Journal of Quantitative Description: Digital Media 5(2025), 1–37 10.51685/jqd.2025.006

Pressing Play on Politics: Quantitative Description of YouTube

Kevin Munger
European University Institute, Italy

James Bisbee
Vanderbilt University, USA

Omer Yalcin
University of Massachusetts Amherst, USA

Joseph Phillips
Cardiff University, Wales

Matthew Hindman
The George Washington University, USA

We present a large-scale quantitative analysis of anglophone politics channels on YouTube, with three distinct units of analysis: channels, comments, and videos. We demonstrate that although channels have been entering the YouTube system at a roughly constant rate since 2008, there is serious inequality in the attention received by different channels and videos. Furthermore, prolific commenters are responsible for an astonishing amount of activity: 50% of total comments are written by just over 2% of all commenters. The toxicity for which YouTube comments are famous tends to be more pronounced among these super-users than among infrequent commenters. Our findings have important implications for the way in which YouTube viewers interpret what they see as representative of public opinion.

Keywords: YouTube, Comments, Quantitative Description

Munger: kevin.munger@eui.eu
Date submitted: 2024-09-17

Introduction

YouTube remains the most-used social media platform in the United States and one of the most-used in the world. Although some uses of YouTube are more like the traditional broadcast mediums of television and radio, the social element is extremely important. In a 2018 interview, Google CEO Eric Schmidt said that "today we have quite a powerful social network embedded inside of YouTube" (Cowen, 2018).

Academic research has not, however, caught up to this fact. Summarizing the state of the literature, Norton and Shapiro (2024) find major biases in which platforms receive academic attention: "Major platforms with large global user bases—including YouTube, WeChat, and Telegram—remain critically understudied. While there were approximately 22.4 and 1.4 published papers per 100 million active users for Twitter and Facebook respectively, YouTube, WeChat, and Telegram attracted far less research despite their substantial userbases. Some 2.3 billion people actively use YouTube each month. And WeChat and Telegram have 1.3 billion and 550 million active monthly users respectively."

Aridor et al. (2024) provide a more granular analysis of the distribution of academic attention to different social media platforms, narrowed to Economics papers up to the year 2022. They find that around sixty percent of global social media users use Facebook, and that around fifty percent of Economics papers that study social media study Facebook; a reasonable match. However, every other social media platform receives an unreasonable amount of attention. Only fifteen percent of users use Twitter, while over forty percent of papers study it. In contrast, YouTube is second only to Facebook in global usage at just over fifty percent, while receiving well under five percent of the attention.

We thus undertake an exercise in quantitative description of political YouTube. YouTube is extremely large and heterogeneous, and we thus take up Gerring's (2012) call for research that consists of "mere description." It is currently possible to invent and test some causal hypotheses involving YouTube, but in the absence of descriptive knowledge of *what it is*, it is impossible to know whether these causal relationships are important. We build on the work of Rieder et al. (2020) in "Mapping YouTube" and the earlier statistical analysis by Bärtl (2018) with the more specific goal of mapping *political* YouTube.

We have the further advantage of a much larger temporal window through which to analyze YouTube. Digital social phenomena tend to be fast-moving, necessitating the curation of time series data in order to understand not simply where they are but where they have been and where they are going. Our approach looks squarely at the challenges of doing social science research in internet time (Karpf, 2012), particularly that of low temporal validity (Munger, 2023).

The scope of our quantitative description is necessarily circumscribed by the limits of data availability and computational power. YouTube sees more than 500 hours of content uploaded every minute. In addition, gathering data from the platform is throttled by limits placed by YouTube on its API.

As such, for this characterization of the platform, we focus on a subset of the overall ecosystem—channels pertaining to politics in the United States. This list of channels was constructed by [author name redacted] who hand-curated it in the summer of 2019 through a combination of extensive experience with the platform and training in qualitative coding. This methodology means that there are some smaller channels that are particularly active in prominent YouTube political subcultures that are included, while some political channels with equivalent aggregate viewership are not. However, given the inequalities observed, we can be confident that we have captured a sufficiently large percentage of (English-language) activity that the patterns we document should be robust to missingness.

This approach ensures that we have not missed any of the largest political channels and that we have sufficient variation in channel size and topic while minimizing false positives. Several of the channels are no longer active, underscoring the selection biases inherent to any analysis of a dynamic social media platform at a given time.

We divide our analyses into three categories, defined by the unit of interest: channels, videos, and users. For each unit, we provide basic summary statistics and temporal summaries of when videos were uploaded and when comments were written. In addition, we use NLP methods to characterize what the videos are about (using video descriptions) and what the comments are saying. Finally, we combine the information on channels and users to build network representations of which videos and channels are commented on by the same users.

This analysis leads us to several conclusions. First, there is significant inequality on both the supply and the demand side: The YouTube politics ecosystem is dominated by a small number of channels producing the vast majority of content, and a small number of users providing the vast majority of comments. As a topline point of comparison, Bärtl (2018) found that 3% of YouTube channels receive 85% of the views; in our specifically curated sample, the number is 72.8%.

Second, small minorities of users define the conversation: Super users define the conversation by providing the bulk of the comments (50% of total comments are written by just over 2% of all commenters), but also because their comments are more likely to be replied to and liked.

Third, short attention spans: Engagement with political videos drops off precipitously after a few days, with most videos receiving almost 100% of their lifetime comments within a week of being posted. Exceptions to this rule are those whose content went viral well after publication.

Fourth, patterns of cross-channel participation show clear ideological divides in the YouTube network, with most commenters participating exclusively within communities. Yet there are a few channels, such as Joe Rogen's podcast and NowThis News, that attract comments from users otherwise active on just one side of the political spectrum.

Our descriptive results deepen our understanding of participatory inequality on dominant digital platforms. Much previous work on digital participatory inequality has relied on crude rules of thumb (Nielsen, 2006), self-reports from survey data (Khan, 2017), or examinations of niche platforms like health forums (Van Mierlo et al., 2014) or Wikipedia (Panciera et al., 2009; Wu et al., 2009). By contrast, we capture individual-level participation patterns at scale on one of the world's largest digital platforms.

Our findings here raise questions about whether the political discussion on digital platforms lives up to the ideals of deliberative democracy. On political YouTube, at least, political debates are dominated by a small, unrepresentative, and generally insular minority.

Situating Political YouTube

YouTube was first launched on February 14, 2005 as a website that allowed people to share videos up to 10 minutes long. The site was an rapid success, with Google announcing the purchase of YouTube for a $1.65 billion only eight months later. YouTube was a significant locus of political discourse almost immediately. Some of the earliest presence of politics on YouTube took the form of mainstream news channels posting clips.

To demonstrate the importance of YouTube, we can do no better than McGrady et al. (2023)'s recent summary in these pages:

"It has given rise to new forms of entertainment (Burgess, 2018), new modes of communication (Tolson, 2010), new economies (Ørmen and Gregersen, 2023), new forms of marketing (Mowlabocus, 2020), new educational tools (Duffy, 2008)...It has transformed older media industries in profound ways, from music (Cayari, 2011) to news and television (Al Nashmi et al., 2017)."

However, research on political YouTube has long lagged behind that of higher-latency, networked-based social media platforms like Facebook and particularly Twitter. The political effects of the latter were faster and more obvious, as they were used by activists to coordinate on street actions that overthrew several sclerotic regimes and changed electoral politics in many developed democracies (Tufekci, 2017; Gerbaudo, 2018; Karpf, 2016; Barberá et al., 2015).

The political effects of YouTube have taken longer to be felt, and even longer to be embraced by political scientists. Communications research adapted faster, developing an understanding of YouTube as a space for participatory creation (Burgess, 2018). This is very different from the way that Twitter was used for coordination and rapid iteration on activist messaging.

In keeping with this distinct theoretical understanding of YouTube, the most important political effects of the platform can be understood through Munger and Phillips (2022)'s "supply and demand framework." There are distinct socio-technical factors on both sides of this information economy that have historically determined how YouTube is used, politically.

For example, conservatives and the far-right tended to adopt fledgling communication technologies in their early stages, due perhaps to an imbalance in the supply of mainstream media (Lewis, 2018). Liberal and progressive creators tended to exist within the Young Turks network or similar networks of professional creators. Some progressive channels existed, but tended to post sporadically. Furthermore, because of harassment from right-wing audiences, these creators often disabled comments (Chae and Lee, 2024), limiting creator-audience interaction.

However, in the latter half of the 2010s, two transformations occurred. First, far-right content declined in viewership. This is sometimes attributed to changes in the YouTube recommendation system meant to de-prioritize politically extreme content, but this period also coincided with an increase in mainstream conservative content (Hosseinmardi et al., 2020). Second, the rise of "BreadTube," a loosely-connected set of left-wing streamers and video-essayists ranging from progressive to Maoist, filled a vacuum for grassroots left-of-center content creators (Maddox and Creech, 2021). By the end of the 2010s, YouTube had come to cater to a wide range of political perspectives (Lai et al., 2022).

Our account therefore agrees with Munger (2024), which argues that politics on YouTube has evolved considerably in the nearly two decades since the platform was founded. Although fringe politics of various strands flourished on YouTube in the early-to-mid 2010s, it was only with the Trump campaign for the 2016 US Presidential Election that anglophone YouTube politics began to receive mainstream attention. And it was not until the following election cycle that it played a significant role in shaping rather than merely reflecting the process:

> The 2020 presidential election was the first major crossover between YouTube politics and traditional electoral politics in the US. Some the most successful members of YouTube Politics are now household names. The insurgent candidacy of entrepreneur and political novice Andrew Yang was only possible because of central YouTube Political figures. Yang said that "what launched us was Sam Harris....Joe Rogan was the game changer...We raised tens of thousands of dollars a day for awhile there and a million bucks in a week." (Weiss, 2020).

The advent of the COVID pandemic dramatically changed the landscape of political YouTube.

During the pandemic, YouTube, like other social networking sites, invested heavily in attempting to combat misinformation. Such misinformation tended to focus on the intentions of governments regarding viral containment measures, the efficacy and safety of the COVID vaccine, and the results of the 2020 United States presidential election (Topinka, 2024; Kim et al., 2020). Several channels, mostly far-right, were banned for posting such misinformation (Klinenberg, 2024). Other channels, anticipating bans, moved to alternative websites such as Rumble and BitChute (Klinenberg, 2024; Rauchfleisch and Kaiser, 2024). As a result, right-wing audiences are now more divided between YouTube and alternative websites. These substantively important changes to both the supply and demand side of YouTube – as well as changes to platform policies around content moderation and monetization – mean that YouTube post-COVID is significantly different than it was previously.

