

Graphika

Cross-Platform Spam Network Targeted Hong Kong Protests

“Spamouflage Dragon”
used hijacked and
fake accounts to amplify
video content

9.2019

By Ben Nimmo, C. Shawn Eib, L. Tamora



Cross-Platform Spam Network Targeted Hong Kong Protests

“Spamouflage Dragon” used hijacked and fake accounts to amplify video content

An active and prolific, but ultimately low-impact, cross-platform political spam network in Chinese boosted attacks on the Hong Kong protesters by using hijacked or fake accounts on YouTube, Twitter, and Facebook well into September 2019, an investigation by Graphika reveals. The behavior appeared designed to support the Chinese government and discredit its critics, both at home and abroad.

In August 2019, Twitter, Facebook, and YouTube all took down assets belonging to a Chinese state-linked information operation. Based on analysis of these networks and their activity, Graphika discovered an amplification network, far less professional, that boosted attacks on critics of the Chinese regime, among other spam activities. This network appeared to have been set up in some haste, built from commercially available online assets that often bore little visible relation to China, rather than dedicated accounts.

Elements of the network were already active in the second half of 2018, when it primarily attacked Chinese émigré billionaire Guo Wengui (郭文贵, also known as Miles Kwok). It struggled to achieve impact or to reach audience eyes; nevertheless, it shines further light on the manipulation of political dialogue in the Chinese-language space. It also demonstrates how spam networks can be used to amplify political messaging while staying below the threshold of larger and better coordinated human operations.

The network’s assets interspersed their political posts with high volumes of spam-like content, including TikTok videos and landscape photographs. These spam posts appeared designed as camouflage, diluting the political content with more human-interest posts; Graphika has therefore nicknamed the network “Spamouflage Dragon.”

Twitter Assets Lead to YouTube

Graphika’s investigation began after Twitter published over [3.5 million](#) tweets posted by assets linked to the Chinese state, following the August takedown. Analysts at the Atlantic Council’s [DFRLab](#) and the [Australian Strategic Policy Institute](#) concluded that the assets had largely been spam accounts, bought in bulk to amplify attacks on the Hong Kong protests.

Graphika found that several of the Twitter accounts repeatedly shared links to a YouTube channel called [Rumor Shredder](#) (谣言碎纸机), which primarily posted attacks on Guo. It was taken down sometime in the second quarter of 2019 for unspecified violations, and all its videos deleted, but a channel with the [same name](#) began posting anti-Guo content on April 9, 2019. The [first video](#) posted by the new channel was identical to a [video](#) posted by the old channel on or before October 23, 2018.



Left, share of a video from the original "Rumor Shredder" on [Facebook](#), October 23, 2018. Right, video posted by the new "Rumor Shredder" on [YouTube](#), April 8, 2019. Google translated, the headline reads, "Full of ridiculous words, how long do you see?"

Graphika concluded that the two accounts were connected and that the April 2019 Rumor Shredder was a recidivist version of the earlier one.

Searching for the Chinese version of the headline on YouTube, the Twitter data dump, and Facebook turned up an unusual pattern of amplification. On Twitter, five different assets in the government takedown shared the headline or URL seven times between October 23 and October 26, 2018. All appeared to be high-volume spam accounts with low follower numbers; one had earlier posted English-language spam.

On YouTube, seven accounts posted a video with the identical name and visuals in late October 2018. Bizarrely, four of those accounts only posted two-second clips; the other three posted the full version.

