ABSTRACT

YOUTUBE'S ADPOCALYPSE: PATREON'S PERSPECTVE

by Sidney Marie Hutson

From its inception, YouTube has grown tremendously, and with it, the scale of its advertisers. Throughout the years to keep this immense revenue stream, YouTube has continued to update their content guidelines and restrictions to appear more ad friendly. With more restrictions however, come negative impacts on the earnings potential for content creators. Using a difference-in-differences specification we are able to pinpoint four instances where policies were updated and had significant impacts on creators. We find that in order to supplement lost ad-revenue from YouTube, content creators have migrated some of their focus to Patreon resulting in increases of patrons of up to ~27 more compared to those who did not experience the YouTube shock (i.e., result for a policy change in April 2017, controlling for several fixed effects and the choice to keep earnings private). In our sample the average patrons per creator is ~13, emphasizing the significance of this change.

YOUTUBE'S ADPOCALYPSE: PATREON'S PERSPECTIVE

A Thesis

Submitted to the

Faculty of Miami University

in partial fulfillment of

the requirements for the degree of

Master of Arts

by

Sidney Marie Hutson

Miami University

Oxford, Ohio

2021

Advisor: Mark Tremblay

Reader: Deborah Fletcher

Reader: Riley Acton

This Thesis titled

YOUTUBE'S ADPOCALYPSE: PATREON'S PERSPECTIVE

by

Sidney Marie Hutson

has been approved for publication by

The Farmer School of Business

and

Department of Economics

____

Mark Tremblay

____

Deborah Fletcher

____

Riley Acton

Table of Contents

See an equation illustrating our empirical strategy on page 14.

List of Tables

List of Figures

Dedication

This paper is dedicated to my family (including Patrick) for always encouraging me to reach for the stars. I love you always.

Acknowledgements

I want to thank Dr. Mark Tremblay for his guidance during this project and allowing me the opportunity to conduct economic research first-hand. Additionally, Dr. Deborah Fletcher and Dr. Riley Acton for their tremendous feedback.

Beyond my advisor and committee, I want to thank the Economics Department at Miami University for instilling a sense of curiosity within me and the tools to above all, think critically.

1 Introduction

Over the course of the past 15 years, the introduction of YouTube has in essence created a new job title—"YouTuber." This phenomenon of posting "sit-down" videos, vlogs, mukbangs, and everything in-between has created an entirely new industry for "content creators" to make their living off of, and a lot of money at that.

But as quickly as YouTube has grown, so has the scale of its biggest advertisers, returning immense profits for the platform. In efforts to keep these companies and ad revenue, YouTube has gone through many transition phases of updating and intensifying content guidelines and restrictions. Through trying to make YouTube's image more advertiser-friendly, several YouTubers have suffered the consequences through demonetization, the act of YouTube's algorithm rendering a video or segment of a video, not ad friendly. Thus, relinquishing the creator's potential to receive profits from their video which is the main way these creators are able to have YouTube be their full-time job. Although there are instances where a video's content is truly not ad or family-friendly, there are many more where the algorithm glitches and a minor error in editing deems an entire video un-advertisable.

YouTube's strategy to keep advertisers comes after a rude awakening during the so-called "Adpocalypse" starting in 2016 in which big advertisers like Starbucks, PepsiCo, and Verizon pulled ads back to 70% of their previous expenditure (Madio & Quinn 2020). The justification given by the Association of National Advertisers is that in the potential light of any scandals or non-family friendly content being associated with brands, "reputation...[could] be damaged or severely disrupted" (ANA 2017). However, YouTube's efforts to keep the peace with brands have caused a commotion amongst creators—those most at-risk for detrimental decreases in revenue.

Although YouTube's data is notoriously private, in this paper we are able to observe the impacts of these "shocks" (instances of YouTube changing ad-friendly content guidelines) on content creators though the use of data from Graphtreon. This Graphtreon data captures data from

Patreon, a “neighboring” platform to YouTube, that several creators have resorted to using in hopes to supplement the lost income.

As previously mentioned, YouTube provides income to creators through ad-cents on videos which is a sum of money per view of ad. Or also the occasional direct sponsorship from brand to creator who promotes a product within their video. In contrast, Patreon uses a subscription profit model in which their equivalent of YouTube subscribers, “patrons,” pay a monthly subscription fee to a creator to engage with or view their content. It is up to the creator’s discretion how much their monthly fee is, and several opt to have a tiered subscription fee in which patrons can pledge more per month to gain more access.

The Graphtreon data contains information on all creators on Patreon over the course of three years, April 2016 to April 2019, including the number of YouTube videos, views, and subscribers the person has since creators typically link the two channels (through personal promotion) to allow for maximum exposure to followers. Using this information, we employ a difference-in-differences estimation strategy among four YouTube shocks to illustrate the impacts the content creators face.

We find that across a variety of specifications, we see a surge of activity on Patreon, a platform that YouTube would view as a competitor, following the policy changes on YouTube.

1.1 Related Literature

In general, literature surrounding user-generated content (YouTube and many other platforms’ model), has mainly focused on sites’ ability to provide the world with news and other types of content at the click of a button (Yildirim et al. 2013; Zhang & Sarvary 2014; Luca 2015; de Corniere & Sarvary 2020; Madio & Quinn 2020). This paper’s contribution on this front is how the Adpocalypse forces content creators to leave the site and with it their ability to spread information and opinions to the public.

Another aspect of the literature that is relevant to this paper is media bias which has been broken down between supply and demand side effects. Supply-side research on media bias primarily refers to private companies or "big brands" or government/lobbyist pressures (Ellman & Germano 2009; Besley & Prat 2006). In contrast, demand-side research on media bias refers to the innate biases that the audience holds and precludes them from truly digesting the information set before them to its fullest extent (Gentzkow & Shapiro 2006; Mullainathan & Shleifer 2005; Xiang & Sarvary 2007; Gal-Or et al. 2012). The Adpocalypse is an overt example of media bias at play. This paper serves as a step towards quantifying censorship as it can impact those who make YouTube their full-time means of income.

Prior economics literature surrounding YouTube specifically has mainly focused on advertising and content differentiation on the website itself (Kerkof 2020). Thus, work has been done to use estimation strategies to capture the algorithm, but not much to quantify how large the Adpocalypse impacts are on content creators themselves at an individual level.

Since YouTube has become an integral part of so many facets of the internet and companies alike (Stanford 2018), we hope to contribute to the literature a cautionary tale of content creators being forced to migrate to other platforms, primarily Patreon, and what that could mean for all parties involved.

2 Data

The data we use comes from a website called Graphtreon which has collected data on Patreon from March 2015 to February 2021 (at the time of this paper's completion) for all creators on the platform who had 1+ patrons for the month specified. Its server generates CSV files for each month pre- and post-monthly payment processing has occurred for users. The process to start tracking creators took a little while to get up and running so it wasn't until March 1st, 2016 that Graphtreon was able to collect data on ~99.9% of Patreon content creators (again, who had 1+ patron at the month of interest). On April 11, 2016, Graphtreon's server was powerful enough to start collecting additional social media data. For this reason, our sample has been limited to April 2016 to April 2019, a full three-year period including social media data.

The main variables of our dataset include category (of content), '*isnsfw*' (denoting "Is Not Safe for Work" content that is considered 'adult'), patrons, earnings (from Patreon, not YouTube), Twitter followers, YouTube subscribers, YouTube videos, and YouTube views. Our data points capture month-end values. However, through the data scraping process there were some instances in which there were duplicate observations for end-of-month values per a single creator. These discrepancies were small, but for the sake of consistency and simplicity, we use averages of these duplicate observations for patrons, earnings, Twitter followers, YouTube subscribers, YouTube videos, and YouTube views to create distinct "creator-at-the-end-of-the-month" measures. Average patrons/Twitter followers/YouTube subscribers/YouTube videos/YouTube views are cumulative measures by construction, whereas average earnings are per month of interest. *Do note that we explored regressions in which we dropped any content creators from the sample that possessed a "duplicate" value, but the estimates were nearly identical and thus we continued to use these "average" values.*

Below is a summary statistics table of our sample. As you can see, '*average patrons per creator*' contains the most observations at 8,291,367 since this measure was guaranteed every month of observation. When rectangularizing the dataset, if the number of patrons was missing, we were able to assume that the creator had zero patrons in that month, thus creating a strongly balanced panel dataset. Additionally, on Patreon content creators have the choice to either display or hide their earnings from the public whereas the number of patrons is always public knowledge. As you can see in the table below, there are a lot less observations for earnings on Patreon than there are for number of patrons, meaning people are choosing not to display their earnings on their page. The number of observations for the social media measures vary dependent on whether the individual content creator uses the platform.

Table 1: Summary Statistics for Entire Sample

	Mean	Std. Dev.	Min	Max	Obs
Average Patrons Per Creator* (avg_patrons)	13.00496	146.2852	0	65848	8,291,367
Average Earnings in Dollars from Patreon Per Month** (avg_earnings)	141.9574	807.7279	0	124945	2,481,510
Average Twitter Followers Per Creator* (avg_twitterfollowers)	18944.81	847234.4	0	7.16E+07	1,797,509
Average Number of YouTube Subscribers Per Creator* (avg_youtubesubscribers)	91464.71	1254110	0	1.06E+08	1,336,057
Average Number of YouTube Videos Per Creator* (avg_youtubevideos)	271.0363	2071.001	0	339191	1,336,057
Average Number of YouTube Views Per Creator* (avg_youtubeviews)	1.90E+07	2.88E+08	0	3.19E+10	1,336,057

Notes:
*These measures represent cumulative values captured at the end of the month.
**This measure is not cumulative but is still captured at the end of the month.

3 Motivation

As mentioned in our introduction, motivation from this paper comes from the "Adpocalypse" YouTube began to experience in 2016. In efforts to combat advertisers pulling participation from the cite, YouTube altered their content guidelines as either a proactive or reactive measure (in some cases following scandal), which we refer to as a YouTube "shock" to content creators. In this section we explore the primary four shocks that occur within our sample of April 2016 to April 2019.

