A View into YouTube View Fraud

Dhruv Kuchhal
dkuchhal@gatech.edu
Georgia Institute of Technology
USA

Frank Li
frankli@gatech.edu
Georgia Institute of Technology
USA

ABSTRACT

Social media platforms are driven by user engagement metrics. Unfortunately, such metrics are susceptible to manipulation and expose the platforms to abuse. Video view fraud is a unique class of fake engagement abuse on video-sharing platforms, such as YouTube, where the view count of videos is artificially inflated. There exists limited research on such abuse, and prior work focused on automated or bot-driven approaches. In this paper, we explore organic or human-driven approaches to view fraud, conducting a case study on a long-running YouTube view fraud campaign operated on a popular free video streaming service, 123Movies. Before 123Movies users are allowed to access a stream on the service, they must watch an unsolicited YouTube video displayed as a pre-roll advertisement. Due to 123Movies' popularity, this activity drives large-scale YouTube view fraud. In this study, we reverse-engineer how 123Movies distributes these YouTube videos as pre-roll advertisements, and track the YouTube videos involved over a 9-month period. For a subset of these videos, we monitor their view counts and metrics for their respective YouTube channels over the same period. Our analysis reveals the characteristics of YouTube channels and videos participating in this view fraud, as well as the efficacy of such view fraud efforts. Ultimately, our study provides empirical grounding on organic YouTube view fraud.

CCS CONCEPTS

• **General and reference** → **Empirical studies**; **Measurement**; • **Information systems** → *Online advertising*; **Social networks**; • **Security and privacy** → **Social network security and privacy**.

KEYWORDS

View Fraud, Fake Engagement, Online Abuse

ACM Reference Format:
Dhruv Kuchhal and Frank Li. 2022. A View into YouTube View Fraud. In *Proceedings of the ACM Web Conference 2022 (WWW '22), April 25–29, 2022, Virtual Event, Lyon, France.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3485447.3512216

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France

ACM ISBN 978-1-4503-9096-5/22/04…$15.00
https://doi.org/10.1145/3485447.3512216

1 INTRODUCTION

As social media takes center stage in modern discourse, protecting online platforms from abuse and manipulation is essential. Social media platforms, such as YouTube, Facebook, Vimeo, and Dailymotion, are driven by engagement metrics, including the view count of videos and the number of likes on posts. Manipulation of these metrics is a form of online abuse.

Video view fraud is one distinct class of fake engagement abuse, which entails artificially inflating the view count of videos (i.e., when the total view count exceeds the number of solicited views by legitimate users). Beyond the content popularity and visibility manipulation also offered by other fake engagement vectors, video view fraud uniquely provides a direct monetization mechanism as many platforms (including YouTube) pay video creators based on video views (specifically, platforms share a portion of the revenue derived from ads displayed when users view a video). Furthermore, the bar for perpetrating video view fraud is low; while other fake engagement efforts require logged-in users to execute an action (e.g., liking or commenting on content, or sending friend requests), video view fraud simply requires an arbitrary user to visit a web page, where the video can start playing automatically.

While other forms of fake engagement have been studied previously [7, 12, 13, 19], there exists limited empirical characterization of real-world video view fraud. The prior work that does exist [19] focuses on bot-driven automated view fraud. In this study, we expand our understanding of video view fraud by investigating organic or human-driven approaches. Organic view fraud relies on real human users (rather than bots) to generate views, presumably because manually-generated fake views may be more challenging to detect and block. Such activity is still classified as view fraud as users do not intentionally request to watch the videos [16].

In this paper, we conduct a case study of a large-scale, long-running, organic YouTube view fraud operation on one of the most popular free video streaming services [15], 123Movies[1] [21]. Before users can watch a stream on 123Movies, a YouTube video is displayed as a pre-roll advertisement and automatically played, thus generating a fake view for that video. Given 123Movies' immense popularity, this activity results in large-scale YouTube view fraud. In our study, we reverse-engineer how 123Movies distributes YouTube videos as pre-roll ads, and collect the YouTube videos it distributes over a 9-month period. For a subset of these videos, we gather detailed metadata about video characteristics (e.g., view count, content category), and track their dynamics over time. We similarly collect metadata about the YouTube channels associated with the videos and analyze their participation in the scheme.

[1] In a 2018 investigation, the Motion Picture Association of America (MPAA) found 123Movies to be the world's "most popular illegal site" serving 98 million visitors a month [15].

Our analysis sheds light on the characteristics of this view fraud ecosystem. We find hundreds of thousands of videos associated with tens of thousands of channels participating in this view fraud effort. These videos and channels skew heavily towards music, gaming, and entertainment content, hinting at certain industries that may engage in such activity. In fact, we observe extremely popular music videos (with millions of views) involved. We also assess the outcome of participating in view fraud, observing that while videos gain views in the short term, the majority of videos eventually lose a substantial amount of views (sometimes even within a week), presumably due to YouTube's detection and removal of abusive activity. Thus, the long-term success of this abuse seems to be limited (although our constrained visibility into YouTube and the view fraud campaign's operations prohibits establishing a definitive relationship between the view fraud and outcomes). By analyzing the channels that participate in view fraud, we identify that most participate only once or over a short period of time (e.g., a few days), suggesting that the participating channels also observe limited returns on investing in view fraud. Nonetheless, this view fraud operation has survived for years, raising questions about the economics at play and mechanisms that can combat view fraud, beyond detecting and removing fraudulent views.

Ultimately, this work provides insights on an organic form of video view fraud, expanding on the limited prior work. Our findings provide further empirical grounding on attacker behavior and the outcomes of online abusive activities in practice.

2 BACKGROUND

Here we provide background context about YouTube view fraud, as well as on related prior work.

