COMPUTATIONAL COMMUNICATION RESEARCH 7.1 (2025) 1–42
HTTPS://DOI.ORG/10.5117/CCR2025.1.4.HORN

Supply and Demand on Alt-Tech Social Media: A Case Study of BitChute

Benjamin D. Horne
School of Information Sciences, University of Tennessee, Knoxville, Knoxville, TN, USA

Myles Bowman
School of Information Sciences, University of Tennessee, Knoxville, Knoxville, TN, USA

Milo Z. Trujillo
Complex Systems Center, University of Vermont, Burlington, VT, USA

Maurício Gruppi
Department of Computing Sciences, Villanova University, Villanova, PA, USA

Cody Buntain
College of Information Studies, University of Maryland, College Park, MD, USA

Abstract
As media platforms continue to develop content moderation policies, alternative platforms have emerged as safe havens for deplatformed content. As these alternatives to major media platforms grow, the importance of understanding their role in the media ecosystem grows too. In this paper, we perform a longitudinal study of the content dynamics of one such alternative media platform, BitChute. BitChute is an alternative video-hosting site similar to YouTube. We first theorize what technological affordances may drive the supply and demand of content on BitChute. We then test those theories through an analysis of 6,363,596 videos from 82,162 channels, which were viewed 2,868,117,905 times, over 54 months. We find that BitChute's minimal content moderation drives much of the content supply and demand. Videos which were more offensive, certain, and covered commonly deplatformed topics were most popular. In particular, we find that BitChute fills a demand gap created by moderation policies on major media platforms around COVID-19 and - to a lesser extent - elections fraud. The most popular videos on the platform were re-uploaded videos that were banned by YouTube and Facebook. As a whole, our results suggest that BitChute's current role is less as a town square and more as a backup for deplatformed video content.

Keywords: alt-tech, new media, social media

Introduction

With increasing pressures on media companies to supress misinformation, hate speech, and far-right extremism on their platforms, the far-right has started a movement developing new infrastructure that facilitates banned and moderated content. This Alternative Technology (Alt-Tech) movement is a concern among researchers and practitioners due to its threat to intervention effectiveness, its potential to radicalize information consumers, and its ability to support the organization of extremist groups (Buckley & Schafer, 2022; M. Childs et al., 2022; Dehghan & Nagappa, 2022; M. Trujillo et al., 2020). As described by Buckley and Schafer (2022): "These alt-tech platforms have attempted to market themselves as unfettered, unmoderated areas which prioritise unlimited free speech... signalling an invitation for users to voice everything from unpopular opinions, to misinformation, to hate speech."

While studies of Alt-Tech platforms have been done - both quantitative description studies of content (Aliapoulios et al., 2021; M. Childs et al., 2022; Israeli & Tsur, 2022; Jakubik et al., 2023; M. Otala et al., 2021; Peucker & Fisher, 2023; Sipka et al., 2022; M. Trujillo et al., 2020; Zannettou et al., 2018) and more in-depth qualitative studies of policies, branding, and content (Buckley & Schafer, 2022; Dehghan & Nagappa, 2022; Donovan et al., 2019; Jasser et al., 2023; Kor-Sins, 2023; Wilson & Starbird, 2021) - examining these platforms longitudinally has received less study. Yet, understanding how content on these platforms has changed over time is necessary for defining their role in the broader media ecosystem and for clearly discerning the threat they pose to the spread of malignant content.

To this end, we study longitudinal content dynamics on one particular Alt-Tech platform, BitChute. Launched in 2017, BitChute is a social-video platform branded as an alternative to YouTube (M. Trujillo et al., 2020). Like research on other Alt-Tech platforms, early research on BitChute found that the platform contains large amounts of hate speech (M. Trujillo et al., 2020), discussion topics often centering around politics, and the platform has lax content moderation (Mahl et al., 2023; M. Trujillo et al., 2020). While interest in BitChute has been growing (M. Z. Trujillo et al., 2022), relatively little is known about how the platform has changed over time. BitChute has been studied considerably less than other Alt-Tech platforms like Gab and Parler, partially due to the difficulty of collecting data from the platform (M. Z. Trujillo et al., 2022).

To study this platform, we analyze the production of videos (called *supply* here forward) and views on those videos (called *demand* here forward) using a unique dataset of nearly all videos published on BitChute between June

2019 and January 2024, totalling 6,363,596 videos from 82,162 channels. In particular, we ask **RQ1:** *How has the supply and demand of videos on BitChute changed over time?* and **RQ2:** *What factors predict popularity of videos on BitChute?*

Before performing this exercise in large-scale quantitative description, we first ask why and how people use the platform to theorize what may drive popularity on the platform. As BitChute is a self-proclaimed alternative to major media platforms, it provides a different set of affordances that change how it is used and who uses it. The term *affordances* encompasses both specific features of the platform (i.e., ability to comment on a video), as well as "the kinds of communicative practices and habits they enable or constrain" (Bucher, Helmond, et al., 2018; Gibson, 1977; Jasser et al., 2023). As put by Jasser et al. (2023), affordances of individual platforms are important because "they help explain the shared practices of platform users". Using this framework, we explore, both theoretically and empirically, content supply and demand on BitChute. Inspired by Munger and Phillips (Munger & Phillips, 2022) theory of supply and demand on YouTube and Guinaudeau et al.'s theory of supply on Tikkok (Guinaudeau et al., 2022), this approach helps us better target and explain our descriptive analysis.

From this analysis, we found that there were broadly two types of popularity dynamics among content producers: *rich-get-richer* and *one-hit wonders*. Videos from channels that were older and were more popular in the past received more views. Overall, the distribution of views per channel was highly skewed, with only 1.8% of channels receiving 85% of the total views over 4.5 years. However, many of the most highly viewed videos on the platform were from *unpopular* channels. That is, the most viewed videos were re-uploads of deplatformed videos (i.e. permanently removed videos from major social media platforms) by channels that did not sustain audiences outside of a couple of popular mirrored videos. Mirroring deplatformed content was common among both groups of content producers, with some of the well-established channels receiving the most views on mirrored videos rather than originally created videos.

The content itself played a crucial role in popularity. We found that videos with titles that were more offensive and novel were viewed more. Further, videos with titles that expressed certitude were viewed more and videos that expressed confusion were viewed less. With regards to the popularity of certain topics, we found that BitChute filled a demand gap created by content moderation policies on major media platforms around COVID-19 and elections fraud. In October 2020, videos covering U.S. election fraud received approximately 25% of the total views on the platform. Similarly, in September 2021, COVID-19 videos received over 25% of all views on the platform. Across time, these two topics dominated both the number of videos produced (*supply*) and the number of views on videos (*demand*).

This study and its results are important for several reasons. First, this study solidifies BitChute's role as a home for deplatformed content. While this result is perhaps predictable given other smaller studies of Alt-Tech platforms, the results provide large-scale evidence of this role. Second, the results of this study highlight the impact of offline events on online engagement, adding to the growing literature about the relationship between the offline and online world (Hebbelstrup Rye Rasmussen & Petersen, 2023; Rice et al., 2022). Lastly, the results of this study add to the broader strand of research on the impact of social media affordances on content engagement and production (Bucher, Helmond, et al., 2018).

The rest of this paper is organized as follows: First, we theorize how the affordances of BitChute may influence the supply and demand of content. Second, we describe the unique dataset used and our methods of analysis. Third, we describe our empirical results and how those results align with our theory on supply and demand on BitChute. Lastly, we reflect on what these results tell us about the role of Alt-Tech media platforms in the larger media ecosystem.

What technological affordances make BitChute distinct from other social media platforms?

Here we theorize what makes BitChute distinct and how these distinctions may influence content production and popularity (summarized in Figure 1). We argue that BitChute represents two unique affordances: 1. it has *minimal content moderation*, and 2. it has *no personalized recommendation algorithms.*

Minimal Content Moderation

The lack of content moderation is baked into the branding and culture of the platform. BitChute's original mission statement was: "to put people and free speech first[1]", and many of the early content creators reflected this motto. While just *how* lax the content moderation is has changed over time due to pressures from companies controlling supporting infrastructures and governmental legal systems (e.g., such as removing terrorist recruitment

[1] https://web.archive.org/web/20170120061049/https://www.bitchute.com/

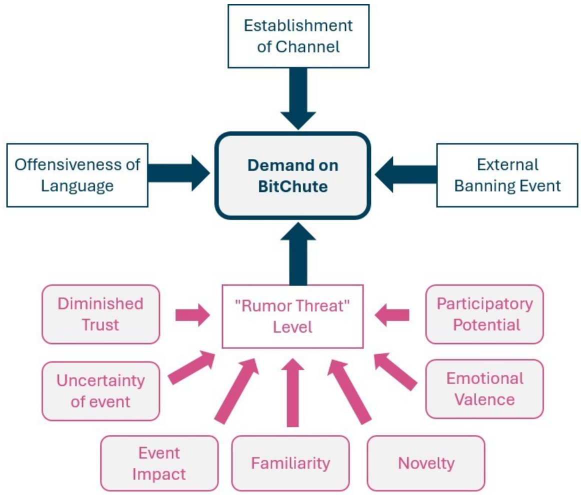

Figure 1: Theoretical model of content demand on BitChute, where the "rumor threat" level of content is borrowed from the event conditions and contextual features described in Spiro and Starbird (2023)'s rumor threat framework.

videos), the platform still harbors extremists (Buckley & Schafer, 2022; Marshall & Tanfani, 2022; M. Trujillo et al., 2020). In 2022, BitChute defended this practice in an interview with Reuters, stating "BitChute's North Star is free speech, which is the cornerstone of a free and democratic society" (Marshall & Tanfani, 2022).

This affordance has clear implications on the users and content demand on the platform. Broadly, the removal of content and content producers on Big-Tech platforms is done to curb the spread of offensive speech or false/misleading information (Innes & Innes, 2023; Jhaver et al., 2021; Thomas & Wahedi, 2023). Past research has provided some evidence that when content creators are deplatformed from YouTube, BitChute is a common alternative for these creators due to its minimal content moderation (Buckley & Schafer, 2022; Rauchfleisch & Kaiser, 2024). Rauchfleisch and Kaiser (2021) found that between 2018 and 2019 similar channels to those that were banned from YouTube were created on BitChute. Similarly, earlier research on BitChute noted that several of the most viewed channels on the platform were previously banned from YouTube, although it was also noted that some other popular content creators maintained channels on both BitChute and YouTube (M. C. Childs & Horne, 2023; M. Trujillo et al., 2020). Further, specific banning events have been associated with movement to

BitChute (Buntain et al., 2021; M. Childs et al., 2022). Given that minimal content moderation is a defining feature of the platform and a key reason for content creators to use BitChute over Big-Tech platforms like YouTube, it is likely that popular BitChute content will break Big-Tech policies and norms, and YouTube's policies in particular.

YouTube's content moderation policies cover a variety of areas, ranging from hate speech to misinformation to violence, and these policies have changed over time. Importantly, two areas of YouTube's moderation policy and practices were implemented during the peak of public interest in BitChute (i.e., see Figure 5 in (M. Z. Trujillo et al., 2022)): the Medical Misinformation Policy and the Elections Misinformation Policy.