The monetization angle is particularly important for understanding YouTube, as it was for a long time the only platform with an explicit revenue-sharing arrangement with its biggest creators. A recent book covers this topic in more detail (Munger, 2024), but we want to emphasize the crucial difference in the two types of participation we study below. The actual *creators*, the ones making videos, are actually able to make money and perhaps even a career. The *commenters,* however, are unlikely to engage for financial reasons. Their participation – sometimes, their extremely prolific participation – comes from another source.

Miller (2021) provides one case study of commenters on QAnon posts, emphasizing the importance of the “likes” that comments receive. In particular, posts mentioned China recieve the most “likes,” on average, perhaps pointing towards the evolution of sentiments in this community. The supply-and-demand process is thus taking place even interior to the larger process of video production and consumption (Inwood and Zappavigna, 2023).

Methods

We rely on the YouTube API to gather data on our list of politics channels. The application process for the API is quite simple; unlike other commonly-studied platforms, YouTube does not have a separate Research API, so we used the same API that YouTube creators use to programmatically interact with the site. The API uses a quota system in which different calls cost different amounts of quota. Up until the end of 2019, free public API keys had a maximum limit of 1,000,000 points daily. Changes to the API dramatically reduced this amount to 10,000 daily points. Generally, 1 point returns 1 piece of information; in our case, each point returns either a channel, video or comment.

The API has several endpoints, corresponding to channels (identified by a `channel.id`), videos (identified by a `video.id`), and comments (identified by a `comment.id`). The API also facilitates searching by keywords. Our pipeline is visualized in Figure 1.

As illustrated, we start with a given channel ID taken from the curated list.We use this to extract a variety of statistics from the API's channel endpoint, including the full list of video IDs. These IDs are then pointed at the API's video endpoint, allowing us to obtain a number of statistics including the list of all comments made on the video, each of which has their own comment ID. Finally, these comment IDs are used to extract data on the comment itself, including a unique identifier for each author.

YouTube gives every user their own channel, meaning that these author IDs are simply the channel ID for each user. We do not restart the pipeline using the `author.id` for the sake of computational tractability. Even with this small subset of the YouTube ecosystem, we have over 320 million comments written by more than 17 million users across 2.5 million videos published by 2,940 channels between June of 2020 and February of 2023. In the sections that follow, we provide rich summary descriptions of each of these units.

Data

Starting in the summer of 2020, we began collecting data on 1,712 YouTube channels identified over the course of 2019, of which 1,438 were active when we started the collection. This selection of channels represents a snowball sample. The 'seed' channels were those catalogued

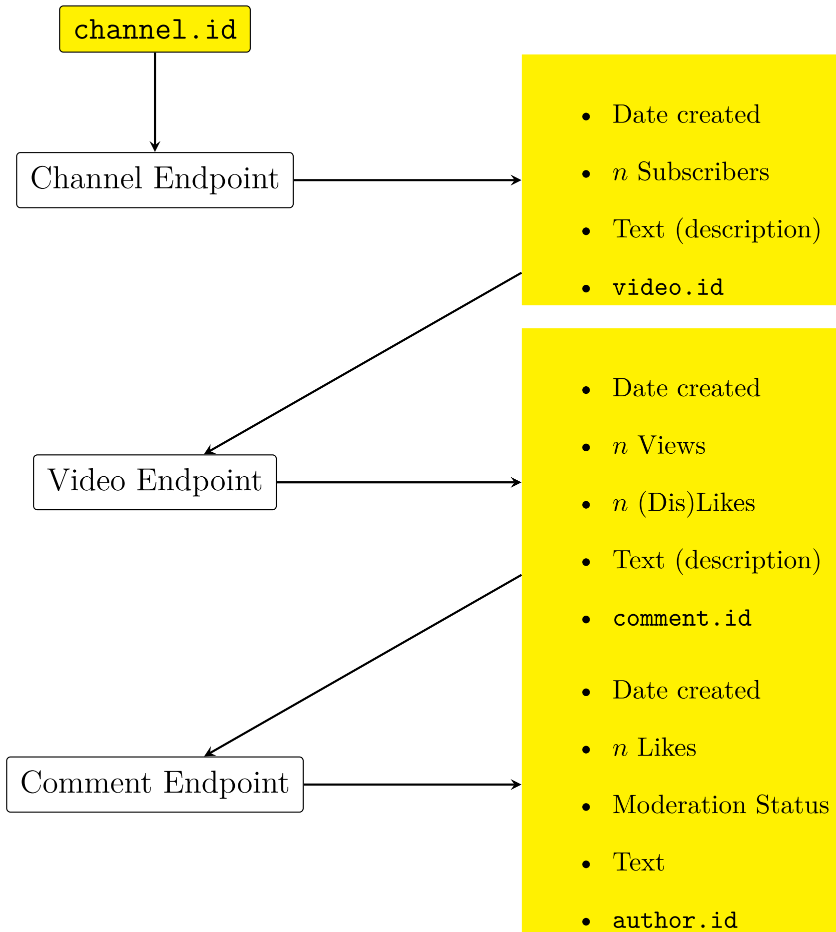

Figure 1. Data collection pipeline.

in Lewis (2018). We first went to each channel's page to find other channels owned by the same creators or channels the creators directly promoted. Then, we sifted through each user's video back-catalogue to find other content creators the channel mentions or invites as a guest. We introduced that content creator into the sample if we determined the channel's content to be "political" in nature. We decided to leave this definition intentionally broad, but often encapsulated discussing current events, ideologies, the culture wars, and/or elections. This sample included larger channels such as The Young Turks and Turning Point USA, as well as smaller channels operated by a single person. Once we had this set of channels, we repeated this exercise for all of our newly collected channels.

The goal was to re-collect data intermittently in order to (1) display the trajectory of channel and video popularity over time, and (2) identify channels that were taken down, due to either their owner's decision or due to violating YouTube's terms of service. As illustrated in Appendix Figure 14, we started off gathering these data in fits and spurts over the remainder of 2020, before settling on a daily cadence in the spring of 2021. (Days on which we gathered the data multiple times correspond mainly to code testing and bug fixing.)

Our initial interest was in the 1,438 channels that were active from our original list, although the preliminary attempts to collect these data consistently at scale were hampered by quota limitations associated with YouTube's API. By the fall of 2020, we began successfully collecting these data, although continued to face challenges until the spring of 2021. On some days, the daily scrape only collected a subset of these, while on other days the scraper failed entirely. On October 17, 2021, we more than doubled the target list to 2,940 channels through repeating the aforementioned exercise of snowball sampling once more.

Notwithstanding the occasional drops in coverage due to technical issues, there is also a subtle but persistent downward trend, even when we successfully scraped all channels on a given day. These declines represent channels that were taken offline, either due to the channel's owner's decision, or due to YouTube policy. By the final full day scraping, 2,629 of the target 2,940 channels remained. Although it is not perfect, we can examine the date at which a given channel last appeared in scraped data. Overall, attrition appears to be low and steady at around 1 or 2 channels per day. The exception is October 17th, 2021, which saw 139 channels record their last appearance in the scraped data, corresponding to our decision to drop many of the channels from the refreshed list because they had ceased publishing.

Descriptively, the channels that disappeared from our data collection exhibit hallmarks we would associate with a "natural causes" disappearance. They have fewer views, fewer subscribers, and fewer videos. They are slightly more likely to change their title, but slightly less likely to revise their channel description, and substantially less likely to change their profile image. To provide a richer understanding of these channels, we turn to natural language processing and run a topic model on the channels' titles and descriptions. We include three covariates of interest when doing so: whether the channel disappeared prior to the end of our data collection, the logged number of subscribers, and the logged number of videos. We estimate the model with 20 topics based on examining elbow plots ranging from 7 to 30 total topics. We plot the association between channels that do and do not survive to the end of our data collection period, and the probability of each of the 20 topics in Figure 2.

As illustrated, there is little to suggest that the content of the channel's metadata differs significantly between those that did and did not disappear over the period of our analysis. We note that our topic model analysis has some limitations. We fit the model to channel metadata and the metadata may vary from channel to channel in terms of its relationship with the actual content and quality of the videos produced by the respective channel. However, we do think our choice of STM is more appropriate for channel metadata than transformer-based approaches as the former's bag-of-words approach is better suited for shorter text like titles and descriptions that are not always made up of full sentences.

While we encourage future research on this point, we do rule out with our application the existence of systematic, large-scale differences in channel attrition by topic. The simplest explanation is the channels that disappeared did so for valence reasons – in other words, the channels that dropped out of our sample were simply uploading worse videos, so the channel creators decided to shut down their accounts. Some of them may also have been suspended by YouTube for violating their content policy; we are unfortunately unable to distinguish these causes from the data provided by the YouTube API.

