

Advertising-funded Attention Markets and Antitrust: Evidence from the YouTube Platform

H. Tai Lam*

July 8, 2024

[Click here for the most recent version](#)

Abstract

In many digital markets (e.g., YouTube, Google Maps, LLMs), the “price” consumers pay is with attention to advertising. Consumer welfare in attention product markets is understudied, and we lack ad-attention elasticities necessary for welfare calculations. The Department of Justice (DoJ) is currently suing Google to improve competition in the \$224b upstream advertising market, but this has unclear effects on downstream consumer welfare. Lower ad prices/revenue may increase ad load and/or reduce the supply of ad-funded products, harming consumers. However, ad quality improvements from stronger competition could benefit consumers. Using data on YouTube and its \$29b ad market, I exploit discontinuities (Fuzzy RDD) and sponsor budget exogeneity (IV) to obtain causal estimates of ad-attention elasticities. Creators choose two types of ad load: ad-rolls (on-platform), and sponsor-reads (off-platform), and I find that an additional unit of ad reducing viewership by 14% and 11%, respectively. This difference is robust across samples and implies off-platform ads are higher quality, lending evidence to DoJ’s concerns of stifled innovation for on-platform ads. Using a simple structural model of video consumption and content creation, I show that innovation gains (on-the-support of the data) that improve ad-roll quality to match that of sponsor-reads can offset a 16% increase in ad load or 38% decrease in creator ad revenue.

*University of California, Los Angeles, Anderson School of Management, h.tai.lam@anderson.ucla.edu

1 Introduction

Valuable goods and services in the online economy (e.g., YouTube, Google Maps, LLMs) are provided to consumers at no charge (or at a heavily subsidized price). The lack of a price is part of why these markets have been understudied, particularly when it comes to consumer welfare. The substantial funding required to develop and supply these products comes from online advertising (“ad”) revenue. Thus, the “price” consumers pay is the time and attention they give to ads, and these settings are called attention markets (Calvano and Polo, 2021).

Understanding consumer welfare in these markets is critical, as the \$224 billion in advertising revenue that Google collects each year has attracted intense political and antitrust scrutiny. The US Department of Justice (DoJ) has launched multiple lawsuits against Google for “monopolizing digital advertising technologies”.¹ The DoJ notes in its complaints that “Google has caused great harm to online publishers and advertisers and American consumers”. However, it is unclear whether successful antitrust action that improves competition in upstream markets would actually benefit or harm downstream consumers. This paper is focused on assessing DoJ’s claims of harm on consumers, with this side of the market having received little attention relative to the upstream advertising market.

The DoJ claims that “Google . . . increased the price of advertising” and “a more competitive market would have fostered greater innovation”. Thus, under DoJ’s desired outcome, the price of ads charged to advertisers would fall and additional competing ad platforms would bring innovation to the upstream ad market. Consider the partial equilibrium effects. On the one hand, lower ad prices leads to lower ad revenues, and this may reduce the quantity of the attention market products produced, harming consumers. Further, to the extent that ad quantity increases when ad prices fall, consumers may face a higher ad load. If consumers have net distaste for ad load, consumer welfare will fall. On the other hand, if competitive innovation increases ad quality (e.g., better ad matching or relevance), this would benefit consumers. To determine the net effect on welfare, it is necessary to quantify two key economic forces: (i) consumers’ ad-attention elasticity of viewership and (ii) firms’ production response to advertising revenue. In standard markets, price elasticity of demand and supply provide the basis for calculating welfare effects of policy changes, and ad-attention elasticities play the analogous role in attention markets.

This paper studies the potential impact DoJ’s upstream antitrust action will have on downstream attention market consumers in the YouTube setting, the largest online video platform in the US. Each year YouTube collects around \$29 billion in advertising revenue and

¹<https://www.justice.gov/opa/pr/justice-department-sues-google-monopolizing-digital-advertising-technologies> and <https://www.justice.gov/opa/pr/justice-department-sues-monopolist-google-violating-antitrust-laws>.

disburses \$14.5 billion (55%) to video creators to fund the creation of videos on the YouTube platform. YouTube is an ideal setting for understanding welfare and antitrust effects in attention markets. YouTube accounts for a substantial 13% of Google’s advertising revenue, there is significant variation in advertising load (the quantity of ads viewers are subjected to), and both on-platform and off-platform ads are served. Content creators on YouTube choose from two main sources of advertising: (i) ad-rolls—Google intermediated dynamic ad breaks distributed throughout a video; and (ii): sponsor-reads—an off-platform typically decentralized market of bilateral sponsorship agreements between sponsors and creators, in the form of in-video ad reads by creators. Sponsor-reads are less common and more lucrative than ad-rolls, and the bilateral nature of sponsor-reads makes them costlier and less efficient than ad-rolls. A key argument in DoJ’s lawsuits is that Google’s anticompetitive conduct extinguished competing advertising platforms, and this has led to a lack of innovation in ad quality. The presence of both on- and off-platform advertising on YouTube provides a unique setting to study the potential innovation gains from improvements in ad quality.

I collect data on 3.5 million videos published by 10,107 video creators (i.e., YouTube channels) meeting a sample criteria of at least 1 million subscribers. Using a combination of official APIs, scraping, third-party data and natural language processing, I obtain a panel on video characteristics (views, likes, comments), video-level ad load (ad-rolls and sponsor-reads) and ad prices (cost-per-mille). My scraping collects novel data on each video’s placement of ad-rolls and I parse video transcripts through local LLMs to obtain a high quality measure of each video’s sponsor-reads.

First, I provide model-free causal estimates of ad-attention demand elasticities for the two types of ads. For ad-rolls, I exploit a discontinuity in the eligibility of mid-roll ads (ad-rolls in the middle of a video). The Fuzzy Regression Discontinuity regressions show that the introduction of an additional ad-roll reduces viewership of a video by around 14%. The results are robust to bandwidth selection and falsification tests. For sponsor-reads, I exploit the exogeneity of advertiser budgets over time to instrument for the probability of a sponsor-read. Two-stage least squares estimates show that a sponsor-read reduces viewership by around 11%. While ex-ante ambiguous, these results show consumers have less net distaste for sponsor-reads and implies that sponsor-reads have higher ad quality than ad-rolls. This lends credence to the DoJ’s arguments that innovation has been stifled on the Google’s vertically integrated ad platform. This may be because creators can exert finer control on their choice of advertiser or engage in novel ad revenue contracts, options not available on YouTube’s on-platform ad network.

Motivated by these estimates, I construct a simple structural model to understand how the supply of videos may change when creators receive less ad revenue, and how consumers

may change their video consumption behavior when ad load increases and the availability of videos shift. The structural model provides additional evidence that sponsor-reads are higher quality, suggests limited tastes for variety within categories and shows that video creators have a strong intrinsic motivation to produce videos. In arguments of antitrust it is simple to claim innovations will benefit consumers, but it is typically difficult to quantify. Using my results, I can see how consumers will be affected by a plausible innovation that is observed in the data, namely where the ad quality of sponsor-reads matches that of ad rolls. To help policymakers contextualize the potential effects on consumers of improved upstream competition, the results show that the plausible innovation will offset a 16% decrease in ad prices. Heterogeneity in video creator responses and viewer taste preferences mean that entertainment and lifestyle categories see the greatest declines in viewership, with education categories being surprisingly robust to shifts in the advertising market.

1.1 Literature and Contributions

There is a growing literature recognizing the importance of attention markets ([Calvano and Polo, 2021](#)) and this paper contributes to the less studied consumer side of these multi-sided markets.