3.1 YouTube Shock #1: August 2016

As far back as 2012, YouTube has tightened up their content guidelines to better appeal to big advertisers. However, starting in August 2016, YouTube began to notify their creators of policy changes regarding video monetization, initializing the "Adpocalypse." Now, if YouTube's algorithm categorized a single part of a creators' video as not adhering to "advertiser-friendly content guidelines," the video would automatically become ineligible to receive ad revenue. Big names on YouTube such as Philip DeFranco started to take notice. YouTube received backlash but stuck to their guidelines as the ad revenue was more important (Pottinger 2018; Business Insider Nederland 2020).

3.2 YouTube Shock #2: April 2017

In April 2017, YouTube changed their algorithm in that it put even stricter policies on content worthy of ads and raised the requirements of eligibility for receiving income from videos. This was due to the massive ad boycott following hate speech videos released on YouTube (Pottinger 2018; Business Insider Nederland 2020).

3.3 YouTube Shock #3: November 2017

In November 2017, YouTube's algorithm went under backfire yet again when content deemed "family-friendly" contained disturbing and abusive content. Advertisers appalled by the videos playing their ads, pulled back. As a solution, YouTube decided to update its policies around age-restricted content (Pottinger 2018; Business Insider Nederland 2020).

3.4 YouTube Shock #4: February 2019

In February 2019, a YouTube content creator by the name of Matt Watson exposed a "softcore pedophilia ring" found in the comment section of YouTube videos featuring children causing even more advertisers to pull content for fear of being associated with the scandal. This time YouTube responded by automatically disabling the comments on videos featuring children (Pottinger 2018; Business Insider Nederland 2020).

3.5 Overall Impact on Content Creators

These shocks could have potentially made it harder for content creators to make a living off of YouTube, specifically through obtaining ad revenue. Which, apart from sponsorships, is the primary form of income for many YouTubers. Migrating to another platform like Patreon is a very accessible option for content creators to make supplemental income.

Patreon is a unique platform in that creators do not make money off of ad revenue, but rather through their followers, called “patrons” paying a subscription fee to view their content. Most content creators on the platform have a tiered subscription platform with a lot of variation between offerings. Some creators allow their patrons to pay a fee per video, or to pay a monthly subscription fee with varying levels of access based on the amount they pay.

Since content creators can operate on both YouTube and Patreon simultaneously (promoting each page on the other platform), we would expect during the months of (and after) these shocks to see a surge in Patreon activity and revenue, as mentioned before, as an effort by creators to make additional income in a “slow” period on YouTube where (potentially) many of their videos would have been demonetized via YouTube’s algorithm changes.

4 Conceptual Framework: DID Trend Assumption

In order to quantify the effects of the Adpocalypse, we employ a difference-in-differences framework in which we compare those on Patreon who would have experienced the Adpocalypse via having a YouTube channel with enough videos and subscribers to obtain monetization to those creators on Patreon that do not engage with YouTube or would not have felt the effects of the YouTube shocks. This comparison leans heavily on the assumption that had the treatment of the shocks not occurred, the treated creators would have continued to display a parallel trend to those who did not. Thus, our DID estimators are measuring the treatment group's deviation from a parallel trend to that of the control. Naturally, those who experienced the shocks are bigger content creators spanning both platforms with more of a following and larger earnings, however, as exhibited in Figure 1 and Figure 2, the two groups do follow trends that we would expect them to.

In Figures 1A and 1C we are unable to see a definite spike in the number of patrons in the treatment group compared to the control during the shocks (which are a month long) since patrons are already trending upwards prior to the shock. Though, in Figures 1B and 1D, we do see an elevated slope in patrons for the treatment compared to the control. This is especially apparent in Figure 1D where the number of patrons is trending downwards for the treatment prior to YouTube shock #4.

Figures 2A and 2B tell a similar story in that post- the windows of the shocks, the treatment groups' earnings do increase from their new followers, as compared to the control group. Figures 2C and 2D don't quite illustrate as much of a jump as we'd like to see.

Overall, even though these are basic average trends with nothing fancy, we do see that our story has some merit. It is also important to note that in our sample Patreon experiences 3 negative shocks that impact the entire platform (both treatment and control groups) that would decrease their content creators' number of patrons and earnings. This happens in t = 20 (Patreon altered their fee structure), t = 26 (the EU GDPR) and t = 32 (Patreon removed certain creators). These basic graphs do display those dips as expected (*see pg. 20 for further descriptions of the shocks*).