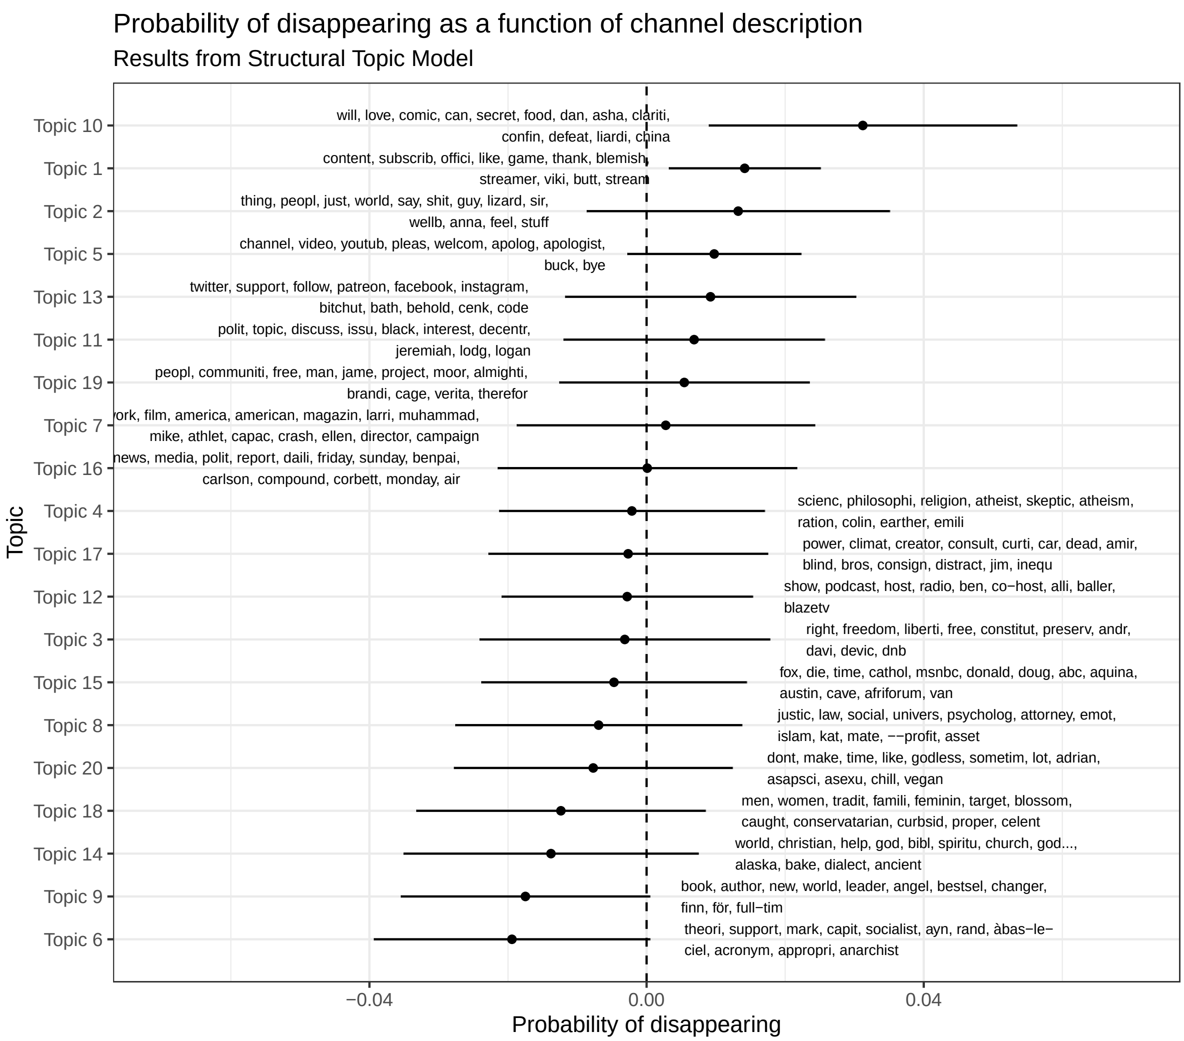

Figure 2. Topic probability by channel attrition. Positive coefficients indicate that the associated topic is more likely to be associated with channels that disappear before the end of data collection.

Descriptive Results

Channels

We begin with the data returned at the first endpoint of our pipeline: the politics channels themselves. As illustrated in Figure 3, these channels vary dramatically in size.

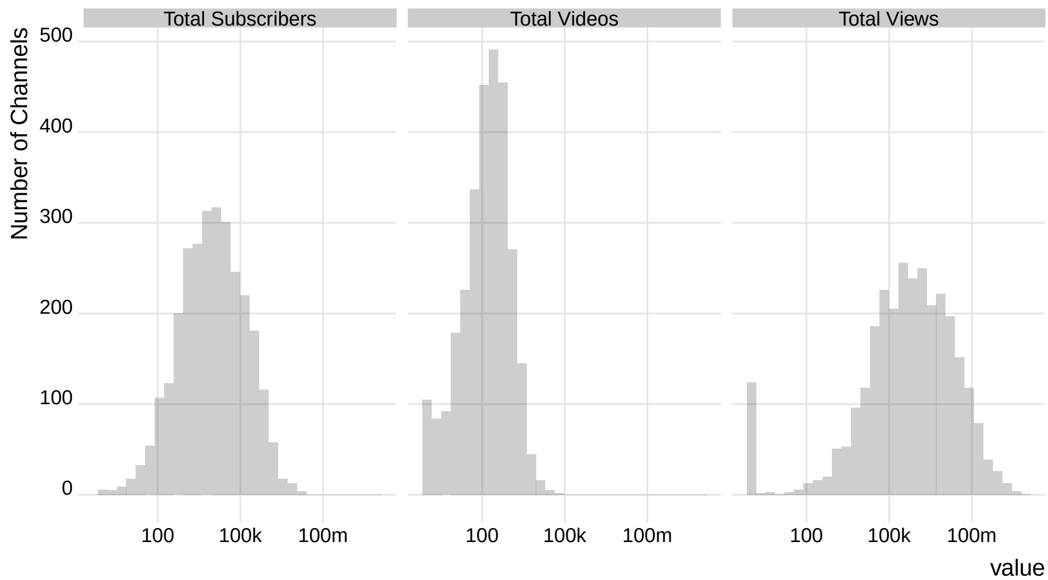

Figure 3. Histograms of channels by size metric, including total views (right), total subscribers (left), and total videos (center).

The majority of them are small, whether measured in terms of views, subscribers, or videos. These patterns are consistent with other descriptions of YouTube, such as those released by the company itself, wherein the platform is described as being comprised of channels of sizes ranging from "dinosaurs" (> 1m subscribers) down to "bugs" (10 to 100 subscribers), in the memorable typology proposed by commentator Matthias Funk.

We are able to observe the creation date for each channel, allowing us to plot the accumulation of our sample over time (Figure 4). As illustrated, there is a strikingly linear accumulation of these channels over time, with a slight lag between the first created channels and the proliferation later in 2005. After this point, politics channels are created at a rate of approximately 2 per week. There is subtle evidence of ebbs and flows in this accumulation of channels, specifically in 2008-2009 and then again most recently in the end of 2019. But these aberrations are so small as to be trivial. Writ large, politics channels on YouTube have been proliferating consistently since the platform was bought by Google in 2006.

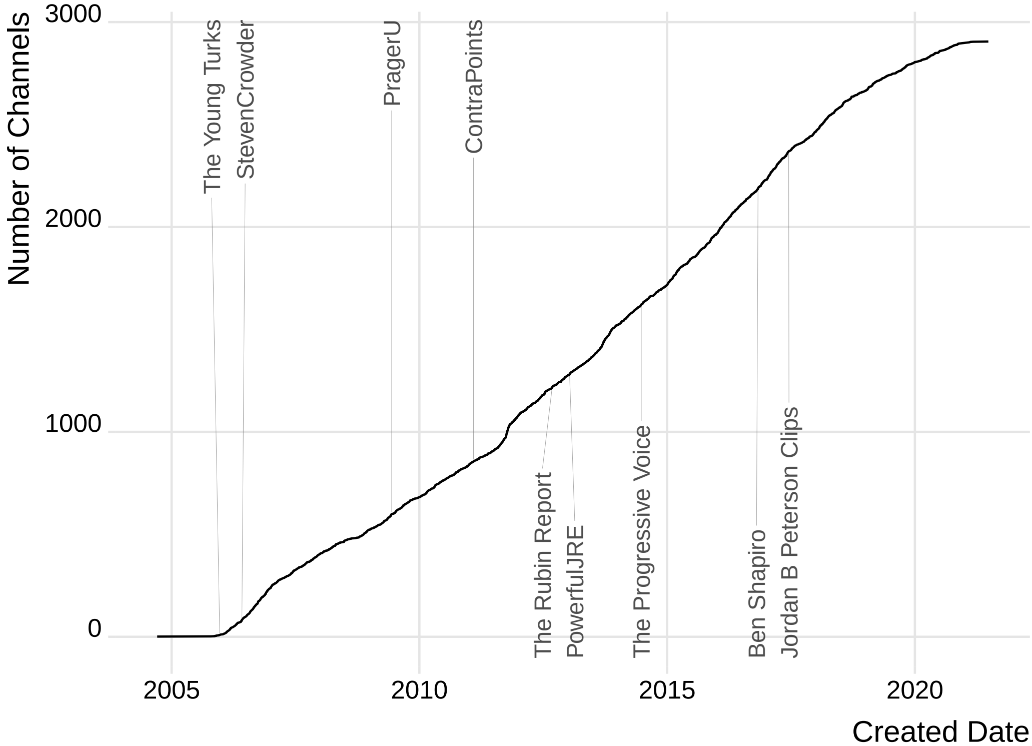

Figure 4. Accumulation of channels over time.

Videos

However, the steady growth of this ecosystem masks the decline and decay of certain channels. We characterize the different trajectories of channel growth or decline as a function of the rate at which they accumulate new videos. Overall, Figure 5 presents evidence of a hyperbolic growth rate between 2005 and mid-2015, after which the increase takes on a more steady linear growth pattern, albeit one much sharper than the accumulation of channels. But how do these trajectories appear when broken out by channel?

Figure 6 plots the accumulation of videos by scaling both the x and y-axes to fall between 1 and 100 where 1 is the first date (x-axis) / fewest videos (y-axis) and 100 is the last

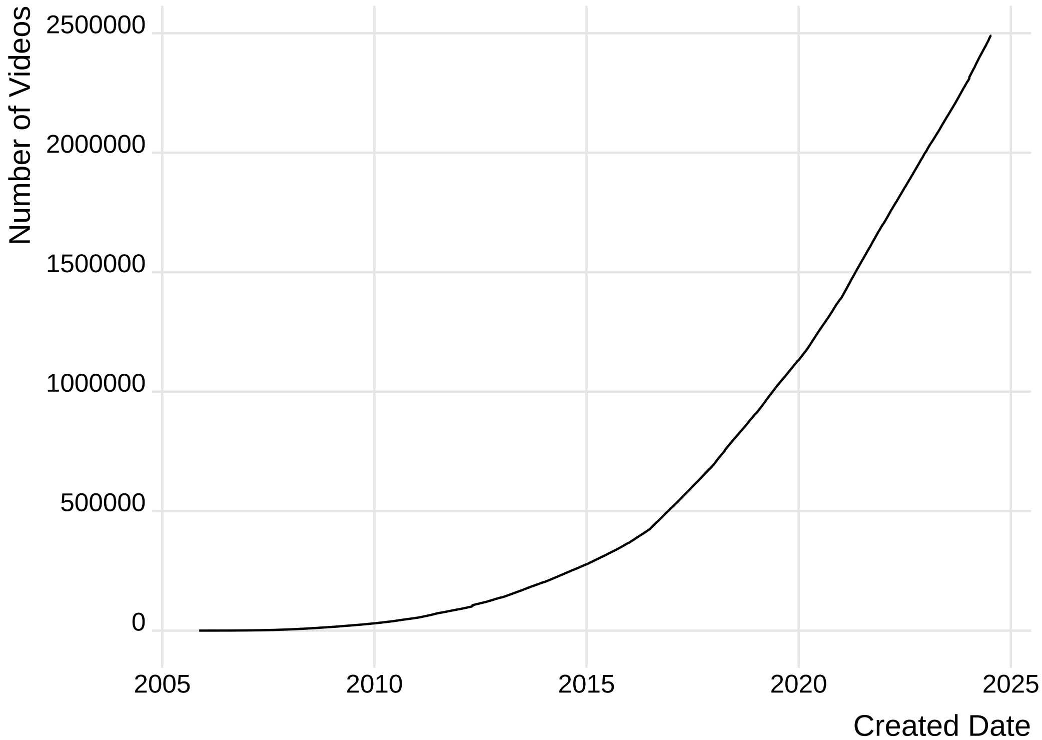

Figure 5. Accumulation of videos over time.

date (x-axis) / most videos (y-axis). The bottom line that curves increasingly upward is an example of a growing channel (Turning Point USA) whose rate of content production increases over time. The center line that is roughly linear is an example of a channel (The Young Turks) whose content is produced consistently. The top line that flattens out over time is an example of a channel (A Swede Speaks) in decline—the rate at which it produces new videos has fallen off dramatically since its most active period.

