2.2 discusses the consequences of a potential measurement error.

[22] I further discuss the issue in Appendix G.1.

[23] If a video is not provided with tags, I generate tags from its title; this concerns 13.69% of videos in my main sample. Appendix B.2 shows that tags and video titles partially overlap. Also, my results are unaffected when I consider only the subsample of videos where tags are provided, or when I generate tags from video titles for all videos in my sample.

It is crucial to distinguish between video *keywords* and video *tags*. *Video keywords* are the most relevant and central terms and topics of a video in general and can (but need not) occur in the video's title, thumbnail, description, and tags. *Video tags*, in contrast, correspond to the descriptive terms that YouTubers can specify through a template when they upload a video. Thus, keywords and tags could, but need not, coincide. Since video tags are more likely to provide an accurate description of video content and are less likely to be strategically chosen than the keywords in a video's title or description (see Appendix B.1 for an extensive discussion), I use video *tags* throughout the empirical analysis.[24] Moreover, Appendix D.2.3 reports the results from an online survey experiment, where human coders verified the accuracy of video tags.

For each month and video category, I compute how many views a certain tag has attracted and rank them in descending order; the upper one percent of this distribution is then classified as "mainstream."[25] Note that it is important to consider each month and each video category separately. First, what is mainstream is likely to change over time, second, different video categories attract very different audiences, whose preferences need to be considered separately. Moreover, it is crucial to define mainstream content based on the universe of *all* active German YouTubers, i.e., before I exclude observations to construct the final dataset. Otherwise, I would compute the most mainstream tags within the sample of YouTubers selected for the main analysis (see Section 5.3), which is conceptually different.

Given the set of mainstream tags per month and category, I generate an indicator variable for mainstream content that is equal to one if a video is equipped with at least one mainstream tag, and zero otherwise. This is a plausible procedure for three reasons. First, tags reveal nothing about the proportion of time that videos devote to particular topics; e.g., a video with three tags is unlikely to spend one third of its time on each of the three topics. Hence, mainstream tags are informative about the extensive, but not the intensive margin of mainstream content. Second, YouTubers could combine more general with more specific tags to describe the same topic. Consider, for instance, a video on German shepherds with the tags *dog* and *German shepherd*. Even if *dog* is mainstream and *German shepherd* is not, it would be more natural to consider the entire video as mainstream, and not just half of it.[26] Third, the procedure is more likely to yield false positives than false negatives, which may result in estimates that are too conservative. In particular, if I can still document that advertising *reduces* YouTubers' incentives to duplicate mainstream content, the result is likely to represent a lower bound on the effect

[24] Appendix B.2 shows that my results are robust to using keywords from the videos' titles, too. See `https://support.google.com/youtube/answer/146402?hl=en` (Dez 2021) for further information on video tags.

[25] I ignore trivial tags that appear in the video categories' titles. E.g., I ignore *people* and *blog* for videos in the category "People & Blogs" and *science* and *technology* for videos in the category "Science & Technology."

[26] A similar point applies to compound tags. Consider, for instance, a video on making tinfoil hats with the tags *diy* and *diy tinfoil hat*. Although a tutorial on do-it-yourself tinfoil hats is quite obscure, the activity "diy" is not. Thus, if *diy* is a mainstream tag, I would classify the video as mainstream.

size. See Appendix B for further discussion and a battery of robustness checks on my measure for mainstream content.

5.2.2. Descriptives

Take the category "Science & Technology" in April 2015 as a concrete example. Videos are given $13,555$ different tags; the three most viewed ones are *diy*, *homemade*, and *selfmade*. The distribution of views over tags is heavily skewed: e.g., while the upper one percent of tags accounts for 45.1% of all views, the lowest ten percent account for just 0.02% of all views in that category and month (Figure A.20). The numbers are similar for other categories and other points in time.

Tables A.21 to A.23 illustrate the measure more generally. Table A.21 displays the top five tags that are most often classified as "mainstream" in each video category (e.g., the most frequent mainstream tags for the category "Sports" are *fitness*, *training*, and *soccer*). In addition, Table A.21 illustrates that the top tags are classified as mainstream in (almost) every month; e.g., *fitness* is mainstream in all 49 months that I consider in my analysis.

Consistent with that, I find that what is mainstream content partially persists over time. Table A.22 shows which fraction of mainstream tags in month t overlaps with the mainstream tags in months $t-1$, $t-2$, and $t-3$, respectively. On average, around a third of the mainstream tags were also classified as mainstream in the previous month, where the category "Nonprofit & Activism" exhibits the smallest, and "Let's Play" the largest overlap. Interestingly, the overlap in mainstream tags hardly diminishes over time, suggesting that each video category features a certain set of evergreen tags. In contrast to that, Table A.23 shows that the overlap of mainstream tags *across* video categories in a given month t is extremely small on average, further supporting my approach to consider each video category separately.

5.3. Final dataset

In a last step, I construct my final dataset. First, I define an appropriate observation period. The central event – the launch of the new ad break tool – took place in Nov 2015. Including videos uploaded between Jan 2013 and Jan 2017 into the final dataset yields a sufficient number of before and after observations. At the same time, this choice excludes both videos that are too old – and therefore not well comparable to more recent ones in terms of content or duration – as well as videos that were too "recent" on the date of data collection. By leaving at least nine months between the latest upload of a video and the data collection process (that started in Oct 2017) all videos in my dataset can be considered as "old", which minimizes any potential bias that may arise through the accumulative nature of some descriptive variables such as likes, dislikes, and views. Moreover, an observation period from Jan 2013 to Jan 2017 excludes the two big demonetization waves from 2017 and 2018 (see Section 3) that could have affected the YouTuber's content choice. Robustness checks on my main results using other observation periods and a summary of minor events that occurred between Jan 2013 and Jan 2017 are presented in Appendix F.1 and Appendix G.5.

Second, I determine which YouTubers to include. Following the outline from Section 4, I restrict the analysis to YouTubers below the ten minutes threshold before Nov 2015 (YouTubers A and B in the example), where I use a YouTuber's median video duration before Nov 2015 to define her "position" on the x-axis in Figure 1. Thus, I include only YouTubers whose median video duration before Nov 2015 is smaller than ten minutes into the final dataset. In addition, I include only YouTubers who uploaded at least one video before and after Nov 2015.[27] Finally, due to the "adpocalypse" (see Section 3), I exclude all videos from the category "News & Politics", since many of these videos were forcefully demonetized by YouTube. The final panel dataset includes $15,877$ YouTubers with $1,349,267$ videos over a time period of 49 months. Table A.19 summarizes all variables used in the main paper; Table A.20 summarizes all variables that only appear in the Appendix.

5.4. Illustrative evidence

Based on the final dataset, this section provides illustrative evidence confirming that the YouTubers whom I consider were unaware of the ten minutes trick before Nov 2015, that the launch of the new ad break tool made the ten minutes trick more apparent, and that YouTubers who were closer to the ten minutes threshold before Nov 2015 were more likely to exploit it. In Appendix E.1, I also provide video level evidence of an increase in the *actual* (not feasible) number of ad breaks per video.

5.4.1. Launch of the new ad break tool

Figure A.21 shows how the fraction of videos between ten and fourteen minutes develops for advertising and non-advertising YouTubers.[28] If the launch of the new ad break tool made the ten minutes trick more apparent, this fraction should increase for advertising, but not for non-advertising YouTubers after Nov 2015, which is indeed the case.

An increase in the number of ad breaks per video is likely to be noted and discussed by viewers. Hence, I screen the comment section of each video for the terms *advertising* ("Werbung" in German) and *mid-roll* on the one hand, and different spelling patterns of *10 minutes* on the other.[29] Then, I compute how often these terms appear per video per month. In line with my empirical strategy, Figures A.22 and A.23 demonstrate that debates on advertising and the ten minutes trick grow sharply after Nov 2015 for advertising,

[27] Appendix F.2 shows several robustness checks based on different selections of YouTubers.

[28] YouTubers who exploit the ten minutes trick to benefit from the increase in the feasible number of ad breaks per video are likely to produce videos that are *just* longer than ten minutes. Considering the share of videos between ten and fourteen minutes is thus more informative than considering the share of all videos that are ten minutes or longer.

[29] To be precise, I search for "10 minuten", "zehn minuten", "10min", and "10 min".

but not for non-advertising YouTubers.[30]

5.4.2. Closeness to the ten minutes threshold

Time trends Since further comparisons of advertising and non-advertising YouTubers are likely to raise selection issues, the remainder of the paper focuses on *advertising* YouTubers.[31] Figure A.24 illustrates how the fraction of videos between ten and fourteen minutes develops for advertising YouTubers that were close to and further away from the ten minutes threshold before Nov 2015. Since "closeness" – in terms of a YouTuber's median video duration before Nov 2015 – is a continuous measure, I cannot compare the trends of two distinct groups. Instead, Figure A.24 shows that the increase in videos just above the ten minutes threshold is steeper for YouTubers around the 75^{th} percentile of "closeness" (YouTubers like B in the stylized example) than for YouTubers around the 25^{th} percentile (YouTubers like A in the stylized example).

Distribution of video duration Next, I examine the distribution of video durations before and after Nov 2015 for the same two groups of YouTubers. First, if advertising YouTubers increase their videos' duration to benefit from the ten minutes trick, there should be bunching *just* behind the ten minutes threshold, and the bunching should be more pronounced for YouTubers closer to the ten minutes threshold. Second, if YouTubers were unaware of the ten minutes trick before Nov 2015, bunching should only be evident afterwards. Figures A.25 to A.28 show that this is the case. In addition, the histograms confirm that it is appropriate to focus on videos between ten and fourteen minutes: YouTubers who exploit the ten minutes trick start uploading videos that *just* enable them to increase the feasible number of ad breaks per video.[32]

While the evidence for bunching in Figure A.26 is quite clear, the bump behind the ten minutes threshold is also relatively small. A possible interpretation of this finding is that relatively few YouTubers expand the duration of their videos and differentiate from their competitors after Nov 2015, and that the majority of YouTubers remains unaffected. Although Figure A.26 presents illustrative evidence for only a small subgroup of YouTubers close to the ten minutes threshold, the evidence is in line with the results from Section 7.1, where I document that the reduced form estimates – which are conceptually similar to an intention-to-treat (ITT) effect – are considerably smaller than the Local Average Treatment Effect (LATE), suggesting that many YouTubers remain unaffected by the

[30] The ten minutes trick is also discussed in pertinent blogs and forums for YouTubers, see, e.g., `https://lited3m.com/en/learn-about-youtube-mid-roll-ads-here/`, `https://aliabdaal.com/triple-your-youtube-ad-revenue/`, `https://digiday.com/future-of-tv/creators-making-longer-videos-cater-youtube-algorithm/`, `https://www.quora.com/Why-do-YouTubers-try-make-their-videos-last-10-minutes` (March 2022).

[31] E.g., the prospect to generate ad revenue is likely to be a major motivation for advertising YouTubers, while non-advertising YouTubers are likely to participate on YouTube for the joy of video creation alone. See Appendix G for further discussion on the differences between advertising and non-advertising YouTubers.

[32] Appendix F.3 demonstrates that my main results are robust to considering *all* videos that are ten minutes or longer.

launch of the new ad break tool in Nov 2015. Thus, the evidence that I present in this paper should be understood as a clean documentation of the effect of advertising on content choice for YouTubers who are willing to adapt their content, and not as a general effect. Section 7.1 provides further discussion on this issue.

6. Empirical strategy

6.1. Baseline regression

This section formalizes the empirical strategy outlined in Section 4. I start by classifying a YouTuber as treated if she could increase her feasible advertising quantity after Nov 2015; in other words, the treatment corresponds to gaining the *option* to show more ads than before. In particular, I compute each YouTuber's share of videos between ten and fourteen minutes[33] before and after Nov 2015; if this share has increased by at least five percentage points, YouTuber i is assigned to the treatment group ($2,513$ YouTubers), and to the control group otherwise ($8,086$ YouTubers).[34]

The baseline difference-in-differences regression is given by

$$Mainstream_{vit} = \beta D_i * post_t + \phi_i + \phi_t + \phi_c + \tau_1 t_{it} + \tau_2 t_{ct} + \epsilon_{vit}, \tag{2}$$

where D_i indicates the treatment group, $post_t$ indicates all months after Nov 2015, ϕ_i, ϕ_t and ϕ_c are YouTuber, monthly, and video category fixed effects, t_{it} is a YouTuber specific and t_{ct} a video category specific linear time trend. The dependent variable $Mainstream_{vit}$ is a dummy variable equal to one if video v of YouTuber i in month t is given a mainstream keyword, and zero otherwise. Thus, I estimate a Linear Probability Model, and the parameter β measures the average percentage point change in the probability to duplicate mainstream content for YouTubers in the treatment relative to the control group.

6.2. IV regression

6.2.1. Model

An OLS estimation of equation (2) is unlikely to yield a causal estimate of the effect of an increase in the feasible advertising quantity on the YouTubers' probability to duplicate mainstream content for three reasons. First, YouTubers can self-select into the treatment group. This applies, for instance, to particularly money-loving YouTubers. If these YouTubers are at the same time more likely to duplicate mainstream content, the OLS estimate for β would be upward biased. Second, omitted YouTuber specific time-varying factors that are neither captured in the YouTuber specific linear time trend nor in YouTuber or monthly fixed effects may drive $Mainstream_{vit}$ and D_i at the same time. To stick with the example, some YouTubers may develop a taste for money over time.

[33]See Section 5 for a discussion on why considering the interval from ten to fourteen minutes is appropriate.
[34]See Appendix F.3 for robustness checks that use other classifications of treatment and control group.

If these YouTubers are more likely to duplicate mainstream content, the OLS estimate of β would, again, be upward biased. Finally, reverse causality may generate a spurious relationship between $Mainstream_{vit}$ and D_i. If, for instance, YouTubers who produce more mainstream content are more likely and more willing to increase their number of ad breaks per video, the OLS estimate for β would be upward biased, too.

To account for the endogeneity in the YouTubers' treatment status, I use YouTubers' *median video duration before Nov 2015* – denoted by $close_i$ – as an instrument for D_i. The first stage equation is given by

$$D_i * post_t = \pi close_i * post_t + \phi_i' + \phi_t' + \phi_c' + \tau_1' t_{it} + \tau_2' t_{ct} + u_{vit}. \tag{3}$$

Equations (2) and (3) are estimated by 2SLS.

Note that the instrument $close_i$ is likely to affect different YouTubers in different ways. In particular, some YouTubers' treatment status may be entirely unchanged. On the one hand, many YouTubers have no interest in increasing their feasible number of ad breaks per video; these YouTubers remain untreated, no matter how close to the ten minutes threshold they are.[35] On the other hand, some YouTubers are desperate to increase the feasible number of ad breaks per video; these YouTubers pursue the treatment, no matter how far away from the ten minutes threshold they are. Thus, the 2SLS estimate for β measures a Local Average Treatment Effect (LATE, see Angrist and Imbens, 1995), i.e., a weighted average of the individual treatment effects.

6.2.2. Instrument validity

The validity of $close_i$ as instrument for D_i hinges on two main requirements: relevance and exogeneity. First, $close_i$ must be correlated to D_i. It is plausible to assume that YouTubers who were closer to the ten minutes threshold before Nov 2015 can more easily produce videos that are ten minutes or longer afterwards, e.g., because they do not have to spend much additional effort. Figures A.24, A.26, and A.28 provide illustrative evidence for this claim. Moreover, a bivariate regression of D_i on $close_i$ yields a t-statistic of around 15, and the first stage F-statistic is above 140 throughout all specifications (see Section 7.1).

Second, $close_i$ must be exogenous to the dependent variable $Mainstream_{vit}$ and only operate through the single, known channel D_i. In other words, a video's duration – and especially its distance to the ten minutes threshold – must not correlate with whether the video covers mainstream topics or not. To support the plausibility of this assumption, I exploit the panel structure of my data for an event study. To this end, I consider the reduced form of equations (2) and (3) and interact the instrument $close_i$ with each monthly dummy, using Oct 2015 ($t = 34$) as baseline. This specification allows me to treat the coefficients of the interaction terms as the effect of $close_i$ on $Mainstream_{vit}$ relative to a base month just before the launch of the new ad break tool. The event study regression

[35]See Appendix G.2 for a related discussion on why some YouTubers do not monetize their content at all.

equation is given by

$$Mainstream_{vit} = \sum_{t=1}^{33} \gamma_t close_i * pre_t + \sum_{t=35}^{49} \gamma_t close_i * post_t + \phi_i'' + \phi_t'' + \phi_c'' + \tau_1'' t_{it} + \tau_2'' t_{ct} + v_{vit}. \quad (4)$$

If $close_i$ has no impact on $Mainstream_{vit}$ except through D_i, then all OLS estimates for $\gamma_{t\leq 33}$ should be close to zero and statistically insignificant.

Figure A.29 shows that this is indeed the case. The estimates for $\gamma_{t\leq 33}$ fluctuate around zero without a visible trend, and most of them are not statistically significant at the 5%-level. In contrast to that, the estimates for $\gamma_{t\geq 35}$, are negative and downward trending; moreover, most estimates are statistically significant at the 5%-level. See Appendix C for further validity checks of my empirical strategy.

7. Results

7.1. Main results

Table 1 presents the main results. Columns 1 to 3 show the results from the potentially biased OLS estimation of equation (2). The estimates are close to zero and not statistically significant despite the large sample size (10,599 YouTubers). In contrast to that, the estimates obtained by a 2SLS estimation of equations (2) and (3), displayed in columns 4 to 6, are negative and statistically significant at the 1%-level. According to these estimates, an increase in the feasible advertising quantity decreases the probability to duplicate mainstream content by about 20.9 to 23.7 percentage points. The effect size is considerable: it corresponds to 40% of a standard deviation in the dependent variable and to around 50% of its baseline value 0.448. The large difference between the OLS and the 2SLS estimates confirms the endogeneity concerns expressed earlier: YouTubers' self-selection into treatment, omitted YouTuber specific time-varying factors, as well as reverse causality may lead to an upward bias in the estimate for β if not taken into account.

The first stage diagnostics in columns 4 to 6 support the validity of my empirical strategy. Having been closer to the ten minutes threshold before Nov 2015 leads to a higher treatment probability: an additional unit of $close_i$ (i.e., an additional minute) increases the treatment probability by about 2.9 percentage points on average. The estimate is highly statistically significant. Moreover, F-statistics ranging from 144 and 150 demonstrate the strength of the instrument (Stock and Yogo, 2002; Kleibergen and Paap, 2006).

Finally, column 7 displays the estimate from a reduced form estimation of equations (2) and (3), i.e., the impact of a one unit increase in $close_i$ on the probability to duplicate mainstream content. Consistent with the results from the 2SLS regressions, the estimate is negative and statistically significant at the 1%-level. However, compared to the estimates in columns 4 to 6, the reduced form estimate is relatively small.[36] This is consistent with the illustrative evidence that I discuss in Section 5.4 and suggests that my empirical

[36]Note that the reduced form estimate corresponds to an ITT, while the 2SLS estimates correspond to a LATE.

strategy isolates a strategy change for a relatively small proportion of YouTubers, while the majority of YouTubers in my sample are unaffected. Similarly, the evidence is consistent with the fact that only about a fifth of all YouTubers whom I consider in the analysis are classified as treated. As argued above, the evidence that I present should ultimately be understood as a clean documentation of the effect of advertising on content choice for YouTubers who are willing to adapt their content, and not as a general effect that involves the entire market.

In sum, Table 1 leads to the conclusion that an increase in the feasible advertising quantity reduces the probability to duplicate mainstream content for YouTubers who are willing to comply. A detailed discussion of competition as a plausible economic mechanism follows in Section 8. Appendix B.6 shows that the main results are robust to alternative measures of mainstream content; Appendix F provides further robustness checks, including alternative observation periods, alternative selections of YouTubers, alternative classifications of the treatment group, and alternative definitions of the instrument.

Table 1: Main results

	OLS			2SLS			Red. Form
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$D_i * post_t$	.009	.004	.004	-.209***	-.237***	-.226***	
	(.008)	(.008)	(.008)	(.050)	(.050)	(.048)	
$close_i * post_t$							-.007***
							(.001)
First stage				.029***	.029***	.029***	
				(.002)	(.002)	(.002)	
F-statistic				144.13	143.29	150.65	
Time FE	X	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X	X
Category FE		X	X		X	X	X
Category Time Trend		X	X		X	X	X
YouTuber Time Trend			X			X	X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v by YouTuber i in month t covers mainstream content. Columns 1 to 3 display OLS, columns 4 to 6 2SLS, and column 7 reduced form estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7.2. Mechanics

This section demonstrates that the results from Section 7.1 are mostly induced by additional video uploads on behalf of the treated YouTubers, where the additional videos cover non-mainstream content. To this end, I consider four alternative dependent variables – the absolute number of video uploads per YouTuber per month, the absolute number of videos with mainstream content per YouTuber per month, the absolute number of keywords per video, and the absolute number of unique keywords per YouTuber per month – and re-run the 2SLS analysis from above.

Table 2 shows the results. Column 1 documents that an increase in the feasible ad quantity induces YouTubers to upload around 1.87 more videos per month; the effect is statistically significant at the 1%-level and corresponds to about 34% of a standard deviation in the dependent variable. The absolute number of mainstream videos per month does not increase, however: column 2 shows that though the estimate is positive, it is relatively small and not statistically significant. Taken together, this implies that YouTubers who could increase their ad quantity produce additional non-mainstream videos, and thereby reduce their probability to duplicate mainstream content on average.

A diminishing probability to duplicate mainstream content could also be driven by YouTubers reducing the number of tags per video. Column 3, however, shows that this is not the case; the estimate is positive and not statistically significant, suggesting that the absolute number of tags per video remains constant on average.

Finally, I consider the absolute number of unique tags per month in column 4, which I interpret as a measure for within-YouTuber diversity. The estimate is positive and weakly statistically significant; in particular, the absolute number of unique tags per month increases by about 7.33, which corresponds to about 12% of a standard deviation in the dependent variable. In sum, the results in Table 2 illustrate that YouTubers who could increase their feasible ad quantity become more productive and use these additional videos to differentiate from the mainstream. Further results on the mechanics of my main results are provided in Appendix E.3.

Table 2: Mechanics

	uploads per month (1)	mainstream per month (2)	no. tags per video (3)	unique tags per month (4)
$D_i * post_t$	1.87***	.268	.200	7.33*
	(.578)	(.411)	(.803)	(3.73)
First stage	.028***	.029***	.028***	.028***
	(.002)	(.002)	(.002)	(.002)
F-statistic	199.76	199.76	150.65	199.78
Time FE	X	X	X	X
YouTuber FE	X	X	X	X
Category FE			X	
Category Time Trend			X	
YouTuber Time Trend	X	X	X	X
YouTubers	10,599	10,599	10,599	10,599
Observations	241,905	241,905	1,067,542	241,905

Notes: Robust standard errors in parentheses. All estimates are 2SLS estimates. The dependent variable in column 1 is the number of video uploads of YouTuber i in month t. The dependent variable in column 2 is the absolute number of videos that cover mainstream content of YouTuber i in month t. The dependent variable in column 3 is the absolute number of tags of video v of YouTuber i in month t. The dependent variable in column 4 is the absolute number of unique tags that YouTuber i uses in month t. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7.3. Effect heterogeneity

A particular strength of my data is its size, which allows me to conduct several subgroup analyses. To this end, this section illustrates effect heterogeneity along three dimensions: First, I show that the average effect from Section 7.1 is driven by YouTubers with many subscribers. Second, I demonstrate that YouTubers who already lived in a niche before Nov 2015 are even more likely to differentiate their content afterwards. Third, I document that some video categories are more flexible regarding their typical video duration, which leads to heterogeneity on the first stage.

7.3.1. Heterogeneity along the subscriber count

Adapting video content entails (non-monetary) costs; e.g., YouTubers must deviate from the content they were producing before, which may force them to focus on topics that they are less intrinsically motivated to cover. The larger a YouTuber's audience, however, the higher her benefit from additional ad breaks and hence the probability that the additional ad revenue covers these costs.

To confirm that the effect of an increase in the feasible number of ad breaks on the probability to duplicate mainstream content is stronger for YouTubers with a high subscriber count, I split my sample at the $1,000$ subscriber threshold – which roughly corresponds to the median number of subscribers – and consider YouTubers with at least $1,000$, and YouTubers with fewer than $1,000$ subscribers separately.[37] Note that reverse causality prohibits including the subscriber count as an interaction term. If, for instance, YouTubers who upload much mainstream content have a larger audience, I would overestimate the subscribers' impact.

Table 3 shows the results from 2SLS regressions of equations (2) and (3). In line with the above considerations, the estimates are smaller (and less statistically significant) than their counterparts in Table 1 when considering YouTubers with few, and larger when considering YouTubers with many subscribers.

7.3.2. Heterogeneity along mainstream content

Next, I examine whether YouTubers who lived in a niche even before Nov 2015 are more or less likely to adapt their content when increasing the feasible number of ad breaks per video. To this end, I split the sample at the median monthly proportion of mainstream content before Nov 2015 (29.6%) and estimate equations equations (2) and (3) on the two subsamples, respectively.

Table 4 shows the results. The 2SLS estimates in columns 1 to 3 are negative, slightly larger than the average effects reported in Table 1, and highly statistically significant. The estimates in columns 4 to 6, in contrast, are close to zero and not statistically significant.

[37] YouTube has also recently disabled all YouTube channels with fewer than $1,000$ subscribers from monetization, arguing that this is a meaningful threshold for a channel to be considered "eligible" for ad revenues (Abou El-Komboz et al., 2022).

Table 3: Effect heterogeneity: subscribers

	$< 1,000$ subscribers			$\geq 1,000$ subscribers		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	-.050	-.144*	-.143*	-.259***	-.265***	-.255***
	(.082)	(.083)	(.083)	(.069)	(.066)	(.065)
First stage	.025***	.024***	.024***	.028***	.028***	.029***
	(.004)	(.004)	(.004)	(.003)	(.002)	(.003)
F-statistic	48.26	46.08	46.93	83.08	83.74	86.72
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	5,416	5,416	5,416	5,183	5,183	5,183
Videos	389,952	389,952	389,952	677,590	677,590	677,590

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v by YouTuber i in month t covers mainstream content. All columns display 2SLS estimates. The estimates are based on using the advertising YouTubers only. Columns 1 to 3 consider only YouTubers with fewer than $1,000$, columns 4 to 6 with more or equal than $1,000$ subscribers. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The first stage estimates are positive and highly statistically significant for all specifications. Thus, YouTubers who produced less mainstream content before Nov 2015 are likely to reduce that proportion even further, while YouTubers who have always produced much mainstream content tend to stick with it.

To interpret this result, note that what is mainstream is largely determined by the YouTube superstars who accumulate a large proportion of the overall number of views (Appendix B.4 elaborates on this). The type of content that the superstars produce is therefore extremely likely to be classified as mainstream. In other words, YouTubers with a large proportion of mainstream content before Nov 2015 tend to be well-known YouTubers with a large subscriber base (Figure A.30 illustrates). Relatively well-known YouTubers, in turn, are less likely to adapt their content when increasing their advertising quantity, because these YouTubers run a smaller risk of losing their audience to a close competitor. Competition as a potential driver of my main results is further discussed in Section 8. Note that the results do not contradict the findings from Section 7.3.1: when I restrict that analysis to the upper ten percent of YouTubers in terms of subscribers ($27,643$ subscribers) – i.e., those who are plausibly well-known – the impact of an increase in the feasible ad quantity is smaller and not statistically significant.

7.3.3. Heterogeneity along video categories

Finally, I demonstrate that some video categories are more flexible regarding their typical video duration, leading to further effect heterogeneity.[38] To this end, I estimate equa-

[38]E.g., a music clip typically takes between three and five minutes and cannot be easily extended to ten minutes. Similarly, a comedy video becomes boring if it does not get the gag across.

Table 4: Effect heterogeneity: mainstream content

	below median			above median		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	-.235***	-.250***	-.247***	-.012	-.055	-.048
	(.076)	(.075)	(.074)	(.069)	(.066)	(.065)
First stage	.028***	.028***	.028***	.027***	.027***	.027***
	(.003)	(.003)	(.003)	(.003)	(.003)	(.003)
F-statistic	66.66	66.08	67.86	64.30	64.79	68.45
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	5,295	5,295	5,295	5,304	5,304	5,304
Videos	501,083	501,083	501,083	566,459	566,459	566,459

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v by YouTuber i in month t covers mainstream content. All columns display 2SLS estimates. The estimates are based on using the advertising YouTubers only. Columns 1 to 3 consider only YouTubers below the median monthly proportion of mainstream content before Nov 2015 (29.6%), columns 4 to 6 consider only YouTubers above and equal to the median monthly proportion of mainstream content before Nov 2015. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

tions (2) and (3) by 2SLS on fourteen subsamples comprising all videos from a particular category, respectively.

The results in Table 5 reveal effect heterogeneity in terms of the first and also in terms of the second stage. Intuitively, the first stage estimate is close to zero for the categories “Music”, “Comedy”, and “Let’s Play.”[39] The first stage estimate is largest for the categories “Cars & Vehicles”, “Pets & Animals”, and “Sports.” Hence, videos from these categories can either be most easily extended to ten minutes or more, or YouTubers who have the strongest desire to increase their feasible number of ad breaks self-select into these categories. The first stage estimate for the remaining categories is similar to the results from Section 7.1.

For the discussion of the second stage estimates, I focus on the categories with a first stage F-statistic above 10. Consistent with the main results from Section 7.1, all estimates are negative; their size ranges from -0.0762 (“Cars & Vehicles”) to -0.922 (“Film & Animation”). The estimates are statistically significant for the categories “Film & Animation”, “People & Blogs”, and “Entertainment”, which are also the categories with the largest number of observations. Hence, in addition to heterogeneity on the first stage, the video categories differ in the extent to which the video content is adapted. There are, again, two plausible explanations. First, it could be easier to create videos that cover non-mainstream content for some categories; in other words, the effect heterogeneity is driven by category specific differences (that are not captured by X_{vit}). Second, YouTubers who

[39] “Let’s Play” videos are often based on how YouTubers complete video game levels, many of which include a time constraint.

are more creative or more willing to try out something new might self-select into the video categories "Film & Animation", "People & Blogs", and "Entertainment" whose second stage effect is strongest.

8. Competition

This section studies competition as a plausible economic mechanism behind the results from Section 7. In particular, I show that mainstream content is provided by many competing YouTubers; hence, viewers could easily switch to a competitor if a YouTuber increased her advertising quantity. Switching is less likely, however, if the YouTuber differentiates her content from the mainstream, moves to a niche, and thus softens competition in the ad "price."[40] The empirical results are thus in line with the predictions from the stylized theoretical framework in Appendix A.