The Medical Misinformation Policy states that YouTube removes content that "denies the existence of COVID-19 or that people have died from COVID-19" and "[c]laims that contradict health authority and World Health Organization guidance on safety, efficacy and ingredients of currently administered and approved vaccines" (YouTube, 2023b). While the current Medical Misinformation Policy was not formalized until 2023, YouTube began moderating medical content, particularly around COVID-19, much earlier. For instance, YouTube began showing fact-checks in search results in April 2020, CDC content was placed on the homepage in March 2020, and the infamous *Plandemic* anti-vaccination documentary was removed in May 5th, 2020 (Buntain et al., 2021; Yadav, 2021).

YouTube enacted moderation of elections content around the same time. In particular, following the 2020 U.S. presidential election, YouTube put a policy in place to remove videos that falsely claimed the election was fraudulent (Bond, 2023). The platform also prohibits false claims about election integrity in other countries like Brazil and Germany (Bond, 2023; YouTube, 2023a). The elections misinformation policies cover misinformation besides election integrity, such as misinformation on voter suppression, candidate eligibility, and "incitement to interfere with democratic processes" (YouTube, 2023a).

Taken together, the minimal content moderation of BitChute, the culture and branding of the platform, and the content demand gaps created by YouTube's content removal, we hypothesize that:

H1: *The more offensive content is, the more popular it will be.*

H2: *The most popular content will cover conspiracy theory topics commonly deplatformed by YouTube, such as COVID-19 and elections fraud.*

No Personalized Recommendation Algorithms

BitChute's relationship to other parts of the online ecosystem may not be the only driver of popularity on the platform, particularly later in the platform's life when content producers and communities have been established. Internally, BitChute is unique from other social video platforms in that it has no personalized content recommendation algorithms.

A key feature of Big-Tech televisual platforms, like YouTube, TikTok, and Instagram, are their algorithmic recommendations (Buntain et al., 2021; Davidson et al., 2010; Guinaudeau et al., 2022; Mosseri, 2023; Munger & Phillips, 2022; Zieringer & Rieger, 2023). All three platforms' algorithms can recommend content from anywhere on the platform, not just from subscribed-to content producers. Recommendations on each platform are personalized based on individual user's activity. While these platforms are different in many ways (e.g., see (Guinaudeau et al., 2022)), they all facilitate effortless content discovery, which can influence content popularity.

BitChute is notably different from YouTube, TikTok, and Instagram as it has no personalized content recommendation algorithm. Rather, there are a few simple pathways for users to find content on BitChute. Notably, the lack of recommendation algorithm likely changes how audiences and communities are built on the platform. For example, as phrased by Jasser et al. (2023): "the recommendation features of YouTube afford the far-right a vital tool to disseminate their ideology and build a community." These features do not exist on BitChute, which suggests that audiences may be built externally first. Other Alt-Tech video platforms have had similarly minimal features but have recently rolled out personalized recommendation algorithms of their own. For example, Rumble - another alternative to YouTube and arguably BitChute's primary rival in the Alt-Tech video space - moved from a homepage of staff-picked videos to personalized recommendations in early 2024. Other smaller Alt-Tech video competitors - such as Odysee - still do not have sophisticated personalized recommendation algorithms but do claim to use trending algorithms on their front page.

There are three main mechanisms to find content on BitChute. First, is the video stack. The video stack is a set of the most recently uploaded videos to BitChute featured on the default homepage. This stack uses no sorting or filtering, it is simply a ranking of videos by upload time. Hence, any channel and content from any topic can be featured at the top of the stack by sheer timing (i.e. uploading a video seconds before a user visits the page). Similarly, above the stack is a side-scrolling set of channels that are recently active.

There are two variations of the video stack available through clicking other tabs on the homepage: *popular* and *trending*. According to BitChute, the popular stack uses a combination of video views and channel subscription numbers that rotate every 24 hours. The trending stack is a ranking of videos by number of views over a day, week, or month time-frame, where the time-frame is selected by the user. Note that both of these variations are not personalized recommendations but instead are broad-based rankings of views and subscriptions over the whole platform (similar to what is used on Odysee). That is, all users receive the same recommendations (although broad-based trending feeds *can* increase engagement with specific content (Chan et al., 2024)). The stack that is selected by default when visiting the homepage has changed over time. Originally, in 2017, the default stack was the simple upload-time stack. By 2018, the default stack was changed to the popular stack. Across these changes in the defaults, the side-scrolling list of recently active channels at the top of the page has not changed.

The second mechanism to discover content on BitChute is search. The search mechanism is similar to, but again simpler than, search on platforms like YouTube and TikTok. On YouTube, search utilizes autocompletion, which makes query predictions about what you're looking for by utilizing what other people have already searched for. On TikTok, the search page is accompanied by suggested searches, again using the personalized recommendation system. On the other hand, BitChute's search does not use autocompletion or recommendations. In fact, currently BitChute's search appears to only do keyword matching in video titles, ranked by the number of views.

Other than the video stack, the side-scrolling channel stack, and search, there are no other internal mechanisms to discover new content from new producers on BitChute. One can of course subscribe to a channel's content, but subscriptions do not facilitate new content discovery outside of the channel. When watching a video on BitChute, there are no personalized recommendations on a side bar like on YouTube. The side bar does show other videos to watch, but these lists are videos from the same channel currently being watched. The "Playing Next" side bar automatically shows the next most recent video produced by the channel.

Given how limited these discovery paths are without recommendations, we hypothesize that the popularity of content producers on BitChute will be more skewed towards channels who already have many views than what we would expect from Big-Tech platforms like YouTube, and that videos from older channels will be more popular than videos from newer channels.

To the best of our knowledge, on YouTube, approximately 85% of the views go to only 3% of the channels (Bärtl, 2018). Further, we know that older channels have a higher probability of capturing a large number of views on YouTube (Bärtl, 2018). Theoretically, on BitChute, these two traits should be exacerbated by the limited discovery paths (i.e., a channel must first have an audience in order to be highly viewed). Saliently, personalized recommendation systems may have different impacts on channel viewership on different platforms, and platforms may utilize multiple recommendations algorithms and features. Evidence suggests that parts of YouTube's recommendation algorithm can create rich-get-richer effects, making already popular channels even more popular, while other parts equalize the view distribution (R. Zhou et al., 2016). On the other hand, from what little we know about TikTok's recommendation system, it occasionally recommends videos from posters with little prior engagement (Guinaudeau et al., 2022), which could widen the channel viewership distribution and allow young content producers to quickly build audiences. Still, we expect that the lack of discovery mechanisms on BitChute will facilitate higher viewership inequality than platforms with recommendation systems.

H3: *The inequality of views per channel on BitChute will be higher than YouTube.*

H4: *Videos from older, established channels on BitChute will be more popular than videos from newer, less established channels on BitChute.*

Content Characteristics Stemming from Affordances

Lastly, while the platform's affordances should play some role in what content is produced and engaged with, as argued above, we would be remiss to assume that the more general properties of the content do not play a role. Given that the "less-moderated nature" of alt-tech social media provides a favorable environment for rumors and conspiracy theories (Bellemare et al., 2020; Dehghan & Nagappa, 2022; Mahl et al., 2023; M. Trujillo et al., 2020), that rumors are highly engaged with on social media (Friggeri et al., 2014; Vosoughi et al., 2018), and that there is a content demand gap for rumors created by Big-Tech content moderation practices, it is likely that content in this category is more engaged with than other content.

Note, we are using the term *rumor* to encompass a larger set of terms used in the literature, such as misinformation, disinformation, and conspiracy theories. While there are differences between these terms, the frameworks built in the decades-old rumors literature can also capture these more granular terms (Spiro & Starbird, 2023). In particular, the "Rumor Threat" framework from Spiro and Starbird (2023) defines well-studied dimensions of viral rumors, which also can apply to misinformation and disinformation. In this framework, there are two groups of features that are particularly applicable to the case of popularity on BitChute: event conditions and contextual features of the content. Events where there is *uncertainty* and *diminished trust* in official information providers are likely to generate highly-engaged-with rumors (Allport & Postman, 1946; Shibutani, 1966; Spiro & Starbird, 2023). Further, there is evidence that external offline events can drive online hostility and engagement with hostile content (Hebbelstrup Rye Rasmussen & Petersen, 2023). If those rumors *impact* the lives of those who consume and spread them, as well as have *participatory potential* - that is, the ability for people to add their own evidence, experiences, and interpretations to the rumor - they are likely to spread further (Spiro & Starbird, 2023; Starbird et al., 2023). Additionally, rumors that convey strong *emotions* and are *novel*, yet also *familiar*, will likely spread further (Hasher et al., 1977; Kapferer, 2013; Spiro & Starbird, 2023). As explained by (Spiro & Starbird, 2023), "online communities that actively engage in conspiracy theorizing are poised to project a common set of ideas onto events' causes and impacts" - hence while the events may be new, familiar concepts may still be applied to rumor content. We hypothesize that these event, context, and content features will predict popularity on BitChute.

While we cannot operationalize all parts of this framework with the data used in this study (described in more detail in the Data Section below), we focus on the event-based topics of videos and the emotions portrayed in the content.

H5: *Content with a high "Rumor Threat" level will be more popular. More specifically, videos covering high "Rumor Threat" events with strong emotions will be more popular.*

For a more detailed description of the rumor literature and this specific framework, please see Spiro and Starbird, 2023.

Data

Now that we have established some expectations for content popularity and production on BitChute, we take a data-driven approach to describe changes in both over time. To do this, we utilize an extended version of the MeLa BitChute dataset (M. Z. Trujillo et al., 2022), which contains **6,363,596 videos** from **82,162 channels**, which where **viewed 2,868,117,905 times**, between **June 2019** and **January 2024**. The unique trait of this dataset is that it is not a random sample of data from the platform, but a near complete dataset over the time frame. This allows us to paint a complete, longitudinal picture of content dynamics. The data was collected using a custom collection engine that scrapes video metadata every 5 minutes for every newly uploaded video that the engine did not already collect. This live collection was made possible by the newly uploaded video page on BitChute, which was a stack off all videos uploaded to the platform. While we cannot precisely confirm the completeness of this dataset, the collection engine was shut down for only 8 days across the 54 months of collection.

There are several parts of this dataset that we use in our analysis. First, we use video **titles** to the describe and model the types of content that are produced and viewed on the platform. We group those titles under the **channels** which produce them. Channels on BitChute are much like channels on YouTube, representing a single content producer. In addition we use the **postdate** of videos to show how supply and demand change over time. The postdates are scraped directly from BitChute and correspond to the time in which a video is uploaded to BitChute. Lastly, we use video **views** which correspond to the number of times a video has been watched. There are a few important caveats to how views are captured on BitChute and are collected in this dataset. First, like YouTube, views capture the number of times a video has been played but not the number of unique users who viewed that video. Second, views are collected one week after the video is uploaded (M. Z. Trujillo et al., 2022). According to M. Trujillo et al., 2020, approximately 56% of views happen during the first week a video is on the platform. Thus, views after one week is a reasonable approximation of engagement with content but should be thought of as an under-estimate. If a video is no longer available, this is noted in the dataset. In preparation for our analysis, we filter out all videos that were deleted by the uploader or removed by BitChute within one week, as noted in the MeLa BitChute dataset.

While this data does allow us to describe changes in content supply and demand, we cannot perfectly operationalize every factor described in the theoretical model above (Figure 1). In particular, the factors that are external to the platform (i.e., “underlying conditions” like the diminished trust in official information sources and the uncertainty surrounding an event (Spiro & Starbird, 2023)) would require more granular qualitative study rather than the large-scale study that we undergo in this paper. However, using this data and large-scale quantitative analysis, we can explore the majority of the factors in the theoretical model, as well as describe how the platform has changed over time.