We can treat the area under the curve (AUC) as a rough proxy for the degree to which different channels are growing or shrinking relative to their historical production of content. Specifically, we measure the difference between the empirical AUCs for each channel and the 45 degree line, re-orienting the measure such that those above the line are negative (shrinking) and those below the line are positive (growing). We plot these as summary statistics for each channel with more than 50 videos in its lifetime, and arrange them in Appendix Figure 15.

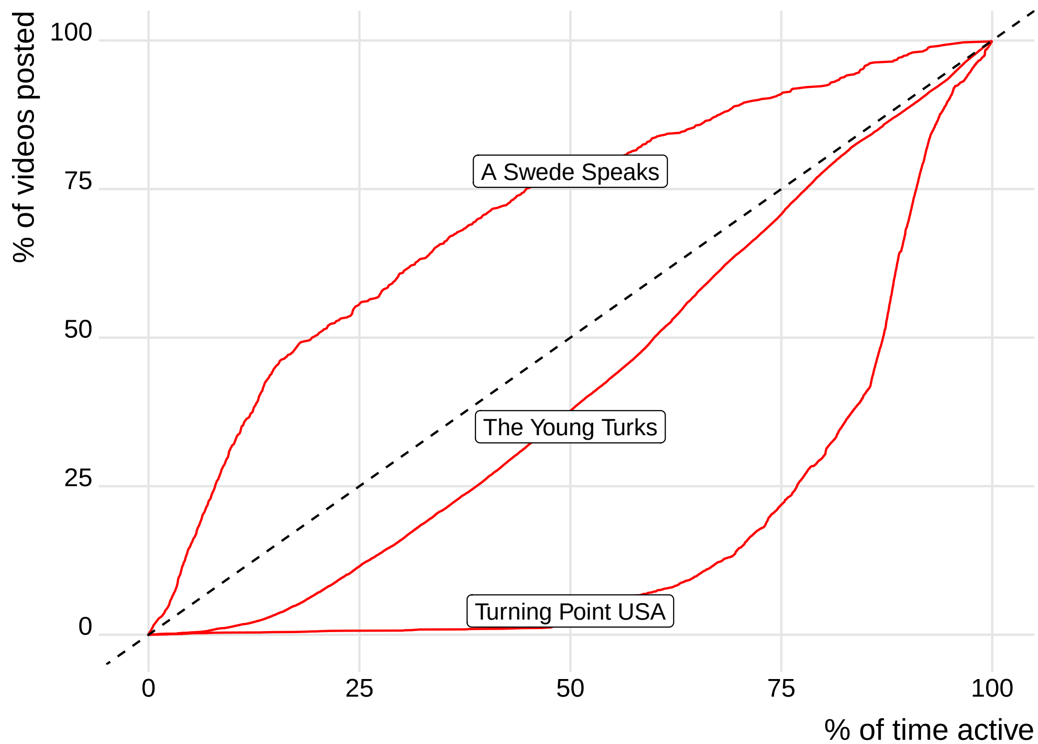

Figure 6. Accumulation of videos over time by channel, scaled.

The politics channels on YouTube are generally growing, with over 60% of all channels in the data ramping up production.

While useful for visualizing whether channels are waxing or waning, the normalization in Appendix Figure 15 masks the dramatically different rates of productivity in our sample. Figure 7 plots the same accumulation of videos over time using the raw dates and cumulative videos (logged for visual clarity). As illustrated, The Young Turks channel has been in production much longer, and has an order of magnitude more videos than either of the other examples, as well as above the 95th percentile of videos for the entire period (given in gray). Overall, we see a mix of big and small channels with a steady accumulation of videos over time. However, this visualization of the data also underscores that the growth metric provided by the AUC in Figure 15 penalizes channels with viral beginnings, such as A Swede Speaks, and emphasizes those that bloom late, such as Turning Point USA.

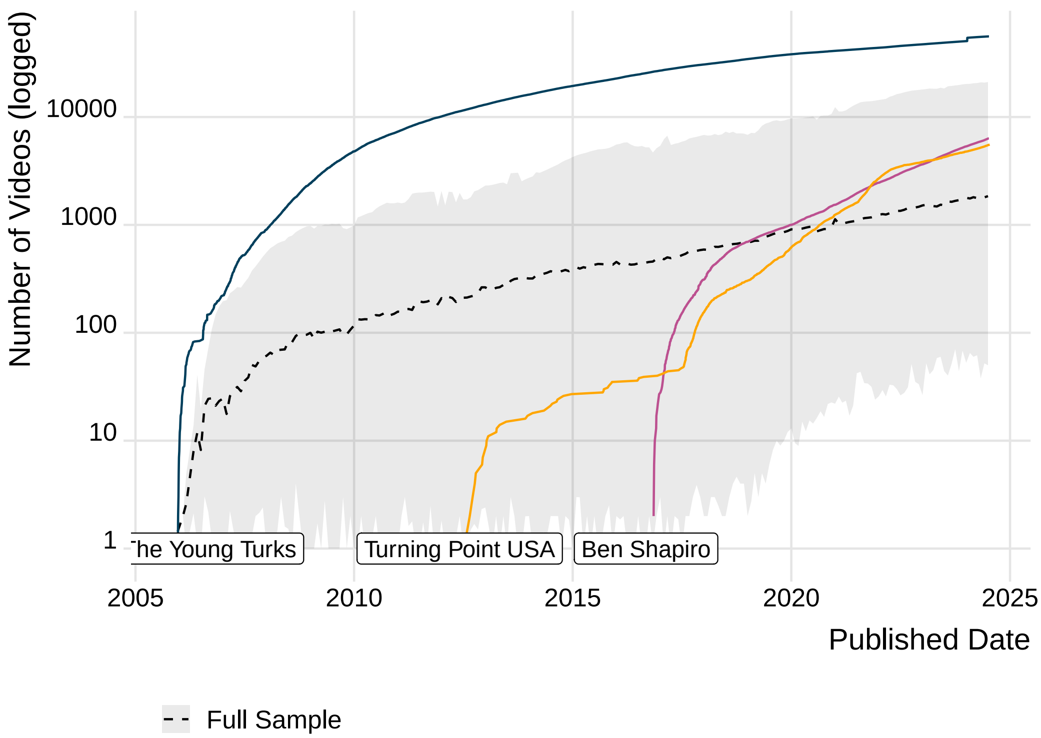

Figure 7. Accumulation of videos over time by channel, logged. Gray density represents the 95% interval of videos by day. Dashed black line represents the average. Labeled lines in blue (The Young Turks) gold (Turning Point USA), and purple (Ben Shapiro) highlight three examples across that were established at different times.

Comments

Finally, we turn our attention to the comments on the videos. Here we see striking evidence suggesting that the majority of a video's comments are written in the first few days after the video is posted. As illustrated in Figure 8, while different videos have more or fewer total comments, almost all videos exhibit a sharp curve in cumulative comments after the first few days. This flattening is consistent with the vast majority of comments being written immediately after the video is published.

However, there are some exceptions to this pattern where videos become popular well after they are published. This is highlighted by the "Asking Cops the Same Silly Questions

They Ask Us" video, which experienced jump in comments approximately 4 months after it was released. Similarly, the video titled "HOW NEOPRONOUNS HURT TRANS PEOPLE" had no comments whatsoever in the first few weeks after it was posted, at which point it suddenly became heavily commented upon, experiencing a sharp linear increase in attention thereafter. Unfortunately, it is generally difficult to characterize the causes of virality based on the content of media, as Salganik et al. (2006) demonstrates, so we cannot say what causes some videos to experience this lagged explosion in popularity.

In general, the overall pattern wherein the majority of comments are written in the first day or two after publication is the norm.

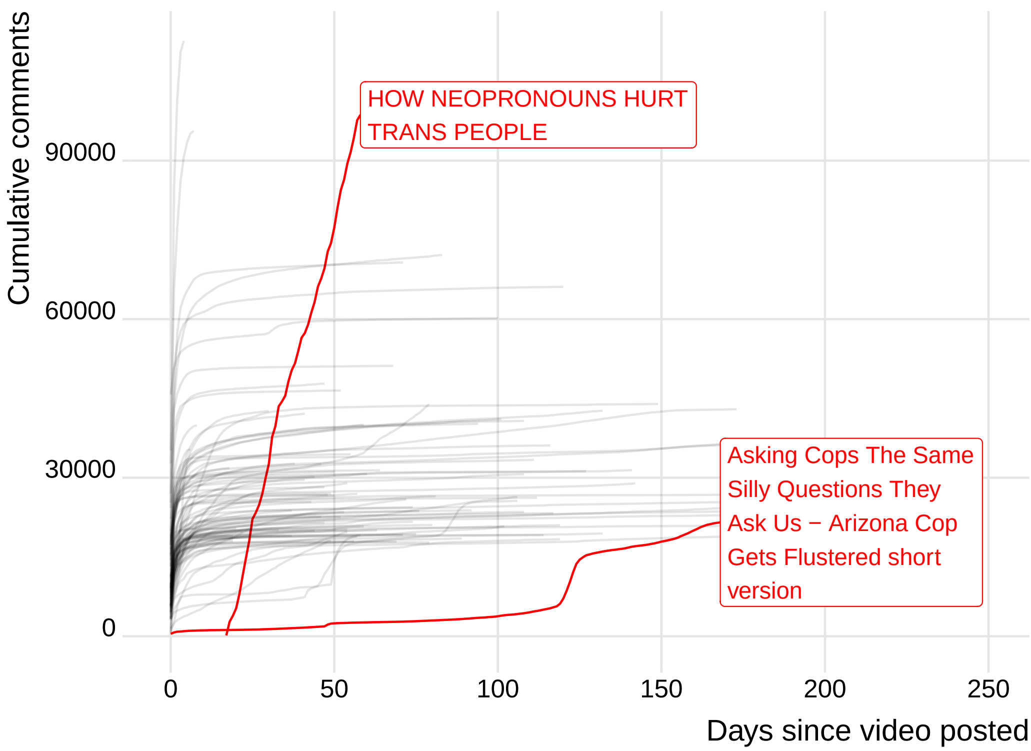

Figure 8. Accumulation of comments over time for top 100 videos. The majority of videos have the majority of their comments written in the first or second day. However, there are some notable exceptions to this rule, highlighted in red.

