

ISD

Powering solutions
to extremism
and polarisation



alliance for
securing
democracy

Reply All: Inauthenticity and Coordinated Replying in Pro-Chinese Communist Party Twitter Networks

Raymond Serrato and Bret Schafer



22nd July 2020

Reply All:

Inauthenticity and Coordinated Replying in pro-Chinese Communist Party Twitter Networks

Author: Raymond Serrato

The world has witnessed an ongoing crisis of disinformation and online manipulation during both the COVID-19 pandemic and the racial justice movements sweeping across communities around the globe. Alongside outright false stories, misleading cures and conspiracy theories virulent on social media, a heady mix of state actors, extremists, political activists and commercial interests are also using deceptive behavior to fool people online. Since the infamous Kremlin-backed campaign targeting the US elections in 2016, social media companies have tried and often failed to get to grips with a torrent of inauthentic accounts and covertly coordinated networks.

The Institute for Strategic Dialogue (ISD) and Alliance for Securing Democracy's (ASD) new report sheds light on the tactics used to manipulate information online through the case study of a pro-Chinese Communist Party (CCP) Twitter network targeting both Chinese and English-language audiences online. A central tactic of this network is the use of replies on Twitter to amplify pro-CCP narratives and to attack anti-CCP individuals or institutions – an understudied aspect of network behaviour to date.

The report highlights the importance of transparency and caution in assessing inauthenticity online. ISD used seven analytical signals to build a picture of a likely coordinated, inauthentic network promoting CCP propaganda and attacking anti-CCP spokespeople or institutions. Activity rates are often relied on as the way to assess possible inauthenticity of accounts or networks, but this approach misunderstands the complexity of disinformation actors and authentic human behaviour online. ISD has used a multilayered research approach to provide a much more comprehensive picture of what coordinated networks can look like in 2020 and how difficult it can be to sift them out from authentic online activity. It is hoped this combined methodology can contribute to the burgeoning field of information operations research. This research was combined with the insights from ASD's monitoring of official CCP social media networks to better understand the parallels in overt state activity and pro-state covert networks online.

ISD analysed the data from 828,646 followers of two CCP officials and has identified a network of suspicious Twitter accounts engaged in propagating pro-CCP narratives. A large number of accounts were created in March 2020 and are characterised by several indicators of inauthenticity and coordinated messaging on issues that include COVID-19 conspiracies, Taiwan independence, the Hong Kong demonstrations, and the recent antiracist protests in the US.

The accounts engaged in high rates of replying, with more than 800 accounts replying at least 100 times, and dozens of accounts replying hundreds of times to single accounts in short bursts of activity. This “reply targeting” focused on President Trump, Secretary of State Mike Pompeo, Taiwanese President Tsai Ing-wen, and other diplomatic accounts. The replies also targeted the Chinese media accounts of the New York Times, US public broadcaster Voice of America, German public broadcaster Deutsche Welle, the BBC, and others. The replies varied in their content, but often contained insults or messaging that was broadly supportive of the CCP’s interests. Although these narratives were pro-CCP, ISD cannot attribute this activity to any specific state or non-state actor.

ISD worked with ASD to compare its findings to data collected by ASD’s Hamilton 2.0 dashboard, an open-source tool that monitors official CCP government and state-backed media accounts. ASD and ISD’s joint analysis found numerous similarities in the behavior, narratives, and specific messages promoted by the pro-CCP follower network and CCP officials. In addition, ASD noted several occurrences of CCP diplomats engaging with suspicious and, in some cases, now-suspended accounts, though ASD cannot gauge whether this engagement was incidental or a sign of coordinated inauthentic behavior.

Key Findings

- ISD identified numerous indicators of inauthenticity in the Twitter followers of the two CCP officials, including:
 - A significant proportion of accounts with a high number of digits in their handles or handles whose naming conventions suggest computer generation or setup in haste
 - Accounts created in short time intervals or a specific time block
 - Accounts using profile photos found elsewhere online or unassociated with the account’s purported identity
- Tweeting in multiple languages, varying mainly between three and five languages per account.
- Accounts engaged in high rates of replying, making up nearly 65% of sampled tweets, and targeting accounts that tweeted negatively about China’s response to COVID-19, media reporting about Taiwan’s bid to join the World Health Organization (WHO), and the demonstrations in Hong Kong.
- The creation of pro-CCP followers coincided with a surge in official CCP activity on Twitter. Of the 162 government and diplomatic accounts ASD tracks on Hamilton 2.0, roughly 75% were created in 2019 or 2020. In addition, CCP government and diplomatic accounts increased their aggregate output from roughly 5,000 tweets in January 2020 to nearly 20,000 tweets by April 2020.¹
- Attempts to inflame US domestic tensions over police brutality while criticising the Trump administration’s response to the protests in relation to its stance on Hong Kong were observed in both the pro-CCP follower and official CCP networks.

¹ This is due in part to new CCP officials joining Twitter after January 2020. In addition, ASD’s ability to pull historic data from Twitter’s open API is limited, meaning some CCP accounts may have incomplete data from January 2020 depending on when they were added to ASD’s official account list.

Methodology

ISD collected 828,646 followers and 179,122 friends from two CCP official accounts and retrieved the associated metadata for these accounts. Additional tweet-level data was retrieved from 5,156 accounts that had tweeted more than 10 times and had been created between 13-23 March, resulting in 445,570 tweets. ISD used a combination of R, a statistical software used for data analysis and visualization, and a custom-built tool to retrieve the data.

As part of its Hamilton 2.0 dashboard, ASD continuously collects data on 238² Twitter accounts belonging to individuals or entities connected to the Chinese government (162 accounts) or Chinese state-funded media outlets (76 accounts). The dashboard tracks accounts that primarily target foreign audiences, namely embassies, consulates, foreign ministries, ambassadors, key government figures, and international media outlets and their affiliated channels and programs. Most of these accounts openly note their affiliation with the Chinese government, and thus represent overt rather than covert influence.³

ASD publicly launched the China section of its dashboard in March 2020; however, ASD's data collection began with an initial seed list of official CCP accounts in late 2019. ASD also regularly adds newly created or previously unidentified accounts to its official account list. ASD's dataset therefore includes data on some accounts dating back to mid-2019, while data on other accounts is only available from early to mid-2020. In total, ASD has collected 425,560 tweets connected to official Chinese government and state-funded media accounts.⁴

² Account totals accurate as of July 31, 2020

³ ASD's complete methodology for the Hamilton 2.0 dashboard, see: <https://securingdemocracy.gmfus.org/hamilton-2-0-methodology-faqs/>

⁴ Tweet totals accurate as of 6:00pm ET on July 31, 2020

1. When were the followers of these CCP accounts created?

Nearly 300,000 of the followers, or 27%, were created in 2020; 16% in 2019; 10% in 2018; and 10% in 2017, with decreasing percentages for all years thereafter. The creation of these accounts generally tracks the increasing number of CCP officials and diplomatic missions that have joined Twitter since 2017.⁵ Of the 162 Chinese diplomatic and government accounts tracked on ASD’s Hamilton 2.0 dashboard, 35% were created in 2020; 42% in 2019; and 8% in 2018, with decreasing percentages thereafter.

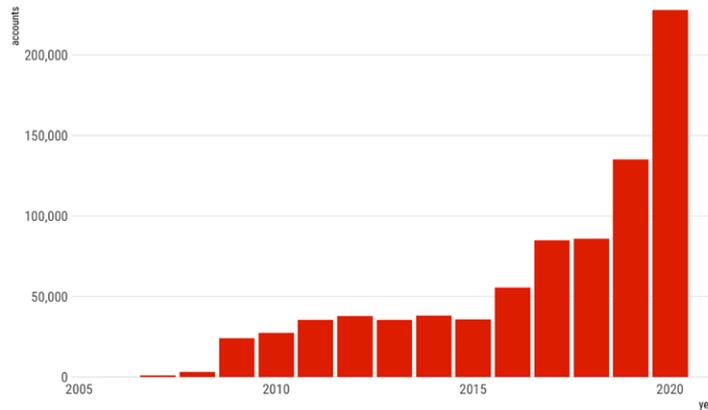


Figure 1: Volume of CPP accounts and year created.

In 2020, March and April saw the highest number of new accounts created in the pro-CCP follower network. In March, in particular, there were spikes of new accounts created between 13-24, with the sharpest rise in accounts on 17 March, the same day Chinese officials rebuked President Trump for referring to COVID-19 as the “China Virus.” This period also coincides with the first instance of a key CCP official promoting a conspiracy theory about the origins of the coronavirus, which received significant international media coverage.⁶



Figure 2: Sharpest rise in CPP accounts created in March/April 2020.

⁵ BBC News, China and Twitter: The year China got louder on social media (29 December 2019), <https://www.bbc.co.uk/news/world-asia-china-50832915>, [Accessed on 30 June 2020].

⁶ See, for example: <https://www.cnn.com/2020/03/13/asia/china-coronavirus-us-lijian-zhao-intl-hnk/index.html>

2. Where do the accounts pretend to be located?

Nearly 20% of accounts report their locations as China, with Pakistan appearing as the next cited location. ISD researchers plotted the top 70 locations by the total number of users and visualised their distribution on a map. It is important to note that locations are self-reported and although Twitter is blocked on the Chinese mainland, many Chinese users reportedly use virtual private networks (VPNs) to access the service.

Note: General locations (“United States”) are assigned a default latitude/longitude within the country and labelled as such.



Figure 3: Top 70 CPP account locations.

