

Public Sentiment on Chinese Social Media during the Emergence of COVID-19

YINGDAN LU

JENNIFER PAN

YIQING XU

Stanford University, USA

When COVID-19 first emerged in China, there was speculation that the outbreak would trigger public anger and weaken the Chinese regime. By analyzing millions of social media posts from Sina Weibo made between December 2019 and February 2020, we describe the contours of public, online discussions pertaining to COVID-19 in China. We find that discussions of COVID-19 became widespread on January 20, 2020, consisting primarily of personal reflections, opinions, updates, and appeals. We find that the largest bursts of discussion, which contain simultaneous spikes of criticism and support targeting the Chinese government, coincide with the January 23 lockdown of Wuhan and the February 7 death of Dr. Li Wenliang. Criticisms are directed at the government for perceived lack of action, incompetence, and wrongdoing—in particular, censoring information relevant to public welfare. Support is directed at the government for aggressive action and positive outcomes. As the crisis unfolds, the same events are interpreted differently by different people, with those who criticize focusing on the government’s shortcomings and those who praise focusing on the government’s actions.

Keywords: *China, COVID-19, crisis, social media, public opinion, deep learning*

Lu: yingdan@stanford.edu

Pan: jp1@stanford.edu

Xu: yiqingxu@stanford.edu

Date submitted: 2021-04-02

Introduction

Between December 2019 and February 2020, a virus outbreak in central China emerged as a once-in-a-century health crisis. As the novel coronavirus disease 2019 (COVID-19) spread, there was speculation that the outbreak and its impact on the Chinese economy would pose an existential threat to the Chinese government, or at the very least, damage the legitimacy and diminish the power of China's top leader, Xi Jinping.¹ Commentators suggested that similarities between the Chinese government's suppression of information regarding COVID-19 and the 2002-2003 SARS outbreak would inflame public anger and reveal the fundamental failures of China's political system.² News articles highlighted discontent fomenting on Chinese social media, including comparisons of China to the USSR before the Soviet collapse and ridicule aimed at China's top leaders.³

This paper provides a look at public sentiment on Chinese social media during the emergence of the COVID-19 public health crisis. When COVID-19 first emerged in China, what was the content of public discussions pertaining to the outbreak? Was there an outpouring of online dissatisfaction, anger, and ridicule directed at the Chinese government, as many commentators suggested at the time? Or did public support accompany stringent measures adopted by the Chinese government to control the disease, as research on political support during crisis would predict?⁴ We answer these and other related questions by quantitatively describing online sentiment related to COVID-19 between December 1, 2019, and February 27, 2020. We analyze data from Sina Weibo, China's largest microblogging platform, using approximately 1 million original Weibo posts from the Weiboscope dataset and 4.3 million Weibo posts from the Weibo-COV dataset.⁵ We rely on human annotation

¹See, for example, <https://wapo.st/3gO59s7>, <https://bit.ly/2LGkRdh>, <https://bit.ly/3miyJHn>, <https://on.wsj.com/3gR5WZb>, and <https://bit.ly/2OiuZURPDF>.

²What we now know of the emergence of COVID-19 in December 2019 and January 2020 suggests that the state-of-the-art disease management systems China put into place after the SARS outbreak of 2002–2003 did not overcome problems in the incentive structure of the Chinese political system that cause lower-level officials to conceal and manipulate information (Gao, 2016; Tsai, 2008; Pan and Chen, 2018; Wallace, 2016).

³See <https://wapo.st/3gTERhQ> and <https://on.wsj.com/3gR5WZb>.

⁴For literature on crisis and political support, see Ashworth et al. (2018); Fair et al. (2017); Gasper and Reeves (2011); Healey (2002).

⁵We downloaded the Weiboscope dataset from <https://bit.ly/2KwK8WQ> and the Weibo-COV dataset from <https://bit.ly/3gRTKAY> in August 2020. Replication materials for this paper can

and deep learning—specifically the Chinese Bidirectional Encoder Representations from Transformers (BERT) with the Whole Word Masking model—to describe the content and sentiment of Weibo discussions.

Content of Discussions on Weibo

Sina Weibo is a Chinese microblogging platform founded in 2009, a few months after Twitter was blocked in mainland China. Like Twitter, Sina Weibo allows users to broadcast short messages publicly. The platform skyrocketed in popularity after its founding, becoming one of China's most popular social media platforms in the early 2010s (Pan, 2017). In 2013, the Chinese government arrested a number of influential Weibo users who had exposed corruption and stirred up interest in social issues using the platform.⁶ That same year, Tencent's WeChat, a multi-purpose social media platform, launched its payment system and eclipsed Sina Weibo as China's top social media platform.⁷ Though Sina Weibo is no longer China's dominant social media platform, it remains highly popular for broadcast communication and public deliberation, and government crackdowns on influential Weibo users continue. Both Weibo and WeChat are used for interpersonal communication and news consumption, but unlike WeChat, Weibo enables public, multi-way communication between strangers, allowing for broad dissemination of information (Guo and Zhang, 2020). At the beginning of 2020, Sina Weibo had 550 million monthly active users.⁸

The use of online social networks after crisis events is well documented around the world (Heverin and Zach, 2012; Stieglitz et al., 2018). When COVID-19 emerged, we would expect individuals to turn to online spaces such as Weibo to try to make sense of the outbreak, to seek information, and to make social connections.⁹ At the same time, we would expect the Chinese government to try to make its voice heard on social media. The Chinese government has stated that social media is the preferred channel for direct government-to-public communication (Jiang and Fu, 2018; Pan, 2020; Schlaeger and Jiang,

be downloaded from the Harvard Dataverse: <https://doi.org/10.7910/DVN/ZIIQUG>.

⁶See <https://nyti.ms/3cbki4L>.

⁷See <https://bit.ly/3t3rFCb> and <https://bit.ly/38pyeao>.

⁸See <https://bit.ly/2OB8YGW>.

⁹Wei et al. (2021) show online information seeking behavior during the emergence of COVID-19.

2014; Wang and Dickson, 2019).¹⁰ The Chinese Communist Party (CCP) describes social media platforms as crucial channels for the regime to guide public opinion, especially in response to “sudden and unexpected events,”¹¹ and the key social media platforms the government has emphasized for this work are Weibo and WeChat (Lu and Pan, 2021). COVID-19 no doubt qualified as a “sudden and unexpected” event. These factors motivate the first three research questions:

RQ1: When did discussions of COVID-19 begin on Weibo?

RQ2: What is the content of Weibo posts related to COVID-19 during this initial period?

RQ3: How does the content of Weibo posts vary between different types of users?

Sentiment of Discussion on Weibo

Although a myriad of expressions and types of information likely accompany emerging crises, we are particularly interested in criticism and support. In times of crisis, researchers have documented a process of social representation where the public criticizes specific targets, attributing blame and responsibility, to make sense of events (Malhotra and Kuo, 2008; Mayor et al., 2013; Washer and Joffe, 2006). At the same time, theories also suggest people may increase their support for institutions that offer an actual and/or symbolic sense of safety in times of crisis.¹² Because people want to view the world as secure and predictable, they double down on support for existing institutions when faced with a salient threat such as an unexpected disaster. How the media frames responsibility can influence these processes (Hallahan, 1999; Holton et al., 2012; Iyengar, 1996; Sun et al., 2021). Understanding the contours of criticism and support is thus of theoretical importance, which leads to the next two research questions:

¹⁰See 2013 State Council “Opinions of the General Office of the State Council on Further Strengthening Government Information Disclosure in Response to Social Concerns and Enhancing Government Credibility,” <https://bit.ly/3c3rTSI>.

¹¹See 2018 State Council “Opinions on Promoting the Healthy and Orderly Development of New Media in Government Affairs.”

¹²Examples include terror-management theory (Greenberg et al., 1997), motivated-social-cognition theory (Jost et al., 2003), and uncertainty-management theory (Van den Bos et al., 2005).

RQ4: How do the shares of critical and supportive posts pertaining to COVID-19 vary in this initial period?

RQ5: Who are the targets of criticism and support?

Our primary interest resides in the levels of criticism and support for the Chinese government. An older body of research argues that crises threaten the stability of governments, especially authoritarian regimes (Carlin et al., 2014; Diamond et al., 1989; Drury and Olson, 1998; Huntington, 1991; Olson and Drury, 1997; Pelling and Dill, 2010; Przeworski and Limongi, 1997). Crises can change the behavior of the core supporters and coalitions autocrats need to stay in power (Haggard and Kaufman, 1997; Magaloni, 2006; Pepinsky, 2009), in turn weakening those in power or forcing powerholders to take steps to consolidate their rule (Geddes et al., 2018; Landry and Stockmann, 2009; Manion et al., 2020). These theories predict that crises would be accompanied by criticisms of those in power. More recent research, however, suggests that whether crises affect the stability of governments depends on the governments' response. Natural disasters have been found to increase votes and support for incumbent politicians when voters see the government responding strongly (Ashworth et al., 2018; Fair et al., 2017; Gasper and Reeves, 2011; Healey, 2002). Likewise, we can imagine authoritarian governments gaining support if their response to unexpected disasters is perceived as strong. We may expect to see such gains in support in a context like China where there are no viable institutions other than the CCP—no organized political opposition, no national-level religious institutions, no prominent foreign or international institutions. The regime is the only source of safety, so even if crises reveal fundamental flaws in the political system, such events may be accompanied by calls to support the existing political regime. This motivates the final research question:

RQ6: How does the share of critical and supportive posts targeting the Chinese government vary in this initial period?

Data and Methods

To answer these questions, we analyze two publicly available COVID-19 datasets containing Weibo content, whose key characteristics are summarized in Table 1. The first dataset,

the Weiboscope COVID-19 dataset, includes Weibo posts collected pre-censorship (Fu and Zhu, 2020). If censors decide that a discussion on Weibo needs to be removed, there exists a window of time before posts are fully removed (King et al., 2013; Zhu et al., 2013), and the Weiboscope data takes advantage of this time gap to generate its dataset.

Table 1: Data Sources.

	Weiboscope	Weibo-COV
Pre-censorship	Yes	No
Checks post-censorship status	Yes	No
Number of seed users for data collection	118,394	20,000,000
Seed users with verified accounts	45-55%	8%
Number of COVID-19 keywords	40	179
Data collection method	Weibo Open API	Scraping
Number of original posts	1 million	4.3 million
User information masked	Yes	Yes
User metadata available	No	Yes

Researchers who maintain the Weiboscope dataset follow the Weibo accounts of 66,126 high-profile users and 52,268 randomly chosen users (Zhu and Fu, 2020). High-profile users in this dataset include individuals and organizations with verified, or VIP, status on Weibo as well as other opinion leaders.¹³ Randomly chosen users in the Weiboscope dataset were identified by generating random user IDs (Fu and Chau, 2013). Overall, we estimate that 45-55% of the Weiboscope sample are verified users such that the composition of the sample is heavily tilted toward online elites and organizational accounts, including government and state-controlled media accounts. Researchers collect all posts made by these accounts in 15- to 20-minute intervals using the Sina Weibo Open API (Fu et al., 2013). After initial data collection, Weiboscope makes repeated calls to the Weibo Open API to

¹³We estimate that more than 80% of high-profile users in the Weiboscope dataset have verified account status and more than 10,000 followers each. This estimate is based on descriptions of Weiboscope data from 2013—see Fu et al. (2013)—and may have changed since. The remaining high-profile users in this dataset include Chinese dissident writers, journalists, and scholars.

check the status of posts and identify censorship. The Weiboscope COVID-19 dataset is a subset of the 11 million Weiboscope posts collected from December 1, 2019, to February 27, 2020, that contain keywords related to the coronavirus outbreak in China.¹⁴ The dataset contains the text of posts, the date and time of posting, the censorship status, the hashed post ID, and the hashed original post ID if the post is a repost.