These patterns are reinforced in Figure 9 which plots the proportion of a video's total comments that are written each day after the video was first published. In this plot, we average over all videos with more than 100 comments and plot the mean as a point and the 95% interval as bars. As illustrated, the majority of the conversation around a video occurs in the first day or two after it is posted and drops off rapidly thereafter.

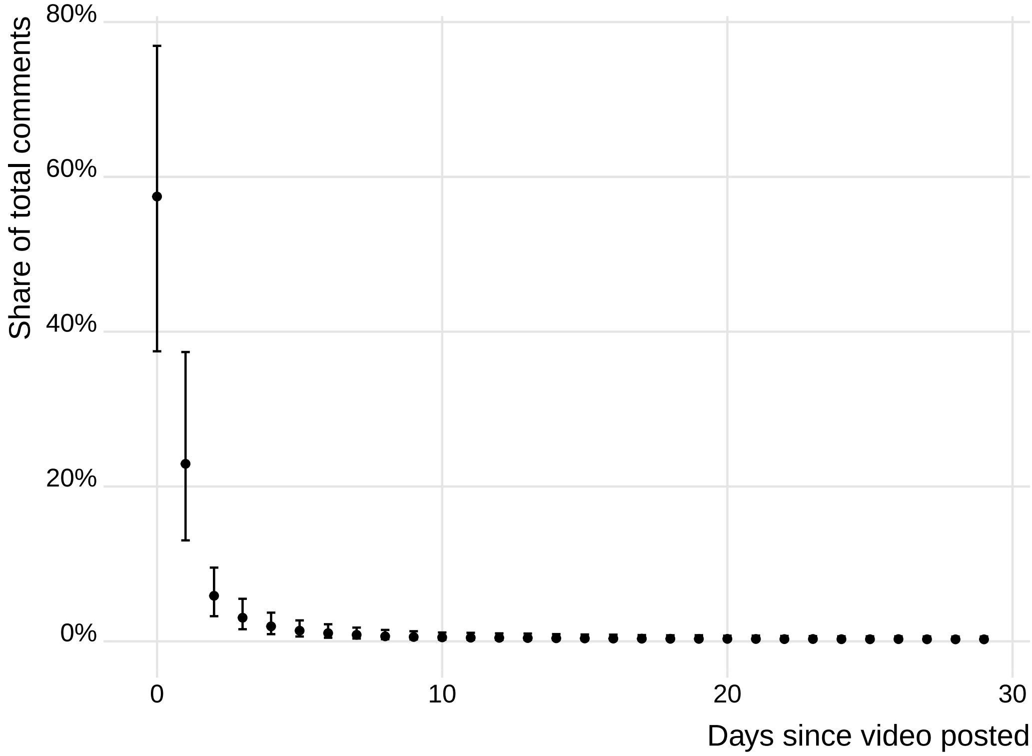

Figure 9. Comment decay visualized as the share of total comments a video receives by day since posting, calculated over a random sample of 20,000 videos. Points indicate median and bars indicate interquartile range.

Commenters

Thus far our descriptive analysis of politics channels on YouTube has consistently reinforced one simple conclusion—inequality. Channel size, video production, and comments are all highly skewed. Do these patterns persist when we look at the users themselves? More specifically, what share of total comments are written by what share of total commenters?

With over 320 million comments observed since January 1st, 2020, and more than

17 million unique commenters, a completely equal distribution of comments to commenters would be roughly 19 comments per user, corresponding to the red 45 degree line in Figure 10. Instead, we observe a very unequal distribution, with roughly 40% of users commenting only once, and 50% of total comments being written by just over 2% of all commenters. At the risk of repeating an obvious but important point: the overwhelming majority of viewers never comment at all, and we can only observe their existence in aggregate. Munger (2024) finds a gradually increasing ratio of comments to views over the lifespan of YouTube, but even in 2023, the median ratio of comments to views is only 2%.

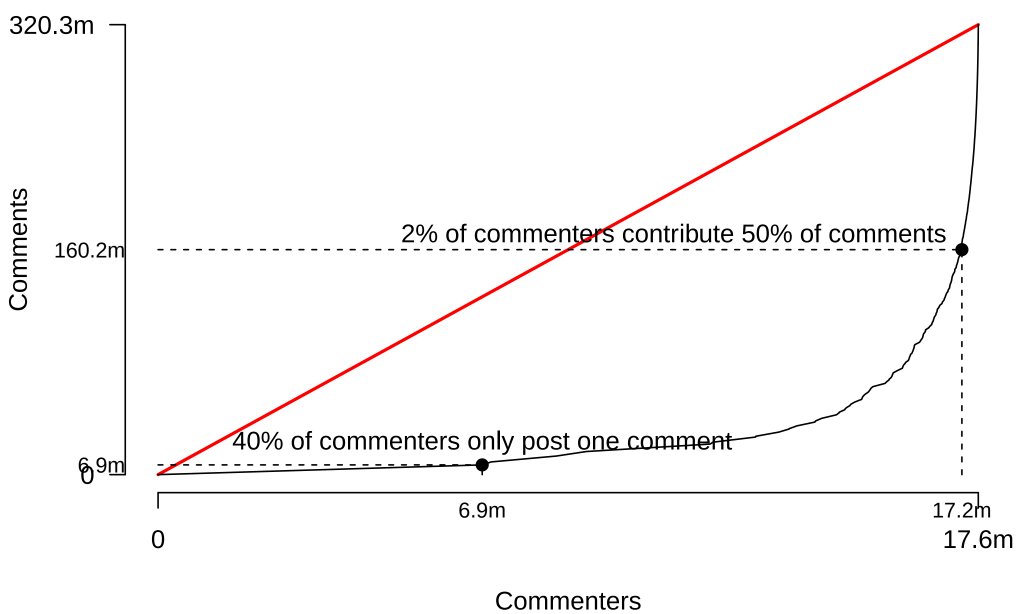

Figure 10. Observed comment Lorenz Curve.

Of course, unequal participation in commenting does not necessarily translate into unequal comment content. If those super-users who comment hundreds of times say the same things as those users who comment only once, the nature of the conversations in YouTube comments is theoretically unaffected by this unequal participation. However, if the super-users differ substantially in what they discuss and how they communicate, one might imagine that the passive majority of YouTube users are subjected to the unrepresentative conversations of a vocal minority.

We characterize how users communicate by combining a random subset of comments with the `peRspective` API. This algorithm scores comments by how likely they are to be reported as toxic. The platform is built from the model developed by Wulczyn et al. (2017); they define toxicity as "a rude, disrespectful, or unreasonable comment that is likely to make you leave a discussion." The model is a neural network originally trained on Wikipedia comments that had been hand-labeled by editors. While the model has some limitations, it is an industry standard and has been incorporated into applied social science work (Munger, 2020). Still, no classifier can be applied to a context distinct from where it was trained without experiencing some degradation in performance, and there is work illustrating certain limitations of peRspective (Muddiman et al., 2019). We hand-label about 500 comments as toxic (1) or not toxic (0) to validate the scores from the API and find a fairly high AUC metric of 0.84. Appendix C provides more details and additional metrics.

We show that toxicity is higher among more active commenters, highlighting that the vocal minority of users influence the tone of conversations had in YouTube's comments. Figure 11 categorizes the sample of almost 7 million commenters into how many comments they make, ranging from the 2.4 million users who comment only once in our data, to the 863 who comment more than 2,000 times. As illustrated, the average toxicity among those who rarely comment is substantially lower than the average toxicity of those who are super users. Importantly, the total comments written by both these groups is approximately equal: 2.46 million comments written by the 2.46 million commenters who do so only once, versus to 2.61 million comments written by the 863 users who comment more than 2,000 times.

Commenter-Channel Network

Thus far we have described each unit in the YouTube ecosystem—channels, videos, comments, and commenters—in isolation. However, we know that each of these units interacts with each other, and that it is these interactions that breath life into this platform. To characterize the platform writ large, we turn to network analysis, where channels are connected to one another by sharing the same commenters on their videos. Or in the language of network analysis, nodes

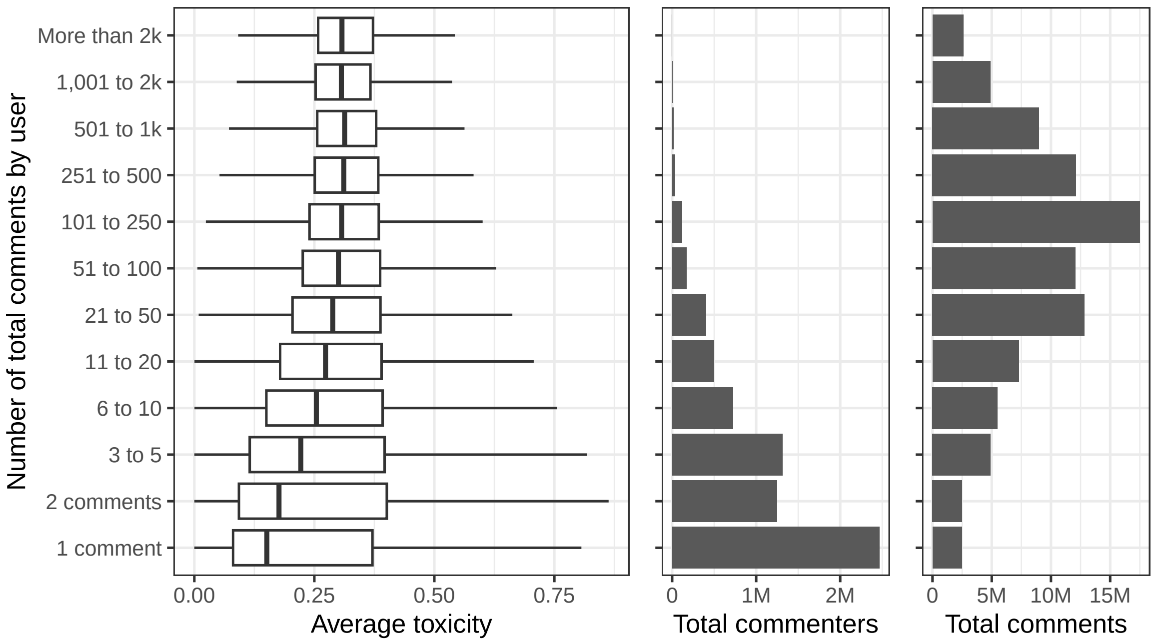

Figure 11. Distribution of comment toxicity (left-facet, x-axis) by the number of comments written by author y-axes, with distributions of total commenters and total comments provided for reference.

are channels linked by undirected edges weighted by the number of comments written by the same commenter. For example, there are 268,778 comments written by users who comment on both Ben Shapiro and Steven Crowder's videos. This constitutes 40% of all comments written on Steven Crowder's videos, and 25% of all comments written on Ben Shapiro's videos. We illustrate a sample of 6 channels ranging from conservative (Steven Crowder and Ben Shapiro) to liberal (David Pakman Show and Sam Seder) in Figure 12. Cells are colored by the raw number of comments that are written by overlapping commenters, and labeled according to what proportion of total comments these represent on the channels indicated on the x-axis.

