

# Tweet dreams

## Looking for evidence of Chinese Government influence on Twitter

---

*Quantitative and qualitative analysis was conducted on large samples of Twitter data collected following two points of tension in the Australia-China relationship in 2020 – Australia’s call for an independent investigation into the origins of COVID-19, and a Chinese Government representative’s retweet of an image of an Australian soldier killing an Afghan child. There was no evidence of a major Chinese Government-backed campaign to influence discussion.*

### Authors

Daniel Whelan-Shamy, Digital Media Research Centre (DMRC),  
Queensland University of Technology

Timothy Graham, DMRC

Daniel Angus, DMRC

Rod Campbell, Australia Institute

March 2022

## About The Australia Institute

The Australia Institute is an independent public policy think tank based in Canberra. It is funded by donations from philanthropic trusts and individuals and commissioned research. We barrack for ideas, not political parties or candidates. Since its launch in 1994, the Institute has carried out highly influential research on a broad range of economic, social and environmental issues.

## About the Centre for Responsible Technology

The Australia Institute established the Centre for Responsible Technology to give people greater influence over the way technology is rapidly changing our world. The Centre will collaborate with academics, activists, civil society and business to shape policy and practice around network technology by raising public awareness about the broader impacts and implications of data-driven change and advocating policies that promote the common good.

## Our philosophy

As we begin the 21st century, new dilemmas confront our society and our planet. Unprecedented levels of consumption co-exist with extreme poverty. Through new technology we are more connected than we have ever been, yet civic engagement is declining. Environmental neglect continues despite heightened ecological awareness. A better balance is urgently needed.

The Australia Institute's directors, staff and supporters represent a broad range of views and priorities. What unites us is a belief that through a combination of research and creativity we can promote new solutions and ways of thinking.

## Our purpose - 'Research that matters'

The Institute publishes research that contributes to a more just, sustainable and peaceful society. Our goal is to gather, interpret and communicate evidence in order to both diagnose the problems we face and propose new solutions to tackle them.

The Institute is wholly independent and not affiliated with any other organisation. Donations to its Research Fund are tax deductible for the donor. Anyone wishing to donate can do so via the website at <https://www.tai.org.au> or by calling the Institute on 02 6130 0530. Our secure and user-friendly website allows donors to make either one-off or regular monthly donations and we encourage everyone who can to donate in this way as it assists our research in the most significant manner.

Level 1, Endeavour House, 1 Franklin St  
Canberra, ACT 2601  
Tel: (02) 61300530  
Email: [mail@australiainstitute.org.au](mailto:mail@australiainstitute.org.au)  
Website: [www.australiainstitute.org.au](http://www.australiainstitute.org.au)  
ISSN: 1836-9014

# Summary

State-backed influence on social media has become an important issue around the globe, while in Australia rarely a week goes by without headlines relating to Chinese Government influence. This report combines these issues to ask whether there is evidence of Chinese state-backed influence campaigns on Australia's social media platform of choice for political discussion – Twitter.

We examine Twitter datasets from two flashpoints in the Australia-China relationship – Australia's call for an independent inquiry into the origins of COVID-19 and a Chinese Government spokesman's retweet of an image depicting an Australian soldier killing an Afghan child. If there was a Chinese Government-backed campaign to influence Australian social media users, these events would have been likely targets for it.

## **Call for independent inquiry into origins of COVID-19**

We analysed 54,441 tweets sent by 36,968 accounts during June 2020 that included "Australia" or "#Auspol" and "China" or "racism". While previous studies have highlighted the importance of reply activity in disinformation campaigns, the top 100 most frequent replying accounts in this sample sent just 947 reply tweets, with 127 coming from an automated account posting anti-left content of little relevance to China. Most of the top 100 repliers sent less than 10 replies, suggesting there was limited automation or co-ordinated activity within the sample.

The content of the top 100 replying accounts was examined using structural topic modelling, identifying 20 topics including both pro and anti-China examples. Over a third of the reply tweets were on anti-CCP topics, while a quarter were on pro-China topics with the remainder spread across various topics including trade, anti-US, and 'other'.

In assessing whether some of the pro-China reply accounts might have been linked to the Chinese state, further analysis was based on account creation date, tweets since creation, engagement metrics and time when tweets were posted. This analysis yielded little evidence of suspicious activity, with no clear patterns of activity or posting at particular times of day.

The connected component extracted from the reply network had two relevant clusters, the main one around the Chinese state-owned Global Times News. Overall, we find that replies @globaltimesnews have a mixture of pro and anti-China categories, including spam-like behaviour, resulting in a messy information environment with multiple, often oppositional, perspectives and narratives all competing for attention in the reply threads.

The second relevant cluster is based on reply activity around the topic of racism in Australia. It includes several anonymous pro-China accounts that post a real and historically-rooted issue (genocide of Indigenous Australians) to magnify tensions and push a particular narrative. Many have a particular focus on #AboriginalLivesMatter and the larger global #BlackLivesMatter movement. It is difficult to ascertain whether these accounts are backed by the Chinese state. Given the relatively small number of them, it seems unlikely. One responded to a direct message, stating that they did work for the Chinese Government, but that their Twitter account expressed their own views and was not linked to their work. It seems more likely that most of these accounts are simply Chinese people with generally pro-Chinese Government views.

We also conducted co-ordination network analysis on this sample, finding little suspicious or inauthentic activity in the coordination networks examined. Co-retweet analysis focused on pro-Hong Kong/anti-mainland China themes and non-malicious automated activity. Co-link and co-tweet analysis showed little activity of interest.

### **Zhao Lijian's "Don't be afraid" tweet**

On November 30, 2020, Chinese Foreign Affairs spokesman Zhao Lijian posted a tweet with an image of an Australian soldier cutting the throat of a young Afghan boy. The context was the release of the Brereton Inquiry report into potential war crimes by Australian soldiers in Afghanistan. This tweet was condemned by Australian politicians, with Prime Minister Morrison branding it "truly repugnant".

Analysis of 14,927 replies to Zhao's tweet identified just 30 pro China accounts that were active across this and the earlier data set on responses to Australia's call for an independent COVID-19 inquiry. These accounts did exhibit some suspicious activity – their output increased from near zero to high levels through 2020, peaking at Zhao's tweet. Some appeared to be automated and many ceased activity soon after. However, Australia represented only a small fraction of the content tweeted by these accounts.

Coordination network analysis of the Zhao reply sample identified accounts that are pro-China and anti-Australia, often using misogynistic language. However, various clusters in the analysis focused on anti-Chinese Government themes, particularly human rights abuses in Xinjiang, Taiwanese and Tibetan independence and Falun Gong.

\*\*\*

In conclusion, this analysis provides no evidence of a major Chinese Government-backed campaign to influence Australian social media in 2020, even through low points in the Australia-China relationship. There are pro-Chinese Government accounts and some of these exhibit suspicious behaviour, but they are not numerous enough or active enough to give the impression of a concerted, state-backed campaign.

This is not to downplay the presence of, or potential for, malign Chinese Government influence in Australia. Espionage and other influence operations have been regularly reported for many years. Indeed, such reports are in the news at time of writing. But commentary that overstates the extent of Chinese Government interference, including on social media, is counterproductive for national security.

# Introduction

State-backed influence on social media has become an important issue internationally, while in Australia rarely a week goes by without headlines relating to Chinese Government influence. This report combines these issues to ask whether there is evidence of Chinese state-backed influence campaigns on Australian politics' social media platform of choice – Twitter.

We examine datasets from 2020 relating to two flashpoints in the Australia-China relationship:

- The period following the Australian Government's call for an independent inquiry into the origins of COVID-19.
- Chinese Ministry of Foreign Affairs spokesman Zhao Lijian's tweet of an image depicting an Australian soldier killing an Afghan child following the release of the Brereton report into potential war crimes by Australian soldiers in Afghanistan.

If there was a concerted Chinese Government-backed campaign to influence Australian social media users, these events would be likely targets for it.

Our analysis uses both quantitative and qualitative methods including statistical analysis of 'coordination networks' and other network analysis techniques. Despite extensive analysis of social media reaction to these events, we find no evidence of a major campaign to influence discussion of them on Twitter. There is evidence of pro-China accounts behaving suspiciously, such as being started around these events and ceasing activity soon after, or behaving inauthentically, such as seemingly automated behaviour, this is not on a large scale and is often dwarfed by similarly inauthentic activity from anti-China accounts.