Methods

Figure 2: Example of data clustering within our multilevel model where videos are grouped under channels and the intercept can differ across channels.

Using this data, we compute a number of features which approximate some of the theoretical factors from Figure 1. Some of these features are computed on video titles. We choose to use video titles as our unit of analysis for several reasons. First, while BitChute does have an option for video descriptions by the content creator, these descriptions are often left empty or contain information not directly related to the video (e.g., "see more of my content at my website"). Second, video transcripts are not easily extracted from the data, particularly given the longitudinal timeline. While there are methods to extract transcripts from audio data, the overhead of collecting this audio data (which is not included in the MeLa BitChute dataset) and the computational cost of running transcriptions over that data is significant. On the other hand, video titles are readily available in the data set and appear to be reasonable descriptions of the videos. We perform several validation tests about this choice which are described in Appendix A.

Further, since the vast majority of BitChute content is in English (M. Trujillo et al., 2020) and some of our features can only be measured on English text, we used a reproduction of Apple's bi-directional LSTM models for language identification in short strings to filter the data to english videos only (Toftrup et al., 2021). This process left us with 4,963,243 videos of the 6,363,596 total videos in the dataset. Further validation of this process can be found in Appendix A.

Given the hierarchical structure of our data (videos under channels as shown in Figure 2), we use these features in a multilevel regression model, in which videos are clustered under channels. Mixed effects regression is an extension of a general linear model that takes into account the multi-level structure of the data. In our case, we account for variation of the intercept with this model, meaning that the model can have different intercepts per channel. Our dependent variable in the model is video views, while our independent variables are video-level and channel-level features. For an introduction to mixed effects modeling, please see Harrison et al., 2018.

In addition, we examine two metrics over time: market share and supply share. The market share of a channel is computed as the number of views the channel has accumulated divided by the total number of views on the platform. Similarly, the supply share of a channel is computed as the number of videos published by a channel divided by the total number of videos on the platform. We slice both of these metrics over time, showing how market share and supply shares change per month. Further, we also compute shares for groups other than channels, such as topics. So, a topic's market share in a given month is the number views on videos within that topic posted in that month divided by the total number of views within that month.

Features

Our features approximate the following factors: establishment of channel, offensiveness of language, emotional valence, novelty, and conspiracy language.

Establishment of Channel To approximate how established a channel is, we compute four features. First, we estimate how old a channel is by recording the month in which a channel's first video appears in our dataset (called *seniority* in Table 1). This feature is an integer where earlier in the timeline is higher, hence representing how old a channel is in our dataset by months. Second, we estimate a channel's prior popularity by computing a lagged market share. Specifically, we compute the market share of each channel in each week, month, and three month chunk of our data. For example, the market share of a channel during a given month is the number of views on videos they produced during that month divided by the total views on the platform in that month. In our video views model, for each video, we record the market share of the channel that produced that video last week, last month, and three months ago. In other words, how popular was the channel that produced this video in the past? Together, these features capture how established a channel is.

Offensiveness of Language To approximate how offensive the language use in the video title is, we compute two features. First, we assign each video a probability of being offensive by passing its title through a pre-trained RoBERTa model for offensive language identification (called *offensive* in Table 1)[2] (Barbieri et al., 2020). This model was pre-trained on 58M Tweets and fine-tuned for offensive language identification with the SemEval2019 OffensEval dataset (Barbieri et al., 2020; Zampieri et al., 2019). Given the model is trained on short-form, social media text, we expect reasonable performance on our short-form video titles.

Second, we utilize Linguistic Inquiry and Word Count (LIWC) to capture how prosocial the language in each video title is (called *prosocial* in Table 1). LIWC is an extensively validated, dictionary-based method to measure various psychological states from open-text. Within the 2022 version of LIWC, there are many dimensions captured, one of which is how prosocial the text is, which contains words such as care, help, thank, and please (Boyd et al., 2022). Given the simplicity of the method and its use across various media studies (Coppersmith et al., 2014; Eichstaedt et al., 2018), we expect the method to work well in our context.

Emotional Valence To estimate the emotional valence of content, we use the Paletz pre-trained Demux model for emotion recognition, which has been extensively tested on Twitter, YouTube, and Facebook data (Chochlakis et al., 2023a, 2023b; Paletz et al., 2024). This model has 27 emotion categories, making it more granular than other pre-trained emotion recognition models. These emotions include: anger, hate, contempt, disgust, embarrassment, love, admiration, sexual attraction, cuteness, wonder, pride, sadness, nostalgia, empathic pain, gratitude, envy, fear, confusion, surprise, happiness, excitement, amusement, hope, boredom, and fatigue. Additionally, there are two generic categories: positive other and negative other. More details on the model and the description of each emotion category can be found in Paletz et al., 2024.

In addition to emotion, we capture sentiment using VADER (Hutto & Gilbert, 2014). VADER is a rule-based model for sentiment analysis of social media text, which has been used across many prior studies (Borg & Boldt, 2020; Cheng et al., 2017; Ferrara et al., 2020; Horne et al., 2019). In particular, we use the compound sentiment score provided by VADER, which is a normalized, unidimensional measure of sentiment where -1 is the most negative and $+1$ is the most positive. While sentiment is not necessarily

[2]https://huggingface.co/cardiffnlp/twitter-roberta-base-offensive

emotion, VADER captures positive/negative direction and is built for the social media context.

As a robustness check, we also measure emotion using TweetNLP's pre-trained emotion recognition model (Camacho-Collados et al., 2022). This model is one of several transformer-based language models within TweetNLP, which are built from RoBERTa (Liu et al., 2019) and XLM-R (Conneau et al., 2019) checkpoints and trained on task specific data (Camacho-Collados et al., 2022). In the case of the emotion recognition task, the model uses data from the SemEval 2018 *Affect in Tweets* Task (Mohammad et al., 2018). Within the model version tested, there are nine emotion types captured: anger, anticipation, disgust, fear, joy, love, optimism, pessimism, and sadness. As described in Appendix A, we manually validate all methods and use the methods with the highest agreement among coders to use in our final analysis.

Novelty Next, we estimate the novelty of content using the prior market shares and supply shares of video topics in the dataset. These topics are extracted using a Structured Topic Model (STM), which is a generative model of stemmed word counts that allows for the use of metadata in the model (Roberts et al., 2014). In our case, we fit a model where time (a continuous variable estimated with a b-spline basis over discrete days) is the only covariate in the model, as we expect time to have an effect on video themes. As done in prior studies (Birim et al., 2022; M. Childs et al., 2022; Joseph et al., 2021; Yan et al., 2013), to choose the number of topics in the model (commonly called k), we use a grid search over values of k to find the model with the minimum perplexity and maximum log-likelihood. This grid search process leads us to choose 21 topics in the model.

After the final model is fit, we reduce the 21-dimensional vector into a one-to-one mapping of video to topic, where the topic with the largest proportion is assigned to the video. Given that our documents are short-text video titles, we expect this reduction to be reasonable, as much of the k-dimensional vector is sparse. Once videos are mapped back to single topics, we can compute both the market share and supply share of a given topic within a given time frame. Just as we do with channels, we compute the market and supply share of each topic in each week, month, and three month chunk of our data. In our video views model, for each video we record the market share and supply share of the topic that video is mapped to last week, last month, and three months ago. In other words, how popular and prevalent was the topic of the video in the past? The less prevalent the topic in the past, the more novel the topic is.

Conspiracy Language Lastly, inspired by Fong et al. (2021) and hypothesis **H2**, we use six categories from LIWC to capture conspiracy theorist language in video titles: *certitude*, *allnone*, *they*, *death*, *power*, and *relig*. The *certitude* category reflects "a degree of bravado, boasting of certainty that often reveals an insecurity or lack of truly verifiable, concrete information", with phrases such as really, actually, and of course (Boyd et al., 2022). The *allnone* category captures "a style of thinking that tends to be over-generalized and more extreme" with absolutist language like all, no, never, and always (Boyd et al., 2022). The *they* category captures outgroup identification with terms such as they, their, and them. Ingroup and outgroup identification has been shown in several studies to be important in conspiracy theories (Douglas et al., 2017). In addition, as described by (Fong et al., 2021), "many conspiracy theories heavily feature narratives that play into themes surrounding power, death, and religion", which are captured by the LIWC categories with corresponding names (i.e., LIWC category "relig" maps to religion). Together, we expect these six features to estimate how conspiracy-like content is on the platform and its relationship to popularity. Notably, the measurements for power, death, and religion may be fragile as they are closely tied to the topic of content. We perform additional topic-level analysis to better show what these measures are actually capturing. Furthermore, we are not arguing that these simple dictionary-based features can capture all of the conspiracy theories on the platform, but rather that - based on the work of Boyd et al., 2022 - these features will capture typical conspiracy theory language in the titles of videos. Additional validations for all metrics can be found in Appendix A.

Empirical Results

The overall inequality of views per channel on BitChute was higher than YouTube, but this changed over time.

Across the full dataset (4.5 years), we found that 85% of the views went to only 1.8% of the channels. As expected, the distribution of views on the platform closely followed a power-law distribution (Figure 3a). The average number of views per video was 34,908, while the median was 470. According to (Bärtl, 2018), across 10 years of YouTube data, 3% of the channels received approximately 85% of the engagement. Hence, this result provides some evidence supporting **H3**.

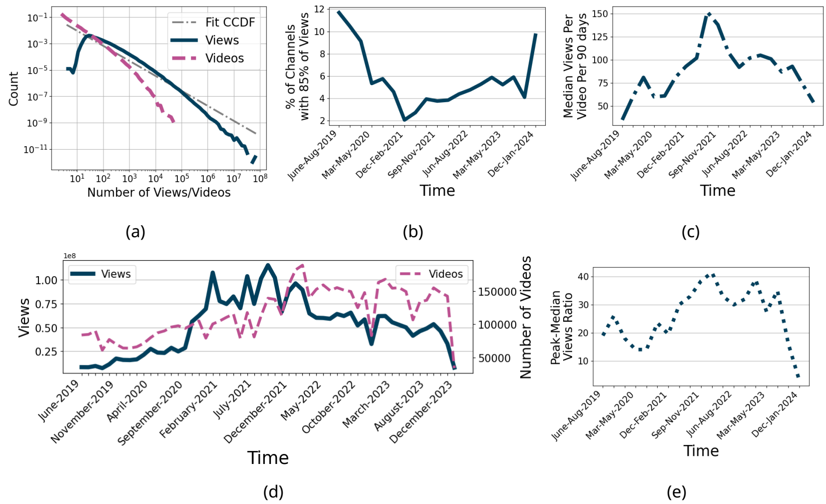

Figure 3: (a) The log-log distribution of views and videos per channel across the full dataset, demonstrating the power-law distribution of the platform as a whole. (b) Skewness of views per channel per 3 months, representing inequality in views across channels. From December 2020 to February 2021, only 2.05% of channels received 85% of the views. Over the full dataset, only 1.81% of channels receive 85% of the views. (c) Median views per video per 90 days, suggesting that while the views per channel were becoming more skewed in mid-2021, the views per video were becoming less skewed. (d) Total monthly views and number of videos per month, where the dashed line and right Y-axis represent the number of videos and the solid line and left Y-axis represent views. (e) Peak-Median Views Ratio per 3 months, as defined by (Guinaudeau et al., 2022). Peak-Median Views Ratio represents the inequality of viewership for videos within channels. The median Peak-Median Views Ratio over the full data is 40. Note the last 3-month time chunk only contains 1.5 months, hence the sudden increase in (b) and decrease in (c) and (e).