The second dataset, the Weibo-COV dataset, contains posts collected retrospectively in April 2020 (Hu et al., 2020). Because content targeted for censorship is usually removed within 24 hours, the Weibo-COV dataset should be considered post-censorship. The Weibo-COV dataset is based on a pool of 20 million active users (Shen et al., 2020).¹⁵ Active users were identified as those with more than 50 followings, followers, posts, and a post within a 30-day window. All uncensored Weibo posts made between December 1, 2019, and April 30, 2020, by the active users were screen scraped. Posts were then filtered to include only those that contained one or more of 179 COVID-19-related keywords (see Table S1 for full list of words and Table S2 for the English translation). The content of the Weibo-COV dataset is not dominated by elite users. As Table 1 shows, only 8% of seed users have verified account status. The Weibo-COV data we obtained contains the text of the post, the date and time of posting and crawling, the hashed user ID, the hashed post ID, the number of likes, the numbers of comments and reshares of each post, the hashed original post ID if the post is a repost, and users' verification status.¹⁶

We rely on the Weibo-COV dataset for our main analysis because it contains data from a broader, more diverse sample of users and a broader set of keywords. We also conduct all analyses on the Weiboscope dataset to ensure that our results are not biased by censorship (See Appendix Section A.4). We focus on the time period covered by both datasets—December 1, 2019, to February 27, 2020—when the virus first emerged in China and before it spread globally. We focus on this time period because after February, COVID-19 quickly spreads around the world, and public attention in China and public attitudes

¹⁴See Table S1 and Table S2 in the Appendix for the keywords used in both datasets.

¹⁵Researchers identified 20 million active users by first starting with 5 million Weibo users, whom the creators of the Weibo-COV dataset had identified from non-COVID-19 projects. Then they expanded the user set by identifying whom the 5 million followed and were followed by, and repeating this procedure with new users, ultimately identifying 250 million unique users.

¹⁶We requested the verification status data.

toward COVID-19 are influenced by what happens abroad. Finally, our analysis is based on original posts. Both datasets also contain reposts, but there is often insufficient information in the content of reposts to determine what content or sentiment is being expressed.¹⁷ In total, we analyze 1,011,275 original posts from the Weiboscope COVID-19 dataset and 4,310,459 original posts from the Weibo-COV dataset.

Content of Weibo Posts

To tackle our first research question of when discussions of COVID-19 begin on Weibo, we train a classifier to identify whether or not a post is about COVID-19. We define a post as being about COVID-19 if its main topic is about the coronavirus, SARS-CoV-2, or the resulting disease, COVID-19.¹⁸ We include posts that talk about the emergence of the outbreak as well as the ramifications of the disease.¹⁹ To train the classifier, we randomly sample approximately 5,000 posts from each dataset (10,541 posts in total), stratified by post creation date. This sample is used for a number of analyses, and we refer to it as the *Hand-Annotated Sample*. Six native Chinese speakers, who trained for one month to achieve higher intercoder reliability, hand annotated posts.²⁰ Specifically, three research assistants annotated whether or not a post was about COVID-19, and the final annotation was based on the majority agreement.

We then use this training data to fine-tune the pretrained Chinese BERT with the Whole Word Masking model (Chinese BERT-wwm-ext). BERT is a deep learning algorithm containing pre-trained deep bidirectional representations that has been shown to outperform other state-of-the-art language models (Devlin et al., 2019). We use the Chinese BERT-wwm-ext model rather than the Chinese BERT-BASE model because the Chinese BERT-wwm-ext model applies upgraded whole word masking on Chinese text and uses more data sources for training (Cui et al., 2020). Fine-tuning is done by testing a range of values for different hyper-parameters (steps, learning rate, class weight, dropout rate)

¹⁷The data collection methods of both Weiboscope and Weibo-COV did not always yield the content of the original post being reposted.

¹⁸See Table S3 in the Appendix for detailed coding rules.

¹⁹If a post only contains a COVID-19-related keyword or hashtag without other COVID-19-related content, we code it as *not* related to COVID-19.

²⁰See Appendix Section A.2 for details on human annotation.

and selecting hyper-parameters to maximize F1 score, which balances precision and recall. Our fine-tuned Chinese BERT-wwm-ext model for classifying COVID-19 content has an accuracy of 0.97 and an F1 score of 0.99.

To address our second and third research questions of the content of Weibo posts pertaining to COVID-19, we identify five mutually exclusive categories of substantive interest: *Government Announcements*, *Domestic News*, *International News*, *Health Advice*, and *Personal Opinion*. *Government Announcements* include all content pertaining to government or CCP regulations, policies, and public health guidance; infections information such as report of cases, sites of outbreak, and deaths; as well as actions and accomplishments attributed to government/Party bureaus and officials at all levels.²¹ *Domestic News* includes reporting on potential treatments, vaccine development, impact on the Chinese economy and society; stories of public compliance and non-compliance; and pieces describing the work of model medical workers, companies, and organizations combating the outbreak. Posts belonging to the domestic news category include posts that contain the full text of news articles and posts that only contain a news headline. *International News* refers to posts pertaining to the spread and control of COVID-19 in other countries as well as Macau, Hong Kong, and Taiwan. *Health Advice* refers to health-related content such as tips on how to dispose of used masks and how to maintain well-being during lockdown. *Personal Opinion* includes personal reflections, updates, opinions (including opinions about government policies and news), expressions of emotions other than criticism and support, and calls for assistance.²²

If a post contains content described in the *Government Announcements* category and content described in *Domestic News*, *International News*, or *Health Advice*, we place it in the *Government Announcements* category. For example, a news post about a local government's decision to close down the city would be a *Government Announcement*. This is because we are particularly interested in the prevalence of government-focused content. If a post contains both *Government Announcements* and *Personal Opinion*, the post is put into *Government Announcements* unless the personal opinion is a substantial portion of the post.²³ We do this to differentiate between more mechanical resharing of government

²¹ Examples of accomplishments include those describing model cadres combating coronavirus.

²² We analyze critical and supportive commentary separately. See Section 2.2.

²³ More than 10 characters are focused on public opinion out of a max 140 characters allowed in a

announcements, which we place into *Government Announcements*, and commentary on government announcements, which we place into *Personal Opinion*. If a post contains some combination of *Domestic News*, *International News*, *Health Advice*, and *Personal Opinion*, then the post is put in the category that is most emphasized in the post. Note that these categories are based on the content and not the source of posts, which is not provided in either dataset. In other words, government announcements are not classified based on whether they are posted by government accounts but rather because their content emphasizes government-related information.

We rely solely on human annotation of the *Hand-Annotated Sample* to analyze how content falls into these categories. The data is imbalanced across the five content categories, which would have necessitated an extremely time-consuming expansion of the human-coded training data in order to increase precision and recall of supervised machine learning methods. We did not use unsupervised machine learning approaches such as topic modeling because they did not yield results that were as substantively precise and as meaningful as the human annotations.

For the third research questions of content by user type, we divide users into ordinary users, celebrities, and organizations. The Weibo-COV dataset contains the verification status of users in their dataset. Organizations are users verified as “Blue V,” which is a designation given to government offices, public institutions, state media outlets, enterprises, non-government organizations, and foreign organizations. Celebrities are users verified as “Yellow V,” which is a designation given by Sina Weibo to influential individuals, including entertainment celebrities, online influencers, reporters, government officials, and other key opinion leaders. All other accounts are what we label as ordinary users.

Sentiment of Weibo Posts

To answer our fourth research question on the shares of critical and supportive posts pertaining to COVID-19, we train two classifiers to identify 1) COVID-19-related posts containing criticism and 2) COVID-19-related posts containing support. We define a post as criticism if it contains a negative evaluation of a target; projects negative emotions toward a target;

denigrates a target in some way; or blames the target for negative characteristics, attributes, and outcomes. Note that our definition of criticism always entails a target, which means that general negative sentiments such as fear or undirected anger are out of scope.²⁴ The same goes for support, where we define a post as supportive if it contains a positive evaluation of a target; projects positive emotions or praise toward a target; or associates the target with positive characteristics, attributes, and outcomes.

The two classifiers are trained based on human annotation of the *Hand-Annotated Sample* of whether posts contain criticism or support.²⁵ Like before, this training data is used to fine-tune two Chinese BERT-wwm-ext models, which predict posts containing criticism and posts containing support, respectively. For classification of criticism and support, we compare the performance of different machine learning and deep learning models, where we find that BERT-wwm-ext outperforms all other models.²⁶ After tuning hyper-parameters, including steps, learning rate, class weight, and dropout rate, the final model for criticism reaches an accuracy of 0.98 and an F1 score of 0.78; the final model for support reaches an accuracy of 0.98 and F1 score of 0.71.

Targets of Criticism and Support

To answer the fifth research question on targets of criticism and support, we identify whether any of eight entities are targets of criticism or support in posts in the *Hand-Annotated Sample*. Any particular post may contain one or more targets. We say a target belongs to the *Central Government* if it refers to the central Chinese government, the Chinese Communist Party, China's political institutions or system, central leaders or officials, or the Chinese Centers for Disease Control (CDC) and National Health Commission (NHC). We say a target pertains to the *Local Government* if it refers to provincial- and lower-level Chinese governments, leaders, and other officials. We consider *Healthcare Workers and*

²⁴We focus on targeted criticism for theoretical and practical reasons. Theoretically, it may be easier for an autocrat to redirect amorphous negativity and harness general negativity to its benefit, whereas that may be harder to do with targeted criticism. Practically, even for our human coders, it is difficult to establish consistent and reliable coding rules for what constitutes non-directed criticism, e.g., should fear be included, should sadness, etc.

²⁵Final annotation based on majority agreement between three research assistants. See Appendix Section A.2 for details on human annotation.

²⁶See Appendix A.3 for comparisons of different classification algorithms.

Organizations to be a target if the post focuses on healthcare workers (doctors, nurses, other medical personnel) or healthcare organizations, including hospitals, pharmaceutical institutes and research centers, and the pharmaceutical industry. We say a target pertains to *Chinese Media* if there is criticism or support for any channel of domestic Chinese media (e.g., newspaper, Weibo) or their affiliated workers (e.g., reporters, editors). We say a target is the *Chinese Public* if the post focuses on persons who are not affiliated with any of the aforementioned categories (government, healthcare, media), including students, patients, enterprise owners, and all other ordinary citizens in China. We say a target falls into the *Other Organizations* category if the criticism or support targets other types of organizations (e.g., enterprises, real estate, airlines, NGOs) in China. We say a target belongs to the *Other Countries and Regions* category if the target is a government, military, political leader or official, or media and people from other nations, Hong Kong, Macau, or Taiwan. We say a post has an *Indeterminant Target* if the post refers to broad concepts such as “people in all walks of life.” We rely solely on human annotation of the *Hand-Annotated Sample* to identify targets because our fifth research question is not focused on trends over time. Because of the large number of the target categories and imbalance between these categories, in order to obtain reasonable accuracy with machine classification, we would have had to substantially increase our training sample.

For our sixth research question, on criticism and support targeting the Chinese government, we are interested in trends over time. To enable overtime analyses, we train two final classifiers to identify 1) posts containing criticism directed at the government, and 2) posts containing support directed at the government. To maximize precision and recall, our classification of posts targeting the government combines central and local levels of government and encompasses the CCP, government officials, and bureaucracies. Similar to our previous trained models, these two classifiers are trained based on human annotation of the *Hand-Annotated Sample*.²⁷ Like before, this training data is used to fine-tune two Chinese BERT-wwm-ext models. The trained models reach an accuracy of 0.85 and an F1 score of 0.74 for criticism toward the Chinese government, and an accuracy of 0.93 and an F1 score of 0.85 for support.

²⁷Final annotation based on majority agreement between three research assistants. See Appendix Section A.2 for details on human annotation.

Finally, to provide additional insight into support and criticism targeting the government, we hand annotate a sample we call the *Event Investigation Sample*. We randomly sample 200 critical posts and 200 supportive posts from each of the two largest spikes of discussion related to COVID-19: January 23, when China's central government issued a lockdown order to temporarily suspend all public transportation in Wuhan and to prevent people from leaving the city;²⁸ and February 7, the day when Dr. Li Wenliang, a prominent critic of the way the government handled the COVID crisis, died. Dr. Li tried to raise the alarm about COVID-19 by sharing information on Chinese social media in December 2019 about growing infections of atypical pneumonia in Wuhan, but he was accused by the Wuhan government of spreading rumors and was silenced. Following the same coding rules as described for research question five, we use human annotation to differentiate whether criticism or support targets the central government or local governments.