满口荒唐言，看你蹦哒到几时？



满口荒唐言，看你蹦哒到几时

Gennadiy Muzin · No views · 10 months ago



满口荒唐言，看你蹦哒到几时？

李维斯 · No views · 10 months ago



满口荒唐言 看你蹦哒到几时

封易 · No views · 10 months ago



满口荒唐言，看你蹦哒到几时？

Mirosław Belotelov · No views · 10 months ago



满口荒唐言，看你蹦哒到几时？

小钱包 · No views · 10 months ago



满口荒唐言，看你蹦哒到几时？

Lew Belchenko · No views · 10 months ago

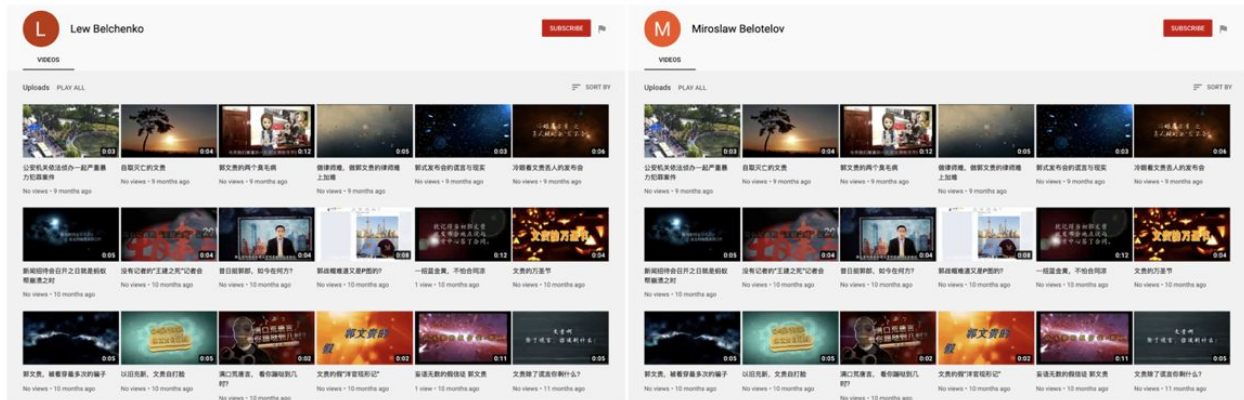


满口荒唐言，看你蹦哒到几时？

yk a · No views · 10 months ago

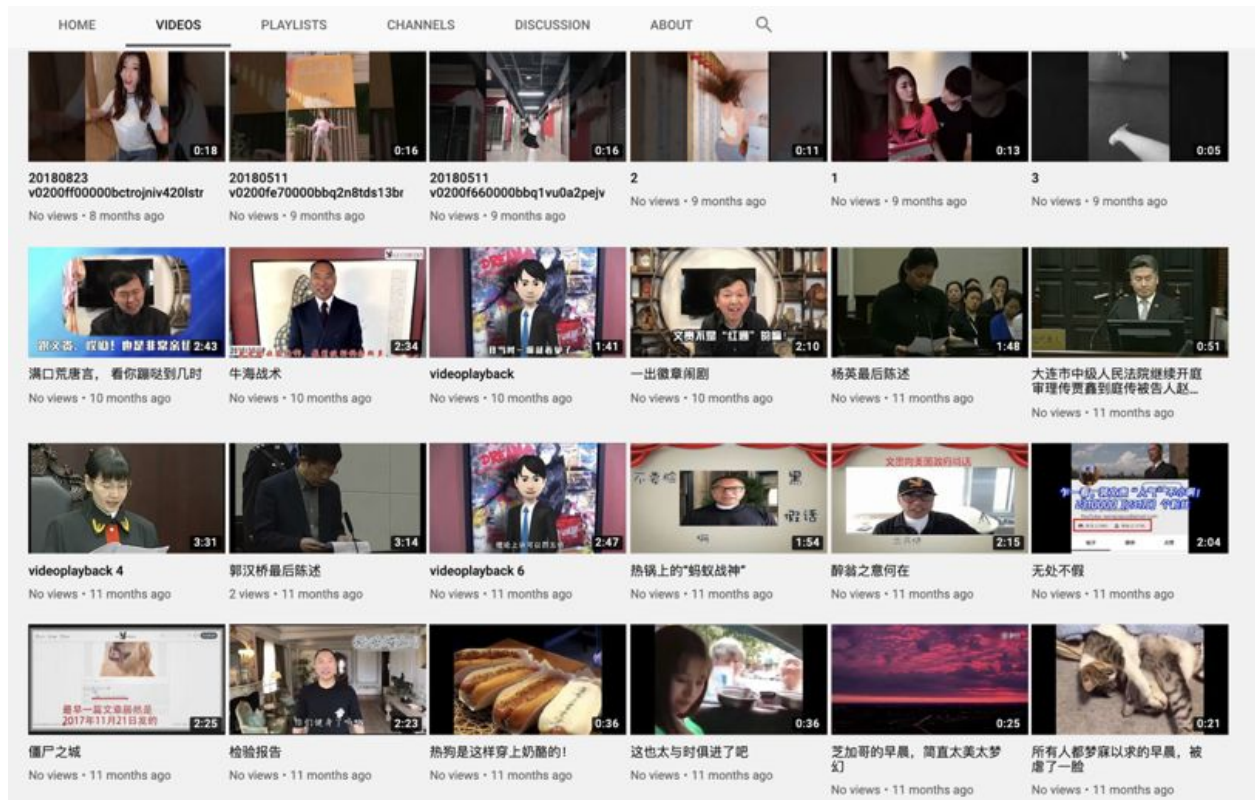
Screenshot of YouTube [search](#) results for *满口荒唐言，看你蹦哒到几时？*, filtered to show videos only. Note the length of the last four clips, 0:02.

The four accounts that only posted short clips were obvious spam assets, most likely automated to replicate videos from other sources, but so poorly programmed that they cut out after a few seconds. Two repeatedly posted the same videos in the same order, as the following screenshot demonstrates.



Video pages for “[Lew Belchenko](#)” and “[Miroslaw Belotelov](#),” screenshots taken on September 19, 2019. Note the shortness of the clips, the lack of views, and the identical videos in the identical order.

The three channels that posted the full version were marginally more competent, in that they at least managed to post the full version, but they were equally spammy. Each one interspersed videos attacking Guo Wengui, usually 2-3 minutes long, with high numbers of shorter clips, often from TikTok, showing cats, dancing girls, and sports. Two had Chinese usernames, one an apparently Russian username written in the Latin alphabet; all three were created between July and September 2017.



Girls and Guo: videos posted by “[Gennadiy Muzin](#)” showing the mix of TikTok videos (top), political posts about Guo Wengui (middle two rows), and a mixture of spam content (bottom).

Of the three, only one, [封易 \(Feng Yi\)](#), was still active in September 2019, posting cat videos. The others fell silent in April. These had every appearance of commercial spam accounts, bought or co-opted in late 2018 to amplify attacks on Guo, and then returned to spam posting, perhaps in a bid to fool YouTube’s detection algorithms.

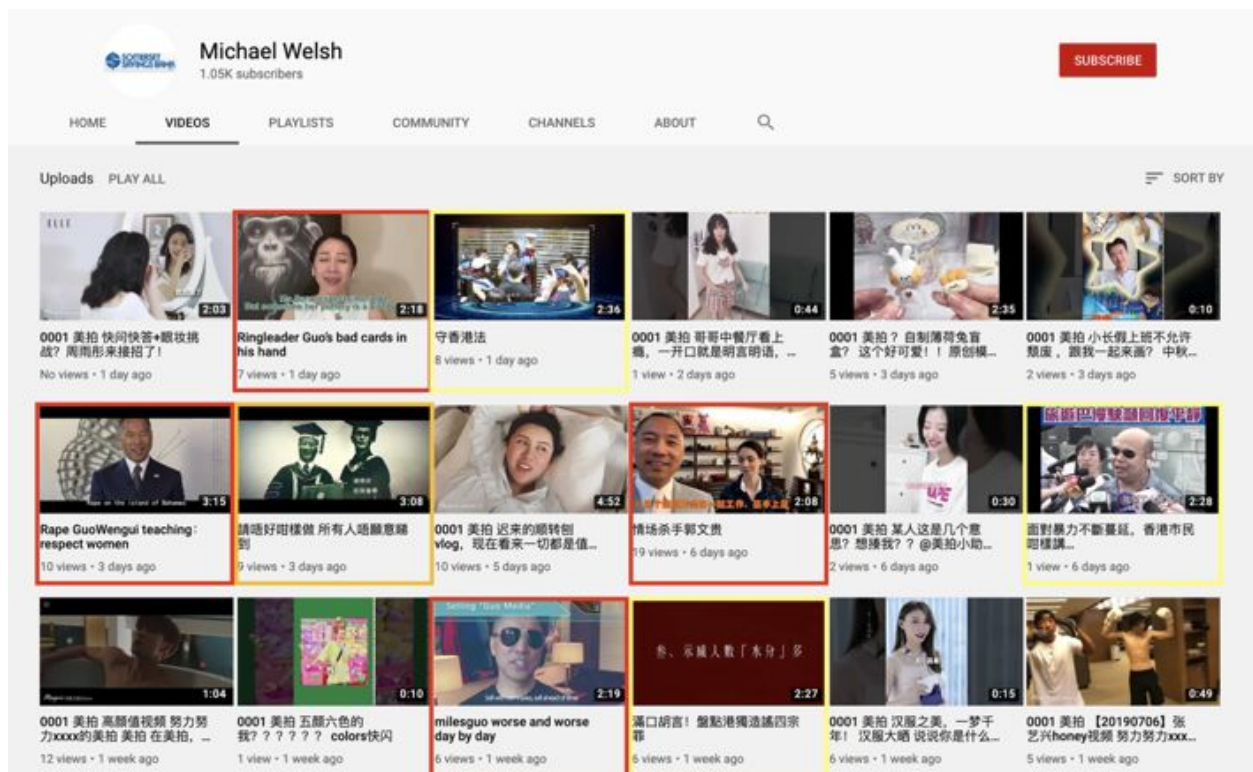
None of these accounts performed well in terms of views: few of their posts scored any views at all. This combination of spam, political messaging, and low impact was typical of the assets in the original Twitter dataset and also characterized the broader network.