This inequality changed across time, just as it has on YouTube (Bärtl, 2018). In Figure 3b, we show the percentage of channels who received 85% of the views in each 90 day chunk of the data. Early in the platform's history (June to August 2019), the platform was much less skewed, with nearly 12% of the channels receiving 85% of the views (matching prior work on the platform's comment distribution (M. Trujillo et al., 2020)). Over the next two years the views inequality consistently grew, reaching a low of just under 2% of channels receiving 85% of the views. Since this point in early 2021, the platform has hovered between 4% and 6% of channels receiving 85% of the views per 90 days. When examining the average of this metric over 90 day windows, we found that on average 5.7% (x 5.2%, σ 2.6) of the channels received 85% of the engagement per 90 days. Note, this results seems paradoxical, as the average over three months windows (5.7%) is much larger than the same metric over the full 4.5 years (1.8%). However, this is due to the change in window size. When examining the same metric per day, there are many individual days where the percentage of channels who received 85% of the views is well below 1.8%.

This picture is made more complex when examining the median views per video over time, which peaks in late 2021 when only 4% of channels were receiving 85% of the views. Note, we are computing median views per video within 90 day windows. While we should expect that the video inequality overall to consistently decrease as new content is posted and consumed, this is not necessarily true within each consecutive window. This suggests that while channel inequality was increasing, video inequality was decreasing. These two trends coincide with the peak in views on the platform as a whole, shown in Figure 3d. On the surface, this results suggests that individual videos are playing an important role in both channel and platform popularity. We explore this implication further below.

Videos from channels that were older and had higher market share in the past were more popular

This high inequality suggests that channels themselves may impact video engagement. Our model results supported this notion. Within the channel feature group in Table 1, we see that the older a channel was, the more popular videos from that channel were. In addition, videos from channels who had higher market share in the prior 90 days were more popular. This result held for both 30 day windows and 7 day windows, reflecting a *rich-get-richer* effect. These two results broadly match results from (Bärtl, 2018) on YouTube, which indicated that older, more popular channels have "significantly higher probability to garner a large viewership". However, in part, these results are different than results from TikTok , where "viewership of videos are less dependent on a given accounts' number of followers (Guinaudeau et al., 2022)". While these are not perfect apples-to-apples comparisons, these results support **H4**.

Although the model suggests that a channel's age and prior popularity impact how much a video is viewed, the Interclass Correlation Coefficient (ICC) was only 0.214. This metric means that the number of times a video was viewed and the channel it's clustered under are only weakly correlated. This weak correlation between individual video views and channels can be explained by examining the channels with the highest market shares over time. In Figure 4A, we show the top six channels ranked by peak monthly market share. In Figure 4B, we show the top six channels ranked by median

Group	Variable	Coef.	Std.Err.	[0.025	0.975]
	Intercept	255.492***	15.431	225.248	285.736
Channel	seniority	3.597***	0.441	2.733	4.460
	priorMS	267372.154***	850.075	265706.038	269038.270
Offensive	offensive	44.226**	12.917	18.910	69.542
	prosocial	-263.171***	56.639	-374.182	-152.159
Novelty	priorMS_topic	1545.903***	72.096	1404.597	1687.208
	priorSS_topic	-1144.402***	81.271	-1303.689	-985.114
Conspiracy	certitude	177.286**	60.896	57.933	296.639
	allnone	195.972***	39.871	117.826	274.118
	they	386.046***	62.728	263.101	508.991
	death	239.752***	44.091	153.335	326.169
	power	-94.266***	19.515	-132.514	-56.018
	relig	-112.055***	29.579	-170.029	-54.082
	Channel Var	1678758.217	4.596	-	-

Table 1: (Part 1) Summary from multilevel, random intercept linear model, where video views are the dependent variable and those videos are clustered under channels. The "Channel Var" is the random effects of the cluster variable. Views could vary across channels by $\pm 1,295.67$ and the Interclass Correlation Coefficient (ICC) (video views within a channel) was 0.214. The mean group size (number of videos per channel) was 66.9. In this model, we only use the 90 day prior market share and supply share windows due to how correlated the features are. However, the same pattern holds at each window size. Emotion features used in the model can be found in Table 2. Significance codes are: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

monthly market share. There are broadly two types of dynamics across the channels. First, there are channels that produce *one-hit wonders*, where their peak market share is entirely driven by a single video (see Figure 4c, 4d, 4e, and 4f). We discuss these one-hit wonders in more detail later in the paper. Second, there are channels that more consistently garner views. These consistent channels include many repeat offenders in the conspiracy theory space, such as *banned-dot-video* and *rongibson* (both Alex Jones/Infowars shows) and *x22report*, a well-known conspiracy theory website.

The number of channels with one-hit wonders changed over time, as reflected by Figure 3e, which shows the Peak-Median Views Ratio (PMVR) over time. This metric captures the inequality in video views within channels. More precisely, as defined by (Guinaudeau et al., 2022), the PMVR is the

Group	Variable	Coef.	Std.Err.	[0.025	0.975]
	vader	6.666	4.344	-1.848	15.180
	anger	3.736	41.384	-77.376	84.848
	hate	-66.539	38.751	-142.489	9.410
	contempt	-104.819*	47.152	-197.236	-12.402
	disgust	102.709	71.219	-36.878	242.296
	embarrassment	412.556*	198.220	24.051	801.060
	love	77.779*	35.841	7.533	148.025
	admiration	164.777**	52.416	62.043	267.511
	sexual attraction	-91.110	101.706	-290.449	108.229
	cuteness	165.047	107.150	-44.964	375.058
	wonder	2283.110**	781.965	750.486	3815.734
	pride	-33.835	48.166	-128.240	60.569
	sadness	42.607*	21.396	0.671	84.542
Emotion	nostalgia	-498.760	427.844	-1337.320	339.800
	empathic pain	314.642*	125.016	69.615	559.668
	gratitude	-231.258*	104.184	-435.455	-27.061
	envy	-16054.528***	1905.075	-19788.405	-12320.650
	fear	66.033**	23.341	20.286	111.780
	confusion	-1205.538***	79.303	-1360.970	-1050.107
	surprise	1686.783***	92.222	1506.030	1867.535
	happiness	-414.275***	41.431	-495.479	-333.071
	excitement	532.741***	78.983	377.937	687.546
	amusement	-191.298**	65.774	-320.213	-62.382
	hope	44.632	56.139	-65.400	154.663
	boredom	-625.473**	189.313	-996.520	-254.426
	fatigue	10071.664***	521.475	9049.592	11093.737
	positive	542.092***	100.515	345.086	739.098
	negative	-1186.950***	203.271	-1585.353	-788.546
	Channel Var	1678758.217	4.596	-	-

Table 2: (Part 2) Summary from multilevel, random intercept linear model, where video views are the dependent variable and those videos are clustered under channels. The “Channel Var” is the random effects of the cluster variable. In this table we show the emotion features from the model. Other features used in the model can be found in Table 1 above. Significance codes are: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

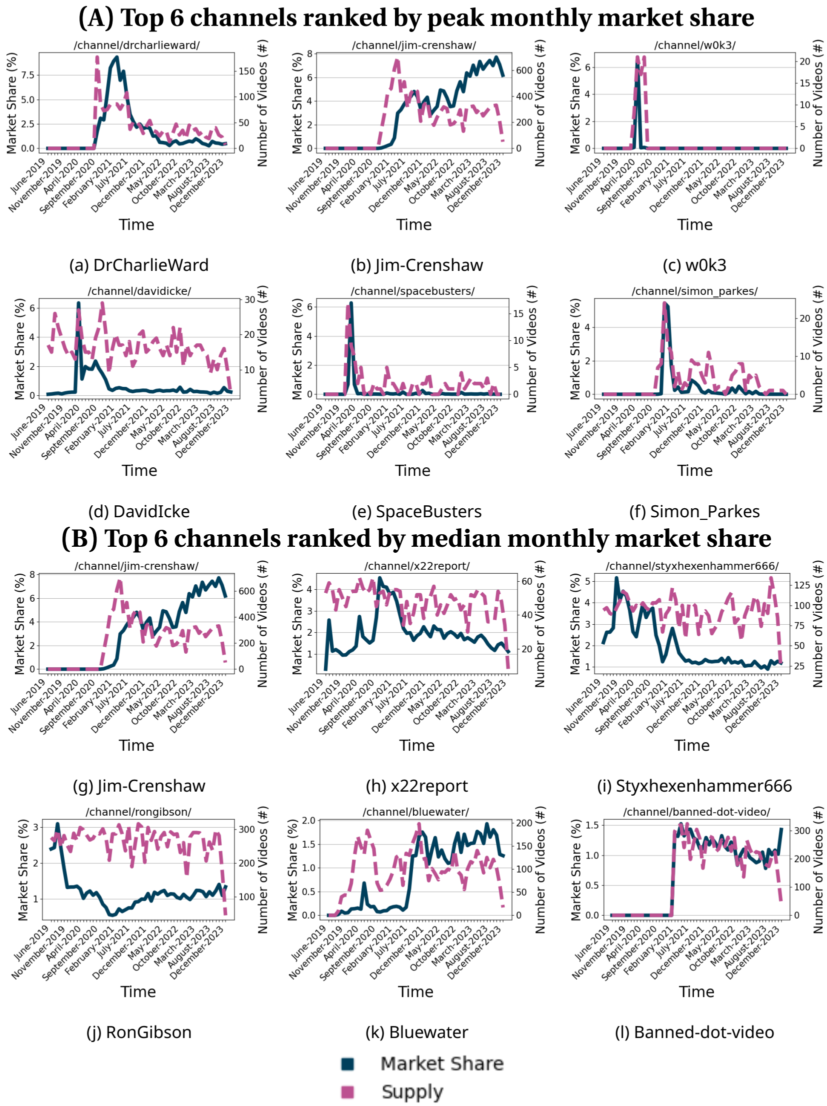

Figure 4: (A) The monthly market share and supply share of the top six channels ranked by peak market share. The left-hand Y-axis is market share, while the right-hand Y-axis is the number of videos posted by the channel (supply). Note the different axis ranges across the plots. (B) The monthly market share and supply of the top six channels ranked by median market share.

median ratio of the number of views on channels' most popular video to the median number of views their videos get. Hence, a higher PMVR indicates a higher discrepancy between the views on a channel's most popular video and its median views. In Figure 3e, we see that early in the platform's life, channels received a more uniform number of views across their videos (i.e., between June 2019 and February 2021 the PMVR ranged between 15 and 25). Just after February 2021, the PMVR rises, peaking at 41 towards the end of 2021. The peak PMVR roughly coincides with the peak in median views per video across the platform (Figure 3c) and the peak number of views on the platform as a whole (Figure 3d).

These results - and the results in the section above - suggest that while typical rich-get-richer effects play some role on the platform, popular individual videos are perhaps more important in attracting audiences. We discuss this implication in more detail later in the paper, in particular how these individual videos relate to external events.

Videos with titles that were more offensive, reactive, and novel were more popular

So far, our results suggest that channels have some influence on video popularity, but this is clearly not the only determining factor. As shown in Tables 1 and 2, features of the content play a role in video popularity. First, videos with a higher probability of being offensive ($p = 0.001$) and a lower number of prosocial words ($p < 0.001$) were viewed more, supporting **H1**.