Findings

We identify 3,142,178 posts (72.9%) in the Weibo-COV dataset related to COVID-19²⁹ and 948,519 posts (93.8%) in the Weiboscope dataset related to COVID-19. It is expected that keyword-based datasets will contain posts unrelated to COVID-19. This is because posts may contain COVID-19-related keywords but not refer to the outbreak. For example, the term “epidemic” is used by both Weibo-COV and Weiboscope to collect data, but in some posts “epidemic” refers to other outbreaks, such as the pneumonic plague. As the number of keywords expands, recall will increase but precision will decrease (Zhang and Pan, 2019), and that is what we observe: Weibo-COV contains more false positives than Weiboscope since Weibo-COV uses a substantially higher number of keywords. This is why it is crucial for the validity of our findings that we apply our machine learning classifier to identify COVID-related posts.

RQ1: When did discussions of COVID-19 begin on Weibo? After we identify posts related to COVID-19, Figure 1 shows that widespread Weibo discussion of COVID-19 started on January 20, 2020 (results based on Weiboscope data are extremely similar;

²⁸See <https://www.bbc.com/news/world-asia-china-51217455>.

²⁹We exclude posts related to fundraising and donations pertaining to COVID-19. During our time period, there were a number of popular COVID-related fundraising drives where Weibo messages were auto-generated after a donation was made as a way of generating new donations.

see Appendix Figure S1). On December 31, 2019, and January 1, 2020, there are small bumps in discussions of COVID-19. December 31 is when the Wuhan Municipal Health Commission announced that cases of atypical pneumonia identified in the city were viral and that an investigation of the cause of infection was underway. On January 1, the Weibo account of the Wuhan police department posted that the police had summoned eight people accused of spreading “rumors” online about atypical pneumonia.

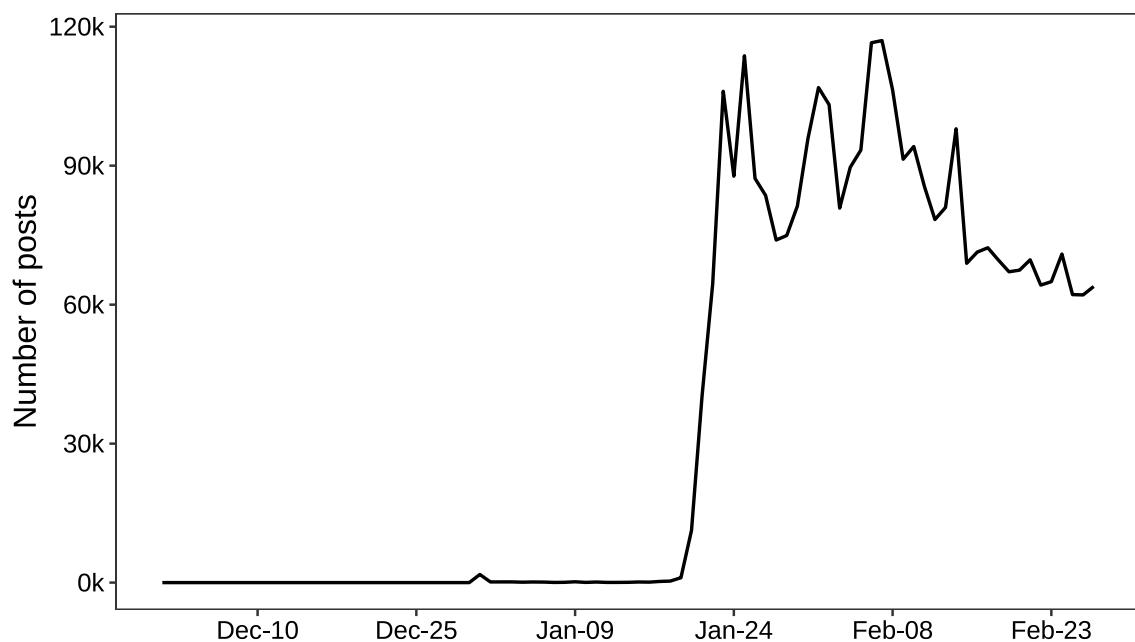


Figure 1. Number of COVID-19 Posts (Weibo-COV).

Note. Number of posts classified as related to COVID-19 from December 1, 2019, to February 27, 2020.

On January 20, 2020, Weibo discussions of COVID-19 went “viral” as Chinese scientists confirmed human-to-human transmission of the virus. This spike grows on January 23, 2020, the day the Wuhan lockdown was established. This large volume of discussion is sustained through early February and decreases slightly in both datasets by the end of February.

RQ2: What is the content of Weibo posts related to COVID-19 during this initial period? Based on the *Hand-Annotated Sample*, Table 2 shows the breakdown of categories of COVID-19-related posts for the Weibo-COV data.³⁰ The largest share of content is personal opinion at 68% of total original posts related to COVID-19. This is followed by government announcements (9%), domestic news (4%), international news (1%), and health advice (1%).

Table 2: Content of COVID-19 Posts (Weibo-COV).

	All users	Ordinary users	Celebrities	Organizations
Govt. announcements	9%	8%	3%	46%
Domestic news	4%	3%	5%	16%
International news	1%	1%	3%	0%
Health advice	1%	1%	1%	4%
Personal opinions	68%	70%	69%	24%

RQ3: How does the content of Weibo posts vary between different types of users?

When we break down content by user type—ordinary users, celebrities, and organizations—we find that ordinary users and celebrities are for the most part sharing their personal opinions (70% of ordinary users, 69% of celebrities), while organizational accounts—e.g., accounts of government offices, state-media outlets—are predominantly sharing government announcements. Organizational accounts are also sharing personal opinions (24%) and domestic news (16%), and to a lesser extent health advice (4%). Eight percent of ordinary users are sharing government announcements and 3% domestic news. Only 3% of celebrities are sharing government announcements and 5% sharing domestic news.

RQ4: How do the shares of critical and supportive posts pertaining to COVID-

³⁰Table S7 in the Appendix shows the content breakdown of COVID-19 posts based on the Weiboscope dataset. Table S7 shows a much higher proportion of government announcements. This is not due to censorship. It may be driven by the high share of verified accounts in Weiboscope, which includes state-controlled media outlets and government accounts more likely to post government announcements. Note that the Weiboscope dataset does not contain any user information, so it cannot be used to answer research question three.

19 vary in this initial period? We find similar proportions of critical and supportive commentary related to COVID-19. In the post-censorship Weibo-COV dataset, 6.4% of posts related to COVID-19 contain criticism, while 8.1% of predicted posts contain support. These results do not appear to be the result of censorship. In the pre-censorship Weiboscope dataset, 5.1% of posts related to COVID-19 contain criticism and 6.8% of COVID-19 posts contain support.

Figure 2 shows the number of critical and supportive posts in both datasets over time. The top panel of Figure 2 shows the number of critical and supportive posts in the Weibo-COV dataset, and the bottom panel of Figure 2 shows the number of critical and supportive posts in the Weiboscope dataset for comparison.³¹ As the pop-out boxes in the bottom left corner of both panels of Figure 2 show, there is a small spike in critical discussion at the end of December 2019 when Wuhan released information about cases of what was then called “atypical pneumonia” and stated that it was investigating. Critical posts either targeted individuals for spreading rumors and misinformation around this outbreak, or blamed the Wuhan government for causing an outbreak by not providing more timely and accurate information. A spike in support coincided with this burst of criticism. Supportive posts expressed appreciation for healthcare workers and called for trust in the central government and the Wuhan government in handling an emerging public health crisis.

A large spike in critical commentary occurs on January 23, 2020, when China’s central government implemented the lockdown of Wuhan. A spike in support appears two days later in the Weibo-COV dataset, on January 25, coinciding with when Xi Jinping set up a Central Committee Leading Group for epidemic response and 30 provinces across China initiated protocols for the highest level of public health emergency response.³²

³¹We examine reshares of critical and supportive commentary, and do not find that critical comments are reshared more than supportive comments. Trends in criticism and support look similar when we examine how diffuse criticism and support are, both in terms of the number of unique users engaged in making critical and supportive posts (see Figure S4 in the Online Appendix) and the number of critical and supportive posts per user per day (see Figure S5 in the Online Appendix).

³²See <https://china.huanqiu.com/article/3y3nz4LCXRB>.

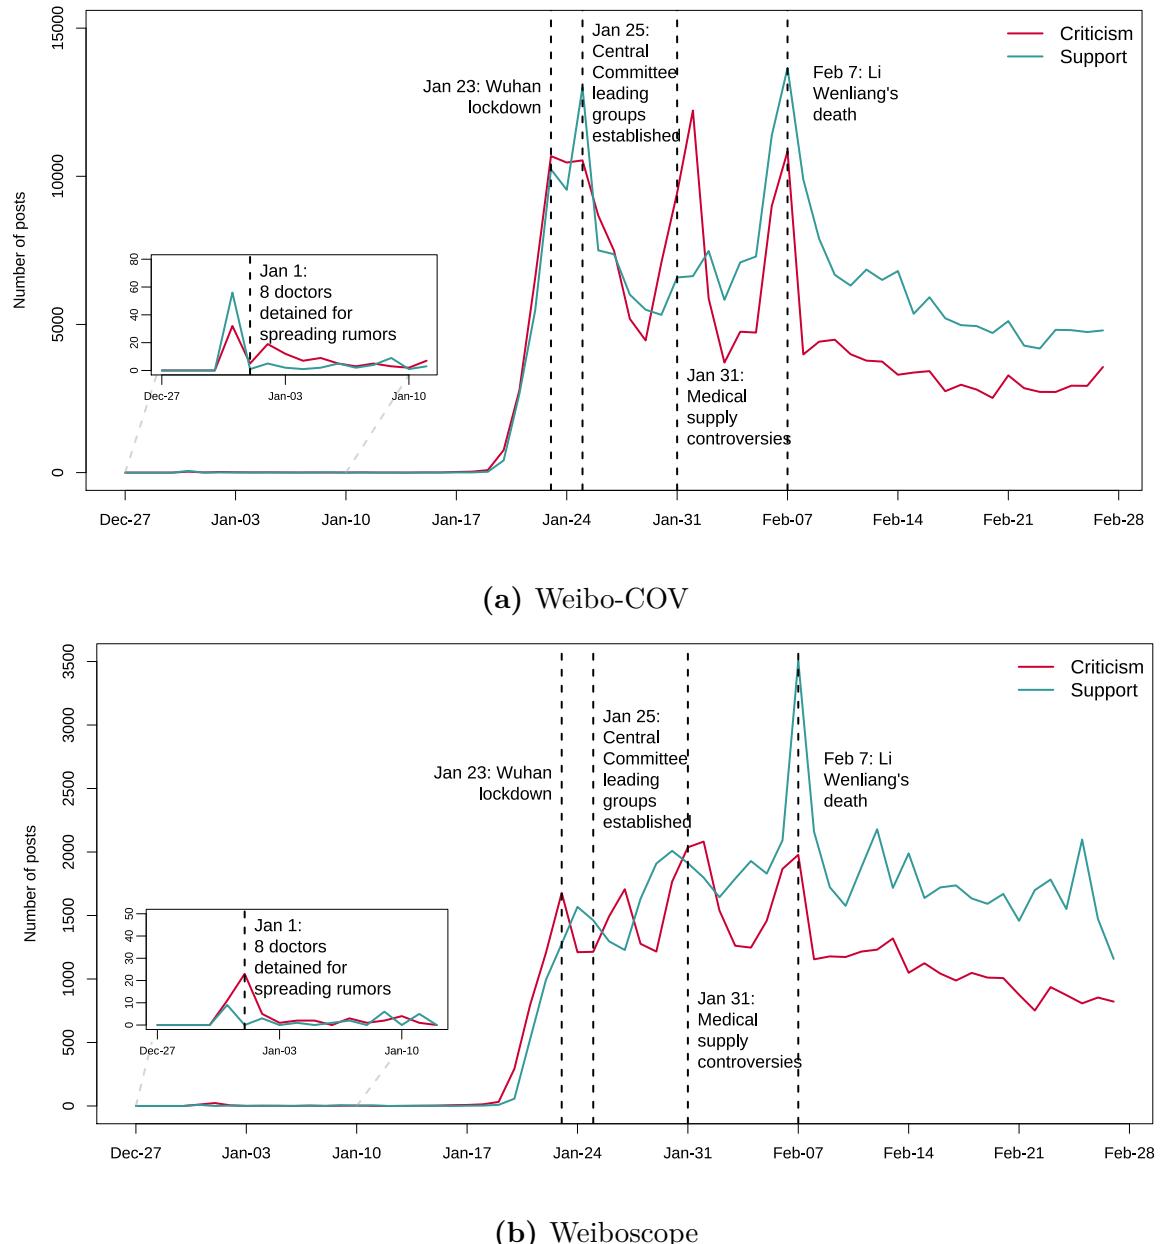


Figure 2. Number of Critical and Supportive Posts.

