Social Media and Government Responsiveness: Evidence from Vaccine Procurement in China

Yixin Mei (HKU) and Yanhui Wu (HKU & CEPR) *

June 27, 2022

Abstract

This paper studies whether and how public opinion on social media affects local governments' procurement of vaccines in China during 2014-2019. We exploit city-level variation in the eruption of social-media opinion on vaccine safety triggered by leaked information on a vaccine scandal, instrumented by the early penetration of social media into each city. We find that governments in cities exposed to stronger information eruption increased the frequency and share of more-transparent procurement. The effect is larger in cities where local officials are ranked lower or have stronger career concerns and where the information environment was more strictly controlled. Interestingly, the effect is muted in a silent real scandal but pronounced in a hot fake scandal. Our findings support that social media enhances top-down monitoring in policy implementation in nondemocracies.

JEL Code: H4, H7, I1, P2

Keywords: Social media, Government responsiveness, Authoritarian regimes, China, Procurement, Public health

*Mei: University of Hong Kong, yixin_mei@connect.hku.hk; Wu: University of Hong Kong, yanhuiwu@hku.hk

1 Introduction

It has been a deep concern that social media may undermine democracies while enhancing autocracies because of its tendency to spur populism and strengthen the power of political leaders (Morozov 2011; Howard 2010; Margetts et al. 2016; Sunstein 2018; Benkler et al. 2018). Whether this concern is valid depends importantly on the impact of social media on citizens' political participation and government responsiveness. Although the former has attracted substantial academic attention in economics, research on how social media affects government behavior remains limited.[1] In the world of traditional media, public opinion on the media plays an important role in enhancing government responsiveness to citizens' needs (e.g., Besley and Burgess 2002; Strömberg 2004; Snyder and Strömberg 2010; Besley and Dray 2022). However, it is less clear how a government may respond to social media information, which is noisy, hard to verify, and susceptible to manipulation. In this paper, we study whether and how information flows on social media affect local governments' implementation of a public-health policy—procurement of vaccines —in China from 2014 to 2019.

As the largest nondemocracy, China is generally viewed as a country with limited media freedom. Information flows are suppressed, and the political function of the media is crippled. On the other hand, it has been argued that, by allowing some vestiges of a free media, top leaders in authoritarian regimes can better acquire bottom-up information for monitoring local officials and correcting policy oversight (e.g., Egorov et al. 2009; Lorentzen 2014). This argument is particularly relevant in the era of social media. The new communication technology substantially facilitates the production and spread of information from bottom up and, at the same time, increases the regime's ability to collect and process massive information for surveillance. Thus, in areas where the regime and citizens share common interests, such as public health and environmental protection, social media has the potential to discipline local politicians' misconduct and improve policy implementation. If this conjecture turns out true, the effect of social media in autocracies may bear some resemblance to that in democracies where politicians address citizens' needs voiced in the media to win voters' support. That is, without democratization, an authoritarian system can improve its governance with the help of social media. To what extent can one draw a parallel between the media effects in democracies and nondemocracies? One important goal of our study is to answer this question and uncover the mechanisms underpinning the interaction between social media and political incentive in nondemocracies.

Our study focuses on government procurement of vaccines for three reasons. First, government procurement accounts for a sizeable share of public expenditure worldwide and constitutes an important aspect of state capacity (Best et al. 2019; Bosio et al. 2022). In developing countries, the lack of transparency and misuse of policy discretion in public procurement is a notable problem in governments' provision of goods and services (Szucs 2017; Gerardino et al. 2017). Second, the safety of vaccine provision has been a salient social problem in China, as witnessed by a series of scandals in the last two decades. Widespread dissatisfaction with such

[1] For surveys on the political effects of the media, see DellaVigna and Gentzkow (2010), Prat and Strömberg (2013), Strömberg (2015a, 2015b), DellaVigna and La Ferrara (2015), and particularly Zhuravskaya et al. (2020), which review recent studies on social media.

a livelihood issue poses a threat to the regime by diminishing citizens' trust of the government and arousing questioning of regime legitimacy. Public opinion can affect local officials' behavior either through their internalization of the regime's political goal of maintaining regime stability or through their fear of top-down intervention which may jeopardize their careers. Third, Chinese social media are abundant in discussions about vaccine issues. This provides an opportunity to investigate the effect of free information flows in a generally strictly controlled media environment.

We assemble a dataset on the procurement of over 20,000 vaccine-related items by Chinese city-level (prefectural) governments during 2014-2019. Since 2005, a national regulation has required that all vaccines and related products and services must be procured by governments before being supplied to health providers. However, the specific arrangement of procurement is decentralized to city governments. This decentralization gives rise to discretionary policy execution and scope for corruption. Notably, many local governments deviate from the regulatory guidance of using an open-bid (competitive auction) format and use the less-transparent procurement formats such as invited bidding and private negotiation. Our main outcome variables are the frequency and share of open-bid procurement made by a city government in a month.