Spam and Hijacking on YouTube

Having identified the spam amplification of one video by Rumor Shredder, Graphika searched for similar amplification of other videos by searching for Rumor Shredder headlines. Repeatedly, these searches led to more spamouflage channels where political messaging was mixed with clickbait. Much of the political messaging started in June 2019, when the Hong Kong protests began. It combined criticism of Guo and criticism of the Hong Kong protesters with praise for the Hong Kong police.

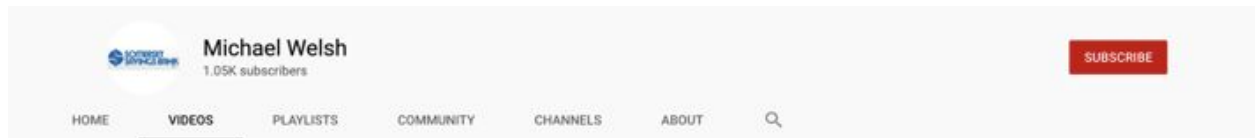
Some of the channels had very high viewing scores, in the hundreds of thousands. On examination, these turned out to be the result of earlier posts – often years earlier – in different languages and on wholly unrelated themes. This suggests that the channels were originally run to build an audience and ultimately came into the possession of the Spamouflage Dragon operators, who used them to promote anti-protest messaging. Such repurposing is characteristic of hijacked or otherwise compromised assets that are sold on the black market, rather than dedicated assets created and run by a single operation; it indicates that Spamouflage Dragon prioritized the quick acquisition of assets over credible appearance.

A channel called Michael Welsh, for example, interspersed Rumor Shredder attacks on Guo and criticism of the protesters with entertainment content. The account's handle was [Somerset Savings](#) and pointed to a bank in the United States.





















[Most recent posts](#) by "Michael Welsh." Attacks on Guo are marked in red, criticism of the protesters is marked in yellow. A video attacking both is marked in orange.

The channel boasted over 796,000 views since it was created in September 2012, but almost all of them came from episodes of the 1970s UK comedy "The Good Life."


 Michael Welsh
 1.05K subscribers

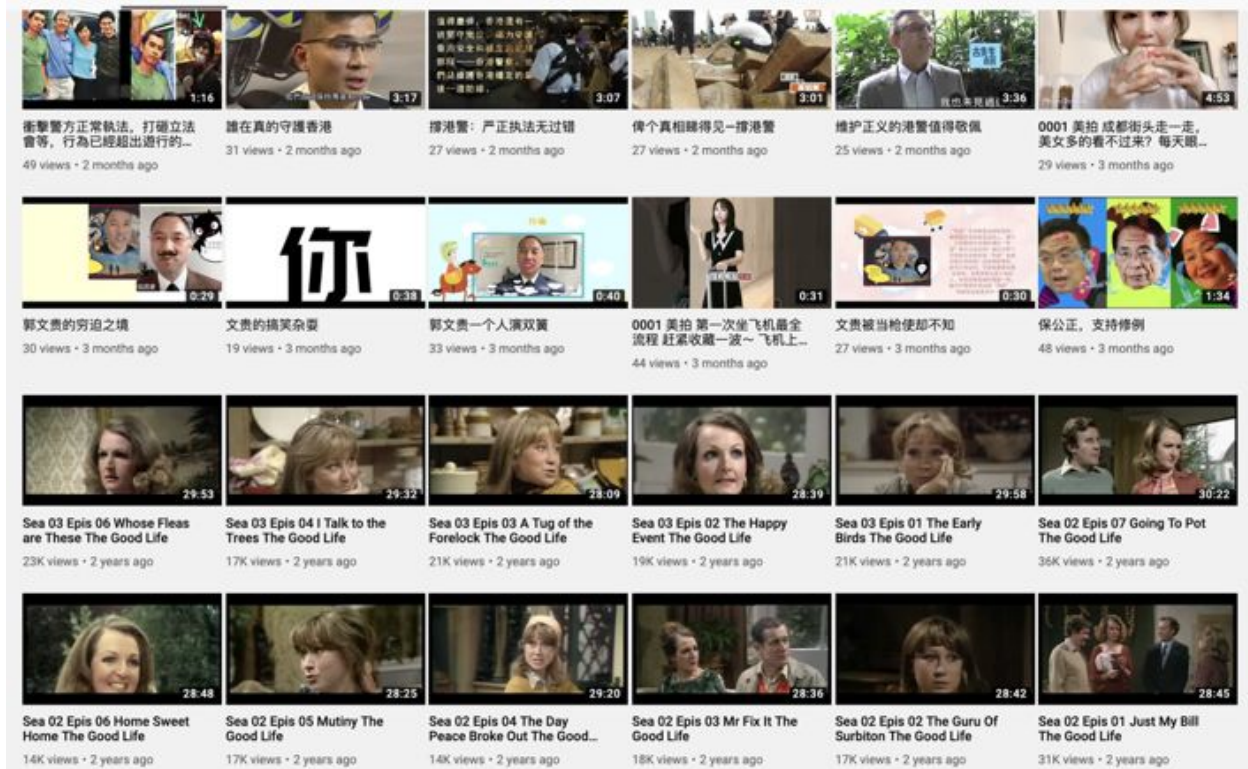
HOME VIDEOS PLAYLISTS COMMUNITY CHANNELS ABOUT

Uploads PLAY ALL SORT BY

 Sea 01 Epis 01 Plough Your Own Furrow The Good Life 109K views • 2 years ago	 Sea 01 Epis 02 Say Little Hen The Good Life 53K views • 2 years ago	 Sea 01 Epis 03 The Weaker Sex The Good Life 49K views • 2 years ago	 Sea 02 Epis 07 Going To Pot The Good Life 36K views • 2 years ago	 Sea 02 Epis 01 Just My Bill The Good Life 31K views • 2 years ago	 Sea 01 Epis 04 Pig's Lib The Good Life 31K views • 2 years ago
 Sea 01 Epis 06 The Pagan Rite The Good Life 24K views • 2 years ago	 Sea 03 Epis 06 Whose Fleas are These The Good Life 23K views • 2 years ago	 Sea 03 Epis 01 The Early Birds The Good Life 21K views • 2 years ago	 Sea 03 Epis 03 A Tug of the Forelock The Good Life 21K views • 2 years ago	 Sea 03 Epis 02 The Happy Event The Good Life 19K views • 2 years ago	 Sea 02 Epis 03 Mr Fix It The Good Life 18K views • 2 years ago
 Sea 02 Epis 05 Mutiny The Good Life 17K views • 2 years ago	 Sea 03 Epis 04 I Talk to the Trees The Good Life 17K views • 2 years ago	 Sea 02 Epis 02 The Guru Of Surbiton The Good Life 17K views • 2 years ago	 Sea 04 Epis 06 Sweet and Sour Charity The Good Life 14K views • 2 years ago	 Sea 02 Epis 06 Home Sweet Home The Good Life 14K views • 2 years ago	 Sea 04 Epis 05 Suit Yourself The Good Life 14K views • 2 years ago

