



**Graphika**

# Spamouflage Dragon Goes to America

Pro-Chinese Inauthentic Network Debuts  
English-Language Videos

Ben Nimmo, Camille François,  
C. Shawn Eib and Léa Ronzaud, Graphika

**08.2020**

**Information Operations**

# Spamouflage Goes to America

## Pro-Chinese Inauthentic Network Debuts English-Language Videos

Ben Nimmo, Camille François, C. Shawn Eib, and Léa Ronzaud, Graphika

### Introduction

Social media accounts from the pro-Chinese political spam network [Spamouflage Dragon](#) started posting English-language videos that attacked American policy and the administration of U.S. President Donald Trump in June, as the rhetorical confrontation between the United States and China escalated.

The videos were clumsily made, marked by language errors and awkward automated voice-overs. Some of the accounts on YouTube and Twitter used AI-generated profile pictures, a technique that appears to be increasingly common in disinformation campaigns.<sup>1</sup> The network did not appear to receive any engagement from authentic users across social media platforms, nor did it appear to seriously attempt to conceal its Chinese origin as it pivoted toward messaging related to U.S. politics.

Spamouflage Dragon's politically focused disinformation campaigns appear to have started in the summer of 2019. It [began in Chinese](#) by attacking the Hong Kong protesters and exiled Chinese billionaire Guo Wengui, a frequent critic of the Chinese Communist Party (CCP). In early 2020, it started [commenting on the coronavirus pandemic](#), praising the CCP's response at a time when it was being accused of covering up the outbreak.

The latest wave of Spamouflage activity differs in two key ways from its predecessors. First, it includes a wealth of videos in English and targets the United States, especially its foreign policy,

---

<sup>1</sup> Graphika first observed this at scale in an unrelated network at the end of 2019; we dubbed this operation [FFS: Fake Face Swarm](#).

its handling of the coronavirus outbreak, its racial inequalities, and its moves against TikTok.<sup>2</sup> This is the first time the network has published substantial volumes of English-language content alongside its ongoing Chinese coverage - a clear expansion of its scope. The network was particularly active, and reactive to current events, in the period of investigation: videos commenting on recent U.S. official statements were created and uploaded in less than 36 hours.

Second, it is the first time that we have seen Spamouflage Dragon use clusters of accounts with AI-generated profile pictures. Other operations are known to have done so,<sup>3</sup> but this is the first time the practice has been adopted by this particular network. Given the ease with which threat actors can now use publicly available services to generate fake profile pictures, this tactic is likely to become increasingly prevalent.

## Spamouflage Again

Graphika uncovered Spamouflage Dragon in [September 2019](#) using signals found in the [Twitter dataset](#), at a time when the network was primarily concerned with attacking two subjects: the Hong Kong protests and exiled Chinese billionaire Guo Wengui, a critic of the Chinese Communist Party. Twitter attributed the original activity on which the dataset is based to operators linked to Chinese state actors, although Graphika has not to date been able to establish a connection between Spamouflage Dragon and government entities.

The network was active and prolific, but ultimately low-engagement. It typically worked by using apparently hijacked or otherwise repurposed accounts on YouTube, Facebook, and Twitter to post politically charged videos in Chinese and then used clusters of fake accounts to share and comment on the posts, creating the impression of an organic community. These accounts all posted in Chinese but often featured names in other languages, notably Russian, Bengali, and English.

Some of these amplifiers interspersed their political posts with large quantities of spammy posts about scenery, poetry, and sports, often taken from TikTok, suggesting that they were either commercial amplifiers rented out for the occasion or using the spam as camouflage. This characteristic is why Graphika named the network Spamouflage Dragon.

After analysis and disclosures by Graphika, Twitter, Facebook, and YouTube disrupted the network with a series of takedowns in August-September 2019, driving it into hiding. After China

---

<sup>2</sup> The set of themes we observed in this network's content echo the recent U.S. Office of the Director of National Intelligence (ODNI) assessment of China's public rhetoric ahead of the U.S. 2020 election: "[China's] public rhetoric over the past few months has grown increasingly critical of the current Administration's COVID-19 response, closure of China's Houston Consulate, and actions on other issues. For example, it has harshly criticized the Administration's statements and actions on Hong Kong, TikTok, the legal status of the South China Sea, and China's efforts to dominate the 5G market"; see August 7, 2020 <https://www.dni.gov/index.php/newsroom/press-releases/item/2139-statement-by-ncsc-director-william-e-vanina-election-threat-update-for-the-american-public>.

<sup>3</sup> See, for example, the [network of fake personas](#) exposed by Adam Rawnsley and Marc Owen Jones.

came under sustained international criticism for its handling of the Covid-19 outbreak in early 2020, Spamouflage Dragon re-emerged, posting videos and memes in Chinese that praised China's response to the virus and denied that it had mishandled its response at the start of the outbreak. Graphika exposed this activity in [April](#), engaging affected platforms in a follow-up wave of enforcement actions.

The latest burst of English-language activity came from accounts that closely matched the earlier behavioral signals.

## English-Language Posts

The key change observed in Spamouflage Dragon's activity since April was its shift to posting videos in English that attacked the United States, initially over its response to the coronavirus pandemic and subsequently over the U.S. closure of China's consulate in Houston, Texas, the Trump administration's crackdown on TikTok, and the administration's expansion of the "clean network" concept.

The earliest case that Graphika has identified was on June 13; after an initial period of sporadic video production, the tempo accelerated sharply in the second half of July. As of August 1, Graphika had identified a resurgence of this network, yielding more than two dozen videos on YouTube in English that criticized various aspects of U.S. policy. YouTube has reviewed and terminated all videos and channels cited in this report.<sup>4</sup>

The first [video](#) in this series was titled "#美国 Lighthouse riots escalated, Public blamed Trump's inaction #USA."<sup>5</sup> It focused on the twin crises rocking the United States: the Covid-19 pandemic and the protests that followed the police killing of George Floyd. This focus on events in the U.S., and by extension on "U.S. hypocrisy" in its criticism of China, echoes [Chinese state messaging](#) in recent months.

"At present, the epidemic in the United States has exceeded 2 million. The people's demonstration never stopped. The riots escalated, and the news of 'premeditated' killing of police recently is even more thrilling. Why the U.S. empire that advocates freedom and democracy in the tide of peaceful protest in our country appears to be more and more dry in the out wall," it began in an uneven woman's voice, over a collage of newsreel footage and rock music. The text was subtitled in English and Chinese.

After describing the Black Lives Matter protests and the police response, the video turned its attention to U.S. President Donald Trump:

---

<sup>4</sup> For previous Google announcements related to this network, see the TAG quarterly bulletin, for instance: <https://blog.google/threat-analysis-group/tag-bulletin-q2-2020/>

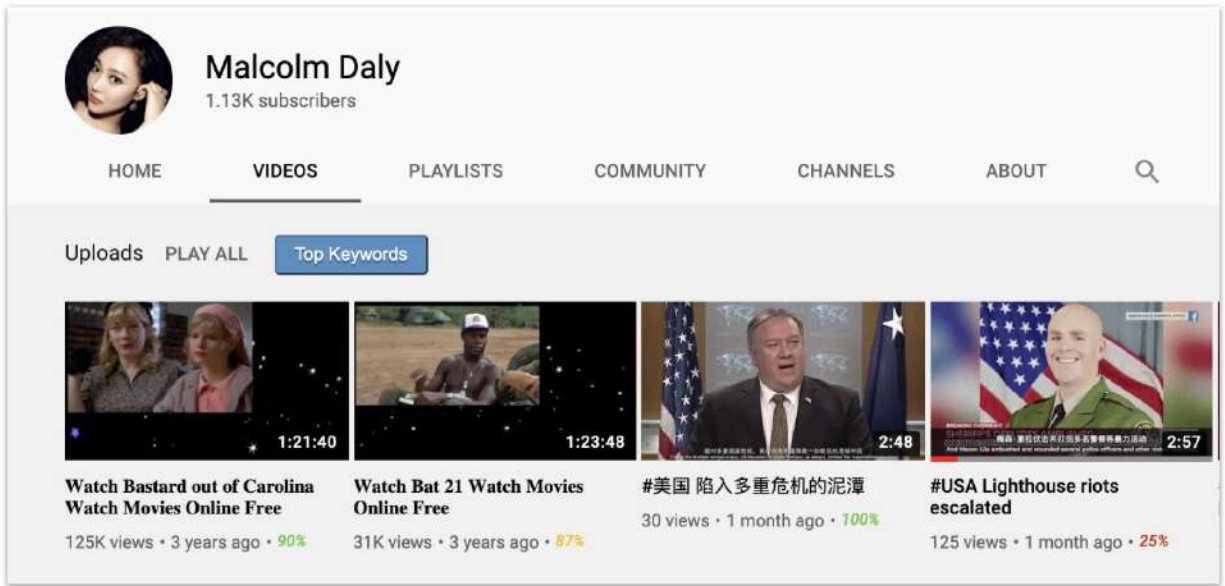
<sup>5</sup> In this and subsequent quotes, all spelling and punctuation, including spacing, will preserve the layout of the original titles and captions.

*But trump is still addicted to his re-election, don't think how to control the epidemic and this riot. Instead, planning to restart the campaign, and even applied "blame pushing" to domestic supporters....Instead of appeasing people's anger, he is devoted to arousing the public anger.*



Still from the video [#美国 Lighthouse riots escalated, Public blamed Trump's inaction #USA](#), timestamp 1:28. Note the bilingual captions.

Graphika has identified [at least three](#) other channels that posted the identical video on June 13-14, including one called [Malcolm Daly](#) that featured a [stock photo of a woman](#) as its profile picture. This account posted two videos - English-language movies - in November 2016, then fell silent until June 2020, when it began posting short videos with Chinese or English headlines, bilingual subtitles, and a focus on Covid-19 and the racial justice protests in the United States. The two movies it posted in 2016 amassed over 150,000 views; the 38 videos it posted in 2020 amassed 472 views between them. This is typical of earlier generations of Spamouflage Dragon and suggests an account that was acquired online and repurposed by the network.