Regarding the social media data, we obtained more than three million posts referencing "vaccine" published on Sina Weibo (Weibo for short)—the Chinese equivalent to Twitter—during our sample period. Previous studies have documented that Weibo is abundant in general discussions about social problems as long as these discussions do not challenge the ruling of the Chinese Communist Party (CCP), target top political leaders, or incite collective action (King et al. 2013, 2014; Qin et al. 2017). Based on the patterns of post eruption and deletion, we are confident that most Weibo posts about vaccines are unlikely to be censored. Even if some posts are censored, the remaining ones are precisely those that are commonly visible and are crucial for social media to have an impact. Because of the central government's direct control over social media, subnational governments have limited influence on information flows on Weibo. Instead, many local governments purchase massive-information-processing software or AI systems to monitor public opinion on social media. We use two common machine learning techniques—perhaps similar to those monitoring systems used by local governments—to extract Weibo opinion about vaccine safety and gauge public sentiment towards vaccine issues in a city.

To identify causal effects, we leverage the drastic change in the information landscape on social media prompted by some sudden events. During a normal time, social media discussion about vaccines is flat and evenly distributed across regions. Upon an unexpected shock (e.g., a scandal), social media discussion is dramatically heated up, and its rugged regional distribution, after controlling for predetermined economic and social conditions and evolving trends, can be singled out to identify the media effect. Specifically, we measure the strength of public opinion on vaccine safety in a city by the per-capita number of vaccine posts originating from that city and classified by our machine learning algorithms as "monitoring posts" (i.e., expressing public grievances and questioning government misconduct). We then construct a city-specific information shock on Weibo (Weibo shock for short) by calculating the change in our strength measure after a vaccine scandal within a narrow time window. Finally, we interact this change with a dummy for the event timing to form a difference-in-differences (DID) estimation with city-month observations. In addition to city fixed effects and year-month fixed effects, we also control for province-specific time trends and general informational flows on social media over time to isolate the effect of the event-induced Weibo shock. To enhance identification, we instrument the Weibo shock with a measure of Weibo penetration into a city during its early entry period, which is arguably random after controlling for technological and educational factors.

Our main empirical analysis focuses on a vaccine scandal in March 2016. A drug ring from Shandong province had sold defective and expired vaccines (worth 88 million US dollars) for years. The scandal was kept secret until it was unexpectedly leaked to the public by an online media outlet on March 18. The news was so sudden that the number of Weibo posts referencing vaccines jumped from 1,039 on March 17 to an average of 11,239 in the next two days and to a spike of 75,250 on March 22. Although the scandal was local, it triggered national discussions of vaccine issues on Weibo. Many posts were noisy, very few targeted at specific officials, firms, or individuals; but a significant number of posts accused governments for procuring vaccines behind the scenes and failing to monitor the circulation of vaccines.

Applying the DID estimation described above to this event setting produces three results. First, in cities experiencing a stronger Weibo shock induced by the event, local governments significantly increased both the frequency and share of open-bid in their vaccine procurement. Putting the effect in perspective, if every citizen in a city tweeted one more post talking about vaccine safety in a month, the local government would have increased the share of open-bid procurement by more than 10%. There are no pretrends, and the effect persists even after half of a year. The IV estimation tends to strengthen the result. Interestingly, we find no significant difference in local governments' responsiveness between cities directly affected by the scandal and those not affected. Second, upon a more intensive Weibo shock, city governments were more likely to hire local procuring agencies so as to better control the procurement and procured more vaccines produced by multinational firms, instead of domestic firms whose products were called in questions. Third, immediately after the information eruption on Weibo, local governments blogged more actively about vaccine issues, changing topics from routine government reporting to public accountability. These results show that Chinese local governments rapidly responded to public opinion expressed on social media in an effort to enhance compliance with the central government's policy (e.g., the use of open-bid procurement and high-quality suppliers) and mitigate distrust of the government (e.g., blogging about public accountability).

We conduct additional analysis to discriminate between two views: a "benevolent-dictator" view in which the Weibo effect is driven by local governments' internalization of the regime's political preferences for addressing citizens' needs and a "strategic-autocrat" view in which the media effect is driven by local officials' fear of being punished by upper-level governments. When a government internalizes the regime's goal of addressing public needs to maintain regime stability, its behavior is analogous, to some extent, to politicians in a democracy who aim at pleasing citizens to maximize votes. In this case, the media effect is likely to be stronger in cities where the supply of vaccines was affected more directly by the scandal and the demand for vaccine quality is larger. Moreover, officials at a higher rank in the government hierarchy are more likely to internalize the regime's political preferences. In contrast, if a local official mostly worries that public opinion on social media triggers intervention from upper-level governments, which could jeopardize his promotion opportunities, the media effect is likely to be larger in cities where local officials are subject to greater top-down pressure and have stronger career concerns.