The centrality of Pakistan in the location data is perhaps not surprising given that the country served as a test bed of sorts for the CCP’s earliest diplomatic forays on Twitter.⁷ Zhao Lijian – one of China’s first and most influential diplomats on Twitter – spent four years honing his craft at the Chinese embassy in Islamabad before his promotion to his current post in Beijing as the deputy director general of the foreign ministry’s information department. The PRC Embassy in Pakistan⁸ was also one of the first Chinese diplomatic missions to create an official Twitter account. Finally, the current Consul General of the PRC in Karachi, Li Bijian (@Libijian2), has been the CCP’s most active diplomat (by a significant margin) since joining Twitter in January 2020.⁹ Pakistan was the fifth most mentioned country by CCP government officials and diplomats tracked by ASD’s Hamilton dashboard between 1 January and 31 July, 2020 (Fig. 4).¹⁰

⁷ Zhao Lijian created his Twitter account (@zlj517) in May 2010. As of 30 July, 2020, he had more than 746,000 followers, more than four times the followers of the official Ministry of Foreign Affairs account (@MFA_china).

⁸ The Chinese Embassy in Pakistan’s Twitter account (@CathayPak) was created in September 2015, making it one of the five earliest PRC embassies on Twitter. As of 30 July 2020 the account boasted over 106,000 followers.

⁹ According to data collected by ASD, @Libijian2 tweeted more than 17,000 times between 1 January and 31 July, 2020, more than double the output of the CCP’s next most prolific diplomat on Twitter.

¹⁰ Note, this data does not include the activity of CCP state media accounts tracked on the Hamilton dashboard.

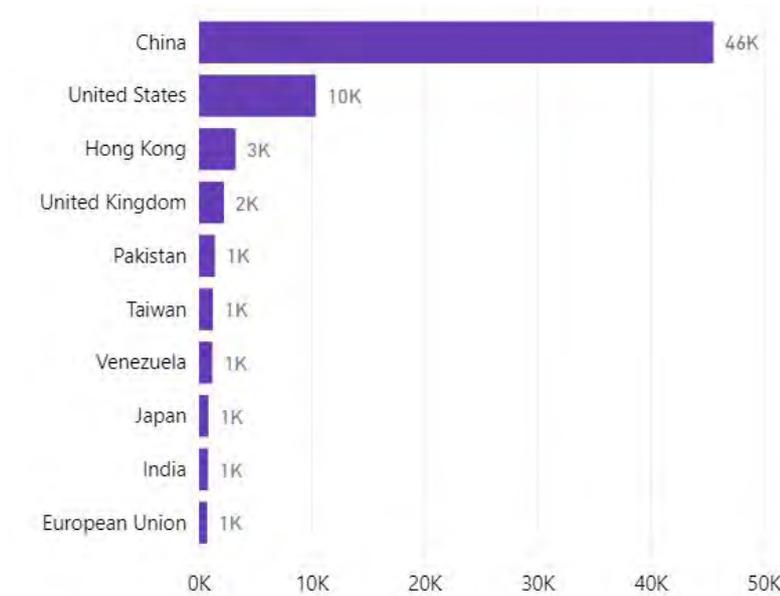


Figure 4: Countries mentioned by CCP government officials and diplomats tracked by ASD's Hamilton dashboard between 1 January and 31 July, 2020.

Retweet data collected by ASD also reflects a strong Pakistan connection. Of the 100 most-retweeted accounts by CCP officials or state media accounts (not including retweets of other government officials or state media accounts tracked by ASD), nine were located in or had a connection to Pakistan. Only China and Venezuela had more connections to the 100 most retweeted accounts.

3. Account creation and tweeting patterns

In 2020, the vast majority of accounts were created almost exactly between 10am-10pm local time in Beijing. Tweeting activity occurred mainly between 10am-12am, with the volume decreasing between 12am-5am. From the tweet-level data, ISD cannot determine whether the pattern of tweeting is in any way automated or scheduled.

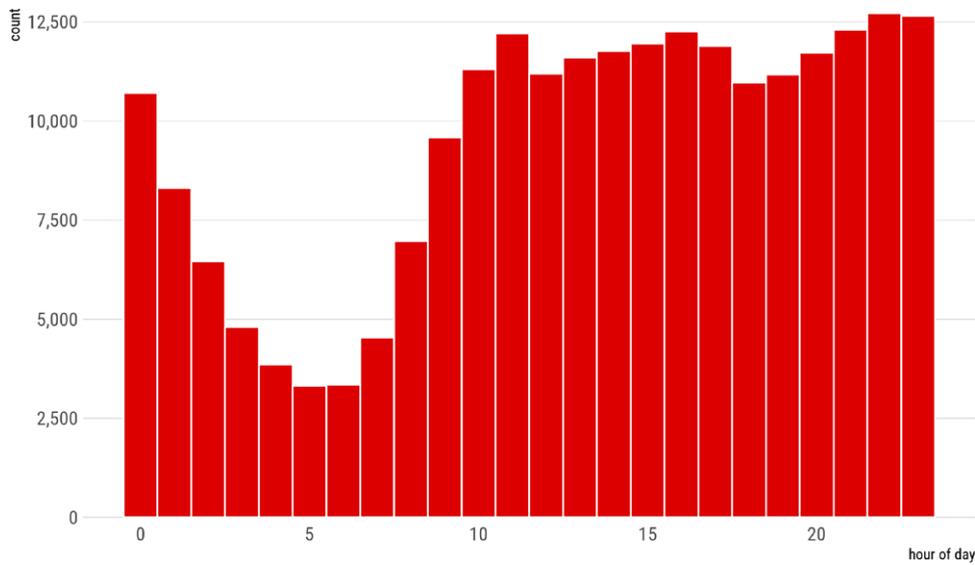


Figure 5: Creation of accounts by the hour of day (Beijing) in 2020.

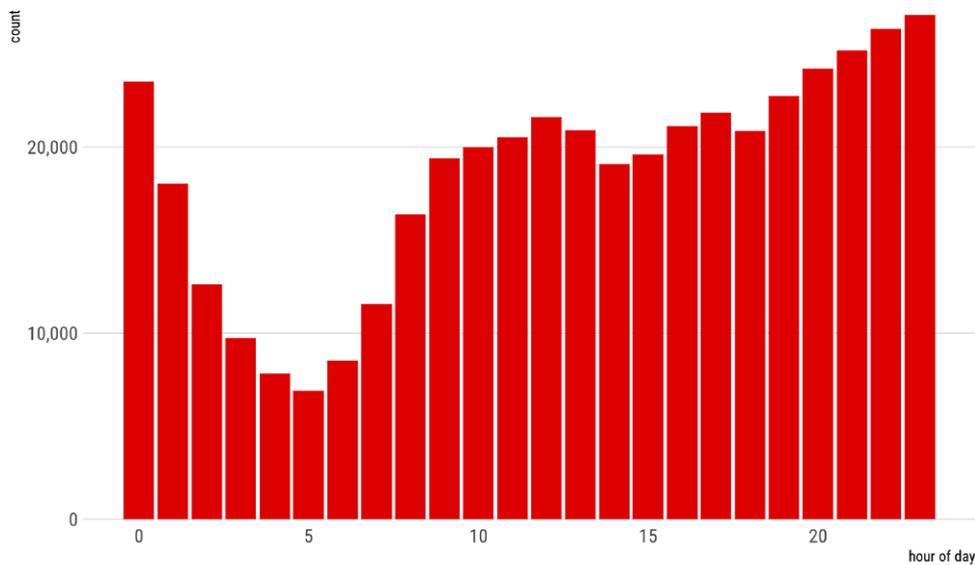


Figure 6: Tweets by the hour of the day (Beijing local time).

4. Account tweet languages

The vast majority of these accounts tweeted in Chinese (38.9%) or English (38%), followed by Japanese (5%) and Urdu (2.7%) (excluding tweets where the language was labelled “undetermined”).

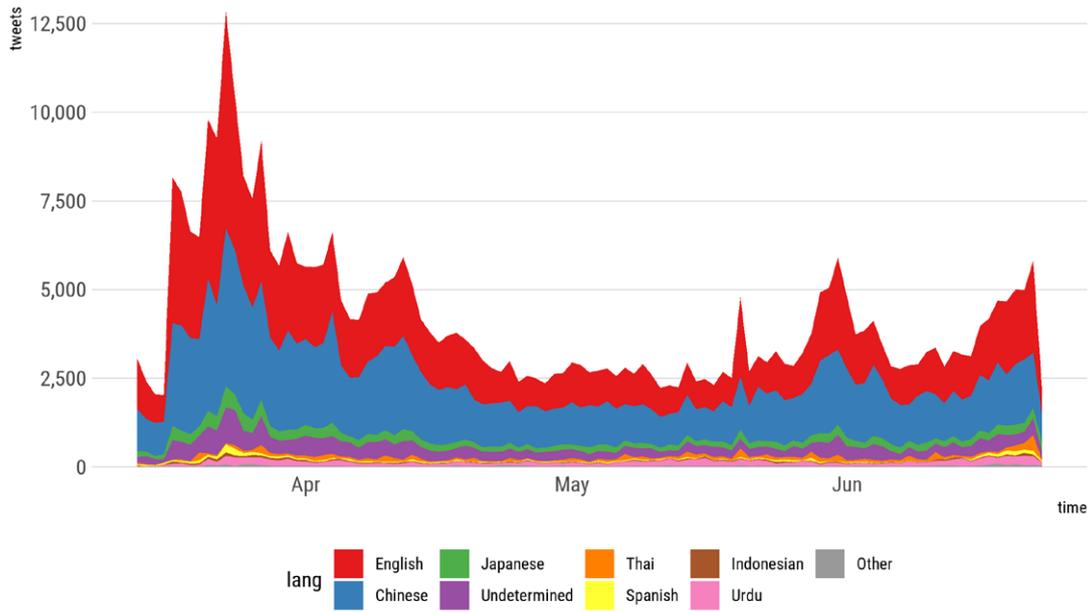


Figure 7: Number of accounts tweeting in different languages.