Unlike price elasticities in standard product markets, estimates of ad-attention elasticities in attention markets are rare in the literature. [Wilbur \(2008\)](#) calculates consumers ad-attention elasticities for TV viewership using Nielsen data as part of a structural model examining the effects of ad avoidance technology. The exercise for this paper is similar, but in the context of digital advertising and the outcomes of antitrust action. In a digital platform setting, Spotify ran experiments varying ad frequency and duration, providing clean causal estimates of consumption and subscription ad-attention elasticities ([Huang, Reiley and Riabov, 2018](#); [Goli, Reiley and Zhang, 2021](#)).

There is an extensive literature in advertising elasticities ([Shapiro, Hitsch and Tuchman, 2021](#)) that is separate to the niche literature on ad-attention elasticities. Advertising elasticities typically refer to how consumers change their purchase of a brand when exposed to ads from that brand. Advertising elasticities are important for understanding wide-ranging effects on consumer purchase behavior in the presence of ubiquitous advertising, but does not shed light on the provision of attention products (e.g., Google search, Google maps and YouTube) nor the welfare consumers obtain from attention products.

There is a rich literature examining various important aspects of the YouTube video platform ([Tang, Gu and Whinston, 2012](#); [Yoganarasimhan, 2012](#); [Kerkhof, 2024](#); [El-Komboz, Kerkhof and Loh, 2023](#)). Notably, YouTube has been the setting for other important legal

policy questions. [Li, Haviv and Lovett \(2024\)](#) provides results on the validity of fair-use doctrine by demonstrating that play-throughs of video games on YouTube can serve as a substitute for purchase. [Johnson, Lin, Cooper and Zhong \(2024\)](#) show that the ad personalization changes brought about by the Children’s Online Privacy Protection Act markedly affected the quantity and quality of video creation on the platform.

More generally, many aspects studied in the online advertising literature ([Choi, Mela, Balseiro and Leary, 2020](#); [Gordon, Jerath, Katona, Narayanan, Shin and Wilbur, 2021](#)) is relevant in my setting, including: advertising non-disclosure ([Sahni and Nair, 2020](#)), ad content quality and consistent ([Kapoor et al., 2022](#)) and the trade-offs between subscription and advertising revenue ([Casner and Teh, 2023](#)).

Lastly, the literature on antitrust of multi-sided markets ([Evans, 2003](#); [Evans and Schmalensee, 2013](#)) cautions that standard models of antitrust may not apply to multi-sided markets. This literature emphasizes the need to consider interrelated demand across different sides of the market when evaluating antitrust policies. My paper quantifies such externalities in a settings of imminent and ongoing antitrust action.

2 Background

This section discusses institutional details of the YouTube platform relevant for understanding demand elasticities and the antitrust questions of the paper.

2.1 The YouTube Platform

It is important to note that the DoJ’s antitrust cases does not cover the YouTube platform specifically. However, the YouTube platform provides an ideal environment to study many of the antitrust issues raised by the DoJ. Firstly, YouTube generates a substantial 13% (\$29 billion) of the total \$224 billion Google collects in ad revenue. YouTube’s share, and video ads across the internet, are a fast growing part of digital advertising. Secondly, DoJ’s concerns of vertical control of advertising markets applies to the YouTube’s ad-rolls as it does to Google’s display ads around the open web. Publishers on the open web (e.g., Washington Post) are analogous to video creators on YouTube. And much like how publishers sell prime advertising space off ad platforms (e.g., banner ads on the New York Times website), creators on YouTube also negotiate bilateral sponsor-reads.

2.2 Two Primary Forms of Ads

There are two primary forms of advertising on YouTube: ad-rolls and sponsor-reads. Video creators have complete control over whether to have any, both or none of these forms of advertising. They can choose the ad load on a video-by-video basis, allow YouTube to choose or bulk update their ad load choices.

Ad-rolls—are ad breaks (similar to TV ad breaks) inserted by the platform at the beginning (pre-roll), middle (mid-roll) or end (end-roll) of a video. Ad-rolls are platform intermediated in that creators are selling their ad space through the Google ad platform and advertisers go to a Google ad auction to purchase the ad space. Ad-rolls are dynamic, with different ads served to different viewers, differing over time depending on who participates at the time of the respective ad auction.

Sponsor-reads—are ad reads delivered by the creators/hosts of the video within the video published on YouTube. They mainly come about from bilateral agreements between creators and advertisers (e.g., advertisers emailing creators with offers of sponsorships).² Sponsor-reads are static, where the ad read is embedded within the video uploaded, does not change over time nor does it differ across viewers. While sponsor-reads are not personalized, creators likely only enter into bilateral negotiations with advertisers they believe their audience values. Creators often appear on-screen speaking in their own voice and will often tailor the content of the ad reads to their audience.

Figure 1: Two Forms of Advertising



Notes: The left image depicts a YouTube video discussing US elections. The center image depicts an ad-roll, which is an interruption that plays over a YouTube video. The right image depicts a sponsorship ad read by the creator of the YouTube video.

In this paper, I consider an instance of ad interruption as the unit of ad load, rather than the length of the ad. This is because essentially all ad-rolls can be skipped after 5 seconds using a skip ad button and sponsor-reads can be fast forwarded away using standard video

²There are some intermediaries that help connect creators with advertisers, but anecdotal evidence suggests agreements are mainly bilateral.

controls.

Ad-rolls and sponsor-reads differ along many dimensions and it is apriori unclear as to which audiences would prefer. This paper’s elasticity estimates help answer this empirically.

2.3 Multi-sided Market

The multi-sided nature of the market and the lack of available data on all sides of the market makes modelling all sides difficult. I choose to focus on consumers because they have been understudied and left out of the modern anti-trust discourse.

Another consideration is the effects on the market of the advertised product. A substantial literature exists for advertising elasticities, the change in consumer purchase of an advertised good after being exposed to advertising. The effects for advertised products is important, but out of scope, particularly since advertisers will continue to be able to advertise through other channels (TV, physical channels, Facebook, etc.).

3 Illustrative Model

This section provides an simple theory model of an upstream advertising market and a downstream video creation and consumption market. The purpose of the model is to help fix ideas for readers seeking more structure for the claims of ambiguous welfare implications from improved competition, and is not a substantive result of the paper. Such ambiguity in consumer welfare from competition is not new and has been shown in related models of TV advertising (Anderson & Coate 2005). The model also motivates the need for estimating demand elasticities to understand welfare.

3.1 Upstream Advertising Market

The upstream market is laid out as a monopolist platform choosing quantity (a) and quality (m) of advertising inventory to sell to a large number of advertisers. I contrast the monopoly conduct outcomes with the social planner’s outcomes to illustrate the DoJ’s claims of monopoly behavior and the potential gap from total welfare maximizing behavior.

Consider a large number of advertisers that together generate a simple linear inverse demand function:

$$p_a(a) = \delta_0 - \delta_a a + \delta_m m$$

where p_a is the price of one unit of advertising paid by advertisers to the platform, a is the

total number of advertising units supplied by the platform (e.g., sum of impressions on a 15 second spots) and m is the overall quality of each unit of advertising (e.g., the match quality of advertising spots to its audience). Thus, δ_m is the degree to which higher quality advertising shifts out the demand curve for advertising spots, incorporating advertisers' willingness to pay for higher quality matching, more data/tracking or other innovations.