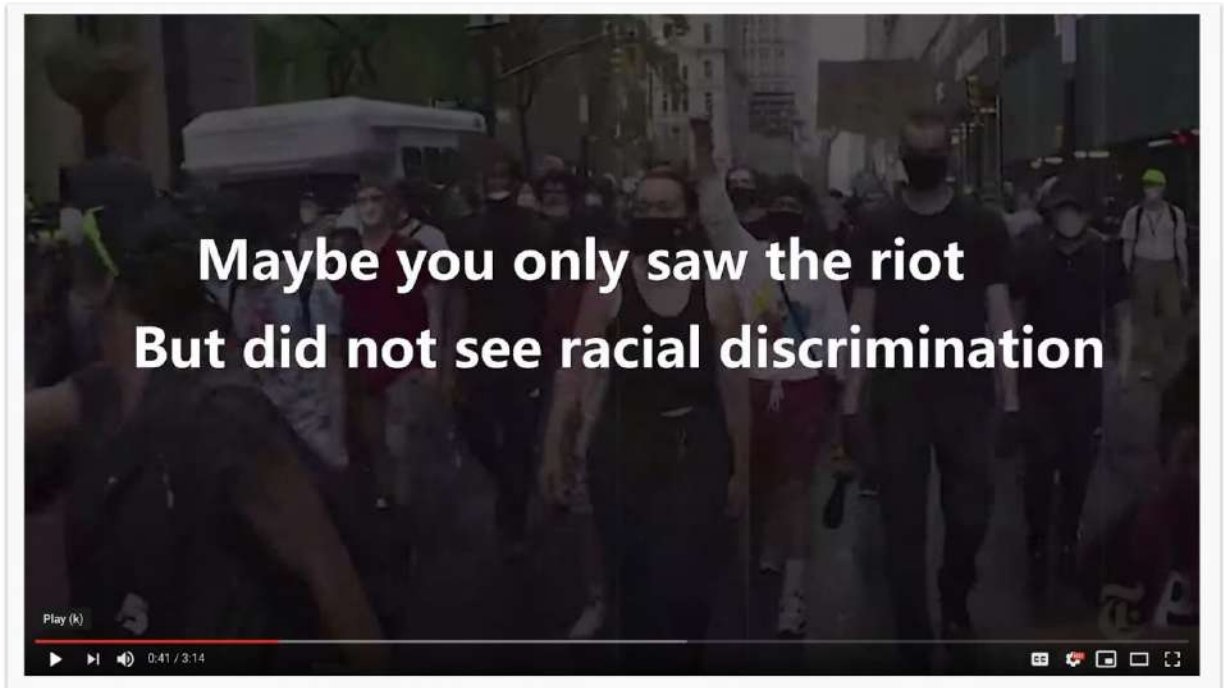
*YouTube profile for Malcolm Daly showing the female profile picture. The videos are arranged in ascending date order and show the two long movies in 2016, followed by the anti-American posts in 2020.*

Further videos on the Malcolm Daly channel followed in the ensuing days, interspersed with short, spammy clips of music, sport, food, and other content. Headlines included [“Who is responsible for this national disaster in the United States”](#) (posted on July 12 and calling Trump a “firm and creative ally” of the virus, it was viewed once); [“The American government ignores science, Politicians turn into’ health experts”](#) (July 12, no views); and [“#USA It is really hard for Americans today #美国”](#) (June 18, 33 views). The tone wavered between scathing criticism of the American government’s handling of the coronavirus outbreak and apparent sympathy for everyday American citizens exposed to the pandemic and the anti-racism protests.

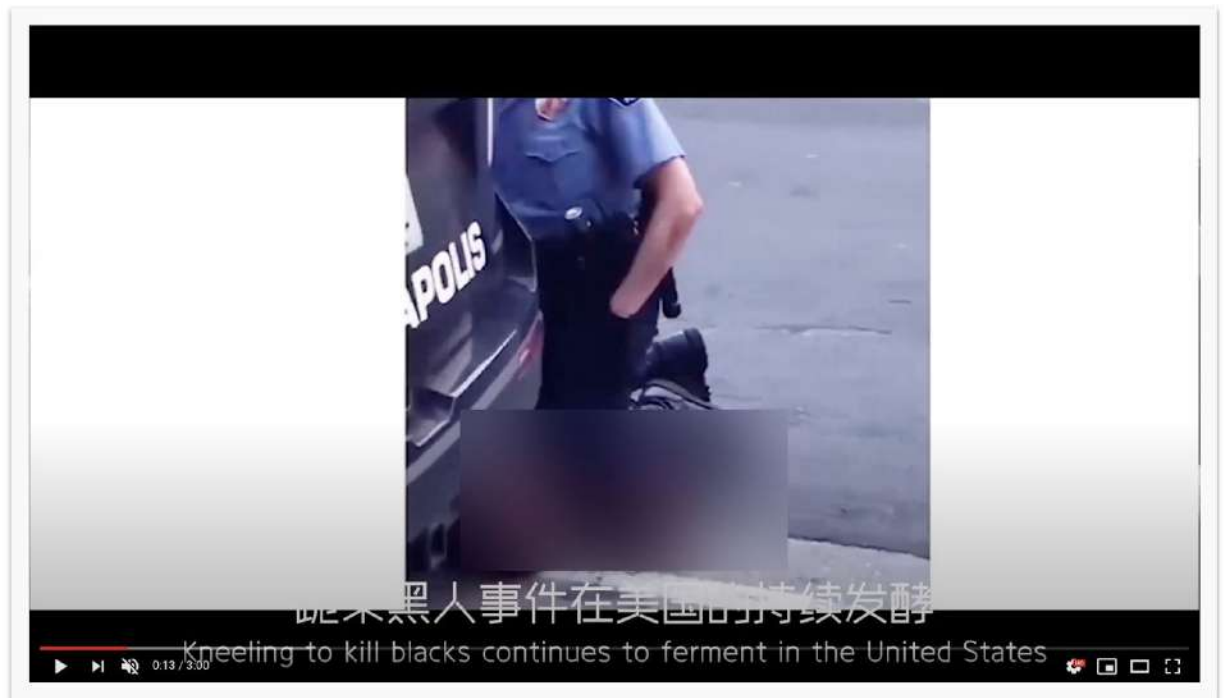


*"Just because the United States doesn't take it seriously." Caption from the video "[#USA It is really hard for Americans today #美国](#)", posted on June 18, 2020, and referring to the pandemic and the racial justice protests.*

The videos themselves appeared to be original productions, based on a collage of news images with a script focusing on the tensions and conflicts within America and the spread of the coronavirus. Especially in late June, some of these videos focused especially heavily on racial tensions in the United States.



Still from the video [“The racial discrimination isn’t solved, the violence will not stop,”](#) posted on June 18, 2020.



Still from the video [当下美国：怎一个“乱”字形容!](#) (“The United States at the moment: how can it be described as ‘chaotic!’”), headlined in Chinese but with English subtitles. Note the claim that “kneeling to kill blacks continues to ferment in the United States.” We have blurred the George Floyd’s face.



Some of these videos focused on individual celebrities and organizations within U.S. society, weaving them into a polarized narrative of racial conflict. One, for example, [argued](#) that “the United States is divided into the rich class dominated by whites and the bottom class dominated by blacks” and went on to accuse former President Barack Obama of failing to address the racial divide or “touch the red line of white interests.”



*Still from the video #美国 暴乱，折射出美国社会的崩塌 (“Riot reflects the collapse of American society”), headlined in Chinese but voiced in English.*

Another [video](#), published on June 25, focused on the Trump administration’s reported criticism of the Centers for Disease Control. Under the headline “#USA Having fought epidemic for half a century, Unexpectedly becoming scapegoat of trump #美国,” the video accused Trump of changing “from indifference to rejection” of the CDC’s expertise because of his own “political ambitions.”



Still from the video "[Having fought epidemic for half a century. Unexpectedly becoming scapegoat of trump](#)," showing its caricature of the President.

As tensions between the United States and China continued to climb in July, the videos took an increasingly hostile tone, with a greater focus on the Trump administration's China policy. Activity spiked after July 21, when the administration ordered the closure of the Chinese consulate in Houston, Texas. Headlines that followed this event and videos that referred to it included "[Does the Us side have any evidence](#)" (July 25, nine views); "[The United States d shamelessly targets China, China's countermeasures should be justified](#)" (July 27, four views); and "[The cold war thinking of US political figures is out of date.](#)" (July 30, three views).



Caption accusing the U.S. government of inciting an "anti-China atmosphere" from video headlined "[The US hatred of China exposed its fears](#)," July 23, 2020.



Caption accusing the U.S. government of "repeated provocations and cries" from video headlined "[Does the Us side have any evidence](#)," July 25, 2020.

On August 4, the network published a video that [criticized Trump's threat to ban TikTok](#) and accused him of "prohibiting the masses from expressing free speech." The video argued that Americans had used TikTok to criticize Trump and that any enforcement action against the platform would result in job and investment losses in America. In typical style, the video appeared on multiple [YouTube channels](#), [Twitter accounts](#), and [Facebook pages](#). In an indication of how prolific the network has become, [three other videos](#) on August 4-5 focused on the same message.