Second, titles that expressed specific emotions were more popular. Most notably, video titles that expressed surprise ($p < 0.001$), excitement ($p < 0.001$), fatigue ($p < 0.001$), fear ($p = 0.005$), wonder ($p = 0.004$), and admiration ($p = 0.002$) were viewed significantly more. To a lesser extent, titles that expressed embarrassment ($p = 0.037$), love ($p = 0.030$), sadness ($p = 0.046$), or empathic pain ($p = 0.012$) were also more popular. On the other hand, titles that expressed happiness ($p < 0.001$), confusion ($p < 0.001$), envy ($p < 0.001$), amusement ($p = 0.004$), or boredom ($p = 0.001$) were significantly less popular. Other emotions associated with less views, but to a lesser extent, were contempt ($p = 0.026$) and gratitude ($p = 0.026$).

Aligning with this finding, videos that used more certitude, outgroup identification, and over-generalizations were more popular (all $p < 0.005$). This combination of features broadly suggests videos with more conspiratorial language were more popular (Fong et al., 2021), partially supporting **H2**. However, the common conspiratorial "themes" of power and religion, as put by (Fong et al., 2021), were not associated with more views. Specifically, although videos with more words related to death were more popular ($p < 0.001$), videos with language related to power and religion were associated with significantly less views (both $p < 0.001$). This differing engagement with conspiratorial themes is likely due to a combination of the in-demand topics on the platform during the time-frame of our data and the limitations of LIWC.

Lastly, videos that covered more novel topics were also more popular. As shown in Table 1, if a video topic had a higher supply share in the past, it was viewed significantly less (variable *priorSS_topic*, $p < 0.001$). In other words, videos that covered topics which were already heavily covered in the past 90 days received less views. Again, we only include prior supply share for the past 90 days in the model, but the results hold across both the 30 day window and the 7 day window.

Together, we can interpret a broad trend of popularity on BitChute. That is, videos that were *reactive* to new events were more popular. For example, surprise is defined as a "reaction to sudden new or unexpected thing or event", wonder is "amazement but not surprise", and excitement is a "high intensity response to novelty and challenge" (Paletz et al., 2024). Those events are reacted to with fear and fatigue, but not with confusion, envy, or happiness. This trend also suggests that - perhaps unsurprisingly based on prior literature (Dailey & Starbird, 2015; Prochaska et al., 2024; K. Zhou et al., 2023) - *sensemaking* during uncertain events played a role in content demand.

Popular content exhibited a high rumor threat level and covered commonly deplatformed topics

The video-level results discussed above begin to suggest that videos with rumor traits were more popular (i.e., emotional valence and novelty). We can explore this notion further by examining the supply and demand of topics over time. Utilizing the STM that was built to estimate content novelty (see Methods Section), we show the monthly market share and supply share per topic in Figure 5. There were a total of 21 topics in the model, as identified by a parameter grid-search. We show the top six topics ranked by peak monthly market share. For these top six topics, the authors, who are experts in conspiracy theories, hand-labeled them using a combination of the keywords provided by the topic model and sampled videos from each topic.

From this figure, there is one immediately clear takeaway: COVID-19 and 2020 U.S. Election fraud conspiracy theories dominated the platform's supply and demand. In the middle of 2021, COVID-19 videos received over 25% of the views and accounted for over 17% of the videos on the platform. In late 2020, election fraud videos received over 20% of the views and accounted for over

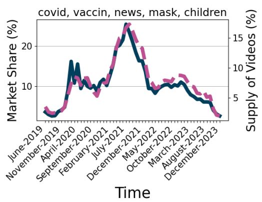

(a) **COVID-19 Conspiracy Theories** (peak 25.41% of views, 17.27% of supply)

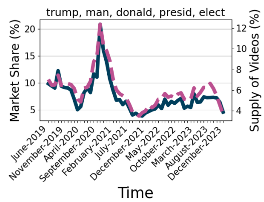

(b) **Trump/Election Fraud** (peak 20.87% of views, 12.36% of supply)

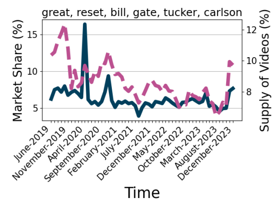

(c) **Great Reset Conspiracy Theories** (peak 16.37% of views, 12.34% of supply)

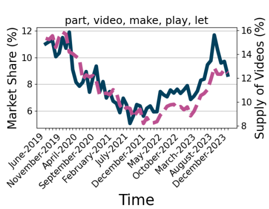

(d) **Video Games** (peak 11.91% of views, 15.84% of supply)

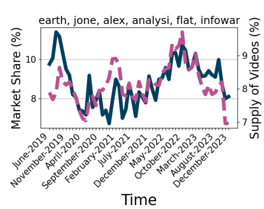

(e) **Talk Shows & Interviews** (peak 11.42% of views, 9.71% of supply)

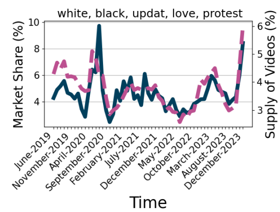

(f) **Race/Anti-Black** (peak 9.71% of views, 6.01% of supply)

Figure 5: The monthly market share and supply share of the top 6 topics ranked by peak market share. The left-hand Y-axis is market share, while the right-hand Y-axis is supply share. Note the different axis ranges across the plots and that, unlike Figure 4, the right-hand Y-axis of these plots show supply share not raw supply.

12% of the videos on the platform. The next highest monthly market share for a topic was 16% for videos about Great Reset conspiracy theories[3], which, upon inspection, also contained videos about U.S. elections and COVID-19. Further, the most viewed video for all channels shown in Figure 4, except for one, were either about about COVID-19 or the 2020 U.S. Presidential Election.

As a robustness check, we also measured the monthly market share of these two topics using two keyword lists. Given that noisy-data in the topic model may over-estimate the number of videos within a topic, this simple, limited keyword filter can provide assurance of the results above. For COVID-19, videos that contained at least one of 10 keywords ('covid', 'coronavirus', 'vaccine', 'vax', 'vaccinated', 'pfizer', 'fauci', 'pandemic', 'plandemic', or 'lock-down') received 24% of views in September 2021. Across all 54 months of

[3]The general premise of Great Reset conspiracy theories is that global elites use crisis events to advance their interests and push a globalist plot to destroy the United States. See https://www.adl.org/resources/blog/great-reset-conspiracy-flourishes-amid-continued-pandemic.

data, COVID-19 videos garnered 10.6% of the total views on the platform (over 305M views). For election fraud, videos that contained at least one of 6 keywords ('election', 'trump', 'biden', 'jan 6', 'vote', and 'stolen') received 25% of views in October 2020. Across all 54 months of data, election-related videos received 6.5% of the total views on the platform (over 187M views). Remember, given that views were only captured a week after each video is uploaded, these numbers are likely underestimates.

To situate these results in the larger media ecosystem, we use the above keyword lists to compare monthly supply share on BitChute to the monthly supply share on Twitter for both topics. To do this, we use an archive of Twitter's public sample stream (approximately 1% of tweets posted), where this content was saved as it is created in real time. For each of the COVID-19 and election-fraud keyword lists, we find all tweets from this sample that include at least one of the relevant substrings in the text, doing a case-insensitive match. This process matches all tweets, retweets, and quoted tweets in the sample stream archive as well as tweets from accounts that mention one of these keywords in their user profiles or screen names. This comparison is done for the months between July 2019 and November 2022. We show the results of this comparison in Figure 6.

Using this keyword list comparison, it is evident that COVID-19 dominated the content production on BitChute compared to Twitter, peaking at a monthly supply share of 11.42% and 9.68%, respectively. While both BitChute and Twitter spiked in COVID-19 content production in March 2020 (8.45% on BitChute, 9.68% on Twitter), BitChute quickly diverges from Twitter in the months to follow. In September 2021, 11.42% of the content on BitChute was about COVID-19 while only 1.92% of the content on Twitter was about COVID-19 (a difference of 9.5%, see Figure 6a). On the other hand, election-related content followed the supply share trend of Twitter closely (see Figure 6b). On BitChute, monthly supply share for election-related content peaked at 9.68% in October 2020 while monthly supply share peaked at 7.57% in November 2020 on Twitter. Throughout the timeline, the maximum difference in proportions between the platforms was 3.3% in October 2020. Hence, although election-related content did have a sizable impact of BitChute, it proportionally was not that different from Twitter.

While it is difficult to demonstrate the full impact of banned content on BitChute's content supply and demand with this data alone, we do see that both highly produced content and popular content on the platform discussed topics and used language that was frequently banned from mainstream social media platforms, suggesting that, perhaps unsurprisingly, the

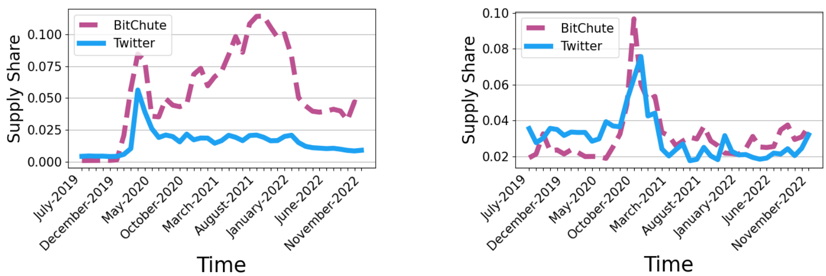

(a) COVID-19 Supply Share

(b) Election Supply Share

Figure 6: Comparisons of monthly supply share of (a) COVID-19 content and (b) 2020 U.S. election content between BitChute and Twitter from July 2019 to November 2022. Content was filtered using the simple keyword lists described in Section . The peak monthly supply share of COVID-19 content was 11.42% on BitChute and 5.63% on Twitter. The peak monthly supply share of election content was 9.68% on BitChute and 7.57% on Twitter.

minimal content moderation on BitChute impacted what types of content were supplied and consumed. Aligning with the commonly deplatformed topics by YouTube and Twitter, BitChute filled a demand gap for COVID-19 and - to a lesser extent - 2020 U.S. elections fraud conspiracy theories. These results support **H2**.

Likewise, it is difficult to characterize external events and their impact on the platform without granular, qualitative study. Yet, it is clear from the analysis above that COVID-19 dominated the supply and demand of content on BitChute. In particular, the supply share of COVID-19 videos greatly differed from that of the Big-Tech platform Twitter, and approximately 25% of the views within one month on BitChute were on COVID-19 videos. Given the high rumor threat level of this long-lasting event, and COVID-19 content's wide spread across many other social media platforms (Cinelli et al., 2020; Thiele, 2024), this outcome should be to some extent expected. As put by (Spiro & Starbird, 2023): "A novel virus with uncertain causes, serious consequences, unknown spreading mechanisms, and few remedies scores high in every category [of the rumor threat framework]." Between the topical coverage and content features discussed above, we argue that these results, at least in part, support **H5**.

The most viewed COVID-19 and U.S. Election videos

To better contextualize these results, we examined the most viewed videos within each topic. In particular, we focused on two of the top six topics:

COVID-19 and election fraud. We choose to only focus on these two topics for two reasons. First, as described above, these two topics dominated the demand on the platform. Second, as described earlier in the paper, they are theoretically interesting topics given the policies enacted by YouTube at that time.