Figure 1: Patrons of Creators Who Experienced A Shock vs. Did Not

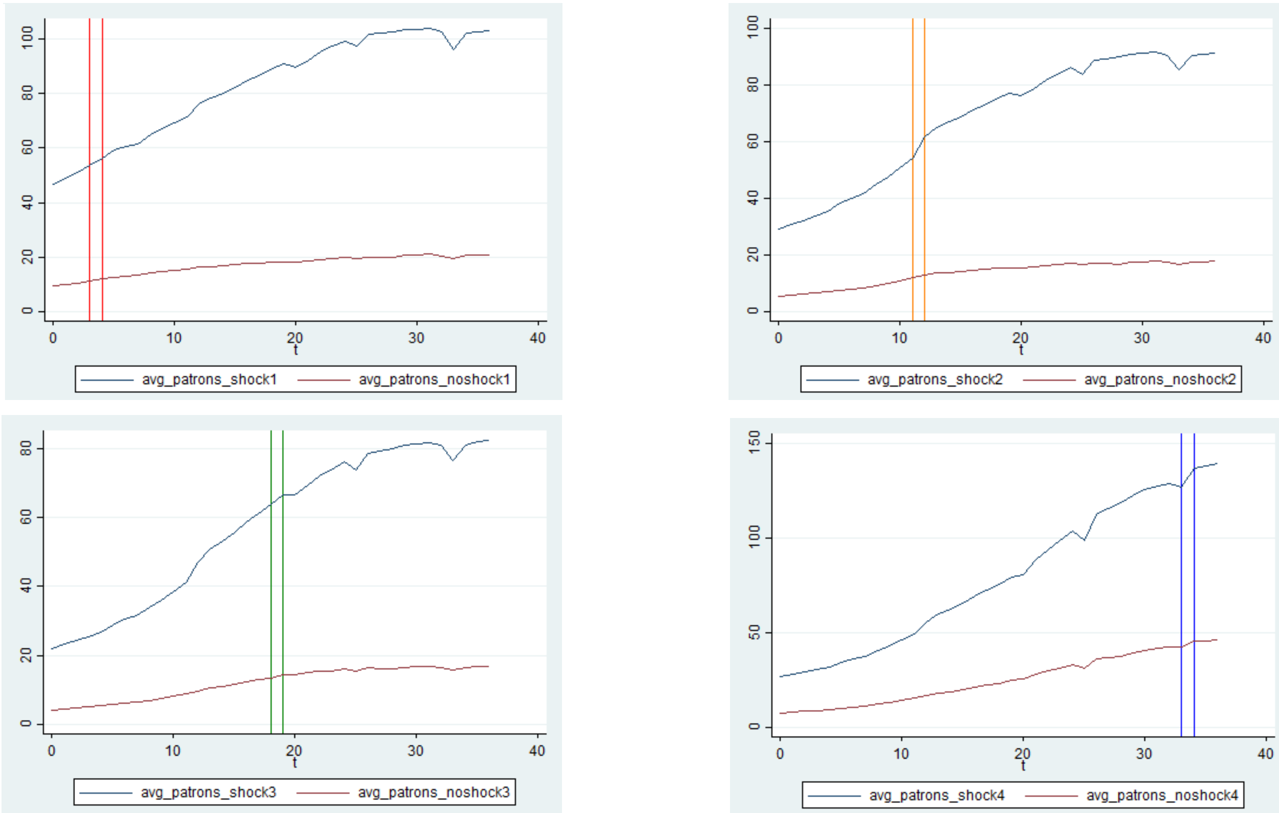

1A 1B 1C 1D

Figure 2: Earnings from Patreon of Creators Who Experienced A Shock vs. Did Not

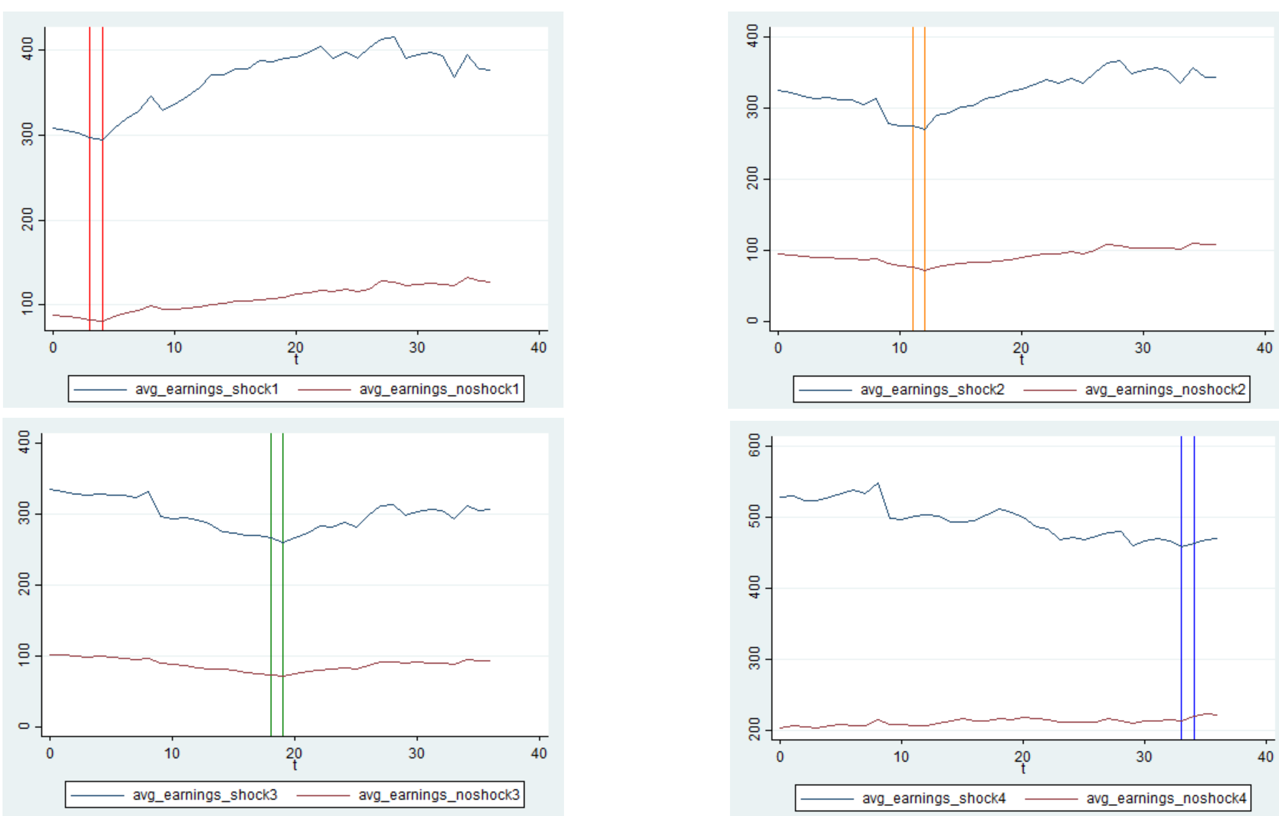

2A 2B 2C 2D

As a way to further contextualize the treated vs. control groups, we additionally created histograms illustrating the number of YouTube subscribers a content creator has. The numbers labeling the histograms correspond to whether it is analyzing YouTube shock #1, 2, 3, or 4. 'A' refers the treated group whereas 'B' refers to the control group. Thus, histograms #1A and #1B (for example) are comparing the control vs. treated for YouTube shock #1.

As expected, the treated groups tend to have a lot more YouTube subscribers and range from ~75,000 subscribers to ~250,000 depending on the shock. In contrast, the treatment groups appear to have less than ~20,000 subscribers overall, clustering more towards the ~0 to ~5,000 range.

These stark contrasts beg the question of whether treatment varies by YouTube subscribers and whether we should take a look at the content creators 'at the margin' of our analysis. In our main regressions, the only YouTube subscriber condition we place on our treatment group is having more than 1,000 subscribers. Though, we come back to this potential issue in a robustness check limiting treatment groups to content creators with 1,000 to 2,000 subscribers (*see Section 7.4 for further details*).

Figure 3: Histograms of YouTube Subscribers of Creators Who Experienced A Shock vs. Did Not

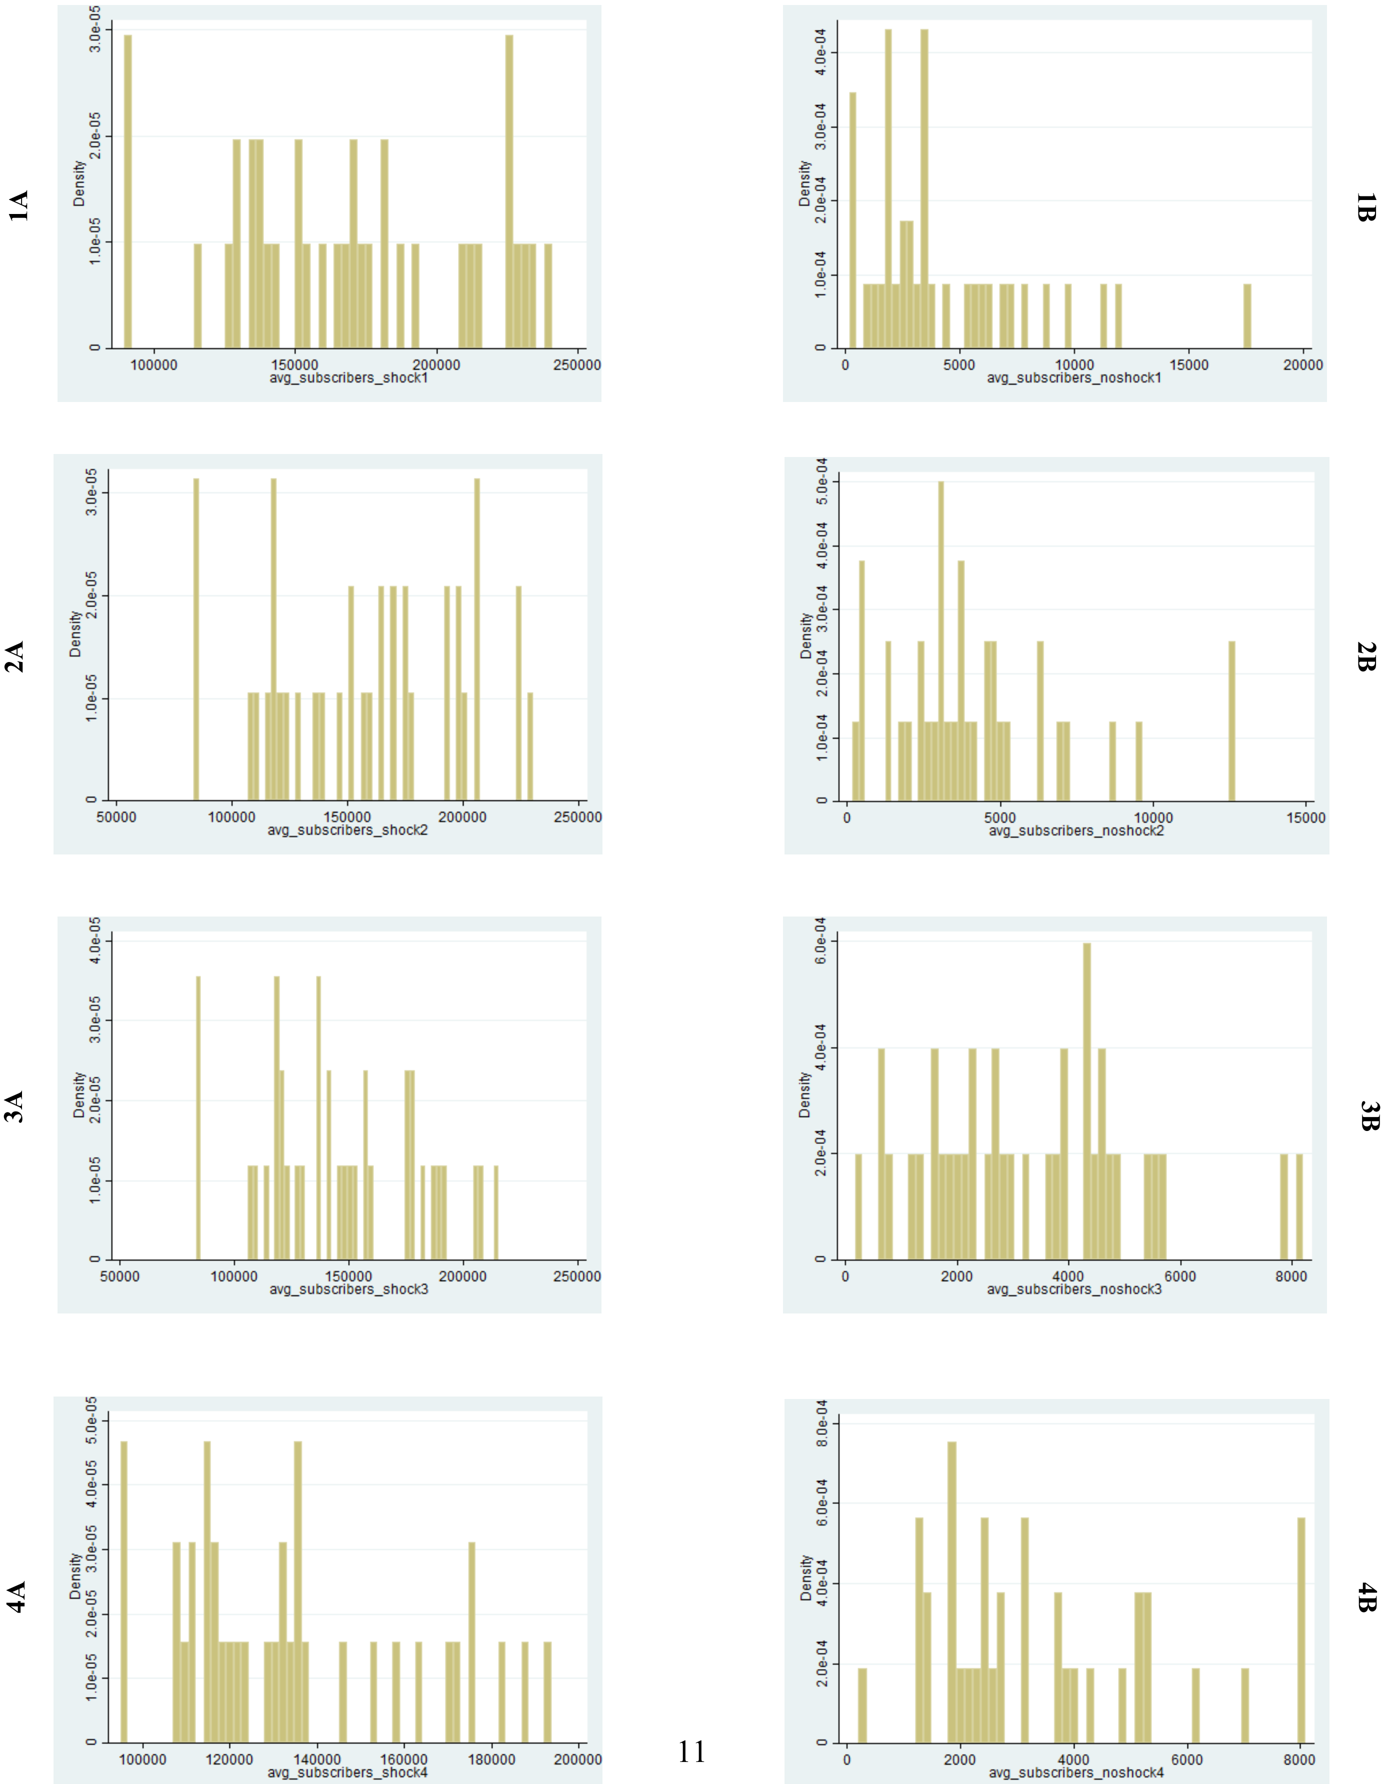

5 Estimation: Measuring the YouTube Shocks

To measure the impact of these YouTube shocks on content creators, we constructed variables called '*interaction_shock#YT*' which measures the interaction between categorical variables '*after_shock#YT*' and '*experienced_shock#YT*'. Note that "#" referrers to the number of the shock (1-4) and "YT" refers to YouTube.