8.1. Definition of competitive pressure

I start by generating a measure for competitive pressure. To this end, I consider each tag per month and video category and compute the absolute number of YouTubers who use it in their videos. Then, I consider each individual video and compute the sum of competing YouTubers over all of its tags; i.e., I calculate the absolute number of competing options to each video.

Take the category "Science & Technology" in April 2015 as a concrete example again. The three most provided tags are *german*, *test*, and *review*, i.e., videos that are equipped with such tags are likely to experience high competitive pressure.[41] Figure A.31 illustrates that the supply of tags is heavily skewed: e.g., the upper one percent of tags is provided by 17.4%, while the lowest ten percent of tags is provided by just 4.4% of all videos. The numbers are similar for other categories and other points in time. Regarding the entire sample of advertising YouTubers, I find that the median number of competing options to each video is equal to 139, and the mean is equal to 629.88, with a minimum of 0 and a maximum of $58,558$ competing options.[42]

Note that the measures for competitive pressure and mainstream content as defined in Section 5.2 are conceptually distinct: while my measure for mainstream content captures *demand* for video content, competitive pressure captures content *supply*. E.g., it is possible that a small number of videos provides some extremely popular and thereby widely viewed content; in this case, videos covering mainstream content would not experience high competitive pressure. Conversely, it could be that many videos provide content that

[40] Appendix G.4 elaborates on viewer switching.

[41] Note that the three most provided tags in this month and video category are different from the three most viewed ones in Section 5.2.

[42] Recall that the average number of tags per video is equal to 11.7, hence, the average number of competing options per tag is roughly equal to 53.83. As argued above, however, the distribution of competing options over tags is heavily skewed, i.e., tags like *test* and *review* exhibit a number of competing options that is much larger, and more obscure tags a number that is much lower than the average.

Table 5: Effect heterogeneity: video categories

	Film and Animation (1)	Cars and Vehicles (2)	Music (3)	Pets and Animals (4)	Sports (5)	Travel and Events (6)	Let's Play (7)	People and Blogs (8)	Comedy (9)	Entertainment (10)	How To and Style (11)	Education (12)	Science and Techn. (13)	Nonprofit and Activ. (14)
$D_i * post_t$	-.922***	-.076	2.71	.170	-.105	.127	-.658	-.243**	-.548	-.255**	-.236	-.100	-.651	-.227
	(.341)	(.095)	(12.38)	(.179)	(.134)	(.139)	(.852)	(.109)	(.936)	(.122)	(.189)	(.258)	(.761)	(.197)
First stage	.022***	.044***	-.003	.037**	.035***	.030***	-.009	.031***	.020	.028***	.028***	.026**	.021	.047
	(.006)	(.010)	(.017)	(.017)	(.009)	(.010)	(.009)	(.006)	(.023)	(.005)	(.009)	(.012)	(.017)	(.031)
F-statistic	13.49	21.08	0.04	4.99	15.34	8.41	0.98	32.15	0.72	28.32	10.99	4.99	1.70	2.32
Time FE	X	X	X	X	X	X	X	X	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X	X	X	X	X	X	X	X	X
YouTuber Time trend	X	X	X	X	X	X	X	X	X	X	X	X	X	X
YouTubers	2,302	1,543	1,382	838	1,776	1,764	827	4,724	769	4,207	1,462	906	868	430
Videos	93,616	87,945	25,512	25,963	96,645	62,041	82,650	200097	15,831	224,739	77,593	43,446	13,951	11,622

Notes: Robust standard errors in parentheses. The dependent variable is a dummy variable equal to 1 if video v of YouTuber i uploaded in month t covers mainstream content. Each column displays the results of a 2SLS estimation including only the observations from one particular video category. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

hardly anybody watches; in this case, videos would experience competitive pressure although they do not cover mainstream content. However, it is more plausible to assume that content in high demand is also widely supplied by the YouTubers – indeed, I find that the correlation between my measure of mainstream content and competitive pressure is equal to 0.27.

Competitive pressure is likely to be influenced by the size of a YouTuber's competitors, i.e., it makes a difference whether the same type of video content is also provided by a relatively unknown YouTuber or by a YouTube superstar. Put differently, I expect that an increase in the feasible advertising quantity induces YouTubers to avoid competition with the big shots, rather than the small fry. To take this into account, I weight my measure for competitive pressure with the number of subscribers of each competing YouTuber.[43] This metric can also be interpreted as the intensive margin, while the baseline measure can be interpreted as the extensive margin of competitive pressure. The correlation between my measure of mainstream content and weighted competitive pressure is equal to 0.35.

8.2. Results

Table 6 shows the 2SLS results from using log competitive pressure and log weighted competitive pressure as dependent variables in equation (2). The estimates are negative and highly statistically significant for all specifications. According to the estimates in columns 1 to 3, an increase in the feasible advertising quantity leads to a 67 to 76 percent decrease in the absolute number of competing options per video relative to the control group; according to the estimates in columns 4 to 6, the size of the competing YouTubers' subscriber bases diminishes by about 143 to 151 percent.[44] Thus, competitive pressure shrinks both on the extensive and on the intensive margin.[45]

The results are in line with the findings that I present in Section 7.3.2 as well as in Appendix B.4 and Appendix E.4. Section 7.3.2 shows that YouTubers who had always produced a large proportion of mainstream content – i.e., relatively well-known trend-setters – were less likely to move to a niche than YouTubers whose proportion of mainstream content was smaller. Plausibly, such trend-setting YouTubers experience relatively low competitive pressure, whereby their incentive to reduce competition and move to a niche after increasing their advertising quantity is lower, too. Indeed, when I split my sample at the same margin as in Section 7.3.2, I find that YouTubers with a small proportion of mainstream content before Nov 2015 substantially reduce their competitive pressure both on the extensive and on the intensive margin, while YouTubers with a large proportion

[43]E.g., consider a video with just one tag, where this tag is also provided by three other YouTubers in that month and category. My baseline measure for competitive pressure would be equal to three. If the three competitors accumulate a total number of 100 subscribers, weighted competitive pressure would be equal to 100; if they accumulated a total of $10,000$ subscribers, weighted competitive pressure would be equal to $10,000$.

[44]Note that a relative decrease by more than 100% is possible if the metric increases for the control group, but diminishes for the treatment group.

[45]Recall that the number of tags per video remains constant (see Section 7.2), so the result cannot be mechanically induced.

of mainstream content are much less affected (Table A.24). Analogously, Appendix B.4 demonstrates that what is mainstream is largely determined by the YouTube superstars. Producing more niche content therefore corresponds to differentiating from the superstars, which is especially in line with the finding that competition on the intensive margin decreases. Finally, Appendix E.4 supports the argument with findings from the YouTubers' audience: if the feasible ad quantity goes up, viewer fluctuation goes down, matching the idea that YouTubers reduce competitive pressure to prevent their audience from switching to a close competitor. Similarly, Appendix E.6 shows that YouTubers who move into a niche tend to increase the proportion of German-speaking videos and thereby reduce competitive pressure from English-speaking alternatives.

Table 6: Competition

	log(competitive pressure)			log(w. competitive pressure)		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	-.672***	-.760***	-.693***	-1.594***	-1.612***	-1.519***
	(.204)	(.199)	(.194)	(.286)	(.271)	(.263)
First stage	.029***	.029***	.029***	.029***	.029***	.029***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
F-statistic	144.35	143.50	150.86	143.65	142.91	150.29
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,597	10,597	10,597	10,591	10,591	10,591
Videos	1,062,993	1,062,993	1,062,993	1,057,360	1,057,360	1,057,360

Notes: Robust standard errors in parentheses. The dependent variable in columns 1 to 3 is the log number of competitors who also use one of the tags of video v in a given month t. The dependent variable in columns 4 to 6 is the log number of competitors, weighted by their respective number od subscribers, who also use one of the tags of video v in a given month t. All estimates are 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

9. Video quality

How does an increase in the feasible advertising quantity affect video quality? Seminal theory on vertical product differentiation illustrates a positive correlation between prices and product quality, where vertical differentiation serves as a tool to soften price competition (e.g., Shaked and Sutton, 1982, 1983; Ronnen, 1991; Armstrong and Chen, 2009).[46] Following this line of thought, one would expect that an increase in the feasible number of ad breaks per video – which is similar to a price increase for viewers – leads to an *increase* in video quality. In addition, having multiple ad breaks per video implies that each viewer

[46] Armstrong and Weeds (2007); Weeds (2013) study vertical product differentiation in media markets in a setting without horizontal product differentiation.

is more valuable than before, whereby the incentive to provide high quality content goes up. On the other hand, Section 8 demonstrates that YouTubers who increase their feasible ad quantity reduce competitive pressure via horizontal product differentiation, and – as argued by Bourreau (2003) – receding competition on the horizontal axis could *reduce* YouTubers' incentive to provide quality content. It is thus a priori unclear whether an increase in advertising quantity enhances or diminishes video quality.

Analyzing the impact of advertising quantity on video quality is challenging for three reasons. First, it is hard to find appropriate measures for video quality. There exist only a few quality measures (e.g., visual and sound quality) that every viewer could objectively agree on – everything else is a matter of taste. Second, objective quality measures are often unobserved; e.g., there exists no information on the visual quality of a YouTube video beyond video resolution, and the variation in video resolution is extremely limited. Third, as argued above, an increase in the feasible ad quantity comes along with a substantial increase in horizontal product differentiation, whereby it is difficult to isolate the impact of advertising on video quality.

I tackle these challenges in three steps. First, I generate two potential measures for video quality that I support with information from an online survey experiment, where I let human coders watch and rate a subset of YouTube videos (see Appendix D for details on the survey). Based on that, I estimate the impact of advertising on video quality in a second step. Finally, to interpret these results, I consider a subsample of videos of which the *actual* number of ad breaks is known (see Appendix E.1.2 for details). Specifically, information on the actual number of ad breaks per video before and after Nov 2015 helps me to distinguish the effect of advertising from the effect of horizontal product differentiation and reduced competitive pressure on video quality.

9.1. Measuring video quality

Likes and dislikes An intuitive way to measure video quality is to exploit likes and dislikes, e.g., in terms of the proportion of positive ratings:

$$proportion\,positive\,ratings_{vit} = \frac{likes}{(likes + dislikes)}. \tag{5}$$

The advantage of this measure is that it is observable, easy to interpret, and independent of a simultaneous increase (or decrease) in videos views.[47] Moreover, liking and disliking is the most common type of user engagement – e.g., there are typically more likes and dislikes than comments – whereby a measure based on these ratings is likely to reflect viewers' opinion very well.

To support the validity of this measure, I compute its pairwise correlation with each of the five dimensions of video quality that I consider in the online survey experiment. Reassuringly, Column 1 in Table A.25 shows that each pairwise correlation is small but

[47] See Appendix E.2 for a discussion on video views.

positive.

Sentiment analysis An alternative approach to measure video quality is to conduct a sentiment analysis on the video comments. To this end, I employ the *SentiWS*, a publicly available and widely-used dictionary for sentiment analyses of German-speaking text (Remus et al., 2010). The dictionary features around 30, 000 terms that express positive or negative sentiment.[48] In particular, each term in the *SentiWS* is assigned a score ranging from -1 to 1, where negative values correspond to negative, and positive values to positive sentiment.

To generate my quality measure, I aggregate all terms from all comments of a video into a "bag of words". Then, I assign each term w within each bag of words the score that it is given by the *SentiWS*. If a term is not in the dictionary, it is given the neutral score 0; analogously, if a video is not commented at all, it is given a neutral score, too.[49] Finally, I sum up all scores within each bag of words, normalize this sum with the total number of terms W within each bag, and hence compute the average sentiment score of a term posted underneath a video:

$$sentiment\ score_{vit} = \frac{1}{W} \sum_{w} score_{w}. \tag{6}$$

Again, I support the validity of this measure by computing its pairwise correlation with each of the five dimensions of video quality that I consider in the online survey experiment; column 2 in Table A.25 confirms that all pairwise correlations are small but positive.

9.2. Regression analysis

To examine the effect of an increase in the feasible advertising quantity on video quality, I replace the dependent variable in equation (2) with expressions (5) and (6), respectively, and estimate equations (2) and (3) by 2SLS.

Table 7 shows the results. When I consider the proportion of positive ratings as dependent variable (columns 1 to 3), the 2SLS estimates are negative and statistically significant at the 1%-level. According to these estimates, an increase in the feasible advertising quantity leads to a four percentage point reduction in the fraction of positive ratings. The effect size corresponds to around 25% of a standard deviation in the dependent variable and to 4.4% of its baseline value 0.91. When I measure video quality in terms of its sentiment score, however, the 2SLS estimates are equal to zero (columns 4 to 6). One potential explanation for this result is that it takes less effort to like or dislike a video than to post a comment; hence, when the effect of an increase in the feasible advertising quantity on video quality is small, the proportion of positive ratings might be more sensitive to this change than a video's comment section.

[48] To be precise, I use the *SentiWS_v2.0*, which contains 16, 401 positive and 17, 807 negative adjectives, adverbs, verbs, and nouns that explicitly and implicitly express sentiment.

[49] The regression results in Table 7 are robust to excluding all videos that are not commented at all.

As argued above, it is unclear if the negative effect on the proportion of positive ratings reflects a decrease in video quality due to receding competition or if it driven by viewers' aversion to ads. E.g., it is possible that viewers dislike a video to express their distaste of advertising, and this action is unrelated to video quality. To disentangle the channels, I consider a random subsample of 500 YouTubers and $52,462$ videos of which the *actual* number of ad breaks is known (see Appendix E.1.2 for details). In particular, I regress my dependent variables (5) and (6) on the *actual* number of ad breaks per video before and after Nov 2015, respectively, where I also control for the observable quality traits that I could capture with my online survey experiment. The idea is that a decrease in competitive pressure did not play a role before Nov 2015; hence, a negative relationship between the actual number of ad breaks per video and video quality only after Nov 2015 – but not before – would suggest that receding competition and a concomitant decrease in video quality drives the results from Table 7. Incorporating observed quality traits as control variables reduces omitted variable bias: a high quality video, for instance, is likely to generate a large proportion of positive ratings, but could also be prone to feature many ad breaks.

Table A.26 shows the results. Panel A demonstrates that the relationship between the proportion of positive ratings and the actual number of ad breaks per video is very small and positive before Nov 2015 (columns 1 to 3), but the OLS estimates are not statistically significant in any specification. In contrast to that, the OLS estimates are negative, roughly twice as large, and statistically significant at the 1%-level when I consider the time period after Nov 2015 (columns 4 to 6). Hence, the results in Panel A suggest that advertising quantity as such plays a minor role for the proportion of positive ratings, and that the results in Table 7 are driven by the receding competition after Nov 2015. Note, however, that Table A.26 only presents correlations, and that the number of observations is small. In line with the results from columns 4 to 6 in Table 7, Panel B in Table A.26 shows that there is no relationship between the actual number of ad breaks per video and a video's sentiment score either before or after Nov 2015. In sum, I find that – if any thing – an increase in the feasible advertising quantity has a negative impact on video quality, where the negative effect is likely to be driven by reduced competitive pressure.

10. Conclusion

This paper demonstrates that an increase in the feasible advertising quantity leads to an increase in content differentiation between several thousand YouTubers. In particular, I show that the option to raise the number of ad breaks per video reduces the YouTubers' probability to duplicate mainstream content by about twenty percentage points, because YouTubers differentiate their content from their competitors to soften competition in the ad "price." The results are in line with recent literature that acknowledges a conceptual similarity of advertising and subscription prices as determinants of content differentiation in media markets, and they provide more general empirical evidence of the interplay of

Table 7: Video quality

	prop. pos. ratings			*sentiment score*		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	-.041***	-.041***	-.040***	-.000	-.000	.000
	(.011)	(.011)	(.011)	(.001)	(.001)	(.001)
First stage	.028***	.028***	.028***	.029***	.029***	.029***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
F-statistic	130.86	130.26	137.13	144.13	143.29	150.65
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,594	10,594	10,594	10,599	10,599	10,599
Videos	990,476	990,476	990,476	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. The dependent variable in columns 1 to 3 is the proportion of positive ratings of video v by YouTuber i in month t as defined by expression (5). The dependent variable in columns 4 to 6 is the sentiment score of video v by YouTuber i in month t as defined by expression (6). All estimates are 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. Note that $77,066$ videos have not received any likes or dislikes and are thus excluded from the analysis in columns 1 to 3. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

price competition and horizontal product differentiation. Crucially, my empirical strategy isolates a strategy change for a relatively small proportion of YouTubers. The results should therefore be be understood as a clean documentation of the effect of advertising on content choice for YouTubers who are willing to adapt their content, and not as a general effect involving the entire platform.

Regarding its tremendous reach, analyses of YouTube are interesting and relevant by themselves. However, the empirical findings may also play a role beyond the specific context of this paper; in particular, the effect of advertising on content choice is likely to occur in other media markets where consumers perceive advertising as a nuisance and many competing media outlets with sufficient scope to adapt their content co-exist. E.g., television and radio broadcasters as well as (online) newspapers might further differentiate their content from their competitors if upper bounds on advertising quantity are lifted. Although YouTubers differ somewhat from conventional media outlets – e.g., channels are often managed by a single person – YouTubers often run their channel like a proper business, whereby the qualitative insights from this paper are likely to be externally valid.

The paper advances persistent debates on the effect of advertising on content differentiation in media markets. In particular, showing that advertising does *not* lead to the duplication of mainstream content entails two implications for present policies. First, advertising quantities are often restricted in an attempt to protect consumers.[50] The Audiovisual Media Services Directive, for instance, requires that the proportion of television

[50] See www.ofcom.org.uk/__data/assets/pdf_file/0021/19083/advertising_minutage.pdf (Dec 2018).

advertising and teleshopping spots within a given clock hour shall not exceed 20% (Article 23 §1). My paper demonstrates that consumers may *benefit* from advertising, because it increases content differentiation. This does not imply that regulators should give up on restricting advertising entirely. Rather, they ought to take this additional effect into account and might thereby derive less restrictive ad policies.[51] Similarly, public interventions in television markets – i.e., public service broadcasters – grow from the claim that advertising funded broadcasting fails to serve all viewers' preferences over content (Armstrong and Weeds, 2007). My results controvert this argument: advertising leads to *more* content differentiation. Thus, while valuable contributions to culture, education, and the public discourse certainly justify public service broadcasting, concerns about content duplication by advertising funded broadcasters do not.

My paper is limited in four respects that open up perspectives for further research. First, although I present competition in the ad "price" as a plausible mechanism for my main results and rule out a YouTuber learning effect, I cannot exclude the possibility that there are other potential mechanisms. For instance, YouTubers might not only soften competition to other YouTubers and acquire a more stable audience when they upload less mainstream content, but the characteristics of their viewers may change, too. Viewers of less mainstream content could be generally less ad averse or have a higher valuation of the video content such that they are willing to endure more ads.

Second, I cannot evaluate the effect of advertising on welfare, because I lack measures for consumer and producer surplus. Although I demonstrate that advertising leads to more content differentiation – which is likely to raise consumer surplus (Brynjolfsson et al., 2003) – the viewers must also pay an increased ad "price" in terms of an increased advertising quantity, which works into the opposite direction. Since I obtain no estimates for the viewers' ad aversion, my setup does not answer which effect overweights. On the producer side, I remain agnostic about the effect of advertising on the surplus of YouTube itself, the YouTubers, and the advertisers. YouTube as a platform is likely to benefit from advertising, though. Advertising leads to more content differentiation, which attracts more viewers; more viewers, in turn, generate more ad revenue. Likewise, the YouTubers' surplus benefits from an increase in ad revenue; it is, however, unclear how their utility from covering different topics than before is affected. Finally, the advertisers' surplus may go up or down. On the one hand, a higher ad quantity makes it more likely that potential customers click on their ads and buy their products. On the other hand, the advertisers cannot influence where exactly their ads appear, whereby it is unclear how well the audience is targeted. Hence, it is possible that the additional costs of advertising surmount the additional revenues.

[51] There is nothing special about the advertising channel per se; in principle, classic price competition between media outlets is likely to induce similar effects on content differentiation (as demonstrated by Seamans and Zhu, 2014). However, many media outlets do not charge monetary prices, and even if they do, it is probably more difficult to justify price than ad quantity regulation. In other words, stimulating competition in advertising quantities is an effective option to induce content differentiation in a media market.

The third limitation of my paper is the implicit assumption that YouTubers perceive the per-view-revenue from advertising as given.[52] In other words, YouTubers neither have an incentive to strategically restrict their ad quantity nor to target a specific audience to jack up the ad price. Hence, my results cannot be extrapolated to an environment where the per-view-revenue increases if a narrow and specific audience is attracted. Plausibly, the effect of an increase in the feasible number of ad breaks was higher in this case, because the YouTubers had an additional incentive to differentiate their content from the mainstream. A downward sloping (inverse) demand curve on behalf of the advertisers, however, might have the opposite effect. If the price per ad declines in quantity, the effect of increasing advertising quantity on content differentiation would probably be reduced.[53]

Finally, I do not discuss any concerns related to commercial media bias, i.e., advertisers exerting pressure on the media outlets' content decisions. As argued, however, there is no direct relationship between YouTubers and advertisers whose ads appear as breaks during the videos, so the issue is of small importance in my application. Yet, it is possible that commercial media bias arises from product placement contracts between advertisers and YouTubers, for instance, if the advertisers want their products to appear within friendly and uncontroversial videos; studying the relationship between product placement and commercial media bias on YouTube would be an interesting question for further research.

[52] See Appendix G.3 for further discussion.

[53] To the best of my knowledge, however, there is no such relationship between ad price and ad quantity on YouTube.

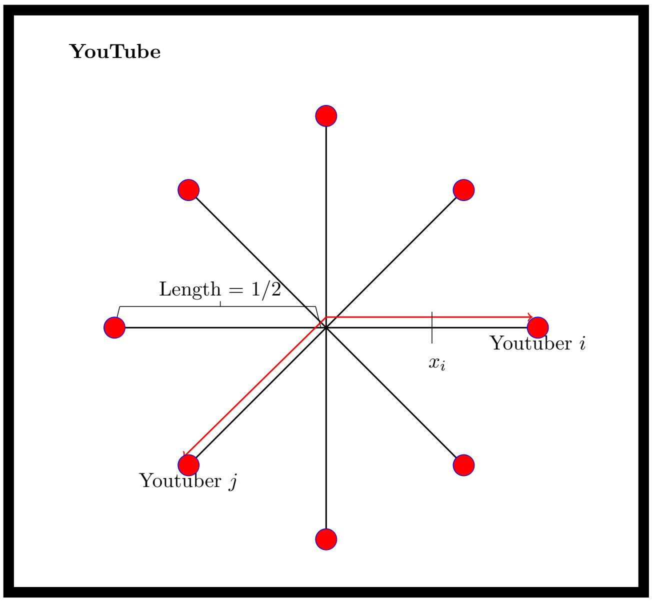

Figure A.1: Spokes model with $N = 8$ spokes (types of video content), $n = N$ YouTubers allocated at the origin of each spoke ($y = 0$), a viewer located on x_i, and the comprising platform YouTube

A. Stylized theoretical framework

This section uses a simple theoretical framework to illustrate the economic incentives of YouTubers, viewers, and YouTube itself. In particular, I show that competing YouTubers differentiate from each other to reduce competition if advertising quantities can be freely determined, and that they tend to duplicate mainstream content when advertising quantities are limited. To this end, I employ the spokes model of non-localized spatial competition as introduced by Chen and Riordan (2007). The spokes model generalizes the traditional Hotelling duopoly to an arbitrary number of firms. In contrast to the Salop circle, firms do not only compete with their neighbors, but with all other firms in the market. The spokes model is thereby very well suited to analyze horizontal product differentiation on YouTube, where thousands of YouTubers compete over viewers.

A.1. Setup

Consider a market that is constituted of $N \geq 2$ spokes (Figure A.1). Each spoke is of length 1/2; it terminates at the center of the radial network and originates at the other end. The spokes represent the N possible varieties of a product, where variety $i = 1, ..., N$ is located on spoke i. The varieties are physically identical, but differentiated by their location. In this setting, a variety can be interpreted as a particular type of video content.

A.1.1. Economic agents

YouTubers There are n symmetric firms (YouTubers), each producing a single type of video content. A YouTuber who is allocated on spoke i produces video content i. There can be at most one YouTuber per spoke. For simplicity, I assume that $n = N$, i.e., each possible type of video content is being produced, and each spoke is occupied by exactly one YouTuber.[54]

In contrast to Chen and Riordan (2007), I do not assume that the location of YouTuber i is fixed at $y_i = 0$ (i.e., the origin of spoke i), but that she can choose any location $y_i \in \left[0; \frac{1}{2}\right]$. Also, the YouTubers do not charge subscription fees for their video content, but finance their operations through advertising revenue instead. In particular, each YouTuber can choose an advertising quantity $a_i \geq 0$ and earns revenue $r > 0$ per ad per view.[55] Advertising quantities are determined after the video has been produced, i.e., after YouTuber i has chosen her location y_i. Each YouTuber incurs fixed costs $F^Y > 0$ and must divert a share s of her advertising revenue R_i to the Platform YouTube. Marginal production costs are normalized to zero. The profit function of YouTuber i is thus given by

$$\pi_i^Y(a_i, a_{j\neq i}, y_i, y_{j\neq i}) = \underbrace{D_i(a_i, a_{j\neq i}, y_i, y_{j\neq i}) * a_i * r}_{=R_i} * (1-s) - F^Y, \tag{7}$$

which she maximizes by choosing her location y_i and her advertising quantity a_i. YouTubers enter the market if $\pi_i^Y \geq 0$ and abstain otherwise.

Viewers Consumers (viewers) are uniformly distributed on the network of spokes. Their total mass is normalized to one. As in Chen and Riordan (2007), viewers have a unit demand for videos and like exactly two types of video content. In particular, a viewer located on spoke i always likes video content i, and each of the remaining $N-1$ varieties is equally likely to be the second type of video content that she likes. Viewers have valuation $v > 0$ for these two types of video content, and valuation $v = 0$ for all others. I assume that v is sufficiently high to ensure that the market is covered.

To watch a YouTuber's video content, viewers must travel through the spokes. Thus, if a viewer on spoke i wants to watch video content from another spoke $j \neq i$, she must travel through the center of the network. Viewers suffer if they have to travel, where a longer distance between the location of a viewer and the location of a YouTuber can be interpreted as a larger mismatch between the viewer's preferences and the type of video content that the YouTuber provides (Figure A.1 illustrates). In contrast to Chen and Riordan (2007), I assume that viewers' transportation costs are quadratic to avoid the well-known problem of multiple equilibria (d'Aspremont et al., 1979). Viewers also suffer from advertising according to their degree of ad aversion γ, which captures the idea that disruptive ad breaks are conceptually similar to a price for watching the video content.

[54] The original spokes model features $n \leq N$ firms, i.e., some spokes can be unoccupied. This allows to study endogenous entry, but also complicates the analysis. To maintain the illustrative purpose of this section, I focus on the case where $n = N$.

[55] I assume that r is independent of a_i and the same for each YouTuber i. See Appendix G for further discussion.

Thus, the utility of a consumer located at x_i is equal to

$$u_i = v - \gamma a_i - t|x_i - y_i|^2 \tag{8}$$

if she watches the video content of YouTuber i located at y_i. If the consumer at x_i watches her other preferred type of video content $j \neq i$, her utility is given by

$$u_j = v - \gamma a_j - t|1 - y_j - x_i|^2. \tag{9}$$

Platform There also exists a Platform (YouTube) that comprises the entire market and provides the necessary digital infrastructure for YouTubers and viewers to come together. One could, for instance, interpret YouTube as the provider of the spokes network on which viewers and YouTubers are allocated. To avoid confusion in the notation, I will refer to "YouTube" as "Platform" for the remainder of the section. Viewers do not have to pay a price to use the Platform. YouTubers, however, must discharge a share $s \geq 0$ of their advertising revenue R_i to use the Platform's infrastructure. The Platform incurs fixed costs $F^P > 0$. Marginal costs are normalized to zero. The Platform's profit function is given by $\pi^P = s \sum_i R_i - F^P$, which she maximizies by choosing s subject to $\pi_i^Y \geq 0$, i.e., she ensures that all YouTubers enter the market. The Platform operates if $\pi^P \geq 0$, which I assume to be the case in every symmetric equilibrium where the YouTubers enter the market.

A.1.2. Timing

The timing of the game is as follows:

1. The Platform determines $s \geq 0$.
2. YouTubers $i = 1, ..., N$ simultaneously choose their location $y_i \in \left[0; \frac{1}{2}\right]$.
3. YouTubers $i = 1, ..., N$ simultaneously choose their advertising quantity $a_i \geq 0$.

I solve the game by backwards induction and focus on symmetric subgame-perfect Nash equilibria.

A.2. Unlimited advertising quantity

I will first consider a setting where YouTubers' advertising quantity is unlimited, similar to the situation on YouTube after Nov 2015. The case of limited advertising quantity is studied in section A.3 below.

Demand specification As a first step, I specify the demand function of YouTuber i by determining the set of indifferent viewers. Consider a viewer located at x_i on spoke i

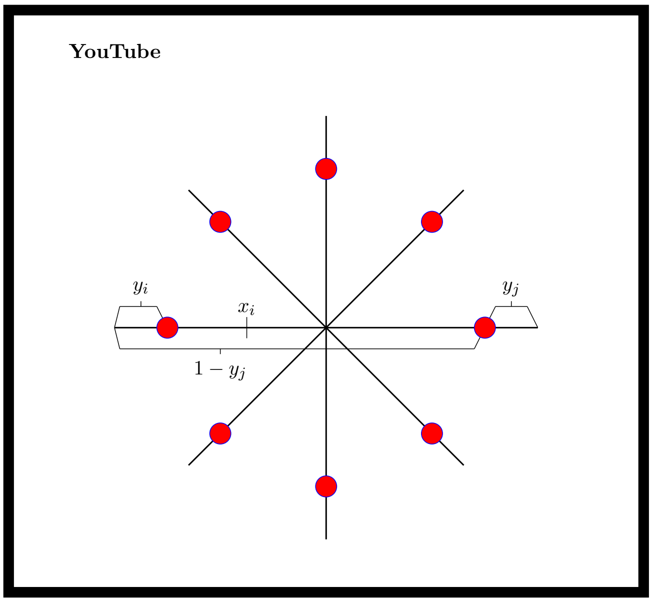

Figure A.2: Spokes model with $N = 8$ spokes (types of video content), $n = N$ YouTubers allocated away from the origin on their spoke ($0 < y_i < 1/2$), a viewer located at x_i, and the comprising platform YouTube.

somewhere between the YouTuber at y_i and an arbitrary YouTuber at $y_{j \neq i}$ (Figure A.2 illustrates). The viewer's utility from watching video content i is given by

$$u_i = v - \gamma a_i - t(x_i - y_i)^2 \tag{10}$$

and her utility from watching video content j is given by

$$u_j = v - \gamma a_i - t(1 - y_j - x_i)^2. \tag{11}$$

Setting $u_i = u_j$ and solving for x_i yields the indifferent viewer

$$\widetilde{x} = \frac{y_i + 1 - y_j}{2} + \frac{\gamma a_j - \gamma a_i}{2t(1 - y_i - y_j)}. \tag{12}$$

Noting that the probability that YouTuber i competes with each of the remaining YouTubers $j \neq i$ is equal to $\frac{1}{N-1}$ and that the density of viewers on the two relevant spokes is equal to $\frac{2}{N}$, her demand is given by

$$D_i(a_i, a_j, y_i, y_j) = \frac{2}{N}\frac{1}{N-1}\sum_{j \neq i}\left(\frac{y_i + 1 - y_j}{2} + \frac{\gamma a_j - \gamma a_i}{2t(1 - y_i - y_j)}\right), \tag{13}$$

and her profit function is

$$\pi_i^Y = \frac{2}{N}\frac{1}{N-1}\sum_{j \neq i}\left(\frac{y_i + 1 - y_j}{2} + \frac{\gamma a_j - \gamma a_i}{2t(1 - y_i - y_j)}\right) a_i r(1 - s) - F^Y, \tag{14}$$

which she maximizes by choosing y_i and a_i. Note that equation (14) is strictly concave in p_i for given p_j, y_i, and y_j.