Three heterogeneous treatment effects regarding the impact of the Weibo shock on the adoption of open-bid vaccine procurement in our event study lend strong support to the strategic-autocrat view and the top-down-pressure mechanism. First, the Weibo effect is weaker in metropolitan areas and provincial capital cities, in which the demand for high-quality vaccines is larger and political leaders have a higher rank. Second, the Weibo effect is stronger in cities headed by political leaders who are younger and in the earlier stage of their tenure. Third, the Weibo effect is stronger in regions where censorship of social media is more extensive during the regular time. This last result is consistent with the view that, in an environment more scarce of public information, an unusual information outbreak is more likely to trigger top-down inspection.

To further verify the general presence of the social media effect and the top-down-monitoring mechanism, we examine the impact of Weibo information on the transparency of vaccine procurement other than the focal event in 2016. First, we investigate two contrasting "placebo" events in 2017: a silent real scandal and a hot fake scandal. The Weibo effect is muted in the former event, which did not spark extensive social media discussions because a concurrent sensational event crowded out public attention to vaccine issues. In contrast, in the latter event, which was just a pure accident but nevertheless triggered considerable social media discussions about vaccine safety, a significant Weibo effect is present. Second, we look into a scandal in July 2018, which spurred heated Weibo discussions about vaccine safely as the focal event. However, this scandal was caused by the discovery of a manufacturer's production of substandard rabies vaccines before they were delivered to the market, and public opinion focused on vaccine production and firm accountability instead of vaccine circulation and government accountability. We find a modest Weibo effect on government vaccine procurement, which is pronounced only in cities where there were vaccine manufacturers. Finally, beyond these events, during the entire period of 2014-2019, we find a positive correlation between Weibo information about vaccine issues originated from a city and the local government's transparency in vaccine procurement, which is driven only by city governments that had acquired mass-information-monitoring software to gauge public opinion on social media. These results show that Chinese local governments strategically respond to public opinion voiced on social media in their policy implementation, and their response is primarily driven by their perceived pressure of top-down inspection.

This paper contributes to the emerging literature on the political effect of social media. The focus of this literature has been on how social media affects citizens' political views and participation (e.g, Bakshy et al. 2015; Bursztyn et al. 2019; Allcott et al. 2020; Yanagizawa-Drott et al. 2021) and grassroots collective action (e.g., Acemoglu et al. 2018; Enikolopove et al. 2020; Qin et al. 2021). Research on the direct impact of social media on government accountability and politicians' behavior remains limited. Our study is among the first efforts that examine the effect of social media on government behavior and public policies. [2] We show that citizens' expression of sentimental opinions on social media, even without revealing any concrete information about specific governments or politicians, causes local governments to increase policy compliance and address public needs. This suggests that understanding the political role of the social media cannot be confined in the wisdom derived from the study of traditional media, which emphasizes that the media influences governments and politicians through informing the public of specific government conduct (Strömberg 2004; Ferraz and Finan 2008) or a politician's type (Besley and Burgess 2002; Snyder and Strömberg 2010) and through providing truthful information about certain social problems (e.g., Besley and Dray 2022). The public-sentiment channel is particularly important for understanding the political effect of social media in a strictly controlled information environment because public sentiment is difficult to silence and easy to spread. This is consistent with the finding in Qin et al. (2021) that the diffusion of public grievances on social media helps spread protests and strikes across Chinese cities.

Our study also contributes to the research on the political economy of nondemocracies. How information is distributed within the regime is key to the functioning of an authoritarian regime (Egorov and Sonin 2020). One particularly notable problem concerns subnational governments' informational advantages, which on one hand encourage local experimentation and adaptation but, on the other hand, create policy noncompliance and corruption (Xu 2011). Social media redistributes information in a way that enables the regime to better control the overall information flow (Edmond 2013) and monitor local officials (Lorentzen 2014), resulting in an enhanced authoritarian regime. The current paper attests to these theoretical arguments. The finding that an information outbreak on Weibo induced by a local event caused better policy compliance (the use of open-bid procurement) across the entire nation suggests that social media is highly useful for the central government to increase policy implementation, but it may bear the cost of policy rigidity and over-reaction. This tradeoff is essential in the theoretical study of the role of information in government organizations (e.g., Maskin et al. 2000; Qian et al. 2006). Moreover, our finding that the social media effect hinges largely on local officials' responses to the perceived top-down pressure is consistent with detailed sociological case studies (Zhou et al. 2012; Zhou 2017) and some experimental evidence (Chen et al. 2016; Anderson et al. 2019; Buntaine et al. 2021, 2022). The interaction between this "fear" mechanism and the mode of information flow helps understand some seemingly puzzling political phenomena in China, such as the insistence of a zero-COVID policy and the rigid implementation of the lockdown policies.