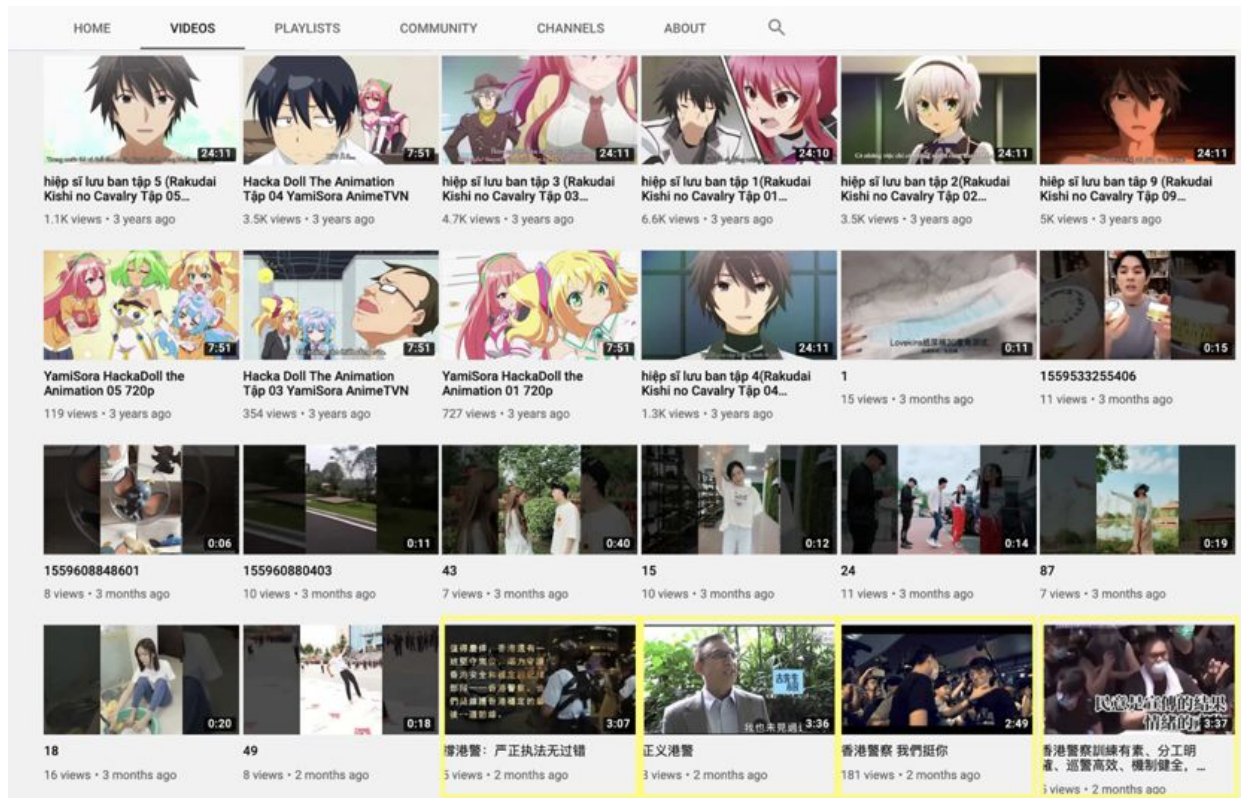
Videos posted by "Michael Welsh," ordered by [popularity](#).

The last time this channel posted a video from "The Good Life" was in September 2016; until then, the channel averaged around 20,000 views per video. The channel then fell silent for more than two years before returning to activity on June 9, 2019, with a video calling on the people of Hong Kong to support the controversial extradition law that triggered the protests. That video scored 45 views; subsequent posts averaged around 30 views and interspersed political content with TikTok and other videos.



The moment of change: videos from "The Good Life" over two years ago give way to posts about Hong Kong politics, and spam.

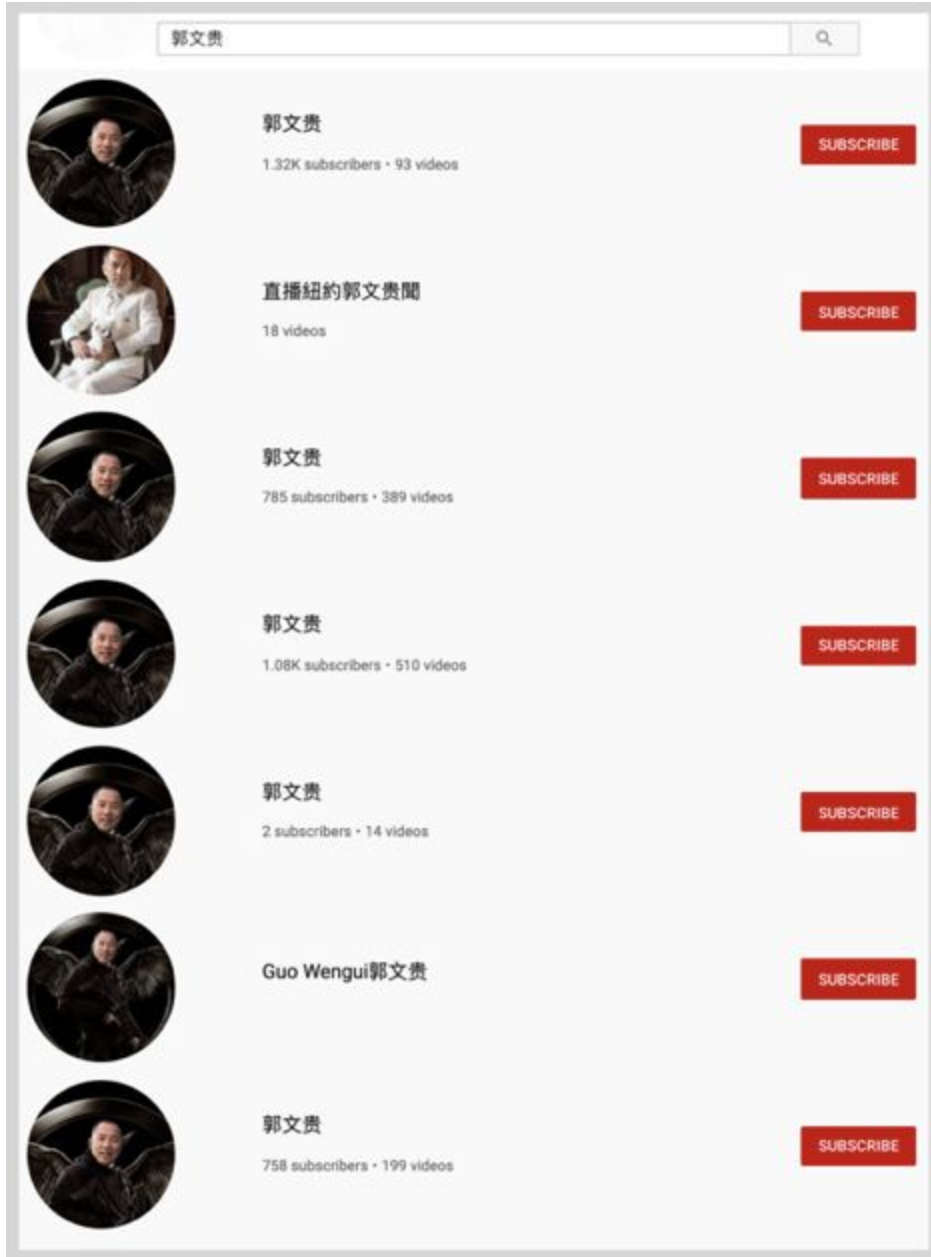
Another channel, called [Miss Aiello](#), shared a separate Rumor Shredder video and boasted even more impressive viewing figures, at 1.6 million. This time, all its [most popular views](#) came from animé cartoons posted on April 19, 2016; it then fell silent until June 2, 2019, when it began posting TikTok videos, followed by criticism of the Hong Kong protests. These seldom scored more than a dozen views each.