Choice of advertising quantity $\mathbf{a_i}$ On the final stage of the game, YouTuber i solves for her profit maximizing advertising quantity a_i, taking locations y_i and $y_{j\neq i}$ as given. Considering only the relevant terms, the FOC for an interior maximum w.r.t. a_i is given by

$$\frac{\partial \pi_i^Y}{\partial a_i} = \sum_{j\neq i}\left(\frac{y_i+1-y_j}{2} + \frac{\gamma a_j - \gamma a_i}{2t(1-y_i-y_j)}\right) - \sum_{j\neq i}\frac{\gamma a_i}{2t(1-y_i-y_j)} \overset{!}{=} 0. \tag{15}$$

In any symmetric equilibrium all competitors $j \neq i$ are identical, so rearranging yields the reaction function

$$a_i(a_j, y_i, y_j) = \frac{t}{2\gamma}(y_i+1-y_j)(1-y_i-y_j) + \frac{a_j}{2} \tag{16}$$

and further solving gives

$$a_i^*(y_i, y_j) = \frac{t}{\gamma}(1-y_i-y_j)(1+\frac{y_i-y_j}{3}). \tag{17}$$

Note that $a_i^*(y_i, y_j)$ is maximal when $y_i = y_j = 0$ (YouTubers locate at the origin of their spoke) and minimal when $y_i = y_j = \frac{1}{2}$ (YouTubers locate at the center of the spokes network).

Choice of location $\mathbf{y_i}$ Anticipating $a_i^*(y_i, y_j)$, YouTubers choose their location $y_i \in [0, \frac{1}{2}]$ by maximizing

$$\begin{aligned}&\pi_i^Y\left(a_i^*(y_i,y_j), a_j^*(y_j,y_i), y_i, y_j\right)\\ &= \frac{2}{N}\frac{1}{N-1}\sum_{j\neq i}\left(\frac{y_i+1-y_j}{2} + \frac{\gamma a_j^*(y_j,y_i) - \gamma a_i^*(y_i,y_j)}{2t(1-y_i-y_j)}\right) a_i^*(y_i,y_j) r(1-s) - F^Y\end{aligned} \tag{18}$$

over y_i. Noting that in any symmetric equilibrium all competitors $j \neq i$ are identical and that

$$a_j^*(y_j, y_i) - a_i^*(y_i, y_j) = \frac{2t}{\gamma}(1-y_i-y_j)(\frac{y_j-y_i}{3}), \tag{19}$$

YouTuber i's maximization problem simplifies to

$$\pi_i^Y(y_i, y_j) = \frac{tr(1-s)}{9\gamma N}(1-y_i-y_j)(3-y_j+y_i)^2 \tag{20}$$

where the FOC w.r.t. y_i yields

$$\frac{\partial \pi_i^Y}{\partial y_i} = 2(1-y_i-y_j)(3-y_j+y_i) - (3-y_j+y_i)^2 \overset{!}{=} 0 \tag{21}$$

and rearranging gives $y_i = -\frac{1}{3}(1+y_j)$. Given that $y_i, y_j \in [0, \frac{1}{2}]$, the LHS of Equation (21) is negative, so YouTuber i chooses $y_i^* = 0$. Thus, when YouTubers can freely determine their location on spoke $i = 1, ..., N$, they choose to distance themselves as far as possible from their competitors and allocate at the origin of their spoke.

Consequently, the profit maximizing advertising quantity is given by

$$a^* = \frac{t}{\gamma}. \tag{22}$$

Thus, a^* increases in t and decreases in γ. In other words, a small substitutability between the different types of video content (high t) allows YouTubers to choose a relatively large advertising quantity, whereas a high degree of ad aversion on behalf of the viewers (high γ) has the opposite effect.

Given a^* and y^*, a YouTuber's profit in the symmetric equilibrium is equal to

$$\pi^Y = \underbrace{\frac{1}{N}\frac{t}{\gamma}}_{=R} r * (1-s) - F^Y \tag{23}$$

which – for a given s – increases in $a^* = t/\gamma$ and r, and decreases in N.

Choice of revenue share s Anticipating a^* and y^*, the Platform chooses s subject to $\pi^Y = 0$ (all YouTubers enter the market), so

$$s^* = 1 - \frac{\gamma N F^Y}{tr} \tag{24}$$

and

$$\pi^P = \left(1 - \frac{\gamma N F^Y}{tr}\right) \sum_i \frac{tr}{\gamma N} - F^P = \frac{t}{\gamma} r - N F^Y - F^P. \tag{25}$$

In other words, the Platform reaps the YouTubers' entire surplus, whereby the Platform's and the YouTubers' incentives w.r.t. generating advertising revenue R are perfectly aligned. It may be surprising that π^P decreases in N, hence, the Platform prefers the number of YouTubers to be low. The result is driven the assumptions that every spoke is occupied by a YouTuber ($n = N$), that the market is fully covered (large v), and that the total mass of viewers is always equal to one, irrespective of the number of spokes. Taken together, these assumptions imply that π^P is maximal when the market is shared among a minimum number of YouTubers and minimal duplication of fixed cost. In a scenario where additional YouTubers could "active" viewers who would otherwise abstain from watching, additional surplus would be generated if the number of YouTubers increases, and N would also have a countervailing positive effect on π^P.[56]

[56]See Chen and Riordan (2007) for a sophisticated analysis of entry.

A.3. Limited advertising quantity

In the next step of the analysis, assume that there exists a binding upper bound on advertising quantity $0 < \bar{a} < \frac{t}{\gamma}$ (e.g., due to technical limitations of the Platform).

Choice of advertising quantity $\mathbf{a_i}$ On the final stage of the game, YouTuber i solves for her profit maximizing advertising quantity $a_i \leq \bar{a}$, taking y_i and $y_{j \neq i}$ as given. The reaction function of YouTuber i is thus given by

$$a_i = \min\{\frac{t}{2\gamma}(y_i + 1 - y_j)(1 - y_i - y_j) + \frac{a_j}{2}, \bar{a}\}, \tag{26}$$

and the reaction function of an arbitrary competitor j is given by

$$a_j = \min\{\frac{t}{2\gamma}(y_j + 1 - y_i)(1 - y_i - y_j) + \frac{a_i}{2}, \bar{a}\}. \tag{27}$$

Note that equations (26) and (27) define the unique equilibrium of the advertising subgame after any location choices of YouTubers i and j.

Choice of location $\mathbf{y_i}$ First, given that $\bar{a} > 0$, note that there is no equilibrium where YouTubers i and j locate at $y_i = y_j = 1/2$ in the center of the spokes network, since deviating towards the origin of her spoke would increase a YouTuber's profit.

Consider the case where $y_i, y_j < \frac{1}{2}$ and suppose that locations are such that in the ensuing advertising game, equilibrium quantities are $a_i^*(y_i, y_j)$ and $a_j^*(y_j, y_i)$ with

$$\frac{t}{2\gamma}(y_i + 1 - y_j)(1 - y_i - y_j) + \frac{a_j^*}{2} > \bar{a}. \tag{28}$$

Then, YouTuber i chooses $a_i^* = \bar{a}$. Moreover, even after a small deviation of YouTuber j to $y_j + \epsilon$, YouTuber i would still choose $a_i^* = \bar{a}$ if $\epsilon > 0$ is sufficiently small.

If, in addition,

$$\frac{t}{2\gamma}(y_j + 1 - y_i)(1 - y_i - y_j) + \frac{\bar{a}}{2} > \bar{a} \tag{29}$$

for the reaction function of YouTuber j, then $a_j^*(y_j, y_i) = a_j^*(y_j + \epsilon, y_i) = \bar{a}$ for sufficiently small $\epsilon > 0$. In other words, YouTuber j would benefit from deviating to $y_j + \epsilon$ as she can increase her market share, so this cannot be an equilibrium.

If, on the other hand,

$$\frac{t}{2\gamma}(y_j + 1 - y_i)(1 - y_i - y_j) + \frac{\bar{a}}{2} \leq \bar{a}, \tag{30}$$

when considering the effect of deviating to $y_j + \epsilon$ on the profit of YouTuber j, one only has to vary j's location: the advertising quantity of YouTuber i is given by $a_i^* = \bar{a}$ irrespective of the deviation, and the effect of a_i^* on YouTuber j's profit cancels because of the envelope theorem. Again, YouTuber j would benefit from deviating to $y_j + \epsilon$ as she can increase her market share, so this cannot be an equilibrium.

Therefore, there is no subgame perfect equilibrium such that, on the equilibrium path,

$$\frac{t}{2\gamma}(y_i + 1 - y_j)(1 - y_i - y_j) + \frac{a_j^*}{2} > \bar{a}. \tag{31}$$

Analogously, there is no subgame perfect equilibrium such that, on the equilibrium path,

$$\frac{t}{2\gamma}(y_j + 1 - y_i)(1 - y_i - y_j) + \frac{a_i^*}{2} > \bar{a}. \tag{32}$$

Also, since $\overline{a} < \frac{t}{\gamma}$, there is no symmetric equilibrium such that the price cap is not binding for both YouTubers.

Hence, in any symmetric equilibrium

$$\frac{t}{2\gamma}(y_i + 1 - y_j)(1 - y_i - y_j) + \frac{\bar{a}}{2} = \bar{a}, \tag{33}$$

and

$$\frac{t}{2\gamma}(y_j + 1 - y_i)(1 - y_i - y_j) + \frac{\bar{a}}{2} = \bar{a}. \tag{34}$$

or

$$\bar{y}_i = \bar{y}_j = \frac{1}{2} - \frac{\bar{a}\gamma}{2t} > 0 = y_i^* = y_j^*. \tag{35}$$

Note that no YouTuber has an incentive to deviate from her equilibrium location. If one YouTuber moves closer to the origin of her spoke, she faces a loss in demand while the advertising cap is strictly binding, so deviating is not profitable. If one YouTuber moves closer to the center of the network, the advertising cap is not binding. Then, moving further away again is profitable, so the YouTuber cannot benefit from moving closer to the center.

A YouTuber's equilibrium profit in case of limited advertising is given by

$$\overline{\pi}^Y = \frac{r\overline{a}}{N}(1 - s) - F^Y < \pi^Y. \tag{36}$$

Choice of s Analogous to section A.2, the Platform anticipates the choice of $\overline{a}$ and $\overline{y}$, and chooses s subject to $\overline{\pi}^Y = 0$ (all YouTubers enter the market), so

$$\overline{s} = 1 - \frac{NF^Y}{r\overline{a}} \tag{37}$$

and

$$\overline{\pi}^P = \left(1 - \frac{NF^Y}{r\overline{a}}\right) \sum_i \frac{r\overline{a}}{N} - F^P = r\overline{a} - NF^Y - F^P < \pi^P. \tag{38}$$

Thus, both YouTubers and Platform earn higher revenues without a cap on advertising quantity.

A.4. Discussion

Though highly stylized, the theoretical framework illustrates the main economic incentives of YouTubers, viewers, and YouTube itself, and thereby supports the economic reasoning from the main part of the paper.

A.4.1. Economic incentives of the YouTubers

First, the model sheds light on the interplay between advertising quantity and content differentiation. Section A.2 demonstrates that YouTubers maximally differentiate from each other if they can choose both advertising quantity and the location on their spoke. If advertising quantity is limited, however, they locate closer to the center of the radial network and content differentiation shrinks. In the extreme, when $\bar{a}$ converges towards 0, all YouTubers would locate near the center of the network and duplicate the same type of video content. The economic intuition behind this result is that YouTubers differentiate from their competitors to avoid ruinous competition in advertising quantities. When competition in advertising quantities is prohibited by upper bounds, YouTubers lack the incentive to differentiate. Key to the result is the assumption that advertising diminishes the utility of viewers. If viewers were not ad averse ($\gamma = 0$), there would be no competition in advertising quantities, and YouTubers would allocate at the center of the radial network even in the absence of a binding upper bound $\bar{a}$.

The results from the theoretical model are in line with the empirical observations in the main part of the paper. In particular, I find that YouTubers' probability to duplicate mainstream content decreased and their incentive to differentiate increased after YouTube made its "ten minutes trick" more prominent, which is comparable to abolishing a cap on advertising quantity. Moreover, the empirical evidence is in line with a decline in YouTubers' competition over especially popular types of video content.

A.4.2. Economic incentives of the Platform

Second, the theoretical framework illustrates the economic incentives of the Platform YouTube. The model shows that both the Platform and the YouTubers wish to maximize advertising revenue R. In particular, Section A.3 demonstrates that the Platform's profit is higher in the absence of a binding cap $\bar{a}$ on advertising quantity, in which case content differentiation is maximal, too.

These results are in line with what is known about YouTube's business strategy. E.g., YouTube offers a broad range of tutorials on how to grow a (commercially) successful channel[57], and it stresses the importance of customizing a brand, being unique, and targeting a well-defined audience.[58] It is thus plausible to assume that YouTube launched its

[57] See, e.g., `https://www.youtube.com/intl/en_ALL/creators/how-things-work/` (Nov 2021).

[58] See, e.g., `https://creatoracademy.youtube.com/page/lesson/brand-identity?cid=bootcamp-foundations\&hl=en#strategies-zippy-link-1`, `https://creatoracademy.youtube.com/page/lesson/niche?cid=great-content\&hl=en`, and `https://www.youtube.com/watch?v=4vjzAi_dUzU` (Nov 2021).

new ad break tool in Nov 2015 in order to nudge YouTubers to increase the advertising quantity in their videos.

A.4.3. Economic incentives of the viewers

Third, the model delineates the behavior of viewers. Although it may appear as a strong assumption that viewers like exactly two types of video content, the idea that they care for a small and finite number of topics and ignore all others is sensible (e.g., a viewer may only be interested in videos on running and yoga).

A.4.4. Limitations

Naturally, the toy model in sections A.1 to A.3 quickly reaches is limits. E.g., it can neither illustrate why non-advertising YouTubers enter the Platform, nor does it capture search engine optimization on behalf of the YouTubers, algorithmic confounding on behalf of the Platform, targeted advertising, or viewer switching behavior. These topics are further discussed in Appendices B.1, E.2, G.2, G.3, and G.4.

B. Measuring mainstream content

This section discusses and extends my measure for mainstream content along several dimensions. First, I consider potential search engine optimization on behalf of the YouTubers and explain why video tags are less likely to be manipulated than video titles. Then, I show that my main results are robust to using various alternative measures for mainstream content, including a measure based on video titles, tags from previous months, YouTube superstars, and an unsupervised topic modeling approach.

B.1. Search engine optimization

Appendix A points out that advertising YouTubers' revenue depends on the number of ad breaks, the price per ad per viewer, and the number of views. While YouTubers can choose the number of ads, they must take the price per ad per view as given (see Appendix G.3), and the number of views is primarily affected by YouTube's recommendation algorithm (see Appendix E.2). Since views are a crucial determinant of revenue, however, some YouTubers pursue *search engine optimization* (SEO) through strategic keyword choice.[59] SEO could confound my measure of mainstream content if it induced YouTubers to select video tags that are not descriptive of the actual video content or, more specifically, if it induced YouTubers to strategically choose tags that are less likely to be classified as mainstream after Nov 2015. Three arguments, however, speak against such concerns.

First, any attempts for SEO – if they exist – are likely to target keyword choice in a video's title, thumbnail, and description, while tags are less prone to be manipulated. As argued in Section 5.2, video tags are descriptive terms and phrases that YouTubers specify through a certain template when they upload a video; keywords, in contrast, correspond more generally to the central terms and topics of a video. Thus, keywords and tags may coincide, but need not necessarily do so. YouTubers, bloggers, and YouTube itself agree that the main purpose of video tags is to help YouTube understand what a video is about, while a video's title, thumbnail, and description are more important for video discovery and SEO.[60] E.g., YouTube states in an explanatory video on "YouTube search results" that videos are ranked based on "how well the title, descriptions and video itself match each query"[61], and in a related video on the "YouTube algorithm" that it considers "things like titles, thumbnails, descriptions".[62]

Second, YouTubers, bloggers, and YouTube itself *strongly* advise against keyword and tag manipulation. E.g., YouTube states that "adding excessive tags" is against their policy on deceptive practices[63]; tricking viewers into "believing the content is something it is not" leads to video removal and, eventually, to a ban from the platform.[64] Similarly, YouTube

[59]See, e.g., `https://creatoracademy.youtube.com/page/lesson/discovery` (Dez 2021).

[60]See, e.g., `https://support.google.com/youtube/answer/146402?hl=en` or `https://blog.hubspot.com/marketing/youtube-tags` or `https://www.youtube.com/watch?v=QA-CcHhsxtI` (Dez 2021).

[61]See `https://www.youtube.com/watch?v=gTrLniP5tSQ` (Dez 2021).

[62]See `https://www.youtube.com/watch?v=hPxnIix5ExI` (Dez 2021).

[63]See `https://support.google.com/youtube/answer/146402?hl=en` (Dez 2021).

[64]See `https://support.google.com/youtube/answer/2801973?hl=en\&visit_id=`

points out that "spammy" keywords may help YouTubers to increase the number of views in the short-run, but that the better strategy is to provide viewers with what they are looking for.[65] YouTubers confirm that videos are ranked worse if YouTube's algorithm finds that keywords and tags are inconsistent with what is being said in the video[66], and bloggers recommend to use "as many tags as possible to accurately describe your content without being spammy".[67] Thus, given YouTube's effective means of prosecution, deliberate manipulation of keywords and tags for the sake of SEO is likely to be a minor concern.

Third, some websites advise to use so-called "long-tail tags" for SEO. E.g., the website `keywordtool.io` distinguishes between short-tail tags that contain only one or two words, and long-tail tags that have more than two words in it.[68] According to the site, viewers who use long-tail terms in their searches know exactly what they are looking for and are thus more willing to watch a specific video if it matches their search intent. Hence, having long-tail tags such as "black thin silicon phone case iphone xs" in addition to "phone case" may increase video views. This SEO strategy could confound my results if long-tail tags were strategically chosen after Nov 2015 and *replaced* more generic short-tail tags. In this case, a video would quite mechanically be less likely to be classified as "mainstream". It is, however, often recommended to have long-tail tags *in addition to* more generic ones.[69] Since I classify a video as mainstream if it is given at least one mainstream tag, *adding* long-tail tags does not affect my empirical strategy. As a further robustness check, I split all tags into single words[70] before I compute my measure for mainstream content and re-estimate equations (2) and (3) using this alternative dependent variable.[71] Table A.1 shows the results. Though smaller, the estimates are qualitatively similar to their counterparts in Table 1. Thus, my main results are not driven by compound long-tail tags replacing more generic short-tail ones, and I conclude that SEO on behalf of the YouTubers is unlikely to confound the results in my paper.

B.2. Titlewords

Next to video tags, there are several other pieces of metadata that YouTubers can supply, including video titles. Unlike video tags, providing a video title is mandatory for each YouTuber. This section shows that my results are robust to using video *titles* instead of video *tags* to construct a measure for mainstream content.

`637725899088506393-102987062\&rd=1` (Dez 2021).

[65] See, e.g., `https://blog.youtube/creator-and-artist-stories/top-tips-for-partners-words-words-words/` (Dez 2021).

[66] See, e.g., `https://www.youtube.com/watch?v=z5Qr7uv54yo` or `https://www.youtube.com/watch?v=QA-CcHhsxtI` (Dez 2021).

[67] `https://www.pmg.com/blog/youtube-video-optimization/` (Dez 2021).

[68] See `https://keywordtool.io/long-tail-keywords` (Dez 2021); note that the website speaks of "keywords" but is actually talking about "tags".

[69] E.g., `https://www.youtube.com/watch?v=QA-CcHhsxtI` (Jan 2022).

[70] Since the splitting exercise leads to the frequent occurrence of single-term tags such as "the" or "very", I remove the most common stopwords as well as punctuation from each video's list of tags.

[71] The correlation between this alternative and my main dependent variable is equal to 0.62.

Table A.1: Split compound tags

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	.013*	.009	.009	-.096**	-.113***	-.114***
	(.007)	(.007)	(.007)	(.043)	(.041)	(.041)
First stage				.029***	.029***	.029***
				(.002)	(.002)	(.002)
F-statistic				144.13	143.29	150.65
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v of YouTuber i in month t covers mainstream content, where the measure for mainstream content is based on split tags. Columns 1 to 3 display OLS, and columns 4 to 6 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

I start by splitting each video title into single words. In contrast to tags, video titles often contain stopwords and punctuation (especially question and exclamation marks), which I remove in a pre-processing step; otherwise, the most mainstream keywords would be "the" and "a".[72] Compared to tags, video titles seem to convey less information on video content; in particular, the number of keywords generated from a video title ("titlewords") is on average about 50% lower than the number of tags (5.96 titlewords vs. 11.7 tags).[73] I also find that titlewords and tags overlap to a certain extent: on average, about a third of the titlewords also occurs in the tags.

Based on the generated titlewords, I construct my measure for mainstream content analogous to Section 5.2 in the main part of the paper. To compensate for the fact that a video has only half as many titlewords as tags and is thus mechanically less likely to be classified as mainstream, I specify the upper two percent of the distribution of most-viewed titlewords as *mainstream*, and classify all videos with at least one mainstream titleword accordingly.[74] Then, I estimate regressions (2) and (3) using this new alternative variable.

Table A.2 shows the results. In contrast to the main results in Table 1, the OLS estimates in columns 1 to 3 are positive and statistically significant, which is consistent with the idea that especially money-loving YouTubers self-select into the treatment and manipulate their video titles to attract a larger audience. The 2SLS estimates in columns 4 to 6, in contrast, are negative and highly statistically significant at the 1%-level, but about 50% smaller than their counterparts in Table 1. As argued in Section B.1, tags

[72] Although I only consider German YouTubers, I remove German *and* English stopwords, since many German teenagers think it is "cool" to blend these two languages and speak some kind of "Germish".

[73] The median number of titlewords is equal to six, whereas the median number of tags is equal to twelve.

[74] The correlation between this alternative and my main dependent variable is equal to 0.27.

Table A.2: Titlewords

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	.025***	.022***	.022***	-.103***	-.120***	-.113***
	(.006)	(.006)	(.006)	(.037)	(.037)	(.036)
First stage				.029***	.029***	.029***
				(.002)	(.002)	(.002)
F-statistic				144.13	143.29	150.65
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. The dependent variable is $MainstreamTitle_{vit}$ which is a dummy variable equal to 1 if video v of YouTuber i uploaded in month t is equipped with a mainstream titleword, and 0 otherwise. Columns 1 to 3 display OLS, and columns 4 to 6 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

may describe video content more accurately than video titles, which have the additional purpose to catch viewers' attention. Thus, titles are more prone to manipulation than video tags, which could explain why the OLS estimates in Table A.2 are positive and larger than in Table 1, and why the magnitude of the 2SLS estimates is smaller. In sum, however, I conclude that my results are robust to using titlewords to construct a measure for mainstream content.

B.3. Time lag in video production

YouTubers may need some time to incorporate current trends into their videos; therefore, I also construct a measure for mainstream content that is based on the mainstream tags from the *previous* month.[75] Table A.3 shows the results from estimating equations (2) and (3) using this alternative dependent variable.

The OLS and 2SLS estimates are nearly identical to their counterparts in Table 1, which is intuitive for two reasons. First, Tables A.21 and A.22 demonstrate that certain tags are always classified as mainstream, whereby it is irrelevant whether I consider t or $t-1$. Second, the main idea of the paper is that YouTubers who could increase their ad quantity deliberately *avoid* mainstream content, which plausibly includes hot topics from current *and* previous months. Thus, my main results are not affected when I consider tags from $t-1$.

[75]The correlation between this alternative and my main dependent variable is equal to 0.72.

Table A.3: Time lag in video production

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	.011	.006	.006	-.164***	-.193***	-.188***
	(.008)	(.008)	(.008)	(.050)	(.048)	(.048)
First stage				.029***	.029***	.029***
				(.002)	(.002)	(.002)
F-statistic				144.40	143.54	150.77
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,060,018	1,060,018	1,060,018	1,060,018	1,060,018	1,060,018

Notes: Robust standard errors in parentheses. The dependent variable is a dummy variable equal to 1 if video v of YouTuber i uploaded in month t is equipped with a tag that was mainstream in $t-1$, and 0 otherwise. Columns 1 to 3 display OLS, and columns 4 to 6 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B.4. Superstars

An important aspect of YouTube is that it is a superstar economy (Rosen, 1981), i.e., a small fraction of YouTubers reaps a large fraction of attention and views. What is truly mainstream at each given moment is thus largely determined by a few trendsetters with millions of subscribers and views. My main measure for mainstream content as described in Section 5.2 is in line with this idea: mainstream tags are those that accumulate the largest number of views in a given month and category. Here, I explicitly consider YouTube superstars and build a measure of mainstream content based on what they do.

To this end, I select the top 25 YouTubers by number of subscribers for each video category and define them as “superstars”. Then, I collect all tags that these YouTubers use in a given month and define these tags as “mainstream”. Finally, I assign a dummy variable equal to one to all videos that are equipped with such a mainstream tag and use this measure as an alternative dependent variable. As a robustness check, I repeat the procedure with the top 50 YouTubers per category.

Table A.4 shows the results from estimating equations (2) and (3) by 2SLS with this alternative measure of mainstream content. The estimates are nearly identical to my main results in Table 1, which is intuitive, given that my main measure and the superstar measure follow the same logic.[76] In addition, the findings are perfectly in line with the heterogeneous effects that I document in Section 7.3.2 and avoiding competition as a plausible economic mechanism as discussed in Section 8.

[76]The correlation between my main measure for mainstream content and the measure based on the top 25 YouTubers per category is equal to 0.56. The correlation between my main measure for mainstream content and the measure based on the top 50 YouTubers per category is equal to 0.58.

Table A.4: Superstars

	Top25			Top50		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	-.182***	-.169***	-.185***	-.205***	-.181***	-.199***
	(.046)	(.045)	(.044)	(.045)	(.043)	(.043)
First stage	.029***	.029***	.029***	.029***	.029***	.029***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
F-statistic	144.13	143.85	150.65	144.13	143.85	150.65
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. The dependent variable in columns 1 to 3 (columns 4 to 6) is a dummy that is equal to one if video v by YouTuber i is equipped with a tag that at least one of the top 25 (top 50) YouTubers in his or her category uses in month t, too. The estimates are based on using the advertising YouTubers only. All estimates are 2SLS estimates. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B.5. Topic model approach: LDA

In the main part of the paper, I interpret each individual video tag as representing a different type of video content. The advantage of this approach is that I do not have to make any assumptions on how tags could be clustered into more general topics. E.g., the tags *dolphin* and *fish food* would be considered as two different topics / types of video content. If in a given month and category, *dolphin* attracts more views than *fish food*, and *dolphin* is identified as a mainstream tag, all videos with the tag *dolphin* (or *dolphin* and *fish food*) would be classified as mainstream, whereas videos with the tag *fish food* (without *dolphin*) would not.

However, it could be that the more general (but unobserved) topic *sea life* is actually the true mainstream topic in that month and category. If *sea life* comprises themes like dolphins *and* fish food (which is plausible), one could argue that videos with the tag *fish food* should be classified as mainstream, too. If, in contrast, *sea life* is rather unpopular in comparison to other general topics and hence not identified as mainstream, neither videos with the tag *dolphin* nor the tag *fish food* should be classified as mainstream.

To take potential clusters of tags into account, this section employs a topic model approach, based on the idea that each video covers one or several general topics (e.g., *sea life*), where topics are defined as a probability distribution over tags (e.g., the topic *sea life* would place high probability on *dolphin* and *fish food*, and low probability on *horses*). More specifically, I apply a *latent Dirichlet allocation* (LDA, Blei et al., 2003), one of the most widespread topic models, to demonstrate the robustness of my main results.

B.5.1. Stylized example

Consider the stylized example in Figure A.3, which displays two hypothetical videos from the category "Pets & Animals". Each video is equipped with 6 tags describing its content. Videos and tags are observed by the researcher. Most of the tags in video 1 are related to *sea life*, but some also describe *plants* and *weather*. Thus, video 1 covers three general topics, but blends them in different proportions; the same is true for video 2. Note that both videos share the same set of potential topics, and each topic comprises a fixed vocabulary of tags. Crucially, the topics and their vocabulary are *unobserved* by the researcher.

The main idea of the LDA is to use the observed videos and tags to infer the unobserved topic structure. Put differently, the LDA asks which unobserved topic structure is most likely to generate the observed videos and tags. To this end, the model "reverses" the data generating process and computes the most likely probability distribution over tags for each topic, and the most likely probability distribution over topics for each video conditional on its tags. Here, an LDA would be likely to find that the topic *sea life* puts a high probability on the tags *dolphin* and *fish food*, and that the topic *horses* puts a high probability on the tags *horse* and *tournament*. Moreover, the LDA would place high probability on *sea life* and low probability on *horses* for video 1, while the reverse is true for video 2.