Note. Number of posts containing critical and supportive sentiments in the Weibo-COV dataset (panel a) and the Weiboscope dataset (panel b).

A second spike in critical commentary occurs between January 31 and February 2, 2020. This spike focuses on two controversies related to medical supplies. In one, the Red Cross Society of Hubei was criticized on Weibo for colluding with private hospitals and causing supply shortages at hospitals designated for COVID-19 treatment.³³ The second controversy erupted over the Chinese herbal remedy Shuanghuanglian, which Chinese state media reported as a possible treatment for the virus. Following this reporting, 29% of Weibo-COV critical posts contained the keyword “Shuanghuanglian.” These posts criticize the Chinese media and medical experts for hawking this “cure” to drive up stock prices and criticize the public for lining up to buy the remedy and thereby spreading the disease.

The last spike of criticism appears in both datasets on February 7, 2020, the day Dr. Li Wenliang died of COVID-19. This critical spike is accompanied by a large burst of support, containing an outpouring of appreciation for healthcare workers and organizations.

RQ5: Who are the targets of criticism and support? Based on human coding of the *Hand-Annotated Sample*, Table 3 shows the breakdown of targets of criticism based on the Weibo-COV data. Aggregating all critical and supportive posts from our time period, we find that the main targets of criticism are the Chinese public (46%) and local governments (25%).³⁴ The Chinese public is criticized for not adhering to public health guidelines (e.g., not staying at home, not wearing masks), for consuming exotic animals, and for spreading misinformation and engaging in discrimination (e.g., against people from Wuhan and Hubei). Local governments, including but not limited to those of Wuhan and Hubei Province, are criticized for not responding quickly or strongly enough to control the outbreak.

The main target of support are healthcare workers and organizations (48%), including doctors, nurses, medical experts, and those working to develop and produce treatments, medical supplies, and vaccines. Local governments are praised for taking steps to mitigate the outbreak (17%), and the central government is praised for pushing out a national plan

³³See <https://bit.ly/34qmWRa>.

³⁴Since one post can contain multiple targets, the sum of the percentages of each column calculated in Table 3 does not equal 100%. For results based on Weiboscope, see Table S8 in the Appendix. In the Weiboscope dataset, the main targets of criticism are also the public (31%) and local governments (32%).

and effectively coordinating across the country (16%). There is also a sizable proportion of praise directed at ordinary Chinese people for complying with public health guidance, and for volunteering and making donations to support those in need (14%).³⁵

Table 3: Targets of Criticism and Support (Weibo-COV).

	Criticism	Support
Central government	6%	16%
Local government	25%	17%
Chinese public	46%	14%
Healthcare workers and organizations	3%	48%
Chinese media	4%	0.3%
Other countries and regions	4%	4%
Other organizations	10%	5%
Indeterminate target	8%	10%

RQ6: How do the shares of critical and supportive posts targeting the Chinese government vary in this initial period? Figure 3 shows how the proportions of critical and supportive posts targeting the Chinese government vary over time.³⁶ Table 4 differentiates between criticism and support targeting the central government and that targeting local governments for the two largest spikes in the data—January 23 and February 7—based on the *Event Investigation Sample* of human-coded data from Weibo-COV.

³⁵ Healthcare workers and organizations are also the main target of support in the Weiboscope dataset (47%), followed by local government (20%), the public (20%), and central government (9%); see Table S8.

³⁶ Results based on the Weiboscope dataset show similar general trends in criticism; the Weiboscope data contains two out of four corresponding bursts of support; see Appendix Figure S2. We examine reshares or critical and supportive commentary targeting the government, and do not find that critical comments are consistently reshared more than supportive comments. Trends also look similar when we look at how diffuse criticism and support are, both in terms of the number of unique users engaging in making posts critical and supportive of the government (see Figure S6 in the Online Appendix) and the number of critical and supportive posts targeting the government per user per day (see Figure S7 in the Online Appendix).

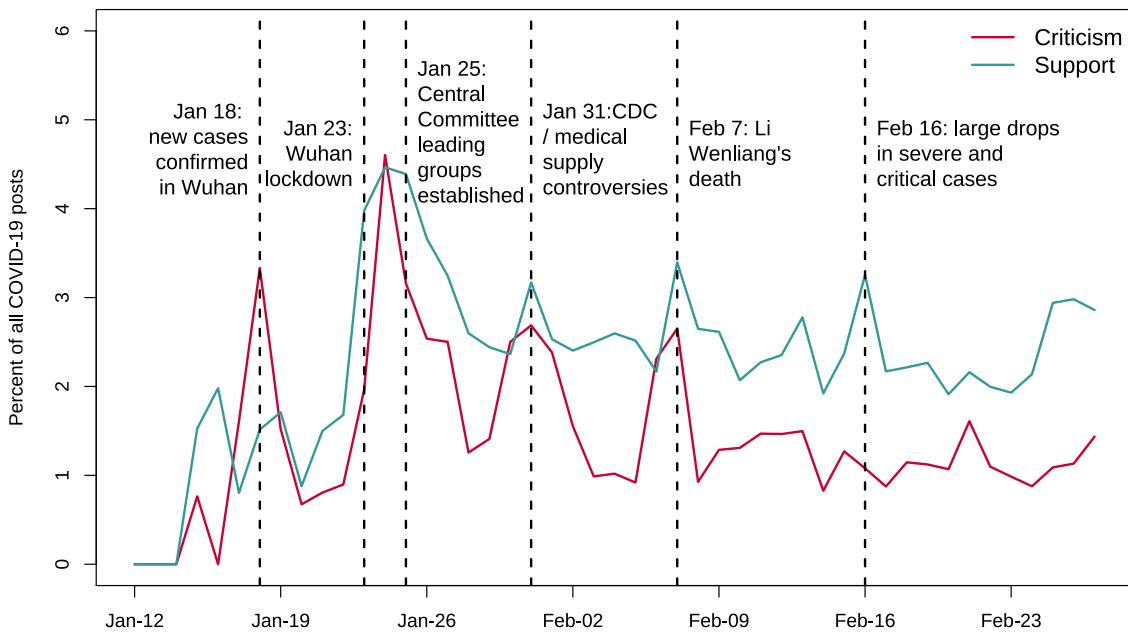


Figure 3. Share of criticism and support targeting Chinese government.

Note. Percentage of critical and supportive posts targeting the Chinese government among all COVID-19 related posts (Weibo-COV).

Table 4: Criticism and Support Targeting Central and Local Governments on January 23 and February 7 (Weibo-COV).

	January 23		February 7	
	Criticism	Support	Criticism	Support
Central government	3%	22%	11%	9%
Local government	18%	13%	25%	16%

Figure 3 reveals four spikes in criticism against the Chinese government on January 18, January 24, January 31, and February 7. On January 18, the Wuhan Municipal Health Commission reported four new infections after not reporting any new cases between January 11 and January 17 when the Hubei Provincial People's Congress and People's Consultative Congress meetings were held. This inconsistency triggered an avalanche of criticism directed

at the Wuhan municipal government—for example:

#FourNewConfirmedCasesInWuhan Isn't there a contradiction between this announcement in the early morning of January 18 and the “no new cases” conclusion on January 15? It was said that [those infected] had received treatment from January 11 to 13...why weren't these new cases announced on January 15?

On January 23, the day of the Wuhan lockdown, criticism focuses on local governments and officials in Wuhan and Hubei Province for not doing enough to control the outbreak. The most-reshared critical post on January 23, 2020, shared 4,741 times, is a post by someone whose elderly relative is awaiting diagnosis in Wuhan:

I want to know when he can get a nucleic acid test for diagnosis. When can he get admitted into the ward for treatment? Are you waiting until he is dead? Maybe if he is dead, then there's no diagnosis, right? After all, one more case is not good news for certain people and certain government departments!
@ChangjiangDaily @WuhanEveningNews

Outside of Wuhan, users also complain about how their local governments are not doing enough to control the spread of the disease. Table 4 shows that on January 23, criticism is primarily directed at local governments (18% of all critical posts) while there is very little criticism of the central government (3%).

On January 31, when medical supply controversies erupted, critical posts blame local governments for ineffectual coordination of medical resources and poor performance. One such post says:

#WuhanUnionHospitalWestCampusDebunkingRumors Hubei government is a disaster at every level!! The party secretary, the provincial head, the deputy mayor, the Municipality Health Commission, and the Red Cross, everything sucks. I'm furious! So angry!

On February 7, criticisms are directed at various levels of government, including the central government, for a persistent culture of censorship and information manipulation and its damaging effects on public welfare. One post, reshared 1,203 times, says:

The letter of rebuke³⁷ is a shameful stain on China's annals of disease prevention. He [Li] almost saved our entire country. Remember this hero, and also remember that bunch of bastards, those public servants with their endless lies. You will be punished with death by a thousand cuts.

Table 4 shows that on February 7, the amount of criticisms targeting local governments is high (25% of all critical posts), but strikingly, the level of criticism targeting the central government is nearly four times higher on February 7 than it was on January 23 (11% of all critical posts).

Figure 3 reveals four spikes in support of the Chinese government: January 23-25, January 31, February 7, and February 16. The January 23-25 burst of support directed at the Chinese government praises a number of government actions, including the implementation of public health emergency protocols and stringent infection control measures in provinces such as Henan and Zhejiang as well as the extremely fast construction of COVID-19 hospitals in Hubei. Supportive posts in this burst also call on the public to trust in the central government and the CCP in defeating the pandemic. For example, hashtags like #LearnFromHenan circulate, praising the Henan provincial government for its extreme measures, such as destroying roads and barricading villages, to control the spread of COVID-19. Table 4 shows that on January 23, the level of support for the central government is high (22% of all supportive posts), and local governments also receive some support (13% of all supportive posts).

On January 31, support directed at the Chinese government is mostly focused on how the World Health Organization (WHO) lauded the Chinese government for “their commitment to transparency, and the efforts made to investigate and contain the current

³⁷This refers to the letter issued by Wuhan police rebuking Li Wenliang for “spreading rumors” and forcing him to promise never to do it again.

outbreak” on January 30.³⁸ One of the most reshared posts, posted by a famous travel blogger on Weibo with more than 5 million followers and shared 1,109 times, said:

The World Health Organization officially declared that [the outbreak in] China has become a #PHEIC [Public Health Emergency of International Concern] in order to reduce the transnational spread of the disease. The PHEIC period is three months and then it expires automatically. I believe China has the capability to contain the disease, and even the WHO praises China’s response to the outbreak as admirable! I hope everyone can work together to overcome this difficulty! Go China!

On February 7, the spike in supportive commentary targeting the government is primarily directed at local governments. In particular, there was much praise of the support provincial governments around China were providing to cities in Hubei Province that were facing escalating case loads and resource shortages, such as sending healthcare workers and providing critical supplies. One such post said:

#16ProvincesSupportHubei I trust in the leadership of Zhejiang government and also its capability. Zhejiang will join hands with Jingmen to overcome difficulties. When spring flowers bloom, we’ll be enjoying ourselves by the West Lake. Go us!

Table 4 shows that on February 7, support directed at local governments increases from the level of support observed on January 23 (16% of all supportive posts). In contrast, the level of support for the central government is less than half of what it was on January 23 and falls below support for local governments (9% of all supportive posts).

On February 16, support for the Chinese government centers around how Tedros Adhanom Ghebreyesus, the WHO Director-General, praised China for “slowing the spread to the rest of the world.”³⁹ Out of 2,322 posts supporting Chinese government on that day,

³⁸See <https://bit.ly/3l1VnEH>.

³⁹See <https://bit.ly/3eljPj6>.

670 (29%) show support for the Chinese government by referencing the headline “WHO Director-General: China’s Measures to Control the Novel Atypical Pneumonia Epidemic from the Source are Encouraging.”

Discussion

Weibo posts about COVID-19 became widespread on January 20, 2020. Expressions of personal opinion—personal updates, reflections, and opinions, as well as appeals and calls for assistance—flooded Weibo, coming from both ordinary users and users with verified celebrity accounts. Government and state-media accounts, as we expected, were actively disseminating government announcements, but this government-related content represented less than 10% of Weibo posts during this early period.