The variable '*after_shock#YT*' takes on a value of 1 when our time indicator (t) is equal to or greater than the month in which this shock occurred, or a value of zero otherwise. Using a "equal to or greater than" condition allows us to pick up on late adoption or creator reactions to the policy shocks.

Generated by unique id for content creator (id), the treatment variable '*experienced_shock#YT*' takes on a value of 1 if at the time of the shock a content creator has 1,000 or more YouTube subscribers and at least one YouTube video. Or a value of zero otherwise. This is because YouTube's policy on video monetization is that at the 1,000-subscriber accomplishment, a content creator may begin to monetize their content to collect ad revenue (Kumar 2020). Having the number of videos be at least one controls for the possibility that a channel deletes all of their videos but not their channel. In that potential case, subscribers would linger but the content creator would not feel the effect of the shock and thus should not be categorized as a part of the treatment group.

Essentially the interaction/difference-in-differences estimator '*interaction_shock#YT*' will "turn-on" when both categorical variables ('*after_shock#YT*' & '*experienced_shock#YT*') have a value of one.

5.1 Naming Scheme for the YouTube Shocks & Treated vs. Control Distribution

For sake of clarity within our regression output, take note of our shock variables naming scheme:

Shock #	Time	'*t*'	Post-Treatment Period	Treatment Indicator	DID Interaction Term
1	August 2016	t = 4	after_shock1YT	experienced_shock1YT	interaction_shock1YT
2	April 2017	t = 12	after_shock2YT	experienced_shock2YT	interaction_shock2YT
3	November 2017	t = 19	after_shock3YT	experienced_shock3YT	interaction_shock3YT
4	February 2019	t = 34	after_shock4YT	experienced_shock4YT	interaction_shock4YT

Additionally, see Table 2 below for a breakdown of observations for the '*experienced_shocktYT*' and '*interaction_shock#YT*' variables. Take note of how the treatment vs. control groups for the '*experienced_shocktYT*' are fairly evenly distributed.

Table 2: Treated vs. Control Distribution

experienced_shock1YT	Freq.	Percent
Control	410,219	54.04%
Treated	348,910	45.96%
Total Observations	759,129	100%

interaction_shock1YT	Freq.	Percent
DID Term Activated	447,939	59.01%
DID Term Deactivated	311,190	40.99%
Total Observations	759,129	100%

experienced_shock2YT	Freq.	Percent
Control	670,810	54.37%
Treated	562,881	45.63%
Total Observations	1,233,691	100%

interaction_shock2YT	Freq.	Percent
DID Term Activated	853,366	69.17%
DID Term Deactivated	380,325	30.83%
Total Observations	1,233,691	100%

experienced_shock3YT	Freq.	Percent
Control	835,645	53.18%
Treated	735,634	46.82%
Total Observations	1,571,279	100%

interaction_shock3YT	Freq.	Percent
DID Term Activated	1,213,403	77.22%
DID Term Deactivated	357,876	22.78%
Total Observations	1,571,279	100%

experienced_shock4YT	Freq.	Percent
Control	346,246	37.17%
Treated	585,229	62.83%
Total Observations	931,475	100%

interaction_shock4YT	Freq.	Percent
DID Term Activated	884,024	94.91%
DID Term Deactivated	47,451	50.90%
Total Observations	931,475	100%

Below is an equation illustrating our empirical strategy:

$$\boldsymbol{Outcome} = \boldsymbol{\beta_0} + \boldsymbol{\beta_1 A_t} + \boldsymbol{\beta_2 T_i} + \boldsymbol{\beta_3 (A_t * T_i)} + \boldsymbol{\beta_n X_{itc}} + \boldsymbol{\varepsilon_{itc}}$$

Where:

- $\boldsymbol{Outcome}$ = cumulative patrons or monthly earnings on Patreon, depending on the regression
- $\boldsymbol{A_t}$ = after_shock1YT, after_shock2YT, after_shock3YT, after_shock4YT
- $\boldsymbol{T_i}$ = experienced_shock1YT, experienced_shock2YT, experienced_shock3YT, experienced_shock4YT
- $\boldsymbol{A_t * T_i}$ = interaction_shock1YT, interaction_shock2YT, interaction_shock3YT, interaction_shock4YT
- $\boldsymbol{\beta_n X_{itc}}$ = all of our additional controls including fixed effects for Patreon shocks, category of content, distinct month-year, whether the content is "safe for work," controls for average Twitter followers, average YouTube subscribers, average YouTube videos, average YouTube views (and fixed effects for each specific content creator in some regressions)

6 Results

In this section we explore our main regression results.

6.1 Patrons Regression Results

To best understand the YouTube shocks' impact on content creators we ran a regression on '*Average Monthly Patrons*'. As illustrated in Table 3, Specification (1), '*avg_earnings*,' '*interaction_shock1YT*,' '*interaction_shock2YT*', and '*interaction_shock3YT*' all have a significant effect at the 99.9+% confidence level. This result can be interpreted as a one-hundred dollar increase in earnings is associated with a ~17.4 increase in the average number of patrons a creator has at any given month, ceteris paribus. Additionally, one of the difference-in-differences estimators (e.g., '*interaction_shock1YT*') can be interpreted as comparing average values from the post-period starting in August 2016 to the pre-period, content creators who experienced shock #1 vs. those who did not have 4.615 more patrons, ceteris paribus. Fixed effects are included and described in the notes on Table 3.

Next, we omit the explanatory variable '*avg_earnings*' in Table 3, Specification (2), and find a significant change between the magnitudes and statistical significance of the coefficients. Now, one of the difference-in-differences estimators (again, e.g., '*interaction_shock1YT*') can be interpreted as comparing average values from the post-period starting in August 2016 to the pre-period, content creators who experienced shock #1 vs. those who did not have ~14.04 more patrons, ceteris paribus. Again, fixed effects are included and described in the notes on Table 3. This significant shift leads us to believe that the variable '*avg_earnings*' may possess some sort of bias which we will later explore in a robustness check and that we should omit earnings from our best regressions.

For now, in order to rule out the possibility that this shift is due to a sample selection issue, in Specification (5) we excluded observations that did not report earnings (i.e., they would have been omitted in Specification (1)) and run a regression without earnings as a right-hand side variable (Specification (6) is the same regression as Specification (5) but in log-terms). Still, we see significant changes in magnitude between our DID estimators comparing Specification (1) to (5) leading us to believe this is not a sample issue.

Additionally, in Table 3 Specification (3), we added fixed effects for each specific content creator. The magnitude of the coefficients and their statistical significance do not change very much from Specifications (2) and (3), but Specification (3) does control for the most variation to give us more accurate measures of the shock. Specification (4) is simply Specification (3) in log terms for patrons. Now, for example, DID estimator '*interaction_shock2YT*' can be interpreted as comparing average values from the post-period starting in April 2017 to the pre-period, content creators who experienced shock #2 vs. those who did not have ~6.34% more patrons, ceteris paribus. **Specifications (2), (3), & (4) are our preferred model for the entirety of this paper**.

The main takeaway of the Table 3 specifications is that those content creators who experienced the YouTube shocks did experience a flux of support of patrons on the Patreon site while going through tough times on their YouTube channels.

Table 3: Patrons Regression Results

	(1) Average Monthly Patrons	(2) Average Monthly Patrons		(3) Average Monthly Patrons	(4) log(Average Monthly Patrons)	(5) Average Monthly Patrons	(6) log(Average Monthly Patrons)
		Preferred Model		***Preferred Model***	***Preferred Model***	*Regressing onto group that reports earnings.*	*Regressing onto group that reports earnings.*
avg_earnings	0.174***						
	(0.0154)						
after_shock1YT	4.582***	20.89***		19.69***	0.376***	19.77***	0.382***
	(1.665)	(2.280)		(4.073)	(0.0152)	(2.571)	(0.0140)
experienced_shock1YT	11.42*	56.41***		-	-	-	-
	(6.896)	(11.44)					
interaction_shock1YT	**4.615*** **	14.04***		13.31***	-0.0148	**10.71*** **	-0.0107
	(1.602)	(3.586)		(3.178)	(0.0122)	(2.027)	(0.0112)
after_shock2YT	1.894**	8.036***		7.568***	0.174***	6.678***	0.156***
	(0.848)	(2.225)		(1.968)	(0.00778)	(1.453)	(0.00776)
experienced_shock2YT	-0.312	-13.07		-	-	-	-
	(6.135)	(18.06)					
interaction_shock2YT	**9.971*** **	28.43***		27.58***	0.0634***	**21.34*** **	0.0497***
	(1.553)	(4.315)		(1.583)	(0.00623)	(1.067)	(0.00617)
after_shock3YT	-0.550	12.67***		12.76***	0.109***	0.823	0.0717***
	(1.217)	(2.955)		(3.052)	(0.00832)	(2.181)	(0.00938)
experienced_shock3YT	-5.719	-10.16		-	-	-	-
	(5.333)	(21.36)					
interaction_shock3YT	**6.228*** **	26.28***		25.00***	0.0395***	**12.45*** **	0.0366***
	(1.473)	(4.665)		(1.245)	(0.00464)	(0.950)	(0.00501)
after_shock4YT	-0.858	-11.87***		-12.14***	-0.0247**	-3.518	-0.0261**
	(1.007)	(2.455)		(3.890)	(0.0113)	(2.572)	(0.0130)
experienced_shock4YT	5.396	49.00***		-	-	-	-
	(5.639)	(17.35)					
interaction_shock4YT	**0.991**	15.36***		15.34***	0.00311	**0.668**	-0.00780
	(1.728)	(3.459)		(2.773)	(0.00886)	(1.890)	(0.0102)
Constant	-0.511	15.76		52.90***	3.012***	66.77***	2.952***
	(5.310)	(12.20)		(7.843)	(0.0870)	(5.594)	(0.0968)
Observations	172,024	205,248		205,242	205,242	172,005	172,005
Number of id	6,061	6,126	R-squared	0.894	0.898	0.938	0.905

Robust standard errors in parentheses: ***p<0.01, **p<0.05, *p<0.1

Notes: Fixed effects for Patreon shocks, category of content, distinct month-year, whether the content is "safe for work”, average Twitter followers, average YouTube subscribers, average YouTube videos, and average YouTube views have been included in these regressions but omitted in this table for the sake of simplicity. Specifications (3), (4), (5), and (6) additionally contain fixed effects for each specific content creator.