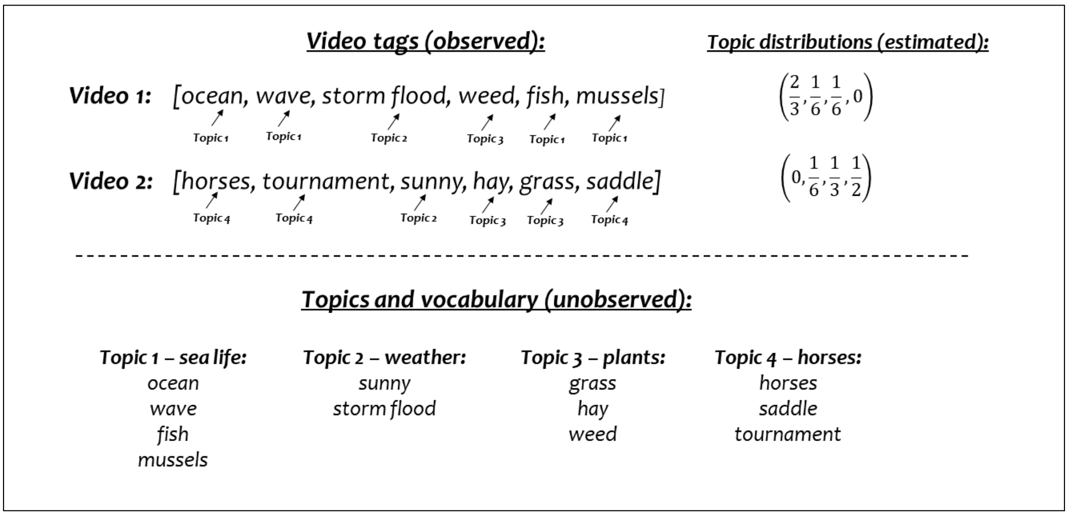

Figure A.3: Stylized example for an LDA topic model.

B.5.2. Implementation

Building on this intuition, I generate an alternative measure for mainstream content in three steps. I start by using separate LDAs to identify the unobserved topic structure for each month and each video category. Crucially, the absolute number of topics in an LDA must be determined by the researcher. A reasonable, feasible, and non-arbitrary approach is to divide the number of unique tags per month and category by 100; e.g., $5,000$ unique tags imply 50 different topics.[77] Given the number of unobserved topics, the algorithm uses the observed videos and tags to compute (i) the most likely probability distribution over tags for each topic and (ii) the most likely probability distribution over topics for each video.

Next, I use the inferred probability distribution over topics to determine which topics are covered by a specific video. A typical video in my sample features about four or five topics with a high, and the remaining topics with a low probability. Based on that, I say that a video covers topic z if the probability $\Pr(z)$ computed by the LDA exceeds a critical value p, which I set equal to the inverse number of topics (e.g., 1/50).[78] Note that a video could cover none, one, or several topics.

Finally, I specify which topics are mainstream and which videos cover mainstream content. Analogous to the procedure in Section 5.2, I first compute how many views a certain tag has attracted in each month and category. Then, I consider each topic along with its inferred probability distribution over tags. In particular, I weight each tag's number of views with its probability to belong to a specific topic and compute the sum of weighted views for each topic. Then, I rank the topics in descending order based on their sum of weighted views and classify the upper 10% as mainstream.[79] Finally, I assign a dummy variable equal to one to all videos that cover a topic that is classified as mainstream.[80]

Example Consider the category "Science & Technology" in April 2015 as a concrete example again. There are 13,555 unique tags, so the LDA algorithm identifies 135 different topics, 13 of which are classified as mainstream (Table A.5). E.g., the first topic seems to be about smartphones in general, the second topic about the Samsung Galaxy S6, and the third topic about the Apple iWatch (note that the topic names are assigned by the researcher). In this month and category, the average number of topics per video is equal to 4.24, and the median is equal to 4.

B.5.3. Results

Table A.6 shows the results from estimating equations (2) and (3) with the alternative dependent variable. The estimates are slightly smaller than their counterparts in Table 1 and confirm my results are robust to using a topic model approach when determining

[77] In applied research, the number of topics in an LDA is often manually determined and fine-tuned after close inspection of the data (e.g., Hansen et al., 2018, p.820). However, this procedure is not feasible in my application, where I implement independent LDAs for 14 different video categories in 49 months. After inspecting the output of the LDA from different categories at different points in time, I concluded that the approach to divide the number of unique tags per month and category by 100 yields the most reasonable results across settings.

[78] As a robustness check, I also employ more restrictive critical values, but the results remain nearly unchanged.

[79] For some of the smaller video categories, the typical number of topics per month often lies between ten and thirty. Classifying the upper 10% of topics as mainstream ensures that at least one topic is specified as mainstream. As a robustness check, I also classify the upper 5% and the upper 2% of topics as mainstream and obtain similar results.

[80] The correlation between this alternative and my main dependent variable is equal to 0.33.

Table A.5: Mainstream topics for "Science & Technology" in April 2015

Topic name	Top ten tags with highest probability
smartphones	'iphone', 'android', 'video', 'app', 'smartphone', 'instructions', 'game', 'tutorial', 'platform', 'on'
galaxy s6	'galaxy', 'samsung', 's6', 'edge', 'test', 'brand', 'led', 'note', 'comparison', 'tips'
apple watch	'apple', 'watch', 'must', 'iwatch', 'wrist band', 'iphone', 'activity', 'with', 'how', 'change',
camera review	'test', 'german', 'review', 'hd', 'german', 'unboxing', 'full', 'camera', '1080p', 'unboxing'
fix computer	'computer', 'battery', 'change', 'category', 'repair', 'product', 'the', 'acer', 'see', 'city'
computer review	'pc', 'unboxing', 'german', 'gaming', 'pro', 'case', 'best', 'air', 'case', 'review',
music fair	'2015', 'fair', 'music fair', 'prolight', 'delamar', 'sou', 'hanover', 'abb', 'am', 'robot'
printers	'sensor', '3d', 'organization', 'cnc', 'online', 'diy', 'printer', 'arduino', 'motion','mechaplus'
media	'media', 'software', 'industry', 'genre', 'tutorial', 'germany', 'language', 'country','new','video'
music equipment	'music', 'music', 'sound', 'music fair', 'audio', 'headphones', 'loudspeaker', 'motor','ip','engineering'
tourism	'german', 'subject', 'city', 'literature', 'museum', 'tourist', 'destination', 'hanover', 'saxon', 'tourism'
drones	'drone', 'drohne', 'quadcopter', 'en', 'helicopter', 'bd', 'flight', 'mode', 'professional','copter'
streaming	'tv', 'airplane', 'plane', 'amazon', 'hobby', 'digital', 'smart', 'fire', 'stick', 'solar'

Notes: Table A.5 shows the ten tags with the highest probability for each of the thirteen topics with the largest number of views (i.e., the mainstream topics) for the category "Science & Technology" in April 2015. All tags are translated into English.

mainstream content.

Table A.6: Latent Dirichlet Allocation (LDA)

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	-.000	-.002	-.002	-.157***	-.142***	-.128***
	(.005)	(.004)	(.004)	(.033)	(.031)	(.031)
First stage				.029***	.029***	.029***
				(.002)	(.002)	(.002)
F-statistic				144.13	143.29	150.65
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v of YouTuber i in month t is equipped with a tag belonging to a mainstream topic with a probability that is larger than the inverse number of topics in that month and category. Columns 1 to 3 display OLS, columns 4 to 6 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B.6. Further specifications of mainstream content

This section shows that my main results are robust to four further robustness checks on mainstream content. First, I assign a dummy equal to one to all videos equipped with one of the upper *half* percent of most-views tags; second, I assign a dummy equal to one to all videos that are given a tag from the upper *two* percent of that distribution. Third, instead of considering a relative share, I classify a *fixed number* of tags per month per category as mainstream – 250 tags for the large categories “Entertainment”, “People & Blogs”, and “Let’s Play”, and 100 tags for all other categories. Fourth, instead of considering views, I use the number of *Likes* a certain tag has attracted.

Table A.7 shows the results from a 2SLS estimation of equations (2) and (3) using the four alternative definitions of $Mainstream_{vit}$. The estimates for β are negative for all specifications; moreover, with the exception of column 1, they are highly statistically at the 1%-level.

Table A.7: Further robustness checks on mainstream content

	2SLS			
	0.5%	2%	Fixed	Likes
	(1)	(2)	(3)	(4)
$D_i * post_t$	-.060	-.368***	-.159***	-.267***
	(.044)	(.054)	(.046)	(.050)
First stage	.029***	.029***	.029***	.029***
	(.002)	(.002)	(.002)	(.002)
F-statistic	150.65	150.65	150.65	150.65
Time FE	X	X	X	X
YouTuber FE	X	X	X	X
Category FE	X	X	X	X
Category Time Trend	X	X	X	X
YouTuber Time Trend	X	X	X	X
YouTubers	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. In column 1, the dependent variable is equal to one if video v of YouTuber i in month t is given a tag from the upper half percent of the distribution of most-viewed tags. In column 2, the dependent variable is equal to one if video v of YouTuber i in month t is given a tag from the upper two percent of the distribution of most-viewed tags. In column 3, the dependent variable is equal to one if video v of YouTuber i in month t is given a tag from a fixed number of the distribution of most-viewed tags. In column 4, the dependent variable is equal to one if video v of YouTuber i in month t is given a tag from the upper one percent of the distribution of most-liked tag. All estimates are 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

C. Validity checks

This section presents additional discussion and validity checks for my empirical strategy. I first present a series of placebo regressions to further support the plausibility of the exclusion restriction. Then, I discuss instrument independence and monotonicity as additional requirements. Third, I show that video duration as such has no impact on video content. Finally, I come back to the non-advertising YouTubers.

C.1. Placebo regressions

To further support the plausibility of the exclusion restriction, I conduct a series of placebo regressions. To this end, I augment the reduced form of equations (2) and (3) to

$$Mainstream_{vit} = \gamma^p close_i * fakepost_t + \phi_i^p + \phi_t^p + \phi_c^p + \tau_1^p t_{it} + \tau_2^p t_{ct} + v_{vit} | \, t \leq 34, \tag{39}$$

where in the first placebo regression, $fakepost_t$ is equal to one if $t \geq 3$, in the second placebo regression $fakepost_t$ is equal to one if $t \geq 4$, and so on; I run 29 placebo regressions in sum. If $close_i$ has no effect on $Mainstream_{vit}$, and YouTubers with different values of $close_i$ were on no different trends before Nov 2015, all estimates for γ^p should be close to zero and not statistically significant.

Of 29 placebo regressions, the estimate for γ^p is in three cases statistically significant at the 5%-level; these estimates are, however, positive. Thus, the results provide additional support for the plausibility of the exclusion restriction.

C.2. Instrument independence

In addition to the exclusion restriction, the instrument must also be as good as randomly assigned such that the first stage equation (3) captures the causal effect of $close_i$ on D_i. Note that reverse causality is of no concern here, because $close_i$ is by definition determined before, and D_i after Nov 2015. Yet, YouTuber specific time-varying factors that drive both $close_i$ and D_i as well as the potential manipulation of $close_i$ on behalf of the YouTubers – in the sense that they choose high values of $close_i$ to increase their treatment probability – may be an issue.

Four facts, however, argue against the manipulation of $close_i$. First, the ten minutes trick was unknown until Nov 2015. Second, YouTube did not announce the new ad break tool before its launch, so the knowledge of the ten minutes trick caught the YouTubers unprepared.[81] Third, YouTubers do not benefit from higher values of $close_i$ before Nov 2015, since the number of ad breaks per video is limited to one, irrespective of how close their are to the threshold. Finally, if a YouTuber chose a high value of $close_i$ to increase her treatment probability, she must know about the ten minutes trick; if she knew about the ten minutes trick, she would either exploit or ignore it, but she would not just move closer to the threshold.

[81] I searched through the YouTube creators blog (`https://youtube-creators.googleblog.com/`) and found no entries announcing the new ad break tool.

It remains to rule out that unobserved YouTuber specific time-varying factors drive both $close_i$ and D_i. Three arguments speak against such concerns. First, t_{it} in equation (3) controls for YouTuber specific linear time trends; in Appendix G.6, I also include higher order polynomials of t_{it} into equation (3). Second, while commercial interests are a plausible driver of D_i, they are unlikely to affect $close_i$, as argued above. Third, YouTubers with a strong commercial interest might self-select into particular video categories that, in turn, require a certain video duration. Yet, any category specific characteristics are captured in the category fixed effects, prohibiting that $close_i$ is indirectly driven by a YouTuber's commercial interest.

C.3. Monotonicity

While $close_i$ may have no effect on some YouTubers, those who are affected must be affected in the same direction, i.e., $\pi_i \geq 0 \; \forall \, i$. This is plausible: it is hard to believe that a high value of $close_i$ prohibits treatment from YouTubers who would have been treated if $close_i$ was low. Figure A.4 provides further illustrative evidence. It plots all values of $close_i$ against the corresponding probability of treatment, $Pr(D_i = 1)$. With the exception of some outliers at the upper left and the lower right corner, the relationship between $close_i$ and $Pr(D_i = 1)$ is monotone.

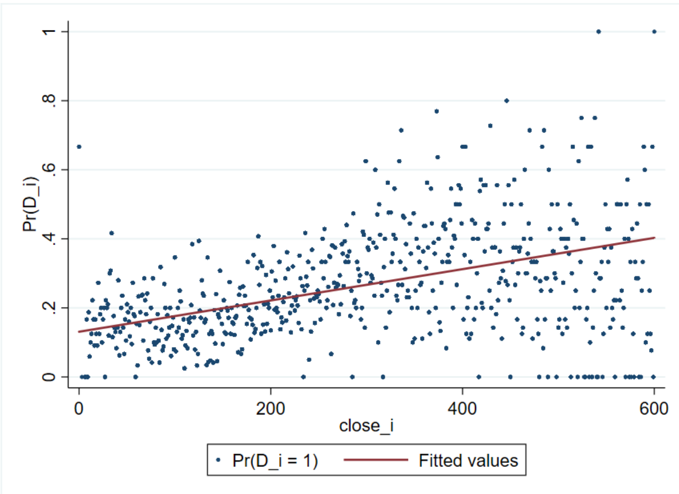

Figure A.4: Illustrative evidence for the monotonicity assumption. The graph plots every value of $close_i$ in the dataset (in seconds) on the x-axis against the corresponding probability $Pr(D_i = 1)$ on the y-axis.

Note that I might violate the monotonicity assumption if I used a continuous measure of treatment intensity – i.e., the *extent* to which a YouTuber increases her share of videos that are ten minutes or longer – instead of the binary treatment status D_i. As argued, YouTubers with high values of $close_i$ have a higher *probability* to increase their share of videos that are ten minutes or longer. At the same time, however, they have *less scope* to do so, because their initial share of videos that are ten minutes or longer is already high. Hence, while the impact of $close_i$ on the *extensive margin* of treatment is monotone and increasing – as shown above – it might follow an inverted U-shape on the *intensive margin*.

C.4. Video duration and mainstream content

To insert additional ad breaks into their videos, YouTubers must make them ten minutes or longer. If mainstream topics generally require a shorter video duration than niche topics, my main results might be spurious (i.e., not driven by advertising). To eliminate such concerns, I demonstrate that video duration as such is not correlated to mainstream content. In particular, I estimate the regression equation

$$Mainstream_{vit} = \delta duration_{vit} + \phi_i''' + \phi_t''' + \phi_c''' + \tau_1'''t_{it} + \tau_2'''t_{it} + e_{vit} \mid t \leq 34 \qquad (40)$$

by OLS, where $duration_{vit}$ corresponds to the video duration of video v by YouTuber i in month t. Crucially, I restrict the analysis to observation periods before Nov 2015 to preclude confounding effects of the ten minutes trick. If video duration as such has no impact on mainstream content, the estimate for δ should be close to zero and statistically insignificant. The results in Table A.8 confirm that this is indeed the case.

Table A.8: Mainstream content and video duration

	OLS		
	(1)	(2)	(3)
$D_i * post_t$	.0000	.0000	.0000
	(.0001)	(.0001)	(.0001)
Time FE	X	X	X
YouTuber FE	X	X	X
Category FE		X	X
Category Time Trend		X	X
YouTuber Time Trend			X
YouTubers	10,113	10,113	10,113
Videos	566,079	566,079	566,079

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v of YouTuber i in month t covers mainstream content. The estimates are based on using the advertising YouTubers only and consider only the time period before Nov 2015. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

C.5. Non-advertising YouTubers

The non-advertising YouTubers, whom I do not consider in the main analysis, allow me to conduct an additional validity check. In particular, the non-advertising YouTubers' decisions w.r.t. video content are unlikely to be driven by commercial considerations, whereby the probability to duplicate mainstream content should be unaffected by the launch of the new ad break tool in Nov 2015.[82]

Figure A.5 presents the results of an event study analogous to Figure A.29 in the main part of the paper. When I consider only the subsample of non-advertising YouTubers, *all* estimates for γ_t fluctuate around zero, and nearly all of them are statistically insignificant, hence supporting the validity of my empirical strategy.

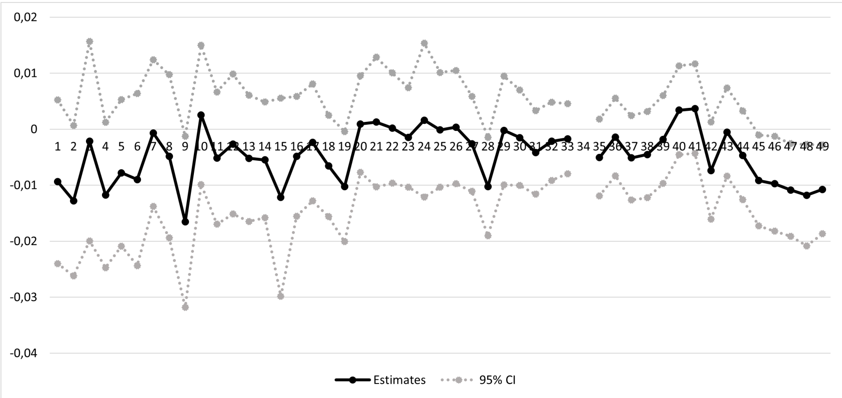

Figure A.5: Event study. The solid line displays the estimates for γ_t, the dashed lines depict a 95% confidence interval. The estimates are based on an OLS regression of equation 4 including YouTuber, time, and category fixed effects, as well as linear YouTuber and category time trends. I consider non-advertising YouTubers only.

[82]See Appendix G.2 for further discussion on the economic incentives of advertising and non-advertising YouTubers.

D. Online survey experiment

To check the plausibility of the automatically generated measures for video content, quality, and language, I conduct an online survey experiment, where I let human coders watch and rate YouTube videos.[83] More specifically, I let the participants assess video quality along five dimensions, asked them which language was used in a video, and I let them check whether a video's set of tags describes its content appropriately. This section illustrates the design and the implementation of the online survey in detail; moreover, I provide descriptive statistics of the participants' responses. Validation of my automatically generated measures is deferred to Section 5.2, where I introduce my measure for mainstream content, Section 9, where I study video quality, and Appendix E.6, where I explore video language.

D.1. Experimental setup

D.1.1. Design

Video selection As it is not possible to let human coders rate every single video in my dataset, I start with a careful pre-selection. Appendix E.1.2 illustrates that there exists a random subsample of 500 advertising YouTubers and $52,462$ videos of which I know the *actual* number of ad breaks per video. Linking the human rating to the actual number of ad breaks per video is likely to be insightful; hence, I extract all videos for the online survey experiment from this subsample.

To balance the selection of videos, I define meaningful strata. To this end, I partition the subsample into videos from before and after Nov 2015, and into videos with zero, one, or multiple ad breaks, which yields a total of six subgroups. Since the number of videos with multiple ad breaks is relatively small, I combine the observations from before and after Nov 2015 into one group, reducing the total number of strata to five. Based on that, I decided that human coding of 500 videos per strata (i.e., $2,500$ videos in total) would be sufficient for my validation exercise.

Within each strata, I sorted YouTubers by their Channel ID, which is a random combination of 24 characters, so sorting by Channel ID corresponds to bringing the YouTubers and their corresponding videos into random order. Many YouTubers deleted some of their videos or vanished entirely since my initial data collection in 2017, so I manually checked the videos' existence and functionality from top to bottom in each of the five strata, and I stopped as soon as I had assembled five times 500 videos that are still accessible.

After some pre-tests with colleagues and research assistants, I found that participants of the survey could watch and rate about five YouTube videos (i.e., one per strata) without becoming inattentive. Thus, when there are $2,500$ videos to be rated, and I want each video to be rated three times on average to mitigate the impact of outliers, the required

[83] The online survey experiment has been officially approved by the Ethics Commission, Department of Economics, University of Munich, under project number 2022-01.

number of participants is equal to $1,500$.

Survey flow The survey starts by informing the participants that they have to watch and rate a total of five YouTube videos. They do not have to watch the videos entirely; rather, I instruct them to watch as much as necessary to answer all questions appropriately. To this end, I let the participants jump back and forth in an embedded practice video before they can proceed.

Next, participants are informed about a potential bonus payment that they can earn. In particular, I tell them that I will randomly draw ten participants, and from each of those ten participants, I will randomly draw one question. If the participant's response to this question is close to the average response, he or she will receive the bonus payment. Questions that are relevant for the bonus payment are marked with a flashy icon. Participants could also click on a "More"-button to see further details. In a drop-down box, I explained that their response must be within a 20%-bandwidth to the average response to qualify them for the bonus payment, and I also gave some concrete examples. Although this incentive scheme could potentially bias responses towards medium values, it is preferred over an experimental setup with no incentive scheme at all.

Then, each participant traverses five screens with identical setup. Each of those screens displays one randomly drawn video from each of the five strata along with several questions about the video:

- First, I inquire the language of the video, where participants could choose between "German", "English", and "Other" (incentivized).
- Next, to generate a measure for a YouTuber's effort, I ask if the video has a "customized intro sequence" and/or a "customized outro sequence" (e.g., a screen that refers to further videos by the YouTuber); both questions are incentivized.
- Then, I ask about "sound quality", "visual quality", and "overall impression" on five-point Likert-scales. The idea is to obtain two relatively objective measures for video quality (sound and visuals) as well as one subjective measure (overall impression). To make this clear, I added the remark "Here, we are interested in your personal opinion" to the question on overall impression. While the questions on sound and visual quality are incentivized, the question on overall impression is exempted from the potential bonus payment.
- Finally, to validate that video tags describe video content appropriately, I displayed four different sets of video tags, where one set is the actual set of tags that the YouTuber assigned to the video, and the other three sets are random draws from the remaining $2,499$ videos in the survey sample. Participants were asked to pick the set of tags that best describes the video (incentivized).

On a final screen, I inquired participants' social media usage, to verify that they are sufficiently experienced with YouTube videos to provide high quality answers to my questions.

D.1.2. Implementation

The online survey experiment was programmed with the survey software Qualtrics and conducted in cooperation with *respondi*, a major German panel provider.[84] Given that the videos in my dataset are mostly German-speaking and targeted at a relatively young German audience, all survey participants had to be German and between 18 and 45 years old; otherwise, I recruited a sample that is representative of the German population and stratified on the federal state level. Participants could use their smartphones, tablets, or desktop PCs to answer the questions; when using a smartphone, I explicitly recommended to use the device in landscape format.

I conducted the survey experiment between March 22 and March 31 in 2022. A total of 1,659 participants completed the survey and passed all attention and quality checks; those who completed the survey received the usual payment by *respondi* plus the potential bonus payment. The median participant spent around 11.36 minutes on the survey.

D.2. Descriptives

First, Figure A.6 confirms that the participants of my online survey experiment are frequent users of social media and thereby sufficiently experienced with the kind of content typically found there: most of them use social media at least several times per week.

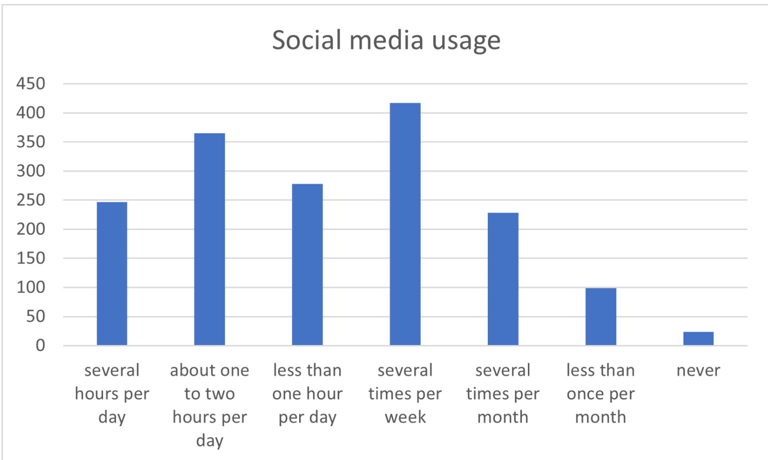

Figure A.6: Social media usage by the participants of the online survey experiment.

Turning towards the YouTube videos, I start by aggregating participants' responses on the video level. I find that the median video was watched three times, with a minimum of 0 and a maximum of 14 times. 76 videos were not drawn at all, whereby the survey yields data on 2,424 videos. More specifically, I obtain information on 480 videos with zero ad breaks before Nov 2015, 491 videos with zero ad breaks after Nov 2015, 488 videos

[84] See `https://www.respondi.com/` for details.

with one ad break before Nov 2015, 485 videos with one ad break after Nov 2015, and 480 videos with multiple ad breaks (both before and after Nov 2015).

D.2.1. Languages

I classify a video as "German" if the majority of participants who watched this video indicate that it is German-speaking. Analogously, I classify a video as "English" or "Other" if the majority of participants agree that it is English-speaking or uses another language, respectively.

Figure A.7 displays the proportions of German-, English-, and Other-speaking videos in the aggregate and for each of the five strata. I find that the vast majority of videos in my sample is German-speaking (74.42% on average), and that the proportions of English- and Other-speaking videos are small (10.60% and 14.98% on average). Moreover, Figure A.7 illustrates that the proportion of German-speaking videos is positively correlated with advertising: while slightly less than 60% of the videos without ad breaks are classified as "German", more than 80% of the videos with one or multiple ad breaks are classified as such, and this proportion seems to increase after Nov 2015. Further analyses of video language can be found in Appendix E.6.

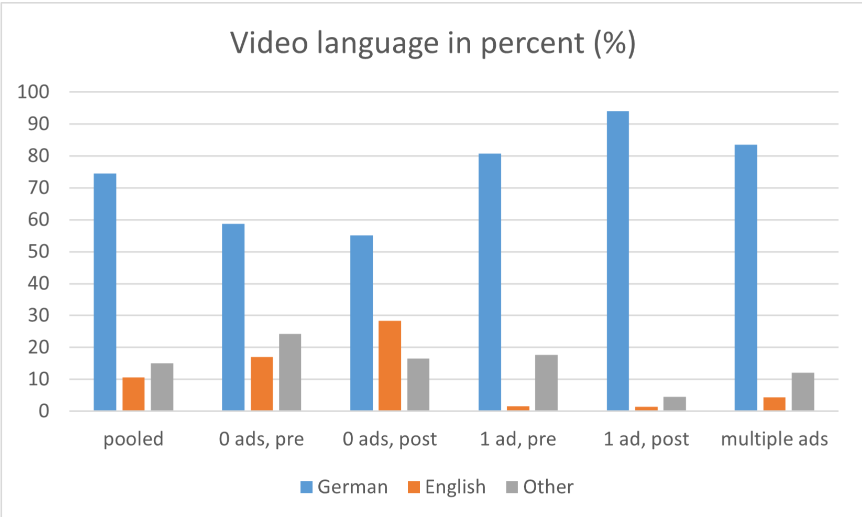

Figure A.7: Video languages in percent, pooled (2,424 videos) and for each of the five strata.

D.2.2. Quality

Next, I examine video quality more closely. To this end, I distinguish between measures that predominantly capture objective facets of video quality, subjective facets of video quality, and general effort put into the video (although it is often difficult to clearly separate those dimensions). Further analyses of video quality can be found in Section 9.

Figure A.8 displays the results of the objective (sound and visuals) and subjective (overall impression) quality evaluation, both in the aggregate and for each of the five strata. The ratings for sound and visual quality are very much aligned (the correlation is equal to 0.75), where the average rating for sound quality is slightly better than the average rating for visual quality. More specifically, the average rating for sound and visuals is slightly above three, where videos with one or multiple ad breaks and videos produced after Nov 2015 score slightly higher. In contrast to that, participants' overall impression of the videos is considerably worse: the average video receives a rating of only 2.55. Both the objective and the subjective quality of videos is larger for videos with at least one ad break and increases over time, although the differences between strata are small.

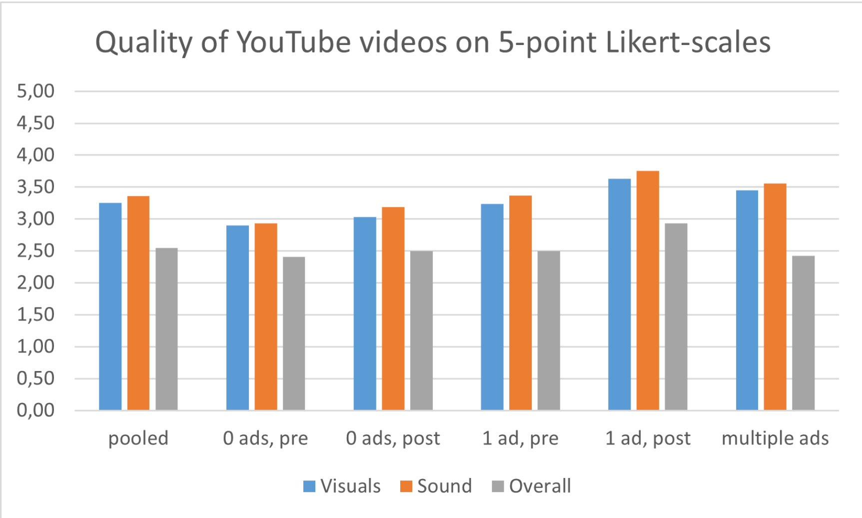

Figure A.8: Quality rating, pooled (2, 424 videos) and for each of the five strata.

Figure A.9 shows the proportions of videos with customized intro or outro sequences. I find that it is more common to feature an intro than an outro, but the majority of videos either displays both kind of sequences (48.31%) or none of them (28.63%). In contrast to the quality measures mentioned above, the differences between strata with regard to intro and outro sequences are relatively large. E.g., while only 18% of the videos without ad breaks and produced before Nov 2015 feature an intro, the same is true for more than 50% of videos with one ad break for the same period of time.

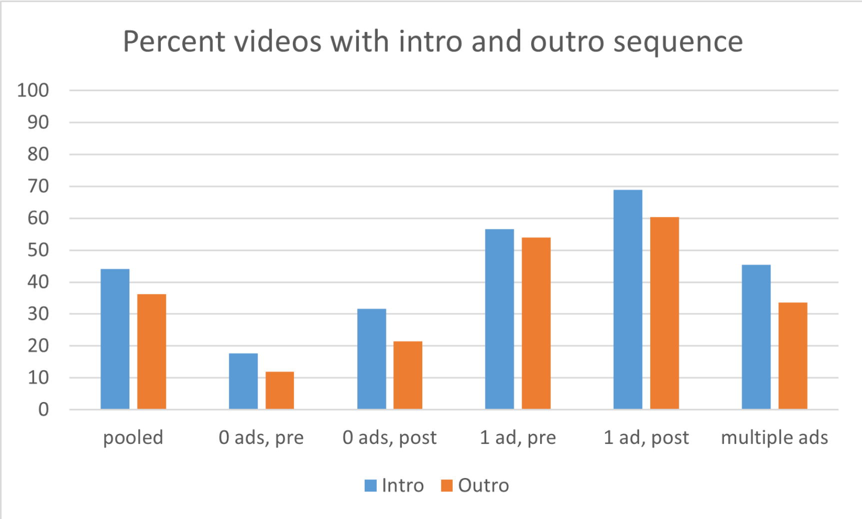

Figure A.9: Percentage of videos where participants could identify the correct set of tags, pooled (2, 424 videos) and for each of the five strata.