# Case Study 1: Call for independent inquiry into origins of COVID-19

Australia-China relations were strained in early to mid 2020, following a call in late April by the Australian Government for an investigation into China's handling of the COVID-19 outbreak in Wuhan.<sup>1</sup> Initially made by Foreign Minister Marise Payne, the call was later backed by Prime Minister Scott Morrison. China's Government viewed this as an insult and the issue became enmeshed in various trade and diplomatic disputes.<sup>2</sup> As the fallout widened, several Chinese ministries, including the Ministry of Culture and Tourism, Ministry of Education and Ministry of Foreign Affairs, made statements that Australia is unsafe for Chinese nationals and students because of racism and discrimination, particularly in relation to the COVID-19 pandemic. These statements were labelled "disinformation" by Minister Payne.<sup>3</sup>

The dataset for Case Study 1 includes 54,441 tweets sent by 36,968 accounts during this period, specifically from 25/06/2020 to 15/07/2020. These tweets included the terms:

1. "china AND australia"
2. "china AND #auspol"
3. "racism AND australia"
4. "racist AND australia"

This dataset is herein referred to as the 'China/Australia collection'. The rationale for these search terms is to gather tweets where: (1) both China and Australia are broadly mentioned together; (2) China is mentioned in Australian political discussions, i.e. with the #auspol hashtag; (3,4) Australia and racism or racist are mentioned together. The focus on racism is a response to diplomatic tensions that flared.

---

<sup>1</sup> Worthington (2020) Marise Payne calls for global inquiry into China's handling of the coronavirus outbreak, <https://www.abc.net.au/news/2020-04-19/payne-calls-for-inquiry-china-handling-of-coronavirus-covid-19/12162968>

<sup>2</sup> Walsh (2021) *Australia called for a COVID-19 probe. China responded with a trade war*, <https://www.abc.net.au/news/2021-01-03/heres-what-happened-between-china-and-australia-in-2020/13019242>

<sup>3</sup> ABC (2020) *China accuses Australia of disinformation and 'political manipulation' in foreign ministry rebuke*, <https://www.abc.net.au/news/2020-06-18/china-says-australia-spreads-disinformation-marise-payne/12366948>

# Methods

We use a mixed methods approach including statistical analysis of ‘coordination networks’ (details below), qualitative content analysis, quantitative text analysis, and digital forensics techniques, including close reading of account profiles, analysis of account activity patterns, and reverse image searches. We draw on, and develop further, state-of-the-art approaches to detecting coordinated disinformation using network analysis techniques.

Firstly, we focus on reply activity to analyse accounts that have high rates of replying. This includes two parts: (1) manual content analysis of the top 100 most frequent repliers in the dataset; and (2) constructing and analysing a *reply network*. Using digital forensics and mixed methods approaches, we closely study clusters of interest in the reply network to report on patterns of interest.

Secondly, we examine temporally coordinated behaviour across four ‘coordination’ network types: (1) co-retweet, (2) co-tweet, (3) co-link, and (4) co-reply. This involves multiple accounts posting or reposting the same content, repeatedly and within 60 seconds of each other:

1. Co-retweet: reposting the same post
2. Co-tweet: posting identical text
3. Co-link: posting the same link
4. Co-reply: replying to the same post

This approach enables us to detect and explore coordinated activity in the dataset, i.e. sets of Twitter accounts that are coordinating their behaviour, which in turn has been used to detect disinformation operations in various settings.

## REPLY ACTIVITY ANALYSIS

Following previous studies, we begin by analysing accounts that have high rates of replying to the tweets in the China/Australia collection. We analyse the 100 most frequent repliers in the dataset, which sent a total of 947 reply tweets. The reply count distribution for these accounts is shown in Figure 1, below:



**Figure 1: Number of replies by top 100 most frequently replying accounts**




Figure 1 shows that only five accounts sent more than 20 replies, with most of the top 100 sending less than ten. This suggests there is limited automation or co-ordinated activity within the sample. The top replying account appeared to be automated, posting an anti-left or anti-globalist duplicated or ‘cypasta’ reply of no apparent relevance to China:

Radical mainstream fake news media,radical leftist/liberals,in European union,democrat party America with racist history,Canada,sweeden, Germany,Belgium,France,Netherlands, UK,Australia,newZealand,devil worshippers/shit wholes/corrupt/arab burka women opressors [sic]

The account, @Jackson87500650, was suspended as of 5th January 2021.

Next, we queried each account’s status in early January 2021 (seven months after the original data collection) to determine if the accounts were still active, suspended, or deleted. We are particularly interested in suspended accounts, which can contain accounts suspected by Twitter of being part of state-backed information operations.<sup>4</sup> Although we cannot determine why each account was suspended, given the length of time that passed

---

<sup>4</sup> Twitter (2021) *Information operations*, <https://transparency.twitter.com/en/reports/information-operations.html>

since data collection (seven months), we assume that the accounts are permanently suspended, which is the most severe enforcement action.

We found that only 16 of the top 100 repplier accounts were suspended by Twitter and 3 were deleted. This includes @Jackson87500650, discussed above and another of the most frequent repplier accounts, @Narsimhan17, which posted India-focused anti-Islam tweets – again, of no relevance to China.

We also examined what topics the top 100 repplier accounts focus on in their reply tweet activity (947 reply tweets in total), including whether there are particular topics correlated with the 16 suspended accounts. Given the size of the data, we use structural topic modelling (STM) to assist this analysis.<sup>5</sup> STM sorts the data into a given number of topics specified by the user, with each document (in this case tweet) having a probability of belonging to (or, more technically, being generated by) each of the k topics. In this way, we can automatically classify each tweet by topic according to the topic with which it has the highest probability. Furthermore, STM enables us to analyse topic covariates, i.e. correlations of topics with variables of interest. Hence, we use STM to assess correlation between tweets by suspended accounts and particular topics. We chose 20 topics to sort the tweet data into an analytically manageable number of topics, and the model converged in X iterations which suggests a reasonable goodness-of-fit with the data (see Appendix for convergence plot).

Table 1 below shows the topics, including a label that we ascribe based on qualitative analysis, number of tweets belong to each topic, and the top 5 terms (words, phrases, emojis, add hashtags) associated with each topic. Additionally, we organise the topics into over-arching ‘themes’. These themes are derived qualitatively through content analysis of similarities between topics and terms within the topics.

**Table 1: China/Australia collection: Frequent repplier tweets by topic**

Topic	Over-arching theme	No. of tweets	Top 5 terms (FREX)
<b>Australia supporting and perpetuating racism</b>	Australia’s racism	138	modern, supporting, level_racism, australia_perpetuated, australia_perpetuated_people
<b>Anti-left cospypasta</b>	Anti-left	132	radical, party, mainstream, germany, america_racist
<b>China and India cospypasta</b>	Anti-CCP	93	, #vietnam, #usa_#india,  _#ccp, #china-
<b>Defending against China’s operations in international water</b>	Anti-CCP	84	uk_china, 🇬🇧, 🇨🇦, canada_china, china_wrong

<sup>5</sup> Roberts et al (2020) STM: R Package for Structural Topic Models, <https://cran.r-project.org/web/packages/stm/vignettes/stmVignette.pdf>

<b>Australia's racist history / trade and the economy</b>	Other	82	australia_severely, australia_severely_racist, banana, banana_nation, banana_nation_onward
<b>China bioweapon conspiracy cypypasta</b>	Anti-CCP	76	amp_many, weak, #corona, #corona_biological, #corona_biological_terror
<b>Alan Tudge criticism cypypasta</b>	Other	63	calling_muscular, calling_muscular_defense, london_calling_muscular, #blacklivesmatteraus, #ccp_propaganda_machines
<b>Colonial countries as racist and incompetent</b>	Australia's racism	54	australia's_level, australia's_level_success, china_manage, china_manage_fool
<b>Criticising Xi Jinping</b>	Anti-CCP	42	china_wuhan, minerals_council, #china_rhetoric, amp_#auspol, anti_#china
<b>Anti-Trump cypypasta</b>	Anti-US	30	donny, donny_trump, america_china, america_china_amp, amp_australia_kahoots
<b>Free trade agreement cypypasta</b>	Trade	26	CN_one, CN_one_says, boo, one_says, one_says_boo
<b>The 'Quad' and countries united against China</b>	Anti-CCP	24	🇨🇦🇨🇦, 🇨🇦, age_expansionism, alliance, alliance_red
<b>Liberal signing trade deal</b>	Trade	23	liberals_signed, liberals_signed_mou, signed_mou, signed_mou_china, tension
<b>Debating Australia's racism</b>	Australia's racism	21	girls, #racism, proof, #australia_#racism, approved
<b>CCP infiltration and interference</b>	Anti-CCP	18	steals, #ccp_#china, anymore, #ccp, private
<b>China imports of Indonesian nickel</b>	Trade	10	di, dari, mereka, dan, punya
<b>Debates about Australia's racism and Chinese students leaving</b>	Australia's racism	8	better, leave_australia, ignoring_science, science, leave
<b>Never trust China cypypasta</b>	Anti-CCP	8	never_trust, never_trust_china, trust_china, trust, never
<b>Debates about Australia's racism towards Chinese people</b>	Australia's racism	7	australia_also, ^, tiny, australia_return, right_chinese
<b>China positive reforms of industry</b>	Pro-China	7	done, picture, bigger_picture, farming, outlawed

As Table 1 shows, over a third of the reply tweets sent by the top 100 frequent repliers fit into the 'Anti-CCP' theme (n=345 reply tweets across 7 topics), while a quarter of tweets

belong to the 'Australia's racism' theme that we might expect from Chinese Government-linked accounts (n=228 reply tweets across 5 topics). These are by far the two most dominant themes of the reply tweets. The remaining tweets are spread across various themes including 'trade', 'anti-US', and 'pro-China', with some not clearly belonging to any particular theme ('other').

