




Authoritarian responsiveness and political attitudes during COVID-19: evidence from Weibo and a survey experiment

Lai Wei^a, Elaine Yao^a and Han Zhang^b 

^aPrinceton University, Princeton, New Jersey, USA; ^bHong Kong University of Science and Technology, Kowloon, Hong Kong

ABSTRACT

How do citizens react to authoritarian responsiveness? To investigate this question, we study how Chinese citizens reacted to a novel government initiative which enabled social media users to publicly post requests for COVID-related medical assistance. To understand the effect of this initiative on public perceptions of government effectiveness, we employ a two-part empirical strategy. First, we conduct a survey experiment in which we directly expose subjects to real help-seeking posts, in which we find that viewing posts did not improve subjects' ratings of government effectiveness, and in some cases worsened them. Second, we analyze over 10,000 real-world Weibo posts to understand the political orientation of the discourse around help-seekers. We find that negative and politically critical posts far outweighed positive and laudatory posts, complementing our survey experiment results. To contextualize our results, we develop a theoretic framework to understand the effects of different types of responsiveness on citizens' political attitudes. We suggest that citizens' negative reactions in this case were primarily influenced by public demands for help, which illuminated existing problems and failures of governance.

Introduction

Does authoritarian responsiveness always improve perceptions of government performance? While a considerable literature focuses on the

existence and sources of authoritarian responsiveness (Chen, Pan, and Xu 2016; Chen and Xu 2017; Su and Meng 2016; Heurlin 2016; Distelhorst and Hou 2017; Jiang, Meng, and Zhang 2019; Kornreich 2019), these studies tend to assume that responsiveness universally improves citizens' attitudes towards the government, or at least placates their dissatisfaction. In comparison, fewer papers spotlight how authoritarian responsiveness is actually viewed by citizens. Does responsiveness always meet with citizen approbation, and if not, why?

Our paper studies the political effects of one exceptionally direct instance of authoritarian responsiveness in China, which occurred during the COVID-19 outbreak. In January 2020, a “special hashtag” was introduced by Sina Weibo (“Weibo”), a highly popular micro-blogging site similar to Twitter, in cooperation with state-run media and local governments. The special hashtag was created in order to structure and facilitate posts seeking help for COVID-19 patients, which had been appearing in increasing numbers on the platform.¹ After the introduction of the special hashtag, the numbers of COVID patients and their families seeking help on Weibo rapidly increased, sparking intense reactions and heated discourse on Weibo. Whereas the state media had downplayed the scale of human suffering, help-seeking posts written directly by patients and their families illuminated the human toll of the outbreak. They described personal details and disease symptoms, often in graphic and heartbreaking detail, and also described unsuccessful attempts to obtain medical treatment which had been foiled by insufficient resources and unhelpful administrators.

We analyze the impact of help-seeking posts on public opinion using a two-part empirical analysis. First, we conduct a survey experiment to determine the causal impact of exposure to a small number of help-seeking posts on political attitudes. Our experiment, which involved about 1,500 respondents, provides robust evidence that viewing help-seeking posts *worsened* subjects' perceptions of the effectiveness of the Hubei government's response to COVID-19, and did not have a positive impact on perceptions of the central government's effectiveness. To explore whether information about the *outcomes* of online help-seekers affected perceptions, we further tested whether differential information about the eventual outcome of help-seekers changed perceptions of government performance. We found that our previous results held irrespective of whether subjects were given positive or negative information. Thus, experimental results imply that government responsiveness does not necessarily improve assessments of government effectiveness, and in some cases demonstrably lowers them.

To test whether our experimental findings were borne out in the real world, we conducted a large-scale social media analysis using a data set

containing over 10,000 real Weibo posts which discuss online COVID help-seekers, and 500,000 Weibo posts which discuss COVID more broadly. First, we estimate the proportions of positive and negative sentiments in the help-seeking discourse, and in particular, proportions of positive and negative *political* sentiments. We find that posts about help-seekers contained significantly more negative sentiments, as well as more *negative political* sentiments, than posts about COVID in general. However, posts expressing positive sentiments and positive attitudes toward the government increased in both popularity and amount over time.

This study is among several pioneering efforts to examine the impact of the COVID-19 crisis on political attitudes in China. It leverages a unique data set of Weibo posts which were scraped in real time to comprehensively describe Chinese netizens' response to this unique government initiative. In comparison to other recent articles, such as Fu and Zhu (2020); Hu et al. (2020); Lu, Pan, and Xu (2021), which primarily seek to document the public opinion on Weibo during COVID-19, we conduct both a descriptive analysis as well as a causal analysis in the form of a survey experiment. This allows us to directly study the causal impacts of reading Weibo posts about COVID-19 and about help-seeking in particular. Furthermore, we provide a theoretical decomposition of the *types* of authoritarian responsiveness. This allows us to relate our argument to the existing studies of responsiveness, which cover a wide range of political phenomena ranging from elite consultations to government interaction with Internet users on social media. Our typology of authoritarian responsiveness better clarifies the relationship between different types of responsiveness and places our case in a comparative perspective with other studies of responsiveness in China and other authoritarian regimes.

Theoretical framework

Typology of authoritarian responsiveness

In democracies, government responsiveness is built into electoral and policy-making institutions, which provide a number of channels for citizens to voice their opinions and influence policy outcomes. In authoritarian countries, institutional channels for public participation are comparatively limited. Insofar as consultative and participatory channels do exist to invite public input for authoritarian policy-making, the focus of existing scholarship has been to explain why they exist at all, and their consequences for authoritarian stability. However, this literature on “consultative authoritarianism” (Chen, Pan, and Xu 2016; Su and Meng

Table 1. Responsiveness schema.

	Citizen-initiated	State-initiated
Formal	Institutional Petition	Institutional Consultation
Informal	Responses to Popular Demands	Public Opinion Detection

2016; Heurlin 2016; Distelhorst and Hou 2017; Jiang, Meng, and Zhang 2019; Kornreich 2019) largely neglects the question of how citizens view these channels, as well as the less well-institutionalized ways in which citizens influence (or attempt to influence) authoritarian governance.

To better understand these differing forms of “authoritarian responsiveness”, as well as to relate our case to existing streams of literature on consultative authoritarianism and responsiveness, we develop a schema to organize the types of responsiveness. We differentiate responsiveness along two primary axes: first, whether responsiveness is *citizen-initiated* or *state-initiated*, and second, whether it is *formal* or *informal*. This schema is summarized in Table 1: Columns reflect whether responsiveness is initiated by demands by civil society actors, or by the state seeking citizen input to respond to. Rows reflect whether or not the state-society interaction is conducted through official channels established or explicitly intended for that purpose, or conducted outside of those channels.

Next, we provide some examples of how each type of responsiveness appears in Chinese politics. The petition (*xinfang* 信访) system epitomizes formal, citizen-initiated responsiveness (Minzner 2006). The petition system consists of an administrative bureaucracy responsible for receiving messages from petitioners, cataloguing and directing those petitions to an appropriate party, and communicating the outcome or response back to the petitioner. Existing literature on the petition system has shown how such institutional arrangements absorb social dissatisfaction and help maintain social stability in China (Lee and Zhang 2013; Minzner 2006). The Chinese government has also been enthusiastically engaged in establishing new formal channels for people to express public demands. The work of Chen, Pan, and Xu (2016) and Su and Meng (2016), who study Chinese provincial officials’ responses to citizen demands on local government website forums, can also be understood as citizen-initiated formal responsiveness.

Consultation exemplifies formal, state-initiated responsiveness, since it focuses on the state’s establishment of channels for popular political participation (Truex 2016; He and Warren 2011; He and Thøgersen 2010). Unlike the petition system, consultation is motivated by state’s efforts to involve civil society in policy discussions. Several authors find that consultation helps achieve higher levels of political stability and legitimacy and may improve citizen satisfaction with the regime, especially for less educated, politically excluded citizens (Truex 2017).

The Chinese state also collects public opinions outside of formal channels and tries to incorporate them in the decision-making process, a practice that we classify as *state-initiated, informal responsiveness*. For instance, Chen and Xu (2017) suggests that the Chinese government deliberately allows the public to voice their dissatisfaction in the purpose of detecting potential threats of collective action. Lorentzen (2014) also stresses the importance of collecting public opinions in authoritarian regimes, which not only helps governments understand popular demands, but also strengthens the central government's control of lower-level officials. Using the case of the 2008 health-care reform in China, Kornreich (2019) shows that public opinions collected by the state do indeed influence the policy content.

Informal citizen-initiated responsiveness, the focus of our paper, occurs when citizens go outside of officially sanctioned channels to voice complaints about and demands of the government. Due to its inherently extra-institutional nature, this is a highly heterogeneous category, including both contentious actions such as street demonstrations (Chen 2012; Zhang and Pan 2019), as well as non-contentious actions such as posting complaints on public platforms in an effort to attract government attention and response (Yang 2009; King, Pan, and Roberts 2013).

Different types of responsiveness imply different processes of state-society interactions, which then have different impacts on public evaluations of the government: Citizen-initiated responsiveness follows a “demand-response” logic, whereas state-initiated responsiveness follows a “consultation-action” logic. In our context, “demands” are expressed in help-seeking Weibo posts, and “responses” are the state's responses to help-seeking claims, or lack thereof. We emphasize the decomposition of citizen-initiated responsiveness into these two steps, since observing demand and responses may have distinctly different impacts on opinions. We discuss these possible different impacts and implications in the Hypotheses section.

Hypotheses

First, we focus on how *observing* help-seeking posts—the demand step of the demand-response logic highlighted above—shapes public opinion. One possibility is that Weibo users view help-seeking posts positively, because users interpret posts as evidence of the government's willingness to set up a channel to collect as many requests for help as possible. By this logic, citizens view the COVID outbreak as an unforeseen natural disaster, rather than the result of government mismanagement, and are optimistic that the government will successfully address the demands of afflicted citizens. This motivates our first hypothesis:

Hypothesis 1a. *Viewing help-seeking Weibo posts improves users' perceptions of government performance, because they are evidence of government responsiveness to citizens' needs.*

However, demands for help may have a negative effect on perceptions of government effectiveness, if users interpret them as evidence of prior government mismanagement or inaction. Indeed, several authors find that authoritarian governments' efforts at transparency, responsiveness, and reform may actually backfire. For instance, Wang and Dickson (2021) find that anti-corruption campaigns in China reduce government support by updating citizens' beliefs regarding the extent of corruption and lowering trust in institutions. Similarly, Malesky, Schuler, and Tran (2012) and Hollyer, Rosendorff, and Vreeland (2015) find that greater transparency may be destabilizing to authoritarian regimes.