D.2.3. Tags

Figure A.10 displays the average proportions of participants who could identify a video's correct set of tags, pooled and for each of the five strata. Reassuringly, the proportions are relatively large – around 80% – and hardly differ between strata.

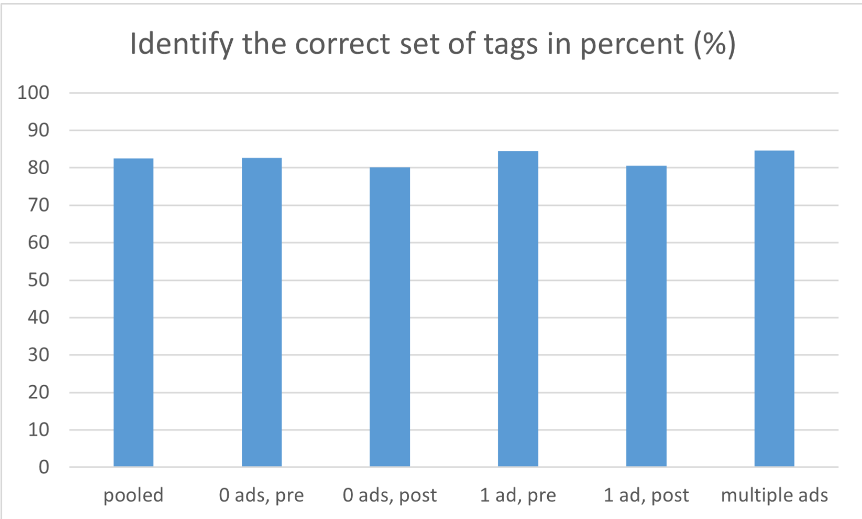

Figure A.10: Percentage of videos with customized intro and/or outro sequence, pooled (2, 424 videos) and for each of the five strata.

E. Further results

This section presents an array of further results that could not be covered in the main part of the paper. In particular, I present anecdotal evidence supporting the main results of the paper, I explore the evolution of video views, differentiation along the tail, viewer fluctuation, and differentiation in the aggregate.

E.1. Anecdotal evidence

This section supports the main story of the paper with anecdotal evidence. I start by exploring the content development of three exemplary YouTubers who have increased their feasible ad quantity after Nov 2015. Then, I analyze evidence from the video-level.

E.1.1. Illustrative examples

This section presents the content development of three exemplary YouTubers from the video categories "Pets & Animals", "People & Blogs", and "Science & Technology" who have increased their feasible ad quantity after Nov 2015. To this end, I illustrate how the YouTubers' proportion of mainstream content per quarter of the year has developed over time, and I consider which tags they use most frequently.[85]

Pets & Animals Figure A.11 illustrates how the proportion of mainstream content for a YouTuber from the category "Pets & Animals" has changed over time. In the beginning of her career, the proportion was relatively small but grew over time until reaching its peak in the first quarter of 2015. The proportion then dropped again and remained relatively low from the forth quarter of 2015, when the new ad break tool was launched (Nov 2015).

Panel A in Table A.9 shows that the YouTuber's most-often used tags have changed accordingly. Apparently, the YouTuber started her career with videos on fancy rats. She eventually started to cover more mainstream animals like cats and hamsters and more mainstream topics like animal shelter and keeping, but stopped covering cats after Nov 2015 to focus again on fancy rats as well as on guinea pigs.

[85] Since many YouTubers only upload one or two (and sometimes zero) videos per month, it is more convenient and more illustrative to aggregate the metrics for each quarter of the year instead of considering them on a monthly basis.

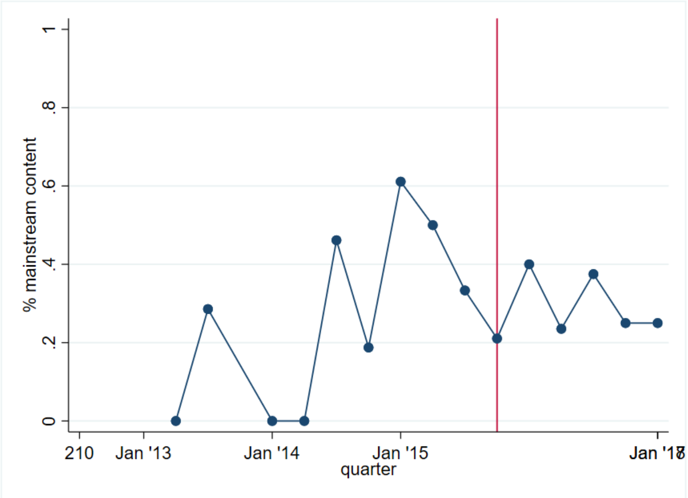

Figure A.11: Development of the proportion of mainstream content over time for an exemplary YouTuber from the video category “Pets & Animals”. The vertical line depicts the forth quarter of 2015, during which the new ad break tool was launched.

People & Blogs Figure A.12 illustrates how the proportion of mainstream content for a YouTuber from the category “People & Blogs” has changed over time. In the beginning of her career, the proportion was relatively high, but it dropped and remained relatively low starting from the forth quarter of 2015, when the new ad break tool was launched.

Panel B in Table A.9 depicts how the YouTuber’s tags have changed along the way. She started with typical mainstream content like make-up tutorials, fashion, hair styles, and cooking, but started to cover more niche topics after Nov 2015; e.g., she started to produce videos on her life as a Filipina in Germany.

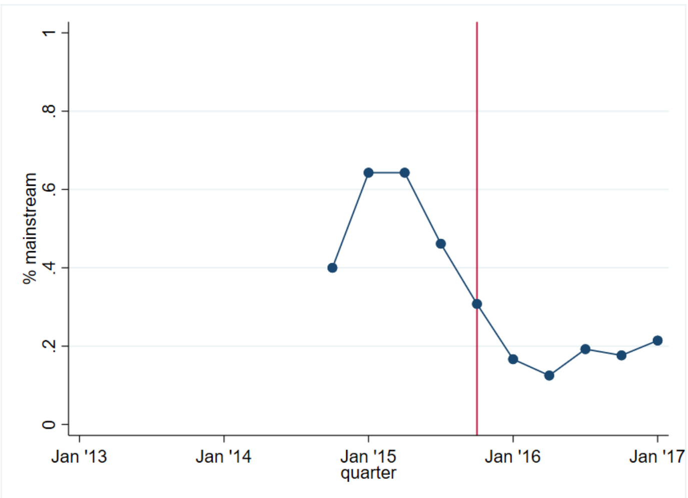

Figure A.12: Development of the proportion of mainstream content over time for an exemplary YouTuber from the video category “People & Blogs”. The vertical line depicts the forth quarter of 2015, during which the new ad break tool was launched.

Science & Technology Figure A.13 illustrates how the proportion of mainstream content for a YouTuber from the category "Science & Technology" has changed over time. While she did not cover any mainstream content in the beginning of her career, the proportion grew sharply in the third quarter and remained high until it dropped again when the new ad break tool was launched in Nov 2015.

Panel C in Table A.9 shows that the YouTuber first produced videos on very specific service tools (e.g., "TS 55R" is a circular hand saw) but eventually became more mainstream by uploading workshops and instructions on woodwork in general. After Nov 2015, however, the YouTuber changed back to her initial focus on tools ("Dewalt", "Festool", and "Bosch" are tool-producing brands).

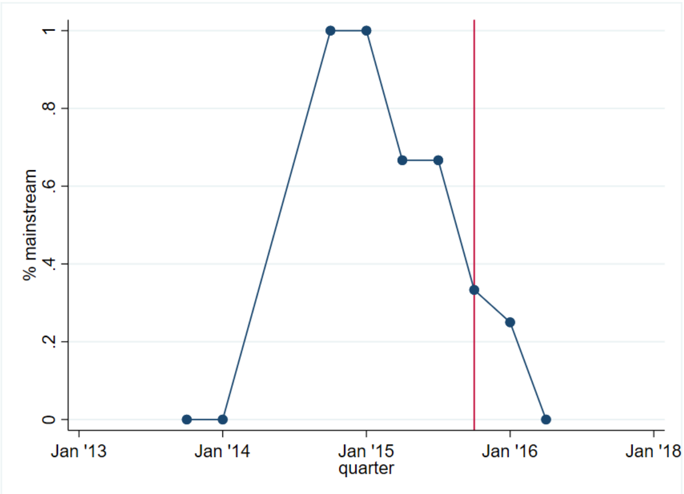

Figure A.13: Development of the proportion of mainstream content over time for an exemplary YouTuber from the video category "Science & Technology". The vertical line depicts the forth quarter of 2015, during which the new ad break tool was launched.

E.1.2. Evidence from the video level

This section provides anecdotal evidence on the *actual* number of ad breaks per video. To this end, I draw a random subsample of 500 advertising YouTubers and collect *video level data* on their monetization settings (52,462 videos); of these, 116 YouTubers (or 23.2%) are classified as treated (see Section 6. Collecting such fine grained information is only feasible for a small subsample of YouTubers; see Section 5.1 for details.

Descriptives I start the analysis with some illustrative evidence. Figure A.14 displays the proportion of videos with zero, one, or several ad breaks. The majority of videos (60.66%) does not permit for any ad breaks at all, and only a small fraction (2.89%) allows for two or more ad breaks, with a maximum number of 52. The average number of ad breaks per video is equal to .466.

Table A.9: Anecdotal evidence

Quarter	**Top 5 tags**
	Panel A: Pets & Animals
2013_02	'fancy rats', 'cute', 'cage', 'nutrition', 'rats'
2013_03	'keeping', 'cute', 'fancy rats', 'rats', 'cage'
2014_01	'keeping', 'cage', 'cute', 'fancy rats', 'rats'
2014_02	'cat', 'cute', 'tomcat', 'keeping', 'animal shelter'
2014_03	'rats', 'cute', 'fancy rats', 'co-housing', 'pups'
2014_04	'rats', 'fancy rats', 'cat', 'diet', 'keeping'
2015_01	'fancy rats' 'rats', 'keeping', 'hamster', 'dwarf hamster'
2015_02	'fancy rats', 'rats', 'keeping', 'hamster', 'dwarf hamster'
2015_03	'rats', 'hamster', 'dwarf hamster', 'fancy rats', 'species-appropriate'
2015_04	'rats', 'fancy rats', 'keeping', 'hamster', 'dwarf hamster'
2016_01	'rats', 'fancy rats', 'keeping', 'species-appropriate', 'correct'
2016_02	'rats', 'fancy rats', 'hamster', 'dwarf hamster', 'species-appropriate'
2016_03	'hybrid', 'hamster', 'dwarf hamster', 'fancy rats', 'rats'
2016_04	'fancy rats', 'rats', 'diet', 'hamster', 'Nager'
2017_01	'medium hamster', 'teddy hamster', 'hamster', 'dwarf hamster', 'multicompartment house'
	Panel B: People & Blogs
2014_04	'gift wrapping', 'how to wrap a gift', 'make-up', 'gift', 'cooking (interest)'
2015_01	'food', 'shoes', 'eating', 'easy cooking', 'food (tv genre)',
2015_02	'ootd', 'hairstyle', 'easy hairstyle', 'make up', 'hair'
2015_03	'look', 'fashion (industry)', 'boot (garment)', 'calvin klein', 'filipino'
2015_04	'makeup', 'outfit', 'ootd', 'cosmetics (quotation subject)', 'german language (language in fiction)'
2016_01	'pinoy', 'pinay', 'filipina', 'filipino', 'pinoy youtuber'
2016_02	'germany vlog', 'pinoy youtuber', 'filipina mom', 'filipino youtuber', 'german vlog'
2016_03	'germany vlog', 'pinoy youtuber', 'filipina', 'pinay mom', 'pinay youtuber'
2016_04	'germany vlog', 'pinay', 'pinay vlog', 'pinoy', 'filipino'
2017_01	'germany vlog', 'pinay', 'germany vlog', 'pinay', 'time lapse'
	Panel C: Science & Technology
2013_04	'fsn', 'ts 55r', 'ts 55', 'makita', 'guard rail'
2014_01	'woodwork, 'stop rails', 'scheppach', 'ts 55', 'bob'
2014_04	'woodworker', 'carpenter', 'workshop', 'joiner', 'pof'
2015_01	'festool', 'bosch', 'carpenter', 'woodworker', 'makita'
2015_02	'gmf', 'routertable', 'woodworker', 'workshop', 'mft'
2015_03	'dewalt', 'festool', 'battery', 'makita', 'first look'
2015_04	'festool', 'bosch', 'woodworker', 'dewalt', 'woodwork'
2016_01	'bosch', 'bessey', 'lamello', 'woodworker', 'woodwork'
2016_02	'woodworking', 'woodwork', 'drill rig', 'woodworker', 'milwaukee'

Notes: Table A.9 displays the top five tags per quarter of the year for three exemplary YouTubers from the categories “Pets & Animals”, “People & Blogs”, and “Science & Technology”. The tags are translated into English.

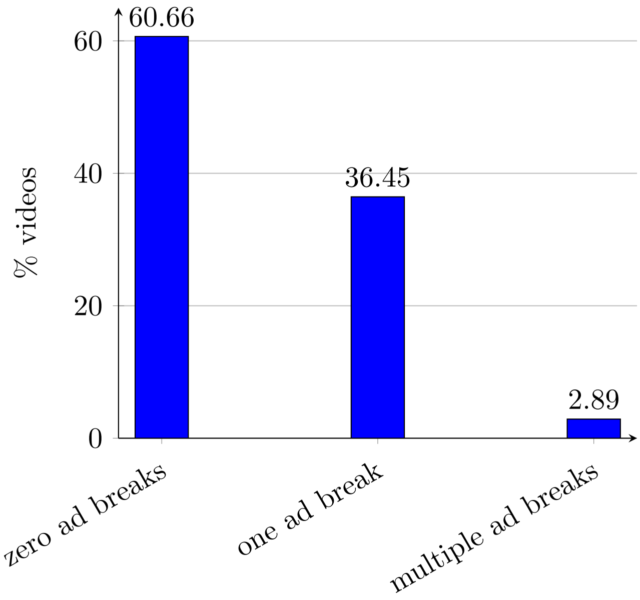

Figure A.14: Proportion of videos with zero, one, and multiple ad breaks based on the entire random subsample of videos.

When I consider only videos that are ten minutes or longer (Figure A.15), I find that the proportion of videos without ad breaks is smaller, both before (52.72%) and after Nov 2015 (51.28%). The proportions of videos with one or multiple ad breaks, in contrast, are larger than in Figure A.14; in particular, the proportion of videos with two or more ad breaks is more than six times as large. The average number of ad breaks in videos that are ten minutes or longer has grown from 0.86 before Nov 2015 to 1.04 after Nov 2015, which corresponds to an increase of 20%. In other words, the actual number of ad breaks has increased both on the intensive and on the extensive margin after the new ad break tool was launched.

Finally, I consider the relationship between video duration and the actual number of ad breaks per video. Specifically, I compute the average number of ad breaks per video for each instance of video duration. Figure A.16 shows the results, cutting the tail at a video duration of 120 minutes (two hours).[86] Plausibly, the number of ad breaks increases with video duration. When I zoom in to the ten minutes threshold in Figure A.17, it also becomes obvious that the actual number of ad breaks sharply increases at the ten minutes threshold. Thus, the illustrative evidence from the video level is in line with the main story of the paper.

[86] An outlier video with 52 ad breaks was dropped here for illustrative purposes.

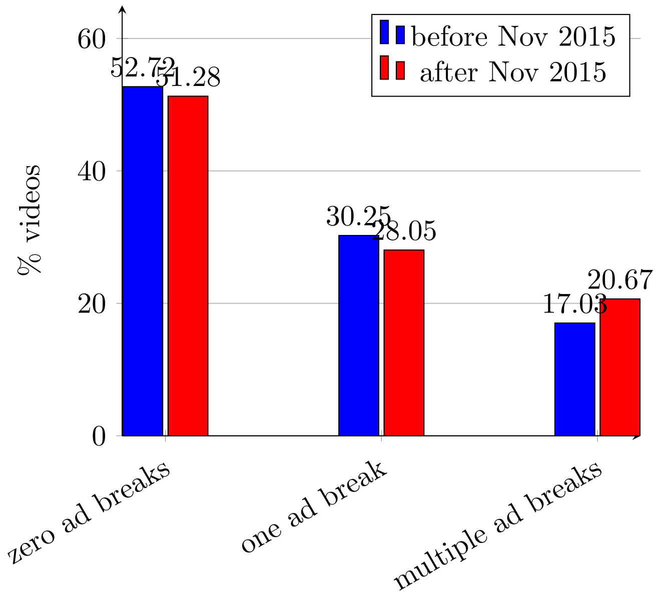

Figure A.15: Proportion of videos with zero, one, and multiple ad breaks before and after Nov 2015 based on all videos that are ten minutes or longer in the random subsample of videos.

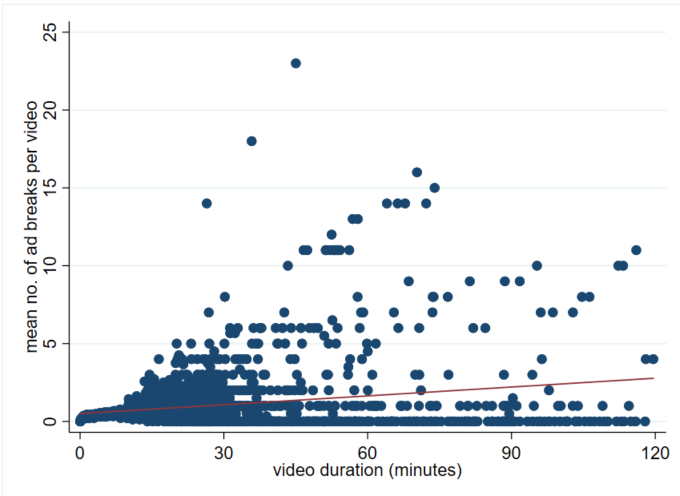

Figure A.16: Average number of ad breaks per video for each instance of video duration below 120 minutes.

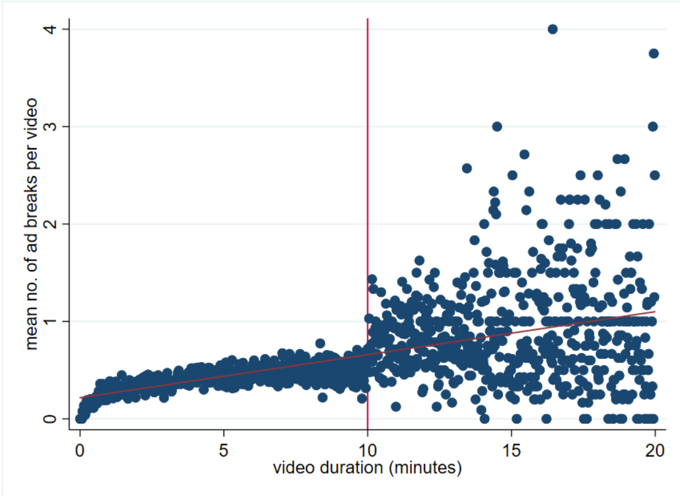

Figure A.17: Average number of ad breaks per video for each instance of video duration below 20 minutes. The vertical line depicts the ten minutes threshold.

Regression analysis Next, I consider the actual number of ad breaks within the framework of my main regression equations (2) and (3). In particular, I replace the treatment indicator D_i with Ads_{vit}, which is equal to the actual number of ad breaks per video. While the main analysis in Section 6 of my paper corresponds to the extensive, the results from this analysis uncover the impact of the intensive margin of advertising on the YouTubers' content choice, conditional on the small sample under consideration.

Table A.10 shows the results from OLS and 2SLS estimations. Analogous to their counterparts in Table 1, the OLS estimates in columns 1 to 3 are close to zero and not statistically significant. In contrast to that, the 2SLS estimates in columns 4 to 6 are negative and (weakly) statistically significant at the 10%- or at the 5%-level. According to these estimates, an additional ad break per video decreases the probability to duplicate mainstream content after Nov 2015 by about 12.3 to 12.7 percentage points; the effect size corresponds to about 25.5% of a standard deviation in the dependent variable and 27.8% of its baseline value. Given the small sample size, the first stage F-statistic is considerably smaller than in Table 1, but still above ten. Similarly, the first stage estimate is positive and highly statistically significant, such that weak instruments are of no concern. In sum, the anecdotal evidence from the video level is absolutely in line with the main results from Section 7.

E.2. Video views

How does the number of video views evolve when YouTubers produce less mainstream content and increase their advertising quantity?[87] Intuitively, one would expect that the number of views goes down, because the videos become ceteris paribus less attractive.

[87] YouTube counts a view if the video is watched for at least thirty seconds; if the video is shorter than that, the viewer must watch it entirely to be counted. See `www.tubics.com/blog/what-counts-as-a-view-on-youtube/` (May 2019).

Table A.10: Evidence from the video-level

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$Ads_{vit} * post_t$	-.001	-.001	-.000	-.123*	-.127**	-.127*
	(.005)	(.005)	(.005)	(.069)	(.063)	(.070)
First stage				.076***	.077***	.071***
				(.016)	(.016)	(.016)
F-statistic				22.02	22.96	20.05
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE		X	X		X	X
Category Time Trend		X	X		X	X
YouTuber Time Trend			X			X
YouTubers	500	500	500	500	500	500
Videos	52,462	52,462	52,462	52,462	52,462	52,462

Notes: Robust standard errors in parentheses. The dependent variable is a dummy equal to one if video v of YouTuber i in month t covers mainstream content. Columns 1 to 3 display OLS, and columns 4 to 6 2SLS estimates. The estimates are based on a sample of 500 advertising YouTubers of whom I could collect video-level information on their monetization settings. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

However, as argued in Section 9, YouTubers might increase their video quality to offset this negative effect. Moreover, YouTube's ranking and recommendation algorithm could favor longer videos with more ad breaks, leading to an increase in video views.

To explore the evolution of video views, I use $log(Views)_{vit}$ as dependent variable for an OLS and a quantile regression estimation of equation (2) as well as a 2SLS estimation of equations (2) and (3). Table A.11 shows the results. The potentially biased OLS estimate in column 1 is positive and statistically significant at the 1%-level. The magnitude of the 2SLS estimate in column 2 is similar, but it is only weakly statistically significant. According to these estimates, an increase in the feasible advertising quantity leads to an *increase* in video views of more than 20% on average. The quantile regression results in columns 3 to 7 show that the effect is qualitatively similar for different quantiles of views, where smaller quantiles benefit more than larger ones. This result is sensible, as less popular videos have more potential to grow.

There are three potential explanations for these results that are not mutually exclusive. First, it is possible that video views increase, because the videos become better. As argued in Section 9, each viewer is more valuable than before if there are multiple ad breaks per video, so YouTubers might increase their video quality to keep viewers (re-)watching. The results from Table 7, however, tend to speak against that. Relatedly, viewers who like niche content could have a higher valuation and/or fewer substitutes for the videos and thus re-watch them more often.

Second, the increased number of views could be reminiscent of YouTube's ranking and recommendation algorithm.[88] Although official information on how the algorithm works

[88]See the white paper by Covington et al. (2016) for a broad idea of how the algorithm operates nowadays.

Table A.11: Views

	OLS	2SLS	Quantile regression estimates				
			.1	.25	.5	.75	.9
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$D_i * post_t$	.215***	.276*	.266***	.245***	.220***	.189***	.158***
	(.026)	(.152)	(.023)	(.022)	(.023)	(.026)	(.030)
First stage		.029***					
F-statistic		150.76					
Time FE	X	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X	X
Category FE	X	X	X	X	X	X	X
Category Time Trend	X	X	X	X	X	X	X
YouTuber Time Trend	X	X	X	X	X	X	X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,081	1,067,081	1,067,081	1,067,081	1,067,081	1,067,081	1,067,081

Notes: In columns 1 and 2, robust standard errors are displayed in parentheses. In columns 3 to 7, standard errors are bootstrapped with 50 replications. All standard errors are clustered on the YouTuber level. The dependent variable is the logarithm of the absolute number of video views, $\log(Views)_{vit}$. The estimates are based on using the advertising YouTubers only. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

is extremely limited, YouTube has made clear in 2012 that it rewards "watch time", i.e., the total accumulated amount of time viewers spend watching a video.[89] If YouTubers upload more videos that are longer than ten minutes, viewers' "watch time" is likely to increase as well; as a consequence, the algorithm is likely to favor those videos when making recommendations.

Third and relatedly, YouTube's algorithm could be commercially confounded. As argued in Appendix A, YouTube retains a share of the YouTubers' advertising revenue. Thus, YouTube has an incentive to navigate viewers to videos with many ad breaks. Anecdotal evidence supports this argument; e.g., Cody Ko – a YouTuber with more than a million subscribers – claims that the algorithm "preferences longer videos, throwing multiple mid-rolls in."[90] Similarly, Bishop (2018) argues that YouTube's algorithm promotes videos that are in line with its commercial aims.

In contrast to that, it is rather unlikely that the results in Table A.11 are driven by deliberate (and successful) manipulation on behalf of the YouTubers. Although some YouTubers attempt to trick the algorithm, it is widely agreed that YouTube is constantly changing, updating, and tweaking the system, making it close to impossible to really understand what is going on.[91] Instead, YouTube recommends to focus on making good videos "rather than trying to find a secret code to these systems"; similar advice is given

[89] In particular, YouTube claims that "watch time" is a better metric for video quality than the number of views, because the former better reflects viewer satisfaction. See `https://blog.youtube/news-and-events/youtube-now-why-we-focus-on-watch-time/` (March 2022).

[90] See `https://digiday.com/future-of-tv/creators-making-longer-videos-cater-youtube-algorithm/` (Jan 2022).

[91] See, e.g., `https://vidiq.com/blog/post/7-reasons-why-youtube-channel-losing-views/` (Dez 2021). Alternatively, type "YouTube algorithm" in YouTube's search bar, which yields plenty of videos trying to explain how the algorithm works at some arbitrary point in time.

in explanatory videos provided by the platform.[92]

E.3. Differentiation along the tail

In Section 7, I consider the effect of an increase in the feasible number of ad breaks per video on the YouTubers' probability to duplicate the *most* mainstream content only. Here, I study content differentiation further along the "tail". In particular, I generate five dummy variables that indicate if a video is given at least one tag from alternative quantiles of the distribution of most-viewed tags: the 1^{st} to 10^{th}, the 10^{th} to 25^{th}, the 25^{th} to 50^{th}, the 50^{th} to 75^{th}, and the 75^{th} to 100^{th} percentile. Note that these categories are not mutually exclusive and videos could be given tags from each of these quantiles. Then, I replace the dependent variable in equation (2) with each of these dummies and estimate equations (2) and (3) by 2SLS.

The results in Table A.12 illustrate a consistent pattern. The estimate in column 1 is similar to its counterpart in Table 1: an increase in the feasible number of ad breaks per video leads to a 20% percentage point reduction in the probability to upload a video that is given a tag from the 1^{st} to 10^{th} percentile of the distribution of most-viewed tags. However, the estimate decreases by half in columns 2 and 3, and by about two-thirds in column 4. Finally, in column 5, the estimate switches its sign and becomes positive. All estimates are statistically significant.

To interpret these results, note that a video is given around twelve tags on average and that this number is constant over time. Hence, a video can be given both mainstream tags from the upper quantiles and niche tags from the lower quantiles of the distribution. Bearing this mind, Table A.12 demonstrates that YouTubers who could increase their advertising quantity likely change the "mixture" of tags in a video: they abandon the more mainstream tags and use a larger number of niche tags instead. However, since they were often using niche tags next to mainstream tags already, the probability to upload a video with at least one niche tag remains unchanged or increases only slightly. Indeed, when I count each video's number of "quantile affiliations" and use this number as dependent variable in equation (2), a 2SLS estimation shows that videos from YouTubers who could increase their feasible ad quantity after Nov 2015 are given tags from fewer different quantiles than before (column 6).

E.4. Viewer fluctuation

This section studies the development of the YouTubers' viewer fluctuation to support competition as a main mechanism behind the results from Section 8. More specifically, if YouTubers differentiate their content from the mainstream to prevent their audience from switching to a competitor, viewer fluctuation should go down, i.e., a given number of views should be generated by a smaller circle of viewers. As precise information on a YouTuber's

[92]See `https://www.youtube.com/watch?v=gTrLniP5tSQ` and `https://www.youtube.com/watch?v=hPxnIix5Ex` (Dez 2021).

Table A.12: Differentiation along the tail

	1^{st} to 10^{th} percentile (1)	10^{th} to 25^{th} percentile (2)	25^{th} to 50^{th} percentile (3)	50^{th} to 75^{th} percentile (4)	75^{th} to 100^{th} percentile (5)	Quantile affiliations (6)
$D_i * post_t$	-.208***	-.101***	-.102***	-.059*	.086***	-.376***
	(.042)	(.038)	(.040)	(.036)	(.032)	(.096)
First stage	.029***	.031***	.031***	0.031***	.031***	.031***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
F-statistic	152.17	165.67	165.67	166.98	168.34	168.34
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE	X	X	X	X	X	X
Category Time Trend	X	X	X	X	X	X
YouTuber Time Trend	X	X	X	X	X	X
YouTubers	10,599	10,591	10,591	10,590	10,589	10,589
Videos	1,064,248	1,033,666	1,033,666	1,031,051	1,028,446	1,028,446

Notes: Robust standard errors in parentheses. Each column displays the results of a 2SLS estimation. In column 1, the dependent variable is an indicator equal to one if video v of YouTuber i in month t is given a keyword from the 1^{st} to 10^{th} percentile of the distribution of most-viewed tags. Analogously for columns 2 to 5. In column 6, the dependent variable is the sum of a video's percentile indicators. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

viewership is not available, I obtain all videos' comments (including the commentators' aliases) from the YouTube Data API, and employ a YouTuber's commentators as a proxy for her viewers.

The commentator fluctuation of YouTuber i is defined as

$$fluctuation_i = \frac{commentators_i}{comments_i}, \tag{41}$$

where $commentators_i$ corresponds to the size of the set of YouTuber i's eventual commentators, and $comments_i$ corresponds to the total number of comments that i receives. If each comment is written by a different user, commentator fluctuation is maximal and $fluctuation_i$ is equal to 1. The more comments are written by the same commentators, the smaller $fluctuation_i$. If YouTuber i never receives any comment, $fluctuation_i$ is not defined.