6.2 Earnings from Patreon Regression Results

Another way to capture the effects of the four YouTube shocks is to take a look at how content creators' earnings on Patreon varied during these time periods. To do this, we ran a regression on 'Average Monthly Earnings on Patreon.' As illustrated in Table 4, Specification (1), '*avg_patrons*' has a significant effect at the 99.9+% confidence level whereas '*interaction_shock2YT*' and '*interaction_shock3YT*' only have significant effects at the 95% confidence level and difference-in-differences estimators '*interaction_shock1YT*' and '*interaction_shock4YT*' do not have a significant effect at all.

These results can be interpreted as a content creator having one additional patron (paid supporter) on Patreon is associated with an ~$3.89 increase in monthly average earnings, ceteris paribus. Also, the DID (difference-in-differences) estimator '*interaction_shock2YT*' can be interpreted as comparing average values from the post-period starting in April 2017 to the pre-period, content creators who experienced shock #2 vs. those who did not make ~$17.62 less, ceteris paribus. This result is not expected. When a content creator has more patrons, they should be making more money... further indicating that there is an issue at play. Fixed effects are included and described in the notes on Table 4.

Next in Table 4, Specification (2), we omit '*avg_patrons*' and three out of four of the DID estimators become statistically significant at the 99.9+% confidence level. Moreover, they all have positive effects on earnings, which aligns with our economic theory. Now, one of the difference-in-difference estimators (e.g., '*interaction_shock1YT*') can be interpreted as comparing average values from the post-period starting in August 2016 to the pre-period, content creators who experienced shock #1 vs. those who did not make ~$35.79 more on Patreon, ceteris paribus. Again, fixed effects are included and described in the notes on Table 4. This flip in significance further leads us to believe that there are biases that we will need to explore later. But for now, it is encouraging to see that more patrons mean more earnings, as expected.

Similar to Table 3, Table 4 includes a third specification that uses fixed effects for each specific content creator, but the change in results is very minor. Plus, Specification (4) is simply

Specification (3) in log-terms. Overall using earnings as our outcome variable biases our results so we prefer to use patrons (as in Table 3).

Additionally, we ran Specification (5) to see how monthly earnings per patron changed due to the YouTube shocks. We find that for YouTube shock #1, comparing average values from the post-period starting in April 2017 to the pre-period, content creators who experienced shock #2 vs. those who did not make ~$0.08 more per patron, ceteris paribus. This result makes intuitive sense that those content creators that were adversely impacted by the YouTube ad shocks would have to make up supplemental income by charging potentially higher prices on their Patreon page, but this result is only significant at the 90% confidence interval and is not consistent across the other shocks. So, on average, it is not clear whether the negative shocks result in content providers increasing prices or simply trying to increase total number of patrons.

Table 4: Earnings from Patreon Regression Results

VARIABLES	(1) Average Monthly Earnings on Patreon	(2) Average Monthly Earnings on Patreon		(3) Average Monthly Earnings on Patreon	(4) log(Average Monthly Earnings on Patreon)	(5) Average Monthly Earnings on Patreon per Patron
avg_patrons	3.885***					
	(0.356)					
after_shock1YT	10.73	88.37***		88.72***	0.358***	-0.547***
	(8.251)	(8.454)		(11.96)	(0.0172)	(0.119)
experienced_shock1YT	31.67	231.3***		-	-	-
	(38.01)	(54.69)				
interaction_shock1YT	-6.369	35.79***		35.94***	-0.0325**	0.00877
	(7.489)	(11.97)		(9.301)	(0.0137)	(0.0875)
after_shock2YT	1.846	28.26***		27.99***	0.130***	-0.269***
	(4.235)	(6.721)		(6.635)	(0.0101)	(0.0575)
experienced_shock2YT	-3.687	-19.67		-	-	-
	(39.27)	(80.64)				
interaction_shock2YT	-17.62**	65.55***		66.07***	0.0327***	0.0821*
	(7.547)	(11.44)		(4.728)	(0.00788)	(0.0436)
after_shock3YT	4.645	7.671		7.473	0.0557***	-0.217***
	(6.544)	(11.07)		(11.93)	(0.0116)	(0.0580)
experienced_shock3YT	18.05	-12.57		-	-	-
	(35.08)	(85.01)				
interaction_shock3YT	-12.34**	36.81***		37.03***	0.0139**	-0.129***
	(5.963)	(11.29)		(3.985)	(0.00643)	(0.0343)
after_shock4YT	-1.317	-14.34		-14.54	-0.0265*	0.0398
	(5.887)	(10.77)		(15.09)	(0.0158)	(0.0723)
experienced_shock4YT	41.95	192.0**		-	-	-
	(33.03)	(76.80)				
interaction_shock4YT	-5.094	-3.632		-2.802	-0.0136	-0.0458
	(9.964)	(17.02)		(11.02)	(0.0124)	(0.0578)
Constant	38.94**	115.5***		353.8***	4.596***	7.116***
	(19.84)	(41.08)		(26.13)	(0.101)	(0.398)
Observations	172,024	172,024		172,005	172,003	172,005
Number of id	6,061	6,061	R-squared	0.922	0.876	0.786

Robust standard errors in parentheses: ***p<0.01, **p<0.05, *p<0.1

Notes: Fixed effects for Patreon shocks, category of content, distinct month-year, whether the content is "safe for work", average Twitter followers, average YouTube subscribers, average YouTube videos, and average YouTube views have been included in these regressions but omitted in this table for the sake of simplicity. Specifications (3), (4), and (5) additionally contain fixed effects for each specific content creator.

7 Robustness Checks

In this section we begin to improve our preferred models through various robustness checks.

7.1 Heterogeneity by Category

Within our sample, the content creators on Patreon fall into several categories to describe the type of work they produce for their viewers. Below in Table 5 is a listing of the 28 potential categories and the distribution amongst them. The difference between '*Adult Writing*' and '*Writing*,' for example, is that '*Adult Writing*' is classified as the aforementioned '*isnsfw*' ("Is Not Safe for Work") type of material.

Important things to note is that the most frequented category is "*Video*" at 26.37%. This is not surprising considering this paper discusses how YouTubers who make video content have been forced to flood to Patreon. Also, **Category #4** and **Category #5** ('*Adult Crafts & DIY*' and '*Adult Dancing & Theater*') make up only 0.03% and 0.10%, respectively, of the sample. Thus, these two categories are automatically omitted in our regressions due to their small sample size and not being able to measure their effect on our output. This omission is systematic through Stata, and does not represent the omitted category for comparison, which in this case happens to be Category #28 ('*Writing*').

We do expect that the differing categories would have a significant impact on both average patrons and average earnings a creator has. Furthermore, we would expect that the differing categories would impact our variables of interest in a heterogenous matter, thus we do control for content category as a part of our fixed effects. In essence, the more popular a category is on Patreon, the more earnings or followers a creator who makes content of that category would have.

See **Section 7.2** of this paper for each categories' fixed effect on our regression of choice, average patrons (as compared to our omitted Category #28, '*Writing*').

Table 5: Breakdown of Categories Within Sample

Type of Content ("Category")	Freq.	Percent
1) Adult Animation	64,487	0.79%
2) Adult Comics	212,893	2.61%
3) Adult Cosplay	108,310	1.33%
4) Adult Crafts & DIY	**2,750**	**0.03%**
5) Adult Dance & Theater	**7,906**	**0.10%**
6) Adult Drawing & Painting	358,821	4.40%
7) Adult Games	153,047	1.88%
8) Adult Magazine	4,576	0.06%
9) Adult Music	35,515	0.44%
10) Adult Other	297,846	3.66%
11) Adult Photography	170,322	2.09%
12) Adult Podcasts	48,290	0.59%
13) Adult Video	388,216	4.77%
14) Adult Writing	144,223	1.77%
15) Animation	76,446	0.94%
16) Comics	335,002	4.11%
17) Cosplay	80,965	0.99%
18) Crafts & DIY	23,385	0.29%
19) Dance & Theater	33,782	0.41%
20) Drawing & Painting	440,317	5.41%
21) Games	570,532	7.00%
22) Magazine	28,143	0.35%
23) Music	522,848	6.42%
24) Other	825,519	10.13%
25) Photography	80,553	0.99%
26) Podcasts	489,586	6.01%
27) Video	2,148,366	26.37%
28) Writing	493,681	6.06%
Total Observations of Categories	8,146,327	100%

Notes: The effects of category #4 (Adult Crafts & DIY) and category #5 (Adult Dance & Theater) become omitted in later regressions due to collinearity because of their very low occurrence within the sample.

7.2 Included Fixed Effects

Aside from content category fixed effects, there are other fixed effects we found important to include: isnsfw, distinct month-year fixed effects, Patreon shocks and individual content creator fixed effects. Table 6 displays three regressions that show the importance of including said fixed effects. In Table 6, Specification (1) does not contain any fixed effects, Specification (2) contains all of them with the exception of individual content creator fixed effects, Specification (3) contains all fixed effects, *including* individual content creator fixed effects, and finally Specification (4) duplicated Specification (3) but in log-terms.