We can use this network to estimate a channel-level measure of ideology via dimension reduction in a manner analogous to methods described in Poole and Rosenthal (1985); Barberá (2013). The underlying assumption is one of homophily, meaning that a liberal commenter will not watch (or comment on) a conservative video. If this assumption holds, and if the underlying

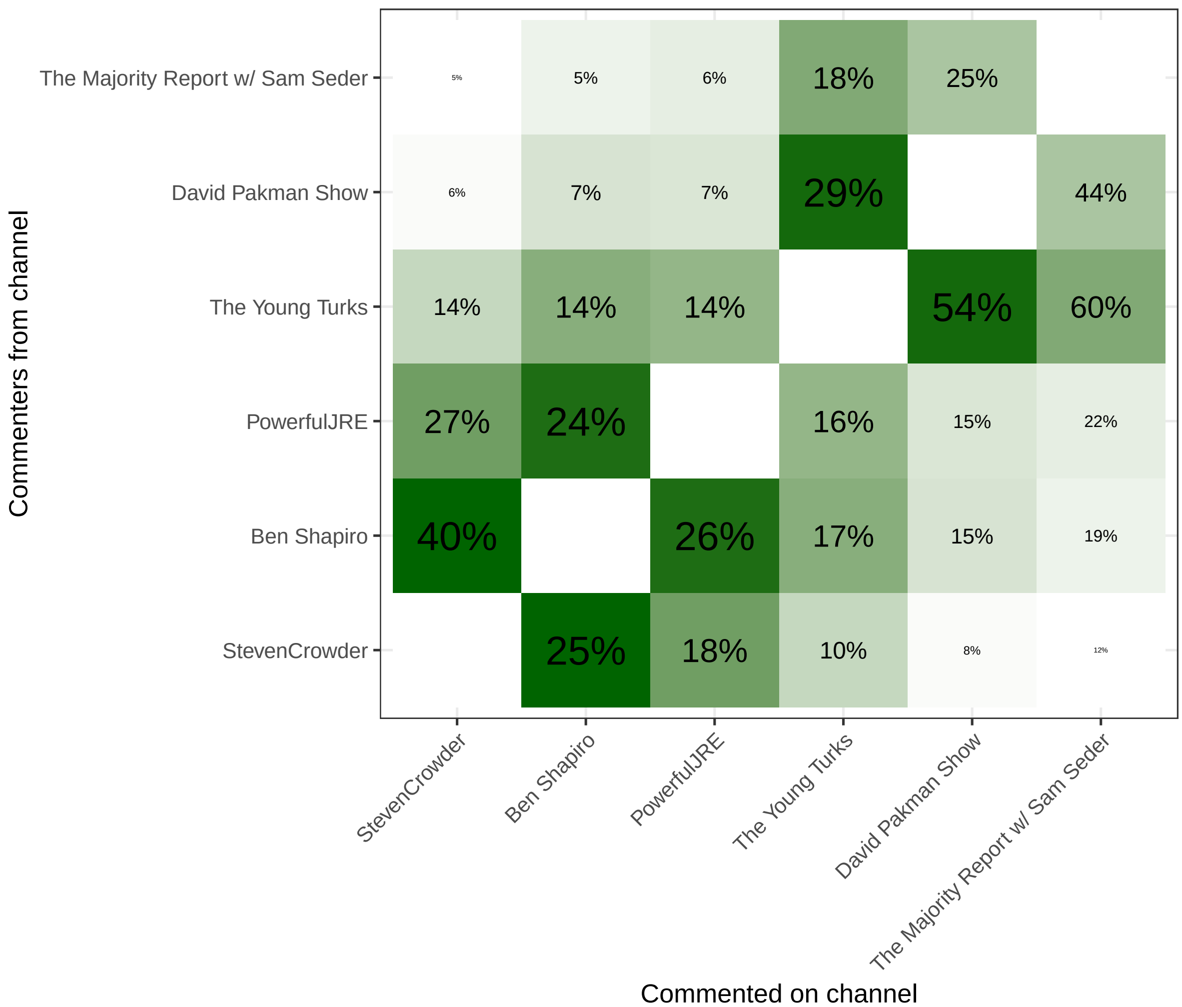

Figure 12. Heatmap of most densely connected channels, shaded by total number of comments written by shared commenters.

latent space is sufficiently low dimensional, we can theoretically use a Bayesian approach to ideology estimation as described in Clinton et al. (2004). However, given the enormous size of our data, even the fast expectation maximization solutions proposed by Imai et al. (2016) are infeasible. Instead, we rely on correspondence analysis as demonstrated in Lai et al. (2022) and visualize the first two dimensions in Figure 13.

To help validate this approach, we labeled all channels prior to analysis according to

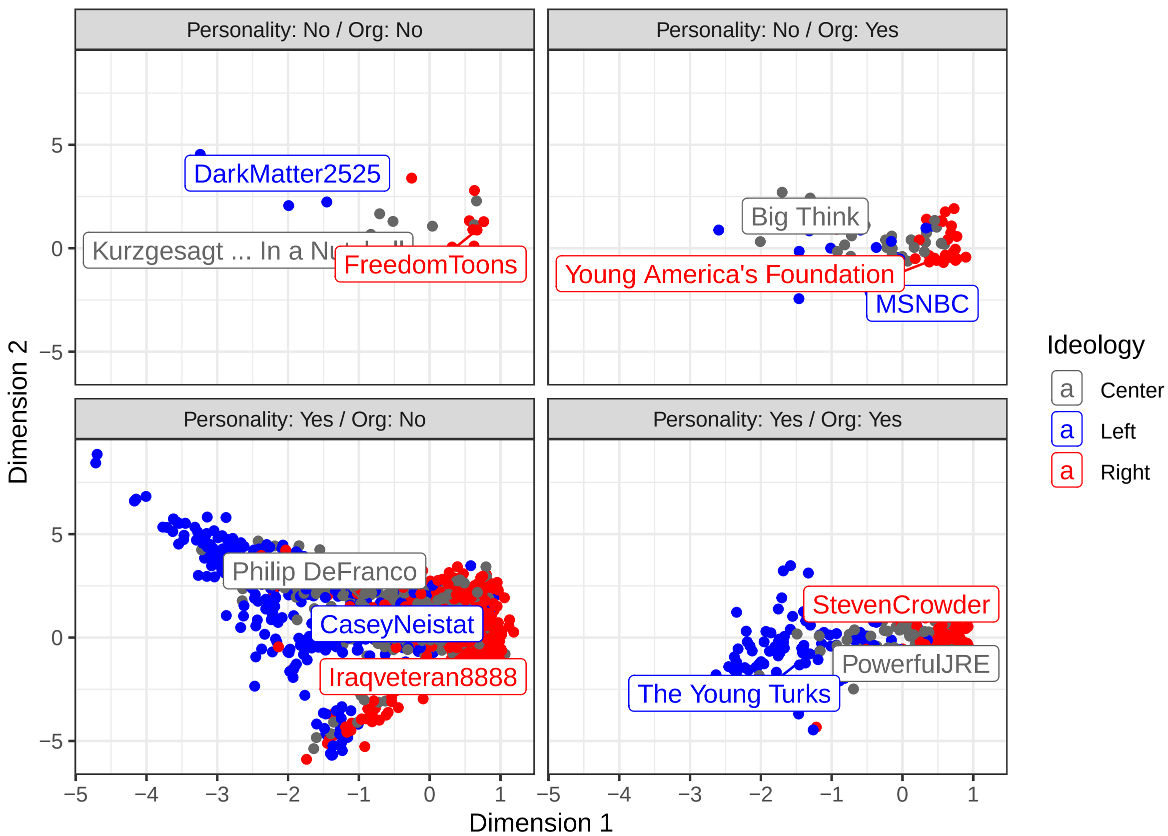

Figure 13. Scatterplot of channels along the first two dimensions of correspondence analysis. Channels were categorized by the authors according to three characteristics based on domain knowledge: whether the channel is personality driven (yes/no); whether the channel is owned by a single user or is part of a larger organization (yes/no); and whether the channel is ideologically progressive (blue), centrist (gray), or conservative (red). Top 3 channels by total views labeled in each facet.

the following three characteristics.

First, we asked whether the channel was "personality" driven or not. By this we mean whether the content features the face and identity of one or more content creators, whose perspective is part of the channel's content. For example, The Young Turks regularly features Cenk Uygur and Ana Kasparian as its primary hosts, who are regularly on camera discussing news and politics. We chose to code this attribute based on Munger and Phillips (2022)'s argument that parasocial relationships and communities can develop more easily on YouTube than other mediums. One or more recognizable hosts makes such a dynamic possible. Through coding this feature, it is possible to examine whether qualitative and quantitative differences in engagement dynamics exist between personality-driven channels and those that are not.

Second, we asked whether the channel was run by a single YouTube user or was part of a larger organization. Such channels are often more capable of uploading large amounts of content in small time intervals, and have resources to pay for production value and entertaining guests. These channels are likely to have different engagement dynamics as well, and coding them makes it possible to examine those dynamics. Despite being associated with Joe Rogan and relying heavily on his personality, the PowerfulJRE channel is run by an organization with employees who handle production, advertising sales, and marketing.

Third, we manually labeled every channel for its ideology, separating each into center, left, and right bins.

These labels help put structure on the results of our correspondence analysis, and serve as validation for our interpretation of the first dimension (x-axis in Figure 13) as capturing ideology. As illustrated, those channels that we label as conservative largely appear clustered together in a red mass on the right of the x-axes, while those we label as liberal appear as a blue smear ranging to the left. The second dimension is slightly harder to make sense of, although we note that the greatest variation along this dimension is found among the single users who have personality-based channels. One potential interpretation of this dimension then is that it captures professionalism or extremism.

One note of caution in interpreting these findings: YouTube, like all social media platforms, has struggled in recent years to control automated comments and other types of "bot" behavior. While early scholarship on bots focused on Twitter (Varol et al., 2017; Ferrara et al., 2016), largely because of better data availability, bot activity is now widespread on all social media platforms, and (with new generative AI technology) increasing hard to detect (Ferrara, 2023). For political content in particular, nation state actors have used bots and other forms of platform manipulation to shape political narratives (Howard et al., 2018; Yang and Menczer, 2024). While YouTube seems to initially to have been a less attractive target for bot-related activity, recent research has found large, sophisticated networks of "social scam bot" commenting activity on YouTube. (Na et al., 2023; Li et al., 2024; Tripathi et al., 2024).