In assessing whether some of the 'Australia's racism' reply accounts might be linked to the Chinese state, further analysis was based on:

- (1) account creation date;
- (2) number of tweets accounts have posted since creation;
- (3) engagement metrics by retweets and likes for tweets authored by these accounts; and
- (4) time when tweets in this dataset were posted.

The three 'Australia as racist' accounts were created on 19/09/2019, 16/06/2016, and 10/09/2010 - no particular patterns of interest. However, the two 'anti-US' accounts were created one week apart (21/07/2019 and 28/07/2019) and the two 'pro-China' accounts created recently and three weeks apart (22/05/2020 and 12/06/2020). Interestingly, one account posted 38,496 tweets since its creation date (21/07/2019) until suspension, which is about one tweet every 12 minutes, if it tweets 16 hours per day (assuming 8 hours 'sleep' time). The accounts received very little engagement for their efforts, with an average of 0.7 retweets (median 0) and 1.5 likes (median 0) per tweet.

Turning attention to temporal characteristics of activity, Figure 2 shows unusual patterns of the time when tweets are authored by these top 100 replier accounts:

**Figure 2: Posting time of reply tweets sent by suspended and deleted pro-China accounts**


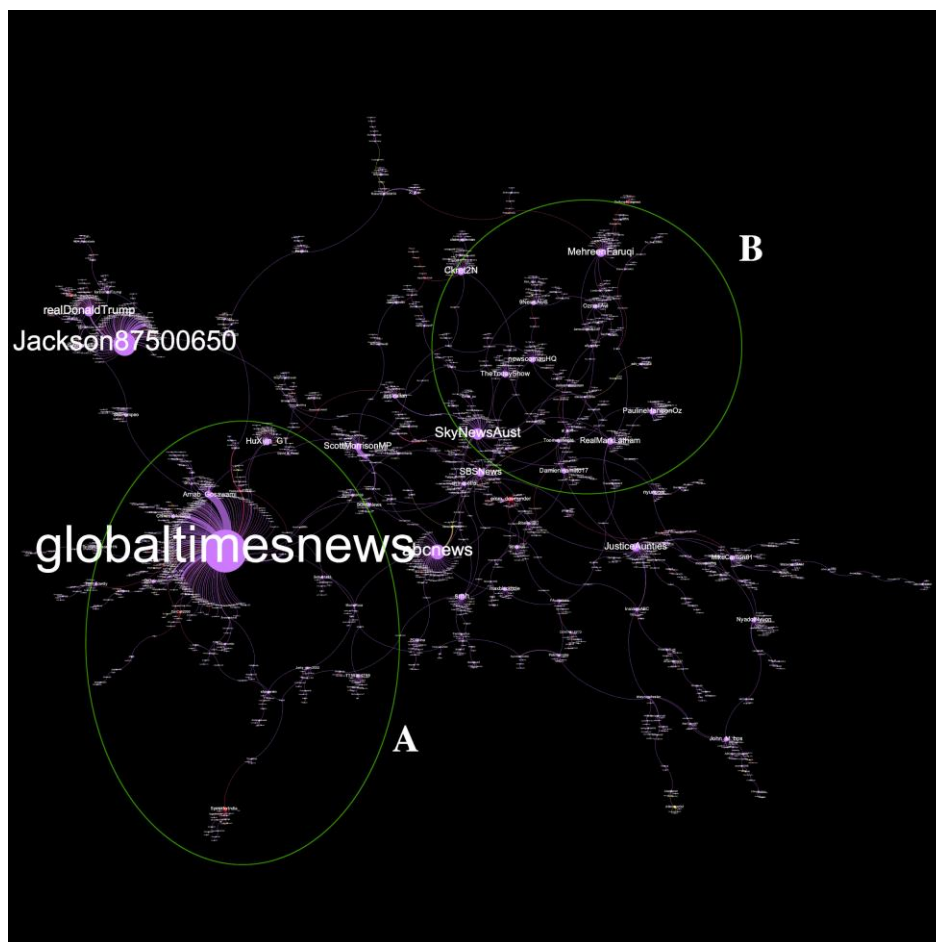


Figure 2 shows that the most active hours are 1am and 5am Beijing time (or 4am and 8am Australian Eastern Standard Time) providing no evidence to suggest these accounts are posting during Beijing work hours as previous studies have identified. A simple, if not particularly insightful, explanation is that these accounts are aligned with multiple time zones including, but not limited to, Australia and China.

We turn attention now to broader reply activity in the dataset. The complete reply network consists of 8,626 nodes and 6,228 edges, where an edge from A to B means that A replied to B, shown in Figure 3:

**Figure 3: Visualisation of the giant connected component extracted from reply network**



The reply network shown in Figure 3 shows active accounts in purple; suspended accounts in red and deleted accounts in yellow. Nodes sized by degree. While it is difficult to see the suspended and deleted accounts in the reproduction above, there little evidence of coordinated behaviour - they are spread out. We therefore focus on the giant connected component in the middle where most activity is concentrated.

Figure 3 shows the giant connected component (GCC) of the reply network, containing 1896 nodes and 2047 edges. We highlight two clusters of interest in this network, A and B, to

which we now turn attention. We also note the cluster around @Jackson87500650, but do not discuss further as it is not relevant to China-Australia discussion.

Cluster A captures reply activity around *Global Times News*, official account of the state-owned Chinese media outlet. This activity includes replies directly to @globaltimesnews and indirectly via reply trees (replies to replies, etc). Irrespective of the topic of each original tweet by @globaltimesnews, the reply activity broadly falls into three categories:

- (1) pro-China / critical of West;
- (2) anti-China; and
- (3) pro-India / anti-China.

Of course, the replies must also contain reference to Australia given the focus of the dataset. Table 2 shows examples of replies that fall into each category:

**Table 2: Examples of reply tweets from each category**

Category	Tweet text
Pro-China / critical of West	<p>“US policy is not to engage with enemies but use other countries for this purpose</p> <p>Nowadays US is using India Australia and UK against China But China was is and will always be strong against them</p> <p>These countries must have some self respect&amp;decrease US influence”</p> <p>Tweet ID1283005772843094021</p>
Pro-China / critical of West	<p>MichaelSchuman @bopinion Full of typical western propaganda and Sinophobia. \U0001f644</p> <ul style="list-style-type: none"> <li>- You skipped western drug trafficking, while talking about Opium War</li> <li>- Never mention racism, colonialism and “Yellow Peril” paranoia in the West</li> <li>- Don’t ever talk about India, Australia etc. being vassal states of the US</li> </ul> <p>Tweet ID1282432112222539776</p>
Pro-China / critical of West	<p>“We must teach lesson to them.</p> <p>They must Stay away from Hong Kong issue. They have nothing to do with China internal affairs. Infact if they (US, UK, Australia) try to continue this non sense then they must wait for Harsh consequences from China</p> <p>Every country must accept OneChina”</p> <p>Tweet ID1283226944029548544</p>