2.1 YouTube View Fraud

On YouTube, a video's view count is its currency. More views can translate to higher video rankings in search results (potentially attracting more organic visits), as well as represent broader popularity and approval for the content creator (particularly relevant for content creators in marketing themselves [11]). Furthermore, video creators make money based on the ad impressions generated while users view their videos. Thus, content creators strive for higher view counts for their videos.

123Movies, a popular free video streaming service, facilitates artificially inflating YouTube video view counts. Before 123Movies displays a video stream to a user, it overlays a pre-roll ad on the stream (similar to legitimate ad-supported video platforms, such as YouTube and Dailymotion). However, unlike pre-roll ads on legitimate platforms, 123Movies' pre-roll videos are regular YouTube videos, not ads. As shown in Figure 1, users must watch a YouTube video for at least 30 seconds before clicking through to the stream, which is the duration that YouTube requires a video to be played before recording a view [9]. While organically driven, these YouTube video views are classified as view fraud, as content creators pay an ad network to obtain views for their videos [17], and the ad network in turn pays 123Movies to show its users the videos [18].

2.2 Related Work

The YouTube platform has been widely studied over the past decade, from analysis of user-generated video content [4] to identification of topic-based communities [8]. More recently, several studies have also detected and characterized various kinds of abuse permeated via YouTube. For example, researchers have investigated the marketing of fraudulent products and services on YouTube [3, 20].

Most relevant to our study, several prior works have considered the manipulation of video engagement metrics. Dutta et al. studied collusion networks that deliver fake collusive likes and subscribers to videos and channels, although they did not consider view fraud [7]. Li et al. investigated a behavior-based clustering approach to detecting automated fake engagement on YouTube [12]. Miriam et al. evaluated the detection systems of five online video platforms (including YouTube, Vimeo, and Dailymotion) and found that YouTube better detected fake views than other platforms [13]. They also observed that YouTube penalizes its videos' public view counters after detecting view fraud activity.

However, there has been little empirical investigation into real-world video view fraud campaigns. What limited work does exist considered detecting and characterizing view bots on live streaming platforms like Twitch [19]. Thus, our work expands upon the literature by providing empirical grounding on video view fraud in practice, in particular considering an organic form of view fraud, rather than the automated view fraud considered in prior work.

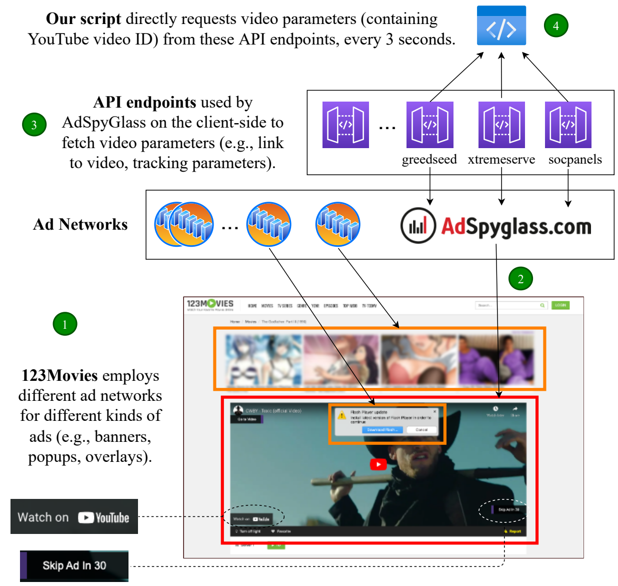

Figure 1: When accessing a stream on 123Movies, (1) various ads are loaded throughout the web page, including a YouTube video as a pre-roll ad overlayed on the video stream. (2) This YouTube video is loaded by first issuing a request to the ad network `AdSpyGlass`, which returns a response containing another subservice endpoint. (3) These subservice endpoints provide the actual YouTube video link. (4) To collect videos participating in view fraud, our milker script directly queries the subservices for YouTube videos.

3 METHOD

Here we describe our study's method for investigating real-world YouTube view fraud.

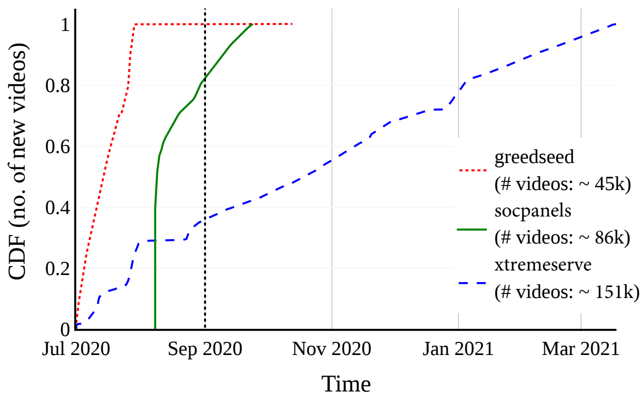

Figure 2: Temporal distribution of newly observed YouTube videos involved in view fraud, collected from the 3 subservices over time. Due to YouTube API quota constraints, our metadata analysis is limited to the initial videos collected between July and September (indicated with the vertical dotted black line).

3.1 Data Collection

To understand the YouTube view fraud occurring through 123Movies, we collect the IDs of the videos involved and their metadata, including information about the associated YouTube channels.

3.1.1 Collecting Videos Receiving Fake Views. To start, in July 2020, we reverse-engineered how YouTube videos are delivered as pre-roll ads on 123Movies. Using `mitmproxy` [5] (as 123Movies uses HTTPS), we monitored the plaintext web traffic when accessing streams on 123Movies[2]. We identified that when a stream is requested, as shown in Figure 1, 123Movies interacts with multiple ad networks to request different types of ads (e.g., popups and banners). We traced the requests that fetched YouTube videos to the API endpoint `deliver.vkcdnservice.com`. We believe that this endpoint belongs to an ad network named `AdSpyGlass` [1], as `deliver.vkcdnservice.com` redirects to `AdSpyGlass.com` when loaded in a browser. Using the Wayback Machine [2], we find that `AdSpyGlass.com` has existed since at least April 28, 2019. Thus, we believe that this service has been operational for years (as it remains online at the time of this paper's publication).