The majority of accounts we sampled sent tweets using between three and five languages and nearly 2,000 accounts tweeted using five or more languages. Fig. 8 shows the distribution of accounts and the number of languages used by those accounts.

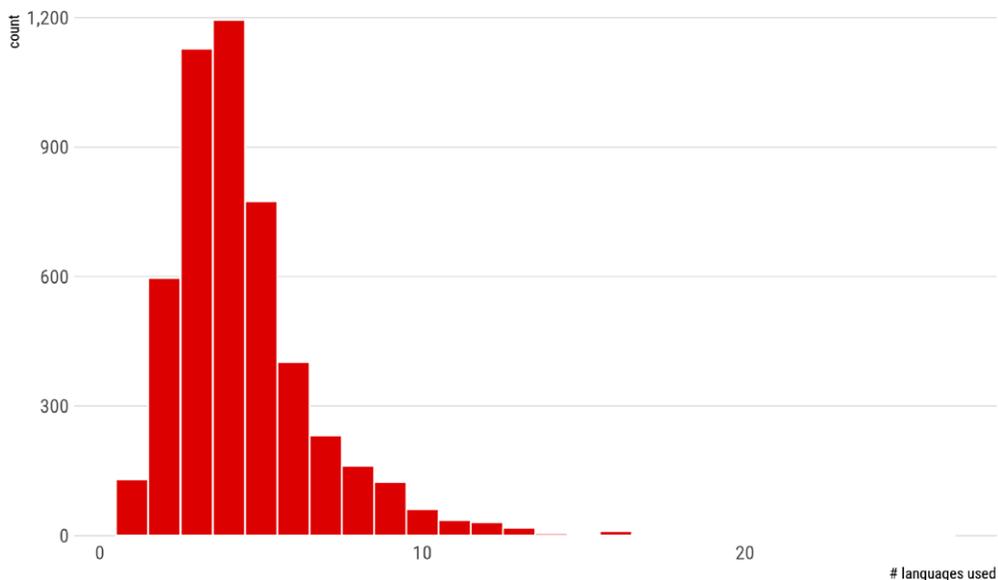


Figure 8: Distribution of accounts and the number of languages used.

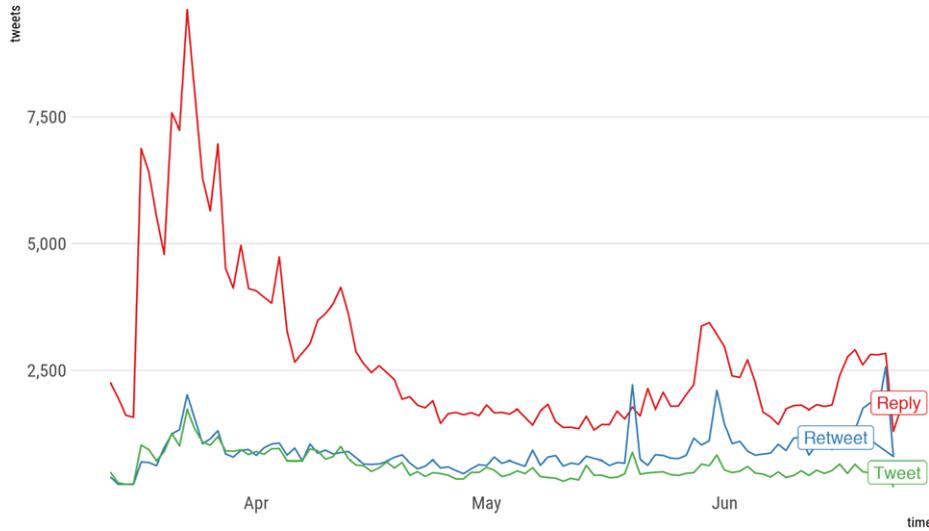


Figure 9: Timeline and types of tweets.

5. Timeline and types of tweets

Of the 445,570 tweets that ISD collected, researchers found that replies made up nearly 65% of tweets, with the largest number of replies coming after President Trump’s use of the phrase, “China Virus”. The number of replies, retweets, and original tweets decreased thereafter, but replies continued to make up a larger proportion of the tweet types.

Replies primarily targeted President Trump, American diplomatic officials, Taiwan’s president, and major media such as the Voice of America, the BBC, Deutsche Welle, the New York Times, and others. The accounts also replied to the official accounts of the CCP and Chinese state-controlled media, largely with praise or in agreement with statements they had made regarding COVID-19 or rebuffs to President Trump.

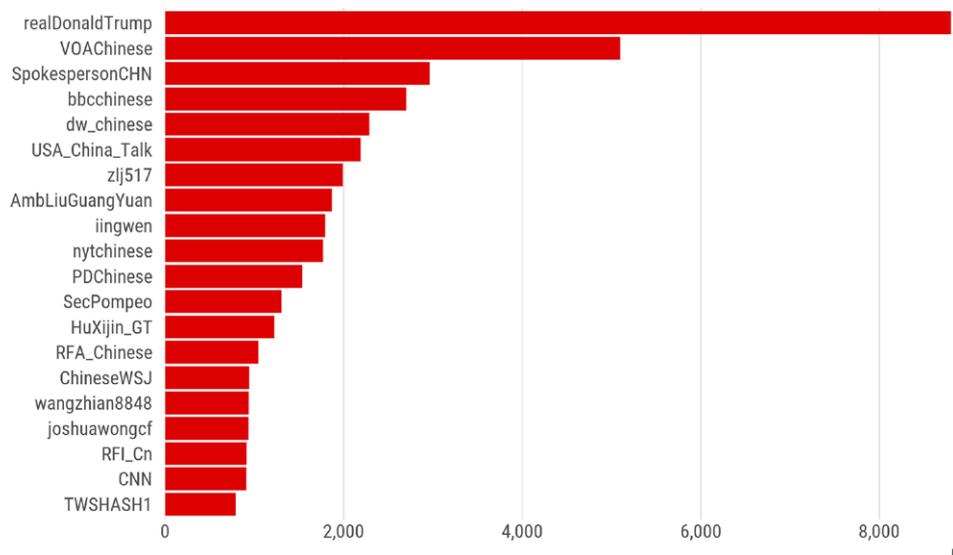


Figure 10: Top 20 Twitter accounts targeted and volume of replies

Replies also often targeted multiple accounts using similar narratives or identical text, as shown in Fig. 11, which suggests coordinated messaging.



Figure 11: Examples of duplicate messaging.

The network graph in Fig. 12 shows how accounts are clustered depending on who targeted them with replies.¹¹ The blue cluster, on the left, consists of major English-language media, whereas the light orange cluster, on the right, consists of Chinese state-controlled media. The green cluster, at the bottom, contains the Chinese-language accounts of English media, as well as Chinese state-controlled media, but it also contains others, such as the account of Taiwan’s president and the account of the US embassy and consulate in China. The top-left cluster, in orange, contains US officials, including President Trump, the official White House account, and secretary of state Pompeo. These clusters suggest that specific groups of CCP followers directed their replies at specific targets.

¹¹ We constructed the graph from a sample of 445,000 tweets, extracting the handles of CCP followers and the accounts to which they replied, filtering out any accounts that sent less than five replies. We additionally removed any accounts that had been replied to less than 50 times. Finally, we determined which accounts tend to be replied to – or targeted – by the same groups of CCP followers by calculating their pairwise correlation.

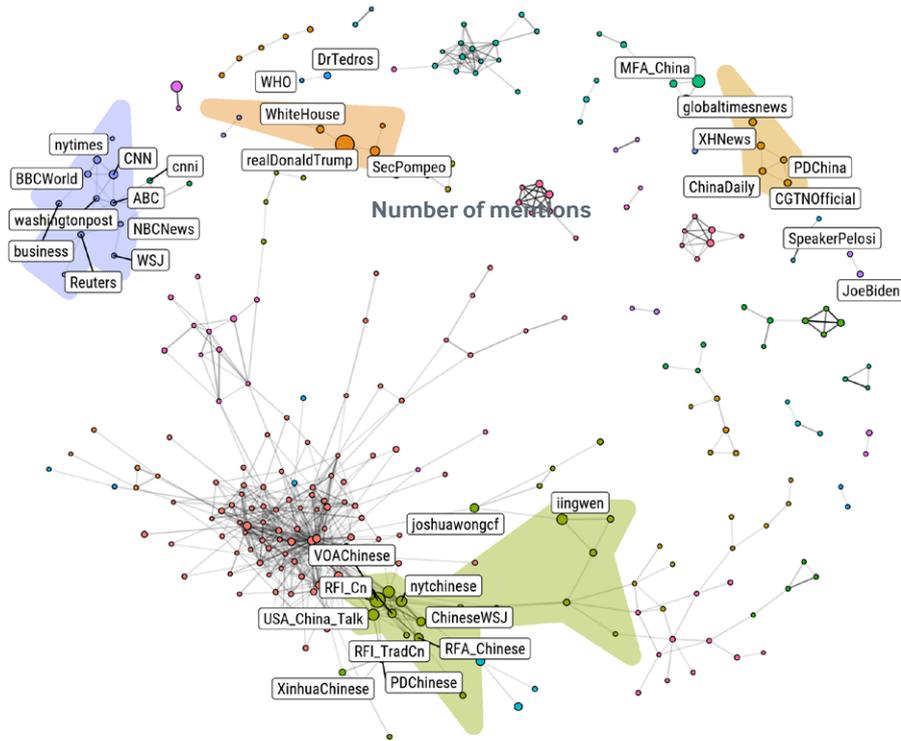


Figure 12: Accounts are clustered depending on who targeted them with replies.