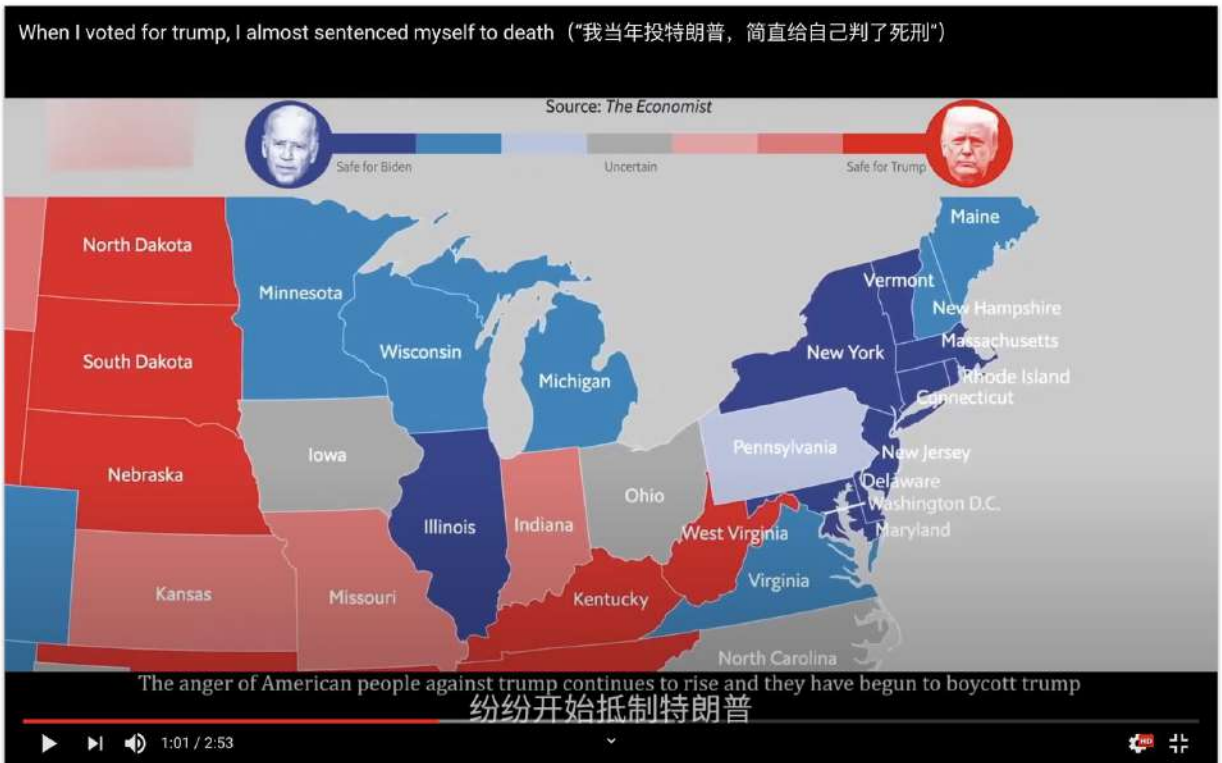
*Still from the Spamouflage Dragon [video](#) about TikTok. The video was headlined in Chinese but voiced in English, with bilingual subtitles.*

On at least one occasion, the network made direct reference to the U.S. election. On August 8-9, dozens of Spamouflage YouTube channels launched a video entitled, "When I voted for trump, I almost sentenced myself to death." Typically for the network, it was voiced in text-to-speech English with Chinese subtitles. At one point, as it featured an American teacher who said that she had voted for Trump and came to regret it, the subtitles continued in Chinese only.



Still from the video, “When I voted for trump, I almost sentenced myself to death,” showing the Chinese-only subtitles at this point in the video. The speaker’s contribution ran, “Just watching the failure of leadership in our country, beginning with the president, over the course of this pandemic, it’s not just my death warrant I might have signed, but there’s 150,000 Americans who are dead.” The clip was taken from a [CNN](#) interview.

The video reported on opinion polls that showed Trump trailing Democratic challenger Joe Biden across a swath of the country, together with a voice-over that spoke of the incumbent’s struggles: “The Trump administration has had the worst of it just before the election.”

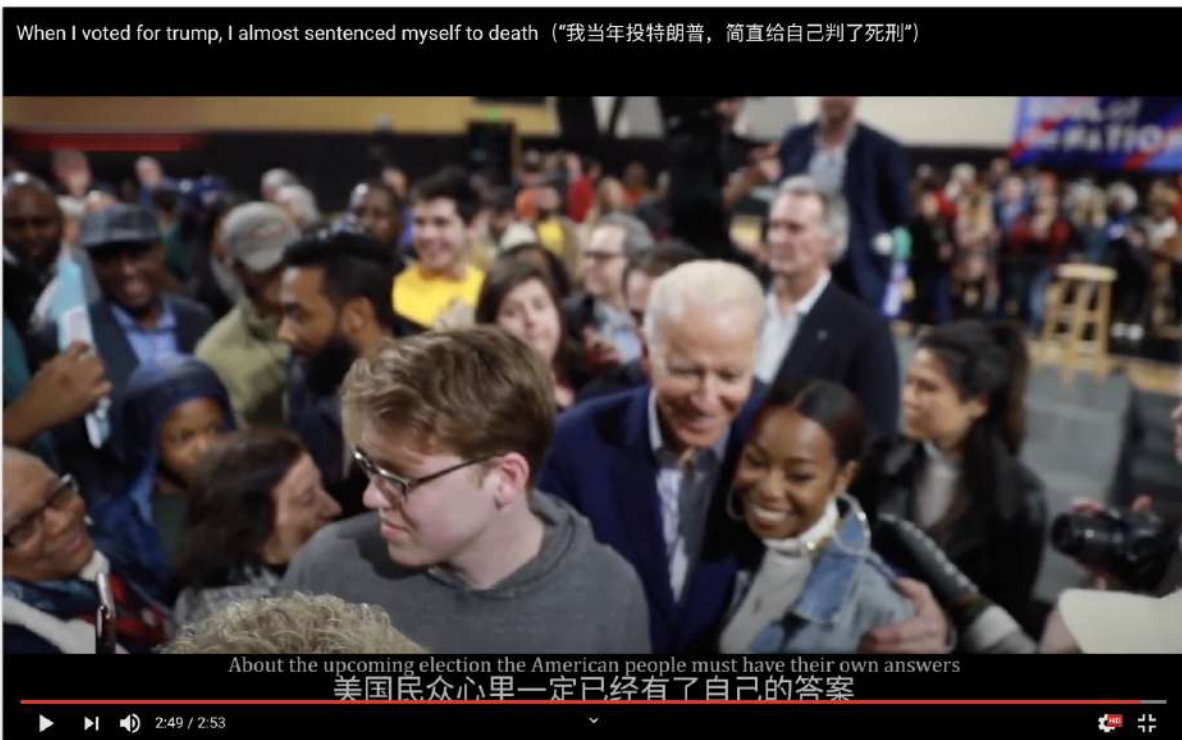


Still from the video, contrasting the polling fortunes of Biden and Trump. Note the caption: “The anger of the American people against trump continues to rise and they have begun to boycott trump.”

For the first time that Graphika has identified on the Spamouflage network, the video referenced Biden directly, showing him in a series of video clips, including over one caption that read, “all the circumstances indicate that [Trump] will not succeed in the election,” and another that read, “About the upcoming election, the American people must have their own answers.”



*Still from the video, showing Biden with the caption, "but all the circumstances indicate that [Trump] will not succeed in the election."*



*Still from the video, showing Biden with the caption, "About the upcoming election, the American people must have their own answers."*

This is the most direct reference to the election that we have seen in this network, but again, couched in a format that made no attempt to hide its Chinese affiliation: between the references to Covid-19 and the election, the video also recited China’s grievances against Trump over issues such as the TikTok ban.

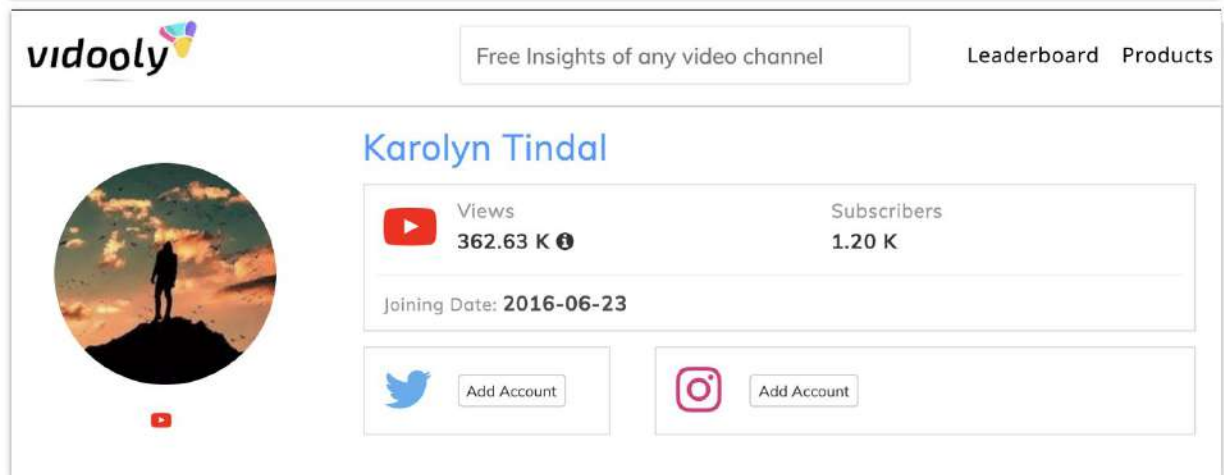


*“He even used his powers as President to ban TIKTOK for no reason.”  
Caption by the Spamouflage network.*

The network appears to have put particular effort into promoting this video: Graphika identified more than two dozen different Spamouflage channels that published it or commented on it. Typically for the network, most of the videos were not viewed at all, or very seldom. However, one post from an account called [风华依旧](#)<sup>6</sup> recorded over 14,000 views, but no comments. The account was created in 2016 and apparently received over 360,000 views in total, but when Graphika accessed it, it only showed five videos with a total of some 26,000 views. According to online tool [vidooly.com](#), the account (listed by its user ID) was originally called “Karolyn Tindal.” These factors suggest that the account was repurposed from a prior user; the viewing figures should be regarded with skepticism.

<sup>6</sup> “still graceful,” user ID UCduP12r1Yebpg0vs3i6lfZg.