Next, I compute each YouTuber's change in $fluctuation_i$ before and after Nov 2015,

$$\Delta fluctuation_i = fluctuation_{i,post} - fluctuation_{i,pre}, \tag{42}$$

where $fluctuation_{i,post}$ is based on the fifteen months after, and $fluctuation_{i,pre}$ is based on the fifteen months before Nov 2015.[93] If the commentator fluctuation of YouTuber i goes down, $\Delta fluctuation_i < 0$.

[93] Since I have 34 observation periods before Nov 2015, but only fifteen observation periods afterwards (including Nov 2015), I restrict the computation of $fluctuation_{i,pre}$ to the fifteen most recent ones to increase the comparability to $fluctuation_{i,post}$.

To demonstrate that an increase in the feasible advertising quantity and the corresponding decrease in the YouTubers' probability to upload mainstream content diminishes commentator fluctuation, I use $\Delta fluctuation_i$ as dependent variable in

$$\Delta fluctuation_i = \rho_0 + \rho_1 D_i + \epsilon_i, \tag{43}$$

where ρ_1 denotes the average effect of having the option to increase the number of ad breaks per video on the change in commentator fluctuation of YouTuber i. To account for endogeneity in D_i, I use

$$D_i = \psi_0 + \psi_1 close_i + e_i, \tag{44}$$

as a first stage and estimate equations (43) and (44) by 2SLS. Since $fluctuation_i$ is sensitive to additional commentators when the total number of comments is small – e.g., if a YouTuber has only received three comments, it makes a big difference if they are written by two or three different commentators – I restrict the analysis to YouTubers who received at least 25 (50, 100) comments before and after Nov 2015.

Table A.13 shows the results for the three different thresholds. All estimates are negative and statistically significant at the 5%- or at the 1%-level. The effect size corresponds to about 38% of a standard deviation in the dependent variable in column 1, 43% in column 2, and 60% in column 3. Hence, the estimates are in line with the idea that viewer fluctuation decreases when YouTubers move to a niche.

Table A.13: Commentator fluctuation

	> 25 (1)	> 50 (2)	> 100 (3)
D_i	-0.048**	-0.050**	-0.066***
	(0.023)	(0.022)	(0.024)
First stage	0.029***	0.031***	0.031***
	(0.002)	(0.002)	(0.002)
F-statistic	159.78	150.96	120.47
YouTubers	5,907	4,924	3,989

Notes: Robust standard errors in parentheses. The dependent variable is $\Delta fluctuation_i$, which is the difference in the commentator fluctuation before and after Nov 2015 for YouTuber i. The estimates are based on the advertising YouTubers only. In column 1 (column 2, column 3), only YouTubers who received more than 25 (50, 100) comments before and after Nov 2015 are included in the analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

E.5. Differentiation in the aggregate

While the main part of the paper studies how an increase in the feasible advertising quantity affects *individual* YouTubers' content choice, this section explores how content differentiation develops *in the aggregate*. In particular, I study if the tail becomes "longer" (i.e., if the total number of unique tags increases), and if the tail becomes "fatter" (i.e., if the concentration of videos on tags decreases). I do not make causal claims here; rather, I pursue a descriptive before-after comparison to put the results from Sections 7 and 8 into a broader context.

There are two options to analyze content differentiation in the aggregate: I could either continue to focus on the subsample of YouTubers whom I selected for the main analysis in Section 5.3, or I could examine the entire population of German YouTubers. Proceeding with the subsample has the advantage of computing aggregate measures that are solely based on YouTubers who have the option to increase their feasible number of ad breaks per video, but the approach does not reveal how the *entire* video supply on YouTube develops after Nov 2015. Even YouTubers who are not directly affected by the launch of the new ad break tool may adapt their video content as a reaction to their competitors' change in content; thus, studying the population of YouTubers might be more informative about aggregate developments. On the other hand, the content choices of YouTubers whom I did not select for the main analysis could be driven by motives that are orthogonal to the launch of the new ad break tool and its consequences; such effects might superimpose the treatment's aggregate effect on content differentiation and complicate the interpretation of the net effect. Since no approach clearly excels the other, I pursue both options and interpret the results accordingly.

E.5.1. The tail becomes longer

To show that the tail of keywords becomes longer both within the subsample and the entire population of YouTubers, I compute the absolute number of unique tags before and after Nov 2015. As I observe 34 months before Nov 2015, but only 15 months afterwards (including Nov 2015), I limit the analysis to the 15 most recent months before Nov 2015.

In the subsample, there exist $607,358$ unique tags before, and $875,503$ unique tags after Nov 2015, which corresponds to an absolute increase of $268,145$ unique tags and to a relative increase of 44.15%. Considering the population of YouTubers, I find that there exist $1,090,355$ unique tags before, and $2,096,373$ unique tags after Nov 2015, which corresponds to an absolute increase of $1,006,018$ tags and to a relative increase of 92.27%. The results match the findings from Sections 7 and 8: it is plausible that the total number of unique tags increases when the YouTubers reduce the probability to upload mainstream or competitive content. The difference in the results between the subsample and the population could stem from entry: by construction, the population also includes YouTubers who entered the platform after Nov 2015, which may further increase the number of unique tags after Nov 2015.

E.5.2. The tail does not become fatter

To study if the tail becomes "fatter", I compute a Gini coefficient for the concentration of videos on tags before and after Nov 2015.[94] Again, I restrict the analysis to the 15 months before and after Nov 2015. Note that the Gini coefficient for the subsample measures the concentration of videos on tags that occur *within the subsample*, while the Gini for the population measures the concentration of *all* videos on *all* tags.

The Gini coefficient for the subsample is high and remains nearly unchanged: it is equal to 0.800 before, and equal to 0.806 after Nov 2015, which corresponds to an increase of 0.75%. This does not contradict the findings from Section 8, however. My measures for mainstream and competitive content are based on *all* active German YouTubers. It is therefore possible that the YouTubers in the subsample decrease their probability to upload competitive content, where competitive content takes *the population of* YouTubers into account, but that the concentration of videos on tags *within* the subsample remains nearly unchanged. In addition to that, the tail of tags becomes longer after Nov 2015 (see Section E.5.1). If many of those additional tags are used by a small number of videos, the Gini coefficient as a *relative measure of concentration* remains unchanged even if the concentration of videos on the remaining tags decreases.

The Gini coefficient for the entire population of German YouTubers increases from 0.848 before to 0.862 after Nov 2015, which corresponds to an increase of 1.65%. Here, too, the increase in the relative concentration measure could be due to the large amount of additional tags. It is also possible that further developments – orthogonal to the launch of the new ad break tool – superimpose the effect of an increase in the feasible number of ad breaks on content differentiation in the aggregate. For instance, the growing popularity of the platform may have led to a large number of entrants who copy from the most popular YouTubers and thereby increase the concentration of videos on tags.

E.6. German vs. English

As argued in Section 5, I collect data on all active *German* YouTubers as of October 2017. Sticking to one language has two advantages. First, articulating a specific term or phrase requires different amounts of time for different languages; e.g., it usually takes longer to say something in German than to say it in English. Given that video duration is a crucial metric in my analysis, focusing on German prohibits potential confounding effects of considering several languages. Second, I study competition as an economic mechanism behind my main results (see Section 8). To conduct this kind of analysis, I must capture *all* potential competitors in my data. Here, it helps to focus on German YouTubers, as it is plausible to assume that the universe of German YouTubers constitutes an isolated market. In particular, non-German speakers are unlikely to perceive German-speaking

[94] I.e., the tags correspond to the households, and the number of videos that use a certain tag corresponds to the income in a conventional Gini computation. Note, also, that I cannot use absolute measures of concentration such as the Herfindahl index, because the number of tags before and after Nov 2015 is different.

videos as a consumption option and vice versa. The only exception to this rule might be English, because many Germans speak English quite well.[95]

Based on this idea, I use the development of video language to further support competition as an economic mechanism behind my main results. Specifically, if YouTubers move to a niche to avoid competition, they might also increase the proportion of German-speaking videos to reduce competitive pressure from English-speaking alternatives.

To identify video language, I let two widely-used language identification algorithms – *langid.py*[96] and *langdetect.py*[97] – classify title and tags of each video in terms of three mutually exclusive categories: *German*, *English*, and *Other*. Table A.14 shows that both algorithms classify the majority of videos as German and about a quarter of them as English. However, both algorithms struggle with some YouTubers' Germish ("unboxing", "haul", "make-up") and with non-translated titles of films and video games ("Avengers", "Assassin's creed"), so the true proportion of German-speaking videos is likely to be higher than the proportions displayed in Columns 1 and 2 of Table A.14.

To further evaluate the performance of *langid.py* and *langdetect.py*, I compare their output to the language classification of human coders from the online survey experiment (see Appendix D for details). Column 3 in Table A.14 shows that the human coders indeed classify a larger proportion of videos as German- and a smaller proportion as English-speaking; given that the human coders actually watched (parts of) the videos, I consider these numbers as more reliable. Note, however, that I could only classify a small subsample of $2,424$ videos by human coders. When I directly compare the human to the algorithmic classification, I find that their output concurs in 73.76% (*langid.py*) and 65.71% of cases (*langdetect.py*); hence, *langid.py* performs better on my data.

Table A.15 shows the 2SLS results from using the language dummies as dependent variables in equation (2), where the estimates in columns 1 to 3 are based on the *langid.py* and the estimates in columns 4 to 6 are based on the *langdetect.py* algorithm. According to both algorithms, the proportion of German-speaking videos increases while the proportion of English-speaking videos decreases when YouTubers increase their feasible advertising quantity; the proportion of videos in other languages remains nearly unchanged. The estimates based on the *langid.py* algorithm are larger and more statistically significant than the estimates based on *langdetect.py*. The effect size is rather modest for both specifications, though (smaller than 10% of a standard deviation in the dependent variable). In sum, the results from Table A.15 are in line with the idea that YouTubers who could increase their feasible advertising quantity after Nov 2015 become more nichy to avoid competitive pressure.[98]

[95] See, e.g., the English Proficiency Index (EPI) at `https://www.ef.com/wwen/epi/` (March 2022).

[96] See `https://github.com/saffsd/langid.py` (April 2022).

[97] See `https://pypi.org/project/langdetect/` (April 2022).

[98] Producing English-speaking videos could also be interpreted as a costly investment into video quality. In this case, the results from Table A.15 could be interpreted as a decrease in video quality and are thereby in line with the results from Section 9, where I also document that – if anything – video quality decreases when YouTubers increase the number of ad breaks per video.

Table A.14: Language classification

	langid.py	***langdetect.py***	***Human***
German	58.31%	51.98%	74.42%
English	25.55%	22.85%	10.60%
Other	16.14%	25.17%	14.98%

Notes: Columns 1 and 2 show the results from letting the language identification tools *langid.py* and *langdetect.py* classify the language of title and tags of each video; the displayed proportions are based on all 1,067,542 videos uploaded by advertising YouTubers during my observation period. Column 3 shows the results from human coders who watched a small subset of videos; the proportions are based on the 2,424 videos that were analyzed in my online survey experiment (see Appendix D for details).

Table A.15: Development of video language

	langid.py			*langdetect.py*		
	German	English	Other	German	English	Other
	(1)	(2)	(3)	(4)	(5)	(6)
$D_i * post_t$	.083***	-.082***	-.001	.039	-.058*	.019
	(.031)	(.029)	(.021)	(.032)	(.031)	(.031)
First stage	.029***	.029***	.029***	.029***	.029***	.029***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
F-statistic	150.65	150.65	150.65	150.65	150.65	150.65
Time FE	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X
Category FE	X	X	X	X	X	X
Category Time Trend	X	X	X	X	X	X
YouTuber Time Trend	X	X	X	X	X	X
YouTubers	10,599	10,599	10,599	10,599	10,599	10,599
Videos	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parantheses. The dependent variable in columns 1 and 4 is a dummy variable equal to one if video v by YouTuber i uploaded in month t is classified as German. The dependent variable in columns 2 and 5 is a dummy variable equal to one if video v by YouTuber i uploaded in month t is classified as English. The dependent variable in columns 3 and 6 is a dummy variable equal to one if video v by YouTuber i uploaded in month t is classified as another language. Coluns 1 to 3 use the language classification by the *langid.py* package. Columns 4 to 6 use the language classification by the *langdetect.py* package. All estimates are based on the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

F. Further robustness checks

This section probes the robustness of my results. In particular, I show that my main results are robust to using an alternative observation period, to alternative selections of YouTubers, to alternative classifications of the treatment group, to alternative definitions of the instrument, and to an alternative definition of competitive pressure.

F.1. Alternative observation period

First, I show that my main results are robust to using an alternative observation period. As argued in Section 5.3, I cannot sensibly extend the analysis to earlier or later points in time; I can, however, select a shorter observation period. Column 1 in Table A.16 shows the results from estimating equations (2) and (3) by 2SLS on observations from Jan 2014 to July 2016 only. Unsurprisingly, the estimates are smaller than in Section 7.1: as the effect becomes stronger over time (Figure A.29), excluding the last six months from the analysis results in a smaller effect on average.

F.2. Alternative selections of YouTubers

Next, I demonstrate that my main results are robust to more restrictive selections of YouTubers. First, I consider only the subsample of YouTubers whose median video duration before Nov 2015 is smaller 7.5 minutes; second, I focus on the subsample of YouTubers whose 90^{th} percentile of the distribution of video durations (not the median) is smaller than 10. Columns 2 and 3 in Table A.16 show that the 2SLS estimates are similar to the effects documented in Section 7.1, and the first stage is even stronger.

F.3. Alternative classifications of the treatment group

This section shows that my main results are robust to alternative classifications of the treatment group. In particular, I show that neither the five percentage point cutoff nor considering only a YouTuber's videos between ten and fourteen minutes drive my results.

Columns 4 and 5 in Table A.16 show the 2SLS estimates from using two alternative cutoffs; YouTubers are classified as treated if their share of videos between ten and fourteen minutes has increased by at least one (column 4) or by at least ten percentage points (column 5). Plausibly, the estimates are even larger than the main estimates in Table 1: the average effect of an increase in the feasible number of ad breaks on the probability to upload mainstream content is stronger for YouTubers who increase their share of videos that are ten minutes or longer to a higher extent.

In column 6, I classify a YouTuber as treated if she increased her share of videos that are ten minutes or longer (instead of the share between ten to fourteen minutes) by at least five percentage points. The 2SLS estimates are negative and statistically significant at the 1%-level, but smaller than their counterparts in Table 1. A potential explanation is that considering *all* videos that are ten minutes or longer leads to more noise in the estimation, for instance, because videos that are more than "just" longer than ten minutes are less likely to indicate that a YouTuber exploits the ten minutes trick.

F.4. Alternative definitions of the instrument

Next, I confirm that my main results are robust to alternative definitions of the instrument. More specifically, while $close_i$ corresponds to YouTuber i's *median* video duration before Nov 2015 in my main analysis, it is equal to the 75^{th} and to the 90^{th} percentile of the distribution of a YouTuber's video durations here. Columns 7 and 8 in Table A.16 shows that the 2SLS are negative, but smaller and not quite as statistically significant as in my main analysis. Also, the first stage is relatively weak in both instances. Hence, the 75^{th} and the 90^{th} percentiles of the distribution of a YouTuber's video durations before Nov 2015 have less power to predict a YouTuber's treatment status D_i than the median.

F.5. Alternative definition of competition

Finally, I demonstrate that the results from Section 8 are robust to using an alternative definition of competitive pressure. More specifically, I consider a measure for "competitive content" that is more analogous to my measure for mainstream content as described in Section 5.2. To this end, I compute how many times a certain tag has been *used* in each month and video category and rank them in descending order; the upper one percent of this distribution is classified as "competitive." Then, I assign a dummy variable that is equal to one to all videos equipped with a competitive tag. Note that a competitive tag is not necessarily a mainstream keyword, too. A tag may attract many views although it is not used by many YouTubers; similarly, a tag may be used by many YouTubers, but does not attract many views. In my sample, the correlation between mainstream and competitive content is equal to 0.57. See Section 8 for further discussion on the difference between my metrics for competition and mainstream content.

Column 9 in Table A.16 shows the 2SLS estimate from using my alternative measure for competitive content as dependent variable in equation (2). The estimate is negative and highly statistically significant at the 1%-level. According to the estimate, an increase in the feasible number of ad breaks per video decreases the probability to upload competitive content by about twenty percentage points; the effect size corresponds to 42% of a standard deviation in the dependent variable and to around 30% of its baseline value. Results qualitatively in line with the results in Section 8. The results very similar to the main results in Table 1, which is no surprise, given the correlation.

Table A.16: Further robustness checks

	Time	YouTubers		Treatment group			Instrument		Competition
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$D_i * post_t$	-.160***	-.276***	-.286***	-.272***	-.386***	-.166***	-.144**	-.384***	-.191***
	(.046)	(.049)	(.058)	(.062)	(.084)	(.035)	(.072)	(.146)	(.046)
First stage	.030***	.040***	.053***	.024***	.017***	.039***	.007***	.002***	.029***
	(.002)	(.003)	(.005)	(.003)	(.002)	(.003)	(.001)	(.001)	(.002)
F-statistic	152.58	166.12	134.34	71.20	100.69	210.21	24.92	14.89	150.65
Time FE	X	X	X	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X	X	X	X
Category FE	X	X	X	X	X	X	X	X	X
Category Time Trend	X	X	X	X	X	X	X	X	X
YouTuber Time Trend	X	X	X	X	X	X	X	X	X
YouTubers	10,513	9,519	6,891	10,599	10,599	10,599	10,599	10,599	10,599
Videos	745,219	923,189	610,496	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542	1,067,542

Notes: Robust standard errors in parentheses. All estimates are 2SLS estimates. The dependent variable is a dummy variable equal to 1 if video v of YouTuber i uploaded in month t is equipped with a mainstream keyword, and 0 otherwise. In column 1, the analysis is restricted to the time periods $t \in [13, 43]$. In column 2, only YouTubers whose median video duration before Nov 2015 is smaller than 7.5 min, and in column 3, only YouTubers whose 90^{th} percentile of the distribution of video durations before Nov 2015 is smaller than 10 min are considered. In column 4, YouTubers who have increased their share of videos between ten and fourteen minutes by at least 1 percentage point, and in column 5, YouTubers who have increased their share of videos between ten and fourteen minutes by at least 10 percentage points after Nov 2015 are classified as treated. In column 6, YouTubers who have increased their share of videos that are ten minutes or longer by at least five percentage points are classified as treated. In column 7, the instrument $close_i$ is defined as the 75^{th} percentile in the distribution of video durations of YouTuber i before Nov 2015, and in column 8 as the 90^{th} percentile in the distribution of video durations of YouTuber i before Nov 2015. In column 9, I use an indicator for competitive content as dependent variable. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

G. Further discussion

This section revisits a number of topics that could not be covered in the main part of the paper. In particular, I discuss conceptual differences between the actual and the feasible number of ads, differences between and advertising and non-advertising YouTubers, I elaborate on the per-view price of advertising, I review how my analysis relates to research on consumer switching costs, I show that no YouTube platform event beyond the launch of the new ad break tool affects my results, and I discard a YouTuber learning effect as a potential economic mechanism behind my results.

G.1. Actual vs. feasible number of ads

Rather than studying the relationship between the actual number of ad breaks per video and content differentiation, my analysis unveils how a YouTuber's *option* to increase her advertising quantity affects her probability to duplicate mainstream content. This approach has three main advantages:

First, the instrumental variable $close_i$ is defined on the YouTuber-level, whereby the IV analysis can isolate exogenous variation in a YouTuber's *feasible* advertising quantity, but not in the *actual* number of ad breaks per video. In particular, while it is plausible to assume that variation in $close_i$ affects whether a YouTuber taken as a whole gains the option to increase her feasible ad quantity or not, it is less clear if variation in $close_i$ affects whether a YouTuber permits for one, two, or even more additional ad breaks in a specific video.

Second, YouTubers can modify their monetization settings at any point in time (e.g., they could permit for two ad breaks today, four ad breaks tomorrow, and three ad breaks next week), whereby using the actual number of ad breaks per video as observed on the day of data collection could entail measurement error. The option to have more ad breaks, in contrast, remains constant over time.

Finally, studying the effect of an increase in the feasible number of ad breaks is likely to yield more interesting policy implications. As argued in the main part of the paper, an increase in the feasible number of ad breaks corresponds to repealing restrictions on the YouTubers' advertising quantity, and the insights from this analysis are likely to have higher external validity than studying the effect of a specific number of additional ad breaks per video.

G.2. Advertising and non-advertising YouTubers

G.2.1. Economic incentives

The theoretical framework in Appendix A – though highly stylized – illustrates the economic incentives of advertising YouTubers: they choose their advertising quantity (if possible) and degree of content differentiation to maximize their profit, and they enter the Platform if their fixed costs are covered. An abundance of anecdotal evidence supports the relevance of money-making for advertising YouTubers. E.g., numerous videos and blogs give advice on "how to make money on YouTube fast"[99], and it appears as though many YouTubers participate on the platform to become rich and famous.[100]

On the other hand, many YouTubers do not permit for ad breaks at all. In my sample, roughly a third of the YouTubers deliberately foregoes advertising revenue. Hence, a substantial share of YouTubers bears the fixed costs of entry without the prospect of being compensated.[101] Why?

First, many non-advertising YouTubers have a relatively small audience.[102] Thus, they would not earn much money in absolute terms even if they chose to monetize their video content, so many prefer not to bother.[103] Importantly, at the time of data collection (late 2017), *every* YouTuber could join the YouTube Partner Program and subsequently permit YouTube to show ads before or during her videos, irrespective of her subscriber count. The infamous $1,000$-subscriber threshold was introduced in Feb 2018, when the data collection for this paper was completed.[104]

Second, it is plausible to assume that many YouTubers derive pleasure from video creation and perceive it as a hobby. E.g., there exists plenty of anecdotal evidence where YouTubers report how YouTube allows them to "unleash their creativity", to "pursue their passion", how they find "fulfillment and purpose" in producing video content, and how much effort they put into their videos even if their audience is small.[105] Thus, the pure pleasure from video creation is often enough to compensate the YouTubers for their fixed costs.

Third, some YouTubers really want to convey a certain message and reach as many people as possible.[106] In this case, they may perceive it as inappropriate to make money with their video. Relatedly, some YouTubers feel that it is wrong to make money with their hobbies.[107]

Finally, it may be a deliberate business strategy to forego advertising revenue. E.g.,

---

[99] See, e.g., https://www.youtube.com/watch?v=I3MeCEwVxB0, https://www.youtube.com/watch?v=SQn8D7B0ef0, https://blog.hootsuite.com/how-to-make-money-on-youtube/, https://blog.sellfy.com/how-to-make-money-on-youtube/, https://www.youtube.com/creators/how-things-work/video-monetization/ (Nov 2021).

[100] See, e.g., https://www.washingtonpost.com/news/the-intersect/wp/2018/06/25/they-became-famous-youtubers-a-new-generation-of-kids-wants-to-take-their-place/. (Nov 2021)

[101] The fixed costs of entry may be small, but a YouTuber needs at least some time and basic equipment like a camera and software to cut and edit the video.

[102] In my main sample, the average number of subscribers of an advertising YouTuber is equal to $25,701$, as opposed to $3,241$ for a non-advertising YouTuber. Similarly, the median number of subscribers of an advertising YouTuber is equal to 928 as opposed 150 for a non-advertising YouTuber.

[103] This is, for instance, one of the reasons why this YouTuber prefers not to monetize her content: https://www.youtube.com/watch?v=mnvRApKqkUk. (Nov 2021)

[104] See Abou El-Komboz et al. (2022) for further discussion on the $1,000$-subscriber threshold and its consequences.

[105] See, e.g., https://www.youtube.com/watch?v=xXXX9pSo9U8, https://www.youtube.com/watch?v=_eBl74SvV_0, and https://www.youtube.com/watch?v=-zwbWRUkQyM (Nov 2021.)

[106] E.g., this YouTuber talks about his depression and how he tried to commit suicide: https://www.youtube.com/watch?v=7qEl6JiffZg (Nov 2021).

[107] See, e.g., https://www.youtube.com/watch?v=cC4fkpuCdVk (Nov 2021).

some YouTubers may choose other channels to make money[108], or decide to first grow a larger audience.[109]

Naturally, these incentives to participate on YouTube are not mutually exclusive and most of them could apply to advertising YouTubers, too. In sum, however, it is plausible to assume that non-advertising YouTubers tend to be more intrinsically motivated, while advertising YouTubers (additionally) seek profit.

G.2.2. Misclassification

As explained in Section 5.1, I cannot retrieve data on the YouTubers' monetization settings on the video level. Instead, I pick twenty randomly drawn videos per YouTuber, and classify her as advertising YouTuber if I detect at least one ad break. In this section, I amplify how measurement errors during this procedure could affect my results. In addition, I discuss the consequences of sample migration between advertising and non-advertising YouTubers.

Potential consequences of measurement error In this section, I illustrate that a potential measurement error would have only minor consequences. First, note that I could erroneously classify an advertising YouTuber as non-advertising, but not vice versa: if a YouTuber never permits for ad breaks, my algorithm cannot classify her as "advertising" by definition. Second, note that I do not use the classification dummy in a regression framework; hence, the regression results do not suffer from an errors-in-variables bias (e.g., Durbin, 1954). Yet, I *split* my sample into advertising and non-advertising YouTubers. Thus, misclassifying some advertising as non-advertising YouTubers might lead to selection bias in the subsamples.

If I misclassified some advertising as non-advertising YouTubers, the estimates in Table 1 may be too large. YouTubers who fall through the grid of the algorithm seldom permit for ad breaks and do not follow strict commercial incentives. Thus, they are on average more reluctant to adapt their content after Nov 2015 than the average YouTuber whom the algorithm detects. On the other hand, the YouTubers whom I missed might not even increase their share of videos between ten minutes and fourteen minutes. Thus, they are not affected by the instrument $close_i$ and their first stage is equal to zero. In this case, the LATE (see Section 6.2) was the same whether or not I classified some advertising as non-advertising YouTubers.

If some advertising YouTubers were included into the subsample of non-advertising YouTubers, the estimates in Figure A.5 may be too large, too. This would, however, strengthen my results: Appendix C.5 demonstrates that there is no effect of an increase in the feasible number of ad breaks on the non-advertising YouTubers' content choice; if

[108] See, e.g., https://www.youtube.com/watch?v=v6rgvg2CAIc (Nov 2021), although he remains quite vague about these other channels.

[109] E.g., https://www.youtube.com/watch?v=mnvRApKqkUk, https://www.youtube.com/watch?v=I3MeCEwVxB0, https://www.youtube.com/watch?v=1b3W_muI2r8 (Nov 2021).

the estimates were even closer to zero, the validity check would be even more convincing.

Potential consequences of sample migration An advertising YouTuber may have been non-advertising in the past and vice versa. Potential sample migration between advertising and non-advertising YouTubers, however, is unproblematic for three reasons. First, I do not directly compare advertising to non-advertising YouTubers. Second, many advertising YouTubers may have started as non-advertising YouTubers in the beginning of their career. If they became advertising YouTubers as a result of the treatment, they may have adapted their content with a delay, which may lead to an underestimation of the effect of advertising on content differentiation. Finally, if former advertising YouTubers have migrated to the subsample of non-advertising YouTubers, I might overestimate the main effect , which would – as argued in the previous subsection – make the validity check more convincing.

G.3. Per-view price of advertising

The stylized theoretical framework in Appendix A assumes that the price per ad per view r is constant, capturing the idea that YouTubers take the so-called "CPM" (cost per mille) as given. In particular, I assume that r does not depend on the advertising quantity a_i as, e.g., in the seminal theoretical framework by Anderson and Jullien (2016), who presume that the price per ad per view decreases in the advertising quantity a_i (p.46). To the best of my knowledge, there is no such relationship between ad price and ad quantity on YouTube. Thus, YouTubers have no incentive to strategically restrict their advertising quantity to drive up the ad price.

Yet, it is possible that the price per ad per view varies across YouTubers, depending on the audience that they target. In terms of the model, each YouTuber would realize a different r_i (but still take it as given). The idea that advertisers' willingness to pay for ads increases if they can target their desired audience is well established in the literature. E.g., Chandra (2009) shows that newspapers who can better segment readers according to their location and demographics can also set higher advertising prices; Chandra and Kaiser (2014) find analogous results for magazines. Bergemann and Bonatti (2011) develop a theoretical model that shows that the equilibrium price of advertisements is first increasing, then decreasing, in the targeting capacity. The latter effect is driven by diminishing competition between advertisers, though, which is likely to be irrelevant in the YouTube setting.

Concrete and reliable information on YouTube's advertising price scheme is rare, however; e.g., YouTube avoids public disclosure of ad prices and refers potential advertisers to a specialist for further information.[110] Moreover, while YouTube displays plenty of information on the billing procedure, it does not say anything about the typical amount of a bill.[111] Anecdotal evidence suggests that advertisers have the option to target spe-

[110] See `https://www.youtube.com/intl/en/ads/pricing/` (Dez 2021).

[111] See `https://support.google.com/google-ads/topic/3119101?hl=en\&ref_topic=3181080,3126923\&_ga=2.62501036.1431473268.1639129110-1155291345.1639129110` (Dez 2021)).

cific audiences based on gender, age, location, and keywords from the viewers' YouTube searches[112], and some blogs claim that certain YouTube audiences are more expensive to advertise to than others, i.e., targeting may increase advertisers' costs.[113] In sum, however, YouTube's advertising price scheme seems to be rather opaque, both to the advertisers and especially to the YouTubers. Thus, it is plausible to assume that the vast majority of YouTubers takes the price per ad per viewer as given, and does not strategically choose a specific type of video content to attract a specific type of audience to realize a higher per-view price of advertising.

G.4. Viewer switching

The central economic mechanism in my paper is that YouTubers differentiate their content to reduce viewers' propensity to switch to a competing channel when they increase their ad quantity. In this section, I consider such viewer switching in more detail. More specifically, I examine the role of switching costs on YouTube, and I discuss the consequences of temporal versus permanent viewer switching.

Switching costs correspond to the disutility of a consumer if he or she switches products or product providers. Following the seminal typology of Burnham et al. (2003), switching costs comprise (i) *procedural switching costs*, such as search, setup, and learning costs, (ii) *financial switching costs*, such as transaction costs from initiating or terminating a relationship, and (iii) *relational switching costs*, including psychological costs of switching and brand loyalty.[114]

Switching costs on YouTube are likely to be low for three reasons. First, procedural switching costs are minimal: YouTube's recommendation algorithm reduces search costs to a minimum, and setup and learning costs do not exist. Second, neither YouTube nor the YouTubers charge a monetary price for video content, whereby financial switching costs do not accrue.[115] Third – although it is difficult to make informed statements about relational switching costs without access to behavioral viewer data – anecdotal evidence suggests that brand loyalty exists, but that is probably not too widespread. E.g., when YouTubers are involved in "scandals", their subscriber count quickly diminishes (sometimes in the thousands).[116]

Viewer switching, potentially facilitated through low switching costs, seems to be a common phenomenon on YouTube. Likewise, YouTubers and blogs expound various reasons leading to the loss of viewers, including infrequent uploads, stale or inconsistent content,

[112]See, e.g., https://www.creatopy.com/blog/youtube-ads-cost/, https://thriveagency.com/news/how-much-do-youtube-ads-cost/ (Dec 2021).