Anti-China	<p>"BIG BREAKIBG - China is weak &amp; surrounded with enemies, for it's #Corona biological terror attack massacre on the world.</p> <p>America, Australia, European Union, India, Japan, Taiwan, Vietnam, South Korea &amp; many others stand united to confront China."</p> <p>Tweet ID1282694571881709568</p>
Anti-China	<p>"US must use its years of Navy experience and train India, Japan, and Australia to stop Chinese expansion. India, Japan, and Australia have been peace loving countries and never participated in battles, but China has changed the scenario, now they must learn naval battles."</p> <p>Tweet ID1282412629445746690</p>
Anti-China	<p>"No we don't trust #china and the #ccp because you lie, coerce Labor politicians, have spies and agents in Australia, covered up the #wuhanvirus #bioweapon and push the confucious rubbish. No trust and friends anymore #Gongfei #ccp #Foffchina #BoycottChina #auspol <a href="https://t.co/nECX2uNwqu">https://t.co/nECX2uNwqu</a>"</p> <p>Tweet ID1279341547440476161</p>
Pro-India / anti-China	<p>Are you a Chinese virus or Kung Flu? Indians leading world's top firms like Microsoft, Alphabet (Google) etc and gave Buddhism to China while Chinese stealing IPR from other nations and giving Chinese virus to world. #Chinesevirus #kungflu #China #USA #India #Japan #Australia</p> <p>Tweet ID 1281875107355516928</p>
Pro-India / anti-China	<p>#China lost #India- Neutral Peaceful Powerful Nation. #GalvanValley would b a Suicide 4 #Chinese #PLA with World hating #CCP &amp; doing #BoycottChina World with #IndianArmy #USA , #India #Taiwan #Australia #Japan #Vietnam #UK #France #Russia</p> <p>Tweet ID1282072818444759041</p>

Across all categories we observe a considerable amount of spam-like behaviour, where the exact same or very similar text is used across multiple tweets in the dataset. Given that some of these accounts are still active at the time of writing (i.e. not suspended or deleted), these duplicated tweets are still publicly available to view and some of the accounts are still actively spamming replies. Overall, we find that replies to Chinese state media, including the case study of @globaltimesnews presented here, have a mixture of all three categories. It is a messy information environment with multiple, often oppositional, perspectives and narratives all competing for attention in the reply threads. The actors range from fringe anonymous 'sockpuppet' type accounts through to Chinese state officials such as Hu Xijin, Editor-in-chief of the Global Times (@HuXijin\_GT).

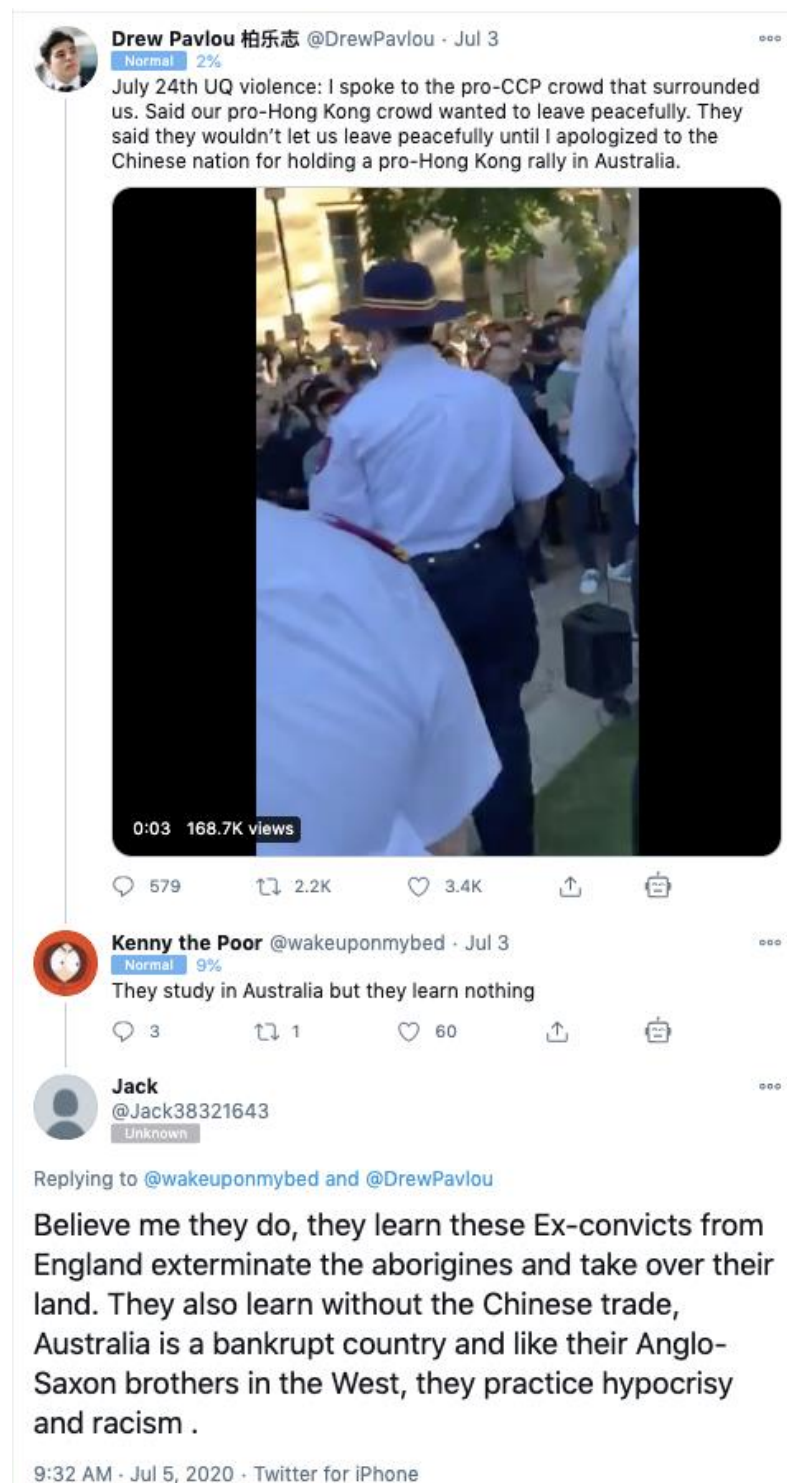
Cluster B in Figure 3 (above) captures loosely clustered reply activity around the topic of racism in Australia, with one side accusing Australia of racism (historically and/or in relation to contemporary policy issues) and the other defending against such accusations and/or promoting racially-motivated policies. The most central accounts in Cluster B are right-wing social media personality and convicted domestic violence offender Avi Yemini (@ozraeliavi), various mainstream news agencies detailed below, two right-wing Australian politicians: Pauline Hanson (@paulinehansonoz) and Mark Latham (@realmarklatham), and Greens Senator Mehreen Faruqi (@MehreenFaruqi).

The main news media accounts in Cluster B are @SkyNewsAus, @newscomauHQ, @TheTodayShow, @9NewsAUS and @SBSNews as well as international agencies such as the South China Morning Post (@SCMPNews). The articles in tweets posted by Australian news accounts are generally about Australia/China relations, and often specifically about two Federal government initiatives at the time of data collection: (1) the \$270 billion defence package announced by Prime Minister Scott Morrison, and (2) Australia's suspension of the Hong Kong extradition treaty and offer to extend visas of some Hong Kongers working and living in Australia.

Cluster B includes a number of anonymous pro-China accounts that emphasise racism issues in Australia on Australian Twitter. Figure 4 shows an example of a pro-China account replying to Drew Pavlou, a controversial Australian student activist who is critical of the Chinese Government and organisations he considers friendly to it:



Figure 4: Anonymous pro-China account reply



In this instance, the anonymous account is using Australian colonial history and genocide as part of its broader argument that Australia is economically dependent on China. This is typical of such accounts in Cluster B which post a real and historically-rooted issue (genocide of Indigenous Australians) to magnify tensions and push a particular narrative. Many have a particular focus on #AboriginalLivesMatter and the larger global #BlackLivesMatter movement.

It is difficult to ascertain whether these accounts are backed by the Chinese state. Given the relatively small number of them, it seems unlikely. One responded to a direct message, stating that they did work for the Chinese Government, but that their Twitter account expressed their own views and was not linked to their work. It seems more likely that most of these accounts are simply Chinese people with generally pro-Chinese Government views.

# Coordination network analysis

Overall, for the Australia/China dataset we find very little suspicious or ‘inauthentic’ activity in the four types of coordination networks examined. For co-retweet analysis, we find that most of the coordinated retweeting is pro-Hong Kong / anti-China activists (7 of the top 10 clusters), as well as commercial spam and non-malicious automated bot activity (i.e. computer enthusiast news bots randomly retweeting China/Australia related IT content).