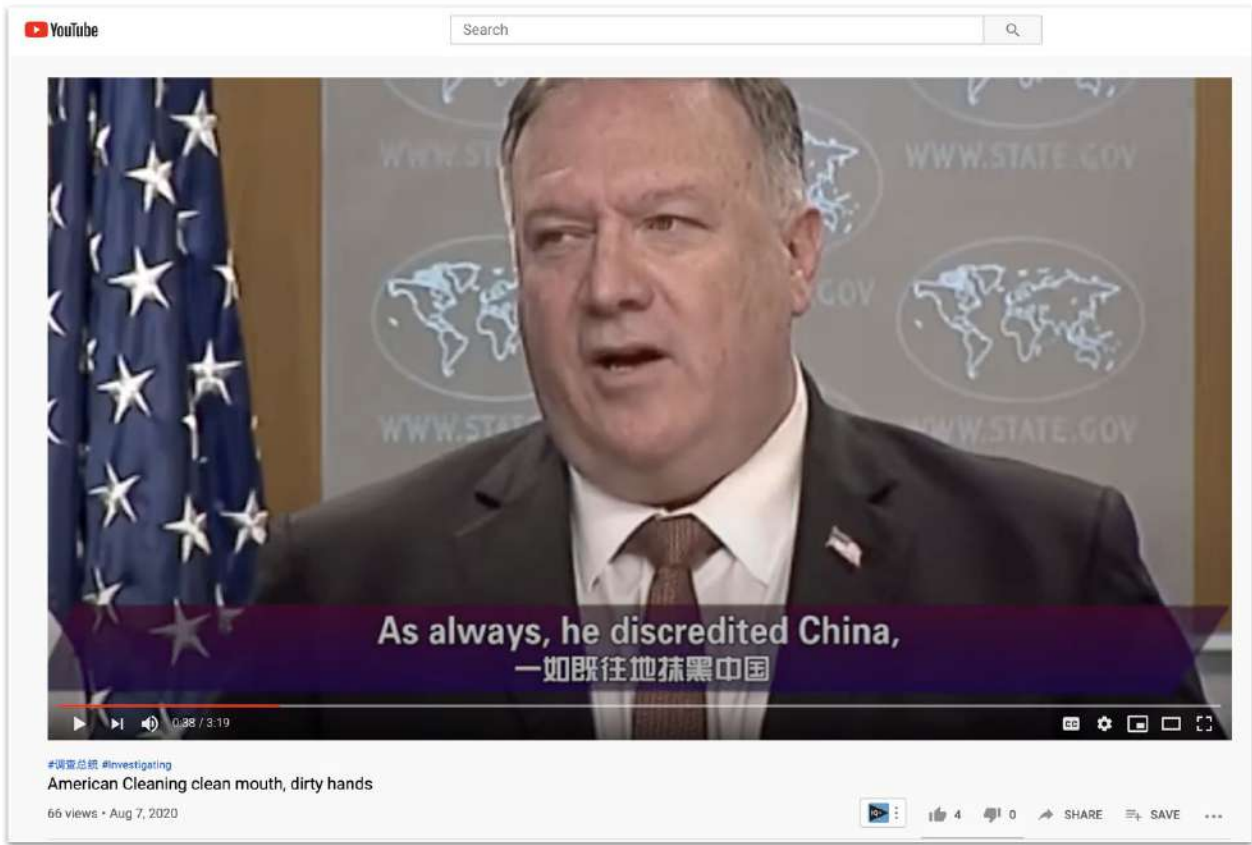




Top, profile for 风华依旧, showing its five videos, subscriber number, and profile picture. Bottom, the results of a search for the same user ID - UCduP12r1Yebpg0vs3i6lfZg - on vidooly.com. Note the identical profile pic and number of subscribers, but different name.

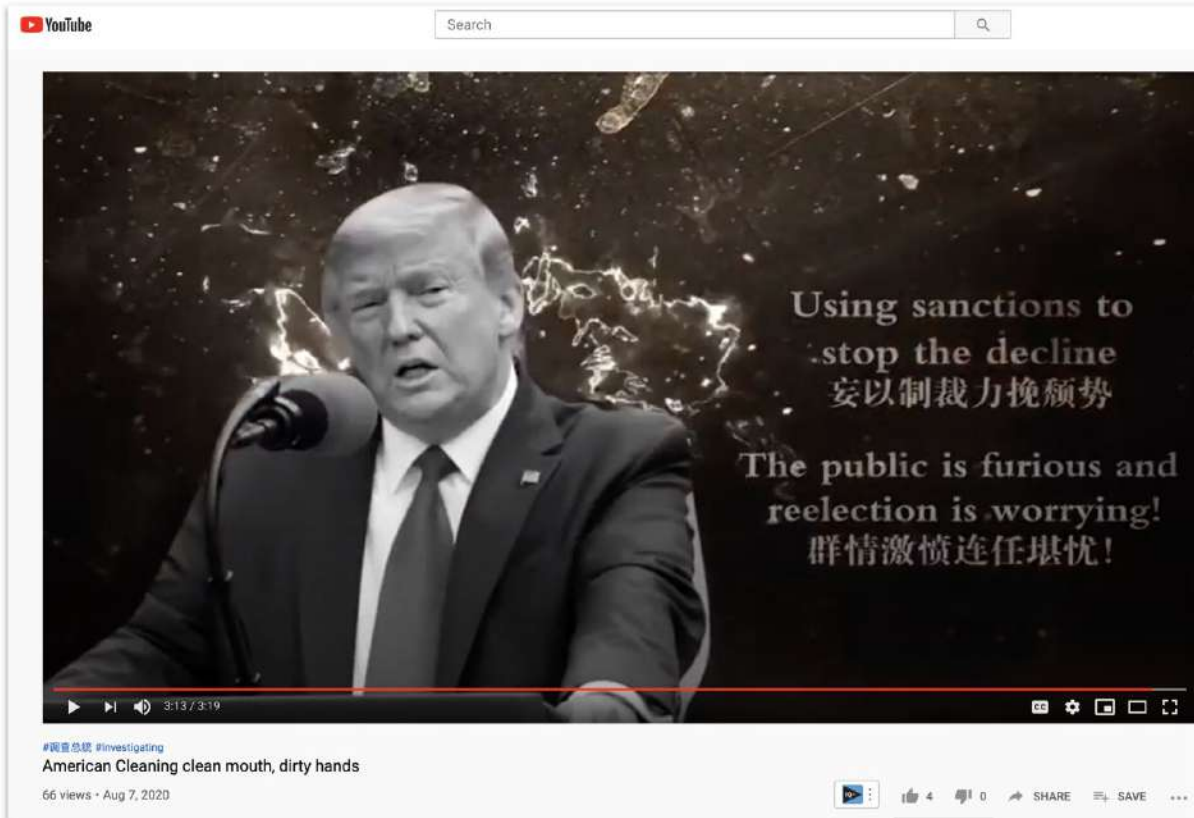
The Spamoouflage network clearly monitored developments in the United States' China policy, and reacted quickly. On August 5, U.S. Secretary of State Mike Pompeo [announced](#) a sweeping series of measures to promote the Trump administration's "Clean Network" program by limiting Chinese access to key parts of the internet in the United States. According to the page source data, the State Department's [announcement](#) was uploaded on August 5 at 20:57 UTC.

Just under thirty-six hours later, beginning at 07:44 UTC on August 7, Spamoouflage YouTube assets began sharing a video titled "[American Cleaning clean mouth, dirty hands.](#)" The video started by showing Pompeo's press conference. It accused the United States of abusing "state power to suppress and contain China's high-tech enterprises" and argued, "It's ridiculous to talk about 'cleaning the internet' when you are covered with dirt."



*Still from the video, "American Cleaning clean mouth dirty hands," posted on August 7.*

The video lasted a little over 3 minutes. It included footage from Pompeo's speech, together with stock news footage of senior American political and tech figures, but without audio of their words to give context. Instead, it featured an English-language voice-over, read by a woman's voice in unidiomatic English (for example, she pronounced "applications" as "apply-cations"), with English and Chinese subtitles. It concluded by linking the TikTok row to Trump's re-election chances.



Still from the video, "American Cleaning clean mouth dirty hands," posted on August 7.

The network's ability to conceive and produce a response to one specific American initiative within 36 hours - not quite within the same news cycle, but while the story was still reverberating - indicates a degree of tactical flexibility and reactivity and suggests a relatively rapid production cycle.

The tone of these videos was angry and partisan, but they were not a viral success. Some were not viewed at all. Some gathered a handful of views. A few showed viewing numbers in the double digits, and a very few were viewed over 100 times. Typically, each video featured on several different channels simultaneously or all but simultaneously; none of the channels that Graphika viewed appeared to have achieved organic virality.

One reason for this is likely to be that the quality of the videos left much to be desired. The voice-overs were clumsy and unidiomatic. Some appeared to have been poorly automated by a text-to-voice system that introduced basic errors, such as pronouncing the initial letters "U.S." as "us" (as in "the us government"). Others used idioms that may have been directly translated from Chinese but seemed out of place in English, such as "[Cast a chestnut in the fire will burn themselves with fire.](#)"



"Cast a chestnut..." Caption from the video, "[The cold war thinking of US political figures is out of date.](#)" The video was posted on July 30, 2020. Note that it was viewed three times (bottom left) but gained 56 likes (bottom right), a likely sign of inauthentic engagement.

"Not afraid of shadow crooked, China never causes trouble but it is not afraid of trouble. In response to the unreasonable behavior of the United States, China will certainly make the necessary response to safeguard its legitimate rights and interests," proclaimed another [video](#) posted on July 24.

As we have already hinted, still others made basic errors in their use of English that marked them out as, at best, incompetent productions with poor quality control. The headline of one video - "[Public blamed Trump's inaction](#)" - became, on the screen, "Public blamed Trump sinaction." Another claimed that the U.S. government "[never has a lower bound](#)" and seems to be "very good at be mischievous."



Caption from the video "[Lighthouse riots escalated, Public blamed Trump's inaction.](#)"



Caption from the video "[The U S government never has a lower bound.](#)" Note once more the numbers of likes compared with the number of views.

One of the earliest videos in this series, headlined “It is really hard for Americans today,” which posted on June 18, focused on accusations of brutality against the American police.