#	Title	Published	Views	Copies
1	Plandemic Documentary: The Hidden Agenda Behind Covid-19	May 6, 2020	1,616,409	20
2	MUST WATCH!!! Funeral Directer John O'Looney Blows the Whistle on Covid	Sep 16, 2021	1,299,935	4
3	Horrific Findings In The Blood Of The Vaccinated	Jul 18, 2021	1,072,612	18
4	Important Information on Coronavirus 5G Kung Flu	Mar 4, 2020	865,882	44
5	Biden's White House Press Corps Admit Covid Was Fake On A Hot Mic	Jan 27, 2021	799,661	0
6	A Final Warning to Humanity from Former Pfizer Chief Scientist Michael Yeadon	May 16, 2021	549,635	36
7	EXPLAINS HOW THE DEPOPULATION mRNA VACCINES WILL START WORKING IN 3-6 MONTHS - DR. SHER	Feb 9, 2021	534,433	10
8	Dr. Bryan Ardis - Hospital Protocol is what is Murdering Covid / Flu Patients	Aug 27, 2021	454,561	13
9	A FINAL WARNING TO HUMANITY FROM FORMER PFIZER CHIEF SCIENTIST MICHAEL YEADON	Jul 22, 2021	433,383	36
10	Dr Bryan Ardis with the most stunning TRUTH about Covid Fauci & Remsesivir you could never imagine	Aug 27, 2021	413,610	3

Table 3: Most viewed COVID-19 videos on the platform. We show the number of times each video was viewed in its first week and the number of times the video was copied with the exact same title.

COVID-19 Content In Table 3, we show the top ten most viewed videos within the COVID-19 topic, along with the dates they were published and the number of verbatim copies of the video that were published on the platform. Interestingly, many of the most viewed videos on the platform, particularly

COVID-19 videos, are not original content but instead were deplatformed content re-published by others to BitChute. For example, unsurprisingly, the most viewed COVID-19 video (and the most viewed video on the platform as a whole) is the infamous *Plandemic* documentary (Buntain et al., 2021; Lytvynenko, 2020; Yadav, 2021). This 26-minute documentary-style video was originally created by filmmaker Mikki Willis and upon being banned from Big-Tech platforms on May 5th was copied to BitChute by many content producers. The most viewed copy is simply the first copy uploaded to the platform.

Similarly, other documentary-style and interview-based COVID-19 videos were highly engaged with on BitChute. Many of these videos are made to look like cable news interviews or professional produced documentaries, where people add their own evidence to and interpretations about the COVID-19 lockdowns and vaccines. Of those being interviewed, they often were with deplatformed media personalities or people who were framed as experts. Within the top ten COVID-19 videos, only one video appeared to be an original creation of the channel. Rather, they were re-uploaded videos from individual conspiracy theorist websites or videos that were deplatformed from Big-Tech platforms. All but one of the "one-hit wonders" in 4A were re-uploaded COVID-19 videos.

Elections Content In Table 4, we show the top ten most view videos within the election fraud topic. While some of the videos within this topic were copies from other content creators, more of them appeared to be original creations of the channels. For instance, five of the top ten videos appear to be original creations. Furthermore, the most popular videos in this topic were often produced by channels that had relatively consistent audiences. For example, three of the top ten election videos were produced by the channel *drcharlieward* (who ironically gained their highest monthly market share by copying the #6 COVID-19 video).

These highly engaged with election videos tend to follow a different format than the top COVID-19 videos, often designed as "evidence collages (Krafft & Donovan, 2020)" of other video clips. However, there is some overlap in style between the COVID-19 videos and election videos, particularly when examining videos outside of the top ten. Both topics have videos interviewing people framed as experts, cable news style talk shows, and evidence colleges.

There are other highly viewed videos discussing election fraud conspiracies that are not captured by the topic model or keyword filter, which have

#	Title	Published	Views	Copies
1	SMOKING GUN: ELECTRONIC VOTE FRAUD CAUGHT LIVE ON CNN!	Nov 7, 2020	672,590	9
2	Michigan Vote Fraud WITNESS assassination attempt	Dec 9, 2020	403,795	0
3	We Are Trump 777!	Jul 2nd, 2021	375,337	6
4	Lieutenant General Thomas Mcinerney blowing the whistle CIA software Hammer score card rig election	Nov 6th, 2020	239,451	0
5	What you didn't know about our President Donald John Trump. Because they hid it from you!	Jan 24th, 2021	229,576	7
6	This Video Proves that Donald Trump Won the 2020 Presidential Election	Nov 12th, 2020	213,073	6
7	BACK TO THE FUTURE WITH DONALD TRUMP	Jan 31st, 2021	207,331	0
8	INCREDIBLE GENIUS PRESIDENT DONALD TRUMP	Apr 3rd, 2021	203,735	0
9	Steve Pieczenik: Triples Down On What Is Coming, 'Mass Arrests, Biden Will NOT Be President' & More	Jan 14th, 2021	201,238	2
10	Gag Order Failed Forensic Results of Dominion Machines Stealing Votes LEAKED!	Dec 12, 2020	185348	0

Table 4: Most viewed U.S. election videos on the platform. We show the number of times each video was viewed in its first week and the number of times the video was copied with the exact same title.

titles formatted as "*Date Update Current News*" (e.g. a video titled "16th January Second Update Current News" had 1,012,761 views, another titled "13th January Update Current News" had 738,911 views). These videos all appear to be a single speaker describing evidence of an election conspiracy theory that they heard from a highly ranking government official. Along with describing evidence, the speaker confidently tells viewers what is going to happen in the future (i.e., Trump will take over the government on day X, etc.), aligning with well-known QAnon conspiracies and our content popularity results above (i.e. high *anticipation* and *certitude*). As best described by (Marwick & Partin, 2022), QAnon conspiracies "demonstrate populist expertise, the rejection of legacy media accounts of current events in favor of the alternative facts constructed through their systematic research programs." This description aligns well with the content in these videos.

Discussion

In this paper, we examined the supply and demand of content on the Alt-Tech video-hosting platform BitChute both theoretically and empirically. Using a near-complete, longitudinal data set of 6M videos from 82K channels, we provide evidence that two key affordances of the platform, at least in part, drove the supply and demand of content: the platform's minimal content moderation and the lack of personalized recommendation algorithms. In addition, characteristics about the content itself and the external events covered by that content played a role.

Our results suggest that there are two groups of content creators on BitChute. First, there are a small number of popular channels that receive *relatively* consistent engagement over time. Second, there are many "one-hit wonder" channels who predominantly mirror videos deplatformed by Big-Tech platforms and receive significant engagement on a few of those mirrored videos. Often these one-hit wonder channels post very little content after gaining popularity. For example, the most viewed copy of the 'Plandemic' documentary was one of the only 20 videos posted by the channel *w0k3*. Notably, even the channels that received relatively consistent engagement over time put concerted effort into mirroring content. This mirroring practice is often explicitly stated in the channel descriptions. For instance, one of the oldest and most established channels stated "Some videos are mine but most are not" and another said their videos were "Intel from various sources". Further, other established channels who produce original content received the most views on mirrored videos rather than their original content.

The sociopolitical events of 2021, along with the content moderation policies of Big-Tech platforms during 2021, appeared to be critical in driving audiences to BitChute. Popular content was not only more offensive than unpopular content, but it was likely to be from a subset of conspiracy theories about COVID-19 or U.S. elections fraud. The overall engagement with the platform peaked in 2021, corresponding with the January 6 United States Capitol attack and the wide deployment of COVID-19 vaccines. Both topics had "high rumor threats (Spiro & Starbird, 2023)" and were heavily moderated by YouTube, Twitter, and Facebook. In fact, the aftermath of the January 6 United States Capitol attack has been referred to as the 'Great Deplatforming' due to the large number of accounts banned by Twitter to limit conspiracy theories and misinformation about the 2020 election (Buntain & Snegovaya, 2024).

BitChute's content popularity is shaped by both its internal affordances, and its dependence on far-right political events: 2021 saw both the platforms highest view inequality across channels, immediately following January 6 United States Capitol attack, and its lowest view inequality across videos, a few months after this peaked. Specifically, between December 2020 and February 2021 only 2% of channels received 85% of the views, and the median views per video was 75 views. In the summer of 2021, just under 4% of channels received 85% of the views, and the median views per video peaked at 150 views. This result highlights the complex relationship between the internal affordances of the platform and the drivers outside of the platform. While audiences on BitChute were viewing a variety of videos, particularly within the topics of COVID-19 and U.S. elections fraud, those videos were increasingly being posted by only a few channels. Without an internal recommendation system, audiences on BitChute may land on the platform due to external drivers but stay within specific channels rather than exploring other content producers on the platform.

Implications

The content supplied on BitChute filled a demand gap created by Big-Tech content moderation policies. This fact supports prior concerns that removing content from major media platforms motivates both content producers and consumers to move to alternative, less moderated, more radical spaces (Ali et al., 2021; Buntain et al., 2023; Monti et al., 2023). However, our results also suggest that the external context and conditions moderate this movement. That is, while banned content surrounding the uncertainty and significance of events like COVID-19 clearly grew engagement with BitChute, that engagement eventually faded away. Without other "high rumor threat" events to replace these past events, the platform as a whole lost engagement (as indicated in Figure 3d). In part, this result should be expected, as rumors lose interest over time (Kapferer, 2013; Spiro & Starbird, 2023).

This result also reveals the current role of BitChute in the larger media ecosystem: *BitChute acts as mirroring system for deplatformed videos*, much more than it facilitates community building or original content creation. This is not to say that there are no communities or original content creators on BitChute. Rather, it is to say that the primary role of the platform appears to be harboring deplatformed videos. While there are channels that consistently post original content on BitChute, even those channels frequently use systems external to BitChute to grow and engage with their communities. For example, the popular conspiracy channel *x22report* maintains its own website and private server to host videos and link to fringe news websites (like The Gateway Pundit). All of the channels under the Infowars profile (e.g., *RobGibson*, *banned-dot-video*, etc.) host live content with audience discussions on the Infowars website. The popular channel *styxhexenhammer666* maintains an active YouTube channel with fruitful audience discussions in the comment section of those videos.

BitChute plays a subtly different role than other Alt-Tech platforms due to its core medium and affordances, providing a content-hosting backbone rather than a community-growing town square. Gab, an alternative to Twitter, facilitates an “ideologically eclectic far-right community” and “networked outrage” due to its “culture of anonymity” and “microblogging architecture” (Jasser et al., 2023; Weigel & Gitomer, 2024). Voat, an alternative to Reddit that is now defunct, facilitated community building through its subreddit-like discussion board structure (Papasavva, Blackburn, et al., 2020). The fringe platform 4chan has a long history of facilitating highly offensive, sometimes illegal, content creation due to its anonymity, ephemeralness, and imageboard structure (Bernstein et al., 2011; Papasavva, Zannettou, et al., 2020; Tuters & Hagen, 2020). By contrast, BitChute serves as a backup for fringe content when that content is deplatformed. This aligns with prior work, which showed that BitChute videos promoting election fraud claims were commonly linked to on Twitter; yet, YouTube videos promoting election fraud claims were linked to even more so (M. Childs et al., 2022).