**Figure 5: Co-retweet network of China/Australia collection (top 10 community clusters)**

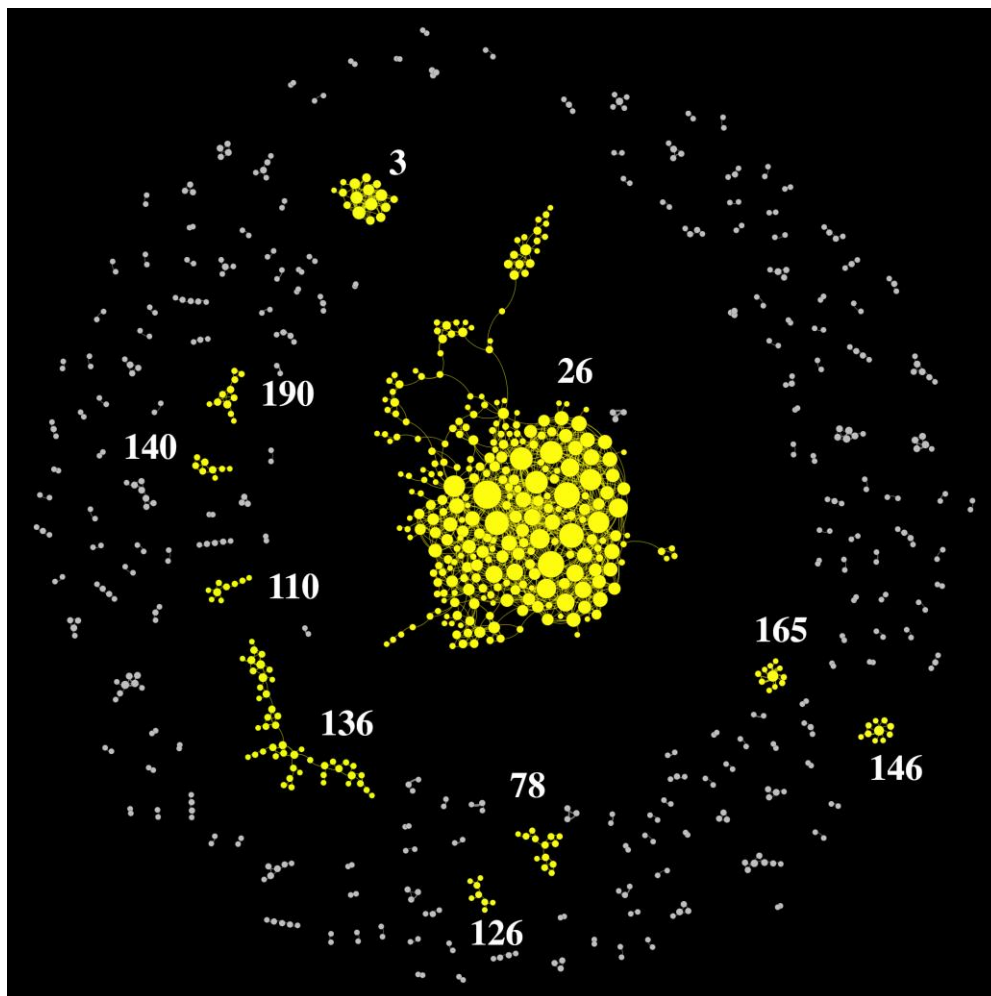


Figure 5 shows the co-retweet network, which contains 886 nodes (accounts) and 1359 edges (co-retweets). The top 10 community clusters by size are highlighted in yellow and labelled according to community number.

The largest community (#26) is a network of fans/spammers promoting Korean popular (k-pop) band BTS (@BTS\_twt). This cluster is a false positive as the tweets are about BTS songs

hitting the top 10 on the iTunes charts across multiple countries, including China and Australia. Most of the accounts engage in high-volume spam.

Community #136 consists of pro-Hong Kong activists tweeting about Australia's suspension of the Hong Kong extradition treaty and offer to extend visas of some Hong Kongers working and living in Australia. Accounts in this cluster display a strongly anti-China stance, including some accounts that retweet in multiple languages. For instance, @Tak47971105 retweets anti-China tweets in English, Japanese, German, French, and Italian.

Community #3 is a (self-identified) bot network that retweets hashtags relating to computer programming and IT, in particular #100DaysOfCode. These accounts retweeted China/Australia related tweets because they also contained hashtags and terms about computer programming and IT.

Community #78 is another pro-Hong Kong / anti-China network, retweeting about the same topic as Community #136.

Community #165 contains Australian news media accounts, including 9 News Sydney, 9 News Australia, 7 News, and The Australian. The cluster centres around an extremely high-volume retweet bot known as @viriyabot, which automatically retweets #china tweets, thus co-retweeting all the news media articles from these outlets. It also includes two high-volume bot accounts that almost exclusively retweet Bloomberg and Business Insider (@DoKscho) and Bloomberg and New York Times (@test11lmagsino). The retweets in this cluster are mainly about the Australia/Hong Kong extradition treaty suspension, visa extension for Hong Kongers, and growing diplomatic tensions between the nations.

Community #190 is another pro-Hong Kong / anti-China network, retweeting about the same topic as Community #136. Likewise, Community #146 has @SheepLiberateHK as the central node, and is another pro-Hong Kong / anti-China network. Notably several of the accounts retweet multiple language tweets as observed in Community #136. Further, we find that communities #110, #126, and #140 are all pro-Hong Kong / anti-China network, retweeting about the same topic as Community #136, and a closer look at these accounts shows that they also post pro-Trump/MAGA content, particularly where such content is critical of China.

Altogether there are 106 unique accounts in the 7 pro-Hong Kong / anti-China clusters, of which six are suspended by Twitter.

# Co-link network analysis

For *co-link* analysis, we find a small amount of activity but nothing of particular interest except for a dense cluster of '7 News' Australian media accounts that coordinate their syndication of China/Australia related news stories, as shown in Figure 6 below:

**Figure 6: Co-link network of China/Australia collection**




Figure 6 shows the co-link network, which is quite small with only 17 accounts (nodes) and 30 co-retweets (edges). The largest cluster consists of '7 News' media outlets, which syndicate their articles nationally and across the state capitals. These are co-retweets of 7 News videos and articles mainly on two topics:

- (1) a Labor MP who was raided by Australian security agencies due to alleged links to Chinese influence operations; and
- (2) Australia's suspension of the Hong Kong extradition treaty and offer to extend visas from Hong Kong citizens.

The second-largest cluster (bottom-left Figure 6) consists of indie music marketing accounts, which appear to be largely automated. The bottom-right cluster in Figure 6 comprise three accounts that share links to a South America-focussed blog The last cluster consists of a bot account known as @grauniad\_news which focuses on tweets with spelling mistakes by The Guardian. Overall, there is little activity of interest in the co-link network.

# Co-tweet analysis

For *co-tweet* analysis, we do not find any evidence of coordinated duplication of content. Similarly, there is no *co-reply* activity, despite previous studies identifying ‘high rates of replying’ as a key strategy of CCP-linked information operations.<sup>6</sup>

---

<sup>6</sup> Serrato and Schafer (2020) Reply All: Inauthenticity and Coordinated Replying in Pro-Chinese Communist Party Twitter Networks, <https://www.isdglobal.org/isd-publications/reply-all-inauthenticity-and-coordinated-replying-in-pro-chinese-communist-party-twitter-networks/>

# Case Study 2: Zhao Lijian's “Don’t be afraid” tweet

On November 30, 2020, Chinese Foreign Affairs spokesman Zhao Lijian posted a tweet with an image of an Australian soldier cutting the throat of a young Afghan boy. The context was the release of the Brereton Inquiry report into potential war crimes by Australian soldiers in Afghanistan. The tweet is reproduced in Figure 7 below:

**Figure 7: The “Don’t be afraid” tweet posted by Zhao Lijian**



This tweet was condemned by Australian politicians, with Prime Minister Morrison branding it “truly repugnant.”<sup>7</sup> The tweet and responses represented another period of particular strain on the Australia-China relationship.

<sup>7</sup> Palmer (2021) *Snarling diplomacy: the ‘wolf warrior’ amping up China’s aggro on social media*, <https://www.smh.com.au/world/asia/great-tweet-forward-the-wolf-warrior-helping-china-discover-its-aggressive-voice-20210719-p58av7.html>

The dataset consists of replies to this tweet and replies between users on this tweet.

In this case study, we collected 14,927 replies to the Zhao tweet. This represents 80% of the total replies to Zhao, it was not possible to collect every reply for technical reasons.<sup>8</sup> This dataset is herein referred to as the Zhao tweet dataset.

We analyse the Zhao tweet dataset using a range of methods including descriptive statistics, network analysis, qualitative close reading of account profiles and tweets, and digital forensics (e.g. reverse image search, behavioural analysis tools).

The main goal for the Zhao tweet collection was to examine whether accounts replying to this tweet were also active in the China/Australia dataset we collected. Having identified the set of accounts that are active across both datasets, the next step involved manual qualitative content analysis of each tweet to classify it as “pro-China”, “anti-China”, “neutral” and “other”. This also involved a close reading of account profiles and general tweet activity and content. This resulted in a set of 30 pro-China accounts that were active across *both* datasets and therefore are of particular interest given the objectives of this study.