Specification (2), (3), & (4) of Table 6 are duplicating Table 3 Specifications (2), (3) & (4), our preferred models. Now, with just showing the coefficients of the specific fixed effects.

7.2A Isnsfw Fixed Effect

As mentioned throughout, '*isnsfw*' has proven a significant control. Based on Table 6, Specification (4), if content is deemed "not safe for work," the average number of patrons who support that channel increase by ~16.8% (significant at the 99.9+% confidence level). Due to YouTube's Adpocalypse and banning non-ad- or family-friendly content, it aligns with our intuition that Patreon would see a surge of that type of content on their platform.

7.2B Distinct Month-Year Fixed Effects

To attempt to control for potential variation not explained by our other covariates, we create distinct '*month_yeart*' variables from '*month_year0*' to '*month_year36*.' The majority of these dummies are significant, and some are omitted due to collinearity with our other fixed time variables.

7.2C Patreon "Shocks" Fixed Effects

Aside from the Adpocalypse shocks on YouTube's side of the data, Patreon's platform also experienced some shocks to their content creators throughout the course of our sample. These shocks differ from the YouTube shocks in that they affect the entire active Patreon platform and are not specific to certain categories.

Patreon Shock #1: December 2017 (t=20)

In December 2017, Patreon altered their fee structure. In an effort to give more money to creators, Patreon increased charges to patrons (content "subscribers") to account for transaction fees rather than that burden falling onto the content creators and eating a cut of their profit.

Patreon's slogan of this new structure was that "creators get to take home exactly 95% of every pledge, with no additional fees" (Patreon 2017) by adding an additional service fee of 2.9% plus $0.35 to patrons for every pledge they have (ongoing subscriptions). A couple of weeks later however, Patreon completely scrapped this new fee structure after receiving backlash and issued an apology to content creators who ultimately lost a lot of followers and income via patrons not wanting to pay extra fees.

As seen in Table 6, Specification (4), at the 99.9+% confidence level, this shock is associated with an ~12.6% decrease in the average number of patrons a creator has, ceteris paribus.

Patreon Shock #2: June 2018 (t=26)

At the very end of May 2018, "the EU General Data Protection Regulation (GDPR), which governs how personal data of individuals in the EU may be processed and transferred, went into effect" (European Union 2020). Since Patreon is a global company, we expect that this limitation may decrease the number of patrons and earnings during this time period.

As seen in Table 6, Specification (4), at the 99.9+% confidence level, this shock is associated with an ~3.05% decrease in the average number of patrons a creator has, ceteris paribus.

Patreon Shock #3: December 2018 (t=32)

In December 2018, Patreon began tightening up on restrictions on creators on their platform. On December 6th, "Patreon kicked the anti-feminist polemic Carl Benjamin, who works under the name Sargon of Akkad, off its site for using racist language on YouTube. That same week, it removed the right-wing provocateur Milo Yiannopoulos a day after he opened an account" (Bowles 2018). This prompted some other famous Patreon creators to leave the platform, thus again having a negative impact on patrons and earnings.

As seen in Table 6, Specification (4), at the 95% confidence level, this shock is associated with an ~2.04% decrease in the average number of patrons a creator has, ceteris paribus.

Constructing the Patreon "Shocks" in Our Data

Generated by unique id for content creator (id), the dummy variables '*patreon_shock#*' ("#" being 1-3) take on a value of one if at the time of the shock, a content creator has at least one Patron pledged to their page. Having the condition of at least one Patron ensures that the content creator did feel the impact of the Patreon shock and thus we would see an effect on the number of patrons or earnings they have.

7.2D Fixed Effects for Each Individual Content Creator

Moreover, in some specifications we also included a fixed effect for each specific content creator to try to control for further variation. Including these fixed effects systematically omitted our '*experienced_shock#YT*' variables since the new individual-level fixed effects already control for people who experienced a shock vs. those who did not.

Table 6: Importance of Fixed Effects

	(1) Average Monthly Patrons	(2) Average Monthly Patrons	(3) Average Monthly Patrons	(4) log(Average Monthly Patrons)
after_shock1YT	5.884***	20.89***	19.69***	0.376***
	(1.538)	(2.280)	(4.073)	(0.0152)
experienced_shock1YT	53.49***	56.41***	-	-
	(11.38)	(11.44)		
interaction_shock1YT	13.96***	14.04***	13.31***	-0.0148
	(3.599)	(3.586)	(3.178)	(0.0122)
after_shock2YT	9.514***	8.036***	7.568***	0.174***
	(2.200)	(2.225)	(1.968)	(0.00778)
experienced_shock2YT	-16.83	-13.07	-	-
	(17.85)	(18.06)		
interaction_shock2YT	28.28***	28.43***	27.58***	0.0634***
	(4.310)	(4.315)	(1.583)	(0.00623)
after_shock3YT	9.300***	12.67***	12.76***	0.109***
	(2.381)	(2.955)	(3.052)	(0.00832)
experienced_shock3YT	-11.31	-10.16	-	-
	(21.39)	(21.36)		
interaction_shock3YT	26.17***	26.28***	25.00***	0.0395***
	(4.662)	(4.665)	(1.245)	(0.00464)
after_shock4YT	0.762	-11.87***	-12.14***	-0.0247**
	(1.461)	(2.455)	(3.890)	(0.0113)
experienced_shock4YT	47.17***	49.00***	-	-
	(17.38)	(17.35)		
interaction_shock4YT	15.28***	15.36***	15.34***	0.00311
	(3.452)	(3.459)	(2.773)	(0.00886)
patreon_shock1		-27.94***	-27.48***	-0.126***
		(3.144)	(3.159)	(0.00797)
patreon_shock2		-10.75***	-10.75***	-0.0305***
		(2.117)	(3.510)	(0.00857)
patreon_shock3		-4.054***	-4.091	-0.0204**
		(0.656)	(4.014)	(0.00925)
category1		3.728	25.65	0.0240
		(37.91)	(32.49)	(0.427)
category2		-45.81***	-61.54***	-0.0886
		(14.97)	(16.28)	(0.151)
category3		-1.762	133.2	-0.364
		(74.91)	(99.39)	(0.261)
o.category4		-	-	-
o.category5		-	-	-
category6		-60.10***	-84.66***	-0.232
		(17.35)	(18.58)	(0.175)
category7		183.0**	210.9***	1.279***
		(74.89)	(54.34)	(0.282)
category8		-142.6***	-	-
		(46.62)		
category9		-98.80***	-	-
		(21.34)		
category10		-41.24	-74.68***	-0.246
		(28.83)	(16.65)	(0.166)

category11	-26.19	72.36	-0.524**
	(45.12)	(66.91)	(0.240)
category12	66.67	-115.8***	-0.103
	(154.7)	(21.59)	(0.184)
category13	-49.73**	-73.54***	0.0828
	(21.77)	(20.77)	(0.182)
category14	-2.457	52.99	0.716***
	(47.60)	(51.26)	(0.256)
category15	-52.77**	-10.32	-0.113
	(21.20)	(9.593)	(0.105)
category16	-8.075	9.593	-0.0558
	(14.49)	(10.23)	(0.108)
category17	-67.21	15.31	-0.427**
	(58.91)	(75.23)	(0.212)
category18	-97.50***	-	-
	(35.19)		
category19	-40.85**	-14.28***	-0.281***
	(18.62)	(4.383)	(0.0563)
category20	-1.877	7.285	-0.0765
	(18.02)	(10.37)	(0.131)
category21	-36.49**	-9.757	-0.132
	(14.54)	(8.294)	(0.116)
category22	32.33	17.14	-0.748***
	(57.83)	(10.99)	(0.148)
category23	-26.40*	-0.785	-0.158*
	(14.11)	(6.922)	(0.0917)
category24	-31.50**	-6.754	-0.321***
	(15.74)	(9.523)	(0.121)
category25	-29.76*	43.74*	-0.132
	(17.53)	(25.70)	(0.145)
category26	56.00***	-0.162	0.0731
	(19.96)	(9.654)	(0.103)
category27	-0.175	13.21**	0.0158
	(12.87)	(6.619)	(0.0884)
month_year1	4.070***	3.814	0.0300*
	(0.921)	(5.460)	(0.0169)
month_year2	5.475***	5.195	0.0334**
	(1.336)	(5.321)	(0.0166)
month_year3	4.994**	4.516	0.0275*
	(2.247)	(5.082)	(0.0159)
month_year4	-21.76***	-20.74***	-0.311***
	(1.532)	(3.213)	(0.0108)
month_year5	-17.85***	-16.94***	-0.228***
	(1.351)	(3.084)	(0.00994)
month_year6	-15.57***	-14.77***	-0.185***
	(1.200)	(2.964)	(0.00942)
month_year7	-15.00***	-14.31***	-0.178***
	(1.076)	(2.852)	(0.00907)
month_year8	-9.764***	-9.329***	-0.118***
	(0.931)	(2.753)	(0.00878)
month_year9	-6.362***	-6.086**	-0.0724***
	(0.783)	(2.670)	(0.00836)
month_year10	-2.811***	-2.696	-0.0374***
	(0.575)	(2.540)	(0.00812)
o.month_year11	-	-	-