Videos posted by "Miss Aiello," ordered from the oldest at the top. Yellow boxes mark posts criticizing the Hong Kong protests.

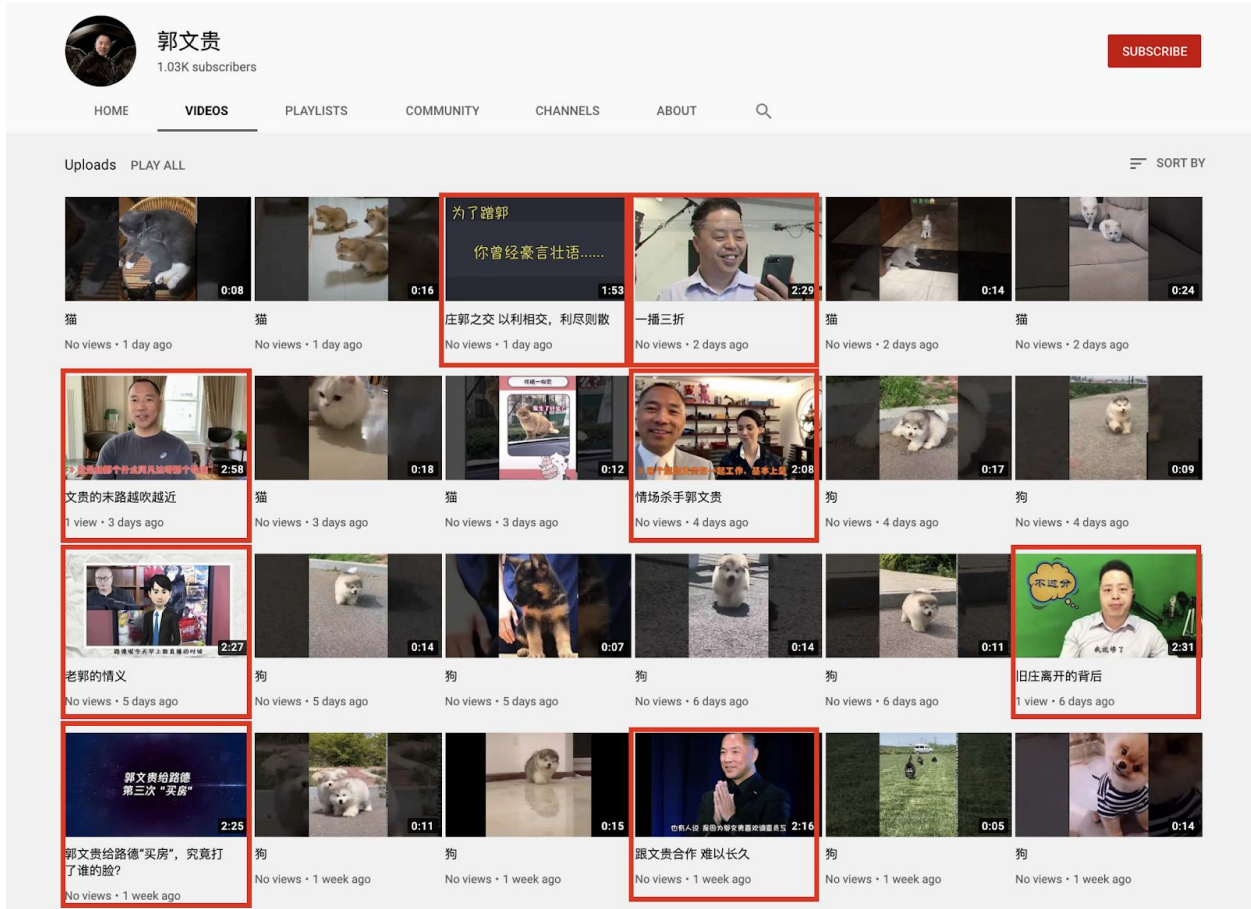
The June timeline, two months after the Hong Kong protests started, is significant; a number of the assets that Graphika reviewed across platforms began posting around that time. Many appeared to have been repurposed from abandoned accounts, either sold on or hijacked, and very few gained more than a handful of views, while a high proportion scored none at all. These factors suggest a hasty operation, gathering assets on the online black market and running them for video amplification without first cleaning out their earlier posting history. This is not the hallmark of a deliberate or long-planned operation; rather, it appears to be characteristic of a rapid and unskilled reaction to unexpected events.

A separate set of assets posted similar content, including videos from Rumor Shredder, but focused much more heavily on Guo Wengui, a controversial Chinese billionaire living in New York. Graphika identified over a dozen YouTube channels that used Guo's name and the same profile picture as his official channel, some of them dating back to early 2017. Visually, the only way to tell these channels apart from the original, and from one another, was by their subscriber numbers, which ranged from 0 to 2,000. (The genuine article has 273,000 subscribers.)



Some of the YouTube channels that used Guo Wengui's name and profile picture.