The monopolist platform's profit maximizing problem and the interior solution to the FOCs are given by:

$$\begin{aligned} \max_{a,m} & p_a(a) \cdot a - \gamma_a \cdot a - \gamma_m m^2 \\ a^{*Mon} &= \frac{2\gamma_m(\delta_0 - \gamma_a)}{4\gamma_m\delta_a - \delta_m^2} \\ m^{*Mon} &= \frac{\delta_m(\delta_0 - \gamma_a)}{4\gamma_m\delta_a - \delta_m^2} \end{aligned}$$

where γ_a is the marginal cost of an additional advertising unit, $\gamma_m m^2$ is the quadratic cost associated with with quality. Note that γ_a is the cost of generating an advertising unit, which comes from the downstream video content market discussed further below. a^{*Mon} and m^{*Mon} are the monopolist's optimal supply of quantity and quality. Assume throughout that the parameters are such that we have an interior solution, $a^{*Mon} > 0, m^{*Mon} > 0$.

A social planner for the upstream market would instead maximizing upstream total surplus, with the resulting optimal choices given by:

$$\begin{aligned} \max_{a,m} & \int_0^a p_a(a) da - \gamma_a \cdot a - \gamma_m m^2 \\ a^{*TS} &= \frac{2\gamma_m(\delta_0 - \gamma_a)}{2\gamma_m\delta_a - \delta_m^2} \\ m^{*TS} &= \frac{\delta_m(\delta_0 - \gamma_a)}{2\gamma_m\delta_a - \delta_m^2} \end{aligned}$$

Under the assumption of interior solutions, it is always the case that $a^{*TS} > a^{*Mon}$ and $m^{*TS} > m^{*Mon}$. That is, the monopolist platform always under-supplies quantity and quality of advertising relative to what a social planner would supply. Note that the functional forms are chosen to obtain under-supply, as this matches the concerns laid out by the DoJ and the paper seeks to explore the consequences for downstream consumer welfare given this baseline. Thus, the upstream market is constructed as a special case of Spence (1975).

3.2 Downstream Video Content Market

While the upstream advertising market is where antitrust action is taking place, this paper is interested in the welfare effects on downstream consumers. Here, a large number of video content creators produce video in return for payment from the platform. The multi-sided nature of the setting becomes apparent when incorporating the downstream market. Content creators do not receive payment from consumers, rather the platform pays content creators for each video consumed. The platform decides the advertising load (advertising units per video) and consumers “pay” with their attention to advertising (e.g., generating an impression). That is, the quantity in the upstream market directly determines the “price” in the downstream market.

Linking the upstream and downstream markets is the platform’s choice of quantity of advertising units (a). Since the upstream quantity is related to the downstream “price”, we introduce α , the advertising load per unit of video (v) such that:

$$\begin{aligned} a &= \alpha * v \\ \alpha &= \frac{a}{v} \end{aligned}$$

Consumers (i.e., viewers) generate a constant elasticity inverse demand:

$$\alpha(v) = (\omega_0 + \omega_m \cdot m) v^{-\frac{1}{\omega_v}}$$

where ω_m is how match quality of advertising shifts out the demand for videos (e.g., advertising being less intrusive), and ω_v is the ad-attention elasticity of demand ($\epsilon = -\frac{\partial v}{\partial \alpha} \frac{\alpha}{v} = \omega_v$).

I assume free entry of content creators, such that supply of videos always meets demand to clear the market. This is because the advertising load (“price”) is already determined by the upstream monopolist platform seeking to generate advertising units. It follows that:

$$\begin{aligned} \pi_{cc} &= \gamma_a * v - \phi_v * v^2 \\ \gamma_a^* &= v \phi_v \end{aligned}$$

The key measure of interest is consumer surplus and how it changes when upstream competition improves. Consumer surplus, assuming $\omega_v > 1$, is given by:

$$\begin{aligned} \text{CS} &= \int \alpha(v) dv - \alpha \cdot v \\ &= \frac{\alpha \cdot v}{\omega_v - 1} \end{aligned}$$

where the second line follows from standard results on constant elasticity demand functions.

Thus, changes in downstream consumer surplus from improvements in upstream competition will depend first-order on consumers' ad-attention elasticity. Second-order changes to the advertising load (α), number of videos (v), and the quality of advertising (m) will also be important to quantify.

4 Data and Descriptives

This section introduces the data and provides an overview of the YouTube video platform.

I construct a sample of the largest YouTube creators and their videos. The sample criteria is English speaking channels from the US, Canada, UK, Australia and undeclared regions with more than 1 million subscribers. I use a combination of the official YouTube API, automated scraping and third-party data sources to compile the working dataset. This results in a sample of 10,107 creators and 3,503,823 videos. [Table 1](#) provides channel-level summary statistics, with the average creator in our sample having published 347 videos (6 videos per month), averaging 12 mins and 2.8 million views. The creators should be thought of as small businesses in the business of producing regular videos, that are consistently consumed by a significant amount of viewers.

This study defines a "view" as a unit of consumption on YouTube. Unlike traditional goods where a purchase represents a unit of consumption, YouTube views are only recorded after a user watches at least 30 seconds of a video, regardless of whether the video was accessed through active selection or autoplay. This definition aligns with YouTube's internal measurement of demand, which is critical for revenue allocation to content creators. While the 30-second threshold may appear restrictive, it is the official metric used for ad revenue calculations and thus represents the most reliable measure of consumption available. Furthermore, small-sample viewing duration data indicates that users who watch at least 30 seconds of a video tend to view an average of 5 additional minutes, suggesting that this threshold serves as a reasonable proxy for meaningful engagement with content.

The data on ad load is worth noting, with 5% of videos featuring an in-video sponsor-read, while YouTube mid-roll ads are featured in about 50% of all videos. This is consistent with the lower cost for creators to access YouTube ad-rolls, which is provided centrally by the platform. In contrast, sponsor-reads are comparatively rare, reflecting the expensive (mostly) bilateral off-platform nature of sponsor-reads. It is important to collect accurate data on ad-rolls and sponsor-reads. For ad-rolls, my scrapers mimic a new user on YouTube, clicking on and viewing a sequence of suggested videos on the home page before proceeding

to the videos to be collected. The scraper plays the videos and records when ad-rolls are served.

For sponsor-reads, there are no official data sources recording whether a video has a sponsor-read within it, and disclosure in video descriptions or using YouTube disclosure flags are known to be severely under reported. I obtain and code sponsor-reads directly from transcripts of the videos using local LLMs. Using LLMs to code sponsor-reads from transcripts will provide the highest quality data as many video creators only disclose their sponsorship verbally in the video. Additionally, the contextual understanding of LLMs are needed. Simply coding a video as sponsored based on keywords (e.g., this video is sponsored) will overstate the incidence of sponsorships. For example, “this isn’t a situation where I’ve been paid money and entered into an agreement with Pepsi to sponsor this video”, contains a critical negation keyword far from the mention of the sponsorship. Specifically, I take roughly 45 second snippets of the transcripts around mentions of specific keywords to parse through the LLM. To ensure replicability, I use a locally-run LLM created by Berkeley (Starling-7B, 4 bit quantized) with 0 temperature. Parsing each snippet of transcript takes around . I use few-shot prompting to address common classification errors like discussions about sponsorship of athletes and congressional bill sponsors. While non-disclosure by influencers is a problem particularly on platforms like twitter, I do not find many instances of non-disclosure, likely due to the sample being large video creators. Out of sample classification error (against research assistant ground truth) is about 3% after successive iterations of the few-shot prompting.