When a video is requested from the `AdSpyGlass` endpoint, it responds with the URL of another service endpoint, which actually provides the final YouTube video link. We observe that across all responses, only a small set of secondary service endpoints are used, which we refer to as *subservices.* We monitored 123Movies through July-September 2020 for active subservices, and found three in total. We observed *greedseed*[3], and *xtremeserve*[4] in July, and *socpanels*[5] in August. To collect the YouTube videos involved in this view fraud, we automatically milked all three subservices for YouTube videos, using a Python script that queried each subservice for a video every 3 seconds[6], and recording the IDs of the YouTube videos returned. To avoid detection and rate limiting, we distributed the milker's requests across a pool of 20 proxy servers (all on our university network). Over a 9-month period (from July 2020 to March 2021), we collected ~45k unique videos from `greedseed`, ~151k from `xtremeserve`, and ~86k from `socpanels`. In Figure 2, we plot the temporal distribution of newly collected videos for the three subservices. While `xtremeserve` remained online throughout our data collection, providing new videos at a consistent rate, `greedseed` and `socpanels` ceased operations by October 2020.

[2]We used `mitmproxy` as 123Movies deploys anti-debugging techniques on its web pages, which prevent stream loads when detecting that a browser's debugger is in use.
[3]`greedseed.world/vpaid/getVideo.php`
[4]`xtremeserve.xyz/add.php`
[5]`socpanels.thevideome.site/feed`

Metadata	Snapshot Frequency	# Videos
Advertised YouTube videos	1 day	~150k
Channels of advertised videos		~75k
Non-advertised videos	5 days	~818k

Table 1: Statistics about the metadata snapshots collected from YouTube's Data API.

3.1.2 Collecting Video Metadata. To better investigate the gathered videos involved in the 123Movies view fraud operation, we used YouTube's Data API v3 [6] to periodically collect metadata snapshots for these videos and their channels, starting from when we first observed a video distributed by a subservice up until March 2021 (or until taken down). The metadata we acquired includes the video and channel titles, author, language, topics, and video and channel statistics (e.g., numbers of views and subscribers). Initially, we collected weekly snapshots starting July 12, 2020, but later transitioned to daily snapshots starting August 3, to monitor video view counts at a finer granularity. As we started collecting videos from `socpanels` on August 8, we recorded daily metadata snapshots for all `socpanels` videos. Since YouTube rate limits access to its Data API per API key, we generated 20 API keys (the maximum number of keys allowed by Google per account) to maximize our daily quota. This quota amount allowed us to access daily snapshots for the first ~150k videos gathered (which we reached in September 2020), which were associated with ~75k channels. Every 5 days, we also collected metadata snapshots for other videos hosted by these channels which we did not record as involved in this view fraud effort (~818k videos). Table 1 summarizes our metadata dataset.

3.2 Limitations

Our study takes a step forward in exploring the ecosystem that facilitates YouTube view fraud. However, we lack full visibility into the ecosystem, which leads to several important limitations:

- Our data collection is centered on `123movies2020.org`, a mirror of 123Movies. We note that 123Movies operates by redirecting multiple domains to a working mirror [11], so beyond direct visitors, our mirror likely received visitors from other domains (e.g., `123movies2020.email`). However, there may exist other mirrors of 123Movies or similar streaming sites that function differently, and for which our findings may not hold. Similarly, as described in Section 3.1.1, we manually monitored 123Movies and initially discovered three active subservices. However, more such subservices may exist that behave differently. Nonetheless,

[6]We determined the querying rate through manual experimentation, aiming to avoid rate limiting and overloading the API endpoints.

Subservice	Video Title	# Views (First Seen) ↑	Δ # Views (Last Seen)	Channel Name	Topic	Duration (mins)	Published (Year)
socpanels	Despacito	6.9B	351M	LuisFonsiVEVO	Music	4.7	2017
	BOOMBAYAH	930M	207M	BLACKPINK		4	2016
	Never Gonna Give You Up	735M	158M	RickAstleyVEVO		3.6	2009
xtremeserve	Saara India!	22M	7.3M	T-Series	Music	2.7	2020
	How To Get Your First 1000 Subscribers FAST!	5.7M	0.25M	Chaos	Video Games	14	2016
	black 45	2.5M	-647	Ramneek Sidhu	N/A	45	2018
greedseed	Mera Mera	1M	29k	Aussie Records	Music	4	2020
	AfrotroniX-Solal	0.9M	0.12M	AfrotroniX		3.4	2020
	One More Day- Jiang Tao	0.8M	0.25M	MIRROR		3.5	2019

Table 2: Top 3 videos by the view count observed in its first metadata snapshot, for each subservice.

Subservice	Channel Name	# Videos (Advertised) ↑	# Videos (Total)	# Subscribers	Key Topic	Country	First Published
socpanels	kim lee	124	309	1240	Affiliate Marketing	Singapore	2018
	countrycampingkorea	118	482	7.4k	Entertainment	Korea	2019
	Ben Tre Oi	102	1210	44k	Lifestyle	Vietnam	2013
greedseed	IndieGamerRetro	27	1300	47k	Video Games	USA	2006
	Emma Moore	22	205	3	Lifestyle	N/A	2020
	Benigna Kennel Fiona Tjhin	18	811	17.6k	Pets	Indonesia	2013
xtremeserve	The Hu Music	64	66	2120	Music	N/A	2015
	Vic Stefanu - Amazing World Videos	50	3.3k	202k	Tourism	USA	2010
	Alioth Club	33	53	478	Society	N/A	2020

Table 3: Top 3 channels by the number of videos observed as advertised on 123Movies, for each subservice.

we believe that our large-scale data provide useful insights into an organic view fraud operation.