Not all these channels appeared to be part of the same online effort. Two had been inactive for over a year, and another had not posted at all. Others closely resembled the activity described above, interspersing political videos, including Rumor Shredder, with cats and other clickbait. Two posted identical content in almost identical order.



Kittens and criticisms: videos posted by one of the [Guo Wengui](#) channels, showing the blend of political (red boxes) and cat content. Note the very low viewing figures.

The activity on YouTube centered around Rumor Shredder, but there is insufficient evidence to show whether it was controlled by the Rumor Shredder operators or unaffiliated supporters trying to amplify the channel.

What is clear is that a group of apparently fake or hijacked accounts posted attacks on Guo mixed with random spamouflage, often taken from TikTok, and pivoted to attack the Hong Kong protesters as the movement gained momentum. Very few of the videos gained any degree of traction, and their viewing figures were extremely low.

YouTube Assets Lead Back to Twitter

Having identified Rumor Shredder as a key node in the YouTube network, Graphika then searched for links to the individual videos on Twitter. This led to a series of clustered accounts that behaved very similarly to the YouTube assets, primarily tweeting spam or inspirational quotes alongside a small number of political tweets with a focus on anti-Guo content.

For example, the URL of the original Rumor Shredder video described above was tweeted by two accounts on October 22 and October 24, 2019. One, [@KimberlyBlackbx](#), gave a UK location and used a drawing of a young Western woman for its profile picture. It was created in 2013, but its earliest visible post dated to September 2018, when it began posting Chinese-language attacks on Guo, including YouTube videos, mixing them in with music videos and TikTok clips. It fell silent in January 2019. This was very similar behavior to the YouTube assets and is likely to indicate an account that was hijacked and repurposed.

Another account that shared content from the network, [@Rulai_RPG](#), was followed by 3,570 other accounts, some of them with strong signs of inauthenticity. The account creation dates for this first cluster centered around September 2017, with a few exceptions having been created in 2012 and 2013. Accounts from this cluster tended to use stolen profile pictures, often of young women. As of September 2019, the followers were largely inactive but remained visible online, without posting.



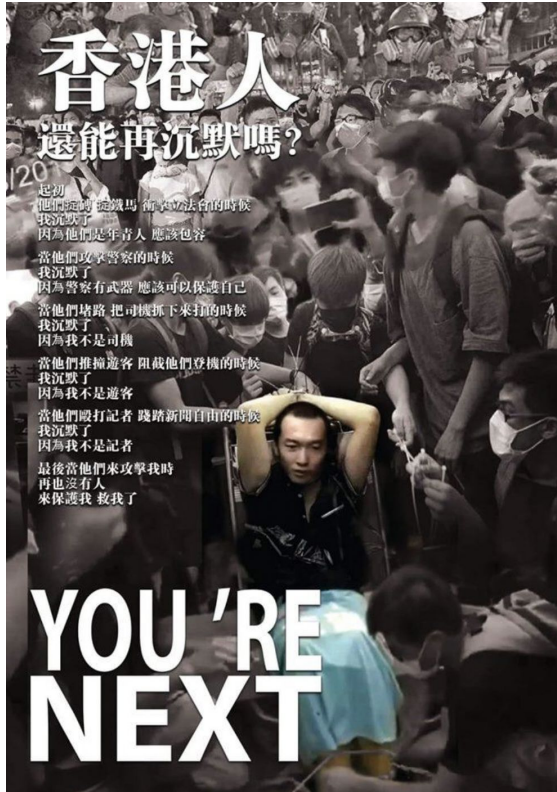
Examples of Accounts and Content Shared from the @Rulai_RPG cluster

A second cluster centered on an account called King (7) [@King767747](#), which was created in March 2019. This account was still active as of September and often posted primarily about food and cooking, but then veered into anti-Guo and anti-protester posts. A second account, Lover (8) or [@Lover7677477](#), was also created in March 2019 and shared the same anti-Guo content while offering travel advice about destinations in China. The third asset identified in this cluster is Ayush (9), or [@Ayush914](#). Like its companions, Ayush's account was created in March 2019 and discussed food, animals, and travel alongside anti-Guo/anti-protester content. Notably, all of these accounts had a bio reading "I love you," or in King (7)'s case, "I love you my India," and the people pictured in Lover (8) and Ayush (9) both appeared to be of South Asian descent.

Graphika identified these assets in late August. Twitter took them down sometime in September for unknown reasons.



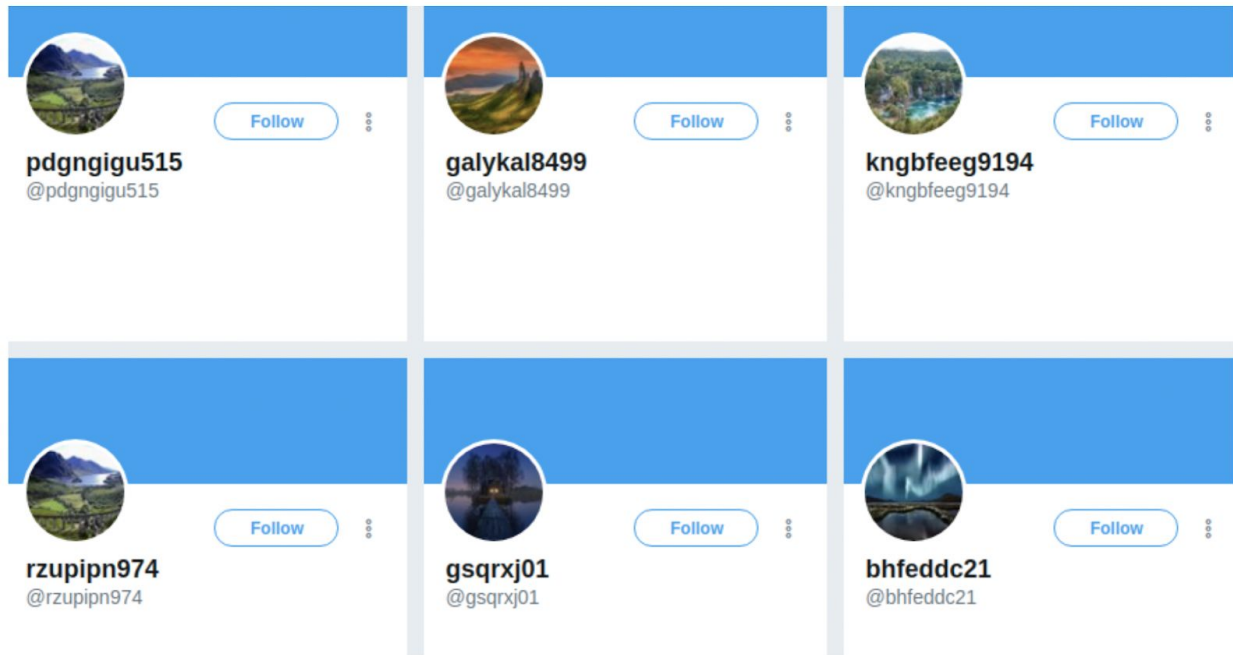
Profile Bios and Pictures of the King (7) Cluster



Sample content from King (7) Cluster

Perhaps the most interesting cluster was centered around an account calling itself Microview Review, or @lengjingwatch, also suspended in September. This account, created August 2019, tweeted 117 times and had 262 followers. Unlike most of the accounts observed in this network, Microview Review seemingly exclusively tweeted content condemning the Hong Kong protestors and only retweeted itself. Its followers largely appeared to be created for this purpose; a large majority were created between June and August, and they tend to use generic nature stock photos as profile pictures, sometimes the exact same image (see example below). Creating a pool of followers gives an account the appearance of legitimacy, but it does not translate into the ability to reach genuine followers: it is best viewed as a camouflage tactic rather than an engagement one.