Our research mostly predates the newer and more convincing generation of AI-powered YouTube scam bots, and we find little evidence of automation among the commenters we have examined closely. But we cannot make definitive conclusions about the impact (or lack thereof) of social bots on the commenting activity and patterns of viewership we observe.

Conclusion

YouTube is a major source of political and social news and ideas for many Americans. Our research provides a rich descriptive analysis of the channels, videos, comments, and commenters that occupied this space from the beginning of YouTube to the present. We find evidence of enormous inequality in participation across all units of analysis. The largest channels are orders of magnitude more popular than the remainder, the most popular videos receive an order of magnitude more views than others, and more than half of all comments are the most active 2% of commenters.

Based on these findings, one might conclude that YouTube is an unimportant component of the social ecosystem of American democracy, a backwater of the internet home to a few hundred thousand devotees. And this conclusion would be true if we were interested in interpreting YouTube as a reflection of *participation* in the discourse of politics online. However, YouTube is first and foremost an information environment in which an increasingly large proportion of the American public receives their news, formulates their opinions, and takes their cues.

These adults (and, despite YouTube's efforts to the contrary, children) are exposed to a vanishingly small minority of voices, be these voices embodied in channels, videos, or comments. These voices, numbering in the hundreds of thousands, are considerably more diverse than traditional media formats like newspapers and especially cable news and talk radio.

But this is part of what makes YouTube, like other network-driven platforms, so concerning for democracy. The veneer of openness and egalitarianism may deceive the viewer into thinking that they are tapping directly into a broader portion of the public than they in fact are. Cable news, for all its flaws, makes clear that the faces on the screen every night are professionals. But while YouTube channels and commenters may be more numerous, in key ways they are *less* representative of the public than traditional media. We find this criticism of YouTube—and of news consumption on social media more broadly—more compelling than some better-publicized critiques based on "fake news" or "echo chambers" that have found little empirical support (Allen et al., 2020; Guess et al., 2018).

We agree with the orientation towards the "Social Media Prism" that Chris Bail advances in his recent book (Bail, 2021). The primary distortions introduced by social media are not in how users perceive the news, but rather in how the audience perceives the *public* and in how the creators perceive themselves. The former problem is based on the "majority illusion" or "false consensus effect" created by the ability to observe so many ostensibly distinct "signals" of belief in the form of likes and comments on the YouTube videos (Lerman et al., 2016; Ortoleva and Snowberg, 2015). The impression that hundreds of thousands of people agree with the content of a given video has the effect of elevating fringe viewpoints. The crucial mechanism — entirely consistent with our finding about the staggering inequality in commenter activity — is what Ortoleva and Snowberg (2015) call "correlational neglect": observers of these signals are unable to adequately discount the ostensibly independent evidence of popularity which is in fact highly correlated, created by the same group of activists.

And this assumes that these commenters are merely enthusiastic. It remains an open question whether or not this vocal minority is already co-opted by strategic political entrepreneurs. Recent descriptions of the ecosystem beyond the scope of this research suggest that the period between 2014 and 2020 saw the rapid growth and subsequent recession of an array of voices organized around a rejection of mainstream politically correct narratives (Munger and Phillips, 2022). With the exception of a few amateur YouTube political commentators who have achieved stable levels fame and cash flow, the early independent vloggers have largely been supplanted by channels with ties to larger media organizations.

References

Al Nashmi, E., North, M., Bloom, T., and Cleary, J. (2017). 'boots on the ground?': How international news channels incorporate user-generated content into their youtube presence. *International Communication Gazette*, 79(8):746–768.

Allen, J., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14):eaay3539.

Aridor, G., Jiménez-Durán, R., Levy, R., and Song, L. (2024). The economics of social media.

Bail, C. (2021). *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing.* Princeton University Press.

Barberá, P. (2013). Birds of the same feather tweet together. bayesian ideal point estimation using twitter data. *Proceedings of the Social Media and Political Participation, Florence, Italy*, pages 10–11.

Barberá, P., Wang, N., Bonneau, R., Jost, J. T., Nagler, J., Tucker, J., and González-Bailón, S. (2015). The critical periphery in the growth of social protests. *PloS one*, 10(11):e0143611.

Bärtl, M. (2018). Youtube channels, uploads and views: A statistical analysis of the past 10 years. *Convergence*, 24(1):16–32.

Burgess, J. (2018). *YouTube: Online video and participatory culture.* John Wiley & Sons.

Cayari, C. (2011). The youtube effect: How youtube has provided new ways to consume, create, and share music. *International journal of education & the arts*, 12(6):n6.

Chae, S. W. and Lee, S. H. (2024). Where do cross-cutting discussions happen?: Identifying cross-cutting comments on youtube videos of political vloggers and mainstream news outlets. *Plos one*, 19(5):e0302030.

Clinton, J., Jackman, S., and Rivers, D. (2004). The statistical analysis of roll call data. *American Political Science Review*, 98(2):355–370.

Cowen, T. (2018). Eric schmidt on the life-changing magic of systematizing, scaling, and saying "thanks". *Medium. https://medium.com/conversations-with-tyler/eric-schmidt-tyler-cowen-google-ec33aa3e6dae.*

Duffy, P. (2008). Using youtube: Strategies for using new media in teaching and learning. In *Enhancing learning through technology: research on emerging technologies and pedagogies*, pages 31–43. World Scientific.

Ferrara, E. (2023). Social bot detection in the age of chatgpt: Challenges and opportunities. *First Monday*, 28(6).

Ferrara, E., Varol, O., Davis, C., Menczer, F., and Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7):96–104.

Gerbaudo, P. (2018). The digital party: Political organisation and online democracy.

Gerring, J. (2012). Mere description. *British Journal of Political Science*, 42(4):721–746.

Guess, A., Nyhan, B., Lyons, B., and Reifler, J. (2018). Avoiding the echo chamber about echo chambers: Why selective exposure to like-minded political news is less prevalent than you think. *Knight Foundation White Paper*.

Hosseinmardi, H., Ghasemian, A., Clauset, A., Rothschild, D. M., Mobius, M., and Watts, D. J. (2020). Evaluating the scale, growth, and origins of right-wing echo chambers on youtube. *arXiv preprint arXiv:2011.12843*.

Howard, P. N., Woolley, S., and Calo, R. (2018). Algorithms, bots, and political communication in the us 2016 election: The challenge of automated political communication for election law and administration. *Journal of information technology & politics*, 15(2):81–93.

Imai, K., Lo, J., and Olmsted, J. (2016). Fast estimation of ideal points with massive data. *American Political Science Review*, 110(4):631–656.

Inwood, O. and Zappavigna, M. (2023). Conspiracy theories and white supremacy on youtube: Exploring affiliation and legitimation strategies in youtube comments. *Social Media+ Society*, 9(1):20563051221150410.

Karpf, D. (2012). Social science research methods in internet time. *Information, Communication & Society*, 15(5):639–661.

Karpf, D. (2016). *Analytic activism: Digital listening and the new political strategy*. Oxford University Press.

Khan, M. L. (2017). Social media engagement: What motivates user participation and consumption on youtube? *Computers in human behavior*, 66:236–247.

Kim, S., Yalcin, O. F., Bestvater, S. E., Munger, K., Monroe, B. L., and Desmarais, B. A. (2020). The effects of an informational intervention on attention to anti-vaccination content on youtube. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 949–953.

Klinenberg, D. (2024). Does deplatforming work? *Journal of Conflict Resolution*, 68(6):1199–1225.

Lai, A., Brown, M. A., Bisbee, J., Tucker, J. A., Nagler, J., and Bonneau, R. (2022). Estimating the ideology of political youtube videos. *Political Analysis*, pages 1–16.

Lerman, K., Yan, X., and Wu, X.-Z. (2016). The" majority illusion" in social networks. *PloS one*, 11(2):e0147617.

Lewis, R. (2018). Alternative influence: Broadcasting the reactionary right on youtube. *Data & Society*.

Li, X., Rahmati, A., and Nikiforakis, N. (2024). Like, comment, get scammed: Characterizing comment scams on media platforms. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, USA. Internet Society.

Maddox, J. and Creech, B. (2021). Interrogating lefttube: Contrapoints and the possibilities of critical media praxis on youtube. *Television & New Media*, 22(6):595–615.

McGrady, R., Zheng, K., Curran, R., Baumgartner, J., and Zuckerman, E. (2023). Dialing for videos: A random sample of youtube. *Journal of Quantitative Description: Digital Media*, 3.

Miller, D. T. (2021). Characterizing qanon: Analysis of youtube comments presents new conclusions about a popular conservative conspiracy. *First Monday*.

Mowlabocus, S. (2020). 'let's get this thing open': The pleasures of unboxing videos. *European Journal of Cultural Studies*, 23(4):564–579.

Muddiman, A., McGregor, S. C., and Stroud, N. J. (2019). (re) claiming our expertise: Parsing large text corpora with manually validated and organic dictionaries. *Political Communication*, 36(2):214–226.

Munger, K. (2020). Don't @ me: Experimentally reducing partisan incivility on twitter. *Journal of Experimental Political Science*, page 1–15.

Munger, K. (2023). Temporal validity as meta-science. *Research & Politics*, 10(3):20531680231187271.

Munger, K. (2024). *The YouTube Apparatus*. Cambridge University Press.

Munger, K. and Phillips, J. (2022). Right-wing youtube: A supply and demand perspective. *International Journal of Press/Politics*.

Na, S. H., Cho, S., and Shin, S. (2023). Evolving bots: The new generation of comment bots and their underlying scam campaigns in youtube. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, pages 297–312.

Nielsen, J. (2006). Participation inequality: Encouraging more users to contribute. Accessed: 2024-11-19.

Norton, S. and Shapiro, J. (2024). How to better study—and then improve—today's corrupted information environment. *Bulletin of the Atomic Scientists*.

Ørmen, J. and Gregersen, A. (2023). Towards the engagement economy: interconnected processes of commodification on youtube. *Media, Culture & Society*, 45(2):225–245.

Ortoleva, P. and Snowberg, E. (2015). Overconfidence in political behavior. *American Economic Review*, 105(2):504–35.

Panciera, K., Halfaker, A., and Terveen, L. (2009). Wikipedians are born, not made: a study of power editors on wikipedia. In *Proceedings of the 2009 ACM International Conference on Supporting Group Work*, pages 51–60.

Poole, K. T. and Rosenthal, H. (1985). A spatial model for legislative roll call analysis. *American Journal of Political Science*, pages 357–384.

Rauchfleisch, A. and Kaiser, J. (2024). The impact of deplatforming the far right: an analysis of youtube and bitchute. *Information, Communication & Society*, pages 1–19.