In the case we study, the expression of demand may also have revealed a great amount of information to previously uninformed citizens. Help-seekers' requests for help on Weibo often contained detailed information about their severe health conditions and personal suffering, which stood in stark contrast to official media accounts at the time. In the early stages of the pandemic, the official media disclosed extremely limited information about underserved patients in COVID-afflicted areas and portrayed the government response as largely competent and sufficient. By this logic, citizens view help-seeking posts as strong evidence pointing to the government's crisis mismanagement and failure to proactively address citizens' needs. Based on this mechanism, our competing hypothesis is that:

Hypothesis 1b. *Viewing help-seeking Weibo posts does not improve users' perceptions of government performance by revealing the large number of people in need of help from the government, and demonstrating lack of state capacity.*

Our next two hypotheses relate to the response step of the "demand-response" logic. We explore whether differential information about the government's response to help-seeking demands changes citizens' perceptions of state performance. In particular, we test whether Weibo users perceive the government to be less effective if they are informed that help-seekers were never reached or given treatment, compared to if they are informed that all help-seekers received help.

There are strong reasons to believe that responsive measures taken by the government positively influence perceptions of government effectiveness. Simple retrospective logic implies that if the government fulfills citizens' needs, then citizens view the government more favorably. Furthermore, in the crisis management literature, governments' responsive efforts to rescue and rebuild after natural disasters tend to earn

public approbation. For instance, You, Huang, and Zhuang (2020) found that political trust toward public officials of all levels increased significantly immediately following the 2008 Sichuan Earthquake in China, and largely attributed this increase to mass media coverage. We hypothesize, therefore, that if the subjects are notified that the government effectively responds to popular demand for medical services, they would have improved perceptions of government effectiveness in the management of the crisis.

Hypothesis 2a. Receiving information about the eventual outcome of help-seekers has significant effects on perceptions of government effectiveness.

However, there may be reasons why this does not hold. If citizen demands provide evidence of state mismanagement that is so strong that even subsequent positive signals of state-capacity are not able to counteract the initial information, viewers may still believe that the government is ineffective. This mechanism may be further exacerbated in our case because the government's response (for example, hospitalizing specific individual patients) would not be easily visible to non-help seekers, particularly on social media, and may have had an even lesser emotional and informational impact than the help-seeking posts. By this logic, our second response-side hypothesis is that:

Hypothesis 2b. Receiving information about the eventual outcome of help-seekers does not have significant effects on perceptions of government effectiveness.

Our final hypothesis describes the empirical implications of the previous hypotheses (particularly 1a and 1b) outside of the survey experiment environment. For example, if our survey experiment provides support for Hypothesis 1b, thus implying that viewing help-seeking posts negatively impacts perceptions of government effectiveness, we should expect that negative sentiments should be more common among Weibo users who viewed help-seeking posts than those who did not. Since we are unable to directly observe whether or not users viewed help-seeking posts, we can approximate this effect by comparing the sentiment composition of posts mentioning help-seeking to posts which discuss COVID, but not help-seeking in particular. Using this approximation, our following hypotheses establish that the causal direction that is found to dominate in our experimental setting should concur with the attitudes we observe in our social media data. Therefore:

Hypothesis 3a. If Hypothesis 1a is true, the proportion of posts expressing government praise to posts expressing government criticism should be higher among posts discussing help-seeking, compared to posts discussing COVID in general.

Hypothesis 3b. If Hypothesis 1b is true, the proportion of posts expressing government praise to posts expressing government criticism should be lower among posts discussing help-seeking, compared to posts discussing COVID in general.

One possible objection is that these final hypotheses neglect to address the empirical implications of the “response” effect, i.e., whether the government successfully or unsuccessfully responded to demands for help. One limitation of our social media data is that citizens’ demands are much more visible on Weibo than the government’s responses. Whereas users could easily read numerous help-seeking posts, there are fewer posts which document the outcome of help-seeking demands. The government’s particular plans and actions were not publicized on social media, and help-seekers did not regularly post “followups” revealing if their demands were met. Even if a help-seeker managed to obtain aid, it would not always be clear to them, or to others, if that was directly caused by their Weibo post. Therefore, while we directly address the response effect in our survey experiment, our social media data is not well equipped to identify this effect.

Background: help-seeking posts on Weibo

Weibo played an important role in distributing information about COVID-19 in China in the early stages of the pandemic. Weibo is one of the largest social media sites in China and globally, with 241 million daily active users and 550 million monthly active users as of March 2020.² It is one of the most diverse and representative platforms of Chinese Internet users (Zhang, Liu, and Wen 2018).

When COVID-19 appeared in Wuhan in December 2019 and quickly worsened throughout early 2020, the Chinese government’s initial response was to censor COVID-related news and downplay its severity. However, as COVID spread first in Wuhan and then across the rest of the country, the government gradually lifted restrictions on traditional and social media, releasing a flood of coverage and information about the outbreak and its human impact.

Many people suspected to have contracted COVID were living in Wuhan, the capital of Hubei province, which faced a shortage of medical resources. This shortage was exacerbated by a highly restrictive lockdown from January 23 to April 8, resulting in many patients being unable to receive tests or treatment in hospitals. They and their families began to voice their struggles and request aid on Weibo, often describing experiences of being denied treatment, and expressing frustration about the chaotic situation and poor management by local authorities.

In response to the wave of COVID help-seekers, Weibo created a “special hashtag” in collaboration with local governments on January 29, providing a sanctioned method to request COVID-related help on the platform. People who tagged their posts with the special hashtag and followed the specified format for information sharing (including the patient’s name, age, city, city district, date of illness, contact information, and description of symptoms) were promised that their information would be shared with local authorities and hospitals who would offer them assistance.

Authorities were enthusiastic about Weibo’s initiative. According to an announcement by Weibo, the Wuhan City Government would be the recipient of the information provided by Weibo and would be responsible for providing medical assistance to patients in difficulty. In the meantime, the most important state-run media in China, the People’s Daily newspaper and China Central Television (CCTV), also opened online platforms in cooperation with Weibo to collect information of patients and promised to deliver them to relevant authorities.³

Because of the deep engagement of the government in this initiative, and the role of various levels of government in publicizing and expanding this initiative, we consider this phenomenon to be an example of informal citizen-initiated government responsiveness, rather than a charity campaign by a private company. By February 26, Weibo collected information from over 10,000 patients and delivered the information of at least 3,000 people to authorities. It remains unclear how many patients eventually received treatment because of this cooperation between Weibo and the government.⁴

Apart from connecting patients with the government and healthcare providers, these Weibo posts also generated enormous publicity around these help-seekers. As Chinese citizens were generally dissatisfied with the government’s slow response to the crisis and information obfuscation in December and early January, their feelings were quickly aroused by the vivid descriptions of suffering and neglect contained in Weibo help-seeking posts. As of late July, the posts under the special hashtag have been viewed a total of 5.1 billion times.⁵

Despite the fact that the state played a positive role in collecting and responding to help-seeking information, Weibo users did not seem to be impressed by the government’s good intentions. Because of the restricted information environment that has persisted thus far, this represented the first time that most Chinese netizens learned of the difficulty and suffering of Wuhan residents due to COVID. As a result, Weibo users seemed to respond with feelings of anger and frustration at the grim situation, rather than feelings of consolation and confidence in the ability of the state to ameliorate the situation. In the next section, we provide

Table 2. Sample descriptive statistics.

Statistic	Mean	St. Dev.	N
Under 40 years old	0.522	0.500	1,452
College educated	0.498	0.500	1,452
Urban resident	0.731	0.443	1,452
CCP member	0.207	0.405	1,451
Income over 200k yuan	0.205	0.404	1,442
Female	0.549	0.498	1,453

experimental evidence supporting the claim that viewing help-seeking posts lowers perceived state effectiveness. In the following section, we investigate the actual sentiments expressed by Weibo users during the time that COVID help-seekers were a national focus.

Research design

Survey experiment

To test Hypotheses 1 and 2, we conduct a survey experiment designed to obtain a causal estimate of the impact of a short exposure to help-seeking posts on subjects’ perceptions of government performance. Survey participants were drawn from a national panel of an international market research company. To strengthen confidentiality, the survey was hosted on an encrypted site, with data stored in the United States. Because this represents a convenience sample where respondents were recruited online, the sample is younger, better-educated, and more urban than the national average, although respondents were recruited to be more demographically representative of the country than a pure convenience sample. [Table 2](#) displays the descriptive statistics for our sample. Throughout our experimental study, we ensured data quality through a combination of attention checks, speed checks, and other data quality verification measures.

Treatment conditions

We randomly divided our sample into two treatment groups and one control group, all of equal size. Visualizations of covariate balance, which prove that our randomization is valid, are displayed in [Appendix A](#). They show that our three groups are balanced on the observed characteristics such as age, education, party membership, income, gender, and Weibo usage.

Both treatment groups viewed the same screen captures of real help-seeking posts on Weibo (details in [Appendix F.2](#)), while the control group users did not see these excerpts. We further divided the treatment groups into two subgroups by providing different information about the success

of the initiative. The “effective state” subgroup was informed that “The help-seekers were able to obtain timely and proper aid and arrangements through the government”, while the “ineffective state” subgroup was informed that “Some portions of COVID-19 help-seekers who requested help through Weibo were not able to obtain timely aid.” Both groups received the same general overview information prior to seeing the posts, which was: “Due to the COVID-19 outbreak, the Wuhan municipal government as well as the central government opened a new channel for COVID sufferers in the Wuhan area and their families to seek help for their illness by posting on Weibo”.

Outcomes

After the treatment, subjects were asked a battery of questions regarding their perceptions of the COVID-19 crisis, governance, emotions, and nationalism. Detailed information about how each outcome reported in our analysis was measured in the experimental instrument is displayed in [Table 3](#). To measure the primary outcomes of interest, namely, subjects’ perceptions of the effectiveness with which different levels of the Chinese government managed COVID, we asked subjects directly to rate effectiveness on a 1-4 scale. Questions asking respondents to evaluate government performance are common in public opinion surveys like the Asian Barometer Survey. In addition to pure perceptions of government effectiveness, we also asked about emotions because we were curious whether the potential treatment effect results purely from rational evaluations of government performance, or also from negative emotions aroused. Our questions regarding emotion were worded in accordance with the Positive and Negative Affect Schedule (PANAS) questions widely used by psychologists to measure emotional states. We asked about nationalism in order to ascertain whether subjects’ perceptions of government effectiveness also influenced nationalist sentiment. Our questions regarding nationalism were adopted from questions in the International Social Survey Program (ISSP) regarding national identity. These questions have been previously used to study nationalist sentiment in the China context (see Huang and Liu (2018), Mattingly and Yao (2020) ([Table 4](#))).