month_year12		-19.57***	-18.70***	-0.161***
		(1.527)	(1.668)	(0.00651)
month_year13		-15.42***	-14.68***	-0.120***
		(1.358)	(1.582)	(0.00621)
month_year14		-13.24***	-12.61***	-0.0932***
		(1.206)	(1.472)	(0.00604)
month_year15		-9.965***	-9.520***	-0.0636***
		(0.986)	(1.391)	(0.00591)
month_year16		-5.940***	-5.663***	-0.0365***
		(0.787)	(1.346)	(0.00584)
month_year17		-3.602***	-3.365**	-0.0176***
		(0.701)	(1.348)	(0.00572)
o.month_year18		-	-	-
month_year19		-27.57***	-26.96***	-0.120***
		(3.234)	(3.172)	(0.00798)
o.month_year20		-	-	-
month_year21		-24.49***	-24.14***	-0.101***
		(2.987)	(3.181)	(0.00800)
month_year22		-20.41***	-20.15***	-0.0766***
		(2.784)	(3.219)	(0.00810)
month_year23		-17.36***	-17.21***	-0.0605***
		(2.532)	(3.323)	(0.00821)
month_year24		-14.19***	-14.19***	-0.0444***
		(2.439)	(3.369)	(0.00832)
month_year25		-17.36***	-17.21***	-0.0605***
		(2.532)	(3.323)	(0.00821)
o.month_year26		-	-	-
month_year27		-7.763***	-7.554**	-0.0220**
		(1.912)	(3.685)	(0.00877)
month_year28		-6.319***	-5.982	-0.0228***
		(1.679)	(3.791)	(0.00879)
month_year29		-5.616***	-5.338	-0.0187**
		(1.534)	(3.833)	(0.00883)
month_year30		-4.925***	-5.069	-0.0131
		(1.460)	(3.903)	(0.00898)
month_year31		-3.680***	-3.591	-0.0145
		(1.174)	(3.950)	(0.00909)
o.month_year32		-	-	-
o.month_year33		-	-	-
month_year34		-0.696	-0.544	0.00927
		(1.257)	(4.137)	(0.0101)
month_year35		0.929	0.985	0.00788
		(1.196)	(4.321)	(0.0104)
o.month_year36		-	-	-
isnsfw		47.85***	67.94***	0.168***
		(5.771)	(11.34)	(0.0402)
avg_twitterfollowers	9.43e-05	9.27e-05	0.000771***	4.61e-07***
	(9.12e-05)	(9.00e-05)	(0.000101)	(1.03e-07)
avg_youtubesubscribers	5.45e-05*	5.31e-05*	4.03e-05***	8.17e-08***
	(2.82e-05)	(2.80e-05)	(1.45e-05)	(2.91e-08)
avg_youtubevideos	0.0149	0.0132	0.0145***	7.38e-05***

	(0.00961)	(0.00950)		(0.00259)	(7.45e-06)
avg_youtubeviews	-8.73e-09	-7.25e-09		-3.43e-09	-2.18e-10***
	(6.30e-08)	(6.29e-08)		(2.53e-08)	(0)
Constant	24.86***	15.76		52.90***	3.012***
	(2.036)	(12.20)		(7.843)	(0.0870)
Observations	205,248	205,248		205,242	205,242
Number of id	6,126	6,126	R-squared	0.894	0.898

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

7.3 Choice to Omit Earnings

Aside from our robustness checks of including more and more fixed effects, we want to address the issue of bias discovered in Section 6 of this paper.

In Section 6.1, when we go from Table 3 Specification (1) to Specification (2 or 3) of the same table, we would expect the two regressions to appear somewhat similar since we are randomly increasing our sample by including people regardless of whether or not they reveal their earnings on Patreon. However, in reality Specification (1) and (2 or 3) vary immensely in the magnitude of their coefficients. This leads us to believe that the choice a creator makes to hide their earnings from public view must *not* be a random choice at all. Thus, regressions that include both average patrons and average earnings are biased.

Additionally, to rule out the notion that this is simply a sample selection issue, in Section 6.1, we included Specifications (5) and (6) (real patrons vs. logs). These regressions were run on the sample of individuals who reported earnings and thus most would have been included in Specification (1). In essence, running different regressions but on effectively the same sample of content creators. Even with this very similar sample, when we omit earnings as a right-hand side variable in Specifications (5) and (6), we still see immense shifts in the magnitudes of our coefficients with, for example, our first DID estimator for YouTube shock #1 changing from 4.615*** to 10.71***.

Thus, to explore the "randomness" of the choice to omit earnings, we created a dummy variable called '*omit_earnings*.' Generated by unique id and time combination, this variable equals one if a value for '*avg_earnings*' is missing but '*avg_patrons*' is greater than zero. This intuitively means that a person is active on Patreon (remember, the number of patrons a person has is public information that is always displayed) and thus making earnings, but if no earnings are shown, then the creator must be choosing to hide them from public view. Otherwise, the dummy variable '*omit_earnings*' equals zero when creators have both their patrons and earnings on public display.

Using '*omit_earnings*' as the dependent variable and including all mentioned fixed effects (see Table 7 notes for further details), we ran a regression that showed that the choice to omit (hide) earnings is not random.

Taking a look at Table 7, Specification (1), we find that '*avg_patrons*' and the DID estimators for YouTube shocks #1,3,4 are significant for at least the 95% confidence level. This results can be interpreted as such: an increase of one-thousand patrons is associated with a 12.2% increase in the probability that a creator omits their average earnings from public view, ceteris paribus. Logically, this holds up. The more popular a creator is, the more patrons they have, the more money they're making, the more likely it is they would want to hide their earnings from the public so that people continue to support them financially.

Alternatively, (e.g., '*interaction_shock1YT*') comparing average values from the post-period starting in August 2016 to the pre-period, content creators that experienced shock #1 are ~1.01% less likely to hide their earnings than those creators who did not experience the YouTube shocks. This last interpretation speaks volume to those YouTubers who were losing earnings on YouTube and wanted to bring that issue to light. To do so, these individuals decided to bring forth full transparency to their earnings on YouTube and their other platforms... in this case, Patreon.

Table 7: Creator's Choice to Keep Earnings Private

	(1) Binary Choice to Omit Earnings		(2) Binary Choice to Omit Earnings
avg_patrons	0.000122***		0.000133***
	(1.79e-05)		(5.60e-06)
after_shock1YT	0.0709***		0.0699***
	(0.00543)		(0.00571)
experienced_shock1YT	-0.00827		-
	(0.0151)		
interaction_shock1YT	-0.0101**		-0.0107**
	(0.00510)		(0.00446)
after_shock2YT	0.0947***		0.0944***
	(0.00658)		(0.00462)
experienced_shock2YT	0.0205		-
	(0.0215)		
interaction_shock2YT	0.00779		0.00690**
	(0.00813)		(0.00321)
after_shock3YT	0.105***		0.105***
	(0.00695)		(0.00478)
experienced_shock3YT	-0.0105		-
	(0.0234)		
interaction_shock3YT	0.0184**		0.0174***
	(0.00786)		(0.00273)
after_shock4YT	0.0135***		0.0137**
	(0.00506)		(0.00599)
experienced_shock4YT	0.0111		-
	(0.0185)		
interaction_shock4YT	0.0141**		0.0135***
	(0.00675)		(0.00472)
Constant	0.0205		0.154***
	(0.0195)		(0.0358)
Observations	205,248		205,242
Number of id	6,126	R-squared	0.599

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Fixed effects for Patreon shocks, category of content, distinct month-year, and whether the content is "safe for work," in addition to controls for average Twitter followers, average YouTube subscribers, average YouTube videos, and average YouTube views have been included in these regressions but omitted in this table for the sake of simplicity. Specification (2) additionally contains fixed effects for each specific content creator.

7.3A Signing the Earnings Omitted Variable Bias

Now that we have determined that the choice to omit earnings is not random, we want to explore how accounting for some of that bias in our preferred models can sign the OVB of earnings.

As mentioned throughout, our preferred models to explain how content creator had to react during the age of the YouTube Adpocalypse are Table 3 Specifications (2), (3), and (4) (or alternatively Table 6 Specifications (2), (3), & (4) which are the same regressions simply showing fixed effect coefficients). Now, in Table 8, we duplicate our preferred models with the addition of the '*omit_earnings*' variable.

The model controlling for the most variation (Table 8 Specification (3)) can be interpreted as follows:

- '*omit_earnings*'
  - "*Significant at the 99.9+% confidence level, content creators on Patreon who choose to hide their average earnings from public view have (on average) ~23.3% more patrons than those who choose to reveal their earnings, ceteris paribus.*"
- '*after_shock1YT*'
  - "*Significant at the 99.9+% confidence level, after August 2016, content creators on Patreon who did not experience YouTube shock #1 had on average ~35.9% more patrons than they did prior to August 2016, ceteris paribus*."

(same pattern of interpretation follows for '*after_shock2YT*,' etc.)

- '*interaction_shock2YT*'
  - "*Significant at the 99.9+% confidence level, comparing average values from the post-period starting in April 2017 to the pre-period, content creators who experienced shock #2 vs. those who did not have ~6.09% more patrons, ceteris paribus*."

(same pattern of interpretation follows for '*interaction_shock1YT*,' etc.)

Table 8: Patron Regression Results with OVB

	(1) Average Monthly Patrons		(2) Average Monthly Patrons	(3) log(Average Monthly Patrons)
omit_earnings	62.52***		60.56***	0.233***
	(7.986)		(2.242)	(0.00523)
after_shock1YT	16.31***		15.30***	0.359***
	(2.227)		(4.040)	(0.0151)
experienced_shock1YT	56.89***		-	-
	(11.38)			
interaction_shock1YT	14.57***		13.85***	-0.0127
	(3.584)		(3.163)	(0.0122)
after_shock2YT	2.055		1.788	0.152***
	(2.309)		(1.998)	(0.00775)
experienced_shock2YT	-14.25		-	-
	(17.88)			
interaction_shock2YT	27.73***		26.94***	0.0609***
	(4.252)		(1.572)	(0.00619)
after_shock3YT	6.011**		6.305**	0.0845***
	(2.819)		(3.035)	(0.00830)
experienced_shock3YT	-9.426		-	-
	(21.10)			
interaction_shock3YT	24.93***		23.75***	0.0347***
	(4.561)		(1.234)	(0.00461)
after_shock4YT	-12.62***		-12.87***	-0.0276**
	(2.511)		(3.879)	(0.0112)
experienced_shock4YT	47.65***		-	-
	(17.22)			
interaction_shock4YT	14.37***		14.39***	-0.000522
	(3.411)		(2.759)	(0.00879)
Constant	10.89		43.14***	2.974***
	(12.53)		(8.055)	(0.0866)
Observations	205,248		205,242	205,242
Number of id	6,126	R-squared	0.895	0.900

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Fixed effects for Patreon shocks, category of content, distinct month-year, whether the content is "safe for work", average Twitter followers, average YouTube subscribers, average YouTube videos, and average YouTube views have been included in both regressions but omitted in this table for the sake of simplicity. Specifications (2) and (3) additionally contain fixed effects for each specific content creator.