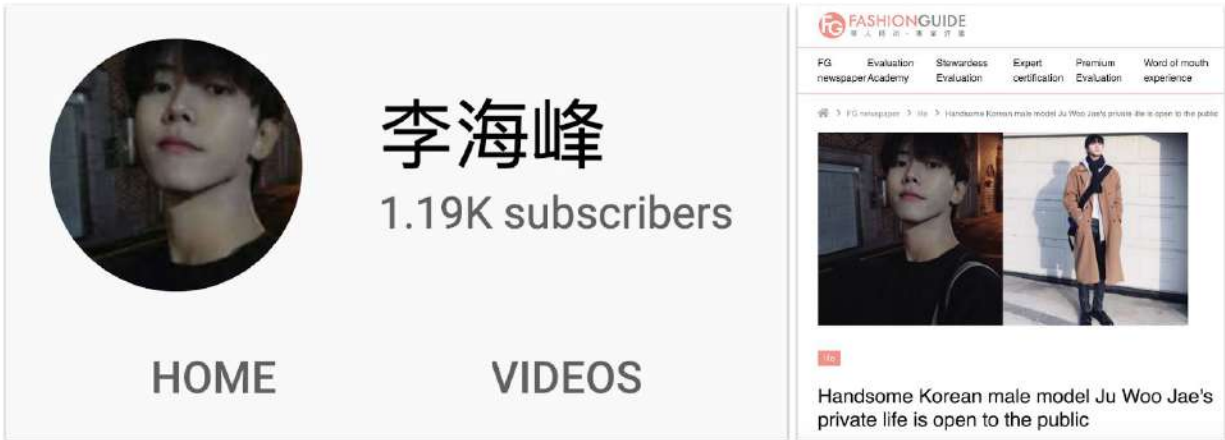
“Police have used tear gas against protesters in 98 cities so far in what is being described as the most widespread attack in half a century,” the video proclaimed over a collage of footage from protests and violent clashes with police. But as it spoke of the “most widespread attack in half a century,” it showed footage of British police officers rather than American ones, identifiable not only by their distinctive black-and-white caps, but by the Union Jack badge on the stab vest one of them was wearing. This may have been a deliberate insertion to highlight the widespread nature of the protests, but the video’s heavy focus on America and the American people elsewhere in the clip makes it at least equally likely that this was a simple error born of carelessness or ignorance.



Left, still from the video headlined “[It is really hard for Americans today](#),” timestamp 0:52. Right, enlarged view of the badge on the policeman’s stab vest, together with a view of the Union Jack at the same angle.

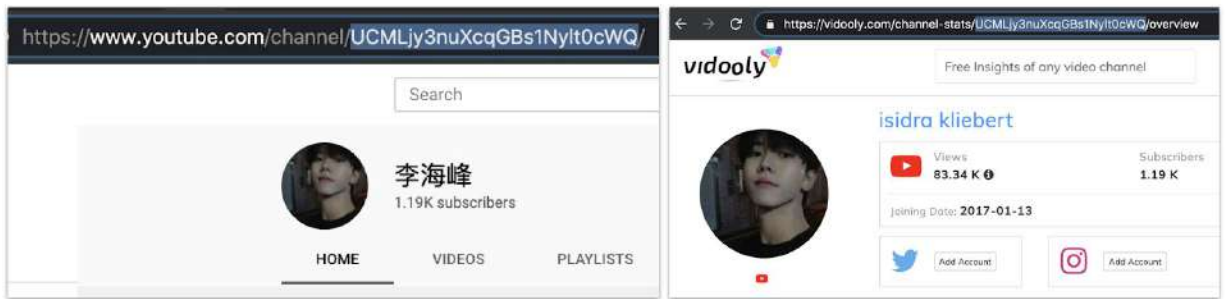
## Cross Eyes, Cross Platforms

We have already mentioned “Malcolm Daly,” the channel that posted the video about the “Lighthouse riots” escalating and whose behavior was typical of earlier Spamouflage Dragon assets. Another channel that posted the same video was equally typical of earlier Spamouflage generations. It was created on January 13, 2017; by August 1, 2020, it had amassed 83,337 views and 1.19k subscribers. Its name in 2020 was [李海峰](#) (Li Haifeng), but its profile picture was taken from Korean model [Joo Woo Jae](#).



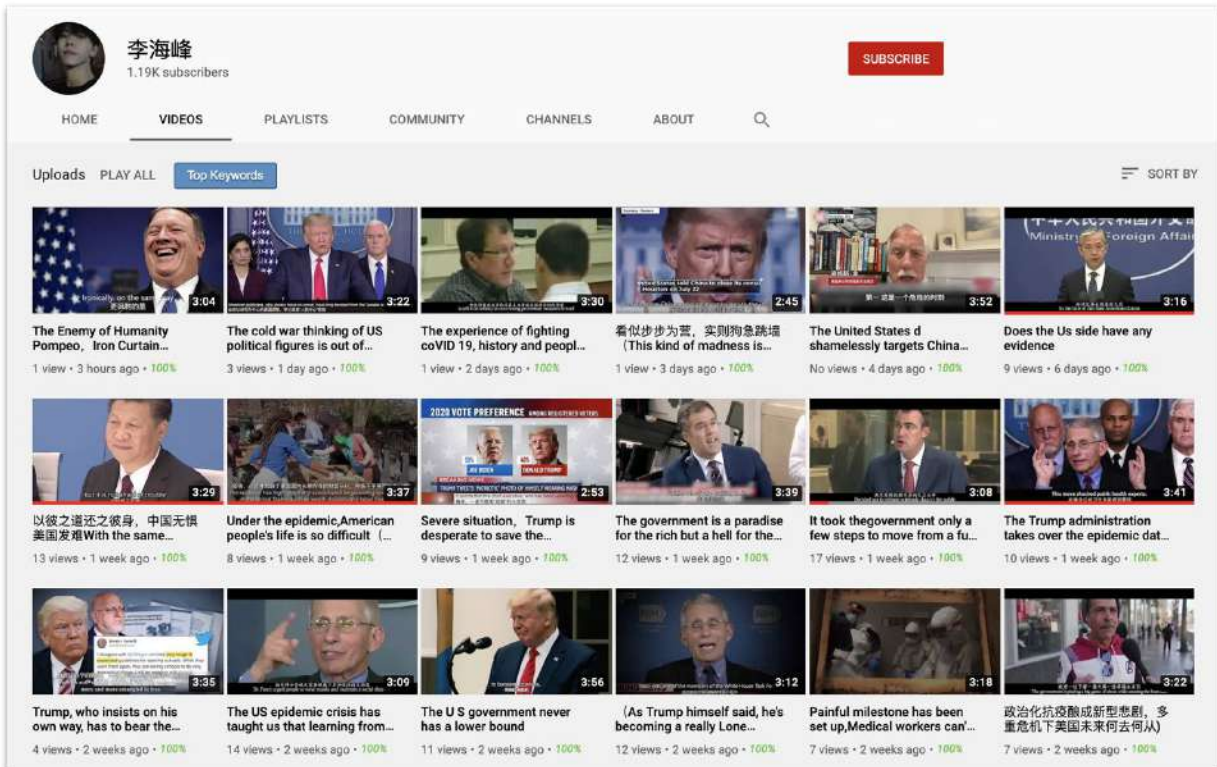
Left, profile details for 李海峰 (Li Haifeng) on YouTube. Right, the same image on the Taiwanese version of [fashionguide.com](http://fashionguide.com), auto-translated and attributed to Joo Woo Jae.

According to the online tool vidooly.com, Li Haifeng was not the channel's original name: it was originally called "isidra kliebert" but changed its name some time before July 20, 2020, when Graphika viewed it under its Chinese name.



Left, the profile page for Li Haifeng on YouTube, with the user ID number highlighted in the URL bar. Right, profile for "isidra kliebert" according to [vidooly.com](http://vidooly.com), showing the same profile picture and the same user ID.

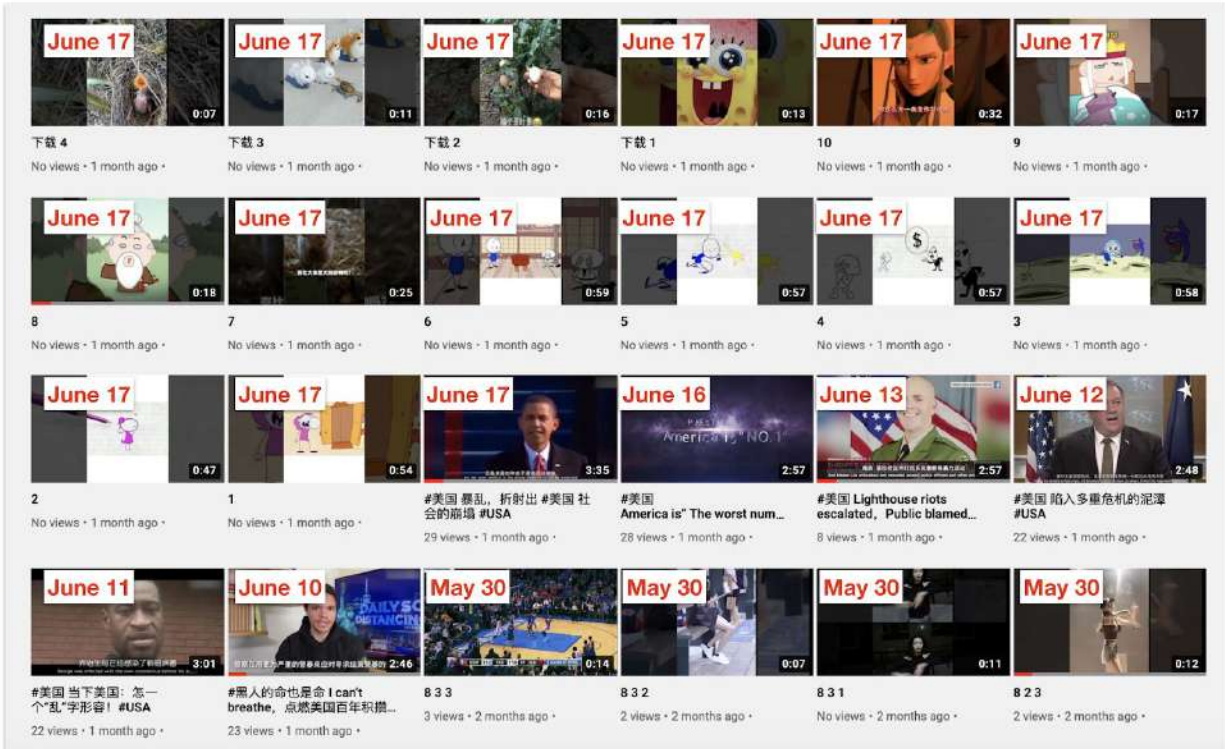
Its posting behavior fell into three distinct periods. The most recent, beginning on June 18, featured an unbroken series of medium-length videos (approximately 3 minutes) in English with bilingual English and Chinese subtitles. All of these criticized the United States in general and the Trump administration in particular; two were headlined in Chinese. None of these videos was viewed more than a handful of times. Many of them were also posted by other accounts that we assess with high confidence as belonging to the Spamuflage Dragon network.