The majority of the data was collected in June 2023, presenting some data challenges. For example, creators who exit the platform and delete their accounts, or specific videos that are deleted, are not included in the dataset. Moreover, cumulative metrics, such as views, are observed at the time of data collection, rather than representing a complete historical record. To mitigate these limitations, we focus on a consistent measure of video views across all videos published over time: the number of views 30 days after publication. This timeframe is also commonly used by sponsors to determine impressions and payouts to creators. Utilizing data from the Internet Archive, we estimate the number of views 30 days after publication for the entire sample. This approach allows us to control for the inherent time dependency of views while providing a consistent measure across all videos in our sample.

Table 1: Summary Statistics - Channel-level

Characteristic	Selected Categories					
	All ¹	<5mil sub. ¹	5mil+ sub. ¹	Entertainment ¹	Politics ¹	Technology ¹
Subscribers (mil)	3.15 (4.91)	1.97 (0.97)	10.93 (10.39)	3.54 (4.92)	2.87 (2.76)	2.45 (2.31)
# of Videos	347 (345)	342 (334)	421 (404)	326 (256)	693 (350)	351 (254)
# of Videos/Month (when active)	6.2 (6.2)	6.2 (6.0)	6.5 (7.5)	7.0 (6.1)	15.1 (19.4)	6.0 (3.4)
Video length (mins)	12 (22)	13 (22)	12 (19)	12 (22)	15 (14)	4 (3)
Views (,000)	2,805 (12,007)	1,811 (6,509)	7,137 (25,209)	2,802 (5,802)	233 (388)	709 (1,221)
Comments (,000)	2.76 (111.98)	2.50 (121.13)	3.97 (8.59)	1.45 (5.05)	1.16 (1.60)	0.36 (0.29)
Likes (,000)	51 (184)	39 (172)	123 (220)	70 (154)	9 (16)	22 (59)
Has mid-roll (%)	48 (31)	48 (31)	53 (29)	49 (32)	25 (22)	21 (26)
Has sponsor (%)	5 (10)	5 (10)	4 (8)	3 (8)	7 (12)	2 (4)
Ad-roll CPM (\$)	3.37 (0.68)	3.37 (0.69)	3.32 (0.61)	3.30 (0.73)	3.48 (1.02)	3.38 (0.38)
Sponsorship CPM (\$)	23 (13)	22 (10)	25 (26)	23 (19)	24 (6)	25 (5)

¹Mean (SD)

Even though our sample of YouTube is constructed from large content creators, the aggregate time series trends match those reported publicly by YouTube and other data providers. YouTube has grown in size exponentially over the last few years. It is also important to note that sponsor-reads are a fairly recent phenomena.

5 Causal Estimates of ad-attention Elasticity

Antitrust action in the upstream advertising market seeks to lower ad prices, which will likely affect ad quantity and ad load. Thus, evaluating consumer welfare requires measuring how consumers respond to changes in ad load. On YouTube, this comes in two forms: (i) ad-rolls; and (ii) sponsor-reads reads. I provide causal estimates of elasticities for both types of ad load in this section.

5.1 Ad-attention Elasticity of Demand - Ad-Rolls

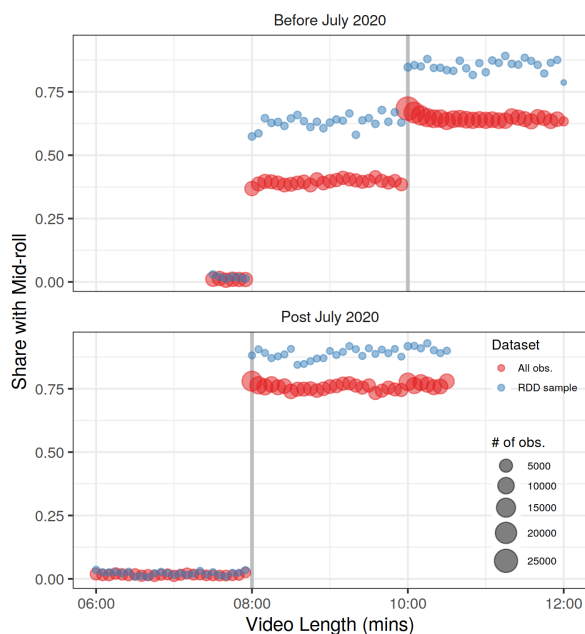
This subsection provides evidence that consumers respond to YouTube’s mid-rolls and estimates the causal effect of an additional mid-roll on viewership (i.e., the ad-roll ad-attention elasticity).

Videos that feature ads may not be comparable to videos without ads, particularly when creators endogenously choose ad load. To address these endogeneity concerns I employ a Fuzzy Regression Discontinuity Design that exploits a discontinuity in the eligibility of mid-roll ads (i.e., ads that play in the middle of a video). YouTube’s centralized ad platform allows content creators to choose when ads play on a video, before the video (pre-roll), during the video (mid-roll) and after the video (end-roll). However, if a video’s length is under a threshold, YouTube does not allow any mid-roll ads to run. The threshold was 10 minutes

prior to July 2020, and was reduced to 8 minutes post July 2020. Having two discontinuities provide an internal check against 8 minutes or 10 minutes being unique in a way unrelated to the additional ad-rolls.

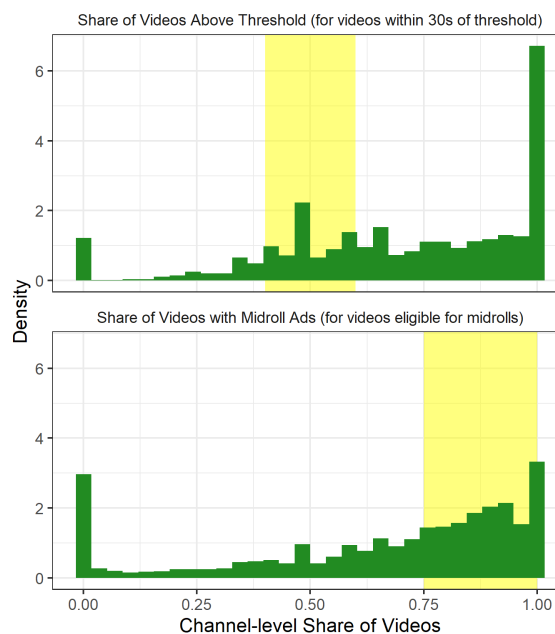
Figure 2 depicts the mid-roll (i.e., treatment assignment) discontinuity across video's length, for two samples: the entire dataset, and an RDD sample that addresses manipulation. Post July 2020, videos shorter than 8 minutes do not feature mid-rolls, while videos longer than 8 minutes have a significantly higher chance of featuring a mid-roll. For video pre-July 2020, there is some noise in the discontinuity for videos between 8 minutes and 10 minutes. This is attributable to creators being able to bulk update their mid-roll ad decisions for all of their videos. These videos did not have mid-rolls but were instead subsequently updated to have mid-rolls post-July 2020. Since I am primarily interested in the effect of mid-roll ads on views accrued in the first 30-days of a videos' life, pre-July 2020 videos under 10mins are re-coded as not having mid-rolls, which would have been the case given YouTube's rules.