Figure 13 shows the creation times of a sample of the accounts involved in reply targeting, with each account represented by a point coloured by the type of account it targeted and sized to total tweets it sent. The accounts are clustered by their creation dates on 13 March, 17 March, 21 and 23 March.

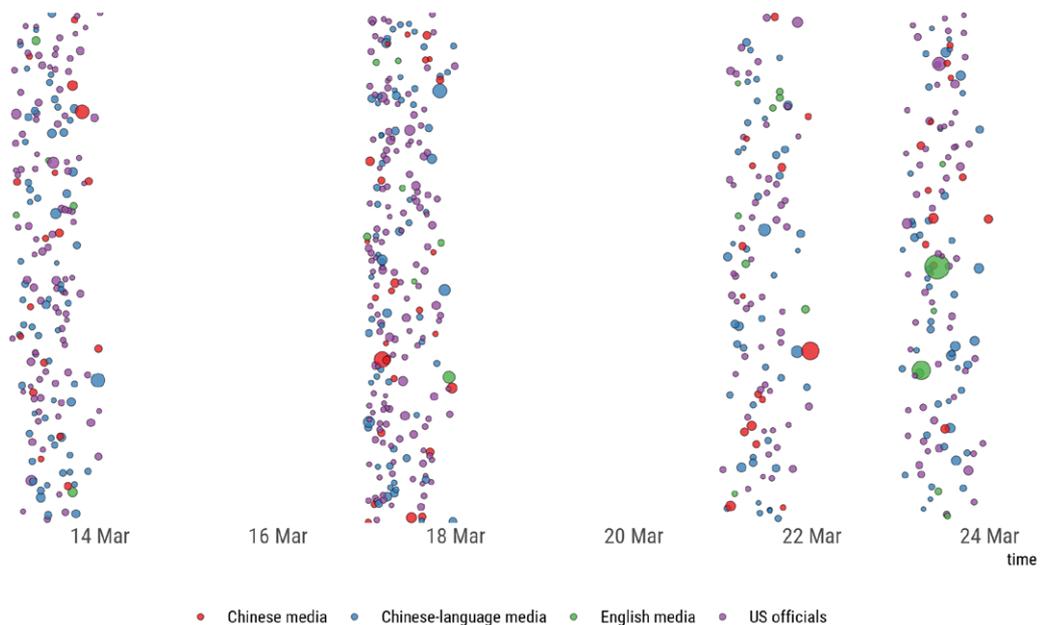


Figure 13: Creation times of accounts involved in reply targeting.

ASD’s data also indicates an effort by official CCP accounts to target key U.S. officials. Between 1 January and 31 July 2020, President Trump and U.S. Secretary of State Mike Pompeo’s accounts were, respectively, the fourth and sixth most mentioned accounts by CCP government officials and state media.¹² The predominance of coronavirus as a topic of conversation during this period likely explains the large number of mentions of @WHO and @DrTedros from the network of official CCP accounts (Figure 14).

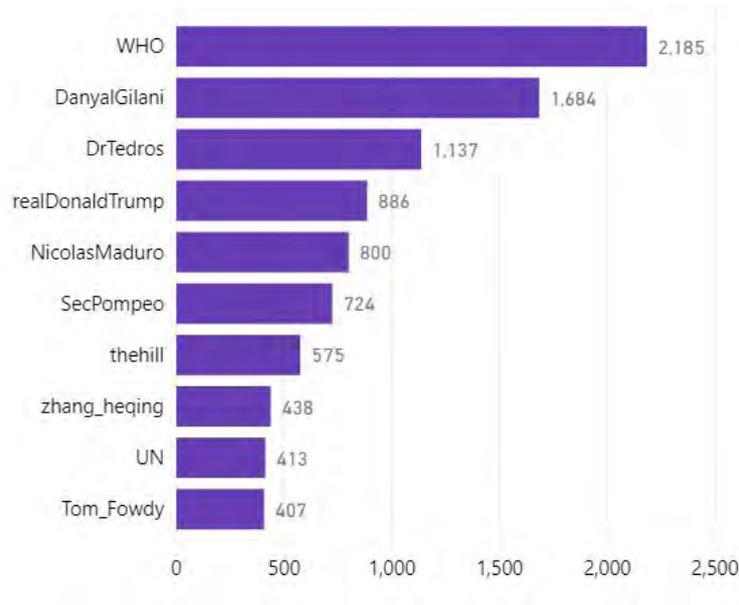


Figure 14: Most mentioned accounts by CCP government officials and state media tracked by ASD’s Hamilton dashboard between 1 January and 31 July, 2020.

5. Account handles with a high number of digits

ISD found that over 150,000 follower accounts contained handles with an unusually high number of digits; accounts with exactly eight digits in their handles are among the highest (Figure 15). There are significant peaks since 2017 when accounts with eight digits were created, but the vast majority were created after March 2020 (Figure 16). Beginning in August 2018, Twitter added a sequence of digits to handles at sign up, requiring users to later change the handle in their settings panel.¹³

¹² ASD filters out CCP accounts monitored on the dashboard from the list of accounts that are replied to.

¹³ For examples of Twitter’s changing onboarding process, see Luca Hammer: <https://twitter.com/luca/status/956916622857572353>, <https://twitter.com/luca/status/1035548732685778944>.

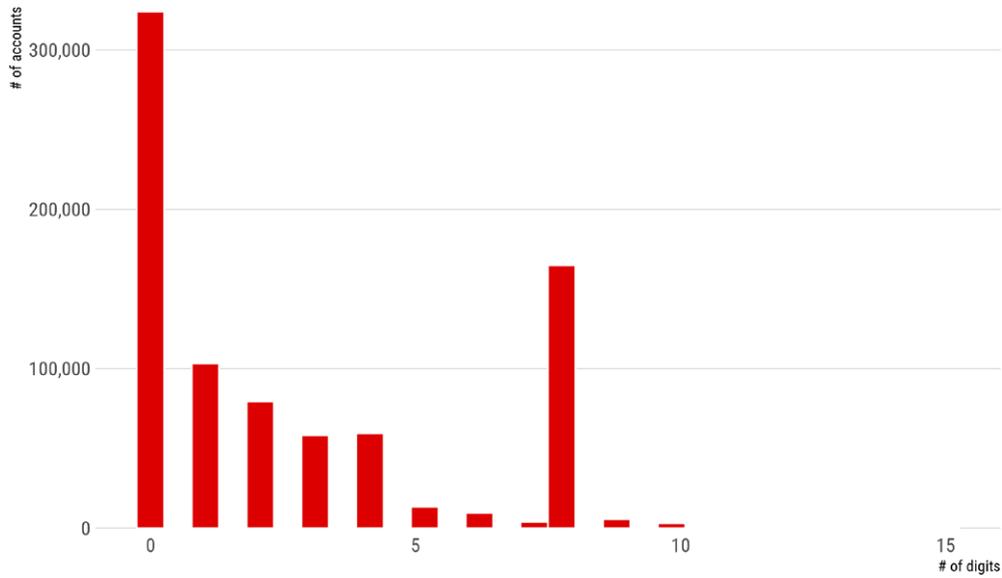


Figure 15: Distribution of accounts with the number of digits in their screen names.

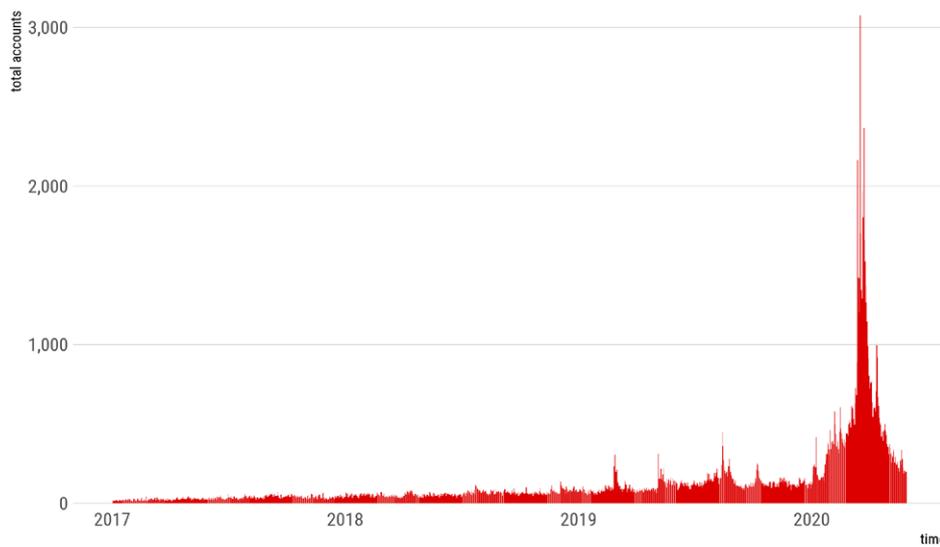


Figure 16: Daily Twitter accounts created with eight digits in their handles.