Recent videos by Li Haifeng as of August 1, 2020, showing the exclusive focus on the United States.

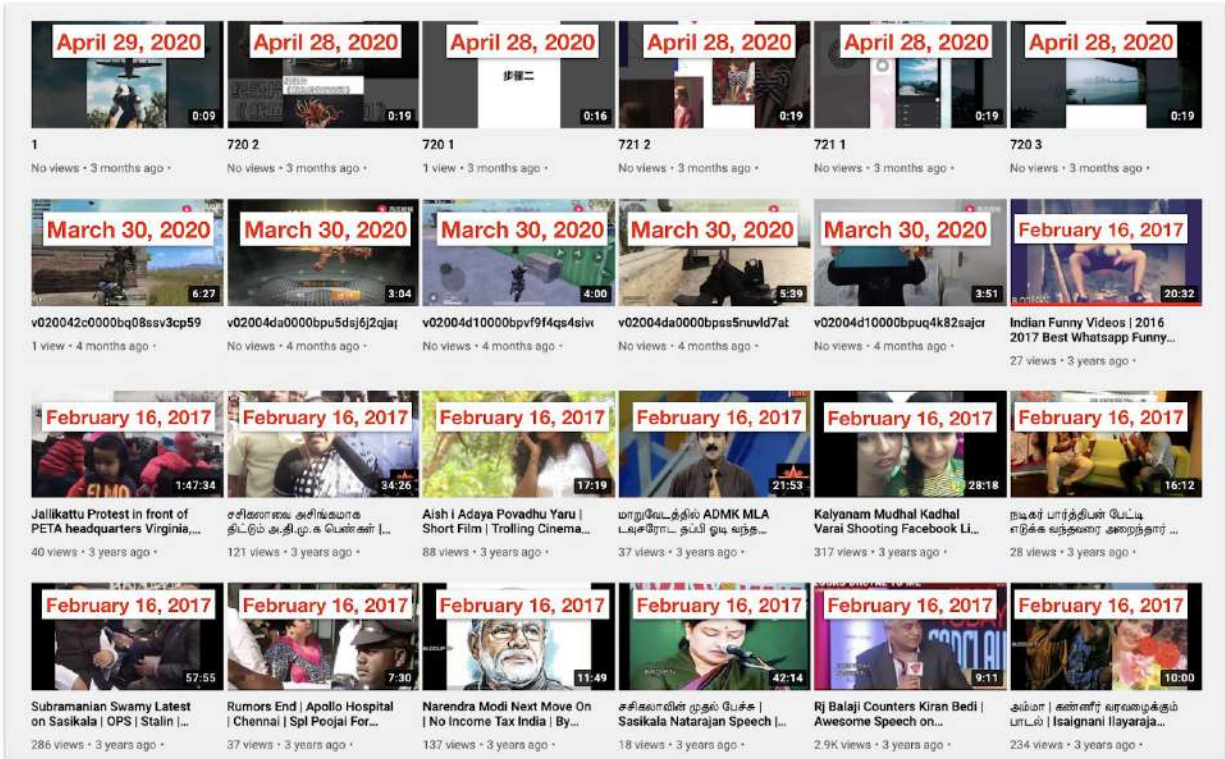
Earlier, between March 30, 2020, and June 17, 2020, the account predominantly posted short videos (typically less than a minute) of cartoons and music clips. From June 10 to 17, it inserted six English-language videos with bilingual subtitles that focused on the U.S. racial justice protests, the United States' history of racial injustice, and the coronavirus pandemic. The posting pattern was not uniform: it posted one new political video per day (with a break on Sunday-Monday, June 14-15) but preceded and followed these with slews of short videos all posted on the same day. Again, none was frequently viewed.





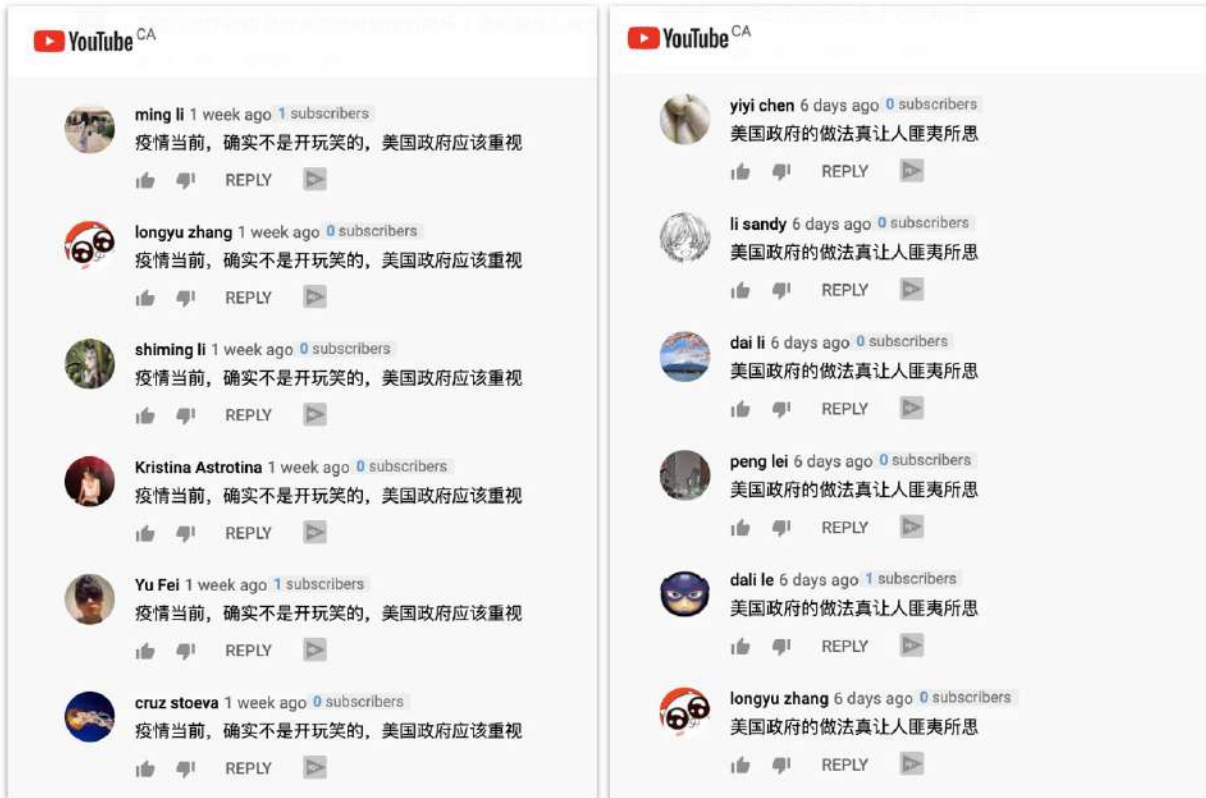
Videos by Li Haifeng between May 30 and June 17, 2020, showing the combination of spammy posts and political ones.

Before this activity, the account behaved very differently. It does not appear to have posted at all between February 17, 2017, and March 30, 2020. On February 16, 2017, however, it posted 45 videos of varying durations (from 1:30 to over an hour), all focused on India. Most of these were viewed a few hundred times, but one video was viewed some 43,000 times. Whatever the reason for this diverse viewing success in early 2017, the account's long silence, followed by its shift to spam-posting short random clips, and then to posting political content about a wholly different subject, suggests that it was compromised or acquired by a new user, and repurposed.



*Early videos posted by Li Haifeng in 2017 and 2020, with a three-year gap between.*

In typical Spamouflage Dragon fashion, such videos were accompanied by large numbers of supportive comments in Chinese - sometimes significantly more comments than there were accounts that viewed the videos. Some of these comments appear to have been automated or outsourced to human users for wholesale re-posting, with the same comment posted by many different headquarters. Some of the accounts that posted in Chinese had names that appeared more Slavic, English, or Spanish - not a conclusive signal on its own, but it does suggest that these accounts, too, had been taken over and repurposed for Spamouflage posting.

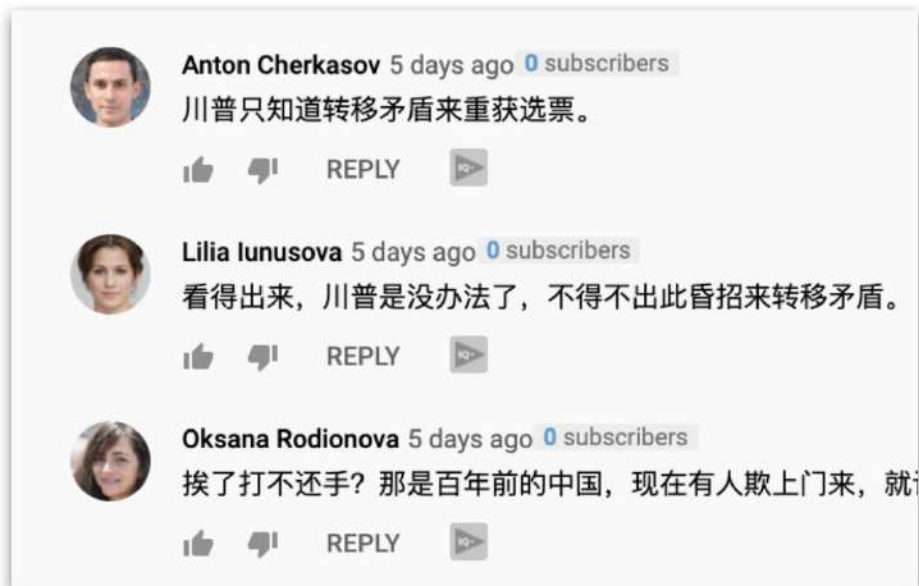
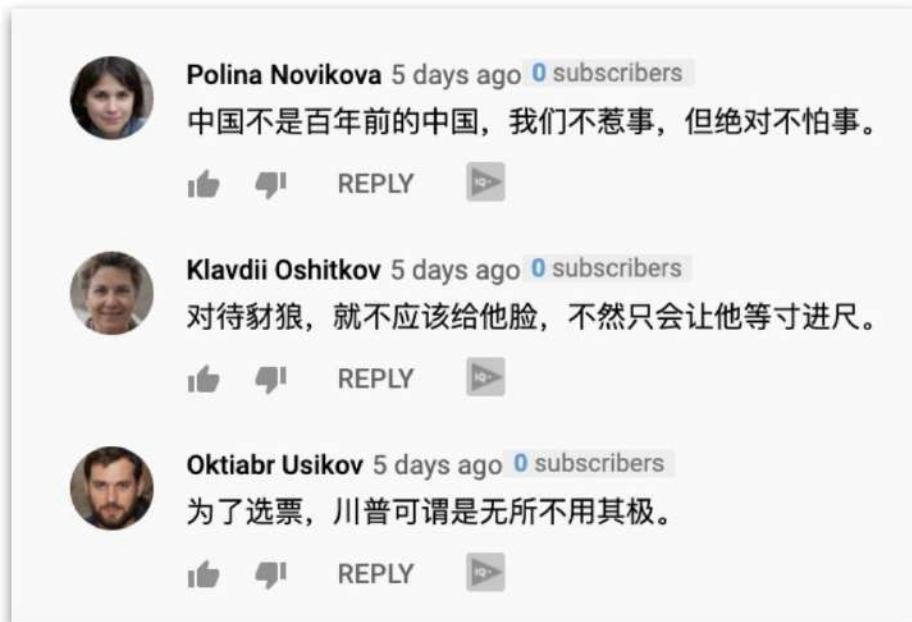