Figure 2: Ads Discontinuity



Notes: None

Figure 3: Channel Manipulation



Notes: None

Two issues must be addressed before utilizing the discontinuity: content creators can manipulate their video length (i.e., the running variable) and content creators make a choice whether to include or exclude mid-roll ads. This manipulation can be observed in Figure 2, with bunching in the number of videos and a slight increase in ad-rolls above the threshold. Unlike standard RDD, where there is only one observation per decision-maker, here we observe many video length and ad decisions from the same creator, allowing us to separate

manipulators from compliers [Figure 3](#). For estimating the causal consumer response, I restrict the RDD sample to complying creators: (i) creators who do not bunch above the threshold (i.e., near random, 40% to 60% share of videos above and below threshold, top-panel shaded in [Figure 3](#)); and (ii) who turn on mid-roll ads for the majority of their videos (i.e., over 75%, bottom-panel shaded in [Figure 3](#)). The RDD sample exhibits no signs of bunching or manipulation [Figure 2](#). The RDD sample is broadly representative and an example of such creators include the official BBC news channel, where TV news content serves as the natural source of exogeneity for video length. Standard caveats with RDD, particularly with respects to locality are important to keep in mind when interpreting the results.

The Fuzzy RDD is implemented with the mid-roll being the treatment, video length as the running variable, the threshold at the appropriate 8 or 10 mins, the instrumental variable being an indicator above the threshold and using a bandwidth of 30 seconds.

The main Fuzzy RDD results are presented in column 3 of [Table 3](#). An additional mid-roll ad leads to around 14% fewer views. I report the change in viewership from one additional ad as opposed to a proper elasticity, since the instance of ad interruption is the primary cost when ads can be skipped. As a sense check, the naive regression in column 1 demonstrates the expected bias and the 1st stage in column 2 (and F-stat in column 3) shows there is relevance in the instrument.

Table 3: Fuzzy RDD of mid-roll Ad on Video Views

	<i>Dependent variable:</i>		
	log(Views ,000) OLS	Mid-roll=1 1st Stage	log(Views ,000) IV
	(1)	(2)	(3)
Has midroll	-0.046 (0.040)		-0.155*** (0.043)
Above threshold		0.849*** (0.007)	
Exp(dep.) mean	1350	1.671	1350
1st Stage F-stat			14500
Controls:			
-Video length quadratic	Y	Y	Y
-Weeks since publication	Y	Y	Y
-Channel FEs + Cl. SEs	Y	Y	Y
Observations	9,102	9,102	9,102

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 4 provides heterogeneity results for the 6 (out of 16) topic areas for which there are statistically significant results.

Table 4: Fuzzy RDD across Topics

	<i>Dependent variable:</i>					
	log(Views ,000)					
	Knowledge	Society	Other	Health	Entertainment	Politics
	(1)	(2)	(3)	(4)	(5)	(6)
Has midroll	−0.428** (0.178)	−0.323*** (0.106)	−0.274*** (0.090)	−0.175** (0.073)	−0.369*** (0.112)	−0.300** (0.125)
Exp(dep.) mean	697.8	2107	220.4	1089	2670	809.8
1st Stage F-stat	290	2400	700	4610	2230	2680
Controls:						
-Video length quadratic	Y	Y	Y	Y	Y	Y
-Weeks since publication	Y	Y	Y	Y	Y	Y
-Channel FEs + Cl. SEs	Y	Y	Y	Y	Y	Y
Observations	980	2,668	418	4,462	2,607	1,961

Note:

*p<0.1; **p<0.05; ***p<0.01

The results are fairly robust to varying the bandwidth, though bandwidths under 30 seconds are significantly under-powered (Table 5).

Table 5: Bandwidth Selection

	<i>Dependent variable:</i>				
	log(Views ,000)				
	60s	45s	30s	20s	10s
	(1)	(2)	(3)	(4)	(5)
Has midroll	−0.259*** (0.044)	−0.220*** (0.043)	−0.150*** (0.041)	−0.086* (0.046)	−0.036 (0.063)
Exp(dep.) mean	1413	1447	1351	1414	1393
1st Stage F-stat	16400	15300	14500	12600	4910
Controls:					
-Video length quadratic	Y	Y	Y	Y	Y
-Weeks since publication	Y	Y	Y	Y	Y
-Channel FEs + Cl. SEs	Y	Y	Y	Y	Y
Observations	17,637	13,297	9,102	6,195	3,261

Note:

*p<0.1; **p<0.05; ***p<0.01

Additionally, the results survive a falsification test assuming the false threshold is at 1 minute plus the true threshold. Additionally, no left digit bias effect was found.

Table 6: Falsification Test

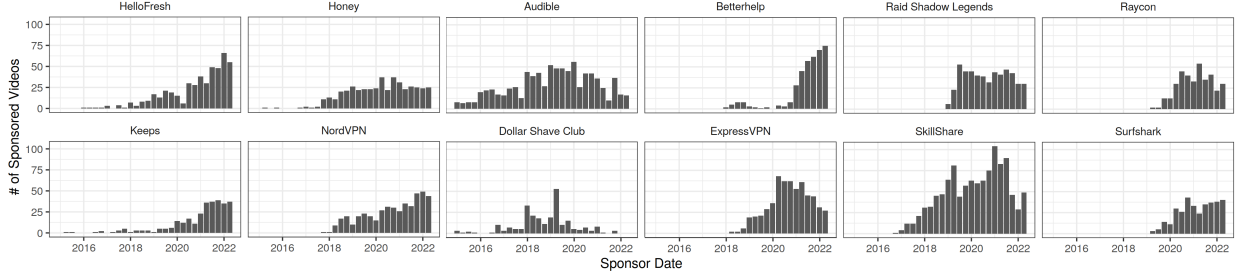
	<i>Dependent variable:</i>
	log(Views ,000)
	+1 min
Has midroll	20.907 (35.302)
Exp(dep.) mean	1269
1st Stage F-stat	0.34
Controls:	
-Video length quadratic	Y
-Weeks since publication	Y
-Channel FEs + Cl. SEs	Y
Observations	7,676
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

5.2 Ad-attention Elasticity of Demand - Sponsor-Reads

This subsection provides evidence that consumers respond to creator’s sponsor-reads and estimates the causal effect of an additional sponsor-read on views (i.e., the sponsor-read ad-attention elasticity).

Sponsorships are predominantly bilateral agreement between video creators and advertisers. Nevertheless, like ad-rolls, creators choose to enter into these agreements, creating an endogeneity problem. Additionally, once a contract is signed, creators may opt to place sponsor-reads in videos they expect to do particularly well. To address these endogeneity issues, I employ an instrumental variables approach that exploits the random variation in advertisers’ budget across time. [Figure 4](#) depicts the largest sponsors on the platform, illustrating the random variation in sponsorship entry and volume of sponsor reads across time. Sponsors enter into hundreds of bilateral agreements in any fiscal year and their budgets are unlikely to be driven by any individual video creator unobservable (i.e., the source of sponsor-read endogeneity).

Figure 4: Sponsor Activity



Notes: None

I construct a leave-1-out sponsor intensity instrument:

$$\text{Sponsor Intensity} = \text{z-score}_{\text{category-year}} \left(\sum_{j \in \text{month}} 1(\text{Sponsored}) \right)$$

to capture the random variation in sponsor entry and budget variation within each category-year grouping of videos. Sponsors generally target entire segments of creators for sponsorships that match their business or product.