Conclusion

In this paper, we demonstrate the complex relationship between internal affordances, content features, and external forces in driving supply and demand on BitChute. We show that videos produced by established channels with titles that were more offensive and reactive, particularly titles that expressed more surprise, excitement, wonder, fear, and fatigue were more popular. Further, videos that expressed *certitude* were viewed more and videos that expressed *confusion* were viewed less, suggesting that *sense-making* during uncertain events played a role in content demand. We also show that popular content covered commonly deplatformed topics from high rumor threat events. While views on BitChute grew between 2019 and 2022, the platform has seen a steady decline in views since 2022. During the platform’s peak engagement, the most popular content were mirrored videos banned by major media platforms. Together, our results suggest that BitChute's current role in the media ecosystem is less as a town square and more as a redundant backup for deplatformed video content.

References

Ali, S., Saeed, M. H., Aldreabi, E., Blackburn, J., De Cristofaro, E., Zannettou, S., & Stringhini, G. (2021). Understanding the effect of deplatforming on social networks. *Proceedings of the 13th ACM Web Science Conference 2021*, 187–195.

Aliapoulios, M., Bevensee, E., Blackburn, J., Bradlyn, B., De Cristofaro, E., Stringhini, G., & Zannettou, S. (2021). An early look at the parler online social network. *arXiv preprint arXiv:2101.03820*.

Allport, G. W., & Postman, L. (1946). An analysis of rumor. *Public opinion quarterly*, *10*(4), 501–517.

Barbieri, F., Camacho-Collados, J., Neves, L., & Espinosa-Anke, L. (2020). Tweeteval: Unified benchmark and comparative evaluation for tweet classification. *arXiv preprint arXiv:2010.12421*.

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

Bellemare, A., Nicholson, K., & Ho, J. (2020). How a debunked covid-19 video kept spreading after facebook and youtube took it down. *CBC News*.

Bernstein, M., Monroy-Hernández, A., Harry, D., André, P., Panovich, K., & Vargas, G. (2011). 4chan and/b: An analysis of anonymity and ephemerality in a large online community. *Proceedings of the international AAAI conference on web and social media*, *5*(1), 50–57.

Birim, Ş. Ö., Kazancoglu, I., Mangla, S. K., Kahraman, A., Kumar, S., & Kazancoglu, Y. (2022). Detecting fake reviews through topic modelling. *Journal of Business Research*, *149*, 884–900.

Bond, S. (2023). Youtube will no longer take down false claims about u.s. elections [Accessed on February 21, 2024].

Borg, A., & Boldt, M. (2020). Using vader sentiment and svm for predicting customer response sentiment. *Expert Systems with Applications*, *162*, 113746.

Boyd, R. L., Ashokkumar, A., Seraj, S., & Pennebaker, J. W. (2022). The development and psychometric properties of liwc-22. *Austin, TX: University of Texas at Austin*, 1–47.

Bucher, T., Helmond, A., et al. (2018). The affordances of social media platforms. *The SAGE handbook of social media*, *1*, 233–253.

Buckley, N., & Schafer, J. S. (2022). 'censorship-free'platforms: Evaluating content moderation policies and practices of alternative social media.

Buntain, C., Bonneau, R., Nagler, J., & Tucker, J. A. (2021). Youtube recommendations and effects on sharing across online social platforms. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW1), 1–26.

Buntain, C., Innes, M., Mitts, T., & Shapiro, J. (2023). Cross-platform reactions to the post-january 6 deplatforming. *Journal of Quantitative Description: Digital Media, 3*.

Buntain, C., & Snegovaya, M. (2024). Individual-level moderating and differential effects of deplatforming. *Available at SSRN 4739537*.

Camacho-Collados, J., Rezaee, K., Riahi, T., Ushio, A., Loureiro, D., Antypas, D., Boisson, J., Espinosa-Anke, L., Liu, F., Martínez-Cámara, E., et al. (2022). Tweetnlp: Cutting-edge natural language processing for social media. *arXiv preprint arXiv:2206.14774*.

Chan, J., Lambert, C., Choi, F., Chancellor, S., & Chandrasekharan, E. (2024). Understanding community resilience: Quantifying the effects of sudden popularity via algorithmic curation. *Proceedings of the International AAAI Conference on Web and Social Media, 18*.

Cheng, J., Bernstein, M., Danescu-Niculescu-Mizil, C., & Leskovec, J. (2017). Anyone can become a troll: Causes of trolling behavior in online discussions. *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, 1217–1230.

Childs, M., Buntain, C., Z. Trujillo, M., & D. Horne, B. (2022). Characterizing youtube and bitchute content and mobilizers during us election fraud discussions on twitter. *14th ACM Web Science Conference 2022*, 250–259.

Childs, M. C., & Horne, B. D. (2023). Examining the production of co-active channels on youtube and bitchute. *arXiv preprint arXiv:2303.07861*.

Chochlakis, G., Mahajan, G., Baruah, S., Burghardt, K., Lerman, K., & Narayanan, S. (2023a). Leveraging label correlations in a multi-label setting: A case study in emotion. *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.

Chochlakis, G., Mahajan, G., Baruah, S., Burghardt, K., Lerman, K., & Narayanan, S. (2023b). Using emotion embeddings to transfer knowledge between emotions, languages, and annotation formats. *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.

Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C. M., Brugnoli, E., Schmidt, A. L., Zola, P., Zollo, F., & Scala, A. (2020). The covid-19 social media infodemic. *Scientific reports, 10*(1), 1–10.

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., & Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

Coppersmith, G., Harman, C., & Dredze, M. (2014). Measuring post traumatic stress disorder in twitter. *Proceedings of the international AAAI conference on web and social media, 8*(1), 579–582.

Dailey, D., & Starbird, K. (2015). " it's raining dispersants" collective sensemaking of complex information in crisis contexts. *Proceedings of the 18th ACM conference companion on computer supported cooperative work & social computing*, 155–158.

Davidson, J., Liebald, B., Liu, J., Nandy, P., Van Vleet, T., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B., et al. (2010). The youtube video recommendation system. *Proceedings of the fourth ACM conference on Recommender systems*, 293–296.

De Bruyne, L., van der Meer, T. G., De Clercq, O., & Hoste, V. (2024). Using state-of-the-art emotion detection models in a crisis communication context. *Computational Communication Research*, *6*(1), 1.

Dehghan, E., & Nagappa, A. (2022). Politicization and radicalization of discourses in the alt-tech ecosystem: A case study on gab social. *Social Media+ Society*, *8*(3), 20563051221113075.

Donovan, J., Lewis, B., & Friedberg, B. (2019). Parallel ports: Sociotechnical change from the alt-right to alt-tech.

Douglas, K. M., Sutton, R. M., & Cichocka, A. (2017). The psychology of conspiracy theories. *Current directions in psychological science*, *26*(6), 538–542.

Eichstaedt, J. C., Smith, R. J., Merchant, R. M., Ungar, L. H., Crutchley, P., Preoţiuc-Pietro, D., Asch, D. A., & Schwartz, H. A. (2018). Facebook language predicts depression in medical records. *Proceedings of the National Academy of Sciences*, *115*(44), 11203–11208.

Ferrara, E., Chang, H., Chen, E., Muric, G., & Patel, J. (2020). Characterizing social media manipulation in the 2020 us presidential election. *First Monday*.

Fong, A., Roozenbeek, J., Goldwert, D., Rathje, S., & van der Linden, S. (2021). The language of conspiracy: A psychological analysis of speech used by conspiracy theorists and their followers on twitter. *Group Processes & Intergroup Relations*, *24*(4), 606–623.

Friggeri, A., Adamic, L., Eckles, D., & Cheng, J. (2014). Rumor cascades. *proceedings of the international AAAI conference on web and social media*, *8*(1), 101–110.

Gibson, J. J. (1977). The theory of affordances. *Hilldale, USA*, *1*(2), 67–82.

Guinaudeau, B., Munger, K., & Votta, F. (2022). Fifteen seconds of fame: Tiktok and the supply side of social video. *Computational Communication Research*, *4*(2), 463–485.

Harrison, X. A., Donaldson, L., Correa-Cano, M. E., Evans, J., Fisher, D. N., Goodwin, C. E., Robinson, B. S., Hodgson, D. J., & Inger, R. (2018). A brief introduction to mixed effects modelling and multi-model inference in ecology. *PeerJ*, *6*, e4794.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of verbal learning and verbal behavior*, *16*(1), 107–112.

Hebbelstrup Rye Rasmussen, S., & Petersen, M. B. (2023). The event-driven nature of online political hostility: How offline political events make online interactions more hostile. *PNAS nexus*, *2*(11), pgad382.

Hershey, S., Chaudhuri, S., Ellis, D. P., Gemmeke, J. F., Jansen, A., Moore, R. C., Plakal, M., Platt, D., Saurous, R. A., Seybold, B., et al. (2017). Cnn architectures for large-scale audio classification. *2017 ieee international conference on acoustics, speech and signal processing (icassp)*, 131–135.

Horne, B. D. (2022). A psycho-linguistic analysis of bitchute. *arXiv preprint arXiv:2204.08078*.

Horne, B. D., Nørregaard, J., & Adalı, S. (2019). Robust fake news detection over time and attack. *ACM Transactions of Intelligent Systems Technology*.

Hutto, C., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the international AAAI conference on web and social media*, *8*(1), 216–225.

Innes, H., & Innes, M. (2023). De-platforming disinformation: Conspiracy theories and their control. *Information, Communication & Society*, *26*(6), 1262–1280.

Israeli, A., & Tsur, O. (2022). Going extreme: Comparative analysis of hate speech in parler and gab. *arXiv preprint arXiv:2201.11770*.

Jakubik, J., Vössing, M., Pröllochs, N., Bär, D., & Feuerriegel, S. (2023). Online emotions during the storming of the us capitol: Evidence from the social media network parler. *Proceedings of the International AAAI Conference on Web and Social Media*, *17*, 423–434.

Jasser, G., McSwiney, J., Pertwee, E., & Zannettou, S. (2023). 'welcome to# gabfam': Far-right virtual community on gab. *New Media & Society*, *25*(7), 1728–1745.

Jhaver, S., Boylston, C., Yang, D., & Bruckman, A. (2021). Evaluating the effectiveness of deplatforming as a moderation strategy on twitter. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW2), 1–30.

Joseph, K., Horne, B. D., Green, J., & Wihbey, J. P. (2021). Local news online and covid in the us: Relationships among coverage, cases, deaths, and audience. *arXiv preprint arXiv:2111.08515*.

Kapferer, J.-N. (2013). *Rumors: Uses, interpretations, and images*. Transaction Publishers.

Kor-Sins, R. (2023). The alt-right digital migration: A heterogeneous engineering approach to social media platform branding. *New Media & Society*, *25*(9), 2321–2338.

Krafft, P., & Donovan, J. (2020). Disinformation by design: The use of evidence collages and platform filtering in a media manipulation campaign. *Political Communication*, *37*(2), 194–214.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Lovejoy, T., Davies, & Prosser. (2021). The truth is in the title? video title generation as a novel training objective for video summarisation.

Lytvynenko, J. (2020). After the "plandemic" video went viral in the us, it was exported to the rest of the world (B. News, Ed.) [Available at"https://www.buzzfeednews.com/article/janelytvynenko/coronavirus-plandemic-translation"].

M. Otala, J., Kurtic, G., Grasso, I., Liu, Y., Matthews, J., & Madraki, G. (2021). Political polarization and platform migration: A study of parler and twitter usage by united states of america congress members. *Companion Proceedings of the Web Conference 2021*, 224–231.

Mahl, D., Zeng, J., & Schäfer, M. S. (2023). Conceptualizing platformed conspiracism: Analytical framework and empirical case study of bitchute and gab. *new media & society*, 14614448231160457.