These findings reveal several patterns in criticism and support for the Chinese regime during this public health crisis. First, bursts of criticism and support often occur simultaneously in reaction to events. When Wuhan was locked down on January 23, when controversies erupted over Red Cross distribution and an herbal remedy on January 31, and when Dr. Li Wenliang died on February 7, we observe simultaneous bursts of general criticism and support as well as simultaneous bursts of criticism and support targeting the Chinese government. In other words, events affected online public sentiment, but not in a single direction.

Why is this the case? Criticisms are directed at the government for perceived lack of action, incompetence, and wrongdoing—in particular, concealing information relevant to public welfare. Support directed at the government rewards government action and good outcomes. This aligns with existing theories that political support in times of crisis that predict strong actions are rewarded and inaction is penalized. However, what is different in this case is that the same event can trigger both sets of perceptions simultaneously. As a crisis is unfolding, the same event is interpreted differently by different people. In the simultaneous bursts of criticism and support, those who criticize point to the government’s shortcomings and those who praise point to the government’s actions and successes.

In these simultaneous bursts of criticism and support, we detect a pattern of public sentiment unusual in the study of China. Between the two largest bursts of discussion

in this time period—the January 23 Wuhan lockdown and February 7 death of Dr. Li Wenliang—criticisms of the central government increase dramatically while support for the central government declines precipitously, falling, by February 7, below the level of support directed at local governments. Over the past decades of China research, numerous studies have consistently found that trust in and satisfaction with the central government is high, and always higher than that for local governments (Jiang and Yang, 2016; Lu and Dickson, 2020; Shi, 2001; Tang, 2005, 2016; Truex and Tavana, 2019; Zhou et al., 2020). With the death of Dr. Li Wenliang, we observe a deviation from this pattern. Weibo users are taking the central government to task for censorship and information manipulation, questioning a core feature of the Chinese political system. However, this deviation does not last; the level of support for the government steadily increases after February 7, while the rate of criticism falls.

This raises the question, relevant to all studies of social media in China, of whether the content and sentiment we observe reflect organic expression or Chinese government efforts to reshape the online information environment. In terms of censorship, because the Weiboscope dataset is pre-censorship and includes measures of post-hoc content removal, we can measure the rate of censorship. Focusing on what is most likely to be censored—critical commentary—we find an extremely low rate of content removal. Out of 47,912 critical posts made during our study period, only 55 posts (0.11%) were censored, and out of 11,901 critical posts targeting the government, only 27 (0.23%) were removed. While the Chinese government may use other strategies to censor online discourse—e.g., intimidating or co-opting key opinion leaders, preventing content containing certain keywords or phrases from successfully posting to Weibo—our data suggest that large-scale, post-hoc removal of critical commentary did not occur.

In contrast, we cannot determine the extent to which the spikes in supportive commentary are organic or artificially generated by the Chinese government or created through government-led mobilization (King et al., 2017; Perry et al., 2020). In our data, it is not the case that a small subset of users are making large numbers of supportive posts. Criticism and support targeting the government are similarly diffuse and there are similar numbers of unique users posting commentary critical and supportive of the government (see Figure S7 and Figure S6). However, this does not rule out government manipulation. Previous re-

search shows that the Chinese government mobilizes government officials and CCP cadres to fabricate social media posts as if they were the opinions of ordinary people. Fabricated Weibo content is often made by real people using their real, personal accounts, making it practically impossible to determine for certain whether content is authentic or inorganically generated. In the Weibo-COV data, we observe organizational accounts—accounts of government agencies and state media—dramatically increasing supportive posts as a proportion of all posts between January 18 and February 7, and increasing the proportion of supportive posts even more after February 7 (see Figure S8). Investigative reports make note of Chinese government efforts to artificially manipulate online sentiment in early to mid-February,⁴⁰ but we do not know from our data whether overt efforts by government-affiliated accounts to broadcast supportive sentiment were accompanied by coordinated, covert efforts to flood and reshape the information environment. It is possible that the low level of support for the central government observed on February 7 goes away because of government information manipulation. It is also possible that the Chinese government's aggressive disease control measures and accompanying drops in cases turned the tide of public opinion. Or it may be some mixture of these and other forces.

Beyond these substantive findings, this study has implications for practices in measuring sentiment on social media. The descriptive findings of this paper suggest that if social media data is used to study public sentiment, researchers should resist an impulse to focus on negative sentiment and proactively measure both positive and negative sentiment. Doing so can have substantive payoffs. For example, scholars of Chinese politics have noted a seeming disconnect between surveys showing high levels of political trust in China and analyses of Chinese social media data showing high frequencies of protest and discontent (Cai, 2010; Chen, 2012; Heurlin, 2020; King et al., 2013; Zhang and Pan, 2019). However, survey-based designs often, by default, incorporate both positive and negative sentiments. When asking about satisfaction with the government, survey response scales range from “very dissatisfied” to “very satisfied,” capturing both ends of the spectrum. This is not the case with computational analyses of social media data. To find positive sentiment, researchers must proactively use or develop classifiers. Our results, based on measurement of positive and negative sentiments, suggest that there may not be as large a disconnect be-

⁴⁰Reporting by the *New York Times* and ProPublica; see <https://nyti.ms/3arKeK1>.

tween social media measures of government support and survey-based measures as previous studies might have implied.

Finally, this study has practical, real-world implications. Whenever a crisis occurs in an authoritarian regime, there seems to be an immediate impulse among commentators to find an outpouring of public anger. While capturing negative sentiment is important to our understanding of these countries in times of crisis, the findings of this paper caution against narrowly focusing on negative sentiment.

Acknowledgements

Our thanks to Mengmeng Guo, Zeyu Li, Bonnie Yue Liang, Jinpeng Wang, Gou Wu, and Yingzi Yang for research support; King-wa Fu and Anfan Chen for facilitating access to data; Melanie Manion, Katerina Tertychnaya, Tian Yang, and participants of the UCL Departmental Research Seminar, Stanford Data Science Lab, and Quantitative China Studies Seminar at NYU for many helpful comments and suggestions; and to the National Science Foundation (Grant # 1934578) and the Stanford Graduate Fellowship for research support. Replication materials for this paper can be downloaded from the Harvard Dataverse: <https://doi.org/10.7910/DVN/ZIIQUG>.

References

Ashworth, S., Bueno de Mesquita, E., and Friedenberg, A. (2018). Learning about voter rationality. *American Journal of Political Science*, 62(1):37–54.

Cai, Y. (2010). *Collective Resistance in China: Why Popular Protests Succeed or Fail*. Palo Alto, CA: Stanford University Press.

Carlin, R. E., Love, G. J., and Zechmeister, E. J. (2014). Natural disaster and democratic legitimacy: the public opinion consequences of Chile’s 2010 earthquake and tsunami. *Political Research Quarterly*, 67(1):3–15.

Chen, X. (2012). *Social Protest and Contentious Authoritarianism in China*. New York: Cambridge University Press.

Cui, Y., Che, W., Liu, T., Qin, B., Wang, S., and Hu, G. (2020). Revisiting pre-trained

models for Chinese natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 657–668, Online. Association for Computational Linguistics.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Diamond, L. J., Linz, J. J., Lipset, S. M., et al. (1989). *Democracy in Developing Countries: Latin America*. Boulder, CO: Lynne Rienner.

Drury, A. C. and Olson, R. S. (1998). Disasters and political unrest: an empirical investigation. *Journal of Contingencies and Crisis Management*, 6(3):153–161.

Fair, C. C., Kuhn, P. M., Malhotra, N., and Shapiro, J. N. (2017). Natural disasters and political engagement: evidence from the 2010–11 Pakistani floods. *Quarterly Journal of Political Science*, 12(1):99–141.

Fu, K.-w., Chan, C.-h., and Chau, M. (2013). Assessing censorship on microblogs in China: discriminatory keyword analysis and the real-name registration policy. *IEEE Internet Computing*, 17(3):42–50.

Fu, K.-w. and Chau, M. (2013). Reality check for the Chinese microblog space: a random sampling approach. *PloS ONE*, 8(3):1–7.

Fu, K.-w. and Zhu, Y. (2020). Did the world overlook the media's early warning of COVID-19? *Journal of Risk Research*, 23(7-8):1047–1051.

Gao, J. (2016). Bypass the lying mouths: how does the CCP tackle information distortion at local levels? *The China Quarterly*, 228:950–969.

Gasper, J. T. and Reeves, A. (2011). Make it rain? Retrospection and the attentive electorate in the context of natural disasters. *American Journal of Political Science*, 55(2):340–355.

Geddes, B., Wright, J. G., Wright, J., and Frantz, E. (2018). *How Dictatorships Work: Power, Personalization, and Collapse*. New York: Cambridge University Press.

Greenberg, J., Solomon, S., and Pyszczynski, T. (1997). Terror management theory of self-esteem and cultural worldviews: empirical assessments and conceptual refinements. In *Advances in Experimental Social Psychology*, volume 29, pages 61–139.

Guo, L. and Zhang, Y. (2020). Information flow within and across online media platforms: an agenda-setting analysis of rumor diffusion on news websites, Weibo, and WeChat in China. *Journalism Studies*, 21(15):2176–2195.

Haggard, S. and Kaufman, R. R. (1997). The political economy of democratic transitions. *Comparative Politics*, 29(3):263–283.

Hallahan, K. (1999). Seven models of framing: implications for public relations. *Journal of Public Relations Research*, 11(3):205–242.

Healey, M. A. (2002). The fragility of the moment: politics and class in the aftermath of the 1944 Argentine earthquake. *International Labor and Working-Class History*, 62:50–59.

Heurlin, C. (2020). Fighting for every inch of land: greed and grievance in petition mobilization in Zhejiang. *Modern China*, 46(4):400–432.

Heverin, T. and Zach, L. (2012). Use of microblogging for collective sense-making during violent crises: a study of three campus shootings. *Journal of the American Society for Information Science and Technology*, 63(1):34–47.

Holton, A., Weberling, B., Clarke, C. E., and Smith, M. J. (2012). The blame frame: media attribution of culpability about the MMR–autism vaccination scare. *Health Communication*, 27(7):690–701.

Hu, Y., Huang, H., Chen, A., and Mao, X.-L. (2020). Weibo-COV: A large-scale COVID-19 social media dataset from Weibo. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.

Huntington, S. P. (1991). Democracy's third wave. *Journal of Democracy*, 2(2):12–34.

Iyengar, S. (1996). Framing responsibility for political issues. *The Annals of the American Academy of Political and Social Science*, 546(1):59–70.

Jiang, J. and Yang, D. L. (2016). Lying or believing? Measuring preference falsification from a political purge in China. *Comparative Political Studies*, 49(5):600–634.

Jiang, M. and Fu, K.-W. (2018). Chinese social media and big data: Big data, big brother, big profit? *Policy & Internet*, 10(4):372–392.

Jost, J. T., Glaser, J., Kruglanski, A. W., and Sulloway, F. J. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin*, 129(3):339.

King, G., Pan, J., and Roberts, M. E. (2013). How censorship in China allows government criticism but silences collective expression. *American Political Science Review*, 107(2):326–343.

King, G., Pan, J., and Roberts, M. E. (2017). How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review*, 111(3):484–501.

Landry, P. F. and Stockmann, D. (2009). Crisis management in an authoritarian regime: media effects during the Sichuan earthquake. In *APSA 2009 Toronto Meeting Paper*.

Lu, J. and Dickson, B. (2020). Revisiting the eastonian framework on political support: assessing different measures of regime support in mainland china. *Comparative Politics*, 52(4):671–701.

Lu, Y. and Pan, J. (2021). Capturing clicks: how the Chinese government uses clickbait to compete for visibility. *Political Communication*, 38(1-2):23–54.

Magaloni, B. (2006). *Voting for Autocracy: Hegemonic Party Survival and its Demise in Mexico*. New York: Cambridge University Press.

Malhotra, N. and Kuo, A. G. (2008). Attributing blame: the public's response to Hurricane Katrina. *The Journal of Politics*, 70(1):120–135.

Manion, M., Rothschild, V., and Zhu, H. (2020). Viral politics: what a public health crisis reveals about government credibility in authoritarian China. *Available at SSRN 3593561*.