- We ultimately are unable to distinguish fake views from real organic ones, and it is possible that the YouTube videos we monitor are involved in other view fraud operations. This lack of visibility prevents us from establishing causal relationships between the view fraud we observe and video outcomes, although we explore correlations to the extent possible.
- There are inherent delays between when a video is first advertised on 123Movies as part of the view fraud campaign, when we first milk it, and when we record its first metadata snapshot. These delays limit our observation of initial video dynamics, which we further discuss in Section 4.3.

4 FINDINGS

Here, we analyze our collected datasets to answer three research questions about the investigated organic view fraud ecosystem:
RQ1: What kind of videos receive the fake views?
RQ2: How extensively do videos/channels participate in view fraud?
RQ3: How effective is the view fraud?

4.1 RQ1: What kind of videos receive the fake views?

Content. We first analyze the general themes of the videos receiving fake views, using the topic tags in the video metadata annotations from YouTube. The distribution of topics among videos, as depicted in Figure 3, shows that the largest portion of videos covers music and entertainment topics (~35% of all videos). This observation corroborates prior reports of the music, media, and entertainment industry's participation in view fraud [11]. Lifestyle and gaming were also popular video topics, each involving approximately 16% of videos.

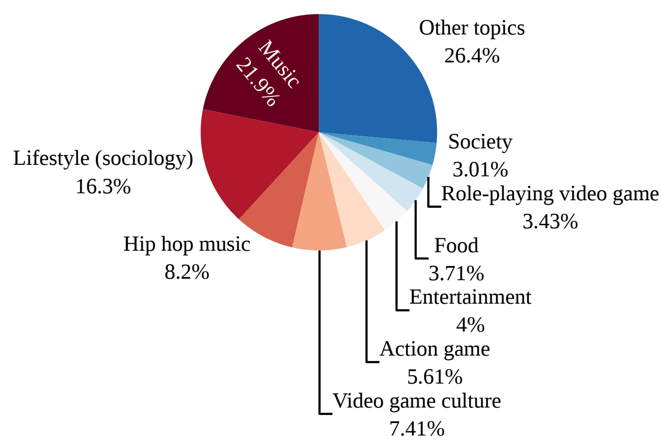

Figure 3: Distribution of topic tags (according to YouTube's Data API) for all view fraud videos. Each subservice exhibits a similar distribution.

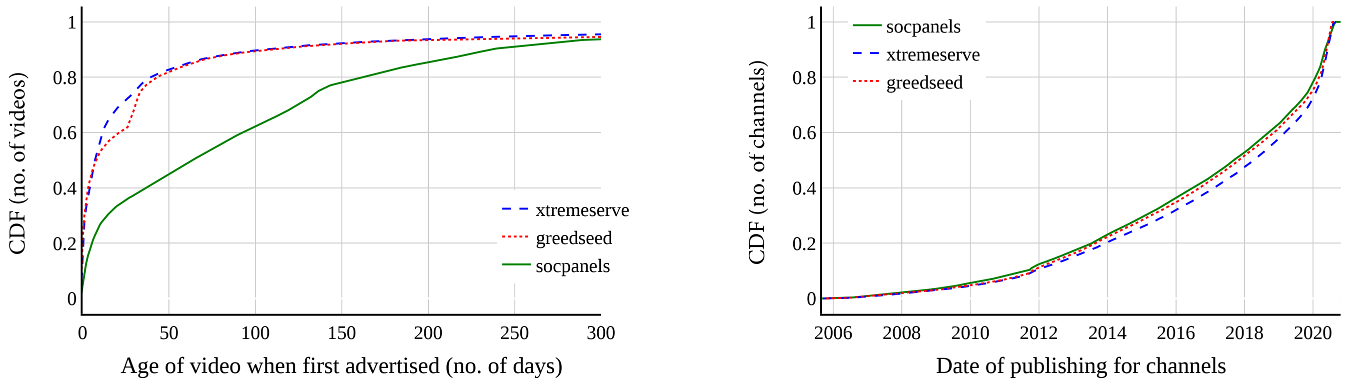

(a) CDF of the age of view fraud videos (in no. of days) when we first observed the videos advertised.

(b) CDF of dates when channels participating in view fraud published their first videos.

Figure 4: Age of videos and channels involved in view fraud.

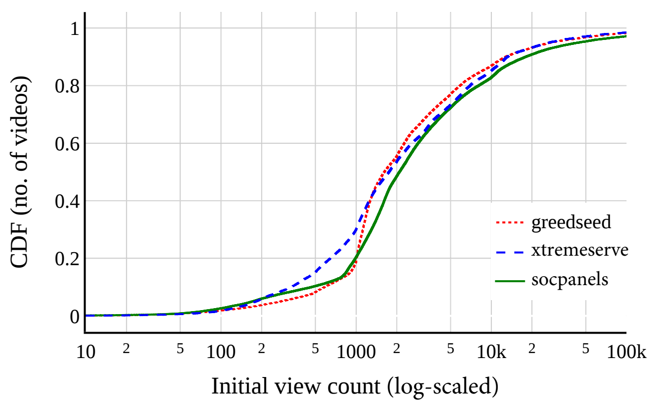

Figure 5: CDF of the view count of videos from their first recorded metadata snapshot.

While YouTube requires content creators to indicate when videos contain "paid promotions", their API does not disclose this information publicly and affiliate marketing disclosures on YouTube have previously been documented as rarely enforced [14]. As a result, we are unable to assess if videos contained commercial content.