The follower accounts showed the same spamouflage behavior as the other identified assets; some retweeted various news sources, often in multiple languages, and others mixed modeling/nude photos into their activity. The percentage of retweets from Microview Review was typically very high. The content from Microview focused on discrediting the Hong Kong government, leaders of the protests, and media figures reporting on the events, much of it with an anti-US angle. Many of the tweets reflected a belief that America was providing funding for the protests. Anti-Guo content was less obvious in this cluster, but the overlap between this and the content being shared on other platforms, in addition to the tactic of hiding its traffic within large amounts of spam posts and overall messaging, indicate its relation to the larger overall network.



Examples of Followers for @lengjingwatch, with Profile Picture Reuse Visible

Spamouflage on Facebook

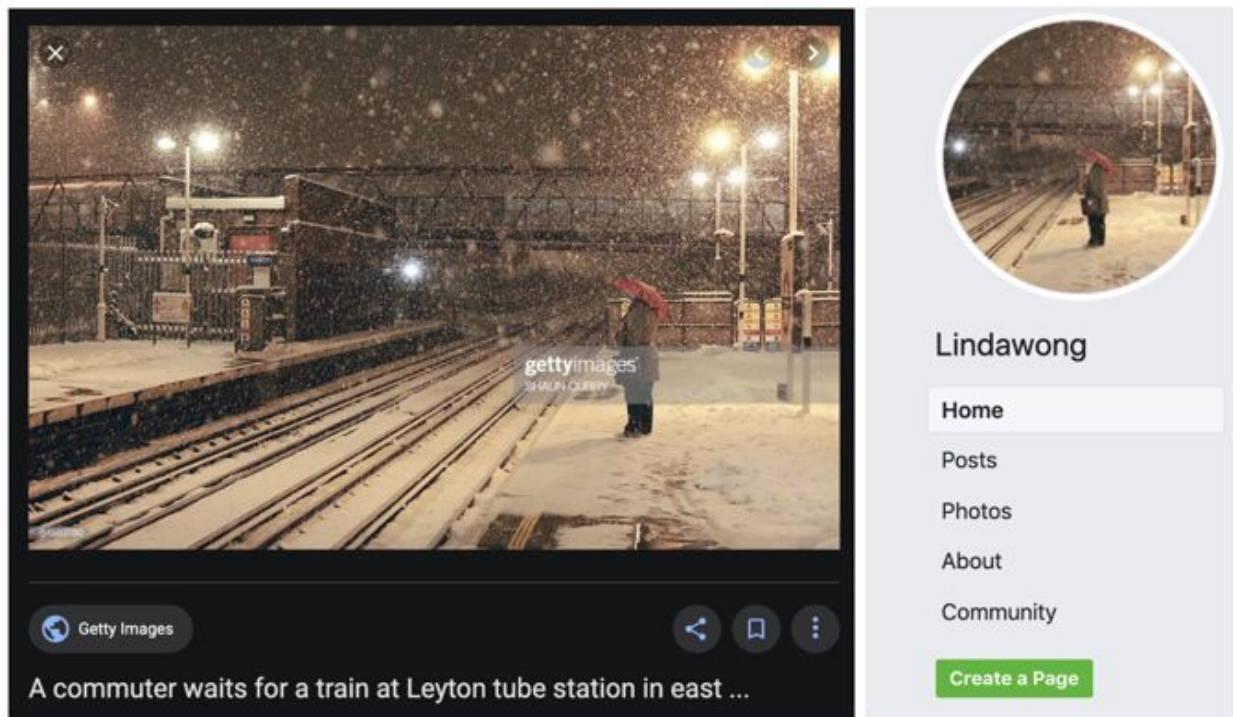
Having identified the spamouflage network on YouTube, and found further assets on Twitter, Graphika searched for similar behaviour patterns on Facebook. The first step was to search for the URL of the Rumor Shredder video that linked the channel's old and new incarnations. This was shared by two assets on Facebook called [Lindawong](#) and [Zhoubian \(周边\)](#).

Despite looking like personal accounts, with a person's name and profile picture, each was actually a page. Pages on Facebook have lower verification requirements than personal accounts, and information operations have begun using pages to masquerade as human users to avoid detection; in August, Facebook took down over 1,500 such pages linked to the government of Honduras.



Screenshot of the Lindawong [page](#), showing the profile picture and follower numbers. Screenshot taken on September 19, 2019. A screenshot of the page's share of the original Rumor Shredder video is also shown.

Lindawong had over 2,500 followers. It mixed spam content with political posts supporting the Hong Kong police, criticizing the protesters, and attacking Guo Wengui; when it was taken down in September, its three latest posts were a video recipe for onion fish fillet, a recipe for steamed tofu rolls, and a video praising the Hong Kong police. The page's profile picture was copied from a Getty Images photo of snow in London.



Left, the Getty Images original, from a Google [search](#). Right, the Lindawong page.