First, we compare subjects who viewed help-seeking posts (pooling the two treatment groups) and subjects who did not. We test whether viewing help-seeking posts led to a higher or lower rating of government performance (Hypotheses 1a and 1b). Then, we compare the two treatment groups to test Hypotheses 2b and 2a. If there is no distinguishable difference between the ratings of government performance in the two treatment groups, we find support for Hypothesis 2b.

Table 3. Measurement of key outcomes in survey experiment. Note that outcomes regarding the “Provincial” government were specifically about subjects’ province of residence.

Variable Type	Variable	Text of question and response options
COVID management	Central government effectiveness	“Do you believe that the (National/ Provincial/Hubei) government’s management of the coronavirus in February was effective? (Very ineffective, Not very effective, Somewhat effective, Very effective)”
	Provincial government effectiveness	
	Hubei government effectiveness	
Economic management	Central government economic effectiveness	“Do you believe that the (National/ Provincial) government’s economic management this year has been effective? (Very ineffective, Not very effective, Somewhat effective, Very effective)”
	Provincial government effectiveness	
Emotions	Pride	“Please indicate how strongly you are feeling the following emotions: (Pride, Anger, Sadness). (Not at all, a little, somewhat, rather strong, extremely strong)”
	Anger	
	Sadness	
Nationalism	Choosing Chinese citizenship	“Do you agree with the following statement: Even if I could choose any other country in the world, I would rather be a Chinese citizen. (Strongly agree, Agree, Neither agree nor disagree, Disagree, Strongly disagree)”
	Supporting country unconditionally	“Do you agree with the following statement: People should support their country even if the country is wrong. (Strongly agree, Agree, Neither agree nor disagree, Disagree, Strongly disagree)”

The wording provided is in English translation from and may have minor grammatical differences from the Chinese wording. Please contact authors if the Chinese-language questions are of interest.

Table 4 Schema of treatment conditions.

	Priming video?	Receive help-seeking overview?	View Weibo help-seeking posts?	Help-seeking outcome info
Control group	Yes	No	No	None
“Effective state” treatment group	Yes	Yes	Yes	All help-seekers received help
“Ineffective state” treatment group	Yes	Yes	Yes	A portion of help-seekers did not receive help

Table 5. Correlation matrix of demographic information and perceived government effectiveness.

	Central gov effectiveness	Provincial gov effectiveness	Hubei gov effectiveness	National gov economic effectiveness	Provincial gov economic effectiveness
Female	−0.02	−0.02	0.02	0.01	−0.01
Non-rural <i>hukou</i>	0.07**	0.08***	0.11***	0.11***	0.05
Years of education	−0.07**	−0.01	0.05	0.05	0.06*
Income	0.14***	0.10***	−0.00	0.05	0.00
Age	−0.08***	−0.07**	0.01	0.05*	0.04
Party member	−0.05	−0.03	−0.02	−0.05*	−0.02

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

To relate our experimental design to the “demand-response” logic of informal citizen-initiated responsiveness, we observe that the total average treatment effect of the demand-response cycle for a given outcome is obtained by the difference between the “Effective state” treatment group and Control group means. To see why, recall that the Control group is given information about neither “demand” by citizens nor “response” from the government, whereas the “Effective state” treatment group receives both. The average treatment effect of information about government “response” is given by the difference between the “Effective state” treatment group and the “Ineffective state” treatment group means. Finally, the average treatment effect of information about “demands” by the citizens is given by the difference between the “Ineffective state” treatment group and Control group means, as the “Ineffective state” treatment group is given information about “demand” by citizens, but ambiguous information about whether the government successfully responded to those demands. Because we use a randomized experiment, we can simply use difference-in-means estimator by comparing the means of treatment and control groups for corresponding outcomes. This strategy is typical for analyzing experimental data (Freedman 2008).

Here we report how perceptions of state effectiveness correlate with important demographic variables. Table 5 shows the Pearson correlation matrix of the main outcomes and demographic information with significance tests. People with non-rural *hukou* (residence registration) and with higher incomes tend to be more satisfied with the government’s management of the COVID crisis, although higher levels of education and older age are correlated with lower satisfaction with state response. Surprisingly, CCP members are less satisfied with the state’s crisis management efforts, although the correlations are generally not significant. To be clear, these correlations do not have causal interpretations and do not necessarily apply to the general Chinese population. Nonetheless, they shed some light on variations in Chinese peoples’ assessment of their government’s COVID response, which has been rarely studied so far.

Design limitations

One challenge of our survey experiment design is related to timing: while help-seeking posts were primarily posted in January–March 2020, our survey experiment was implemented several months later, in July 2020. By July, the COVID situation in China had changed substantially from in early 2020, which have may cause changes in public opinion as well.⁶ This implies that experimentally measured *negative* treatment effects may be significantly muted, since any experimentally-induced negative perceptions of the government may be overwhelmed by positive information about the Chinese government's performance in the months between. On the other hand, experimentally measured *positive* treatment effects may be exaggerated, since experimentally provided information about positive government performance may be buoyed by months of information about positive government performance.

In order to address this concern, we take a number of steps to replicate the ideal setting of February 2020 as closely as possible. All respondents, including those in the control group, were exposed to a short priming video which portrayed the urgent medical situation in Wuhan area hospitals in February. A transcript of the video is contained in [Appendix F.1](#).⁷ Furthermore, questions which pertained to COVID-19 and governance repeatedly specified that the time period of interest was February 2020. The priming video also serves as politically neutral information which was administered to all treatment conditions, including the control group.

An inevitable effect of administering our survey experiment online is that our sample population is not perfectly representative of the general population, and the effects we discover may therefore not apply directly to the general population. Nonetheless, this does not yield the experimentally estimated causal effect useless. On the contrary, prior studies have found that non-representativeness is less harmful in experimental research that aims at estimating a treatment effect, compared to descriptive research that aims at describing the distribution of a variable (Mullinix et al. 2015; Coppock, Leeper, and Mullinix 2018; Lupton 2019; Coppock 2019). These studies compared survey experiments conducted with non-representative samples and a nationally representative sample, and generally do not find noticeable differences in treatment effect estimates. As Coppock and McClellan (2019, p.12) nicely summarizes, these studies do not find significant differences between a convenient sample and a nationally representative sample because these studies found low treatment heterogeneity by covariates, such that even though the representative and unrepresentative sample differ on covariates, the treatment effect does not. In our own experiments, we also noticed that there is little to no treatment effect heterogeneity across different demographic groups

(Appendix B), meaning that the results would not significantly change if we re-weighted the sample based on demographic information. Still, we acknowledge that the validity of our findings beyond the experimental setting is a genuine concern. Therefore, we next seek to validate the findings of our survey experiment using a media study that draws upon non-experimental “real world” data.

Social media study

In order to test whether the experimentally identified causal effect of viewing help-seeking posts is supported by real-world data, we study real Weibo posts about help-seekers and COVID. The advantage of our data is that we are able to analyze real-time behavioral outcomes, namely, users’ decisions to post responses after viewing help-seeking posts. However, we caution readers that our results in this section are not causal evidence, due to the fact that the counterfactual—what users *would have* posted had they not seen help seeking posts—is not observable from our data.

To collect public responses, we scraped Weibo minutely from February 10 to March 17 with the keyword “seeking help” (求助). To establish a quasi-control group to compare these against, we also scrape posts that mention “COVID” (新冠), but do not contain words related to help-seeking. One concern about social media data in China is platform censorship, which could bias our conclusions by reducing the numbers of anti-government posts. Real-time scraping reduces concerns about ex-post censorship, and has been used in other studies to investigate what contents are more likely to be deleted (Fu, Chan, and Chau 2013; King, Pan, and Roberts 2013).⁸

After excluding duplicate posts, posts unrelated to COVID, and posts that were originated from official institutional accounts (e.g. People’s Daily), we collected 10,495 total posts which mention “help-seeking.” After performing further pre-processing steps and filtering out posts which did not display clear sentiments or opinions (using human coders), we were left with 5,234 posts. For our quasi-control group, after performing pre-processing and filtering steps, we were left with 359,223 posts. Our preprocessing steps are described in greater detail in Appendix C.

Our targets of estimation are the proportions of opinions (e.g., positive vs. negative) among the help-seeking posts and the COVID non-help-seeking posts. Each of the 5,234 posts about help-seeking that expressed clear sentiments or opinions were assigned a label by human coders, following a categorization schema which is shown in detail in Table 6. Simply put, posts were categorized broadly into “positive” and “negative”

Table 6. Categorization regime of the discussions about Weibo help-seekers.

Name	Description	Ideal Example
Explicit anti-government	Posts that clearly criticize the way the situation has been handled	Those help-seekers are not getting treated! The government is doing a poor job!
Implicit anti-government	Posts that express dissatisfaction but without targeting the government directly	People cannot get help. The society is totally corrupt.
Negative emotion	Expression of sadness or frustration without criticizing anyone	It is so devastating to see so many people in desperation.
Pro-government	Posts that clearly praise the way the situation has been handled	The situation in Wuhan is improving and the government has done a good job!
Positive emotion	Posts that express encouragement, hope or optimism without making political references	Cheers for the medical staff in Wuhan! They are the true heroes!

sentiments categories. Within each category, posts were further classified based on whether the sentiment they expressed was primarily *political* or *emotional*. The negative political posts are then classified as either “explicit” or “implicit” based on how explicitly their criticism was expressed.

Due to the extremely large quantity of posts which mention COVID, it was not feasible to categorize them manually. Therefore, for these general COVID-related posts, we assigned sentiment labels using dictionary-based automated text analysis (Grimmer and Stewart 2013).⁹ The labels produced by dictionary words could contain errors, which could lead to biased estimates of the proportions of positive or negative sentiment posts. We used the bias-correction procedure proposed by (Hopkins and King 2010) to adjust the estimates of the proportions of positive and negative posts. The bias-correction procedure relies on a random sample of about 1,000 posts that were labelled by human coders¹⁰ which provide information on the bias of machine methods, and then uses this information for correction. Details of preprocessing, dictionary methods, and bias-correction procedures are discussed in Appendix D.