Age. To further characterize the videos receiving fraudulent views, we studied the age of the videos and channels involved, as shown in Figure 4. In Figure 4a, we plot the number of days between when a video was published on YouTube and when we first recorded it as advertised in this view fraud ecosystem. (Here, our results are upper bounds on video age, as we may not have observed initial advertisement for videos already being advertised at the start of our data collection.) We see that ~80% of videos collected from `xtremeserve` and `greedseed` and ~45% of those from `socpanels` were published within 50 days of first observed advertisement. Meanwhile, as seen in Figure 4b, the median year of first video publication for YouTube channels participating in the view fraud was 2018, with only ~20% of channels publishing for the first time after 2020 (for all subservices). These findings indicate that while the videos involved in view fraud tend to be relatively new, the participating channels are often long-established, with some over a decade old.

Popularity. We additionally investigate the initial popularity of view fraud videos by evaluating the number of views in our first metadata snapshot for each video (taken within a day of being first milked). We plot this distribution in Figure 5, seeing that over 65% of videos already had more than a thousand views upon our first metadata snapshot, with a noticeable inflection point at approximately a thousand views for `greedseed` and `socpanels`. We hypothesize that these videos (represented near the inflection point in Figure 5) might have had a negligible number of views initially, and gained at least a thousand views upon being advertised on 123Movies (before our first metadata snapshot).

For each subservice, we also look at the top 3 most popular videos (see Table 2) and the top 3 channels by the number of videos advertised on 123Movies (see Table 3). Interestingly, we see that the top videos receiving view fraud include "Despacito" by "LuisFonsiVEVO", which was the most viewed video on YouTube overall from 2018-2020 [22]. Another popular video was "Never Gonna Give You Up" by "RickAstleyVEVO", a 1987 song that is part of a popular Internet meme known as "Rickrolling" [23]. These videos are already popular, with hundreds of millions of views, and we are unclear about their motivations for participating in view fraud.

4.2 RQ2: How extensively do videos and channels participate in view fraud?

Participation across subservices. In our dataset, we observed three separate subservices that `AdSpyGlass` relied upon to distribute YouTube videos for advertisement on 123Movies. We first evaluate the extent to which videos were distributed by different subservices, as depicted in Figure 6. We find limited overlap in the videos milked from different subservices; no more than ~6% of videos from one subservice was observed for another subservice, and less than ~0.5% of videos was seen for both other services. Channels also exhibit limited overlap across subservices, with only 5% of channels observed with videos advertised by two subservices, and only 0.2% of channels observed for all three subservices.

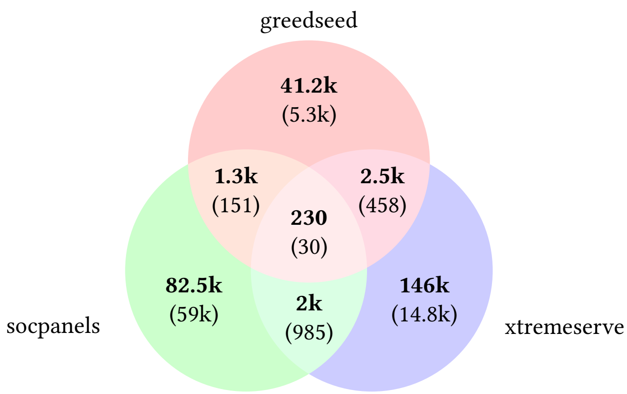

Figure 6: Overlap in the videos provided by the three different subservices. The bold numbers represent all videos collected during our measurement, and the numbers in brackets represent the videos for which we collect daily metadata snapshots (as discussed in Section 3.1).

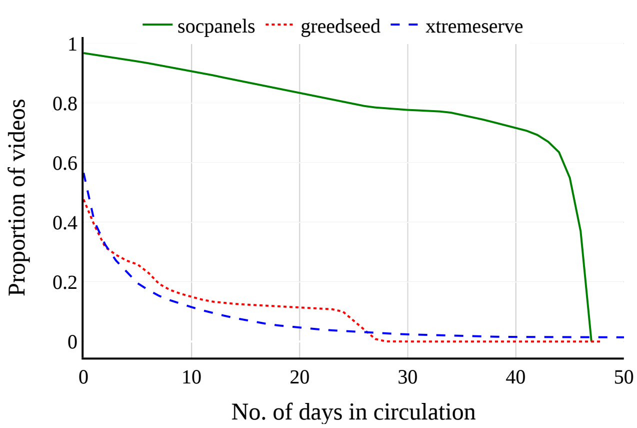

Figure 7: Proportion of videos that remain in circulation (i.e., are being advertised) over time (in days) after their first observed advertisement.

View Fraud Duration per Video. We investigate the duration for which videos are advertised as part of the view fraud operation. In Figure 7, we plot the proportion of videos that remain advertised over different durations, for each subservice. We observe differences across subservices. For `socpanels`, ~80% of videos were advertised for over a month, whereas half of `xtremeserve` and `greedseed` videos were advertised for only a day. We note that for videos already being advertised when we first started milking a subservice, our observed duration lower bounds the true duration. However, only small portions of the videos collected for `xtremeserve` and `greedseed` were from the initial milking periods (as seen in Figure 2), so our high-level observations hold. (For `socpanels`, a larger portion of videos was collected at initial milking, but the advertising duration is already significantly longer than with the other two subservices.) We hypothesize that the differences across subservices may have arisen due to different service offerings. For example, `socpanels`' longer video advertising could be because it

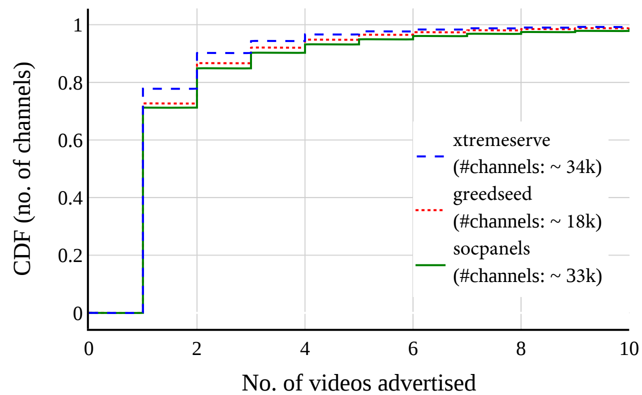

Figure 8: CDF of the number of videos advertised for view fraud per channel, for each subservice.