Table 7: Sponsorship IV

	<i>Dependent variable:</i>				
	log(Views ,000)	Sponsored=1	log(Views ,000)	log(Likes ,000)	log(Comments ,000)
	OLS	1st Stage	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)
Has Sponsor	-0.123*** (0.009)		-0.112*** (0.010)	-0.102*** (0.009)	-0.054*** (0.009)
Sponsor Intensity		0.010*** (0.00001)			
Exp(dep.) mean	482.3	1.119	482.3	28.47	1.605
1st Stage F-stat			4980	4960	4710
Controls:					
-Video length quadratic	Y	Y	Y	Y	Y
-Weeks since publication	Y	Y	Y	Y	Y
Channel-year FEs	Y	Y	Y	Y	Y
Channel Cl.SEs	Y	Y	Y	Y	Y
Observations	836,037	836,037	836,037	829,811	771,433

Note:

*p<0.1; **p<0.05; ***p<0.01

The main IV results are presented in column 3 of Table 7. The regression includes a standard battery of fixed effects and conservative standard error calculations. The main takeaway is that an additional sponsor-read leads to around 11% fewer views on a video. The 1st stage confirms the relevance of the instrument. Similarly but smaller in magnitude effects can be seen for engagement (e.g, likes and comments) metrics.

5.3 Ad Quality

An important claim in the DoJ’s complaints is that Google’s monopolistic behavior stifled innovation in the market. Potential innovations are wide ranging, but can be formalized through ad quality. For this paper, I focus on ad quality from the perspective of consumers.³ A high quality ad for a consumer may be one that provides valuable information, increase utility of the advertised good, or have direct entertainment value (the advertisement is enjoyable in of itself). Innovations could allow better matching of advertisers and consumers (i.e., consumers may enjoy ads that are more relevant to themselves).


It is apriori unclear whether ad-rolls or sponsor-reads are higher quality. Ad-rolls are served dynamically and in theory can be hyper-personalized to each individual viewer. However, sponsor-reads allow creators to carefully curate and select advertisers their audience will respond to. Sponsor-reads allow creators to speak in their own voice and tailor the advertising message to their audience (e.g., emphasizing valuable features and disarding ineffective messaging). While such ad-creator congruence is within the reach of current technologies, no such innovation has appeared on YouTube ad networks.

The ad-attention elasticity results in the prior subsections for ad-rolls (14%) and sponsor-reads (11%) suggests that sponsor-reads are in fact higher quality. However, the different methodologies lend themselves to different estimating samples and a natural concern is whether this holds. I use a consistent sample of videos with a bandwidth of 1 min around the ad-roll thresholds and find the results is robust. Additionally, splitting across the major categories of videos also delivers a consistent, if noisier, ordering of ad-rolls and sponsor-reads (Figure 5).

In sum, our results suggest there is validity to the DoJ’s concerns about stifled innovation, with platform participants having to bilaterally negotiate due to the inability to obtain similar ad quality arrangements on the centralized ad platform.

³Note, another relevant measure may be the ad quality as interpreted by advertisers (e.g., return on advertising).

Figure 5: Ad-roll and Sponsor-read Across Categories



Notes: None

6 Simple Structural Model

In this section, I provide an illustrative and preliminary structural model of ad-supported video production and consumption. The structural model predicts how video content creation and will change in response to changes in advertising price and advertising quantity.

The model does not include, nor makes predictions about, the upstream advertising market. The current iteration also treats the creator ad load decision as fixed. Given the consumer welfare focus of this paper, I focus on the consumer and creator responses, and seek to draw antitrust implications under a range of possible upstream outcomes.

The timing of the model is as follows:

1. Video content creators choose effort level in response to their expected earnings. Effort level determines the number of videos produced.
2. Viewers choose between a different categories of videos to spend time consuming, know-

ing the ad load associated with each of the types of videos and the number of new videos available in each category.

6.1 Supply of Videos

I model production decisions at the creator-week (ct) level. The average creator in our sample produces around two videos per week, and anecdotally, creators follow a weekly schedule for videos.

I do not model creators' ad load decision, which I assume is fixed at creator-level and determined prior to the effort decision.

Creators decide on effort (unobserved to the econometrician) taking into account a combination of monetary and intrinsic motivations, where higher effort translates to a greater number of videos produced. I model realized effort as:

$$r_{ct} = \underbrace{\gamma_1 \text{Ad-roll Earnings}_{ct} + \gamma_2 \text{Sponsor read Earnings}_{ct}}_{\text{monetary incentives}} + \underbrace{\text{Creator FEs}_c}_{\text{intrinsic motivation \& cost of production}} + \eta_{ct}$$

and

$$\text{Ad-roll Earnings}_{ct} = 30\text{-day Views}_{ct} \times \# \text{ Ad-rolls}_{ct} \times \text{Ad-roll CPM}_c$$

$$\text{Sponsor read Earnings}_{ct} = 30\text{-day Views}_{ct} \times \# \text{ Sponsor reads}_{ct} \times \text{Sponsored CPM}_c$$

where r_{ct} is realized effort, the γ s capture how monetary earnings influence creator effort levels, creator fixed-effects capture anything that does not vary with views and earnings (e.g., intrinsic motivation, difficulty of producing videos and creator ability), η_{ct} is an i.i.d. standard normal effort shock.

The mapping of effort to the number of videos is operationalized as an ordered probit, where effort levels exceeding latent thresholds determine the number of videos produced. Where

$$\text{Number of videos}_{ct} = \begin{cases} 0 & \text{if } r_{ct} < \tilde{r}_1 \\ 1 & \text{if } \tilde{r}_1 \leq r_{ct} < \tilde{r}_2 \\ 2 & \text{if } \tilde{r}_2 \leq r_{ct} < \tilde{r}_3 \\ \vdots & \\ N & \text{if } \tilde{r}_N \leq r_{ct} \end{cases}$$

For simplicity, I assume creators have correct expectations of their 30-day Views_{ct} and adjust effort accordingly.⁴

This simple formulation of the creator video production process allows for creator heterogeneity in the costs of production and intrinsic motivation for producing videos.

6.1.1 Identification of Supply Parameters

The creator parameters are identified using within-creator variation in the number of videos produced on a weekly basis.

I treat variation in 30-day views as exogenous. Consider a creator providing movie reviews or reporting on news about movies. In a week where there is more than the average number of movie releases, the creator correctly expects that they will receive higher views for their content and will exert higher effort and produce a higher than average number of videos.

A few caveats are worth noting. The model assumes creators are in a steady-state of video production, and that variation in video production observed in the data comes from exogenous variation in 30-day views or idiosyncratic effort shocks. Relatedly, the measure of CPM does not vary across time, which requires that the profitability of a creator’s audience for advertisers remains fixed. Since we focus on video creators with over 1 million views, these creators, while steadily growing, have an established type of video and type of audience that enjoys their videos.

6.2 Demand for Videos

I model viewership behavior at a highly aggregated category-week (jt) level, where consumers choose to spend time between 16 categories of videos to view.

The demand model requires a transformation of the underlying video-level views data in viewing hours data. First I estimate daily views for each video using Internet Archive data to distribute the aggregate 30-day views. I transform daily views into daily viewing hours by multiplying the each videos viewing hours by their respective video lengths. Then viewing hours are aggregated into the 16 categories, week by week, for the category-week level demand model. This procedure results in an average of 3.5 hours spent per week viewing YouTube videos. As a sense check, external surveys of time use put U.S. adult viewing hours at 2.9 hours per week.⁵ To back out outside shares, I use Bureau of Labor Statistics’ reported 35 hours per week of leisure hours as the total share.

⁴Our sample of creators have been on YouTube for an average of 5 years, and should therefore be able to accurately predict their expected viewership and earnings.

⁵eMarketer.com survey of U.S. adults.

Demand is modeled as a simply logit, where indirect utility is given by:

$$\delta_{ijt} = \alpha_1 \text{Ad-roll rate}_{jt} + \alpha_2 \text{Sponsor read rate}_{jt} + \beta_1 \# \text{ Videos} + \text{Category FEs}_j + \text{Year FEs}_t + \epsilon_{ijt}$$

where α s capture consumers' sensitivity to ad load, β captures consumers' taste for variety within a category and ϵ_{ijt} is i.i.d. type 1 extreme value.