Table 1 illustrates a sample of the naming conventions of some of these accounts. Some accounts are made up of a generic name followed by a sequence of numbers, whereas others are an alphanumeric combination, and others obeyed patterns, where specific characters remain consistent and there are only changes in the sequence of numbers. These accounts are also usually clustered by similar blocks of creation dates.

Screen name	Number of digits	Creation date
Michael [redacted] 03955	8	5/22/20 19:53
Michael [redacted] 11640	8	5/28/20 4:07
Michael [redacted] 65958	8	5/26/20 11:35
Michael [redacted] 61395	8	5/26/20 3:49
Michael [redacted] 16301	8	5/26/20 7:35
Michael [redacted] 55113	8	5/25/20 9:50
Alicezh [redacted] 30180	8	4/22/20 11:10
AliceLi [redacted] 87838	8	4/22/20 4:00
Alice [redacted] 79129	8	4/22/20 15:35
Alice [redacted] 12444	8	4/25/20 16:46
Alice [redacted] 31344	8	4/25/20 13:12
Alice [redacted] 94758	8	4/24/20 3:30
980631 [redacted] qqcom1	10	3/21/20 16:20
918608 [redacted] qqcom1	10	3/27/20 14:49
710667 [redacted] qqcom1	10	3/24/20 18:56
408381 [redacted] qqcom1	10	3/23/20 11:28
278835 [redacted] qqcom1	10	3/24/20 9:16
185116 [redacted] qqcom1	10	3/24/20 2:23
2069475 [redacted] qqco1	11	5/24/20 10:02
1945483 [redacted] qqco1	11	5/22/20 2:53
1413213 [redacted] qqco1	11	5/26/20 4:36
1373943 [redacted] qqco1	11	5/27/20 18:22
1208777 [redacted] qqco1	11	4/21/20 4:04
1056790 [redacted] qqco1	11	4/16/20 3:29

Table 1: General naming conventions of accounts.

ASD found that several CCP officials retweeted accounts whose behavior and naming characteristics were consistent with the signals of inauthenticity identified by ISD. Among the 500 most retweeted accounts by CCP officials, ASD found several potentially suspicious accounts, including at least 10 that were suspended and at least 10 others that no longer existed. The removal of the accounts, coupled with other suspicious indicators (namely naming conventions), suggests that at least some of these accounts were removed for violations of Twitter rules. In addition, several of the deactivated accounts showed patterns of both coordination and inauthenticity.

The most glaring example was a series of 12 accounts beginning with “ddpx5w0” that were retweeted between one and 38 times by official CCP accounts. In total, the “ddpx5w0” accounts were retweeted 109 times (Table 2). By comparison, that total is more than the combined number of retweets of @Time (28 retweets), @BBCWorld (27 retweets), @nytimes (22 retweets), @realdonaldtrump (17 retweets), and @EmmanuelMacron (14 retweets).¹⁴

Screen name	Retweets by official CCP accounts
ddPx5w0 [redacted] wvkzF	38
ddpx5w [redacted] 222	14
ddpx5w [redacted] 231	13
ddpx5 [redacted] 11	10
ddpx [redacted] 1	9
ddpx5w [redacted] 211	7
ddp [redacted] 0	5
ddpx5 [redacted] 122	4
ddpx! [redacted] 1221	4
ddpx [redacted] 121	2
ddpx! [redacted] 1223	2
ddp: [redacted] 012	1
Total Retweets	109

Table2: Retweets of accounts with screen names containing “ddpx5w0” by CCP government officials and state media tracked by ASD’s Hamilton dashboard.

When ASD identified the accounts in its late July review of official CCP retweets, 11 of the 12 accounts no longer existed. The only active account (@ddpx5w012231) was restricted, but a manual review of the account found unusual activity, including posting in multiple languages, an odd follower (1,001) versus following (0) split, and an almost singular focus on boosting CCP diplomatic and state media accounts and narratives.



Figure 17: Account identified as part of a network of accounts with “ddpx5w0” in the screen name.

¹⁴ Figures accurate as of 18 July, 2020. Note that quote tweets – a retweet with added commentary – are typically logged as original tweets and are therefore not included in retweet data.

The account’s bio suggested support for LGBT issues, “black people,” and human rights. ASD noted variations on these themes in the bios of several suspicious accounts retweeted by CCP officials and state media outlets. A reverse image search of the profile photo uncovered hundreds of other social media accounts on Twitter and other platforms using the same photo. As of 31 July, the @ddpx5w012231 account had been suspended by Twitter.



Figure 18: Social media accounts using the same profile photo as Twitter account @ddpx5w012231

It is certainly possible that CCP officials amplified these accounts based on the content of their tweets rather than as part of a coordinated effort to manipulate or manufacture support for CCP viewpoints. At the same time, the fact that clearly suspicious accounts with few followers managed to earn repeated engagement from CCP ambassadors and consul generals from Venezuela to South Africa suggests, if nothing else, a concerning lack of critical digital literacy from such officials.

7. Narrative Analysis

ISD uncovered four primary themes in the tweets: COVID-19, Hong Kong, Taiwan, and the recent antiracist protests in the US. Figure 19 shows the volume of activity for each narrative, with grey representing the volume of activity for all narratives combined.

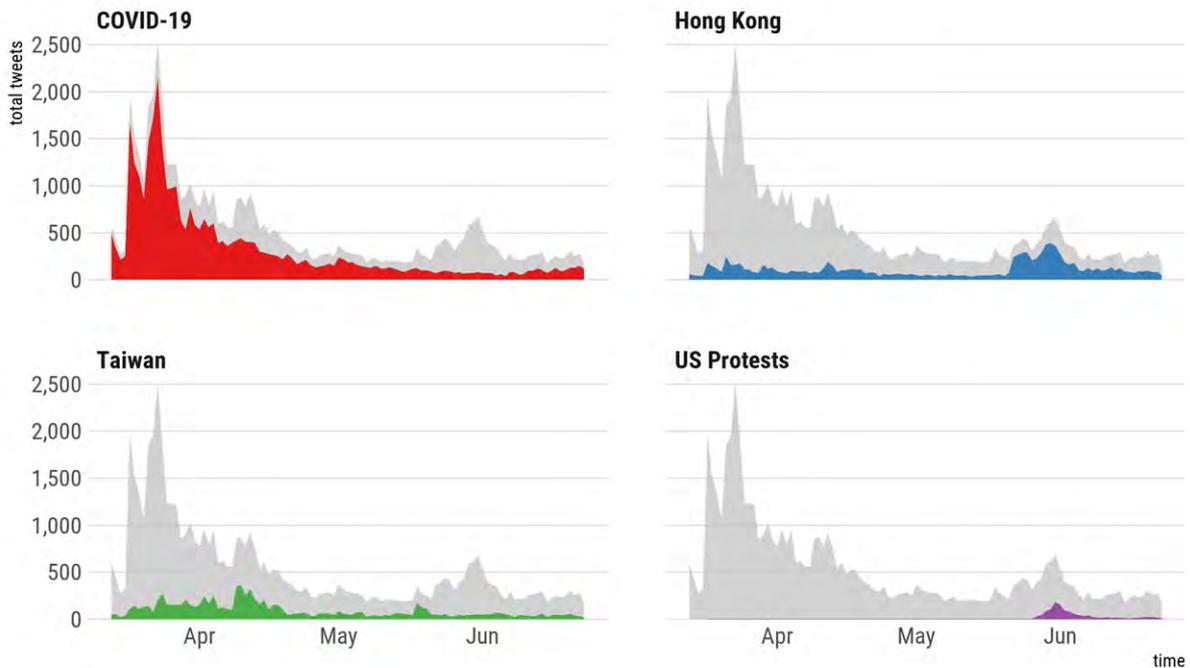


Figure 19: Volume of activity for each narrative.

Tweets about COVID-19 targeted President Trump and other accounts for calling COVID-19 a “Chinese Virus” or defended the CCP’s response to the pandemic while highlighting the US’s failure to contain it. The phrase, “American virus” was used in more than 1,000 tweets, with accounts calling Trump an “American virus” or with claims that COVID-19 had originated in the US. Alongside these tweets were references to various conspiracy theories forwarded by pro-CCP media, including 274 mentions of “Fort Detrick” in Maryland, as a “biological weapons lab” and 283 references to “e-cigarettes” and a “vaping pneumonia” reported in the US last year. Other tweets suggested that COVID-19 was similar to H1N1, which spread rapidly throughout the US when it first emerged.

These conspiracies were bolstered by YouTube videos shared in the tweets, including the most frequently shared video, “The Coronavirus CONSPIRACY - Did COVID-19 Come from America?,” published by Nathan Rich, an expatriate American who shares and promotes pro-CCP views. Two other frequently shared videos attacked the New York Times for its “Coronavirus propaganda” or “Western media” for “lies about China”.

Similarly, ASD found at least 80 different uses of the term “Fort Detrick” in its monitoring of activity surrounding official CCP accounts. Those totals do not include the initial promotion (or any retweets) of the conspiracy theory by Zhao Lijian, who linked to an article suggesting a Fort Detrick connection but did not specifically name it in his original tweet. In a subsequent July tweet, Zhao returned to the Fort Detrick narrative, this time including a reference to vaping cases in the United States (Fig. 20).