For each of the 30 accounts, we collected their recent tweet history (up to a maximum of 3,200 tweets due to Twitter API restrictions) to understand their volume of activity, how much engagement they get, the temporal patterns of activity, and what topics they talk about (with a particular focus on tweets that include reference to “Australia”). Next, we collected their follower and friend lists to examine how they are socially connected, and what this tells us about their patterns of behaviour.

The tweet history dataset for these 30 accounts consists of 77,455 tweets, of which 8,340 are original tweets (10.8%), 27,333 are retweets (35.3%), and 41,782 are replies (53.9%).


Over half of all tweets received zero retweets (56.9%;  $n=44,051$ ) and two-thirds received 0 likes (66.2%;  $n=51,249$ ). Per *tweet*, the median number of retweets is 0 (average = 430) and median likes is 0 (average 2.9). The median number of followers is 350 (average 28,228) and 11 accounts have less than 100 followers.

---

<sup>8</sup> The problem is recursively collecting replies, which is slow due to API rate limits, and ultimately caused errors with the software we used (the ‘twarc’ open source command line toolkit).



**Figure 8: Volume of tweet activity (excluding retweets) and engagement metrics**




Next, for the 30 accounts that both replied to Zhao’s tweet and were in the Australia/China dataset, we examine the volume of activity and engagement for tweets that mention ‘Australia’ and/or ‘Morrison’. There are 1,848 tweets about Australia/Morrison, with an average of 31 per account. Nearly half the activity (47.9%) is reply tweets (n=885). We also observe 208 original tweets (11.2% of total) and 755 retweets (40.9%).

Figure 9 below shows the volume of Australia/Morrison tweets (excluding retweets) along with the total engagement metrics for these tweets by number of retweets and likes. Almost half these tweets received zero retweets (n=737; 46.8%) and over two-thirds received zero likes (n=1098; 69.8%).

**Figure 9: Volume of tweet activity for tweets mentioning 'Australia' or 'Morrison' (excluding retweets) and engagement metrics, ordered by number of tweets**

Most active accounts tweeting about Australia (exc. retweets), along with engagement metrics



Note that the account dominating engagement statistics is the account of Zhao Lijian.

Turning attention to temporal patterns of 'Australia' tweet activity, Figure 10 shows the volume of Australia/Morrison tweets per day sent by these 30 accounts during 2020.

**Figure 10: Tweet volume per day and timeline of events for pro-China accounts mentioning 'Australia' or 'Morrison'**




Figure 10 shows that there is very little activity until late April 2020, when a low-volume but steady stream of tweets begins. This coincides with Prime Minister Scott Morrison calling for a global review of China's handling of the COVID-19 outbreak in Wuhan. The volume increased until mid-June when Australia announced the suspension of the Hong Kong extradition treaty, which also drew criticism from the Chinese government. On 9 September, we observed a small spike of activity when the story broke about ASIO questioning Chinese journalists in Australia. The volume steadily increased until 30 November when Zhao Lijian sent the "Don't be afraid" tweet, coinciding with a massive spike of activity.

By comparison, Figure 11 shows the tweet volume over time for tweets about Australia/Morrison versus all content. It is clear that Australia is a relatively minor topic for these accounts compared to their broader tweet activity:

**Figure 11: Tweet volume per day for pro-China accounts, Australia vs all content**

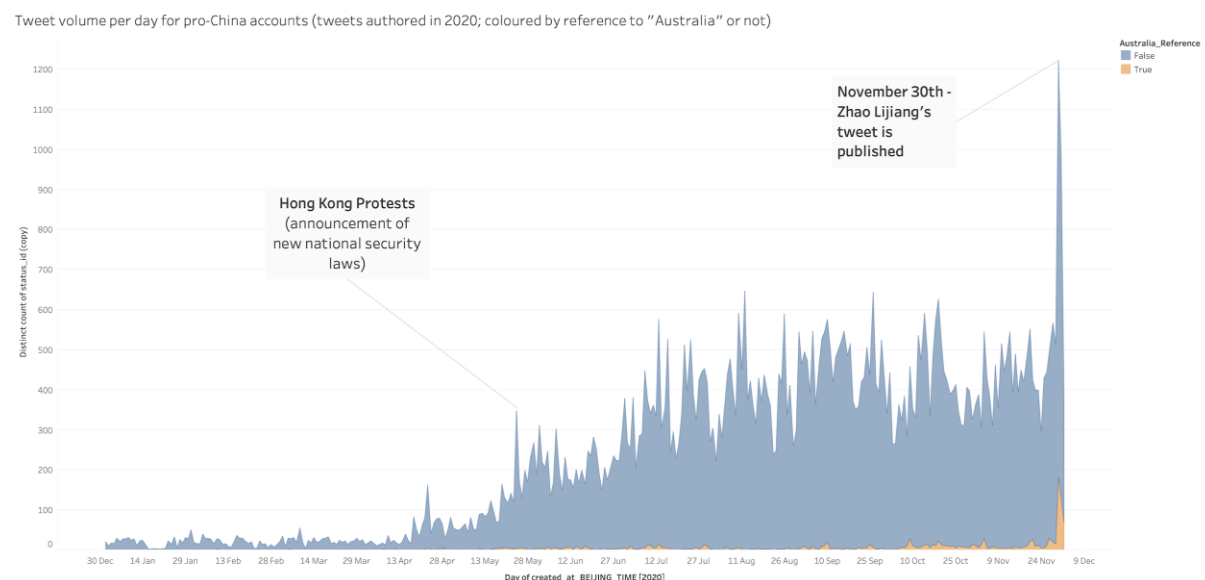


Figure 11 shows that activity in these accounts increases markedly from about 24th May onwards, coinciding with the announcement of new national security laws in relation to the Hong Kong protests. Likewise, there is a massive spike of activity on November 30th coinciding with Zhao's tweet, so the accounts tweeted more than usual about this event.

Next, we break down the main topics of the reply tweets sent by these accounts. We examine specifically the subset of tweets that mention Australia and/or Morrison. To do this, we used an approach known as Latent Dirichlet Allocation (LDA) topic modelling to automatically sort the reply tweets into categories or 'topics'. This produced 67 topics, with the top ten shown in Table 3 below:

**Table 3: Top 10 topics by pro-China accounts of interest in tweets re Australia/Morrison**


Topic	Label	Top 10 words
25 (n=48)	Murder of two Afghani boys / artwork in Zhao's tweet	australian, government, throats, @zlj517, chinese, forces, now, julian, assange, apology
4 (n=43)	Anti-Trump / Anti-Morrison cospypasta	china, australia, trump, trade, wrong, answer, minister, morrison, australian, america
35 (n=40)	Condemning the Australian military	world, australian, forces, australia, army, afghanistan, shame, worst, unarmed, killed
53 (n=38)	Failure of Western democracies	australia, countries, like, open, western, china, point, trump's, @financialreview, country
6 (n=37)	Chinese/Australia trade and economy	@australian, chinese, children, australia, can, just, time, @bennpackham, little, every
13 (n=37)	Australia's historical war crimes and atrocities	australia, war, like, china, us, crimes, wine, australian, australia's, real
55 (n=36)	'Australia really needs China market' cospypasta	china, australia, new, zealand, us, australian, exports, australia's, canada, wine
42 (n=35)	Pro-Daniel Andrews / anti-Morrison cospypasta	morrison, must, furious, sure, people, right, away, amp, hypocrisy, mr
20 (n=33)	Pro-China / CCP and denying Chinese war crimes and/or human rights issues	years, people, chinese, australia, morrison, know, good, china, australian, criminal
30 (n=31)	Criticising Australia media and journalists	really, ask, want, australia, better, response, australian, @australian, saying, people

The distribution of tweets per topic is shown in Figure 12.<sup>9</sup>

---

<sup>9</sup> We used a data-driven approach to determining the number of topics, specifically the 'spectral initialisation' method in the R 'stm' package. The model converged after 11 iterations.

**Figure 12: Number of tweets that mention Australia/Morrison per topic (k=67 topics)**



The 10 most frequent topics are shown in Table 3, along with a label and summary description. The number of tweets that belong to each topic is shown in brackets. The most frequent topic (#25) is about Zhao’s tweet and the artwork, arguing that it helps Australia’s war crimes to be held to account and drawing attention to other examples of Australian military ‘disgrace’. Likewise, topic #35 (‘Condemning the Australian military’) and topic #13 (‘Australia’s historical war crimes and atrocities’) draw attention to this issue.