Results

Survey experiment

Figure 1 shows the effects of help-seeking treatment on citizens’ assessments of the government’s effectiveness at managing the COVID-19 crisis

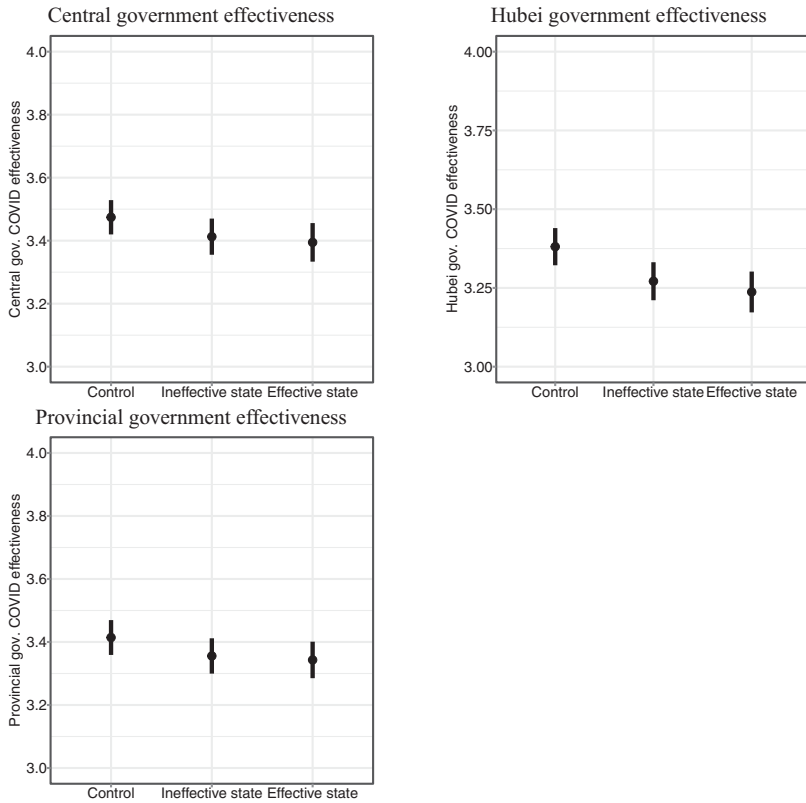


Figure 1. Treatment effects of viewing help-seeking Weibo posts on the effectiveness of governments' COVID-19 management in February 2020. The outcomes were measured on a 1–4 scale: very ineffective (1); somewhat ineffective (2); somewhat effective (3); very effective (4). Bars show 95% confidence intervals for the mean for each condition. (a) Central government effectiveness. (b) Hubei government effectiveness. (c) Provincial government effectiveness.

in February 2020 at three levels: the national government, the provincial government of the respondent's residence, and the Hubei provincial government.

We observe a strong negative effect on assessment of the Hubei government's effectiveness, statistically significant at the $p < 0.05$ level, as well as weaker negative effects on assessments of the national government and subjects' provincial government effectiveness, which are not statistically significant.

Overall, we find support for Hypothesis 1b, since our results show that viewing help-seeking posts does not improve the assessment of government performance at *any* level, and indeed worsens assessments of the Hubei provincial government (in whose jurisdiction the crisis originated

and was most severe). We also explore heterogeneous treatment effects on all outcomes to see if treatment effects varied by various demographic variables, and find no discernible patterns (see [Appendix B](#) for details).

We observe no significant differences between the perceived effectiveness of different levels of government by the two treatment groups. As long as respondents viewed help-seeking initiatives, their evaluations of government effectiveness were lowered, regardless of the outcomes of help-seeking initiatives. This provides evidence in favor of Hypothesis 2b, i.e. that information about the results of the help-seeking initiative does not affect assessments of government performance. This suggests that the primary mechanism driving perceptions of responsiveness in this instance is observation of the “demand” step, and information regarding the “response” step does not contribute observably to this effect (to either moderate or exacerbate it). To put this in the context of our identification strategy, the total effect of the demand-response cycle (comparison between “Efficient state” treatment and Control groups) is driven primarily by the “demand” effect (comparison between “Inefficient state” treatment and Control groups), rather than the “response” effect (comparison between the treatment groups).

We also investigated whether treatment conditions impacted subjects’ perceptions of government competence in policy areas presumably unrelated to the COVID-19 crisis. We asked subjects about their perceptions of the effectiveness of the central government’s and their provincial government’s *economic governance* in February 2020. The results we find are similar to the results for COVID-19 crisis management in direction: viewing help-seeking posts, regardless of the outcomes of help-seeking behaviors, reduces the confidence of government competence in economic management. We note, however, that the size of these effects is small, and they are not statistically significant at $p < 0.05$ level ([Figure 2](#)).

We are also interested in experimentally verifying the observation that Weibo users’ reactions to help-seekers’ posts have a strong emotional component. The questions in our survey that pertained to subjects’ emotions yielded strikingly strong results, as displayed in [Figure 3](#). Both treatment groups experienced a significant decrease in pride and a significant increase in anger and sadness. To be more specific, on a 1–5 scale, pride was on average a full point lower for each treatment group compared to the control group.

Based on these strong results about emotions, it is reasonable to ask if there were accompanying changes in nationalist sentiment—for example, if a decrease in pride was accompanied by a decrease in national pride. However, we found almost no results for questions that pertained to nationalism. [Figure 4](#) displays the treatment effects on two measures of nationalism: whether subjects would choose to be a Chinese citizen given

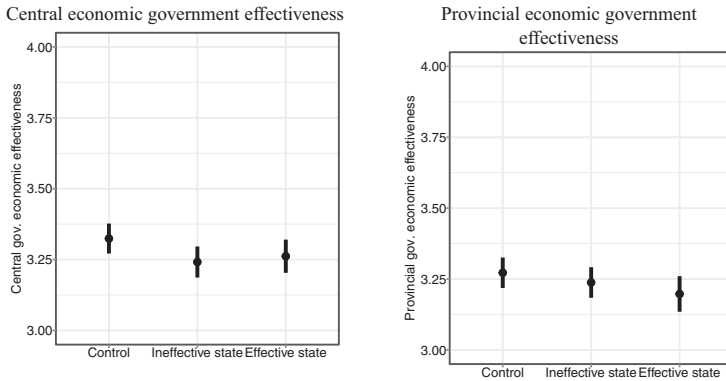


Figure 2. Treatment effects of help-seeking posts on the effectiveness of governments' economic management. The outcomes were measured on an 1–4 scale: very ineffective (1); somewhat ineffective (2); somewhat effective (3); very effective (4). Bars show 95% confidence intervals for the mean for each condition. (a) Central economic government effectiveness. (b) Provincial economic government effectiveness.

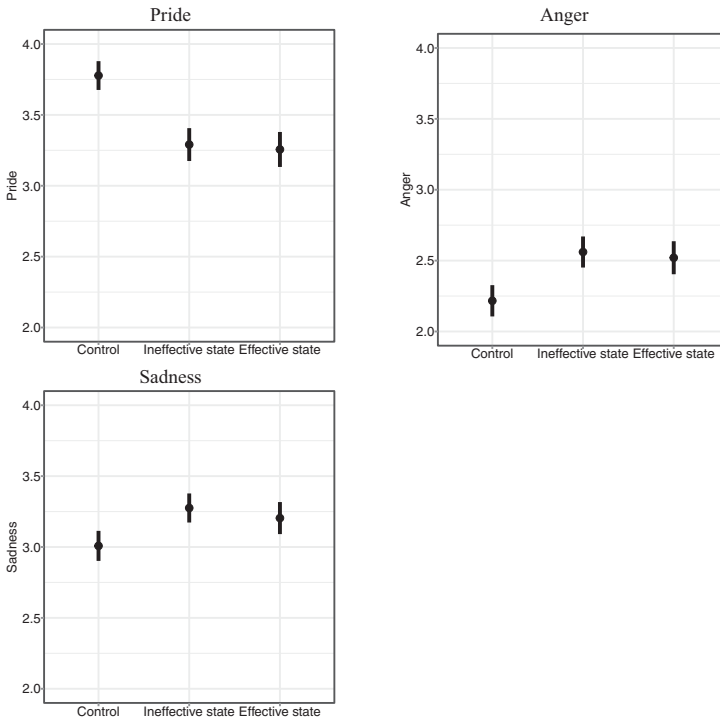


Figure 3. Treatment effects of help-seeking posts on self-reported emotions. The outcomes were measured on an 1–5 scale: almost none (1); a little (2); somewhat (3); somewhat strong (4); extremely strong (5). Bars show 95% confidence intervals for the mean for each emotion. (a) Pride. (b) Anger. (c) Sadness.

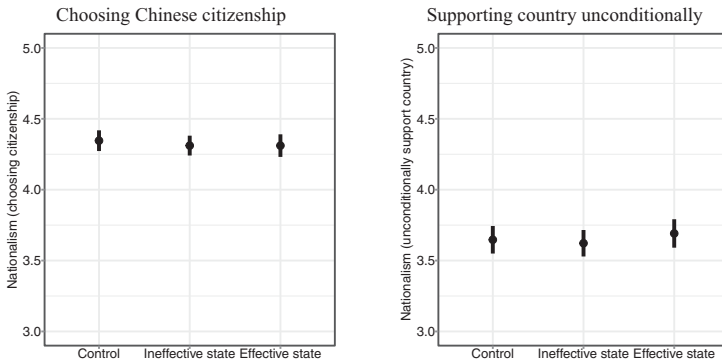


Figure 4. Treatment effects of help-seeking posts on nationalist attitudes, on a 1–5 scale. Bars show 95% confidence intervals for the mean for each condition. The outcomes were measured on a 1–5 scale, where (1) represents strong disagreement and (5) represents strong agreement with the statement. (a) Choosing Chinese citizenship. (b) Supporting country unconditionally.

the choice, and whether they would still support their country if their country was wrong. There is no discernible effect on either measure of nationalism, suggesting that while our treatments had significant emotional effects, they did not immediately spill over into nationalism and remained restricted to the domain of government competence and effectiveness. We do not claim that this is evidence that the COVID-19 epidemic had no effect on nationalism. Rather, we simply find no evidence that viewing help-seeking posts has a direct impact on nationalist sentiment.

These experimentally identified effects (and null effects) are among the first to shed light on the causal linkages between public calls for aid during COVID-19 crisis and perceptions of government performance. It is worth reiterating that the magnitude of the results may be affected because we could not go back to January 2020 and implement our experiment when COVID-19 first hit China. A second limitation is external validation. Although we did our best to mimic what people would see on social media by showing them priming videos and real help-seeking posts, these may not translate naturally into people's experience of browsing content naturally on Weibo. Precisely to address this external validity concern, we next presented our results using descriptive social media data, which lacks the strong internal validity of the survey experiment, but has broader implications beyond the experimental subjects.