Figure 20: Tweets from Zhao Lijian in March 2020 and July 2020

Additionally, ASD found at least seventeen instances of official CCP accounts retweeting or mentioning Nathan Rich, including at least one retweet of Rich’s Coronavirus conspiracy video.

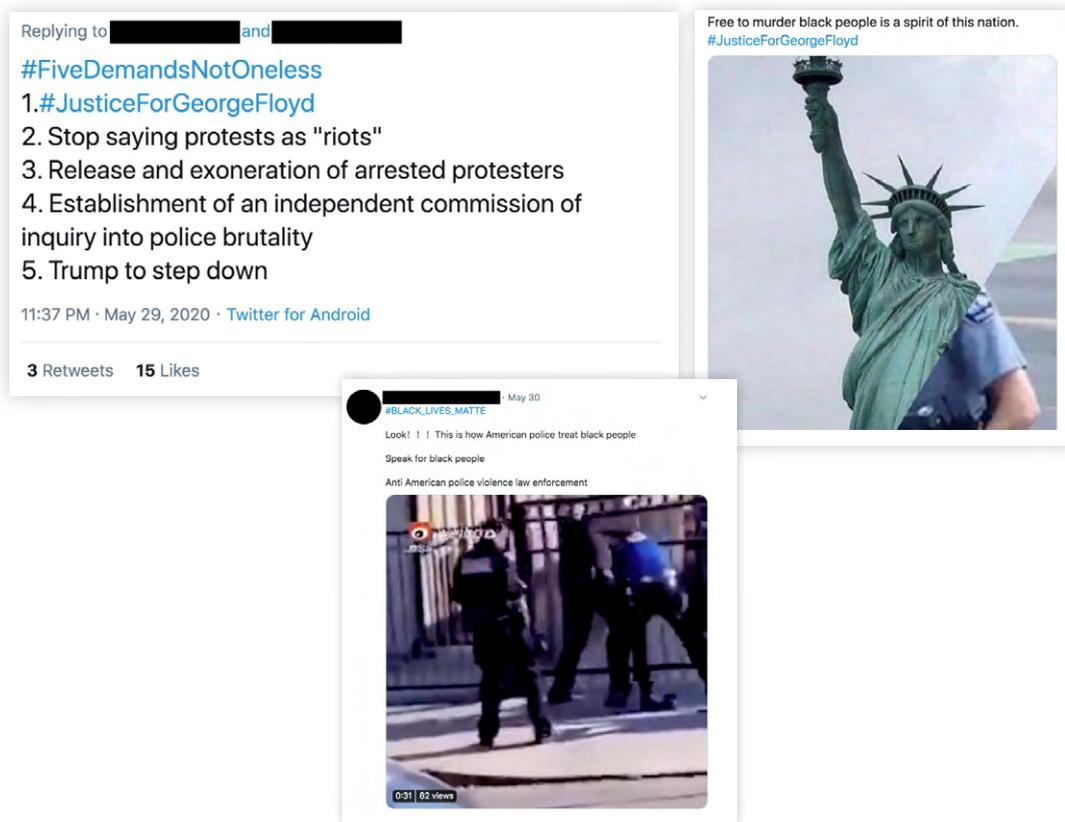


Figure 21: Tweets depicting police violence and “demands” of “protesters”

Antiracist Protests

ISD also found that the CCP followers started focusing on the demonstrations in the US, sparked by the killing of George Floyd on 25 May. The tweets within this narrative have several aims: exacerbating domestic tensions over police brutality against Black Americans; video and images illustrating the often-violent suppression of protests by police and; comparing the Trump administration’s response to the demonstrations to the protests in Hong Kong. Some accounts used, “#FiveDemandsNotOneLess,” a hashtag that has been previously used by Hong Kong protestors, modifying it to reference George Floyd’s murder and the removal of Trump. This hashtag was retweeted by at least two official CCP accounts: @CGCHINA_CPT and @jindurban1.

#minneapolisriots Why not go to the police station to make trouble, are some supermarkets cowards?

Look at this video! The white peoples rob the shop is ok, police did nothing to white people. Why? Because they are racist. #blacklivesmatternl #protests2020 #america #blacklivesmatter #racism #trump @realdonaldtrump

nypd cop pulls down peaceful protestor’s mask to pepper spray him. this video is being removed all over twitter, they are trying to hide this. <https://t.co/hf0b39wlzc>

ASD noted a similar focus on race and racism in the United States after the killing of George Floyd. Official CCP accounts used the hashtags #BlackLivesMatter, #BLM, and #GeorgeFloyd more than 500 times combined in the weeks after Floyd’s death. Ministry of Foreign Affairs spokesperson Hua Chunying also used Floyd’s dying words as a rebuttal to a U.S. State Department tweet condemning China’s actions in Hong Kong (Fig. 22).



Figure 22: Tweet from Ministry of Foreign Affairs spokesperson Hua Chunying in May 2020.

Taiwan

ISD found 10,672 references to Taiwan in both English and Chinese. Tweets referencing Taiwan dismissed its bid to join the WHO, asserted that it remains a province of China, attacked Taiwan’s development, or accused the country of being a lackey of the US. Tweets often replied to President Tsai Ing-wen, Taiwan’s ministry of foreign affairs, or media accounts reporting about Taiwan.

ASD found 3,633 mentions of Taiwan from government and state-backed media accounts, making it the seventh most-referenced country, region, or special territory in its dataset (Fig. 23).¹⁵ Overt narratives about Taiwan broadly coincided with those noted above, with particular animus directed at the suggestion of Taiwanese independence (Fig. 24).

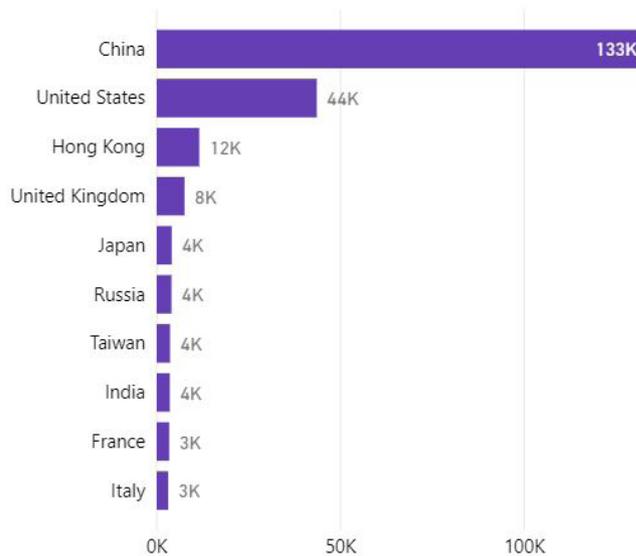


Figure 23: Most mentioned countries by CCP government officials and state media tracked by ASD’s Hamilton dashboard between 1 January and 31 July, 2020.

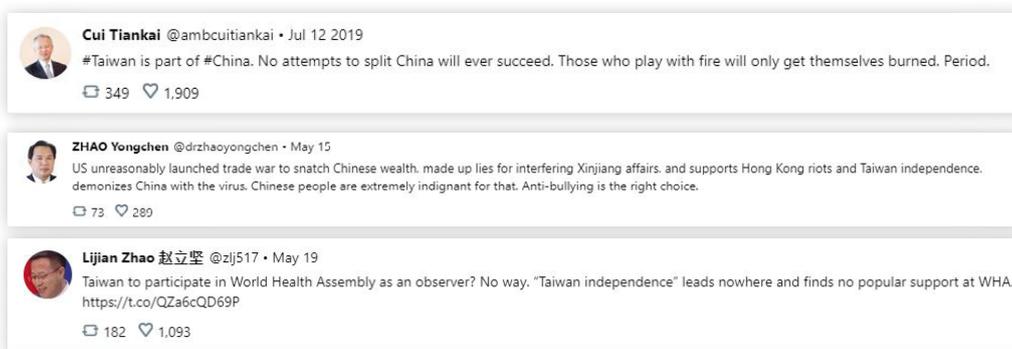


Figure 24: Examples from ASD’s CCP dataset of CCP officials’ Tweets about Taiwan.

¹⁵ ASD uses key phrase extraction to determine the countries most mentioned in tweets. Special territories and regions (e.g., Taiwan, Hong Kong, and the European Union) are tabulated as unique entities. Alternative country names are included under the official country name (e.g., U.S., USA, and America are all included under the United States).

Hong Kong

There were 11,185 references to Hong Kong in English and Chinese, with a large proportion of tweets characterising Hong Kong protesters as violent or likening the protests to those in the US as a way to highlight US hypocrisy.

Secretary Pompeo, please stand with the angry people of Minneapolis, just like you did with the thugs of Hong Kong.

When Hong Kong rioters used catapults to lob Molotov cocktails, US media like CNN glorified it and gleefully shared videos. Let's hope that #riots2020 don't use this strategy, for example, to attack Trump's White House. #BlackLivesMatter #WashingtonDC #protests

That's what is happening in Hong Kong, but you called them "Protestors"! Lol double standard and contradicting huh? Cmon Trump, these people are your people, they are the "Protestors", NOT rioters! You can't shoot innocent people just because they wanna justice!

There were 11,584 mentions of Hong Kong in ASD's dataset. It was the third most referenced country, region, or territory, behind only the United States and China. ASD found 98 examples of the term "double standard" and 29 examples of the term "hypocrisy" in tweets related to Hong Kong (Fig. 25).