[113]See, e.g., https://ppcprotect.com/blog/display-ads/youtube-advertising-cost/ or https://www.digitalmarketing.org/blog/how-much-does-youtube-advertising-cost (Nov 2021).

[114]See Klemperer (1995) for an alternative, though similar, classification.

[115]As argued in Appendix G.5, the premium service YouTube Red was only later introduced in Germany.

[116]Logan Paul, for instance, lost several thousand subscribers after uploading a video showing a suicide victim; see, e.g., https://www.youtube.com/watch?v=CUI_PGpE_ls or https://nextshark.com/logan-paul-losing-subscribers-second-posting-japanese-suicide-video/ (Dez 2021).

competition, and over-promotion (i.e., too much advertising).[117] Many YouTubers recognize the trade-off between ad revenue and viewer alienation, stating, e.g., that "while most users understand that ads are how their favourite creators make money, most people also have a point where it becomes too much, and that can be a reason to leave for many"[118], that "noone came to watch your ads"[119], and that "viewers may unsubscribe to watch a better channel".[120]

In the context of my analysis, potential viewer switching as a consequence of advertising is crucial. More specifically, YouTubers risk viewer switching for two reasons. First, as argued throughout the paper, an increase in advertising quantity corresponds to a price raise for viewers, whereby it becomes worth to switch from one YouTuber to a competitor despite the existence of switching costs. When switching costs are low, this effect is likely to be accelerated. Second, it is possible that an increase in advertising quantity additionally decreases psychological switching costs and dilutes viewers' brand loyalty, it is exists. It is often argued that many YouTubers' success is based on their authenticity and a feeling of community (e.g., Tolson, 2010; Cunningham and Craig, 2017). If YouTubers raise their ad quantity, they may appear as overly commercial, thus destroying their authenticity and weakening the bond between viewer and YouTuber. These two channels are not mutually exclusive, and without access to behavioral viewer data, it is difficult to quantify their relative importance. Note, however, that a potential reduction in viewer switching costs is not a necessary requirement for my results, but likely to support the plausibility of the central mechanism in the paper, because YouTubers have an additional incentive to avoid close competition.

Whether viewers switch permanently or just temporarily to a competitor plays a large, but no decisive role. Naturally, the effect of permanent switching is way more destructive, as it erodes a YouTuber's subscriber base and thus impedes the channel's monetization potential. As argued above, it is plausible to assume that viewers switch permanently: when they are overly annoyed by the advertising quantity on a particular YouTube channel, they invest some (minimal) search costs once to find a suitable alternative, and lack a reasonable incentive to return. Such permanent switching may, but need not necessarily, coincide with unsubscription; e.g., many regular viewers never subscribe to a channel and may just stop watching, and some subscribers may never unsubscribe (since subscriptions are costless), but permanently switch to a competitor nevertheless. Temporal switching, in contrast, seems less likely, and its consequences would be less severe. If viewers' brand loyalty was strong, for instance, they might just reduce the consumption of video content from a particular YouTuber and partially substitute for it with content from a close competitor.

[117]See, e.g., `https://alanspicer.com/why-am-i-losing-subscribers/`, `https://vidiq.com/blog/post/losing-youtube-subscribers/`, or `https://www.youtube.com/watch?v=45uzksHgDdQ` (Dez 2021).

[118]See `https://alanspicer.com/why-am-i-losing-subscribers/` (Dez 2021).

[119]See `https://www.youtube.com/watch?v=45uzksHgDdQ` (Dez 2021).

[120]See `https://vidiq.com/blog/post/losing-youtube-subscribers/` (Dez 2021).

G.5. Platform events during the observation period

Next, I provide a systematic review of all platform "events" during my observation period, i.e., technical novelties or changes in YouTube's monetization policy beyond the launch of the new ad break tool. Note that an event can only affect my results if it is correlated to a YouTuber's probability to upload mainstream content *and* to her value of $close_i$ – no such event exists during the observation period. Since YouTube has no serious competitors, I remain agnostic about events at competing video sharing platforms.

G.5.1. Data collection

I collect information on all events from the YouTube Creators Blog, which announces YouTube news, introduces technical features, and gives general advice to YouTubers.[121] In a first step, I retrieve all blog posts from Jan 2013 to Jan 2017. Next, I manually exclude any post that does not deal with a platform event, such as YouTube promotion for academies, awards, (real world) events, and YouTuber portraits. The remaining 42 posts are listed in Table A.17. In a last step, I review all posts from Table A.17 and indicate if a YouTuber's monetization options or her probability to upload mainstream content could be affected. Thirteen events require further investigation; I discuss them chronologically.

G.5.2. Platform events in 2013

First, in March, YouTubers' access to their financial data changed. This event applies to all YouTubers equivalently, has no effect on their content choice, and is therefore unproblematic.

In May, selected YouTubers from the U.S., and in October, selected YouTubers worldwide were given the option to raise a subscription fee of 0.99$ per month. The pilot was, however, extremely limited: not even 100 YouTubers worldwide participated.[122] Thus, my results are unlikely to be affected by these events.

Next, YouTube launched its "Fan Finder": a YouTuber could let the platform turn one of her videos into an "ad" and show it to viewers of a different channel in place of a conventional ad; this was supposed to enlarge a YouTuber's fan base. Since YouTubers were asked to produce special videos that advertise their channel, the event may have affected their content choice. Yet, all YouTubers with at least $1,000$ subscribers could participate and there were no restrictions on the advertising video's duration. Hence, the event is not correlated to $close_i$ and thereby unproblematic.

Finally, live streams became technically feasible in December and may have influenced YouTuber's content choice. The feature is open to all YouTubers, though. Hence, the

[121] See `youtube-creators.googleblog.com/` (May 2019).

[122] E.g., `www.fastcompany.com/3020553/the-most-popular-youtube-channels-might-start-charging-you-to-watch`, `www.bbc.com/news/business-22474715`, or `searchenginewatch.com/sew/news/2267170/youtube-launches-paid-channels-subscription-fees-start-at-usd099-per-month` (May 2019).

event is not correlated to $close_i$ and cannot affect my results.

G.5.3. Platform events in 2015

In March, 360 degree videos became technically feasible. Similar to the live streams, the event may have influenced YouTubers' content choice, but since it is open to all YouTubers, there is no correlation to $close_i$.

YouTube Red, a paid subscription service that provides advertising-free streaming of all videos and exclusive original content was launched in October. The availability of YouTube Red is, however, limited to the US. Since my dataset includes only German YouTube channels, the event cannot affect my results.

In November, several virtual reality tools became available. Again, YouTubers' content choice may have been affected, but since the features are open to all YouTubers, there is no correlation to $close_i$.

G.5.4. Platform events in 2016

In January, YouTube launched a "Donate Button": users who click on the button can donate to a YouTuber after watching her video. As with the technical novelties from above, this may have influenced YouTubers' content choice. In addition, their monetization options were affected. Still, the feature is open to all YouTubers and thereby not correlated to $close_i$.

Next, in April, YouTube announced that it would withhold (not block) all ad revenue generated during copyright disputes. This event applies to all YouTubers equivalently, has no effect on their content choice, and is therefore unproblematic.

Mobile live streams became technically feasible in June, i.e., YouTubers could stream from their mobile devices. Similar to the "stationary" live streams from 2013, the event is not correlated to $close_i$ and cannot affect my results.

In October, YouTube launched an optional feature for paid promotion disclosure: by checking the "video contains paid promotion" box in their settings, YouTubers can inform their audience about paid product placement and endorsements by third parties. This may influence their videos' content, but is unrelated to $close_i$.

Finally, in October, video end screens, that allow YouTubers to promote up to four different videos or playlists, became technically available. Although the event may have affected the YouTubers' content choice, the feature is open to all YouTubers, thereby not correlated to $close_i$, and hence unproblematic.

G.6. YouTuber learning effect

Here, I discuss a YouTuber learning effect as an alternative explanation for the results from Section 7: YouTubers copy the most mainstream content in the beginning of their career, but deviate from the mainstream when they become more experienced and start to develop a personal style. If such a learning effect was positively correlated to $close_i$, it

Table A.17: YouTube platform events

	Date		Summary of the event	Monetization	Content choice
1	2013	Jan	The channel view count only includes views from publicly available videos from now.		
2	2013	Feb	It is now technically feasible to update several video updates at the same time.		
3	2013	Mar	YouTube changes the interaction with AdSense: a YouTuber's financial overview is now available at YouTube Analytics	X	
4	2013	Mar	The new channel design "YouTube One" is available for all YouTubers.		
5	2013	Apr	Users see more videos in their homepage feed.		
6	2013	May	YouTubers receive an e-mail once a video upload has finished.		
7	2013	May	The new channel design "YouTube One" is mandatory for all YouTubers.		
8	2013	May	Selected YouTubers from the US may raise a subscription fee of 0.99$ per month.	X	X
9	2013	June	Mobile users (Android and iOS) may follow links embedded into videos from now.		
10	2013	July	YouTubers may now connect multiple channels via a Google+ page.		
11	2013	Aug	Improved mobile features for users.		
12	2013	Sept	Launch of the YouTube Audio Library (150 royalty-free tracks).		
13	2013	Sept	Improved tools for moderating comments.		
14	2013	Sept	New tools to identify and interact with one's top viewers.		
15	2013	Sept	YouTubers may now feature playlists from other channels.		
16	2013	Oct	Selected YouTubers from outside the US may also raise a subscription fee of 0.99$ per month.	X	X
17	2013	Nov	A YouTuber may let the platform turn her video into an ad that is then shown to viewers from different channels.		X
18	2013	Dec	Live streams are now technically feasible.		X
19	2014	Feb	YouTube validates a video's view count repeatedly from now on.		
20	2014	Feb	Users can create their own playlists.		
21	2014	Apr	Enhanced playlist tools in YouTube Analytics are launched.		
22	2014	June	New messaging and commenting features for YouTubers.		
23	2014	June	YouTube removes blocked users from a channel's subscriber count.		
24	2014	Nov	New YouTube homepage for music videos.		
25	2015	Mar	360 degree videos are now technically feasible.		X
26	2015	May	60fps for live streams is now technically feasible.		
27	2015	June	New data tool Music Insights is available: shows the cities where an artist is most popular, top tracks by artist, and views from both artists' official music videos and fan uploads claimed using Content ID.		
28	2015	July	A new design for YouTube mobile app is launched.		
29	2015	Oct	YouTube Red is launched in the US.	X	X
30	2015	Nov	New language and translation tools are available.		
31	2015	Nov	New virtual reality tools are available.		X
32	2016	Jan	Users can donate to the YouTuber after watching a video.	X	X
33	2016	Feb	A new blurring tool (to blur faces etc.) is available.		
34	2016	Apr	YouTube withholds any ad revenue generated during content ID disputes from now.	X	
35	2016	June	Mobile live streams are now technically feasible.		X
36	2016	Sept	YouTube Analytics becomes easier to understand for YouTubers.		
37	2016	Sept	New tools for YouTubers to engage with their community.		
38	2016	Oct	An optional feature for paid promotion disclosure is available.	X	X
39	2016	Oct	Special video end screens are available.		X
40	2016	Nov	New comment features are available for users.		
41	2016	Dec	Launch of a new URL system that is independent from Google+.		
42	2017	Jan	User messages in a chat stream may be highlighted.		

Notes: Summary of YouTube platform events during my observation period Jan 2013 to Jan 20217.

could be the driving force behind the decrease in the probability to upload mainstream content after Nov 2015 rather than an increase in the feasible number of ad breaks per video.

Three arguments, however, speak against a YouTuber learning effect. First, there exists no plausible reason why YouTubers with a high value of $close_i$ would experience a stronger learning effect than YouTubers whose value of $close_i$ is low. See Section 6.2.2 for a detailed discussion on the independence of $close_i$.

Second, t_{it} controls for a YouTuber's average change in the probability to upload mainstream (or competitive) content over time. Columns 1, 4, and 7 in Table A.18 replicate the 2SLS results from Tables 1 and 6, respectively, and illustrate that a linear YouTuber learning effect is of minor importance: the estimates for t_{it}, though negative, are extremely small.

Third, allowing for a more flexible YouTuber specific time trend by adding t_{it}^2 and t_{it}^3 does not affect the estimates of interest, either. However, it becomes obvious that the YouTuber specific time trend is not linear. For instance, columns 2, 5, and 8 illustrate that a YouTuber's probability to upload mainstream or competitive content increases in the beginning, but decreases from around her 160^{th} video, which is consistent with the story from above. Note that the average number of videos per YouTuber is 99.3 and the median number of videos is 64. Thus, many YouTubers in my sample do not reach the turning point of 160. In sum, even though I find small evidence for a YouTuber learning effect, it is definitely not the driving force behind my main results.

Table A.18: Learning

	Mainstream			log(competitive pressure)			log(w. competitive pressure)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$D_i * post_t$	-.226***	-.233***	-.229***	-.693***	-.742***	-.722***	-1.52***	-1.56**	-1.54***
	(.048)	(.049)	(.048)	(.194)	(.195)	(.194)	(.263)	(.267)	(.264)
t_{it}	-.000**	.000***	.001***	-.001***	.002***	.004***	-.001***	.002***	.004***
	(-.000)	(.000)	(.000)	(.000)	(.000)	(.001)	(.000)	(.000)	(.001)
t_{it}^2		-9.79e-07***	-2.93e-06***		-7.56e-06***	-.000***		-6.92e-06***	-.000***
		(2.56e-07)	(9.04e-07)		(1.10e-06)	(3.81e-06)		(1.31e-06)	(4.72e-06)
t_{it}^3			3.51e-09**			1.86e-08**			2.49e-08***
			(1.71e-09)			(7.46e-09)			(8.85e-09)
First stage	.029***	.029***	.029***	.029***	.029***	.029***	.029***	.029***	.029***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
F-statistic	150.65	148.10	149.67	150.86	148.40	149.95	150.29	147.88	149.37
Time FE	X	X	X	X	X	X	X	X	X
YouTuber FE	X	X	X	X	X	X	X	X	X
Category FE	X	X	X	X	X	X	X	X	X
Category Time Trend	X	X	X	X	X	X	X	X	X
YouTubers	10,599	10,599	10,599	10,597	10,597	10,597	10,591	10,591	10,591
Videos	1,067,542	1,067,542	1,067,542	1,062,993	1,062,993	1,062,993	1,057,360	1,057,360	1,057,360

Notes: Robust standard errors in parentheses. In columns 1 to 3, the dependent variable is a dummy variable equal to 1 if video v of YouTuber i uploaded in month t is equipped with a mainstream tag, and 0 otherwise. The dependent variable in columns 4 to 6 is the log number of competitors who also use one of the tags of video v in a given month t. The dependent variable in columns 7 to 9 is the log number of competitors, weighted by their respective number od subscribers, who also use one of the tags of video v in a given month t. All estimates are 2SLS estimates. The estimates are based on using the advertising YouTubers only. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

H. Omitted figures and tables

Figures and tables that were omitted from the main part of the paper.

H.1. Omitted figures

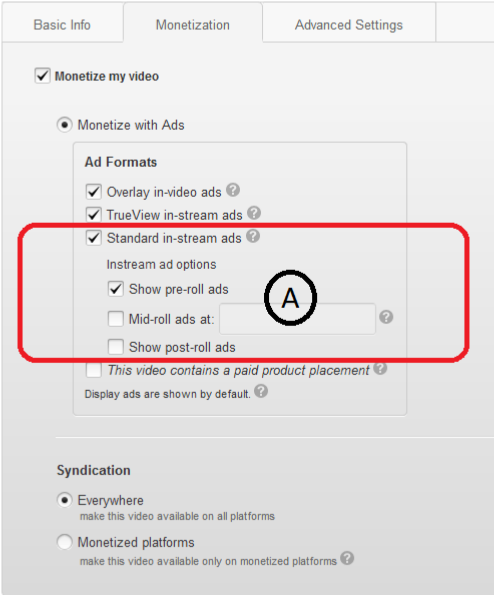

Figure A.18: Old ad break tool (before Nov 2015).

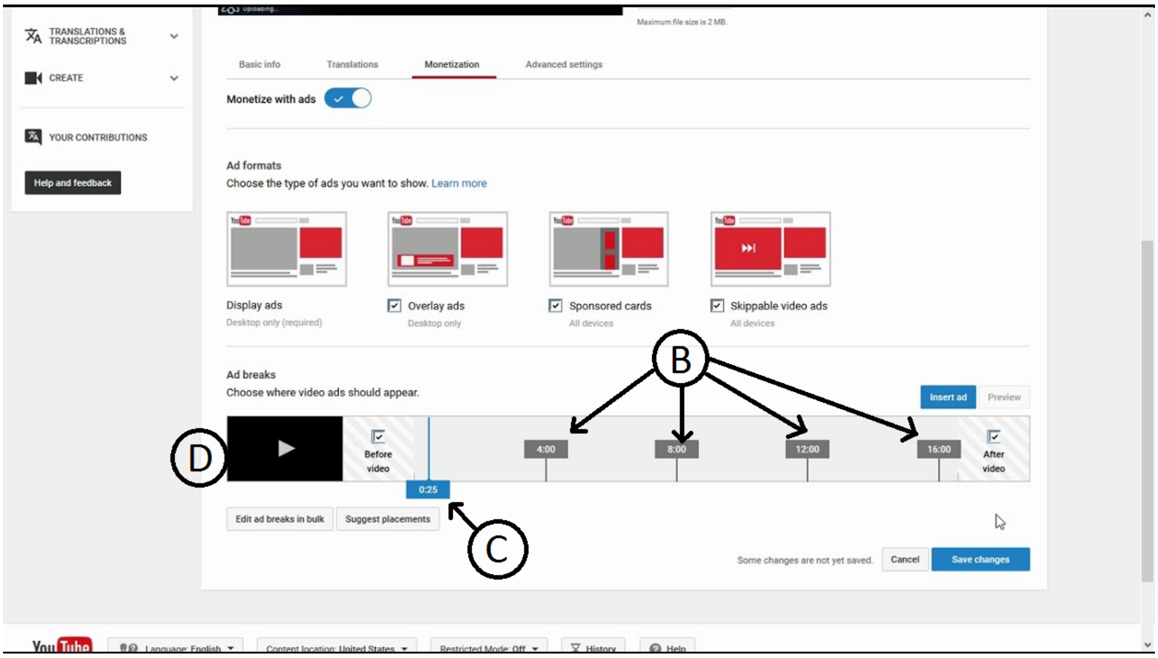

Figure A.19: New ad break tool (after Nov 2015).

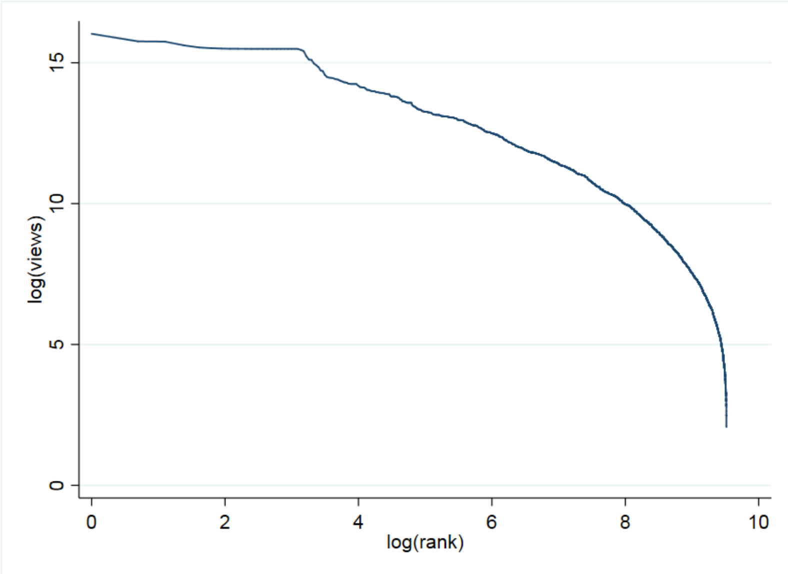

Figure A.20: Log-log plot of the number of views a certain tag attracts and its associated rank in the category "Science & Technology" in April 2015.

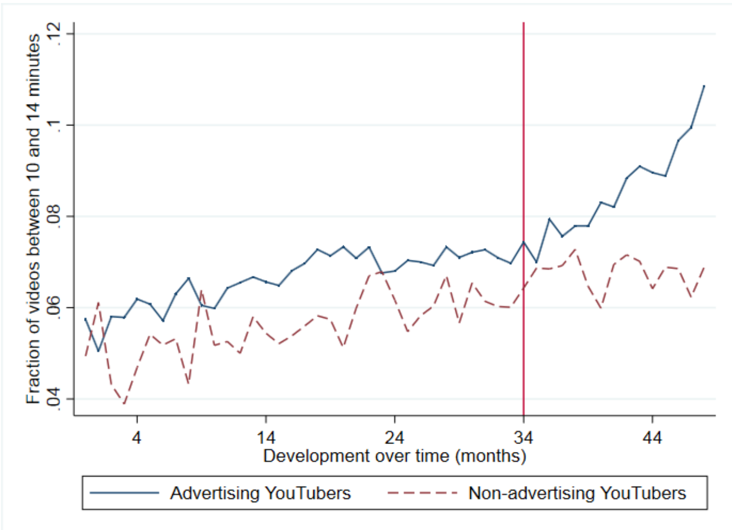

Figure A.21: Development of the fraction of videos between ten and fourteen minutes for advertising and non-advertising YouTubers. The vertical line depicts Nov 2015, where the new ad break tool was launched.

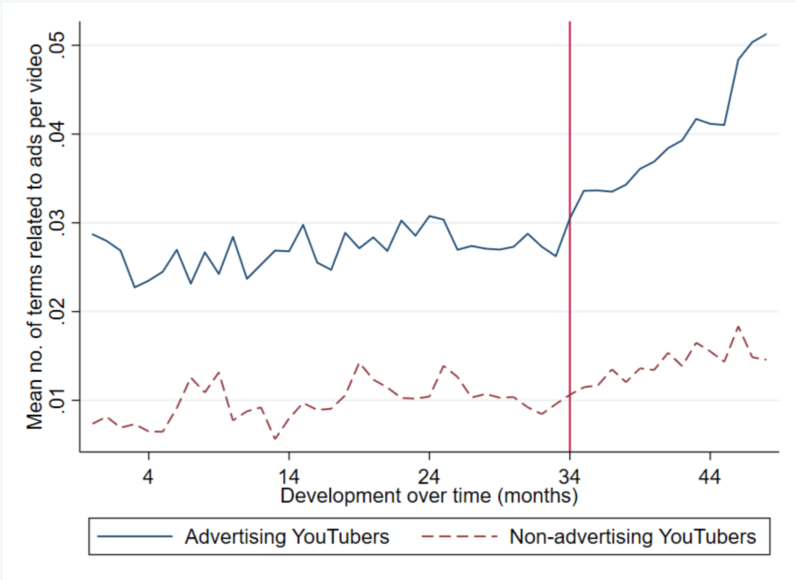

Figure A.22: Development of the average number of terms related to advertising in the comment sections of advertising and non-advertising YouTubers over time. The vertical line depicts Nov 2015, where the new ad break tool was launched.

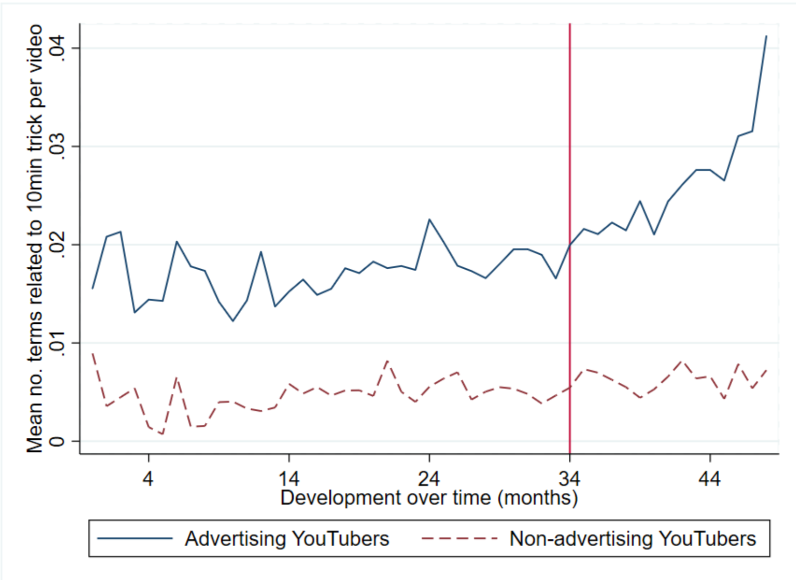

Figure A.23: Development of the average number of terms related to the ten minutes trick in the comment sections of advertising and non-advertising YouTubers over time. The vertical line depicts Nov 2015, where the new ad break tool was launched.

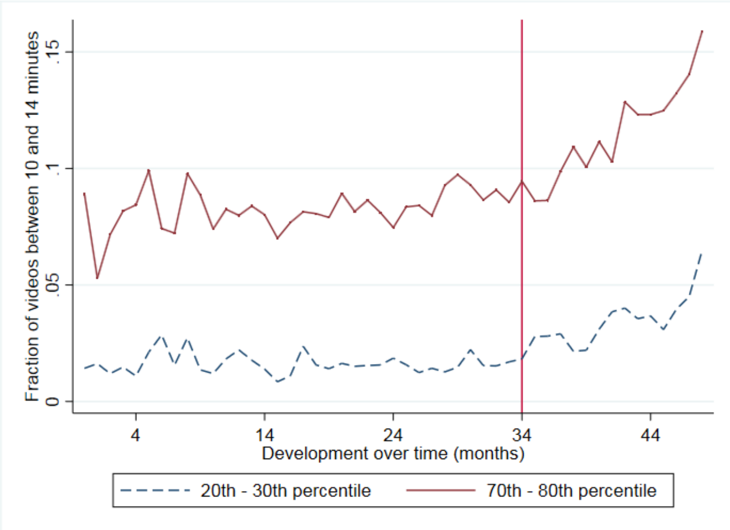

Figure A.24: Development of the fraction of videos between ten and fourteen minutes for advertising YouTubers close to and further away from the ten minutes threshold before Nov 2015. The vertical line depicts Nov 2015, where the new ad break tool was launched.

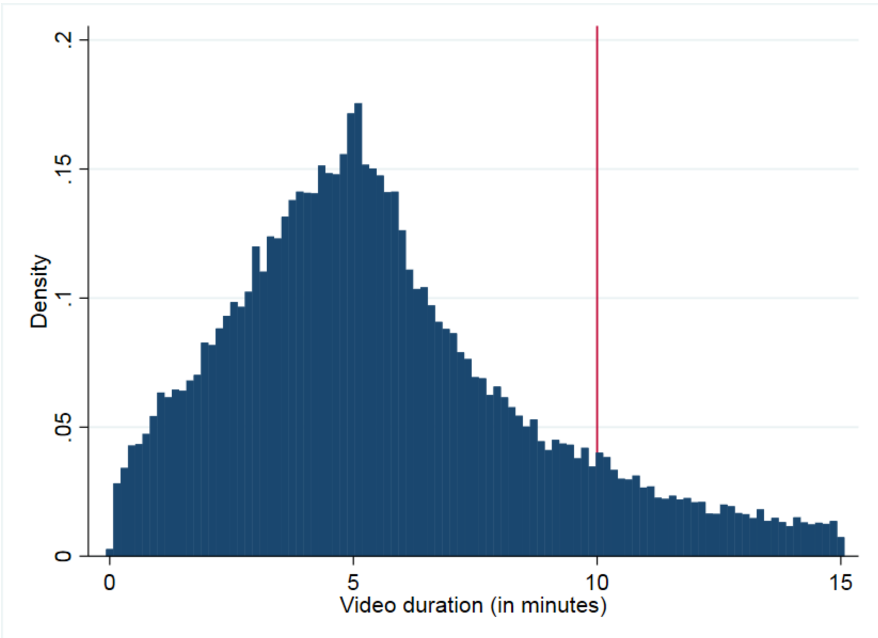

Figure A.25: Histogram of the distribution of video durations *before* Nov 2015 for YouTubers around the 75th percentile of "closeness" to the ten minutes threshold. The vertical line depicts the ten minutes threshold.

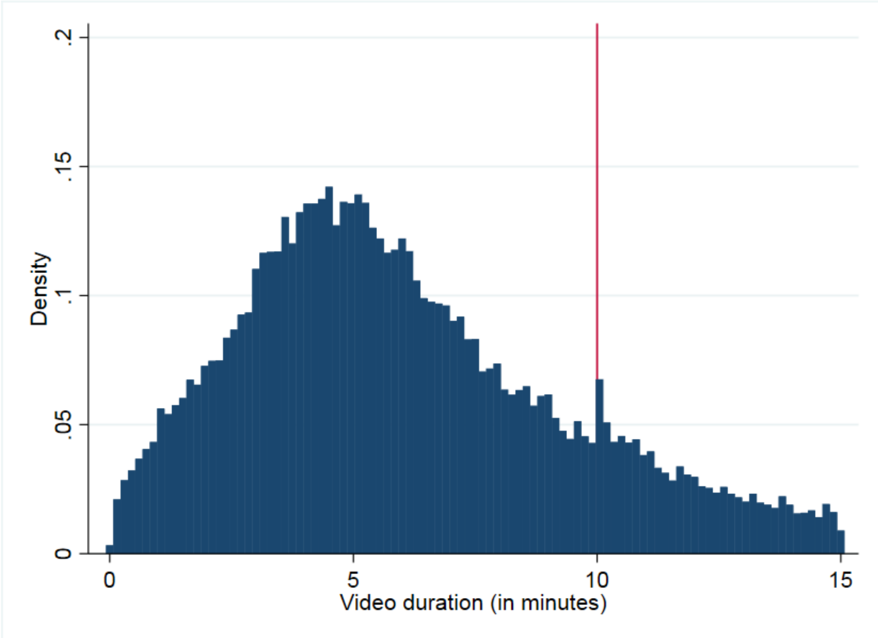

Figure A.26: Histogram of the distribution of video durations *after* Nov 2015 for YouTubers around the 75th percentile of "closeness" to the ten minutes threshold. The vertical line depicts the ten minutes threshold.