Social media results

Having established that seeing help-seeking posts does not enhance citizens' perceptions of government performance in managing COVID, we

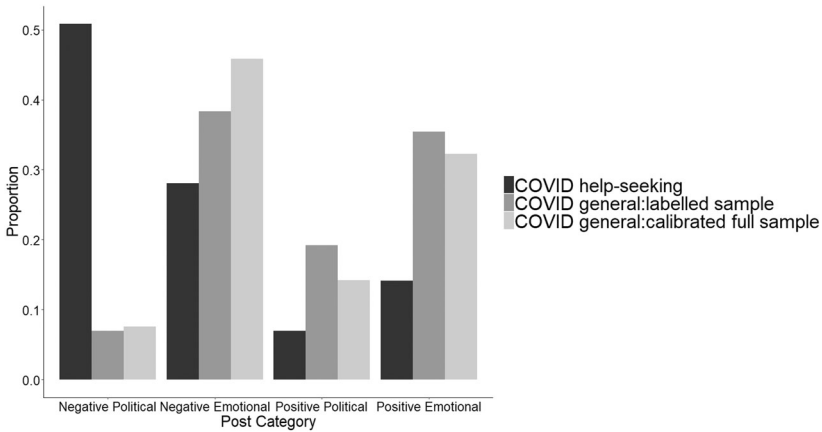


Figure 5. Distribution of opinions among Weibo posts mentioning help-seeking vs. mentioning COVID but not seeking for help.

now examine if this causal finding is consistent with actual social media behavior. The findings below show whether the causal effect of seeing help-seeking posts translates into observable behavioral outcomes on Weibo.

Figure 5 compares the distribution of sentiments for help-seeking posts (the “treated” group) and COVID non-help-seeking posts, which we use as a quasi-control group for comparison purposes.¹¹ Sentiment distributions are displayed for three categories: the “COVID help-seeking posts” (posts that mentioned help-seeking), the “Labelled COVID general sample” (the 1,000 COVID-relevant posts which were manually coded by human labellers), and the “Calibrated COVID general full sample” (which shows the machine-predicted distribution of sentiments calibrated by human labels among the 359,233 COVID-relevant posts). A more detailed explanation about the collection and categorization of posts in the research design section is available in [Appendices D and E](#).

We find that far more help-seeking posts—roughly four times as many—express *overall* negative sentiments than positive sentiments toward the government. Among posts with *political* sentiments, negative political opinion posts were nearly eight times as frequent as positive opinion posts. In comparison, negative political sentiments were by far in the minority among general discussions about COVID. Among these general discussions of COVID (which excluded mentions of help-seeking), any posts that *did* express political sentiments were largely positive. Furthermore, the proportion of non-political to political posts was much higher in general COVID discussions compared to help-seeking related discussions. In summary, COVID posts related to help-seeking tend to express more negative sentiments and, in particular, more negative

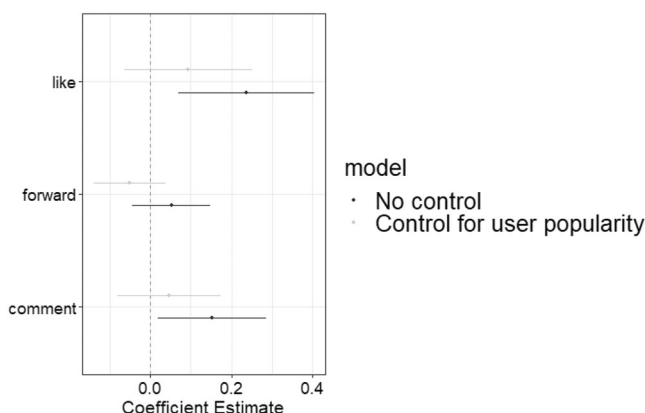


Figure 6. OLS estimates and 95% confidence intervals of post sentiment on popularity metrics. Negative sentiment posts are used as the reference group.

political sentiments, compared with posts discussing COVID but not seeking for help.

Given that responses to help-seeking posts on Weibo predominantly expressed anti-government sentiments, did anti-government posts also garner more attention? To address this question, we regress the number of likes, “forwards” (analogous to retweets on Twitter), and comments on a sentiment dummy. Our results, displayed in Figure 6, demonstrate that *positive sentiment* posts received more likes and comments compared with negative posts, and this result is statistically significant. Positive posts were also forwarded more often than negative posts, but the difference is not statistically significant at $p < 0.05$ level. However, once we control for user popularity, the coefficient estimate is smaller and is no longer significant. In summary, Figure 6 suggests that positive posts attract more “likes” and “comments” on Weibo because they are posted by users who have more followers and post more frequently. Although there were eight times more anti-government posts than pro-government posts in help-seeking discussions, the pro-government posts attracted more attention because they were posted and endorsed by more influential users.

Figure 7 displays the temporal trend of the opinion composition of the Weibo discussion. From the figure, we can see that in early February, anti-government sentiments and negative emotions dominated online discussions. Posts expressing pro-government sentiment and positive emotions constituted less than 20% of all posts.

However, there was a clear decline in the proportion of negative posts over time. The proportion of pro-government posts gradually increased in February and surged in early March. However, since the number of posts about help-seekers decreased significantly over time (see Figure A2 in the Appendix for the temporal trend of the number of posts), the later

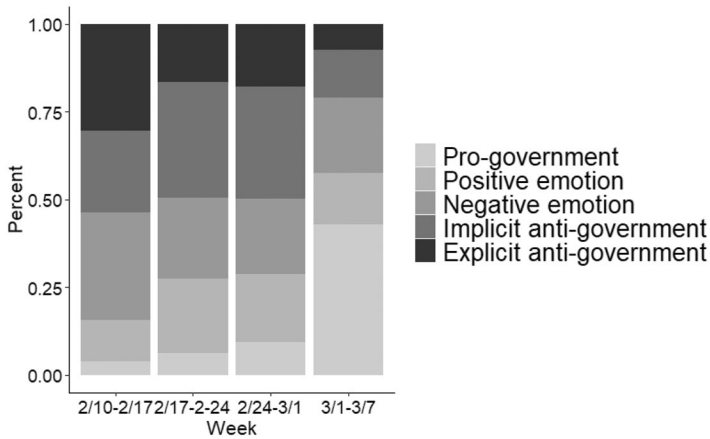


Figure 7. Temporal trend of opinion composition of Weibo help-seeking discussion.

increase in the proportion of positive posts does not significantly affect the overall distribution of opinions for all posts in our sample. Furthermore, [Figure A4](#) in the [Appendix](#) shows that pro-government posts were the least popular among all categories at the beginning of our data collection but became the most popular at last.

In summary, anti-government posts *far outnumbered* pro-government posts, but pro-government posts were originated by more influential users and hence attracted more attention. Over time, pro-government posts, which were initially in the minority and less popular, increased in both number and popularity.¹²

Discussion

In the early period of the COVID-19 outbreak in China, Weibo quickly became a hotbed of information sharing, discussion, and personal opinion. Furthermore, as we demonstrate, it became a significant interface for citizens to interface directly with the state. Our paper represents a deep dive into one aspect of this interaction; namely, the state's effort to be responsive to the help-seeking demands of residents in the hardest-hit areas. By leveraging the causal design of the survey experiment as well as observational data scraped from thousands of Weibo posts, we demonstrate that the government's effort at responsiveness did not immediately improve citizens' perceptions of government performance, and indeed may have worsened them. However, our findings demonstrate that the dynamics of social media are somewhat more nuanced, since Weibo posts that praised the government also more popular and widely circulated. Furthermore, the negative trend in the public opinion of the government

that is evident in the discourse around help-seekers is not persistent in the long term.

These findings are complementary to those of Lu, Pan, and Xu (2021), who investigated the Weibo discourse around COVID-19 prior to the time period of this study. They find that “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” (Lu, Pan, and Xu 2021, 24). While the government’s attempts to salvage the medical situation in Wuhan may have caused many to criticize perceived incompetence and failure to service citizen needs, later events caused public opinion to shift in generally more positive directions.

Additionally, our study takes steps to advance the study of authoritarian responsiveness, particularly in an era in which much of the information acquisition and communication between the government and citizens occurs on the internet. By developing a typology of authoritarian responsiveness, we distinguish our case from earlier literature that focuses on responsiveness that takes place either through well-institutionalized channels (such as the petition system) or is initiated by government request (such as consultation). In this sense, our contribution resembles that of Zheng and Meng (2021, 1), who write that “institutionalized and noninstitutionalized channels [of responsiveness] are not substituting but complements, and improving responsiveness to institutionalized participation may ironically lead to more protests.” Our findings also echo the work of Wang and Dickson (2021), who find that corruption investigations in China undermine regime support by revealing information about the magnitude of corruption in institutions.

We believe that our theoretical framework is highly useful for making sense of these previous studies, as well as the empirical phenomena we focus on, in the context of the existing literature on responsiveness, which typically portrays responsiveness as a tool to secure authoritarian stability. Our focus on understanding the mechanics and dynamics *responsiveness* further differentiates our work from that of Lu, Pan, and Xu (2021), since their study is focused on public sentiment in a time period, whereas our study aims at discovering how responsiveness was perceived and received in a narrower discourse.

We caution that our study should not be taken as evidence that authoritarian regimes should cease to be responsive to citizen input. As others have pointed out, regimes may adopt a “consultative authoritarianism” where citizen demands and complaints are incorporated in the policy making process, but are not exposed to the general public

(Distelhorst and Hou 2017). Even in instances when citizen demands and complaints are publicly known, regimes may still enjoy long-term benefits.

We urge readers to remain conscientious of the limitations of our study. We conducted our study during the period that the COVID outbreak spread from China to the rest of the world. Our contemporaneous approach offers many advantages, such as collecting social media responses in real time, but also creates several research design challenges. While the ideal environment for our survey experiment would have been during the initial COVID-19 outbreak in China, we did not conduct the experiment until several months later. Although we used a priming video to mitigate this problem, recall and confirmation biases were still likely at play. Furthermore, although the direction of the experimentally identified causal impact is predictive of the predominantly negative evaluations of the state on social media, social media posts are likely contaminated by a variety of other factors, such as international comparisons and state interventions in media coverage. Lastly, Weibo restrictions on real-time data collection strengthened over the period of our data collection, so we were unable to collect additional data after mid-March 2020 in order to assess longer-term effects.