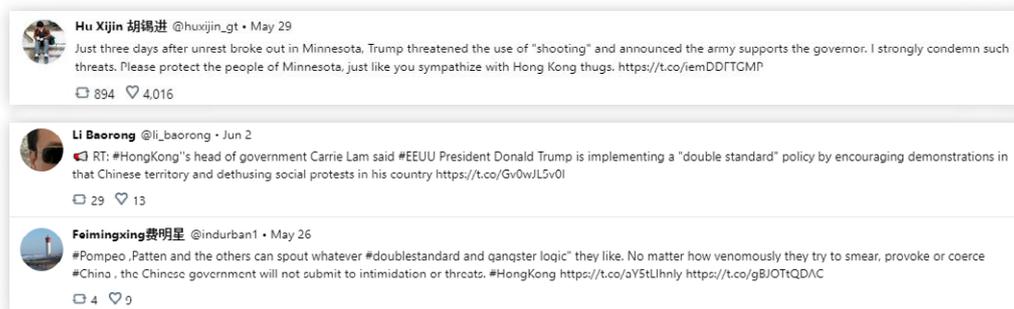


Figure 25: Example tweets from ASD's dataset of CCP officials' tweets about Hong Kong.

Methodological Note

ISD used a variety of signals in an effort to examine authenticity and coordination among CCP followers on Twitter. Some of the signals we used have been utilized in other investigations or disinformation-related literature, with researchers using high numbers of digits in account handles,¹⁶ account creation dates,¹⁷ narrative analysis, retweet networks,¹⁸ and other coordinated patterns (co-tweeting, co-retweeting, and co-mentioning)¹⁹ to identify potentially inauthentic networks of accounts.

We also note that some features, like bursts of account creation, continue to be a useful clue for investigating information operations, as Twitter’s own data illustrate (Fig. 17). Principals often need to create new accounts to respond to their objectives in real-time and where they use commercial or repurposed accounts, they are easily uncovered because of what amounts to weak tradecraft.

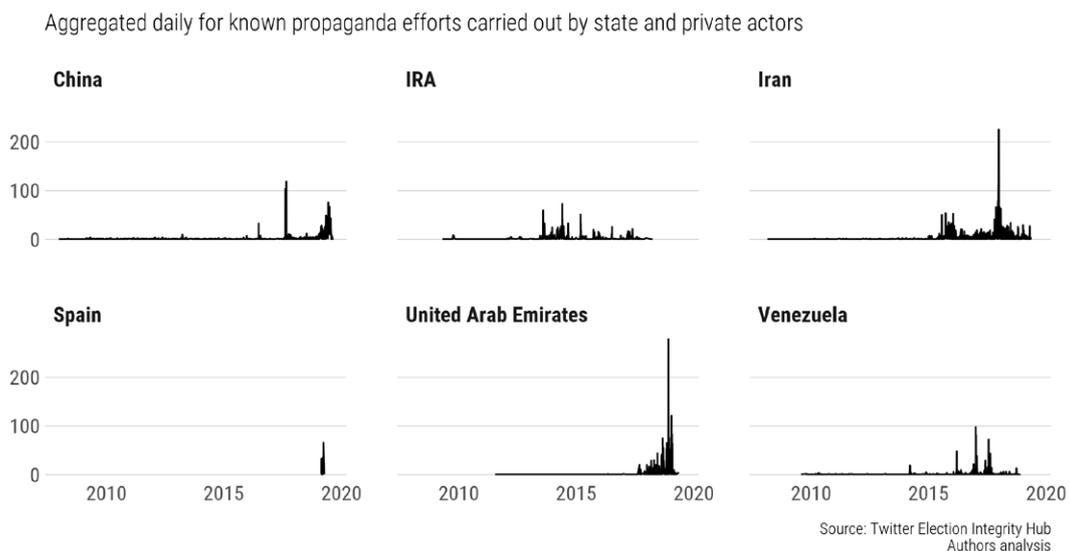


Figure 26: Creation dates of Twitter accounts for known information operations

But these signals are not definitive. Recent research shows that coordination patterns resembling those of information operations can manifest even in benign activity, particularly that found in political communities.²⁰ Given this finding, we compare data from our investigation of CCP followers with the activity from two “baselines” of political and non-political users on Twitter, analysing their features for similar evidence of coordination and inauthenticity. We follow the reasoning used by other researchers comparing known information operations against legitimate Twitter activity.²¹

¹⁶ Charlotte Godart and Johanna Wild, “Case study: Finding evidence of automated Twitter activity during the Hong Kong protests,” Verification Handbook 3: Investigating Disinformation and Media Manipulation. Accessed at: <https://datajournalism.com/read/handbook/verification-3/investigating-actors-content/3a-case-study-finding-evidence-of-automated-twitter-activity-during-the-hong-kong-protests>

¹⁷ Franziska B. Keller, David Schoch, Sebastian Stier and JungHwan Yang, “Political Astroturfing on Twitter: How to Coordinate a Disinformation Campaign, Political Communication,” 37:2, 256-280, Accessed at: <https://www.tandfonline.com/doi/abs/10.1080/10584609.2019.1661888>

¹⁸ Stewart, L. G., Arif, A. & Starbird, K. “Examining trolls and polarization with a retweet network.” In Proceedings ACM WSDM, Workshop on Misinformation and Misbehavior Mining on the Web (2018). Accessed at: <https://faculty.washington.edu/kstarbi/examining-trolls-polarization.pdf>

¹⁹ Diogo Pacheco, Pik-Mai Hui, Christopher Torres-Lugo, Bao Tran Truong, Alessandro Flammini, and Filippo Menczer, “Uncovering Coordinated Networks on Social Media,” [Preprint], 16 January 2020 Accessed at: <https://arxiv.org/abs/2001.05658>

²⁰ Luis Vargas, Patrick Emami, and Patrick Traynor, “On the Detection of Disinformation Campaign Activity with Network Analysis,” [Preprint], 27 May 2020. Accessed at: <https://arxiv.org/abs/2005.13466>

²¹ Vargas, Emami, and Traynor 2020.

For the political baseline, we used Twitter’s standard Search API to collect 8,958 tweets for a recent trending hashtag, #KeepTexasRed. From that data, we identified 6,534 unique accounts. For the non-political baseline, we retrieved 1,000 accounts for a hashtag associated with the R programming community, #rstats. From each baseline set, we took a random sample of 500 accounts and used Twitter’s API to collect the most recent 600 tweets from each account’s timeline.

Account set	Accounts ²²	Tweets	Retweets (%)	Replies (%)
CCP Followers	4,922	445,570	21%	64%
Non-political (#rstats)	1,000	31,687	41%	40%
Political (#KeepTexasRed)	6,534	169,907	70%	19%

Table 3: Tweet types in the CCP followers, #rstats and #KeepTexasRed datasets

Timelines and tweet types

The CCP Followers set is characterized by much higher volumes of replies, whereas the #KeepTexasRed set has higher levels of retweets. The tweets from #rstats users are almost equal in distribution. The #KeepTexasRed set is highly influenced by politically active accounts retweeting content from President Donald Trump and other conservative figures.

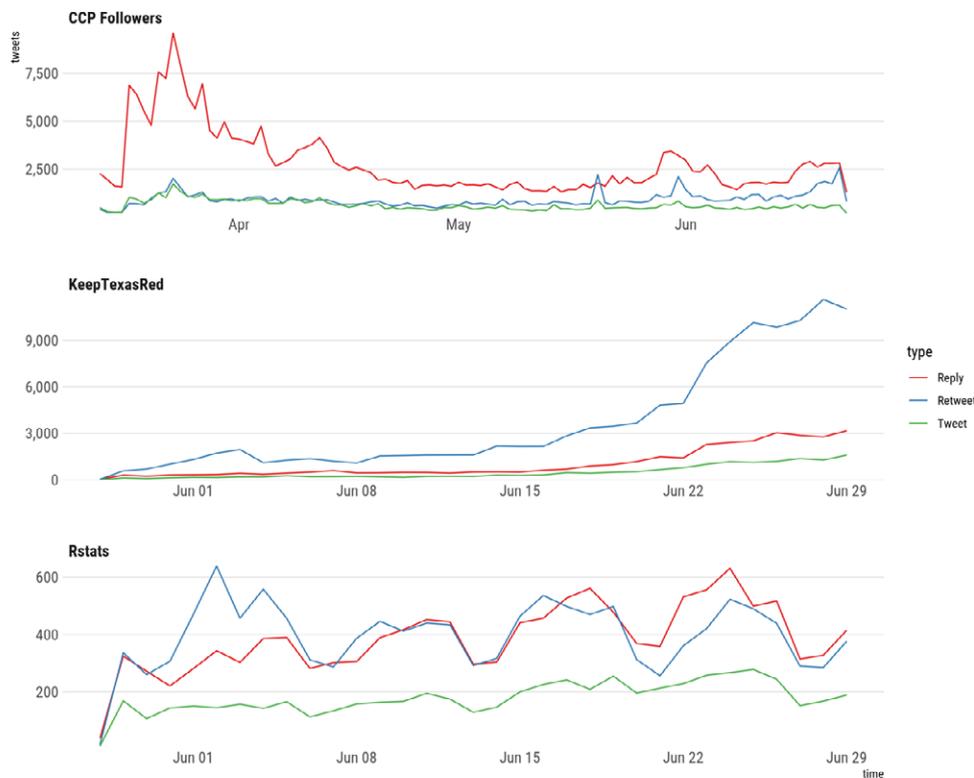


Figure 27: Time-series of tweet types for each set

²² Account numbers refer to total accounts obtained from either searching for users whose profile description contained the query or tweets containing the query.