There are three topics that are broadly anti-Scott Morrison (topic #43 and #42) and Western liberal democracies (topic #53), often with a special focus on anti-US sentiment and a (perceived) problematic relationship between Australia and US President Donald Trump. We also find a number of duplicated or ‘copypasta’ reply tweets where single or multiple accounts post the same text multiple times. Likewise, an overarching theme of topic #53 and topic #20 is the hypocrisy of Australia and failure of Western liberal governments as well as minimising and/or denying criticisms of Chinese war crimes and/or alleged human rights issues in Xinjiang.

Topic #6 and topic #55 focus on Chinese/Australia trade and the economies of both countries. An overarching narrative is that the Chinese economy is strong and Australia needs to keep it on side to avoid a one-sided economic impact. This topic includes a copypasta tweet that argues “Australia really needs China market” and posts statistics about China’s import of wine from Australia, fitting with the narrative that Chinese trade sanctions are a problem for Australia’s economy and it should therefore cooperate:


According to Australian Wine Authority, Australian wine exports totalled A\$2.89B from October 2018 to September 2019. China imported A\$1.2B, more than the US, the UK, Canada, New Zealand and Japan add together. So, your show is

meaningless.\U0001f61c\nAustralia really needs China market.  
<https://t.co/QcPBxm8ax9> <https://t.co/MHx4moYOUj>

# Coordination network analysis - Zhao tweet dataset

We now turn attention to coordinated reply activity to Zhao Lijian's tweet. Figure 13 below shows the coordinated reply network, consisting of 4,199 nodes and 4,681 edges. Accounts (nodes) are connected if they replied within 60 seconds of each other, at least twice. Size of nodes indicates the volume of co-reply activity (larger nodes co-replied more). The findings suggest that accounts replying to this tweet are not all acting independently - Figure 13 reveals a large number of co-reply clusters, connected into a large connected component at the network's core. The top 10 community clusters are coloured and labelled in Figure 13 and we now analyse these.

**Figure 13: Coordinated reply activity on Zhao Lijian's tweet**



The biggest Cluster (#9) consists of cypypasta reply tweets that are pro-China and anti-Australia, often using misogynistic language to troll other accounts that are critical of the tweet and/or Zhao Lijian or the CCP in general. Figure 14 and Figure 15 below show screenshots of example tweets from the main accounts in this cluster. The tweets are a mix

of English and Chinese, often with the same or very similar content cypasted. For instance, Figure 14 shows tweets duplicating the phrase “澳大利亚政府就像个婊子” (rough translation: “The Australian government is like a whore”) and Figure 15 shows tweets with exact duplicates of “This is just a painting, but it reveals the real behavior of Australia!”. These two accounts were created in November 2020 and August 2020 and have 2 and 0 followers respectively. Both largely ceased activity following the Zhao tweet.

**Figure 14: Tweets starting with “The Australian government is like a bitch” in Chinese**



**Figure 15: Copyasta tweets pushing the narrative that the image represents Australia**



The small nodes connected to the larger ones in Cluster #9 are a mix of accounts that ‘take the bait’ and are arguing with the trolls or criticising Zhao and/or China, as well as other trolls that are praising Zhao and denying any criticisms. We note that the literal translation of the example tweet in Figure 16 is “fuck your mother one time for 50 cents”.



**Figure 16: Example troll tweet from Cluster #9 with misogynistic language**



Cluster #22 is the second largest and centres around a single account, @MattNA12, which has 0 followers and through this period only replied with photos and videos of Chinese government human rights abuses and state oppression (see example tweet in Figure 17 below). The other main nodes in Cluster #22 are also critical of China and Zhao's tweet, including counter-trolls that only post photos of the Taiwan flag or post gifs with phrases such as "China is starting to remind me of a crazy girlfriend I once dumped".

**Figure 17: Anti-China replies from the main account in Cluster #22**



Cluster #20 is led by an anti-China account, @Julzy62081062, who replies to pro-China accounts with criticisms of the Uighur Muslim camps (Figure 18) and taunts Zhao directly with tweets such as “Come down to austrlia little man” [sic].

**Figure 18: Anti-China tweets highlighting Xinjiang human rights abuses**



The other main accounts in cluster #20 are also critical of China and try to bring attention to a range of issues including Taiwan and Tibetan independence, oppression of Falun Dafa, and the Tiananmen Square protests. The cluster also includes pro-China/anti-Australia accounts posting counter-replies and/or trolling the anti-China accounts. As shown in Figure 19, some of these counter tweets are cypasta duplicates.

Figure 19: Copypasta troll reply tweets in Cluster #20



Cluster #54 consists of anti-China/pro-Australia and pro-China/anti-Australia accounts arguing and trolling each other. Figure 20 provides an example from Cluster #54 of a pro-China/anti-Australia account (@luxuryoem) posting an image of Indigenous Australians in chains, along with the reply from an anti-China account that tries to downplay and make light of Indigenous genocide.

Figure 20: Pro-China and anti-China accounts trolling each other



As Table 4 shows, brief investigation of @luxuryoem’s recent tweet history reveals multiple copy-pasted tweets that use inflammatory language and imagery to criticise Australia and get an emotional reaction out of other users.

**Table 4: Repeated content from Cluster #54 account of interest**

Tweet text	Number of times repeated
Perfect! Drink this Afghan Blood Make Australia great again!	13
Drinking Afghan blood makes Australia great again! Cheers!	10
Like this?	6
Perfect! Drink this Afghan blood and make Australia great again!	6
Nice picture, but I like this one.	5
Just like this?	3
Cute picture, but I like this one better.	3
かわいそうなゴキブリ (Japanese for “poor cockroach”)	2
ばかやろう ... (Japanese for “idiot”)	2
I am Tibetan, and when you come to Tibet I will make you white pigs suck my dick	2
Evil Country!	2
How many bathtubs does the Austrian Navy have? Makes you so confident? Sorry Australia.	2

Cluster #54 includes tweets accusing China of the Uyghur abuse and detention centres, with pro-China accounts replying to deny these accusations and offer rebuttals. Figure 21 provides an example - this particular pro-China account was created on November 10th 2020 and has 7 followers. It still exists, but has not been active since early 2021.

**Figure 21: Example of newly created pro-China accounts countering tweets about human rights abuses in Xinjiang**



Cluster #85 centres around one account, @democracy\_forhk, that replies with photos and videos of Hong Kong police assaulting activists during the Hong Kong protests, as shown in Figure 22:

**Figure 22: Reply tweets highlighting police brutality against Hong Kong protesters**



The small nodes surrounding this large node in Cluster #85 are mostly a mix of pro-China accounts posting anti-Hong Kong sentiment, pro-China accounts endorsing Zhao's tweet, and anti-China accounts criticising Zhao's tweet.



Figure 23: Anti-Taiwan and anti-Hong Kong reply tweet



Cluster #31 reply activity is about the hypocrisy of the tweet given China's human rights abuses. The largest node is an account spamming replies that reference various issues such as China's dependence on Australian iron ore imports, detainment of journalist Cheng Lei, and oppression of Falun Gong. It also contains a number of smaller nodes referencing various events and issues in China's human rights track record. Cluster #31 also includes accounts replying with inflammatory and racist content directed at Chinese people.

Cluster #66 consists of two anti-CCP accounts posting reply tweets that include the photo of "Tank Man" in Tiananmen Square and also screenshots of pornographic content and accounts that Zhao Lijian's account has "liked" or "followed" (Figure 24). This account was created in November 2020 and ceased activity in mid-2021. The cluster also includes pro-China accounts that criticise Australia for the Afghan war crimes and arguing that China should not apologise for the tweet.

Figure 25: Anti-China accounts claiming Zhao engaged with pornographic content



Cluster #176 is a mix of accounts posting about China's human rights abuses, military actions and detention camps in Xinjiang. Several of these accounts post racist content as shown in Figure 26. There are also a number of pro-China accounts in this cluster that argue against accusations of China's human rights record, for example in Figure 27.

**Figure 26: Anti-China racist tweets in Cluster #176**



Finally, both Cluster #39 and Cluster #2 contain a number of pro-China accounts that post replies to tweets accusing China of human rights abuses, which try to downplay and/or dismiss such accusations (see example in Figure 27) or post memes and photos as a strategy to counter images of Tank Man and other sensitive Chinese events and issues (see example in Figure 28 and 29).



**Figure 27: Tweets in Cluster #39 on human rights abuse in Xinjiang**



Like many suspicious accounts, the profile in Figure 27 is taken from the web and is an old photo of John Lone, a Hong Kong-born American performance artist, singer and actor. The account was created in June 2020 and is still active on various topics.