Bearing in mind these limitations, our work speaks to questions posed by many international observers regarding the risks that COVID-19 poses to the legitimacy of the CCP and the stability of its rule. Our study illuminates one facet of the regime's approach to manage COVID-19—namely, soliciting direct public requests for aid on a popular public platform—that stands in contrast to the rest of the regime's management strategy, which has been portrayed as heavy-handed and lacking compassion. While this strategy evidently engendered a short-term backlash, which our paper has investigated in depth, it has evidently not led to a long-term complete collapse of citizen trust in the CCP, nor any major restructuring within the CCP itself. Nonetheless, our work highlights some of the unique risks and possibilities generated by governments engaging directly with citizens online. For instance, the fact that a *localized* crisis—the inability of residents of one province, mostly in one city to access medical care—generalized a *national* backlash is a product of the fast-paced and open information environment of social media. In contrast to previous studies which tend to focus on explanations for the existence of responsive institutions, this finding spotlights the complicated ways that responsiveness feeds back into public opinion. Our findings call for future research that investigates how, why, and when authoritarian governments use the internet to both manage citizen complaints and address citizen needs in both times of crisis and business as usual.

Notes

1. Posts tagged with the special hashtag can be found at: https://www.weibo.com/p/1008084882401a015244a2ab18ee43f7772d6f/super_index
2. <https://finance.sina.com.cn/stock/usstock/c/2020-05-19/doc-iircuyvi3942770.shtml>
3. The platform of People's Daily is still accessible as of August 1, 2020: <https://activity.peopleapp.com/qiuzhutongdao/>
4. https://weibo.com/p/1008084882401a015244a2ab18ee43f7772d6f/super_index; website accessed on July 29, 2020
5. https://weibo.com/p/1008084882401a015244a2ab18ee43f7772d6f/super_index; website accessed on July 29, 2020
6. The COVID crisis peaked in China in February 2020 with the highest daily new cases numbering 14,108. In July 2020, the daily number of new cases in China was virtually zero. See: <https://www.worldometers.info/coronavirus/country/china/>
7. Please contact the authors if other aspects of the video are of interest.
8. We note that our real-time data collection strategy does not address possible biases that resulted from of ex-ante filtering, namely preventing posts about sensitive topics in the first place.
9. Posts that have more positive words than negative words are considered positive, and vice versa.
10. The categorization regime used by human coders is identical to the one used in help-seeking posts, with the exception that we do not differentiate between "explicit negative political" and "implicit negative political" posts. This was done to make them matchable to labels automatically generated by the dictionary method, which cannot separate explicit and implicit sentiments.
11. Posts which did not show any clear sentiment were excluded from our analysis, and therefore are not included in this figure.
12. We cannot rule out the possibility that this change in composition was driven at least in part from improving censorship tactics, though our real-time scraping has evaded some effects from ex-post censorship.

Acknowledgment

We thank Yu Xie, Rory Truex, Dean Knox, participants of the Center on Contemporary China summer workshop, POL574 class at Princeton University, and participants of QCSS 2021 conference for useful advice. We thank Jingyi Li, Xinrou Cai, Jiaqi Gao, Qing Yuan, Kai Zhang, Haoyu Zhang, and Yunming Yang for research assistance.

Funding

The research is conducted with the funding provided by the Data-Driven Social Science Initiative at Princeton University.

ORCID

Han Zhang  <http://orcid.org/0000-0003-2912-8780>

References

- Chen, J., J. Pan, and Y. Xu. 2016. "Sources of Authoritarian Responsiveness: A Field Experiment in China." *American Journal of Political Science* 60 (2): 383–400. doi: [10.1111/ajps.12207](https://doi.org/10.1111/ajps.12207).
- Chen, J., and Y. Xu. 2017. "Why Do Authoritarian Regimes Allow Citizens to Voice Opinions Publicly?" *The Journal of Politics* 79 (3):792–803. doi: [10.1086/690303](https://doi.org/10.1086/690303).
- Chen, X. 2012. *Social Protest and Contentious Authoritarianism in China*. Cambridge, NY: Cambridge University Press.
- Coppock, A. 2019. "Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach." *Political Science Research and Methods* 7 (3):613–628. doi: [10.1017/psrm.2018.10](https://doi.org/10.1017/psrm.2018.10).
- Coppock, A., T. J. Leeper, and K. J. Mullinix. 2018. "Generalizability of Heterogeneous Treatment Effect Estimates across Samples." *Proceedings of the National Academy of Sciences of the United States of America* 115 (49): 12441–12446. doi: [10.1073/pnas.1808083115](https://doi.org/10.1073/pnas.1808083115).
- Coppock, A., and O. A. McClellan. 2019. "Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents." *Research & Politics* 6 (1):205316801882217. doi: [10.1177/2053168018822174](https://doi.org/10.1177/2053168018822174).
- Distelhorst, G., and Y. Hou. 2017. "Constituency Service Under Nondemocratic Rule: Evidence from China." *The Journal of Politics* 79 (3):1024–1040. doi: [10.1086/690948](https://doi.org/10.1086/690948).
- Freedman, D. A. 2008. "On Regression Adjustments to Experimental Data." *Advances in Applied Mathematics* 40 (2):180–193. doi: [10.1016/j.aam.2006.12.003](https://doi.org/10.1016/j.aam.2006.12.003).
- Fu, K.-W., C.-H. Chan, and M. Chau. 2013. "Assessing Censorship on Microblogs in China: Discriminatory Keyword Analysis and the Real-Name Registration Policy." *IEEE Internet Computing* 17 (3):42–50. doi: [10.1109/MIC.2013.28](https://doi.org/10.1109/MIC.2013.28).
- Fu, K.-W., and Y. Zhu. 2020. "Did the World Overlook the Media's Early Warning of COVID-19?" *Journal of Risk Research* 23 (7–8):1047–1051. doi: [10.1080/13669877.2020.1756380](https://doi.org/10.1080/13669877.2020.1756380).
- Grimmer, J., and B. M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21 (3):267–297. doi: [10.1093/pan/mps028](https://doi.org/10.1093/pan/mps028).
- He, B., and S. Thøgersen. 2010. "Giving the People a Voice? Experiments with Consultative Authoritarian Institutions in China." *Journal of Contemporary China* 19 (66):675–692. doi: [10.1080/10670564.2010.485404](https://doi.org/10.1080/10670564.2010.485404).
- He, B., and M. E. Warren. 2011. "Authoritarian Deliberation: The Deliberative Turn in Chinese Political Development." *Perspectives on Politics* 9 (2):269–289. doi: [10.1017/S1537592711000892](https://doi.org/10.1017/S1537592711000892).
- Heurlin, C. 2016. *Responsive Authoritarianism in China*. New York: Cambridge University Press.
- Hollyer, J. R., B. P. Rosendorff, and J. R. Vreeland. 2015. "Transparency, Protest, and Autocratic Instability." *American Political Science Review* 109 (4):764–784. doi: [10.1017/S0003055415000428](https://doi.org/10.1017/S0003055415000428).
- Hopkins, D. J., and G. King. 2010. "A Method of Automated Nonparametric Content Analysis for Social Science." *American Journal of Political Science* 54 (1):229–247. doi: [10.1111/j.1540-5907.2009.00428.x](https://doi.org/10.1111/j.1540-5907.2009.00428.x).

- Hu, Y., H. Yan Huang, A. Fan Chen, and X. L. Mao. 2020. "Weibo-COV: A Large-Scale COVID-19 Social Media Dataset from Weibo." In *Proceedings of the 1st Workshop on NLP for COVID-19@ EMNLP 2020*, Online. Association for Computational Linguistics 2020.
- Huang, H., and X. Liu. 2018. "Historical Knowledge and National Identity: Evidence from China." *Research & Politics* 5 (3):205316801879435. doi: [10.1177/2053168018794352](https://doi.org/10.1177/2053168018794352).
- Jiang, J., T. Meng, and Q. Zhang. 2019. "From Internet to Social Safety Net: The Policy Consequences of Online Participation in China." *Governance* 32 (3): 531–546. doi: [10.1111/gove.12391](https://doi.org/10.1111/gove.12391).
- King, G., J. Pan, and M. E. Roberts. 2013. "How Censorship in China Allows Government Criticism But Silences Collective Expression." *American Political Science Review* 107 (2):326–343. doi: [10.1017/S0003055413000014](https://doi.org/10.1017/S0003055413000014).
- Kornreich, Y. 2019. "Authoritarian Responsiveness: Online Consultation with "Issue Publics" in China." *Governance* 32 (3):547–564. doi: [10.1111/gove.12393](https://doi.org/10.1111/gove.12393).
- Ku, L.-W., Y.-T. Liang, and H.-H. Chen. 2006. "Opinion Extraction, Summarization and Tracking in News and Blog Corpora." In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, Vol. 100107, 1–167.
- Lee, C. K., and Y. Zhang. 2013. "The Power of Instability: Unraveling the Microfoundations of Bargained Authoritarianism in China." *American Journal of Sociology* 118 (6):1475–1508. doi: [10.1086/670802](https://doi.org/10.1086/670802).
- Lorentzen, P. 2014. "China's Strategic Censorship." *American Journal of Political Science* 58 (2):402–414. doi: [10.1111/ajps.12065](https://doi.org/10.1111/ajps.12065).
- Lu, Y., J. Pan, and Y. Xu. 2021. "Public Sentiment on Chinese Social Media During the Emergence of COVID19." *Journal of Quantitative Description: Digital Media* 1. doi: [10.51685/jqd.2021.013](https://doi.org/10.51685/jqd.2021.013).
- Lupton, D. L. 2019. "The External Validity of College Student Subject Pools in Experimental Research: A Cross-Sample Comparison of Treatment Effect Heterogeneity." *Political Analysis* 27 (1):90–97. doi: [10.1017/pan.2018.42](https://doi.org/10.1017/pan.2018.42).
- Malesky, E., P. Schuler, and A. Tran. 2012. "The Adverse Effects of Sunshine: A Field Experiment on Legislative Transparency in an Authoritarian Assembly." *American Political Science Review* 106 (4):762–786. doi: [10.1017/S0003055412000408](https://doi.org/10.1017/S0003055412000408).
- Mattingly, D., and E. Yao. 2020. "How Propaganda Manipulates Emotion to Fuel Nationalism: Experimental Evidence from China." *Available at SSRN* 3514716.
- Minzner, C. F. 2006. "Xinfang: An Alternative to Formal Chinese Legal Institutions." *Stanford Journal of International Law* 42 : 103.
- Mullinix, K. J., T. J. Leeper, J. N. Druckman, and J. Freese. 2015. "The Generalizability of Survey Experiments." *Journal of Experimental Political Science* 2 (2):109–138. doi: [10.1017/XPS.2015.19](https://doi.org/10.1017/XPS.2015.19).
- Su, Z., and T. Meng. 2016. "Selective Responsiveness: Online Public Demands and Government Responsiveness in Authoritarian China." *Social Science Research* 59:52–67. doi: [10.1016/j.ssresearch.2016.04.017](https://doi.org/10.1016/j.ssresearch.2016.04.017).
- Truex, R. 2016. *Making Autocracy Work: Representation and Responsiveness in Modern China*. New York: Cambridge University Press.
- Truex, R. 2017. "Consultative Authoritarianism and Its Limits." *Comparative Political Studies* 50 (3):329–361. doi: [10.1177/0010414014534196](https://doi.org/10.1177/0010414014534196).