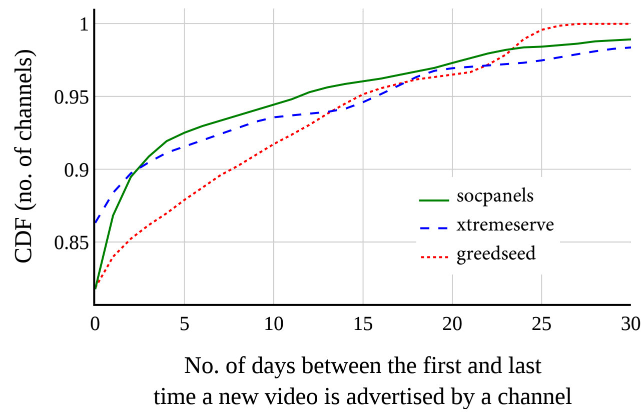

Figure 9: CDF of the duration (in days) over which a channel advertises new videos.

offers monthly view fraud services, instead of a daily offering. Alternatively, if `socpanels` provided guarantees on video view growth, videos may remain actively advertised by `socpanels` for longer periods (see Section 4.3.1). However, we lack further visibility into the different subservices' operations to validate our hypothesis.

View Fraud Participation per Channel. To understand how extensively channels participate in view fraud, we analyze the number of videos we observed receiving view fraud per channel, as depicted in Figure 8. We find that over 70% of channels only advertised one video on 123Movies. In Figure 9, we also plot the duration over which channels advertised new videos. We see that more than 90% of channels only advertised new videos over a period of up to 10 days. We note that channels possibly participated in this view fraud campaign prior to our data collection, and our results are lower bounds. However, given the short periods that channels were observed actively participating with new videos, we believe it is unlikely that we missed observing significant amounts of prior channel activity. Thus, channels appear to only participate in this view fraud activity to a limited extent, either involving few videos or participating only over a brief duration. This observation naturally raises a question about the efficacy of this ecosystem, our final research question.

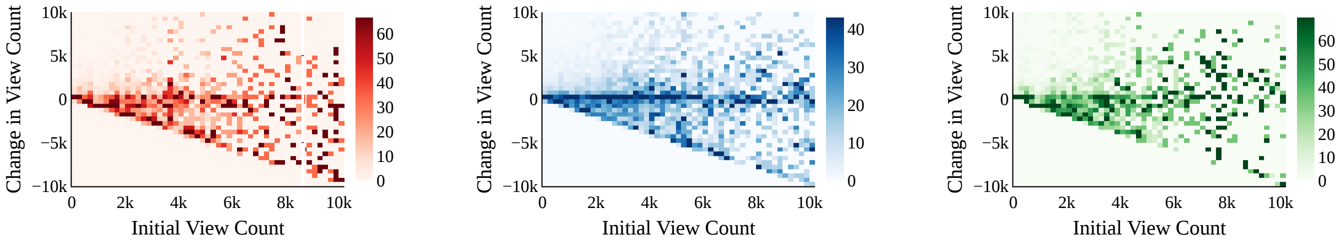

(a) `greedseed` (b) `xtremeserve` (c) `socpanels`

Figure 10: Heatmaps of the distribution of net view count changes for videos, conditioned on their initial view counts. A video's net view count change is calculated by comparing its first and last metadata snapshot.

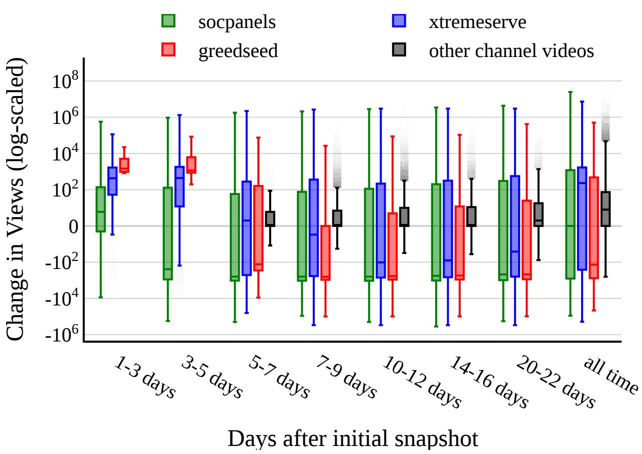

Figure 11: Boxplots of the net view count changes for videos, over different time periods after their first metadata snapshot. We consider videos across the three subservices, as well as the other videos of channels participating in the view fraud campaign that were not observed as advertised. The *all time* x-axis label represents the long-term view count change, comparing a video's first and last metadata snapshot (regardless of time period). Note that outliers beyond the boxplot whiskers are marked as individual data points.

4.3 RQ3: How effective is the view fraud?

Finally, we aim to evaluate the efficacy of the view fraud. To do so, we require analyzing videos for which we observed the beginning of their advertisement in the view fraud campaign. We identify such videos through multiple filtering steps.

First, we filter out the small portion of videos that appear across multiple subservices, to avoid convoluting the influences of different subservices. We then filter out videos that were already advertised when we began milking a subservice, as we may not have observed the start of their involvement in view fraud. For each subservice, we analyzed the growth in videos collected and identified the inflection point when the rate of new video collection reached a steady state, reflecting the natural rate of new videos being distributed by a subservice (rather than the gathering of already advertised videos that we had simply not yet observed). This initial data collection period lasted for 3, 11, and 15 days for `xtremeserve`, `greedseed`, and `socpanels`, respectively, and we filtered out the videos observed during those periods.