Accounts in this cluster use a visual form of 'whataboutism' to counter accusations of Chinese human rights issues, posting images such as US police suppression of Black Lives Matter activists during 2020 or Australian soldiers in Afghanistan, as shown in Figure 28 and 29:

**Figure 28: Responses to images of Chinese human rights abuses**



**Figure 29: Responses to images of Chinese human rights abuses**



**Daniel Bryar** @BryarDaniel · Nov 30

Unknown

Replying to @zlj517

Were these soldiers held accountable?



501

94

2.8K



**Clement C** @Clement33968103 · Dec 7

Unknown





# Conclusion

This analysis finds no evidence of a major Chinese Government-backed campaign to influence Australian social media in 2020, even during low points in the Australia-China relationship. There are pro-Chinese Government accounts and some of these exhibit suspicious behaviour, but they are not numerous enough or active enough to give the impression of a concerted, state-backed campaign.

Some of these accounts propagate a ‘whataboutism’ that equates Chinese human rights abuses as being no worse than those that have occurred in Australia or elsewhere. This arguably suits the Chinese Government and the circulation of this kind of content makes it more difficult to hold it to account for its abuses and territorial provocations.

This conclusion seems at odds with other studies by the Australian Strategic Policy Institute (ASPI) that claim to have found evidence that “indicates that actors linked to the Chinese Government may have been running covert information operations on Western social media platforms for at least two years”,<sup>10</sup> or that “large-scale influence campaigns linked to Chinese state actors on Twitter and Facebook” represent an “effort to influence global public opinion [that] is ongoing and is unlikely to relent.”<sup>11</sup>

However, closer reading of these studies finds several points of explanation. First, the headlines of the ASPI studies seem to overstate what their analysis revealed. Their initial study on Hong Kong is based on just 112 accounts and 1,600 tweets, which the authors point out represented “a relatively small and hastily assembled [campaign] rather than a sophisticated information campaign planned well in advance.”

Second, the ASPI studies are largely based on samples of already suspended accounts supplied by Twitter. Instead of working with data already known to be suspicious, we collected a full sample of data around events that were rapidly unfolding at the time. Given the different nature of the samples, it is perhaps unsurprising to draw different conclusions.

Finally, the ASPI studies are based on analysis of topics that more directly involve social stability in China, particularly the 2019 Hong Kong protests. These are more likely to directly threaten social stability in China and therefore be of more interest to Chinese state-backed actors. By contrast, our case studies relate to Australian COVID-19 diplomacy and the shocking contents of the Brereton Report. These events in turn riled the Chinese

---

<sup>10</sup> Uren et al (2019) *Tweeting through the Great Firewall: Preliminary analysis of PRC-linked information operations against the Hong Kong protests*, <https://www.aspi.org.au/report/tweeting-through-great-firewall>

<sup>11</sup> Wallis et al (2020) *Retweeting through the great firewall: A persistent and undeterred threat actor*, <https://www.aspi.org.au/report/retweeting-through-great-firewall>

Government and drew a distasteful response, but they hardly compare with huge street protests in a major city in terms of impact on Chinese society.

This is not to downplay the presence of, or potential for, malign Chinese Government influence in Australia. Espionage and other influence operations have been regularly reported for many years.<sup>12</sup> Indeed, such reports are in the news at time of writing.<sup>13</sup>

But overstating the extent of Chinese Government interference is counterproductive for national security, and seemingly becoming more common.<sup>14</sup> Research like the ASPI China Twitter reports, or at least the headline reporting of them, similarly overstates a problem that according to this research finds no evidence for.

---


<sup>12</sup> See for example Greene (2016) *Chinese spies in Australia on the rise, former diplomat Chen Yonglin says*, <https://www.abc.net.au/news/2016-11-20/how-many-spies-does-china-have-in-australia/8041004>

<sup>13</sup> Rice (2022) *ICAC finds true source of \$100,000 cash donation to NSW Labor was gambling account of Chinese billionaire*, <https://www.theaustralian.com.au/nation/icac-finds-true-source-of-100000-cash-donation-to-nsw-labor-was-gambling-account-at-star-city-of-chinese-billionaire-huang-xiangmo/news-story/86e94329c9654d3cca99792d162d309d>

<sup>14</sup> Belot (2022) *ASIO boss says politicising national security 'not helpful' after PM accuses Labor MP of being 'Manchurian candidate'*, <https://www.abc.net.au/news/2022-02-16/prime-minister-labels-labor-mp-manchurian-candidate/100837380>

# Appendix

Convergence plot for topic model of top replying accounts



## COORDINATION NETWORKS

Drawing on the methods established by Keller et al. (2019), we construct ‘co-tweet’ and ‘co-retweet’ networks. Co-tweets are defined as two accounts that both tweet the exact same text within a one minute window of time. Likewise, co-retweets are defined as two accounts as both retweet the same tweet with a one minute window. In this way, the nodes in such networks are *accounts* and edges or links in the network are weighted by the number of times both accounts co-tweeted (for a co-tweet network) or co-retweeted (for a co-retweet network). Additionally, for the co-retweet network we set the minimum edge weight to two, meaning that accounts must have co-retweeted *at least twice* in order to be included. This greatly reduces the prevalence of false positives, i.e. accounts that randomly happened to retweet the same tweet within one minute of each other. We do not apply this to co-tweets - with some exceptions (e.g. very short tweets that contain only a hashtag or a common phrase), it is unlikely that two accounts post exactly the same text.

Secondly, we construct ‘co-link’ networks, following the methodology of Giglietto et al. (2020)<sup>15</sup>. The objective is to detect patterns of coordinated link sharing, where for example multiple accounts all share the same news article in a short window of time. In the co-link network, nodes are accounts and two accounts are connected if they both shared the same URL within a one-minute window, *at least twice*. As with the co-retweet network, setting a minimum edge weight of 2 greatly reduces random chance co-links and focuses more specifically on accounts that show a pattern of coordinated behaviour.

Thirdly, we construct ‘co-reply’ networks, which to our knowledge is the first time this approach has been proposed. In this network, nodes are accounts and two accounts are connected if they both replied to the same tweet within a one-minute window, at least twice.

For each network (16 coordination networks in total; four network types by four datasets), we calculate the community clusters using the Louvain method (Blondel, Guillaume, Lambiotte, & Lefebvre, 2008). Following the approach of Freelon and Lokot (2020), we then filter each network to only include the top 10 communities by number of nodes. This lets us focus on exploratory analysis of the most important clusters of coordinated activity, and the key accounts in each cluster measured by degree (i.e. sum of co-retweets).

The subsequent exploratory analysis focuses on the accounts and tweets contained in the top 10 communities for each network. For accounts, we focus on the accounts’ tweet activity in the dataset, profile picture, header picture, bio description, as well as their typical tweeting behaviour through a close-reading of their tweet history (tweets and replies timeline). Where necessary, we collect additional account-level data using the Twitter API, in order to undertake more in-depth analysis of particular accounts and/or clusters of accounts.

---

<sup>15</sup> We implement our own approach as the *coornet* R package is only suited to CrowdTangle Facebook data.



# References

Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10), P10008.

Bruns, A. (2016). Big Data Analysis. In T. Witschge, C. W. Anderson, D. Domingo, & A. Hermida (Eds.), *The Sage Handbook of Digital Journalism* (pp. 509–527). Sage.

Bruns, A. (2017). Making Audience Engagement Visible: Publics for Journalism on Social Media Platforms. In B. Franklin & S. A. Eldridge II (Eds.), *The Routledge Companion to Digital Journalism Studies* (pp. 325–334). Routledge.

Bruns, A., & Stieglitz, S. (2012). Quantitative Approaches to Comparing Communication Patterns on Twitter. *Journal of Technology in Human Services*, 30(3–4), 160–185.  
<https://doi.org/10.1080/15228835.2012.744249>

Giglietto, F., Righetti, N., Rossi, L., & Marino, G. (2020). It takes a village to manipulate the media: coordinated link sharing behavior during 2018 and 2019 Italian elections. *Information, Communication and Society*, 1–25.  
<https://doi.org/10.1080/1369118X.2020.1739732>

Molter, V., & DiResta, R. (2020). Pandemics & propaganda: how Chinese state media creates and propagates CCP coronavirus narratives. *Harvard Kennedy School Misinformation Review*, 1(3).

Sauter, T., & Bruns, A. (2015). #auspol: The Hashtag as Community, Event, and Material Object for Engaging with Australian Politics. In N. Rambukkana (Ed.), *Hashtag Publics: The Power and Politics of Discursive Networks* (pp. 47–60). Peter Lang.

Wallis, J., Uren, T., & Thomas, E. (2020). Retweeting Through the Great Firewall: A Persistent and Undeterred Threat Actor. Australian Strategic Policy Institute, Policy Brief (No. 33, pp. 2020-06). Report.