Note that the micro-foundations of the demand are restrictive. Each decision maker is not an individual, but rather an individual-time instance, obtaining i.i.d. shocks to their preferences. This rules out individual level persistence in tastes for this simply structural model. The setup of the demand model also rules out any competition between firms that a typical demand model seeks to express. and competition is not a focus of the paper.

6.2.1 Identification of Demand Parameters

The same endogeneity issues discussed in the model-free section apply here. I use similar instruments here, except aggregated to the category-week level, to address any endogeneity. The number of videos within a category are treated as exogenous to demand taste shocks.

7 Structural Model

This section reports the results of the structural model estimation. First, I discuss the model estimates and provide an interpretation of the parameters. Second, I use the model estimates to provide counterfactual predictions of video consumption under a range of changes to ad load, ad quality and video production. The results allow us assess DoJ's claims of consumer harm and to quantify the potential consequences of antitrust interventions in the upstream advertising market.

7.1 Model Estimates

The raw estimates of structural model (described in [Section \(6\)](#)) are presented in [Table 8](#). These estimates, discussed below, form the basis for the counterfactual predictions.

Table 8: Structural Parameter Estimates

Demand Side			Supply Side		
Ad-roll load	-0.579441	***	Ad-roll earnings	0.002100	***
Sponsor-read load	-0.374587	*	Sponsor-read earnings	0.002172	*
# of videos	0.000013	***			

The consumer ad-attention sensitivity estimates appear sensible. Importantly, we observe the same ordering in net distaste for ad-roll and sponsorship as in the model-free results, which we interpret as sponsorship quality being higher than ad-roll quality. The model results imply consumers’ taste for having a large number of videos within a category is low. This reflects the fact that, within the data, weeks with a large number of available videos do not have a correlated increase in viewership. In other words, viewers are time constrained and have sufficient variety.

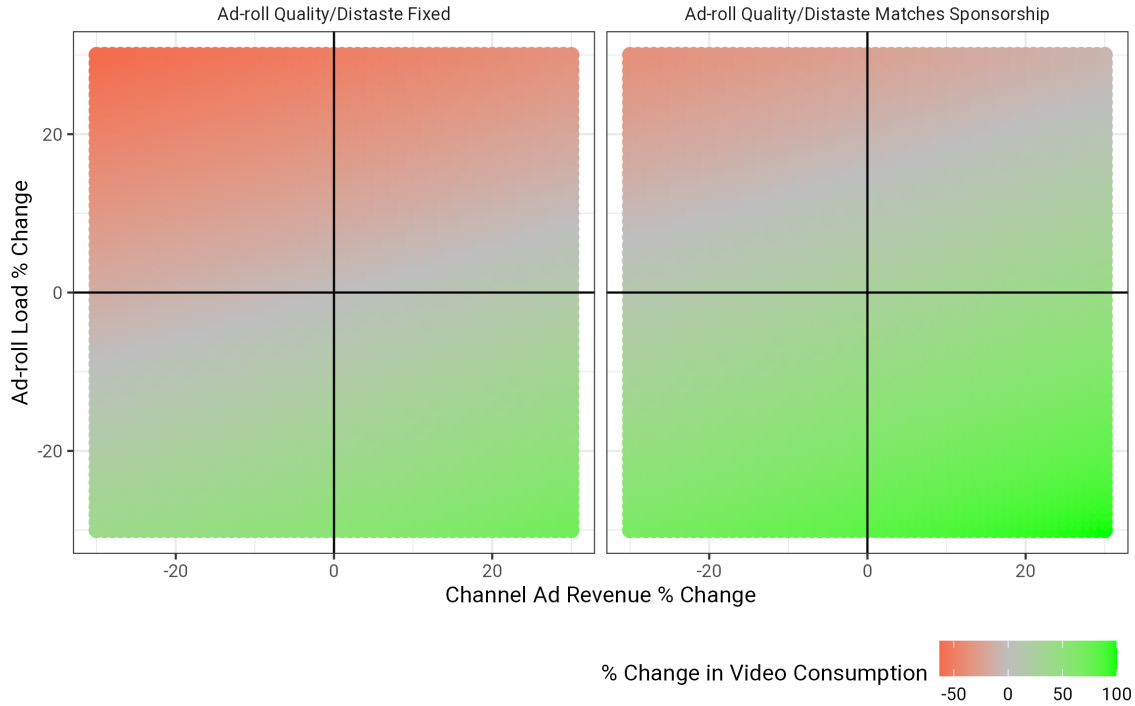
The supply side estimates also appear plausible. Video creators treat ad-roll earnings and sponsor-read earnings fairly similarly, possibly reflecting an efficient mix of ad-roll and sponsorship revenue in the market.

7.2 Counterfactual Video Consumption

This section examines the counterfactual implications of our structural model for video consumption. We focus on three key dimensions of change: (i) ad load, (ii) creator ad revenue, and (iii) ad quality. Our model allows us to predict how changes in these factors would affect YouTube video consumption, taking into account both the supply side (video creators) and the demand side (viewers).

The structural model predicts counterfactual leisure time consumption of YouTube videos. This is analogous to using a price sensitivity coefficient to convert changes in consumption to welfare dollars for goods with explicit prices. In attention markets, the analogous welfare measure would be welfare time. For now, I restrict my analysis to reporting changes in consumption, rather than taking a stance on the value of leisure time.

The counterfactual analysis holds fixed the estimated parameters of the structural model and varies creator ad revenue and ad load. Importantly, the model allows video creators to adjust their production based on their expected ad revenue, and consumers to adjust their video consumption based on the amount of content available and the associated ad load.



Notes: None

Figure 6 depicts the counterfactual video consumption for a range of outcomes. The status quo with no changes is represented by the center of the left-panel graph. Moving towards the upper-left, increases in ad-roll load (vertical axis) and decreases in creator ad revenue and video production (horizontal axis) predictably leads to lower consumption. Consumption is more sensitive to ad load changes than creator ad revenue, reflecting the intrinsic motivations and production function of creators, and consumers' taste for variety.

A move from the left-panel to the right-panel depicts how consumption changes when ad-roll ad quality increases to match that of sponsorships ad quality. With no changes to ad load or creator ad revenue, video consumption increases by 20%, reflecting the significant distaste of ad-rolls. A scenario relevant for policymakers seeking to improve competition is how increases in innovation can offset the direct costs of higher ad load. These results show a plausible innovation on the support of the data can offset a 16% increase in ad load, demonstrating the potential for better ad quality to mitigate the consumer harming aspects of improved competition.