Through economic theory, we would expect the sign of the omitted variable bias for earnings to be positive in that the correlation between patrons and earnings is positive (the more patrons, the more earnings) and the correlation between earnings and '*omit_earnings*' could be positive (the more earnings you make, the higher chance a person will want to keep them private). This positive omitted variable bias would mean that the Table 8 Specification (3) results are too positive.

In order to see this OVB in action, we compare the sign and magnitude of the coefficients from Table 3 Specification (4) and Table 8 Specification (3). Table 8 Specification (3) contains that OVB through '*omit_earnings*.' Our story of positive omitted variable bias does not seem to align for any of our shocks, but that could indicate there are other biases at play in Table 8 Specification (3). Or, alternatively, that a person who has less earnings chooses to omit their earnings whereas a person who has more would like to keep them public (now plausibly negative OVB). This notion could also make sense in that the biggest YouTubers were impacted by the Adpocalypse the most so they would want to bring transparency to their earnings for change to happen (as illustrated in the coefficient for '*interaction_shock1YT*' in Table 7 Specification (2)).

Table 9: Signing the Earnings Omitted Variable Bias

Main Variables of Interest	***Table 3 Specification (4)*** *(Unbiased by Earnings)*	***Table 8 Specification (3)*** *(OVB of Earnings)*	*Probable Sign of Earnings Bias*
interaction_shock1YT	-1.48%	-1.27%	N/A
interaction _shock2YT	6.34% ***	6.09% ***	Negative (-)
interaction _shock3YT	3.95% ***	3.47% ***	Negative (-)
interaction _shock4YT	0.311%	-0.0522%	N/A

7.4 Alternative Treatment Group

As an additional robustness check, we explore the question posed in Section 4 regarding content creators who lie 'at the margin' of treatment.

Originally for our YouTube shocks, our treatment groups consist of creators that have 1,000 or more YouTube subscribers and at least one YouTube video at the time of the shock. However, within our histograms we noticed that a lot of the treated individuals have much more than 1,000 subscribers, making us wonder if these smaller YouTubers with 1,000 to 2,000 subscribers were impacted differently.

From comparing DID estimator coefficients within Table 3 Specification (4) to Table 10 Specification (2), we see the "smaller" content creators (less subscribers) being adversely impacted on Patreon by YouTube shocks #1-3, until YouTube shock #4 where they do see an increase of ~6.14%*** in patrons. One potential theory is that during these YouTube shocks, there is increased competition on Patreon and these small players in Table 10 lose the attention of their supporters to the bigger players included in Table 3 since both groups are wanting to supplement lost YouTube revenue with Patreon earnings.

Overall, this model is consistent with the main results in Table 3 which suggest that the effects of the YouTube shocks have different effects across different levels of content providers.

Table 10: Alternative Treatment Group

	(1) log(Average Monthly Patrons)		(2) log(Average Monthly Patrons)
after_shock1YT	0.340***		0.340***
	(0.0185)		(0.0196)
experienced_shock1YT	-0.0502		-
	(0.189)		
interaction_shock1YT	-0.0918*		-0.0922***
	(0.0520)		(0.0323)
after_shock2YT	0.154***		0.153***
	(0.0102)		(0.00983)
experienced_shock2YT	0.226		-
	(0.216)		
interaction_shock2YT	-0.144***		-0.143***
	(0.0401)		(0.0148)
after_shock3YT	0.144***		0.143***
	(0.0130)		(0.0116)
experienced_shock3YT	0.0150		-
	(0.185)		
interaction_shock3YT	-0.0934***		-0.0934***
	(0.0351)		(0.0105)
after_shock4YT	-0.0288***		-0.0300**
	(0.00645)		(0.0149)
experienced_shock4YT	0.285***		-
	(0.110)		
interaction_shock4YT	0.0581**		0.0614***
	(0.0293)		(0.0192)
Constant	2.570***		2.304***
	(0.162)		(0.171)
Observations	75,861		75,860
Number of id	2,276	R-squared	0.885

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: These results were generated by limiting the treated group to content creators who have 1,000 to 2,000 subscribers. Fixed effects for Patreon shocks, category of content, distinct month-year, whether the content is "safe for work", average Twitter followers, average YouTube subscribers, average YouTube videos, and average YouTube views have been included in both regressions but omitted in this table for the sake of simplicity. Specification (2) additionally contains fixed effects for each specific content creator.

8 Conclusion

Throughout this paper we have captured the migration of some of content creators' focus from YouTube to Patreon. Relying on Patreon to supplement lost income during the four YouTube Adpocalypse shocks (August 2016, April 2017, November 2017, February 2019), many creators self-promoted their YouTube channels and Patreon channels simultaneously to ask their viewers for support. Using data from Graphtreon and a difference-in-differences identification strategy, we observe creators gaining a statistically significant number of patrons during these time periods as seen in our best model (Table 3 Specification (3) or Table 6 Specification (3)).

Additionally, we are able to control for differences in effects among different content categories on Patreon, whether content is "safe for work," distinct month-year effects, fixed effects for each specific content creator and arguably most importantly, shocks on Patreon's side of this interaction. As extensions, we find that including the binary choice to omit earnings on Patreon from public view does bias results.

Future work in this area should seek to further quantify how large of an impact the YouTube Adpocalypse has had on content creators, especially in terms of revenue or subscribers lost instead of patrons and earnings gained, which was the primary focus of this paper. Furthermore, potential policy implications may arise surrounding the right to protect free speech and limit censorship on any user-generated platform.

9 References

Alexander, J. (2019, February 19). YouTube still can't stop child predators in its comments. The Verge. https://www.theverge.com/2019/2/19/18229938/youtube-child-exploitation-recommendation-algorithm-predators

Besley, T. and Prat, A. (2006). Handcuffs for the grabbing hand? Media capture and government accountability. American Economic Review, 96(3):720–736.

Bowles, N. (2018, December 24). Patreon Bars Anti-Feminist for Racist Speech, Inciting Revolt. The New York Times. https://www.nytimes.com/2018/12/24/technology/patreon-hate-speech-bans.html?searchResultPosition=1

de Corniere, A. and Sarvary, M. (2020). Social media and the news: Content bundling and news quality. TSE Working Paper, n. 20-1152.

Ellman, M. and Germano, F. (2009). What do the papers sell? As model of advertising and media bias. The Economic Journal, 119(537):680–704.

European Union - Data Privacy and Protection. (2020). International Trade Administration | Trade.Gov. https://www.trade.gov/european-union-data-privacy-and-protection#:%7E:text=The%20EU%20General%20Data%20Protection,to%20companies%20of%20all%20sizes

Gal-Or, E., Geylani, T., and Yildirim, T. P. (2012). The impact of advertising on media bias. Journal of Marketing Research, 49(1):92–99.

Gentzkow, M. and Shapiro, J. M. (2006). Media bias and reputation. Journal of Political Economy, 114(2):280–316.

Graphtreon: Patreon Earnings + Statistics + Graphs + Rankings. (2021). Graphtreon. https://graphtreon.com/

Kerkhof, Anna (2020). "Advertising and Content Differentiation: Evidence from YouTube," VfS Annual Conference 2019 (Leipzig): 30 Years after the Fall of the Berlin Wall - Democracy and Market Economy 204468, Verein für Socialpolitik / German Economic Association.

Kumar, B. (2020, February 17). How to Make Money on YouTube (Without a Million Subscribers). Shopify. https://www.shopify.com/blog/198134793-how-to-make-money-on-youtube#:%7E:text=Whether%20you%20want%20to%20earn,hours%20over%20the%20past%20year

Luca, M. (2015). User-generated content and social media. In Handbook of Media Economics, volume 1, pages 563–592. Elsevier.

Madio, Leonardo and Quinn, Martin, Content moderation and advertising in social media platforms (2021). Available at SSRN: https://ssrn.com/abstract=3551103 or http://dx.doi.org/10.2139/ssrn.3551103

Mullainathan, S. and Shleifer, A. (2005). The market for news. American Economic Review, 95(4):1031–1053.

Pottinger, N. E. (2018). Don't forget to subscribe: Regulation of online advertising evaluated through youtube's monetization problem. Kentucky Law Journal, 107(3), 515-[viii].

Robertson, A. (2017, December 13). Patreon scraps new service fee and apologizes to users. The Verge. https://www.theverge.com/2017/12/13/16772556/patreon-drops-service-fee-plan-apology-update

Stanford, S. (2018). YouTube and the Adpocalypse : How Have The New YouTube Advertising Friendly Guidelines Shaped Creator Participation and Audience Engagement?

We’re Updating Patreon’s Fee Structure. Here’s Why. (2017). Patreon Blog. https://blog.patreon.com/updating-patreons-fee-structure/

Xiang, Yi & Sarvary, Miklos. (2007). News Consumption and Media Bias. Marketing Science. 26. 611-628. 10.1287/mksc.1070.0279.

Yildirim, P., Gal-Or, E., and Geylani, T. (2013). User-generated content and bias in news media. Management Science, 59(12):2655–2666.

YouTube Adpocalypse. (2019). Wiki. https://youtube.fandom.com/wiki/YouTube_Adpocalypse

YouTube is 15 years old.... (2020, May 30). Business Insider Nederland. https://www.businessinsider.nl/history-of-youtube-in-photos-2015-10?international=true&r=US#august-2016-youtube-starts-notifying-creators-when-their-videos-are-demonetized-meaning-the-platforms-algorithm-has-decided-some-part-of-the-clip-doesnt-adhere-to-advertiser-friendly-content-guidelines-so-its-ineligible-to-make-money-although-these-guidelines-have-existed-since-2012-popular-creators-take-notice-of-their-rising-impact-youtuber-philip-defranco-criticizes-demonetization-as-a-form-of-censorship-41

Zhang, K. and Sarvary, M. (2014). Differentiation with user-generated content. Management Science, 61(4):898–914.