Mayor, E., Eicher, V., Bangerter, A., Gilles, I., Clémence, A., and Green, E. G. (2013). Dynamic social representations of the 2009 H1N1 pandemic: shifting patterns of sense-making and blame. *Public Understanding of Science*, 22(8):1011–1024.

Olson, R. S. and Drury, A. C. (1997). Un-therapeutic communities: a cross-national analysis of post-disaster political unrest. *International Journal of Mass Emergencies and Disasters*, 15(2):221–238.

Pan, J. (2017). How market dynamics of domestic and foreign social media firms shape strategies of internet censorship. *Problems of Post-Communism*, 64(3-4):167–188.

Pan, J. (2020). Temporality alignment: how WeChat transforms government communication in Chinese cities. *Chinese Journal of Communication*, 13(3):241–257.

Pan, J. and Chen, K. (2018). Concealing corruption: how Chinese officials distort upward reporting of online grievances. *American Political Science Review*, 112(3):602–620.

Pelling, M. and Dill, K. (2010). Disaster politics: Tipping points for change in the adaptation of sociopolitical regimes. *Progress in Human Geography*, 34(1):21–37.

Pepinsky, T. B. (2009). *Economic crises and the breakdown of authoritarian regimes: Indonesia and Malaysia in comparative perspective*. New York: Cambridge University Press.

Perry, E., Ekiert, G., and Yan, X. (2020). *Ruling by Other Means: State-Mobilized Movements*. New York: Cambridge University Press.

Przeworski, A. and Limongi, F. (1997). Modernization: theories and facts. *World Politics*, 49(2):155–183.

Schlæger, J. and Jiang, M. (2014). Official microblogging and social management by local governments in China. *China Information*, 28(2):189–213.

Shen, C., Chen, A., Luo, C., Zhang, J., Feng, B., and Liao, W. (2020). Using reports of symptoms and diagnoses on social media to predict COVID-19 case counts in mainland China: observational infoveillance study. *Journal of Medical Internet Research*, 22(5):e19421.

Shi, T. (2001). Cultural values and political trust: a comparison of the People's Republic of China and Taiwan. *Comparative Politics*, 33(4):401–419.

Stieglitz, S., Bunker, D., Mirbabaie, M., and Ehnis, C. (2018). Sense-making in social media during extreme events. *Journal of Contingencies and Crisis Management*, 26(1):4–15.

Sun, Y., O'Neill, L., Liu, M., DeLuca, K. M., and Bowen, B. (2021). Mobilizing the public in saving the bonneville salt flats: understanding blame as a psychological construct. *Environmental Communication*, 15(3):301–319.

Tang, W. (2005). *Public Opinion and Political Change in China*. Palo Alto, CA: Stanford University Press.

Tang, W. (2016). *Populist Authoritarianism: Chinese Political Culture and Regime Sustainability*. Oxford: Oxford University Press.

Truex, R. and Tavana, D. L. (2019). Implicit attitudes toward an authoritarian regime. *The Journal of Politics*, 81(3):1014–1027.

Tsai, L. (2008). Understanding the falsification of village income statistics. *The China Quarterly*, 196:805–826.

Van den Bos, K., Poortvliet, P. M., Maas, M., Miedema, J., and Van den Ham, E.-J. (2005). An enquiry concerning the principles of cultural norms and values: the impact of uncertainty and mortality salience on reactions to violations and bolstering of cultural worldviews. *Journal of Experimental Social Psychology*, 41(2):91–113.

Wallace, J. L. (2016). Juking the stats? Authoritarian information problems in China. *British Journal of Political Science*, 46(1):11–29.

Wang, Y. and Dickson, B. (2019). How corruption investigations undermine regime support: evidence from China. *Available at SSRN 3086286*.

Washer, P. and Joffe, H. (2006). The “hospital superbug”: social representations of MRSA. *Social Science & Medicine*, 63(8):2141–2152.

Wei, L., Yao, E., and Zhang, H. (2021). Authoritarian responsiveness and political attitudes during COVID-19: Evidence from Weibo and a survey experiment. *Working Paper*.

Zhang, H. and Pan, J. (2019). Casm: A deep-learning approach for identifying collective action events with text and image data from social media. *Sociological Methodology*, 49(1):1–57.

Zhou, Y. J., Tang, W., and Lei, X. (2020). Social desirability of dissent: an IAT experiment with Chinese university students. *Journal of Chinese Political Science*, 25(1):113–138.

Zhu, T., Phipps, D., Pridgen, A., Crandall, J. R., and Wallach, D. S. (2013). The velocity of censorship: high-fidelity detection of microblog post deletions. In *22nd USENIX Security Symposium*, pages 227–240, Washington, DC.

Zhu, Y. and Fu, K.-w. (2020). Speaking up or staying silent? Examining the influences of censorship and behavioral contagion on opinion (non-) expression in China. *New Media & Society*, pages 1–22.

Appendix

Dataset Keywords

Table S1: Keywords Used by Weiboscope and Weibo-COV

	Weiboscope dataset	Weibo-COV dataset
Number of keywords	40	179
Overlapping keywords	疫情, 冠状, 确诊, 新冠, 李文亮, 疾控中心, N95, 钟南山, 防护服, 管轶, CDC, 高福, 疑似病例, 潜伏期 (14 keywords)	
Non-overlapping keywords	口罩, 病毒, 肺炎, 感染, 隔离, 防疫, 传染, 封城, 非典, 蝙蝠, 卫健委, 世卫, 重症, 李兰娟, 流行病, 华南海鲜市场, 人传人, 世界卫生组织, 消毒液, 洗手液, 危重, 张文宏, 穿山甲, 粪口传播, WHO, 飞沫传播 (26 keywords)	Cov-19, 感染人数, 大众畜牧野味店, 华南野生市场, 武汉病毒所, 中国疾病预防控制中心, #2019nCoV, 双黄连 AND 抢购, 双黄连 AND 售罄, 武汉卫健委, 湖北卫健委, #nCoV, PHEIC, 火神山, 雷神山, Coronavirus, Remdesivir, 瑞德西韦, 感染 AND 例, 武汉 AND 封城, 王延轶, 舒红兵, 协和医院, 武汉 AND 隔离, 李文亮 AND 医生, 云监工, 武汉仁爱医院, 黄冈 AND 感染者, 孝感 AND 感染者, 居家隔离, 隔离 14 天, 潜伏期 AND 24 天, 潜伏期 AND 14 天, 国际公共卫生紧急事件, 方舱医院 AND 武汉, 一省包一市, 晋江毒王, 超级传播者, 湖北 AND 王晓东, 蒋超良, 千里投毒, 武汉病毒研究, 武汉 AND 李医生, 国家疾控中心, 武汉 AND 疫苗, 武汉 AND 征用宿舍, 周佩仪, 武汉中心医院, 张晋 AND 卫健委, 张晋 AND 卫生将康委员会, 刘英姿 AND 卫健委, 刘英姿 AND 卫生健康委员会, 王贺胜 AND 卫健委, 王贺胜 AND 卫生健康委员会, 复工, 中小企业 AND 困境, 武汉 AND 死亡病例, 武汉 AND 感染病例, 湖北 AND 死亡病例, 湖北 AND 感染病例, 中国 AND 死亡病例, 中国 AND 感染病例, 北京 AND 病例, 天津 AND 病例, 河北 AND 病例, 辽宁 AND 病例, 上海 AND 病例, 江苏 AND 病例, 浙江 AND 病例, 福建 AND 病例, 山东 AND 病例, 广东 AND 病例, 海南 AND 病例, 山西 AND 病例, 内蒙古 AND 病例, 吉林 AND 病例, 黑龙江 AND 病例, 安徽 AND 病例, 江西 AND 病例, 河南 AND 病例, 湖北 AND 病例, 湖南 AND 病例, 广西 AND 病例, 四川 AND 病例, 贵州 AND 病例, 云南 AND 病例, 西藏 AND 病例, 陕西 AND 病例, 甘肃 AND 病例, 青海 AND 病例, 宁夏 AND 病例, 新疆 AND 病例, 香港 AND 病例, 澳门 AND 病例, 台湾 AND 病例, ECOM, sars-cov-2, 复学, 护目镜, 核酸检测, COVID-19, 2019-nCoV, 疑似 AND 病例, 无症状, 累计病例, 境外输入, 累计治愈, 绥芬河, 舒兰, 健康码, 出入码, 返校, 美国 AND 例, 西班牙 AND 例, 新加坡 AND 例, 加拿大 AND 例, 英国 AND 例, 印度 AND 例, 日本 AND 例, 韩国 AND 例, 德国 AND 例, 法国 AND 例, 意大利 AND 例, 新增 AND 例, 人工膜肺, 双盲测试, 疫苗, 小区出入证, 战疫, 抗疫, 湖北卫健委 AND 免职, 发热患者, 延迟开学, 开学时间 AND 不得早于, 累计死亡数, 人口排查, 武汉 AND 肺炎, 新型肺炎, 不明原因肺炎, 野味肺炎, 出门 AND 戴口罩, 3M AND 口罩, KN95 AND 口罩, 新肺炎, 2019nCoV, 新型肺炎 AND 死亡, 新型肺炎 AND 感染, 武汉 AND 肺炎 AND 谣言, 8 名散布武汉肺炎谣言, 黄冈 AND 新肺炎, 孝感 AND 新肺炎, 武汉中心医院 AND 新型肺炎, 武汉肺炎, 企业复工, 围积口罩, 零号病人, 黄燕玲, 病毒源头, 电子烟肺炎 AND 新型冠状, 病毒战, 病毒 AND 实验室泄露, 比尔盖茨 AND 疫苗牟利, 美国细菌实验室, pandemic (165 words)

Table S2: English Translation of Keywords Used by Weiboscope and Weibo-COV.

		Weiboscope dataset	Weibo-COV dataset
Number of keywords		40	179
Overlapping keywords	Epidemic, Coronavirus, Confirmed case, Novel coronavirus, Li Wenliang, Center for Disease Control and Prevention N95, Zhong Nanshan, Hazmat suit, Guan Yi, CDC, Gao Fu, Suspicious cases, Incubation period (14 keywords)		
Non-overlapping keywords		Mask, virus, pneumonia, infected, quarantine, combat the outbreak, infection, lockdown, SARS, bat, Health Commission, WHO (abbreviation in Chinese), severe, Li Lanjuan, epidemic, Huanan seafood market, human-to-human transmission, World Health Organization, bleach, hand sanitizer, critically ill, Zhang Wenhong, pangolin, fecal-oral transmission, WHO (in English), droplets transmission (26 keywords)	Cov-19, Infected cases, Dazhong wildlife restaurant, Huanan wildlife Market, Wuhan Institute of Virology, Chinese Center for Disease Control and Prevention, #2019nCoV, Shuanghuanglian AND Rush to buy, Shuanghuanglian AND Sold out, Wuhan Municipal Health Commission, Health Commission of Hubei Province, #nCoV , PHEIC, Huoshen Shan hospital, Leishen Shan hospital, Coronavirus (in English), Remdesivir (in English), Remdesivir (in Chinese), Infection AND Cases, Wuhan AND Lockdown, Wang Yanyi, Shu Hongbing, Xiehe Hospital, Wuhan AND Quarantine, Li Wenliang AND Doctor, Supervising work on cloud, Wuhan Ren'ai Hospital , Huanggang AND infected cases, Xiaogan AND infected cases, Isolated at home, Isolation AND 14 days, Incubation period AND 24 days, Incubation period AND 14 days, International Public Health Emergencies, FangCang Hospital AND Wuhan, one province gives a hand to one Hubel city, Super spreader of COVID-19 in Jinjiang, Super Spreader, Hubei AND Wang Xiaodong, Jiang Chaoliang, Spread Virus from a thousand miles, Virology research in Wuhan, Wuhan AND Li Wenliang, Chinese Center for Disease Control and Prevention, Wuhan AND Vaccine, Wuhan AND Requisitioned students' dormitory, Zhou Peiyi, The Central Hospital of Wuhan, Zhang Jin AND Health Commission, Zhang Jin AND Health Commission, Liu Yingzi AND Health Commission (abbr.), Liu Yingzi AND Health Commission, Wang Hesheng AND Health Commission (abbr.), Wang Hesheng AND Health Commission, Enterprise work resuming, Small and medium-sized enterprise AND Dilemma, Wuhan AND Death Cases, Wuhan AND Infection Cases, Hubei AND death cases, Hubei AND infection cases, China AND death cases, China AND infection cases, Beijing AND cases, Tianjin AND cases, Hebei AND cases, Liaoning AND cases, Shanghai AND cases, Jiangsu AND cases, Zhejiang AND cases, Fujian AND cases, Shandong AND cases, Guangdong AND cases, Hainan AND cases, Shanxi AND cases, Inner Mongolia AND cases, Jilin AND cases, Heilongjiang AND cases, Anhui AND cases, Jiangxi AND cases, Henan AND cases, Hubei AND cases, Hunan AND cases, Guangxi AND cases, Sichuan AND cases, Guizhou AND cases, Yunnan AND cases, Tibet AND cases, Shaanxi AND cases, Gansu AND cases, Qinghai AND cases, Ningxia AND cases, Xinjiang AND cases, Hong Kong AND cases, Macau AND cases, Taiwan AND cases, ECOM (in English), sars-cov-2 (in English), Resumption of schooling, Goggles, Nucleic acid test, COVID-19 (in English), 2019-nCoV (in English), suspected AND cases, asymptomatic, cumulative confirmed cases, Imported cases of NCP, Cumulative cured cases, Suifenhe, Shulan, health QR code, Community Access Code, Back to school, US AND cases, Spain AND cases, Singapore AND cases, Canada AND Cases, British AND cases, India AND cases, Japan AND cases, South Korea AND cases, Germany AND cases, France AND cases, Italy AND cases, New AND cases, Extracorporeal Membrane Oxygenation, Double-blind test, Vaccine, Community entry card, Anti-COVID-19, Anti COVID-19, Health commission of Hubei Province AND Remove from the position, Fever patients, Postpone the reopening of school, The start time of school AND Not earlier than, Cumulative deaths, Household troubleshoot, Wuhan AND pneumonia, Novel pneumonia, Pneumonia of unknown cause, Wildlife pneumonia, Going out and wear mask, 3M AND mask, KN95 AND masks, Novel pneumonia, 2019nCoV, Novel pneumonia AND death, Novel pneumonia AND infection, Wuhan AND pneumonia AND rumors, Eight people spreading rumors of Wuhan pneumonia, Huanggang AND novel pneumonia, Xiaogan AND novel pneumonia, The Central Hospital of Wuhan AND Novel Pneumonia, Wuhan Pneumonia, Enterprise Work resuming, Hoarding Masks, Zero Patient, Huang Yanling, Origin of Cov-19, E-cigarette Pneumonia AND Novel Coronavirus, Virus War, Virus AND Laboratory leakage, Bill Gates AND Vaccine for profit, US Army Bacterial Laboratory, pandemic (in English) (165 words)