- Wang, Y., and B. J. Dickson. 2021. "How Corruption Investigations Undermine Regime Support: Evidence from China." *Political Science Research and Methods*:1–16. doi:[10.1017/psrm.2021.27](https://doi.org/10.1017/psrm.2021.27).
- Yang, G. 2009. *The Power of the Internet in China: Citizen Activism Online*. New York: Columbia University Press.
- You, Y., Y. Huang, and Y. Zhuang. 2020. "Natural Disaster and Political Trust: A Natural Experiment Study of the Impact of the Wenchuan Earthquake." *Chinese Journal of Sociology* 6 (1):140–165. doi: [10.1177/2057150X19891880](https://doi.org/10.1177/2057150X19891880).
- Zhang, H., and J. Pan. 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. doi: [10.1177/0081175019860244](https://doi.org/10.1177/0081175019860244).
- Zhang, Y., J. Liu, and J.-R. Wen. 2018. "Nationalism on Weibo: Towards a Multifaceted Understanding of Chinese Nationalism." *The China Quarterly* 235:758–783. doi: [10.1017/S0305741018000863](https://doi.org/10.1017/S0305741018000863).
- Zheng, S., and T. Meng. 2021. "The Paradox of Responsiveness and Social Protest in China." *Journal of Contemporary China* 30 (128):212–232. doi: [10.1080/10670564.2020.1790899](https://doi.org/10.1080/10670564.2020.1790899).

About the authors

Lai Wei is a PhD student in the Department of Sociology at Princeton University, where he is also affiliated with the Paul and Marcia Wythes Center on Contemporary China. Prior to Princeton he obtained BA in sociology from Tsinghua University. His current research interests are generally about politics and inequality in China, as well as quantitative and computational methods.

Elaine Yao is a PhD student in Politics at Princeton University. Prior to doctoral studies, she studied Economics at the University of Chicago and worked as a Research Analyst in the Federal Reserve Bank of New York, where she worked on economics research and policy related to consumer finance, bank capitalization, and labor economics. Her current interests lie generally in the political economy of authoritarian regimes, with a regional focus on China.

Han Zhang is an Assistant Professor in Division of Social Science at The Hong Kong University of Science and Technology. He obtained his PhD in Sociology from Princeton University, and his B.S. in Computer Science and B.A. from Peking University. His research interests include computational social science, social movements, social networks, and methodology. His past research have been published in *Sociological Methodology* and have won the Mayer N. Zald Distinguished Contribution to Scholarship Student Paper Award from the Section on Collective Behavior and Social Movements of American Sociological Association.

Appendices

A. Covariate balance

Figure A1 shows the covariate balance between the control group and the two treatment groups respectively. The control group does not show a statistically

significant difference compared with the two treatment groups in any respondent characteristics, showing that the covariates are well balanced by randomization.

B. Treatment effect heterogeneity

Table A1 displays the coefficients of the interaction terms from regressions of government effectiveness on the treatment condition and demographic covariates. Columns refer to specific government effectiveness outcomes, and rows display demographic information. Each number represents the coefficient of the interaction of a single demographic variable with the “ineffective state” treatment. We only display treatment effects for the “ineffective state” treatment group because there is no discernible difference between the two treatment groups. According to the interaction terms, there is barely any discernible treatment effect heterogeneity on all outcomes, with the one exception that the treatment effect is significantly more negative for party members in the central government effectiveness outcome. In plain words, party members’ beliefs about the central government are more negatively impacted when seeing the help-seeking posts compared with non-CCP members. However, we

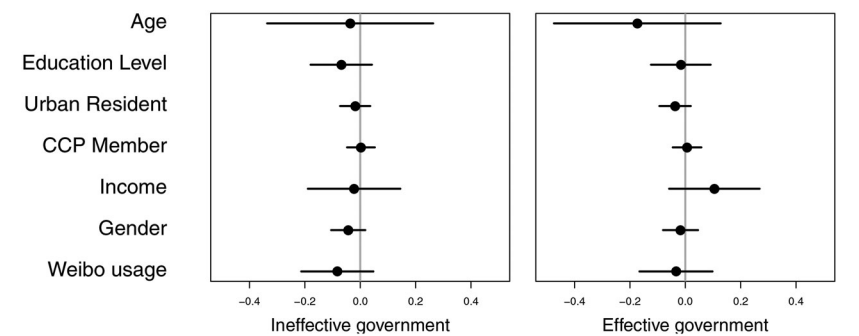


Figure A1. Covariate balance. The mean differences between the treatment and control groups. 95% confidence intervals from two-sided t-test also are plotted.

Table A1. Treatment effect heterogeneity on perceived government effectiveness for different demographic groups.

	Central gov effectiveness	Provincial gov effectiveness	Hubei gov effectiveness	National gov economic effectiveness	Provincial gov economic effectiveness
Female	0.09	0.01	−0.02	0.05	0.08
Urban hukou	0.01	−0.14	−0.25	0.17	0.08
Years of education	0.00	0.00	0.01	0.01	0.00
Income	−0.00	−0.00	−0.00	−0.00	−0.00
Age	0.01	−0.00	0.00	−0.00	0.00
Party member	−0.28***	−0.14	−0.01	0.01	0.01

Table A2. Pre-processing steps of Weibo data of help-seeking discussions.

	Sample size	Description
Original	49,893	
Removing duplicated content	44,298	
Removing official accounts	41,318	Official account identified when the username includes one of the following: 司法 (legislative); 播报 (broadcast); 宣传 (propaganda); 公安 (police); 热线 (hotline); 交通 (transportation); 法院 (court); 局 (bureau); 在线 (online); 交警 (transportation police); 新闻 (news); 快讯 (live news); 直播 (live); 关注 (focus); 百态 (overview); 广播 (radio); 电视台 (TV station); 电视 (TV); 报 (newspaper)
Keeping posts including COVID-relevant keywords	10,495	COVID-relevant keywords include: 疫; 医; 住院; 病; 发烧; 高烧; 肺; 救治; 救助; 谣; 封城; 封锁; 患; 症状; 新冠; 病毒; 武汉; 湖北; 床位; 确诊; 疑似; 感染; 隔离; 去世; 化验; 体温; 发热; 检测; 核酸; 防控
Removing posts that do not express sentiments	5,234	Removed posts include irrelevant content; sentiment-free content; forwarded news or information
Removing posts whose expressed sentiments three labelers disagree on	4,872	

Table A3. Inter-coder consistency of the discussions about Weibo help-seekers.

	Number of categories	Accuracy	Kappa
Original	5	0.62	0.51
Collapsing explicit and implicit anti-government	4	0.73	0.58
Collapsing to only positive and negative	2	0.90	0.70

recommend the readers to treat this finding with caution because it could be an artifact of multiple hypotheses testing.

C. Help-seeking Weibo posts pre-processing

To ensure data quality, each post was labelled twice. Inter-coder accuracy is summarized in Table A3. In order to address inter-coder inconsistencies, a third coder was assigned to all posts whose labels the original two coders disagreed on. If all three labels differed, the posts were dropped from the data set; otherwise, the category was determined by majority vote. Under the finest five-category regime, inter-coder accuracy is only 62% and the Kappa statistic is 0.51. While less than ideal, this is understandable given that social media posts are not structured to express opinion or sentiments clearly. Inter-coder consistency improves substantially for simpler coding regimes: when we collapse explicit and implicit anti-government sentiment, the two most easily confused categories, the accuracy increases to 73%. For the simple binary classification of “positive” and “negative,”

Table A4. Pre-processing steps for the general COVID-related posts.

	Number of Posts	Description
Removing sentiment-free posts and posts by official accounts	413,150	Official accounts defined in the same way in the pre-processing of help-seeking posts; sentiment-free posts are those that contain zero sentiment words as defined by the dictionary
Removing help-seeking posts	405,625	Removing posts with the word “help-seeking” (求助)
Removing irrelevant posts	359,233	The number is estimated from the calibrated distribution

Table A5. Frequency Confusion Matrix of the Sentiments of General COVID posts. The columns are dictionary-based labels and the rows are human labels.

	Negative Political	Negative Emotional	Positive Political	Positive Emotional
Negative Political	24	25	4	4
Negative Emotional	73	193	17	29
Positive Political	42	17	63	23
Positive Emotional	46	83	56	98
Irrelevant	5	53	15	21

inter-coder accuracy is 90%. The inter-coder accuracy, while lower than ideal for finer categorization regimes, does not significantly impact our main conclusions, which depend primarily on the binary positive-negative classification.

D. General COVID Weibo posts pre-processing

To create a baseline comparison group, or a “control group” for the sentiment of help-seeking posts, we collected posts that include the keyword “COVID” (新冠) published in the same time period as the help-seeking posts. These posts reflect the public sentiment in the discussion of the COVID crisis in general, beyond the case of COVID help-seeking.

The preprocessing steps and the number of remaining posts after each step can be found in [Table A4](#). We removed posts containing the keyword “help-seeking,” and posts published by official accounts. We then continued preprocessing using two dictionaries, one about sentiment words and another about political words. The sentiment dictionary is the National Taiwan University Semantic Dictionary (NTUSD), a widely used tool for natural language processing in Chinese (Ku, Liang, and H.-H. Chen 2006), and the political dictionary was compiled by the authors and contains 2,127 Chinese political words. Following the preprocessing steps used on the help-seeking posts, we removed the posts that contain neither positive sentiment words nor negative sentiment words. We also removed posts which had an equal number of positive and negative words.

We classified the remaining posts into positive political, negative political, positive emotional, and negative emotional categories based on dictionary methods. Specifically, we classify one post as positive if positive words outnumber negative words, and negative otherwise. We classify one post as political if there is any political word in the post. To further improve the results, we applied the method in Hopkins and King (2010), which proposed a better method to

Table A6. Top 25 Most Frequent Words in Posts by Sentiment Category. “Percent” means the sum of that word among all posts divided by the number of posts, so it can exceed 1. Non-Chinese characters and generic Weibo words are not included, which include “微博,” “超话,” “收起全文”.