Handles

We also compared the sets' distribution of handles and their number of digits. The CCP follower set has the highest proportion of 8-digit handles (32%) compared to other sets. Indeed, no accounts are found to have 8-digit handles in the #rstats set and only 6% contain them in the #KeepTexasRed set.

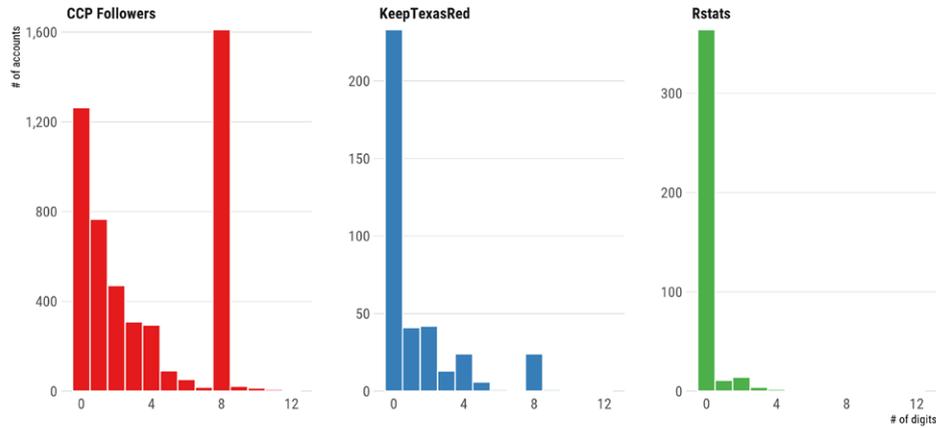


Figure 29: Distribution of handles with digits for all sets

Co-Replying

Finally, we examined co-replying in each set, or when two accounts sent a reply to the same user within a specific time. Fig. 20 shows that the CCP set appears to have more prevalent coordination, with the #KeepTexasRed's coordination patterns increasing only slightly as the hashtag trended.

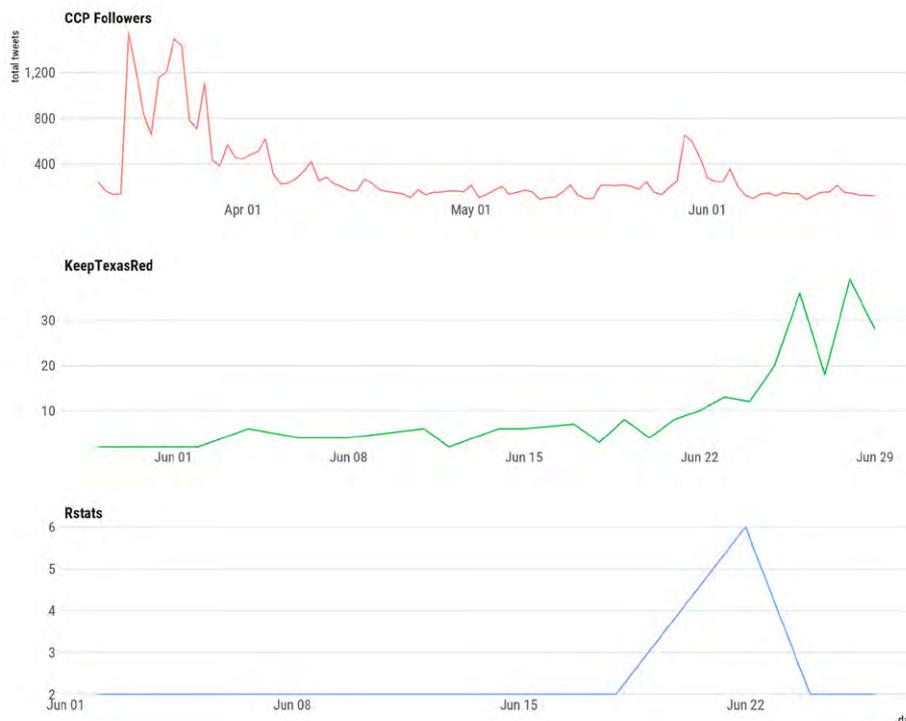


Figure 29: Daily coordination based on co-replying

Conclusion

ISD's analysis suggests that a network of potentially inauthentic accounts appear to be actively promoting the narratives of the CCP. While we cannot attribute this activity to a state (or non-state) actor at this time, previous reporting and disclosures have implicated the PRC, patriotic trolls, and government-contracted social media agencies in recent campaigns supporting CCP interests.²³ That, coupled with ASD's findings that CCP diplomats have repeatedly amplified the messaging of now-suspended and seemingly inauthentic accounts, demonstrates the importance of understanding the overt activities of states in parallel with covert pro-state networks online.

ISD's findings rely on indicators widely used by researchers attempting to establish inauthenticity and coordination online, including high volumes of account activity, creation dates, ambiguous or reused profile photos, account handle patterns, and others. Some indicators, such as handle patterns, should be treated with caution, particularly when they can result from platform changes to the onboarding process. Other indicators, such as account creation, may be more robust, as repeated information operations illustrate how actors continue to use new accounts to respond to real-world events as they emerge.

ISD's analysis also illustrates how messaging tactics – namely, targeted replying – can be used to propagate disinformation and foment political unrest online. The protests in the US have provided ample space for outside actors to amplify internal division and while recent PRC-backed campaigns have also targeted the US protests, the messaging has primarily targeted Chinese-speaking audiences in relation to the Hong Kong protests.²⁴ The tweets in our analysis suggest that pro-CCP campaigns are also increasingly targeting English-speaking audiences, advancing additional narratives whose purpose seems to be to stoke further unrest and polarisation in the US.

In this case study, the number of signals of inauthenticity ISD required to gather a clear impression of network coordination proves the complexity of research approaches that are required to detect contemporary information operations. Simplistic assessments of inauthenticity or coordination can often sensationalise the scale and nature of the disinformation problem, distracting from the very real problems that exist online and the difficult conversations needed to find solutions.

Finally, the broadly similar narratives promoted by the pro-CCP accounts examined by ISD and the official CCP accounts monitored by ASD, points to the need for further research into the interplay between covert pro-CCP networks and overt CCP information activity online.

The case study also demonstrates the critical value of access to social media data for researchers, which Twitter has long understood and enabled. The research sector's understanding of tactics and strategies used by disinformation perpetrators on Twitter is far more developed than parallel efforts mapping activities on Facebook and YouTube, as a result. Other social media companies must provide secure access to data to allow similar studies to take place across the ecosystem of information operations: these activities do not take place in a Twitter vacuum.

²³ Australian Strategic Policy Institute, COVID-19 Attracts Patriotic Troll Campaigns in Support of China's Geopolitical Interests (23 April 2020), <https://www.aspi.org.au/report/covid-19-disinformation>; ProPublica, How China Built a Twitter Propaganda Machine Then Let It Loose on Coronavirus (26 March 2020), <https://www.propublica.org/article/how-china-built-a-twitter-propaganda-machine-then-let-it-loose-on-coronavirus>; Twitter Safety, Disclosing networks of state-linked information operations we've removed (12 June 2020), https://blog.twitter.com/en_us/topics/company/2020/information-operations-june-2020.html. [Accessed on 30 June 2020].

²⁴ Australian Strategic Policy Institute, Retweeting through the great firewall: A persistent and undeterred threat actor (12 June 2020), <https://www.aspi.org.au/report/retweeting-through-great-firewall>.

About The Alliance for Securing Democracy (ASD)

A bipartisan initiative housed at the German Marshall Fund of the United States, develops comprehensive strategies to deter, defend against, and raise the costs on authoritarian efforts to undermine and interfere in democratic institutions. ASD brings together experts on disinformation, malign finance, emerging technologies, elections integrity, economic coercion, and cybersecurity, as well as regional experts, to collaborate across traditional stovepipes and develop cross-cutting frameworks.

About the Institute for Strategic Dialogue

We are a global team of data analysts, researchers, innovators, policy-experts, practitioners and activists - powering solutions to extremism, hate and polarisation.

The Institute for Strategic Dialogue (ISD) is an independent nonprofit organisation dedicated to safeguarding human rights and reversing the rising global tide of hate, extremism and polarisation.

We combine sector-leading expertise in global extremist movements with advanced digital analysis of disinformation and weaponised hate to deliver innovative, tailor-made policy and operational responses to these threats.

ISD draws on fifteen years of anthropological research, leading expertise in global extremist movements, state-of-the-art digital analysis and a track record of trust and delivery in over 40 countries around the world to:

- Support central and local governments in designing and delivering evidence-based policies and programmes in response to hate, extremism, terrorism, polarisation and disinformation
- Empower youth, practitioners and community influencers through innovative education, technology and communications programmes.
- Advise governments and tech companies on policies and strategies to mitigate the online harms we face today and achieve a 'Good Web' that reflects our liberal democratic values

Only in collaboration with all of these groups can we hope to outcompete the extremist mobilization of our time and build safe, free and resilient societies for generations to come.

All of ISD's programmes are delivered with the support of donations and grants. We have the data on what works. We now need your help to scale our efforts.

Support our work: isdglobal.org/donate

General

info@isdglobal.org

Media

media@isdglobal.org

Support

getinvolved@isdglobal.org

ISD

Powering solutions
to extremism
and polarisation



alliance for
securing
democracy

PO Box 75769 | London | SW1P 9ER | UK
www.isdglobal.org