Marshall, A., & Tanfani, J. (2022). New breed of video sites thrives on misinformation and hate.

Marwick, A. E., & Partin, W. C. (2022). Constructing alternative facts: Populist expertise and the qanon conspiracy. *New Media & Society*, 14614448221090201.

Mohammad, S., Bravo-Marquez, F., Salameh, M., & Kiritchenko, S. (2018). Semeval-2018 task 1: Affect in tweets. *Proceedings of the 12th international workshop on semantic evaluation*, 1–17.

Monti, C., Cinelli, M., Valensise, C., Quattrociocchi, W., & Starnini, M. (2023). Online conspiracy communities are more resilient to deplatforming. *PNAS nexus*, *2*(10), pgad324.

Mosseri, A. (2023). Instagram ranking explained.

Munger, K., & Phillips, J. (2022). Right-wing youtube: A supply and demand perspective. *The International Journal of Press/Politics*, *27*(1), 186–219.

Paletz, S. B., Golonka, E. M., Pandža, N. B., Stanton, G., Ryan, D., Adams, N., Rytting, C. A., Murauskaite, E. E., Buntain, C., Johns, M. A., et al. (2024). Social media emotions annotation guide (smemo): Development and initial validity. *Behavior Research Methods*, *56*(5), 4435–4485.

Papasavva, A., Blackburn, J., Stringhini, G., Zannettou, S., & De Cristofaro, E. (2020). " is it a qoincidence?": A first step towards understanding and characterizing the qanon movement on voat. co. *arXiv preprint arXiv:2009.04885*.

Papasavva, A., Zannettou, S., De Cristofaro, E., Stringhini, G., & Blackburn, J. (2020). Raiders of the lost kek: 3.5 years of augmented 4chan posts from the politically incorrect board. *Proceedings of the international AAAI conference on web and social media*, *14*, 885–894.

Peucker, M., & Fisher, T. J. (2023). Mainstream media use for far-right mobilisation on the alt-tech online platform gab. *Media, Culture & Society*, *45*(2), 354–372.

Prochaska, S., Vera, J., Tan, D. L., & Starbird, K. (2024). Misleading ourselves: How disinformation manipulates sensemaking. *arXiv preprint arXiv:2410.14858*.

Rauchfleisch, A., & Kaiser, J. (2021). Deplatforming the far-right: An analysis of youtube and bitchute. *Available at SSRN 3867818*.

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

Rice, N. M., Horne, B. D., Luther, C. A., Borycz, J. D., Allard, S. L., Ruck, D. J., Fitzgerald, M., Manaev, O., Prins, B. C., Taylor, M., et al. (2022). Monitoring event-driven dynamics on twitter: A case study in belarus. *SN Social Sciences*, *2*(4), 36.

Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., & Rand, D. G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, *58*(4), 1064–1082.

Shibutani, T. (1966). *Improvised news: A sociological study of rumor*. Ardent Media.

Sipka, A., Hannak, A., & Urman, A. (2022). Comparing the language of qanon-related content on parler, gab, and twitter. *Proceedings of the 14th ACM Web Science Conference 2022*, 411–421.

Spiro, E., & Starbird, K. (2023). Rumors have rules. *Issues in Science and Technology*, *29*(3).

Starbird, K., DiResta, R., & DeButts, M. (2023). Influence and improvisation: Participatory disinformation during the 2020 us election. *Social Media+ Society*, *9*(2), 20563051231177943.

Thiele, D. (2024). How covid-19 and the news shaped populism in facebook comments in seven european countries.: A computational analysis. *Computational Communication Research*, *6*(1), 1.

Thomas, D. R., & Wahedi, L. A. (2023). Disrupting hate: The effect of deplatforming hate organizations on their online audience. *Proceedings of the National Academy of Sciences*, *120*(24), e2214080120.

Toftrup, M., Sørensen, S. A., Ciosici, M. R., & Assent, I. (2021). A reproduction of apple's bi-directional lstm models for language identification in short strings. *arXiv preprint arXiv:2102.06282*.

Trujillo, M., Gruppi, M., Buntain, C., & Horne, B. D. (2020). What is bitchute? characterizing the free speech alternative to youtube. *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, 139–140.

Trujillo, M. Z., Gruppi, M., Buntain, C., & Horne, B. D. (2022). The mela bitchute dataset. *Proceedings of the International AAAI Conference on Web and Social Media*, *16*, 1342–1351.

Tuters, M., & Hagen, S. (2020). (((they))) rule: Memetic antagonism and nebulous othering on 4chan. *New media & society*, *22*(12), 2218–2237.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *science*, *359*(6380), 1146–1151.

Weigel, M., & Gitomer, A. (2024). Hate-sharing: A case study of its prevalence and impact on gab. *New Media & Society*, 14614448241245349.

Wilson, T., & Starbird, K. (2021). Cross-platform information operations: Mobilizing narratives & building resilience through both'big'&'alt'tech. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW2), 1–32.

Yadav, K. (2021). Platform interventions: How social media counters influence operations. *Partnership for Countering Influence Operations, Carnegie Endowment for International Peace. https://carnegieendowment. org/2021/01/25/platform-interventions-how-social-mediacounters-influence-operations-pub-83698*.

Yan, X., Guo, J., Lan, Y., & Cheng, X. (2013). A biterm topic model for short texts. *Proceedings of the 22nd international conference on World Wide Web*, 1445–1456.

YouTube. (2023a). Elections misinformation policies [Accessed on February 21, 2024].

YouTube. (2023b). Medical misinformation policy [Accessed on February 21, 2024].

Zampieri, M., Malmasi, S., Nakov, P., Rosenthal, S., Farra, N., & Kumar, R. (2019). Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). *arXiv preprint arXiv:1903.08983*.

Zannettou, S., Bradlyn, B., De Cristofaro, E., Kwak, H., Sirivianos, M., Stringini, G., & Blackburn, J. (2018). What is gab: A bastion of free speech or an alt-right echo chamber. *Companion Proceedings of the The Web Conference 2018*, 1007–1014.

Zhou, K., Wilson, T., Starbird, K., & Spiro, E. S. (2023). Spotlight tweets: A lens for exploring attention dynamics within online sensemaking during crisis events. *ACM Transactions on Social Computing*, *6*(1-2), 1–33.

Zhou, R., Khemmarat, S., Gao, L., Wan, J., & Zhang, J. (2016). How youtube videos are discovered and its impact on video views. *Multimedia Tools and Applications*, *75*, 6035–6058.

Zieringer, L., & Rieger, D. (2023). Algorithmic recommendations' role for the interrelatedness of counter-messages and polluted content on youtube–a network analysis. *Computational Communication Research*, *5*(1), 109.

Appendix A Validations

To ensure these automated methods are well-suited for our study, we perform several validations and robustness checks. First, video titles should be of reasonable length for these various pre-trained NLP models to work well - even for those trained on short text. In Appendix Figure 1a, we show the complementary cumulative distribution of title lengths, which shows that approximately 20% of video titles were five words or shorter. While it is difficult to say what a *reasonable* text length is for these models - and prior work has not examined this question to our knowledge, approximately 80% of the videos had *long* titles of greater than five words.

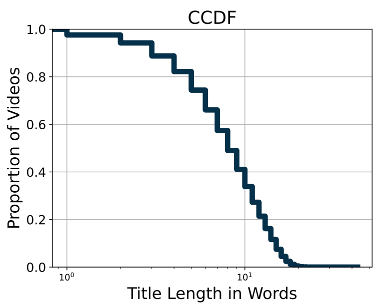

(a) CCDF of BitChute Title Lengths

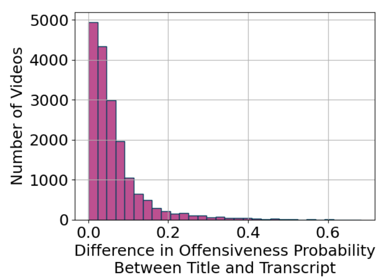

(b) Offensiveness of YouTube Title versus Transcript

Figure 1: (a) A complementary cumulative distribution of video title lengths, in whitespace-delimited words. Approximately 20% of videos have titles five words or shorter, which may be challenging to analyze emotion. (b) Distribution of differences between probability of being offensive measured by the title and probability of being offensive measured by sentences in transcript. On average, the predicted probability of offensiveness of titles differed from the predicted probability of offensiveness of video transcripts by 0.065 (median 0.044).

Second, one can reasonably argue that some information may be lost by using video titles instead of video transcripts. To approximate how much information is lost, we utilized a dataset of titles and transcripts from 17,886 YouTube videos (Hershey et al., 2017; Lovejoy & Prosser, 2021). Using this data, we use the pre-trained RoBERTa model for offensiveness to compute the offensiveness of the titles and the offensiveness of the video transcripts. Specifically, for all videos in the YouTube transcript dataset, we compute the offensiveness of each sentence in the transcript of the video. The median sentence offensiveness is then compared to the title of the video. On average, the predicted probability of offensiveness of titles differed from the predicted probability of offensiveness of video transcripts by 0.065 ± 0.074, suggesting that the offensiveness of video titles is a reasonable proxy for the offensiveness of the videos themselves. The distribution of these differences is shown in Appendix Figure 1b

Lastly, we perform manual validation on samples for each model described above. Similar to the validation done by De Bruyne et al., 2024, two authors (two different combination of authors for each model) examined 600 randomly sampled titles for each model used for feature extraction and marked whether they agreed (A), disagreed (N), or agreed with doubt (D) with the model's label. An acceptance rate for each coder was then computed as follows: $(A + D)/(A + D + N)$ (De Bruyne et al., 2024).

From this analysis we found the following: The acceptance rate of the LSTM for language identification (used to filter the data to English only videos) was between 87% and 88%. The acceptance rate of the pre-trained RoBERTa model for offensiveness was between 84% and 95%. The acceptance rate of VADER sentiment was between 68% and 87%. The acceptance rate of the LIWC categories used in this study was between 88% and 92%.

While the acceptance of the specific LIWC categories was relatively high, its important to note the limitations. Unlike the other models used in this study that make predictions based on training data, LIWC is simply counting words from a dictionary. Dictionary-based methods like LIWC have been validated and used widely, but they can struggle to generalize to new contexts, particularly if new words, slang, or *dog-whistles* are used (like one may find on BitChute (Horne, 2022)). Furthermore, we used LIWC to capture the higher-level concept of "conspiracy language" as was done in prior work (Boyd et al., 2022); however, as suggested by our manual validation, this operationalization is under-counting the number of conspiracy theories on the platform, and thus should not be thought of as capturing what videos are conspiracy theories and which are not.

As described in the methods section of the paper, we used two different emotion recognition models and used the one with a higher acceptance for modeling in the paper. The first model was TweetNLP's pre-trained model, which received an acceptance rate between 35% and 76%. It was noted during the coding process that the model was perhaps not granular enough to capture the correct emotions expressed in the video titles, capturing only 9 emotion categories. Hence, we tried a more recent pre-trained model: the Paletz pre-trained Demux model, which had 27 emotion categories and had been validated on a variety of social media datasets (see Paletz et al., 2024). This model received an acceptance rate between 70% and 78%, and it was used in for analysis in the paper.

Together, these validations suggest that both the NLP methods used and the unit of analysis capture *something* about the content within the videos. However, as expected, they are not perfect. While many of our key results are not dependent on these methods, these limitations should be taken into account when interpreting the results of this work.