Finally, we further filter out videos where their first metadata snapshot was taken more than 12 hours after first observed advertisement, as the initial view counts recorded later than that duration may be too inaccurate. We explored other thresholds but chose the 12 hours window to balance the accuracy of the initial view count observed with the size of the remaining unfiltered video population across subservices. Our remaining unfiltered video population consists of 13.3k videos for `xtremeserve`, 4.7k for `greedseed`, and 3.8k for `socpanels`.

In our analysis, we consider two penalties that YouTube applies to discourage view fraud. First, YouTube discounts views that they identify as fake engagement [13]. Second, YouTube takes down videos that violate its Terms of Service, including those that artificially manipulate engagement metrics [24].

4.3.1 Net Change in View Counts. We analyze the unfiltered video population (for which we observe their initial advertisement by a subservice and obtain a timely first metadata snapshot), comparing the net change in view counts for these videos conditioned on (i) time and (ii) the initial view count.

View Count Change over Time. In Figure 11, we depict how the video view counts change over time by plotting the distributions of the view count changes at each 2-day interval across the first 20 days after the initial video advertisement, as well as the *all time* distribution, which compares the first and last snapshot per video across our full measurement period (using the last valid snapshot, if a video is taken down during our measurement). For each time interval, we display separate boxplots for each of the three subservices, as well as one for all other videos of channels participating in view fraud, where these videos were not observed as advertised.

We see that in the first few days after their initial metadata snapshot, the majority of videos gain views (sometimes in the thousands) for all three services. However, within a week, the median video has negative net growth in view count, across all three subservices. Ultimately, the median long-term change in view count is `0`, `232`, and `-136` for `socpanels`, `xtremeserve`, and `greedseed`, respectively. We also find that the median view count change for other channel videos that were not observed as advertised is small but positive (8). However, a quarter of these other channel videos had net negative view count changes, suggesting that they too

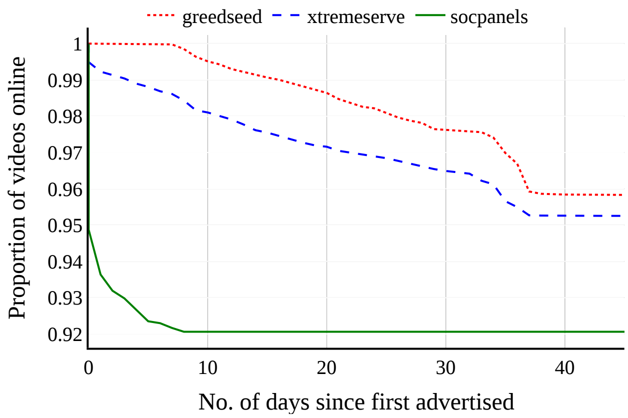

Figure 12: Proportion of videos receiving view fraud that remain online over time after first advertisement.

were involved in other view fraud efforts (or alternatively are being falsely penalized by YouTube). These results suggest that while this view fraud campaign might be effective in the short term, YouTube is able to largely mitigate its visible impact in the long term.

View Count Change by Initial View Count. We next assess the distribution of view count changes, conditioned on the videos' initial view counts. In Figure 10, we plot heatmaps that visualize the distribution of view count changes (across our full measurement period) for videos with up to 10k initial views (which is over 80% of videos for all three subservices, as seen in Figure 5).

We observe that for all three subservices, the majority of videos experience net negative view growth, with the highest densities of videos near the $y = 0$ line, which indicate limited net change in view count, as well as the $y = -x$ line, which indicates that all initial views observed were removed. This result reinforces our conclusion that YouTube is able to mitigate view fraud in the long term. We note that a small portion of videos does experience sizable view count growth, which could be false negatives in YouTube's detection or cases where these videos attract real organic viewers.

4.3.2 Video Takedowns. When a YouTube video is taken down, its metadata snapshot indicates that the video is not found or unavailable. We note that a video may be taken down for various reasons beyond artificially manipulating engagement metrics [24], such as copyright/trademark infringements or inappropriate content, although its metadata does not indicate the specific takedown justification. Nonetheless, we believe that it is valuable to characterize the takedown of the videos observed participating in view fraud.

Figure 12 plots the proportion of videos that remain online over time after the video is first advertised. We find that in total, less than 5% of `greedseed` and `xtremeserve` videos were taken down within 50 days, whereas for `socpanels`, nearly 8% of videos were removed within 10 days of being advertised (with about 5% removed within the first day). We hypothesize that `socpanels` videos may have been more heavily penalized due to a higher rate of video distribution; during our monitoring of 123Movies, we observed that `AdGlassSpy` returned `socpanels` videos more frequently than for the other two subservices. For all other videos of the channels participating in the view fraud that were not observed as advertised, we found only 5 were taken down in total (out of 818k). Thus, videos receiving view fraud are significantly more likely to be removed, although only a small minority of videos receive such penalties, and YouTube appears to primarily rely on eliminating fake views instead of taking down videos.

5 CONCLUDING REMARKS

In this paper, we conducted an empirical investigation of a large-scale organic YouTube view fraud campaign. We monitored the campaign's operations over time to characterize the participants in this ecosystem, their behaviors, and the outcomes of the view fraud. What we found was an expansive ecosystem with hundreds of thousands of videos from tens of thousands of channels.