Comments on two different videos by Li Haifeng, showing the identical comment being posted by half a dozen different users. Note also the names of the apparent authors of these Chinese-language posts, including “Kristina Astrotina” and “cruz stoeva.”

## FFS: 2.0: Fake Faces Swarm Again

In what appears to be a new development for this network, some of the YouTube accounts that commented on these anti-U.S. videos had profile pictures that appear to have been generated by Generative Adversarial Networks (GAN), a class of machine-learning frameworks enabling computers to generate synthetic photographs of people who never were (in short: fake faces!). The website “thispersondoesnotexist.com” popularized this application, making pictures created by StyleGAN2 accessible to any Internet user.

Graphika has already found such GAN-generated profile pictures used *en masse* in another operation that we dubbed [FFS](#), Fake Face Swarm. To date, these images can be recognized as GAN-generated thanks to a number of telltale features, such as asymmetries where there should be symmetries (e.g., the frames of glasses or earrings), vague and misaligned background features (e.g., walls or trees), and the fact that the eyeballs of each GAN-generated profile picture are in the same place. This can easily be demonstrated by superimposing all the profile pictures on one another and then making them opaque to see if the eyeballs line up: it is worth noting that this forensic “hack” may soon be rendered obsolete by the rapid evolution of GANs, leading to improved and harder to detect synthetic photographs



Comments on a [video](#) posted by suspected Spamouflage Dragon asset "Siqi Zheng" on the closure of the Houston consulate, showing the profile pictures of some of the commenters. The video featured almost a dozen comments from accounts with profile pictures such as these.



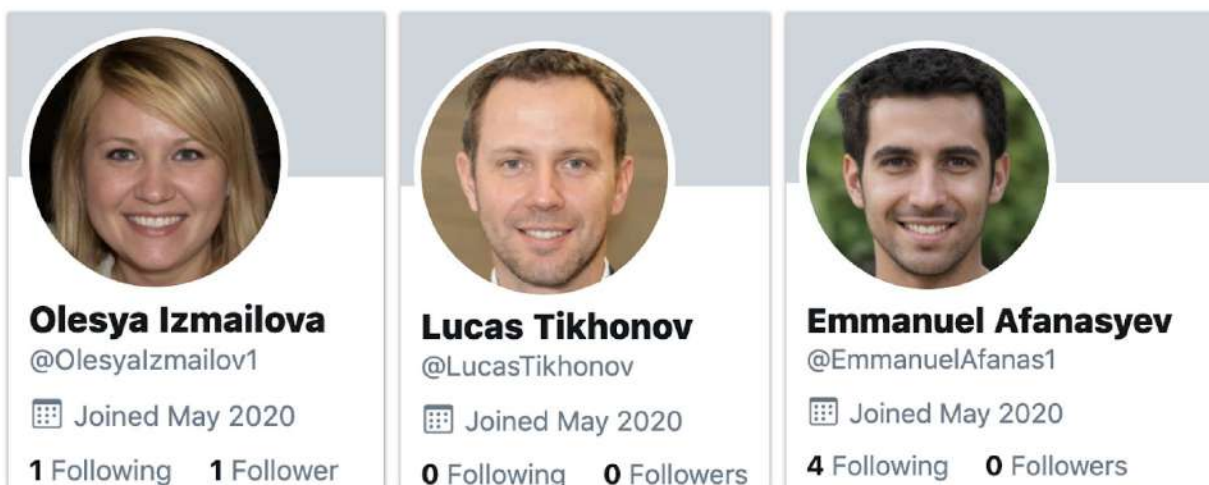
*Profile pictures of the YouTube accounts [Irina Ivanova](#), [Adrian Mishatkin](#), [Larissa Caiden](#), [Klavdii Oshitkov](#), [Oktiabr Usikov](#), [Polina Novikova](#), [Anton Cherkasov](#), and [Lilia Lunusova](#), all of which commented on the above-mentioned video. On the right, the eight profile pictures rendered opaque and superimposed on one another. Note how the eyeballs align, and how each individual profile picture is set against a blurred and indeterminate background, typical of GAN-generated images.*

The Spamouflage Dragon network did not limit itself to YouTube, although this appears to have been its favored platform for posting its video content. The same videos also showed up on Twitter and Facebook, where they were also amplified by accounts and pages that we assess with a high degree of confidence are part of the same network.

For example, on July 24, a YouTube channel with one subscriber called Shiqi Zhang posted a [video](#) called [以彼之道还之彼身，中国无惧美国发难](#), “With the same principle, China is not afraid of the US” (the Chinese portion is a translation of the English). This video focused on the closure of the Chinese consulate in Houston and accused the United States of “unreasonable” and “sadistic” behavior. It was only viewed ten times, but it received 53 likes.

On the same day, a Twitter account called @raojiu541 tweeted the same video.<sup>7</sup> This appears to have been a Spamouflage Dragon asset: its content, behavior, and interaction with other accounts all matched the network’s characteristics. On this occasion, it was retweeted, and replied to, by a collection of recently created accounts, all of which featured GAN-generated profile pictures. Some of the amplifier accounts posted retweets in multiple languages, especially Turkish and Arabic, and featured a range of ads for bitcoin and online gambling services; they may have been either bought or rented for extra engagement. Others only shared or commented on Chinese-language content from assets that we assess belonged to Spamouflage Dragon. These typically had between none and five followers.

<sup>7</sup> The account was suspended before Graphika could archive it, although we managed to screenshot some of its posts; hence no link is available.

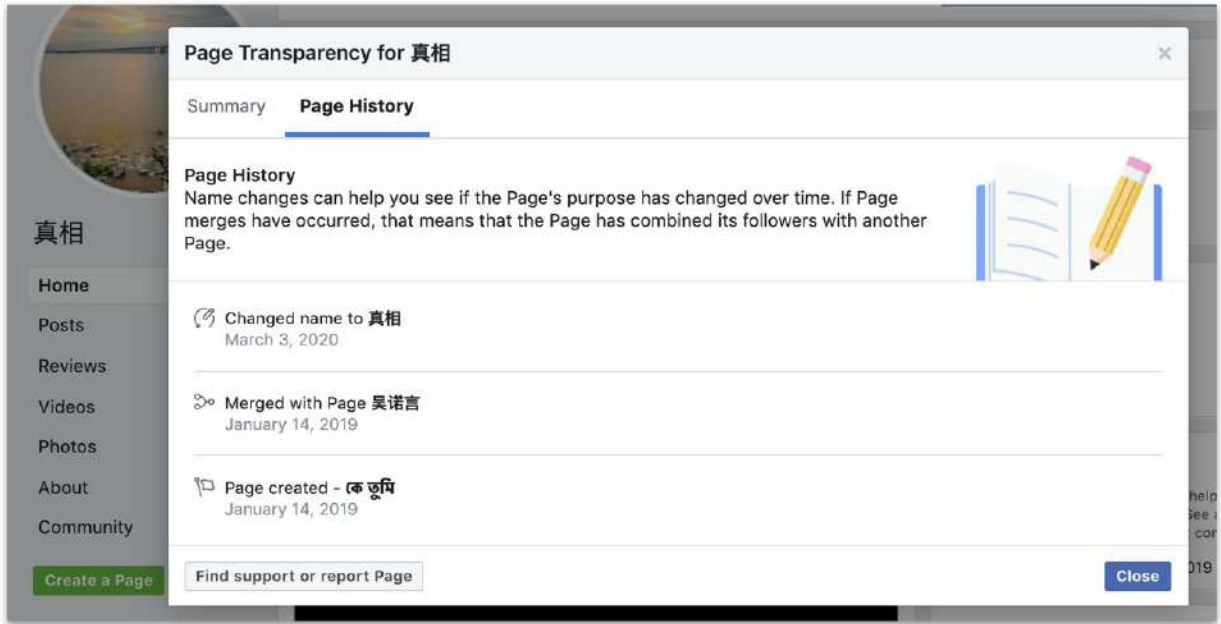


Profile pictures and follower stats for three of the Twitter accounts that posted in Chinese and shared a number of Spamuflage Dragon videos. Note the minimal followings.



Profile pictures of the Twitter accounts [Olesya Izmailova](#), [Lucas Tikhonov](#), [Emmanuel Afanasyev](#), [Jamal Demidov](#), [Myron Bezrukov](#), [Angela Agapova](#), [Lilith Ulyanova](#), and [Albert Zuyev](#), superimposed and opaqued.