Human Annotation

Table S3: Coding Scheme for human annotation.

Variable	Value	Rules
COVID-19 related	0	If the post contains no COVID-19 specific content If the post only contains a COVID-19-related keyword or hashtag without other COVID-19-related content If the post is an advertisement to sell masks or other medical supplies
	1	If the post contains one topic and it is about the coronavirus, SARS-CoV-2, or the resulting disease, COVID-19 If the post contains multiple topics and one of them is about the outbreak If the post describes the emergence of the outbreak If the post describes the impact of the outbreak (such as daily life during lockdown, influence on the economy, and compliance/non-compliance with regulations) If the post targets people or entities affected by the outbreak
Criticism	0	If the post contains only statements of fact If the post only contains an opinion but does not criticize If the post only contains criticism that is not about COVID-19 If the post only contains factual questions raised in a non-negative way If the post only points out problems without making any judgment If the post only contains self-criticism If the post only contains grievance direct at an amorphous target (e.g., the world, technology) or does not specify a target
	1	If the post contains negative evaluations of a target or its behavior If the post shows distrust or dissatisfaction with a target If the post questions a target in a negative, fatalistic, dismissive way, or indirectly suggests problems and failures brought by the target If the post curses, ridicules, or otherwise denigrates the target If the post contains sarcasm to criticize the target If the post implies that the problem could have been controlled or solved if the target of criticism had behaved differently
Support	0	If the post only contains only statements of fact If the post only contains an opinion but does not indicate any approval or support If the post only contains emotions that are low-valence If the post only contains self-approval If the post only contains positive evaluation without specifying a target or toward an insubstantial target (e.g. technology, world)
	1	If the post contains positive evaluation of a target or its behavior If the post praises or shows support to a target or its behavior If the post attributes positive or beneficial outcome to a target If the post shows or calls for a high level of trust toward a target

Table S4: Human Coder Agreement for All Variables.

Variable	%Agreement	Krippendorff's alpha	Total Cases
COVID-19-related	94.96	0.88	337
Criticism	89.32	0.83	337
Criticism toward government	92.88	0.72	337
Support	85.16	0.80	337
Support toward government	87.83	0.69	337

Model Training and Fine-Tuning

Using the human annotated data, we first conducted an experiment to predict critical posts in Weiboscope dataset using different algorithms. As Table S5 shows, Chinese BERT-wwm-ext outperforms other deep learning and machine learning frameworks.

Table S5: Model comparison (on Weiboscope criticism prediction).

Model	Accuracy	Precision	Recall	F1 Score
Chinese BERT-wwm-ext	0.98	0.80	0.77	0.78
CNN	0.98	0.86	0.67	0.75
RNN	0.97	0.64	0.78	0.70
SVM	0.97	0.69	0.50	0.58
Random Forest	0.97	0.75	0.50	0.60

Then, for each variable in each dataset, we fine-tuned the Chinese BERT-wwm-ext by using different hyper-parameters (steps, learning rate, class weight, dropout rate), and selected hyper-parameters to maximize F1 score, balancing precision and recall. Table S6 shows the performance of each finalized model.

Table S6: Performance of BERT Models.

Variable		Accuracy	Precision	Recall	F1 Score
Weibo-COV dataset	COVID-19-related Criticism	0.97	0.98	1.00	0.99
	Support	0.98	0.78	0.78	0.78
	Criticism toward government	0.85	0.78	0.70	0.74
	Support toward government	0.93	0.80	0.91	0.85
	COVID-19-related Criticism	0.98	0.98	1.00	0.99
Weibo-scope dataset	Support	0.98	0.80	0.77	0.78
	Support	0.97	0.88	0.76	0.81
	Criticism toward government	0.86	0.74	0.79	0.77
	Support toward government	0.92	0.76	0.72	0.74

Analysis Based on Weiboscope Data

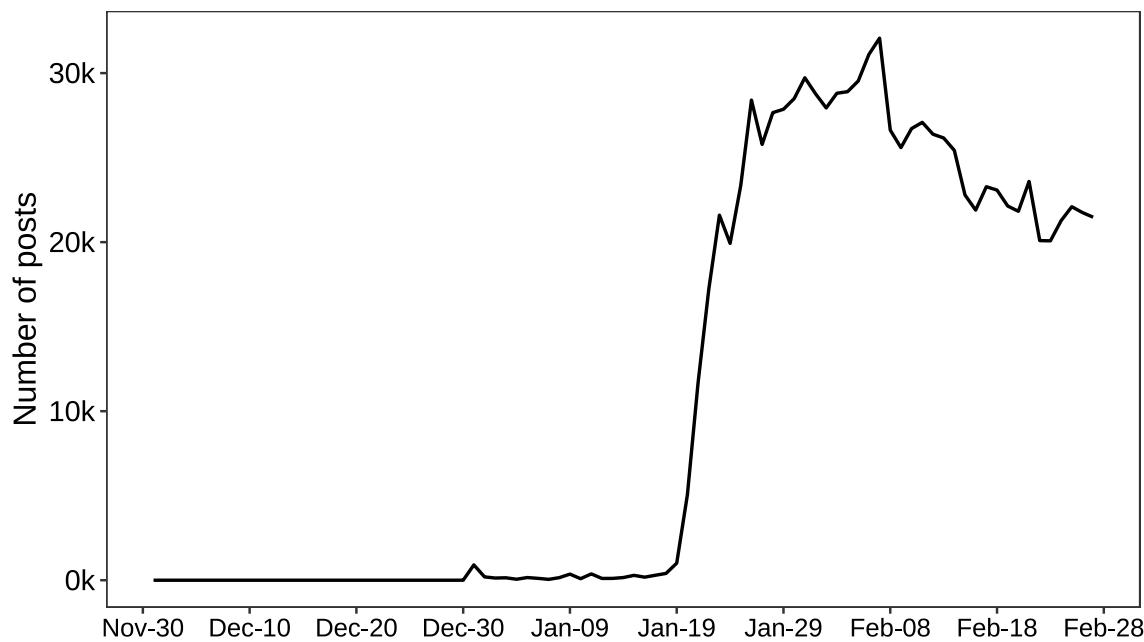


Figure S1. Number of posts related to COVID-19 (Weiboscope).

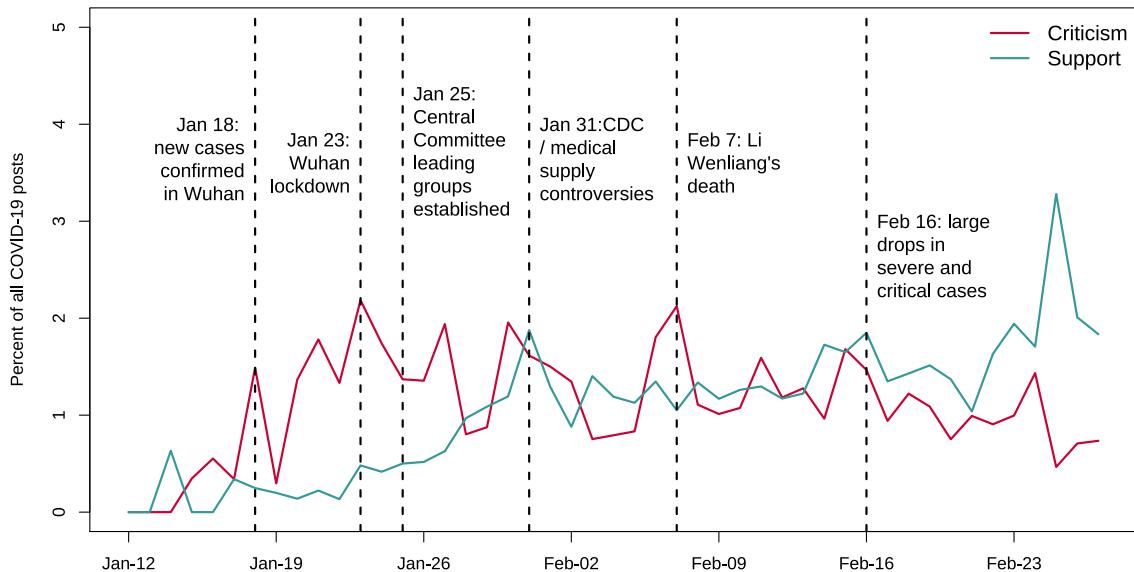
Note. This figure shows the total number of posts classified as COVID-19-related in the Weiboscope dataset from December 1, 2019 to February 27, 2020.

Table S7: Content of COVID-19 Posts (Weiboscope)

	Weiboscope
Government announcements	45%
Domestic news	14%
International news	7%
Health advice	4%
Personal opinions	15%

Table S8: Target of Criticism and Support (Weiboscope).

	Criticism	Support
Central government	5%	9%
Local government	32%	20%
The public	31%	20%
Healthcare workers & organizations	6%	47%
Chinese media	5%	1%
Other countries & regions	14%	3%
Other organizations	9%	7%
Indeterminate target	9%	3%

**Figure S2. Share of criticism and support targeting Chinese government (Weiboscope).**

Note. This figure shows the percentage of critical and supportive posts targeting the Chinese government among all COVID-19-related posts in the Weiboscope dataset from January 12, 2020 to February 27, 2020.