Our investigation into the success of view fraud efforts suggests that benefits are primarily short-term, and that YouTube is able to quickly detect and remove many of the fake views. We also observed that this operation has only a few "repeat customers", who typically participate only for a brief period of time, perhaps recognizing the poor return on investing in the view fraud. This brings into question how this operation continues to thrive over years. It seemingly exhibits "snake oil" properties, where the promised outcome (i.e., growth in views) is not truly delivered. Like with many snake oil scams, we hypothesize that this ecosystem survives by luring in new unsuspecting participants. However, future work can build on our initial results to study how participants are drawn into this scheme and the economics at play. It is interesting to note that over half of participating channels have over 1k subscribers, which is one of the qualifications required for receiving ad revenue from videos via the YouTube partner program [10]. It is possible that ad revenue from videos could change the cost-benefit tradeoff for participants (such as through arbitrage, where fake views are cheaper than the ad revenue received), although the lack of long-term participants suggests otherwise.

Overall, the view fraud campaign remains persistent, despite YouTube's ability to detect and remove fake views. Thus, further work is needed on other methods to combat this type of abuse. For example, we observed that our studied view fraud operation relied on link redirection when loading the YouTube videos (such as using Twitter's `t.co` link shortener and Google Plus's redirection link[7]). This redirection is done presumably to obfuscate the HTTP referer, which would reveal the true location of where the YouTube videos are displayed (i.e., 123Movies) and potentially lead to easier detection and filtering by YouTube. Therefore, one signal for detecting videos involved in such campaigns could be the frequency with which a video is accessed from such redirection links. Other socio-technical directions may also need to be considered, such as more punitive penalties to discourage such view fraud activity or legal action. Ultimately, we consider our study to be a step forward in better understanding video view fraud in practice, informing future exploration.

REFERENCES

[1] AdSpyGlass. 2021. AdSpyGlass.com. https://www.adspyglass.com/.
[2] Internet Archive. 2018. Wayback Machine. https://web.archive.org/.

[7] `https://plus.google.com/url?<youtube-link>`

[3] Vlad Bulakh, Christopher W. Dunn, and Minaxi Gupta. 2014. Identifying Fraudulently Promoted Online Videos. In *The International Conference on World Wide Web*.

[4] Meeyoung Cha, Haewoon Kwak, Pablo Rodriguez, Yong-Yeol Ahn, and Sue Moon. 2007. I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User Generated Content Video System. In *The ACM SIGCOMM Conference on Internet Measurement*.

[5] Aldo Cortesi, Maximilian Hils, Thomas Kriechbaumer, and contributors. 2021. mitmproxy: A free and open source interactive HTTPS proxy. https://mitmproxy.org/.

[6] Google Developers. 2021. YouTube Data API. https://developers.google.com/youtube/v3.

[7] Hridoy Sankar Dutta, Mayank Jobanputra, Himani Negi, and Tanmoy Chakraborty. 2020. Detecting and analyzing collusive entities on YouTube. In *arXiv*. https://arxiv.org/abs/2005.06243

[8] Ullas Gargi, Wenjun Lu, Vahab Mirrokni, and Sangho Yoon. 2011. Large-scale community detection on YouTube for Topic Discovery and Exploration. In *AAAI Conference on Weblogs and Social Media*.

[9] YouTube Help. 2021. How engagement metrics are counted. https://support.google.com/youtube/answer/2991785.

[10] YouTube Help. 2022. YouTube Partner Program overview. https://support.google.com/youtube/answer/72851.

[11] Michael Keller. 2018. The Flourishing Business of Fake YouTube Views. https://www.nytimes.com/interactive/2018/08/11/technology/youtube-fake-view-sellers.html.

[12] Yixuan Li, Oscar Martinez, Xing Chen, Yi Li, and John E. Hopcroft. 2016. In a World That Counts: Clustering and Detecting Fake Social Engagement at Scale. In *The International Conference on World Wide Web*.

[13] Miriam Marciel, Rubén Cuevas, Albert Banchs, Roberto González, Stefano Traverso, Mohamed Ahmed, and Arturo Azcorra. 2016. Understanding the Detection of View Fraud in Video Content Portals. In *The International Conference on World Wide Web*.

[14] Arunesh Mathur, Arvind Narayanan, and Marshini Chetty. 2018. Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest. In *The ACM on Human-Computer Interaction*.

[15] An Nguyen. 2018. World's most popular pirated movie site being run from Vietnam: US trade association. https://e.vnexpress.net/news/news/world-s-most-popular-pirated-movie-site-being-run-from-vietnam-us-trade-association-3722948.html.

[16] Erick Schonfeld. 2011. Yup, YouTube Counts Video Ads As Regular Views. https://techcrunch.com/2011/06/19/youtube-counts-video-ads-regular-views/.

[17] Sergey. 2022. How can I buy ads from Tube Corporate websites? http://feedback.tubecorporate.com/en/articles/2864513-how-can-i-buy-ads-from-tube-corporate-websites.

[18] Sevak. 2022. How to set up In-Video (VAST) ads. https://support.adspyglass.com/en/articles/4141834-how-to-set-up-in-video-vast-ads.

[19] Neil Shah. 2017. FLOCK: Combating Astroturfing on Livestreaming Platforms. In *The International Conference on World Wide Web*.

[20] Ashutosh Tripathi, Kusum Kumari Bharti, and Mohona Ghosh. 2019. A Study on Characterizing the Ecosystem of Monetizing Video Spams on YouTube Platform. In *The International Conference on Information Integration and Web-Based Applications and Services*.

[21] Wikipedia. 2021. 123Movies. https://en.wikipedia.org/wiki/123Movies.

[22] Wikipedia. 2021. List of Most Viewed YouTube videos. https://en.wikipedia.org/wiki/List_of_most-viewed_YouTube_videos.

[23] Wikipedia. 2021. Rickrolling. https://en.wikipedia.org/wiki/Rickrolling.

[24] YouTube. 2021. Terms of Service. https://www.youtube.com/t/terms.