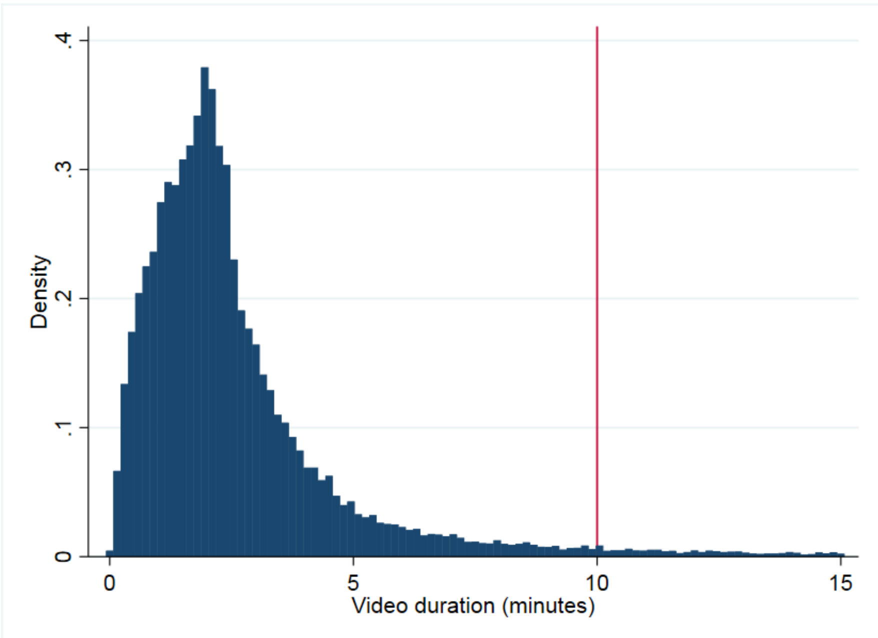

Figure A.27: Histogram of the distribution of video durations *before* Nov 2015 for YouTubers around the 25th percentile of "closeness" to the ten minutes threshold. The vertical line depicts the ten minutes threshold.

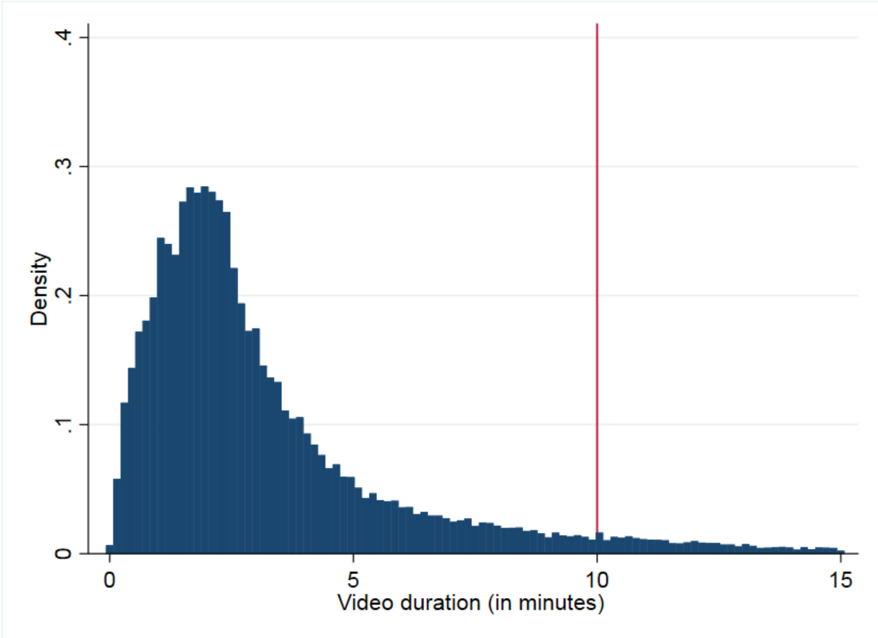

Figure A.28: Histogram of the distribution of video durations *after* Nov 2015 for YouTubers around the 25th percentile of “closeness” to the ten minutes threshold. The vertical line depicts the ten minutes threshold.

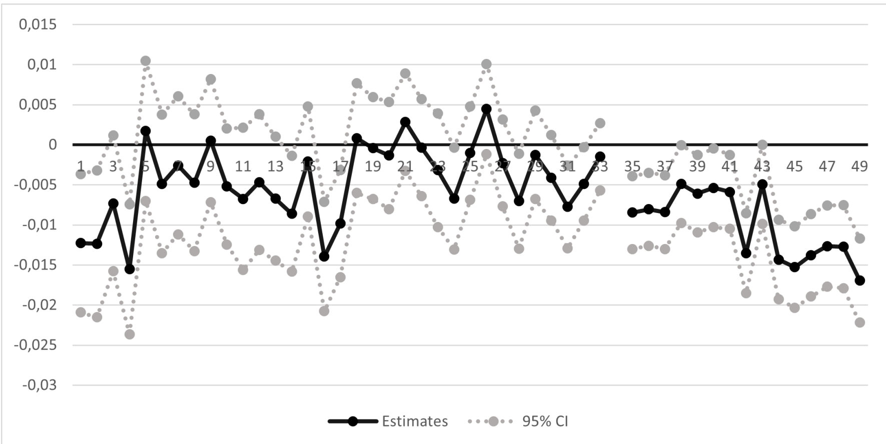

Figure A.29: Event study: The solid line displays the estimates for γ_t, the dashed lines depict a 95% confidence interval. The estimates are based on an OLS regression of equation (4) including YouTuber, time, and category fixed effects, as well as linear YouTuber and category time trends. Standard errors are clustered on the YouTuber-level.

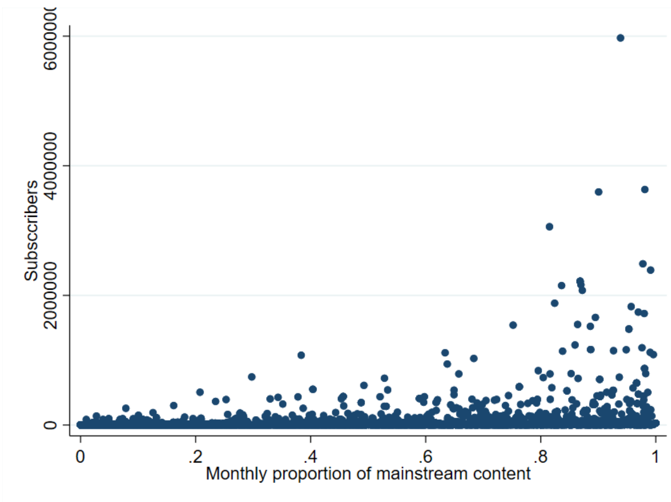

Figure A.30: Scatterplot of subscribers and average monthly proportion of mainstream content before Nov 2015.

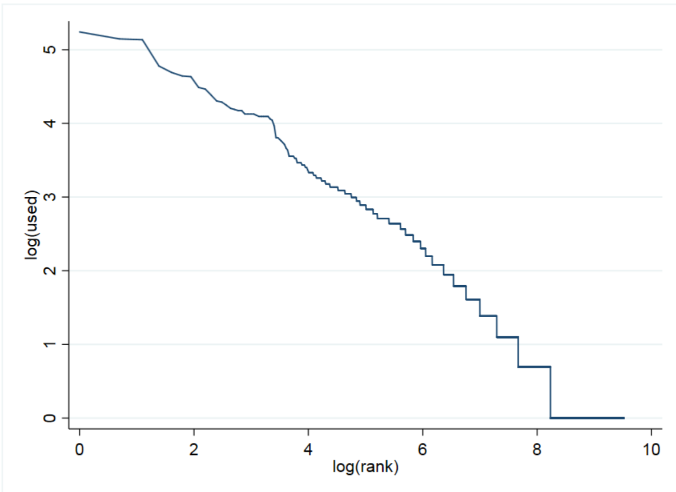

Figure A.31: Log-log plot of the number of videos using a certain tag and its associated rank in the category “Science & Technology” in April 2015.

H.2. Omitted tables

Table A.19: Summary statistics – Main paper

Variable	Mean	Std. Dev.	Min.	Max.	N
Main paper					
$Mainstream_{vit}$	0.425	0.494	0	1	1,397,267
$Advertising_i$	0.668	0.471	0	1	15,877
$post_t$	0.475	0.499	0	1	1,397,267
D_i	0.226	0.418	0	1	15,877
$close_i$	3.718	2.440	0	9.992	15,877
$Duration_{vit}$	6.411	13.341	0	1440.033	1,397,267
$Subscribers_i$	18,234.506	138,282.229	0	6,581,640	15,877
$prop.\ mainstream_i$	0.341	0.323	0	1	15,877
$Film\&Animation_{vit}$	0.086	0.280	0	1	1,397,267
$Cars\&Vehicles_{vit}$	0.081	0.272	0	1	1,397,267
$Music_{vit}$	0.025	0.155	0	1	1,397,267
$Pets\&Animals_{vit}$	0.026	0.159	0	1	1,397,267
$Sports_{vit}$	0.085	0.278	0	1	1,397,267
$Travel\&Events_{vit}$	0.056	0.229	0	1	1,397,267
$Let'sPlay_{vit}$	0.085	0.278	0	1	1,397,267
$People\&Blogs_{vit}$	0.202	0.402	0	1	1,397,267
$Comedy_{vit}$	0.015	0.121	0	1	1,397,267
$Entertainment_{vit}$	0.201	0.401	0	1	1,397,267
$HowTo\&Style_{vit}$	0.064	0.245	0	1	1,397,267
$Education_{vit}$	0.046	0.210	0	1	1,397,267
$Science\&Technology_{vit}$	0.014	0.119	0	1	1,397,267
$Nonprofit\&Activism_{vit}$	0.015	0.120	0	1	1,397,267
$Uploads_{it}$	4.337	5.269	1	249	322,200
$MainstreamMonth_{it}$	1.841	3.603	0	249	322,200
$No.Tags_{vit}$	15.623	8.704	1	67	1,067,542
$UniqueTags_{it}$	34.602	33.376	1	1,059	241,905
$log(competitors)$	4.755	1.935	0	11.111	1,391,210
$log(w.competitors)$	13.734	3.292	0	21.880	1,383,024
$prop.\ pos.\ ratings_{vit}$	0.904	0.169	0	1	1,277,678
$sentiment\ score_{vit}$	0.005	0.021	-1	1	1,397,267
has_intro_{vit}	0.441	0.497	0	1	2,424
has_outro_{vit}	0.363	0.481	0	1	2,424
$visual\_quality_{vit}$	3.248	0.778	1	5	2,423
$sound\_quality_{vit}$	3.36	0.837	1	5	2,423
$overall\_impression_{vit}$	2.553	0.807	1	5	2,418
Ads_{vit}	0.466	0.832	0	52	52,462

Notes: This table presents the summary statistics of all variables used in the main part of the paper.

Table A.20: Summary statistics – Appendix

Variable	Mean	Std. Dev.	Min.	Max.	N
Appendix					
$Titlewords_{vit}$	0.513	0.5	0	1	1,397,267
$SplitCompoundTags_{vit}$	0.574	0.495	0	1	1,397,267
$Mainstream\_t-1_{vit}$	0.406	0.491	0	1	1,387,840
$Superstars25_{vit}$	0.41	0.492	0	1	1,397,267
$Superstars50_{vit}$	0.487	0.5	0	1	1,397,267
$MainstreamLDA_{vit}$	0.820	0.385	0	1	1,397,267
$log(Views)_{vit}$	6.867	2.429	0	19.638	1,396,299
$CorrectSet_{vit}$	0.825	0.248	0	1	2,424
$German\_langid.py_{vit}$	0.589	0.492	0	1	1,397,267
$German\_langdetect.py_{vit}$	0.531	0.499	0	1	1,397,267
$German\_survey_{vit}$	0.744	0.436	0	1	2,424
$English\_langid.py_{vit}$	0.253	0.435	0	1	1,397,267
$English\_langdetect.py_{vit}$	0.227	0.419	0	1	1,397,267
$English\_survey_{vit}$	0.106	0.308	0	1	2,424
$Other\_langid.py_{vit}$	0.158	0.365	0	1	1,397,267
$Other\_langdetect.py_{vit}$	0.242	0.428	0	1	1,397,267
$Other\_survey_{vit}$	0.15	0.357	0	1	2,424
$I(1^{st}to10^{th})_{vit}$	0.673	0.469	0	1	1,397,267
$I(10^{th}to25^{th})_{vit}$	0.581	0.493	0	1	1,397,267
$I(25^{th}to50^{th})_{vit}$	0.548	0.498	0	1	1,397,267
$I(50^{th}to75^{th})_{vit}$	0.390	0.488	0	1	1,397,267
$I(75^{th}to100^{th})_{vit}$	0.284	0.451	0	1	1,397,267
$SumAffiliations_{vit}$	2.472	1.160	0	5	1,397,267
$Competitive_{vit}$	0.641	0.48	0	1	1,397,267

Notes: This table presents the summary statistics of all variables used in the Appendix.

Table A.21: Most common mainstream tags per category

Category	**Top 5 tags**
Film & Animation	'german' (49), 'trailer' (48), 'german' (48), 'cartoon' (47), 'style' (40)
Cars & Vehicles	'tuning' (48), 'test' (47), 'sound' (46), 'bmw' (46), 'vw' (43)
Music	'banger music' (33), 'farid bang' (28), 'official' (26), 'rap' (23), 'hd' (23)
Pets & Animals	'horses' (41), 'horse riding' (38), 'riding' (28), 'pony' (26), 'horses' (26)
Sports	'fitness' (49), 'training' (48), 'soccer' (46), 'football' (46), 'bodybuilding' (46)
Travel & Events	'vacation' (37), 'vlog' (30), 'germany' (25), 'holiday' (24), 'dner' (24)
Let's Play	'walkthrough' (49), 'tutorial' (49), 'playthrough' (49), 'minecraft' (49), 'lets' (49)
People & Blogs	'rick' (49), 'radio' (49), 'tutorial' (47), 'steve' (47), 'german' (46)
Comedy	'freshtorge' (37), 'freshhaltefolie' (36), 'torge' (30), 'sandra' (22), 'funny' (21)
Entertainment	'comedy' (47), 'video' (46), 'music' (46), 'music' (46), 'rap' (45)
How To & Style	'tutorial' (49), 'diy' (49), 'beauty' (49), 'fashion' (48), 'instructions' (48)
Education	'topten' (49), 'top ten' (49), 'top 10' (49), 'most important' (49), 'worst' (49)
Science & Technology	'tutorial' (49), 'test' (49), 'review' (49), 'german' (49), 'german' (48)
Nonprofit & Activism	'islam' (31), 'jesus' (29), 'god' (26), 'bible' (25), 'religion' (21)

Notes: Table A.21 shows the top five most-common mainstream tags per category by number of classifications as *mainstream* (in brackets) over all 49 observation periods. The tags are translated into English. The tag 'german' is both in German and in English among the top 5 tags in the categories 'Film & Animation' and 'Science & Technology' and hence appears twice; the same is true for the tag 'music' in the category 'Music' and 'horses' in the category 'Pets & Animals'. 'freshtorge' (and deducted terms) is the alias of a well-known German comedy YouTuber.

Table A.22: Correlations over time

Category	$t-1$	$t-2$	$t-3$
Film & Animation	.37	.33	.32
Cars & Vehicles	.30	.31	.27
Music	.25	.20	.17
Pets & Animals	.30	.26	.25
Sports	.27	.26	.23
Travel & Events	.18	.12	.12
Let's Play	.44	.37	.34
People & Blogs	.30	.27	.25
Comedy	.30	.24	.21
Entertainment	.34	.30	.28
How to & Style	.41	.38	.36
Education	.35	.31	.29
Science & Technology	.25	.22	.21
Nonprofit & Activism	.17	.17	.15

Notes: Table A.22 displays the average fraction of mainstream tags in month t that overlaps with the mainstream tags in months $t-1$, $t-2$, and $t-3$ for each video category.

Table A.23: Correlations cross

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
(1) Film & Animation													
(2) Cars & Vehicles	.025												
(3) Music	.023	.022											
(4) Pets & Animals	.022	.021	.049										
(5) Sports	.025	.025	.054	.069									
(6) Travel & Events	.024	.024	.053	.067	.024								
(7) Let's Play	.026	.027	.056	.070	.026	.032							
(8) People & Blogs	.025	.026	.055	.069	.025	.031	.010						
(9) Comedy	.020	.020	.047	.060	.020	.025	.007	.012					
(10) Entertainment	.026	.026	.056	.069	.026	.032	.010	.016	.016				
(11) How to & Style	.024	.024	.053	.067	.023	.030	.009	.015	.094	.010			
(12) Education	.024	.024	.053	.066	.024	.030	.009	.015	.094	.010	.025		
(13) Science & Techn.	.025	.026	.054	.068	.025	.031	.010	.016	.094	.011	.026	.030	
(14) Nonprofit & Activism	.021	.022	.048	.063	.022	.027	.008	.013	.088	.009	.023	.026	.019

Notes: Table A.23 displays the average fraction of mainstream tags in month t that overlaps with the mainstream tags in month in each of the other video categories.

Table A.24: Competition by the proportion of mainstream content before Nov 2015

	log(competitive pressure)		log(w. competitive pressure)	
	below median	above median	below median	above median
	(1)	(2)	(3)	(4)
$D_i * post_t$	-.937***	-.277	-1.424***	-.677**
	(.319)	(.263)	(.432)	(.313)
First stage	.028***	.027***	.028***	.027***
	(.002)	(.003)	(.003)	(.003)
F-statistic	68.01	68.59	67.62	68.08
Time FE	X	X	X	X
YouTuber FE	X	X	X	X
Category FE	X	X	X	X
Category Time Trend	X	X	X	X
YouTuber Time Trend	X	X	X	X
YouTubers	5,293	5,304	5,287	5,304
Videos	498,634	564,359	493,973	563,387

Notes: Robust standard errors in parentheses. The dependent variable in columns 1 and 2 is the log number of competitors who also use one of the tags of video v in a given month t. The dependent variable in columns 3 and 4 is the log number of competitors, weighted by their respective number od subscribers, who also use one of the tags of video v in a given month t. All estimates are 2SLS estimates. The estimates are based on using the advertising YouTubers only. Columns 1 and 2 consider only YouTubers below the median monthly proportion of mainstream content before Nov 2015 (29.6%), columns 3 and 4 consider only YouTubers above and equal to the median monthly proportion of mainstream content before Nov 2015. Standard errors are clustered on the YouTuber level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.25: Correlation between measures for video quality

	prop. pos. ratings	*sentiment score*
	(1)	(2)
has intro	0.03	0.06
has outro	0.05	0.06
visual quality	0.04	0.03
sound quality	0.06	0.01
overall impression	0.03	0.02
N	2,331	2,418

Notes: Column 1 displays the pairwise correlations between the proportion of positive ratings and the measures from the online survey experiment. Column 2 displays the pairwise correlations between the sentiment score of a video and the measures from the online survey experiment.

Table A.26: Video quality and actual number of ad breaks

	Panel A: *prop. pos. ratings*					
	before Nov 2015			after Nov 2015		
	(1)	(2)	(3)	(4)	(5)	(6)
Ads_{vit}	.004	.003	.003	-.008***	-.009***	-.009***
	(.003)	(.003)	(.003)	(.002)	(.002)	(.002)
has_intro_{vit}		-.020*	-.021*		.005	-.000
		(.012)	(.012)		(.009)	(0.009)
has_outro_{vit}		.040***	.039***		.009	.007
		(.012)	(.012)		(.009)	(.009)
$visual\_quality_{vit}$			.004			-.004
			(.010)			(.008)
$sound\_quality_{vit}$			.001			.017**
			(.009)			(.007)
$overall\_impression_{vit}$			.001			-0.004
			(.007)			(.006)
Time FE	X	X	X	X	X	X
Category FE	X	X	X	X	X	X
Videos	1,152	1,152	1,149	1,182	1,182	1,179
	Panel B: *sentiment score*					
	before Nov 2015			after Nov 2015		
	(1)	(2)	(3)	(4)	(5)	(6)
Ads_{vit}	-.000	-.000	-.000	-.000	-.000	-.000
	(.000)	(.000)	(.000)	(.000)	(.000)	(.000)
has_intro_{vit}		.001	.001		.002*	.002*
		(.002)	(.002)		(.001)	(.001)
has_outro_{vit}		.002	.002		.001	.001
		(.002)	(.002)		(.001)	(.001)
$visual\_quality_{vit}$			.002			-.001
			(.002)			(.001)
$sound\_quality_{vit}$			-.002			.001
			(.001)			(.001)
$overall\_impression_{vit}$			.001			-.000
			(.001)			(.001)
Time FE	X	X	X	X	X	X
Category FE	X	X	X	X	X	X
Videos	1,204	1,204	1,201	1,217	1,217	1,214

Notes: Robust standard errors in parentheses. The dependent variable in Panel A is the proportion of positive ratings of video v by YouTuber i in month t as defined by expression (5). The dependent variable in Panel B is the sentiment score of video v by YouTuber i in month t as defined by expression (6). Ads_{vit} corresponds to the actual number of ad breaks per video. has_intro_{vit} and has_outro_{vit} are dummy variables that indicate if a video has a customized intro or outro sequence, respectively. $visual\_quality_{vit}$, $sound\_quality_{vit}$, and $overall\_impression_{vit}$ measure video quality on five-point Likert-scales, respectively. All estimates are OLS estimates and based on videos that were rated in the online survey experiment (see Appendix D). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

References

Abou El-Komboz, Lena, Anna Kerkhof, and Johannes Loh (2022) "Complementor Reactions to Platform Control – Evidence from the YouTube Partner Program," Unpublished.

Anderson, Ashton, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec (2013) "Steering User Behavior with Badges," in *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13, 95–106: ACM.

Anderson, Chris (2006) *The long tail: Why the future of business is selling less of more*: Hachette Books.

Anderson, Michael and Jeremy Magruder (2012) "Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database," *The Economic Journal*, **122** (563), 957–989.

Anderson, Simon P., Andre De Palma, and Jacques-Francois Thisse (1992) *Discrete choice theory of product differentiation*: MIT press.

Anderson, Simon P. and Bruno Jullien (2016) "The advertising-financed business model in two-sided media markets," in *Handbook of Media Economics*, **1**: Elsevier, 41–90.

Anderson, Simon P and Joel Waldfogel (2015) "Preference externalities in media markets," in *Handbook of Media Economics*, **1**: Elsevier, 3–40.

Angrist, Joshua D and Guido W Imbens (1995) "Identification and Estimation of Local Average Treatment Effects," Working Paper 118, National Bureau of Economic Research.

Aral, Sinan and Dylan Walker (2012) "Identifying Influential and Susceptible Members of Social Networks," *Science*, **337** (6092), 337–341.

Armstrong, Mark and Yongmin Chen (2009) "Inattentive consumers and product quality," *Journal of the European Economic Association*, **7** (2-3), 411–422.

Armstrong, Mark and Helen Weeds (2007) "Programme quality in subscription and advertising-funded television," *Unpublished working paper.*

Beebe, Jack H. (1977) "Institutional structure and program choices in television markets," *The Quarterly Journal of Economics*, **91** (1), 15–37.

Bergemann, Dirk and Alessandro Bonatti (2011) "Targeting in advertising markets: implications for offline versus online media," *The RAND Journal of Economics*, **42** (3), 417–443.

Berry, Steven T and Joel Waldfogel (2001) "Do mergers increase product variety? Evidence from radio broadcasting," *The Quarterly Journal of Economics*, **116** (3), 1009–1025.

Bishop, Sophie (2018) "Anxiety, panic and self-optimization: Inequalities and the YouTube algorithm," *Convergence*, **24** (1), 69–84.

Blei, David M, Andrew Y Ng, and Michael I Jordan (2003) "Latent dirichlet allocation," *Journal of Machine Learning Research*, **3** (Jan), 993–1022.

Bond, Robert, Christopher Fariss, Jason J. Jones, Adamd. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler (2012) "A 61-million-person experiment in social influence and political mobilization," *Nature*, **489** (7415), 295–298.

Bourreau, Marc (2003) "Mimicking vs. counter-programming strategies for television programs," *Information Economics and Policy*, **15** (1), 35–54.

Brynjolfsson, Erik, Yu Hu, and Duncan Simester (2011) "Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales," *Management Science*, **57** (8), 1373–1386.

Brynjolfsson, Erik, Yu Hu, and Michael D Smith (2003) "Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers," *Management Science*, **49** (11), 1580–1596.

Burnham, Thomas A, Judy K Frels, and Vijay Mahajan (2003) "Consumer switching costs: A typology, antecedents, and consequences," *Journal of the Academy of Marketing Science*, **31** (2), 109–126.

Chandra, Ambarish (2009) "Targeted advertising: The ol of subscriber characteristics in media markets," *The Journal of Industrial Economics*, **57** (1), 58–84.

Chandra, Ambarish and Ulrich Kaiser (2014) "Targeted advertising in magazine markets and the advent of the internet," *Management Science*, **60** (7), 1829–1843.

Chen, Yongmin and Michael H Riordan (2007) "Price and variety in the spokes model," *The Economic Journal*, **117** (522), 897–921.

Chevalier, Judith and Dina Mayzlin (2006) "The effect of word of mouth on sales: Online book reviews," *Journal of Marketing Research*, **43** (3), 345–354.

Covington, Paul, Jay Adams, and Emre Sargin (2016) "Deep neural networks for youtube recommendations," in *Proceedings of the 10th ACM conference on recommender systems*, 191–198.

Cunningham, Stuart and David Craig (2017) "Being 'really real'on YouTube: authenticity, community and brand culture in social media entertainment," *Media International Australia*, **164** (1), 71–81.

d'Aspremont, Claude, J Jaskold Gabszewicz, and J-F Thisse (1979) "On Hotelling's" Stability in competition"," *Econometrica*, 1145–1150.

Davcik, Nebojsa S and Piyush Sharma (2015) "Impact of product differentiation, marketing investments and brand equity on pricing strategies," *European Journal of Marketing*.

Durbin, James (1954) "Errors in variables," *Revue de l'institut International de Statistique*, 23–32.

Easley, David and Arpita Ghosh (2013) "Incentives, gamification, and game theory: An economic approach to badge design," *ACM Transactions on Economics and Computation (TEAC)*, **4** (3), 16.

Economides, Nicholas (1984) "The principle of minimum differentiation revisited," *European Economic Review*, **24** (3), 345–368.

——— (1986) "Minimal and maximal product differentiation in Hotelling's duopoly," *Economics Letters*, **21** (1), 67–71.

Gabszewicz, Jean and Jaques-Francois Thisse (1986) "Spatial competition and the location of firms," in *Fundamentals of pure and applied economics*, **5**: Routledge, 1–71.

George, Lisa (2007) "What's fit to print: The effect of ownership concentration on product variety in daily newspaper markets," *Information Economics and Policy*, **19** (3-4), 285–303.

Goldfarb, Avi and Catherine Tucker (2019) "Digital economics," *Journal of Economic Literature*, **57** (1), 3–43.

Hamilton, James (2004) *All the news that's fit to sell: How the market transforms information into news*: Princeton University Press.

Hansen, Stephen, Michael McMahon, and Andrea Prat (2018) "Transparency and deliberation within the FOMC: a computational linguistics approach," *The Quarterly Journal of Economics*, **133** (2), 801–870.

Herman, Edward S and Robert Waterman McChesney (1997) *The global media: The new missionaries of corporate capitalism*: Burns & Oates.

Hotelling, Harold (1929) "Stability in Competition," *The Economic Journal*, **39**, 41–57.

Huang, Jason, David Reiley, and Nick Riabov (2018) "Measuring Consumer Sensitivity to Audio Advertising: A Field Experiment on Pandora Internet Radio."

Kleibergen, Frank and Richard Paap (2006) "Generalized reduced rank tests using the singular value decomposition," *Journal of Econometrics*, **133** (1), 97–126.

Klemperer, Paul (1995) "Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade," *The Review of Economic Studies*, **62** (4), 515–539.

Luca, Michael (2016a) "Reviews, reputation, and revenue: The case of Yelp. com," Unpublished.

——— (2016b) "User-generated content and social media," in *Handbook of Media Economics*, **1**: Elsevier, 563–592.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier (2014) "Promotional Reviews: An Empirical Investigation of Online Review Manipulation," *American Economic Review*, **104** (8), 2421–2455.

McChesney, Robert D (2004) *The problem of the media: US communication politics in the twenty-first century*: NYU Press.

Netz, Janet S and Beck A Taylor (2002) "Maximum or minimum differentiation? Location patterns of retail outlets," *Review of Economics and Statistics*, **84** (1), 162–175.

Numan, Wieteke and Marleen Willekens (2012) "An empirical test of spatial competition in the audit market," *Journal of Accounting and Economics*, **53** (1-2), 450–465.

Remus, Robert, Uwe Quasthoff, and Gerhard Heyer (2010) "Sentiws-a publicly available german-language resource for sentiment analysis," in *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10).*

Ronnen, Uri (1991) "Minimum quality standards, fixed costs, and competition," *The RAND Journal of Economics*, 490–504.

Rosen, Sherwin (1981) "The Economics of Superstars," *The American Economic Review*, **71** (5), 845–858.

Salop, Steven C (1979) "Monopolistic competition with outside goods," *The Bell Journal of Economics*, 141–156.

Seamans, Robert and Feng Zhu (2014) "Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers," *Management Science*, **60** (2), 476–493.

Shaked, Avner and John Sutton (1982) "Relaxing price competition through product differentiation," *The Review of Economic Studies*, 3–13.

——— (1983) "Natural oligopolies," *Econometrica*, 1469–1483.

Steiner, Peter O. (1952) "Program patterns and preferences, and the workability of competition in radio broadcasting," *The Quarterly Journal of Economics*, **66** (2), 194–223.

Stock, James H and Motohiro Yogo (2002) "Testing for weak instruments in linear IV regression."

Sun, Monic and Feng Zhu (2013) "Ad revenue and content commercialization: Evidence from blogs," *Management Science*, **59** (10), 2314–2331.

Sweeting, Andrew (2010) "The effects of mergers on product positioning: evidence from the music radio industry," *The RAND Journal of Economics*, **41** (2), 372–397.

Tolson, Andrew (2010) "A new authenticity? Communicative practices on YouTube," *Critical Discourse Studies*, **7** (4), 277–289.

Waldfogel, Joel (2007) *The tyranny of the market: Why you can't always get what you want*: Harvard University Press.

——— (2017) "How digitization has created a golden age of music, movies, books, and television," *Journal of Economic Perspectives*, **31** (3), 195–214.

——— (2018) *Digital Renaissance: What Data and Economics Tell Us about the Future of Popular Culture*: Princeton University Press.

Wang, Zhongmin (2010) "Anonymity, social image, and the competition for volunteers: a case study of the online market for reviews," *The BE Journal of Economic Analysis & Policy*, **10** (1), 1–35.

Weeds, Helen (2013) "Programme Quality in Subscription and Advertising-Funded Television."

Wilbur, Kenneth C (2008) "A two-sided, empirical model of television advertising and viewing markets," *Marketing Science*, **27** (3), 356–378.