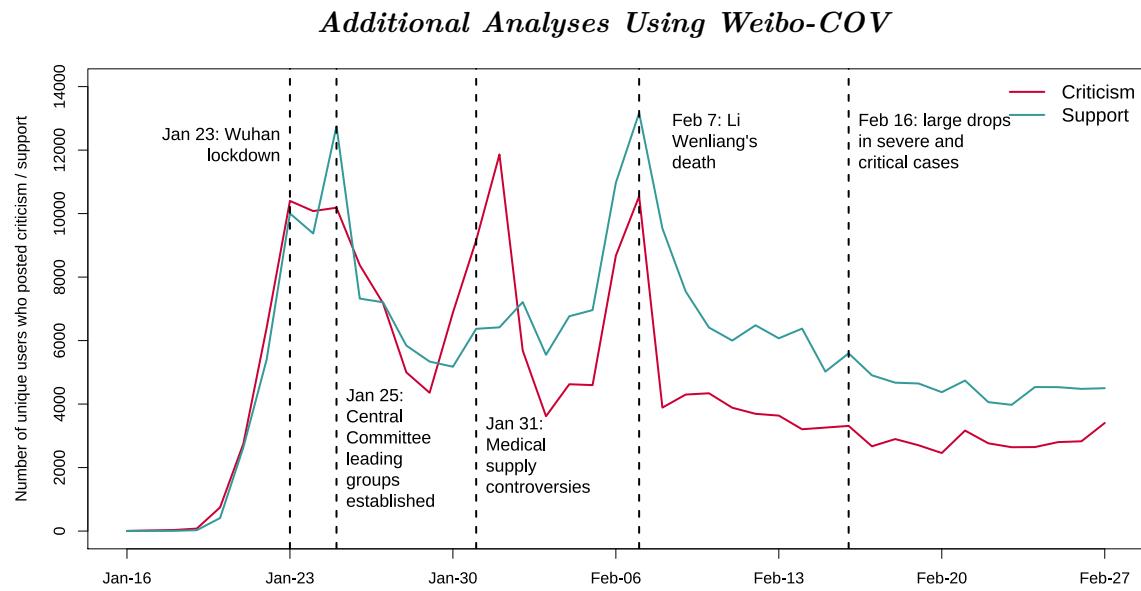


Figure S4. Unique users making critical and supportive posts (Weibo-COV).

Note. This figure shows the number of unique users who made at least one critical post (red line) and who made at least one supportive post (green line) by day from January 16, 2020 to February 27, 2020 based on the Weibo-COV dataset. Prior to February 7, the number of unique users engaging in criticism and support is similar. After February 7, the overall number of unique users declines and the number of unique users engaged in posting support outpaces the number of unique users engaged in posting criticism.

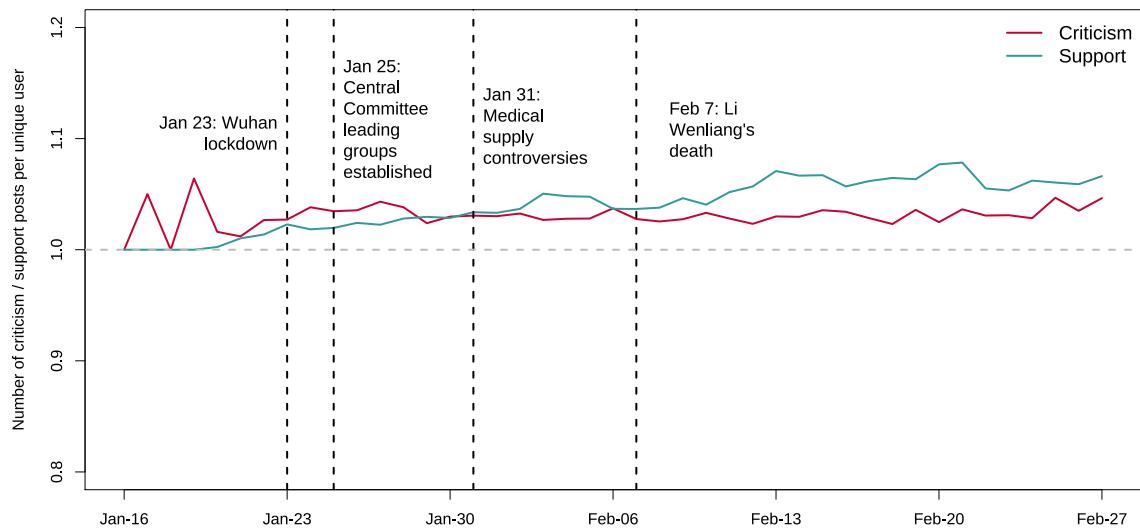


Figure S5. Critical and supportive posts per user per day (Weibo-COV)

Note. This figure shows the number of critical posts divided by the number of unique users who posted at least one critical post by day (red line) and the number of supportive posts divided by the number of unique users who posted at least one supportive post by day (green line), between January 16, 2020 to February 27, 2020 from the Weibo-COV dataset. Both red and green lines are slightly above 1, which means that support and criticism are similarly diffuse, with each post coming, for the most part, from unique users.

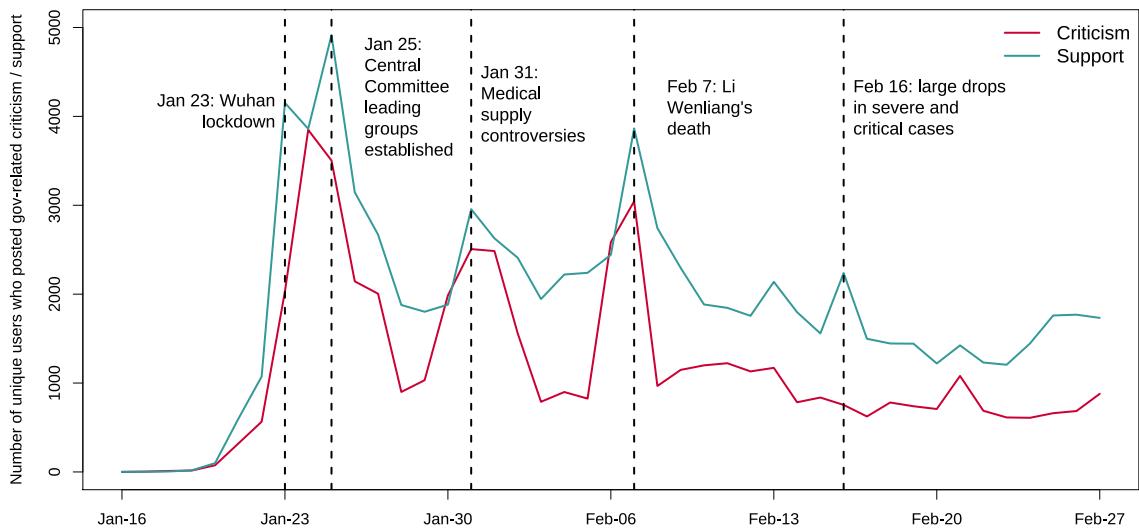


Figure S6. Unique users making critical and supportive posts targeting Chinese government (Weibo-COV)

Note. This figure shows the number of unique users who made at least one critical post targeting the government (red line) and who made at least one supportive post targeting the government (green line) every day from January 16, 2020 to February 27, 2020 in the Weibo-COV dataset. The number of unique users engaging in support of the government consistently exceeds or tracks the number of unique users engaged in criticism of the government over the time period.

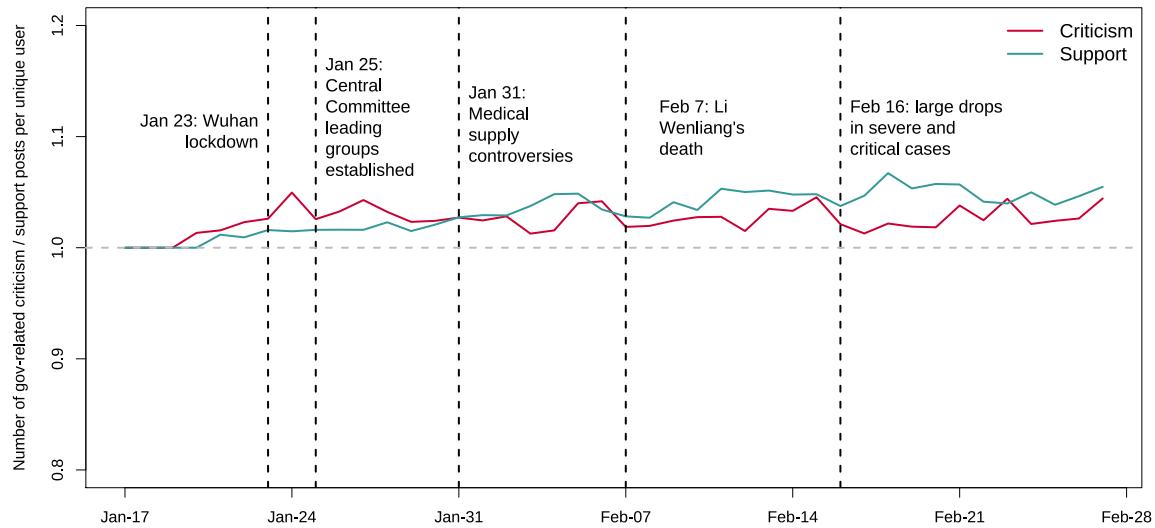


Figure S7. Critical and supportive posts targeting Chinese government per user per day (Weibo-COV)

Note. This figure shows the number of critical posts targeting the government per user per day (red line) and the number of supportive posts targeting the government per user per day (green) from January 17, 2020 to February 27, 2020 based on the Weibo-COV dataset. Critical posts targeting the government per user per day is calculated by dividing the total number of critical posts targeting the Chinese government per day by the number of unique users who posted at least one critical post targeting the Chinese government that day. Supportive posts targeting the government per user per day is calculated by dividing the total number of supportive posts targeting the Chinese government per day by the number of unique users who posted at least one supportive post targeting the Chinese government that day. Both red and green lines are slightly above 1, which means that support and criticism targeting the Chinese government are similarly diffuse, with each post coming, for the most part, from unique users.

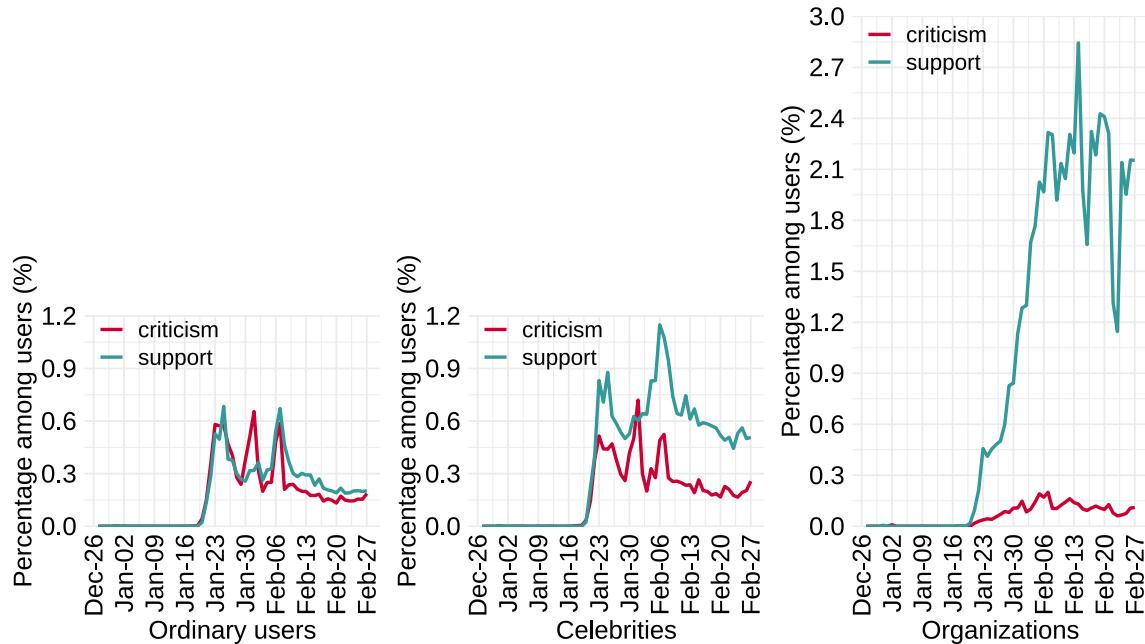


Figure S8. Critical and supportive posts as percentage of total posts by user type (Weibo-COV)

Note. The above figures show the percentages of critical and supportive posts among all posts every day from December 27, 2019 to February 27, 2020 in the Weibo-COV dataset for ordinary users (left), celebrities (center), and organizations (right). The figure shows that celebrity users are more likely to post supportive posts than ordinary users, and organizations are less likely to post critical comments than any other type of user and much more likely to post support than any other type of user.