Rieder, B., Coromina, O., and Matamoros-Fernandez, A. (2020). Mapping youtube: A quantitative exploration of a platformed media system. *First Monday*, 25(8).

Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *science*, 311(5762):854–856.

Tolson, A. (2010). A new authenticity? communicative practices on youtube. *Critical discourse studies*, 7(4):277–289.

Topinka, R. (2024). 'conspiracy theories should be called spoiler alerts': Conspiracy, coronavirus and affective community on russell brand's youtube comment section. *New Media & Society*, page 14614448241237489.

Tripathi, A., Bharti, K. K., and Ghosh, M. (2024). An efficient algorithm for exploitative monetization scam video detection over social media platforms. *Security and Privacy*, page e474.

Tufekci, Z. (2017). *Twitter and tear gas: The power and fragility of networked protest*. Yale University Press.

Van Mierlo, T. et al. (2014). The 1% rule in four digital health social networks: an observational study. *Journal of medical Internet research*, 16(2):e2966.

Varol, O., Ferrara, E., Davis, C., Menczer, F., and Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 280–289.

Weiss, B. (2020). Did i just get yanged? *The New York Times*.

Wu, F., Wilkinson, D. M., and Huberman, B. A. (2009). Feedback loops of attention in peer production. In *2009 International Conference on Computational Science and Engineering*, volume 4, pages 409–415. IEEE.

Wulczyn, E., Thain, N., and Dixon, L. (2017). Ex machina: Personal attacks seen at scale. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1391–1399. International World Wide Web Conferences Steering Committee.

Yang, K.-C. and Menczer, F. (2024). Anatomy of an ai-powered malicious social botnet. *Journal of Quantitative Description: Digital Media*, 4.

Data Collection Overview

We attempted to scrape 2,940 total channels over the course of our data collection. Call this the final “target” at the end of the data collection period. Our initial ”target” was 1,712 when we began in Summer 2020, based on a list from a previous project. 1,438 of these were still active in Summer 2020.

The difference between our final and initial targets is the product of adding many new channels in October of 2021, as well as our decision to drop 139 inactive channels from the original target list.

There is also a difference between our target list and our “effective” list, which comprises 2,629 total channels that were still active from our final target list as of the final day of scraping.

Finally, there is the difference between our target list and the full data, which comprises 2,905 total channels that we were able to collect data on at least once over the course of this three year period.

Table 1: Timing, Target and Effective Samples

Date	Target Sample	Effective Sample
Pre 10/2021	1,712	1,438
Post 10/2021	2,940	2,905
Final day	2,940	2,629

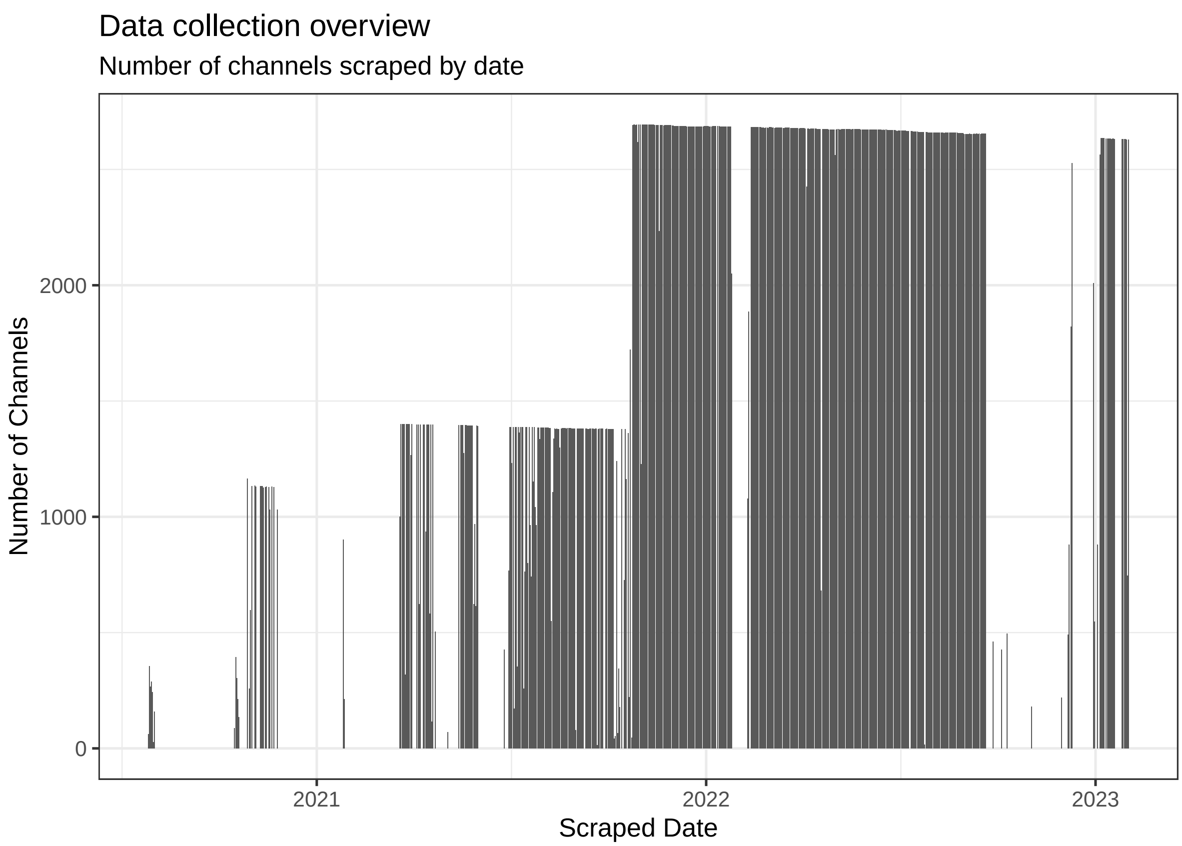

Figure 14. Data collection overview

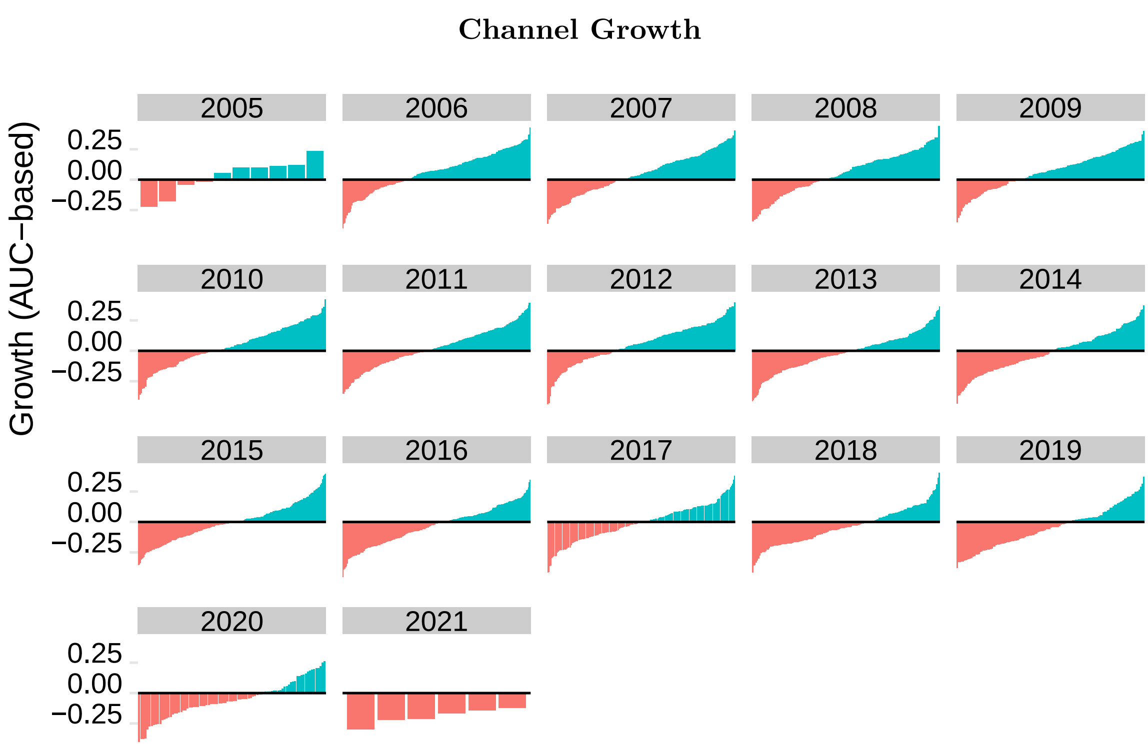

Manual Validation of the Toxicity Metric

As we mention in the main text, we use the Perspective API for characterizing the toxicity levels in comments. Because the algorithm classifies comments by how likely they are to be reported as toxic, it produces a value between 0 and 1 for each input text. While one can convert these probabilities into "labels" for each comment as toxic or not-toxic, we choose to retain the probabilities as a continuous measure and base our analysis on that.

In order to validate the results we get from the API, we—distributed among the authors—manually labeled 500 random comments. We labeled 132 comments as toxic and the remaining 368 as non-toxic. The API returned probability scores for 488 comments; we lost 12 to API limits or other errors. The metrics reported below are based on the 488 comments.

We have a situation in which we have two sets of scores: a probability between 0 and 1 reported by the API and a label of toxic (1) or non-toxic (0) produced by our manual annotation. We first use the area under the ROC (Receiver Operating Characteristic) curve (AUC), as it is a commonly used metric of classifier performance that does not depend on a single threshold. In other words, we can use it without committing to a certain probability threshold—such as 0.5—above which comments should be classified as toxic, since AUC measures performance across all thresholds at once. An AUC score of 1 is a perfect classifier. A score of 0.5 is the performance of a random classifier with no ability to discriminate between the two classes. With our data, we get an AUC of 0.84, which is fairly high. The AUC is also a good measure because none of the analysis we do depend on the choice of a threshold, since we use the actual probabilities directly. To put another way, we need to ask the Perspective model to lower its threshold and more liberally label comments as toxic for the model to have the most agreement with our hand labels.

We can also calculate accuracy, precision, recall, and F1 scores by using a threshold. When we use the most commonly used threshold of 0.5, we get an accuracy of 0.79, precision of 0.71, recall of 0.38, and an F1 score of 0.497. The high precision, but low recall, suggests that our threshold for classifying something as toxic seems to be lower than that of the Perspective API. In fact, the F1 score is maximized—at 0.65—if 0.38 is used as the threshold. Accuracy is maximized—at 0.82—if 0.4 is used as the threshold. So the Perspective model's tendency to label a comment as toxic seems to be slightly lower than ours.

Overall, we find Perspective's toxicity scores to be valid and useful given the fairly high AUC score, which is an appropriate metric for our use case where the probability itself, not the label, is used.