7.2.1 Heterogeneity in Video Categories

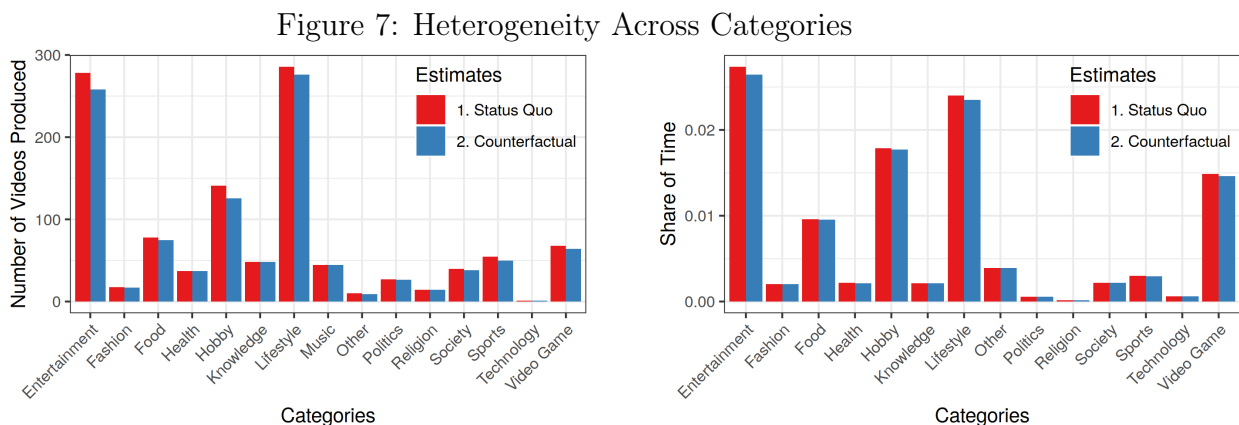
The structural model allows for heterogeneity in the intrinsic motivations and production costs of creators across different video categories. The model also allows for consumer preferences to vary across different types of videos. The interplay of these heterogeneous effects

influences the composition of video production and consumption under different counterfactual scenarios.

For the purposes of illustrating the heterogeneity in outcomes, consider a situation where ad load increases by 10% and creator ad revenue fall by 10%. Figure 7 breaks out how video production and video consumption is affected across 16 categories of videos. Production of entertainment, hobby, and lifestyle videos are most significantly impacted. This reflects a combination of the high cost of these videos (e.g., travel videos sit in the lifestyle category) and the low intrinsic motivation of video creators in these categories. In contrast, there is little change to the production of videos in the knowledge and health. A speculative explanation is that video creators in these categories are driven by strong intrinsic motivations to create such videos. To the extent that knowledge and health videos are socially desirable, this finding may alleviate concerns about antitrust action distorting such provision.

While changes in consumption largely follow the patterns observed in video production, the magnitude of change in consumption is muted due to viewers' relatively low desire for more videos within a given category. For example, despite a significant decrease in production of hobby videos, consumption of these videos sees little to no change. This can be attributed to viewers having strong preferences for this type of content and a lower ad load on hobby videos to begin with.

By explicitly allow for the heterogeneity, the model provides a more complete understanding of the complex interplay between ad load, creator ad revenue, ad quality, and video consumption. These results are relevant for policy-makers seeking to promote competition and innovation in the digital advertising market in a way that maintains a diverse range of video content for consumers.



Notes: None

8 Conclusion

The effects of antitrust actions on ad-funded attention goods are complex due to the multi-sided nature of these markets. The Department of Justice’s lawsuits against Google aim to improve competition in upstream ad markets, but the downstream consequences for consumers remain unclear. This paper quantifies the potential spillovers on consumers under various scenarios stemming from successful antitrust action.

I address a key gap in the attention markets literature by providing new estimates of ad-attention elasticities, which are crucial for calculating consumer welfare in this context. We exploit ad-roll discontinuities and the exogeneity of advertiser budgets on YouTube to estimate causal ad-attention elasticities of demand. Notably, our findings suggest that sponsorships are perceived as higher quality than ad-rolls, lending credence to the Department of Justice’s concerns about stifled innovation.

While increased competition in ad markets could potentially lower ad prices and increase ad quantities, this may lead to higher ad load, diminishing consumer welfare. Additionally, reduced ad revenue could lead to fewer ad-funded products, decreasing consumer choice. Conversely, increased competition could unlock previously stifled innovation, leading to higher-quality ads and improved consumer welfare.

Our analysis demonstrates the potential trade-offs of antitrust action in ad-funded attention markets. Using a simple structural model, we show that a plausible innovation gain, which reduces consumer aversion to ad-rolls, could offset a 16% increase in ad load. This highlights the importance of considering both the direct and indirect effects of antitrust actions on consumer welfare in multi-sided markets.

References

- Calvano, Emilio and Michele Polo**, “Market power, competition and innovation in digital markets: A survey,” *Information Economics and Policy*, March 2021, *54*, 100853.
- Casner, Ben and Tat-How Teh**, “Content-Hosting Platforms: Discovery, Membership, or Both?,” January 2023.
- Choi, Hana, Carl F. Mela, Santiago R. Balseiro, and Adam Leary**, “Online Display Advertising Markets: A Literature Review and Future Directions,” *Information Systems Research*, June 2020, *31* (2), 556–575.
- El-Komboz, Lena Abou, Anna Kerkhof, and Johannes Loh**, “Platform Partnership Programs and Content Supply: Evidence from the Youtube Adpocalypse,” 2023.
- Evans, David**, “The Antitrust Economics of Multi-Sided Platform Markets,” *Yale Journal on Regulation*, January 2003.
- Evans, David S. and Richard Schmalensee**, “The Antitrust Analysis of Multi-Sided Platform Businesses,” January 2013.
- Goli, Ali, David Reiley, and Hongkai Zhang**, “Personalizing Ad Load to Optimize Subscription and Ad Revenues: Product Strategies Constructed from Experiments on Pandora,” June 2021.
- Gordon, Brett R., Kinshuk Jerath, Zsolt Katona, Sridhar Narayanan, Jiwoong Shin, and Kenneth C. Wilbur**, “Inefficiencies in Digital Advertising Markets,” *Journal of Marketing*, January 2021, *85* (1), 7–25.
- Huang, Jason, David Reiley, and Nick Riabov**, “Measuring Consumer Sensitivity to Audio Advertising: A Field Experiment on Pandora Internet Radio,” April 2018.
- Johnson, Garrett, Tesary Lin, James C. Cooper, and Liang Zhong**, “COPPAocalypse? The Youtube Settlement’s Impact on Kids Content,” March 2024.
- Kapoor, Anuj, Sridhar Narayanan, and Amitt Sharma**, “Does Emotional Matching Between Video Ads and Content Lead to Better Engagement: Evidence from a Large-Scale Field Experiment,” October 2022.
- Kerkhof, Anna**, “Advertising and Content Differentiation: Evidence from Youtube,” *The Economic Journal*, May 2024, p. ueae043.

- Li, Nan, Avery Haviv, and Mitchell J. Lovett**, “Opposing Influences of YouTube Influencers: Purchase and Usage Effects in the Video Game Industry,” March 2024.
- Sahni, Navdeep S. and Harikesh S. Nair**, “Sponsorship Disclosure and Consumer Deception: Experimental Evidence from Native Advertising in Mobile Search,” *Marketing Science*, January 2020, *39* (1), 5–32.
- Shapiro, Bradley T., G  nter J. Hitsch, and Anna E. Tuchman**, “TV Advertising Effectiveness and Profitability: Generalizable Results From 288 Brands,” *Econometrica*, 2021, *89* (4), 1855–1879.
- Tang, Qian, Bin Gu, and Andrew B. Whinston**, “Content Contribution for Revenue Sharing and Reputation in Social Media: A Dynamic Structural Model,” *Journal of Management Information Systems*, October 2012, *29* (2), 41–76.
- Wilbur, Kenneth C.**, “A Two-Sided, Empirical Model of Television Advertising and Viewing Markets,” *Marketing Science*, May 2008, *27* (3), 356–378.
- Yoganarasimhan, Hema**, “Impact of social network structure on content propagation: A study using YouTube data,” *Quantitative Marketing and Economics*, March 2012, *10* (1), 111–150.

A [Appendix]

[Appendix]