On Facebook, meanwhile, the video was shared by a page called [真相](#) (the truth). This was not, however, its original name: when it was created on January 14, 2019, it was called কে তুমি (Bengali for “who are you,” according to Google Translate), but on the same day, it merged with a page called [吴诺言](#) (Wu Nuoyan). This pattern of behavior has been observed at multiple instances throughout Spamuflage Dragon’s operations.



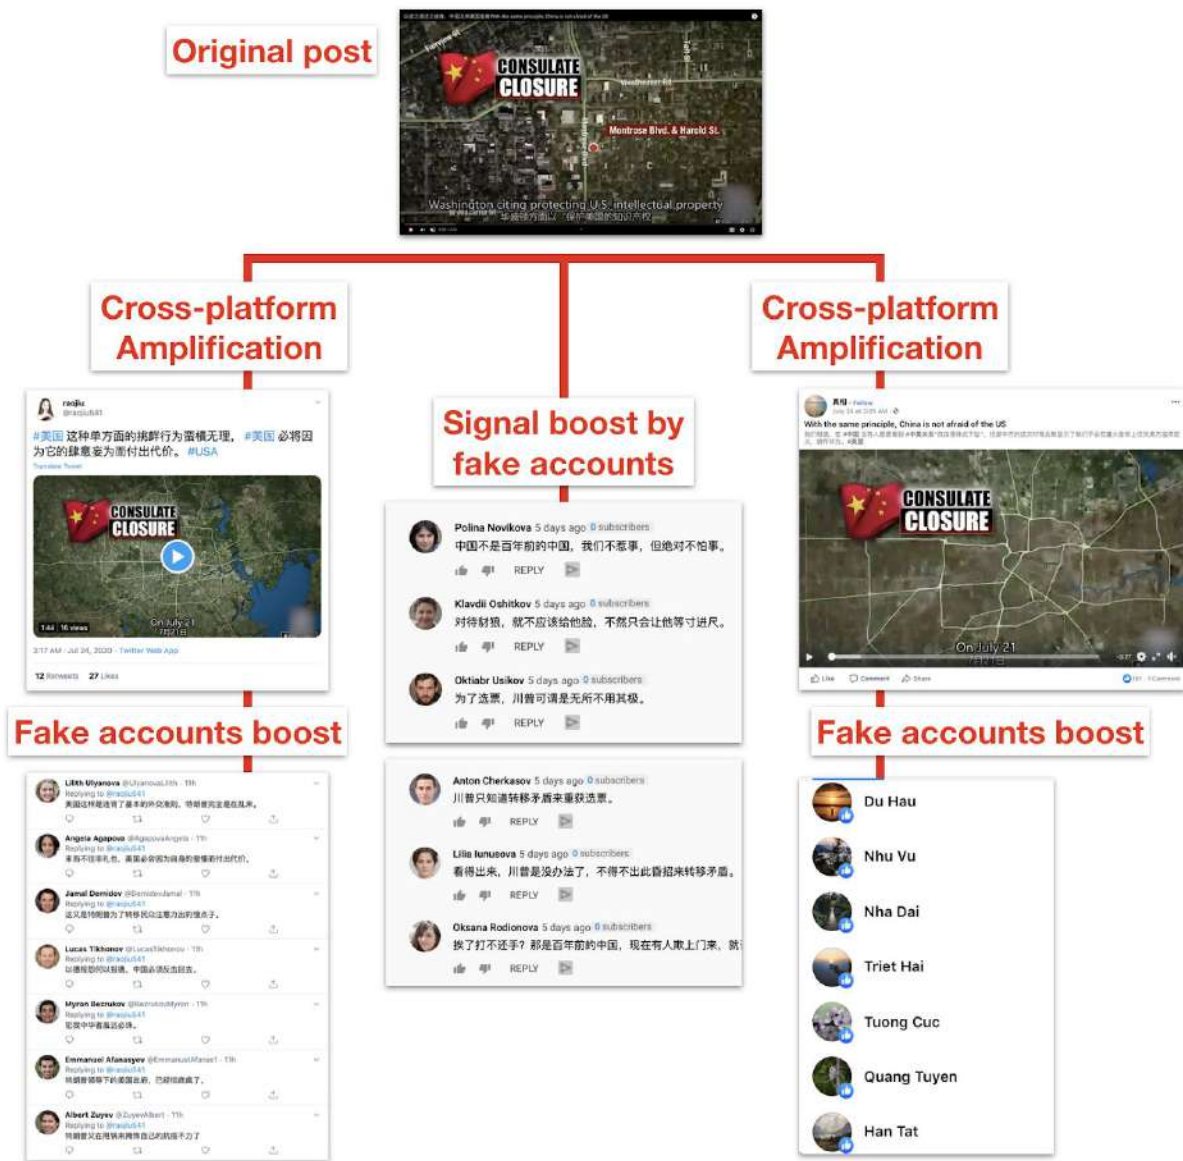
Page Transparency for “the truth,” showing its original Bengali name and the fact that it was merged with another page on the day it was created, a pattern that has repeatedly been observed in Spamuflage Dragon.

Strikingly, when Graphika first analyzed the Facebook post on July 25, it had been viewed five times but had 101 likes, including a large number that came from accounts with no profile picture or posting history. Two days later, the number of likes had dropped to 41, suggesting that 60 accounts had left the platform, likely as a result of Facebook’s enforcement against fake accounts. As of August 2, the number of likes stood at 33, indicating further disappearing accounts. Many of the remaining accounts had Vietnamese names, featured generic landscape shots as their profile pictures, showed no posting history, and uploaded their first profile pictures (a proxy for their creation) in early June, suggesting that they, too, were inauthentic.



The only posts made by Quang Tuyen, Nhu Vu, and Nha Dai: the upload of their profile pictures, all of which are generic stock shots. All the images were uploaded on June 15-16, 2020.

Taken together, these assets created a cross-platform network of inauthentic amplification, providing comments, likes, replies, retweets, and shares across Facebook and Twitter, as well as YouTube. The retweets and shares acted as a straightforward signal boost; the comments both added a veneer of authenticity (as long as one does not look too closely!) and provided validation, in Chinese, to the video's main message.

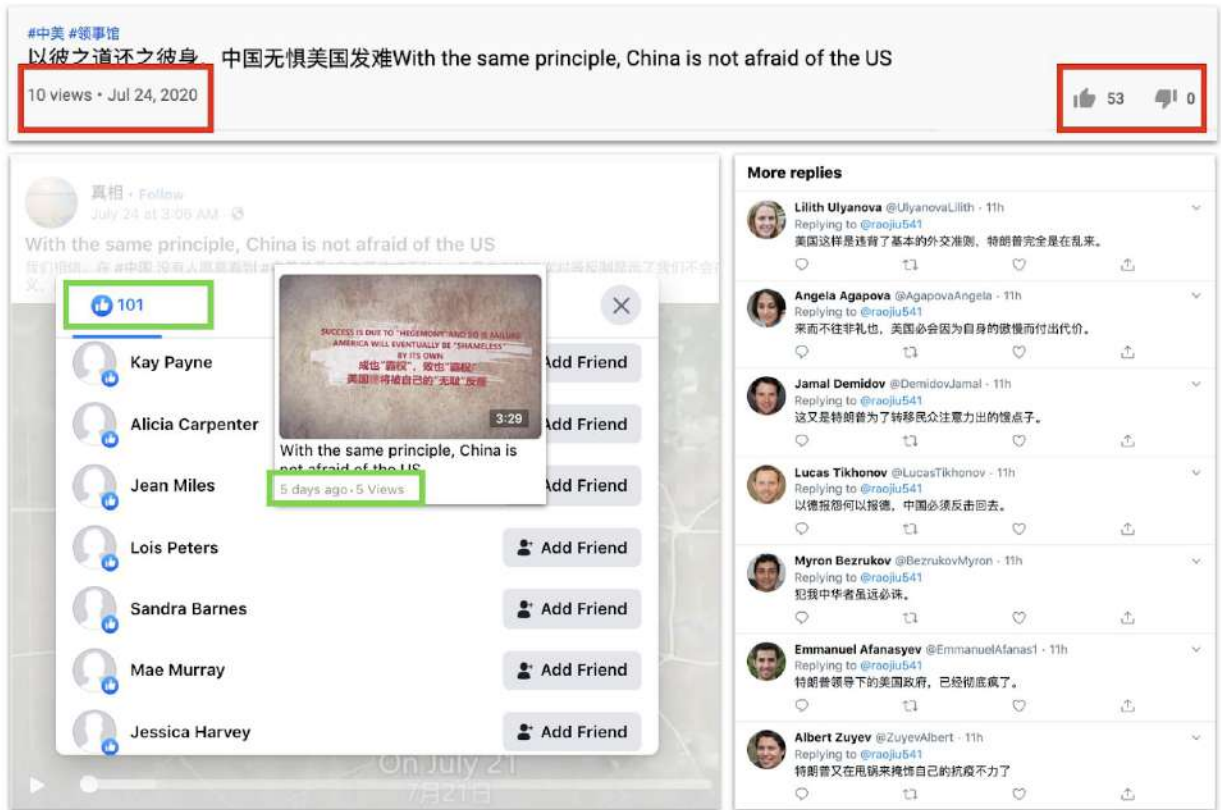


Posting and amplification of the Spamouflage Dragon video on [YouTube](#) (top), [Twitter](#) (left), and [Facebook](#) (right).


However, this was only a veneer. The beauty of Spamouflage Dragon was decidedly skin-deep: the mismatch between viewing figures and reactions and the poorly crafted fake profiles leveraging stock photos and GAN images combined to expose a clear pattern of inauthentic behavior. When the fake accounts were stripped away (and accounts on all three platforms were



taken down, most likely by automated systems, while Graphika was analyzing them), there was no sign of sustained engagement by demonstrably authentic users.



Patterns of inauthentic behavior: top, the headline bar from the YouTube video, showing the number of views compared with the number of likes. Bottom left, screenshot from the Facebook video when it had been viewed five times but still had 101 likes; the faceless accounts that liked it were taken down shortly afterward. Bottom right, replies to the Twitter version from accounts with probable GAN faces.



**Graphika**

© Graphika All Rights Reserved 2020