Negative Political		Negative Emotion		Positive Political		Positive Emotion	
Word	Percent	Word	Percent	Word	Percent	Word	Percent
武汉	0.82	武汉	0.49	中国	1.24	武汉	0.78
没有	0.47	看到	0.45	疫情	0.69	疫情	0.53
说	0.39	真的	0.43	国家	0.62	希望	0.46
患者	0.39	没有	0.39	武汉	0.57	看到	0.43
肺炎	0.34	肺炎	0.39	塞尔维亚	0.43	肺炎	0.42
一个	0.33	一个	0.39	肺炎	0.39	一个	0.42
真的	0.31	疫情	0.36	一个	0.37	患者	0.37
疫情	0.30	说	0.31	看到	0.35	会	0.37
医院	0.28	希望	0.29	说	0.33	说	0.35
现在	0.28	会	0.29	加油	0.30	没有	0.34
看到	0.28	患者	0.28	会	0.29	医院	0.33
知道	0.25	每天	0.26	真的	0.26	真的	0.32
会	0.24	信息	0.25	没有	0.26	信息	0.32
没	0.21	想	0.25	已经	0.26	已经	0.27
湖北	0.21	知道	0.24	医院	0.25	很多	0.27
信息	0.20	现在	0.24	现在	0.25	今天	0.27
物资	0.20	已经	0.19	希望	0.23	加油	0.26
社区	0.18	太	0.19	总统	0.23	感谢	0.24
希望	0.18	转发	0.18	患者	0.22	转发	0.23
很多	0.17	觉得	0.18	政府	0.21	帮助	0.23
隔离	0.17	没	0.17	很多	0.19	现在	0.22
想	0.17	里	0.17	一定	0.19	做	0.21
做	0.16	医院	0.17	确诊	0.19	每天	0.19
确诊	0.16	很多	0.16	人民	0.18	朋友	0.19
排查	0.15	做	0.15	物资	0.18	得到	0.18

aggregate machine predictions on individual text documents into counts and proportions. We manually labelled around 1,000 posts and compared the predictions by the dictionary method with human labels.

The confusion matrix is shown in Table A5. Denote the dictionary-based distribution as a 1×4 vector F , and the transpose of the probability confusion matrix, that is all cells in Table A5 divided by the total number of labelled posts, as C^T which is 4×5 . Then the calibrated distribution \tilde{F} is simply calculated as:

$$\tilde{F} = FC^T,$$

which is a consistent estimator for the true distribution.

E. Other Weibo post results

Table A6 shows the most frequent words for posts of different categories. The frequent words under each category are marginally informative and correspondent with our understanding of the label. For instance, the word “中国” (China) occurs more than once per post in the positive political category.

Figure A2 shows the temporal trend of the number of help-seeking posts and the number of discussions around them. The top panel shows the trend of actual help-seeking posts.¹³ The frequency of help-seeking reached a peak in early

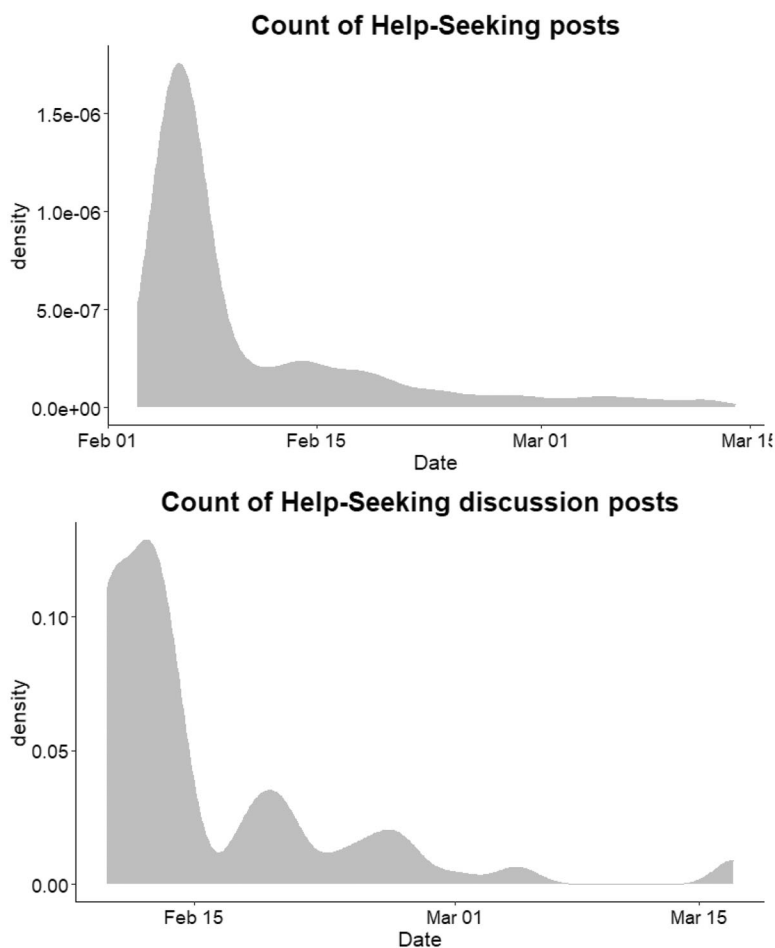


Figure A2. Temporal trend of actual help-seeking and discussion of help-seeking on Weibo.

February, when around 150 posts were published per day, but the numbers of help-seeking posts soon declined and eventually stopped in mid-March. The decline is partly attributable to the improving medical situation in Wuhan, and partly attributed to increasing censorship on relevant posts. The bottom panel shows the temporal trend of discussions around the help-seekers. The discussions follow an almost identical trend with the actual help-seeking posts, with a peak in early February and a steady decline since.

Figure A3 shows the geographical distribution of Weibo posts by sentiment category. Among all posters with geographic location information, Hubei province has the largest number of Weibo users discussing online help-seeking. This makes sense given that Hubei province was the epicenter of the COVID-19 outbreak in China. Apart from Hubei, more developed provinces generally have more users posting on this topic, which likely results from the fact that more

F. Primers and treatment conditions

F.1. Transcript of priming video

The following is an English translation of the text appearing in our priming video, which described the footage being shown.

(Medical workers shouting) “United in will, using scientific prevention, Wuhan must be victorious!”

Strong attention has been paid to the novel coronavirus infectious outbreak. A large number of medical workers are on the front lines of the fight against the virus. On January 22, we visited several hospitals in Wuhan. We took footage of the fevered outpatient clinic, the quarantine ward, and other scenes. The people in Wuhan Union Hospital wear masks, and medical personnel are all wearing personal protective equipment.

At the Renmin Hospital of Wuhan University, the hospital has already invested in placing 500–600 people in disease prevention and control. At Zhongnan Hospital of Wuhan University, human-powered cardiac and respiratory support was able to successfully rescue a coronavirus patient. At No. 6 Wuhan Hospital, medical workers took body temperature measurements routinely. Wuhan Hospital Medical Emergency Center medical workers are currently undertaking disinfecting measures and transporting sick patients. At Hubei Provincial Hospital of Traditional Chinese Medicine, medical workers are currently calibrating intravenous devices. One patient, who had been on maternity leave, had to return to the emergency room with a 6-month-old child.

The workers at Wuhan Jinyintan Hospital loudly shouted, “Wuhan must be victorious!”

F.2. Transcript of treatment

The following is an English translation of the text of the Weibo posts that were used as treatments. Names and other identifying information are redacted.

Post 1:

City: Wuhan

Date of illness: January 21

Description: My mother felt a headache, full-body soreness, and exhaustion on January 21, and thought it was a cold. Starting from January 28, she suffered from persistently high fever of 39 degrees Celsius. We went to Wuhan #1 People hospital in the evening of January 28 and was unable to see a doctor due to the long line of patients waiting to see the doctor. We went home and used a rectal fever reducer for her, which caused her temperature to go down slightly, but remained high at 37°C. At that point, the patient was disoriented and unable to go to hospital. The patient stayed at home and we used physical means (applying ice) to lower her temperature. However, on February 1, her temperature went up to 39.3°C but the patient was too weak to go to the hospital on that day. On February 2, the patient went to the hospital. The CT scan showed that the patient was highly suspected to have coronavirus. The result of the nucleic acid test is “suspicious/probable”. The hospital advised the patient to go back to her own community and seek quarantine arrangements from the community. But the community in charge responded that there was no hurry in getting community quarantine arrangements and that the patient should be isolated at home. Our home has only 40 m² of living space, hence it was impossible to have the patient

isolated. During the time that the patient stayed at home, she suffered from high fevers and was unable to eat or sleep. In the past, she had occasionally had disrupted sleep due to anxiety/depression. Now, the situation is getting worse as she suffered from pneumonia as well. She then had trouble breathing. We desperately went to Wuhan Xiehe Hospital to see the doctor. Her CT scan once again showed that she was extremely highly suspected to have coronavirus, but her nucleic acid test was not able to be performed due to lack of testing kits. The fact is, my mother really cannot wait anymore for treatment. The doctor again asked us to seek quarantine arrangements from the community, and the community again said they were unable to provide quarantine arrangements. My mother was given IV for two days at Xiehe Hospital but her high fever remained the same. We now apply ice at home. There are four people in the family. I also have a persistent high fever. I never expected that I would catch the virus as well. The doctor said my CT scan result was suspicious/probable and prescribed medication and asked me to quarantine at home. I feel like killing myself after taking medication for 8-9 days without effect. Right now I am going to different hospitals and hope I can be admitted, but I have failed to get myself admitted. Luckily, my father and my sister are not showing any symptoms yet. I ask the government, news agencies, and people with compassion to disseminate this information and save our family.

Post 2:

City: Wuhan

Date of illness: Feb 4, 2020

Description: We went to Gutian street clinic on February 4. The X-ray showed that the lung was infected, and the doctor asked us to go to Puai hospital in Gutian Three Road. The CT scan taken there showed the patient was highly suspected to have coronavirus. The nucleic acid test was not performed. The patient has been unable to get a test because we do not have a car to go the assigned hospital, and the community is unable to dispatch a car to take us there either. The CT scan shows white spots and patches in most of her lungs. The patient's condition is very serious. Her symptoms include difficulty breathing, vomiting, diarrhea, exhaustion, and inability to eat. She has suffered from diabetes, high blood pressure, stomach ulcer, and she is in danger of getting complications. We have called the Mayor's hotline, assistance hotline; we have contacted People's Daily, State Council Pandemic Committee; we have asked for help from the community; all to no avail. We are told that our situation will need to be reported and we need to wait. Now we ask the help from any friends who have a connection with the hospital. We are extremely grateful for your help.

Copyright of Chinese Sociological Review is the property of Taylor & Francis Ltd and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.