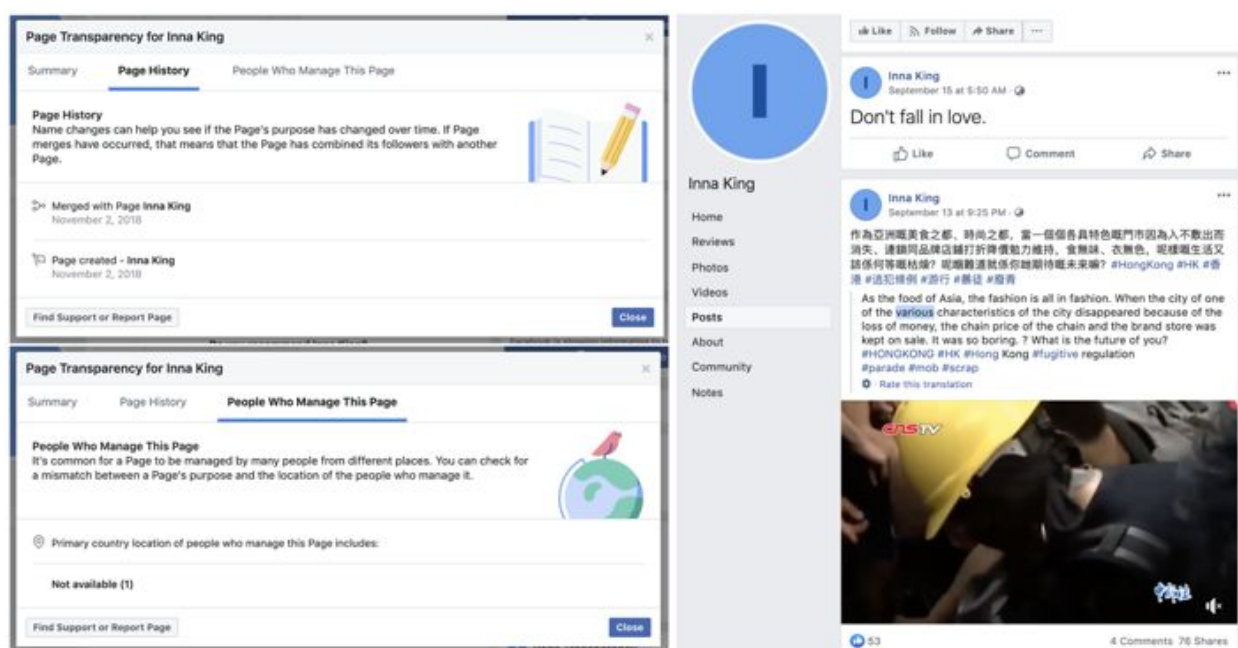
Zhoubian only had one follower. Its profile picture was taken from South Korean actress [Kim Yoo-Jung](#). The page almost exclusively posted criticism of Guo or the Hong Kong protesters, including at least one Rumor Shredder [video](#) from September. Unlike Lindawong, most of its posts were not YouTube links but shares of Chinese-language videos and memes from other Facebook assets, a high proportion of which had non-Chinese names. When Graphika examined these assets, several proved to be accounts attributed to young men in Bangladesh that switched to posting in Chinese early in the year. A friend of one such account was startled enough by the move to ask his friend, in Bengali, what he said. Graphika concluded that both accounts were created by genuine Bangladeshi users and then either hijacked or sold on, ending up in the hands of the Spamoouflage Dragon network.



Transformations: screenshots from the Bangladeshi accounts' timelines: one showing the shift from Bengali to Chinese, and another showing the question from a Bengali friend. Graphika has chosen to obscure all user names and profile images.

Both accounts interspersed political posts about the Hong Kong protesters and Guo Wengui with short Chinese textual posts of apparent wisdom, such as “For Life, acceptance is the best tenderness, whether it is to accept the appearance of a person, or the acceptance of a person has never seen it.” This was another type of spamouflage, interspersing political content with unrelated texts.

[Inna King](#) proved to be a page whose manager had hidden their location, and which posted a mixture of uplifting sayings, life advice (some in English), and political videos attacking the protesters. Curiously, on the day it was created, November 2, 2018, it was merged with another page of the same name, perhaps to hide its earlier posting history, or to increase its apparent following. This pattern recurred frequently across the network.



Inna nutshell: screenshot of the Facebook Page Transparency settings for Inna King, showing the history and manager location. On the right, two posts from the page's timeline, one a life lesson, the other attacking the Hong Kong protesters.

As Graphika examined further pages and accounts that Zhoubian shared, several patterns emerged. Some accounts appeared to originate from Bangladesh but switched to Chinese posting in March-April 2019; one originally Japanese-language account and several Russian-language ones did the same. Some accounts had Western names or profile pictures but appeared to be clumsily automated, posting the same content (a mixture of spam and politics) in the same order.



Joined at the hip: timelines for the accounts Helena Hart and Regina Montgomery, showing identical posts in the identical order.

Many accounts had Chinese names and East Asian profile pictures, often stock photos taken from glamor shoots. Some had obviously made-up names, such as Banana, Abcytgf, and Coffee. All posted a mixture of spam and political content, frequently including Rumor Shredder videos. They often liked, shared, or commented on each other's posts, and small groups of assets shared the same posts in the same order.

Graphika identified over 120 assets, evenly split between accounts and pages, that behaved in the same way and shared subsets of the same political content. While the number of assets was significant, virtually no assets outside the spamouflage network reacted to, shared, or commented on any of the posts outside of its network. These assets constituted a closed network, interacting with one another but not reaching genuine users outside the circle. They did not appear to reach out to other users (for example by tagging them), organize events, or run ads. As such, their behavior was more characteristic of spam operations that boost the numbers on individual posts without reaching genuine users than of sophisticated information and influence operations.

Open-source research cannot estimate the impact of this activity on recommendation systems on Facebook and across platforms, but the network's reach among unaffiliated users is likely to have been marginal at best.

Conclusions

The network described here centered on YouTube and largely existed to promote video content: many of the posts on Twitter and Facebook led back to YouTube videos. Earlier posts focused their criticism on Guo Wengui, especially in late 2018 and early 2019, but progressively from April 2019, they added criticism of the Hong Kong protests to the mix.

Many of the assets appeared to have been taken over from genuine users and repurposed, no doubt in an attempt to evade the platforms' systems for identifying fake accounts. Graphika identified a large number of apparently Bangladeshi Facebook accounts, alongside smaller numbers of Russian, Japanese, and English assets. Other assets appeared custom-made but in relatively small numbers, again presumably to avoid detection.

The network was active and prolific but seemed to have had low impact. Many of the YouTube videos had no views at all; only a few scored views into the hundreds. The Facebook and Twitter assets appeared to be locked in closed circles in which they all reacted to each other's posts but did not manage to reach genuine users outside. They did not appear to practice more aggressive outreach in the manner of other known information operations, such as tagging real users or promoting their posts.

The network was based on a core of assets dedicated to promoting a YouTube channel "Rumor Shredder" and then expanded to take on the Hong Kong protests. The expansion appeared to be hurried and only marginally effective, bringing in large volumes of spam-like assets to boost content without any apparent broader engagement strategy. This suggests that it was not a state-backed operation, such as Twitter and Facebook exposed in August, but a smaller and less professional network, run by unknown individuals to boost anti-Guo and anti-protest messaging.

It was also, most likely, cheap. Other than the investment required to produce the videos in the first place, the network relied on the sort of cheap spam assets that are readily available on the black market. This underlines the continuity between professional information operations and commercial spam posting and one way in which threat actors can leverage apolitical assets for political ends.