The remainder of this paper proceeds as follows. Section 2 provides a brief description of

[2] To our knowledge, several existing studies, without exploring the content and pattern of information circulated on social media as we do, are related to our paper in some certain aspect. Qin (2013) finds a correlation between social media penetration and the number of bad drugs detected by local drug administrations in China. Enikolopov et al. (2018) show that social media exposure of corruption in large state-controlled firms in Russia was associated with improvements in corporate governance of these firms, indirectly suggesting a positive effect of social media on public accountability. Gavazza et al. (2019) study the impact of the internet on local elections and government policies in the UK. Bessone et al. (2020) study how legislators respond when their constituencies have access to 3G mobile technology in Brazil and find that politicians increase their interaction with voters on Facebook but reduce their offline effort. Petrova et al. (2021) show that opening a Twitter account helps politicians compete for campaign contributions in the US.

the institutional background. Section 3 describes the data and the empirical strategy. Section 4 presents the basic results to test the main hypothesis, and Section 5 presents further evidence to shed light on the mechanisms. Section 6 concludes.

2 Background

2.1 Vaccination in China

China is the world second largest vaccine market, only after the US. However, vaccine safety, alongside drug quality, has been constantly ranked by Chinese as a top livelihood problem. The last decade has witnessed dozens of scandals involving the circulation of substandard and defective vaccines that severely affect the health of vaccinated children.

The Chinese government classifies vaccines into two categories. Category-I consists of 14 vaccines, including DPT and MMR. Jabs of Category-I vaccines are compulsory and freely offered by the government. Category-II includes common vaccines, such as chickenpox, flu, and rabies, as well as some substitutes for Category-I vaccines. Vaccination of Category-II is voluntary and paid by consumers themselves. The vaccination rates of Category-I are above 95% across China and close to 100% in big cities. However, the vaccination rates of Category-II are rather low, even in the most-developed regions. [3] In 2019, the sales of the Category-II vaccines in China reached 7.81 billion US dollar, accounting for 20% of the global market with an annual growth rate of 33% in the past five years.[4] During our sample period, Category-II vaccines were supplied by 35 vaccine manufacturers. Among them, 66% were domestic private firms, 25% state-owned enterprises (SOEs), and 9% multinational firms such as Merck and Pfizer.

Between vaccine manufacturers and consumers lies a complex network of distributors, service (e.g., transportation and storage) suppliers, and healthcare providers. Despite the vast overall market, regional markets are highly segmented due to the involvement of local governments. It is difficult for a manufacturer to enter a regional market without cooperation with distributors and service suppliers who have established good relationships with local governments. Because of the high agency costs, it was estimated that the prices of the Category-II vaccines that consumers ultimately pay were more than double the factory prices.

The lucrative business and heavy government involvement provide ample opportunities for corruption. One striking example is in 2007, when more than 100 children in Shanxi province died or became disable after receiving vaccination. Despite protests from angry parents, these incidences were covered up by local governments until 2010 when an investigative journalist from Beijing found out that the death and disability of the vaccinated children were caused by vaccines that were improperly exposed to high temperature but still approved to be used. Defective and substandard products, inappropriate storage and transportation, and illegal acquisition and distribution of vaccines are recurrent problems in the Chinese vaccine market.

[3] In 2017, only 300 to 3,500 doses of Category-II vaccines were given in every 10 thousand people in China. Even in the economically most-developed cities in the coastal areas, the vaccination rates of Category-II was below 20%. (Fu 2021)

[4] Global and China Vaccine Industry Report, Frost&Sullivan, March 2022.

2.2 Public Procurement

In China, public procurement constitutes an important part of government expenditure. In 2019, the amount of government procurement was over 510 billion US dollars, accounting for 10% of the overall fiscal expenses and 3.3% of the GDP. According to a law passed in 2003, goods and services purchased with public money above certain thresholds (determined by local governments) should go through formal procurement. The default format is open-bid procurement, during which a government makes the procurement publicly known and open to any qualified suppliers, and interested suppliers compete in a sealed-bid auction. Other procurement formats include invited-bid in which only a small number of supplies are invited to bid, private negotiation in which a government negotiates privately with several suppliers, and assigned procurement in which the procurement is assigned to a particular supplier. By regulation, these other formats are permitted only in exceptional circumstances. This regulatory emphasis on open-bid procurement is intended to increase transparency and reduce corruption.

In 2005, the Chinese government required all vaccines and related products (equipment) and services (storage and transportation) must be procured by the government. Supplies are not allowed to sell their products and services directly to health providers, most of which are public hospitals. The procurement of vaccines is subject to the general regulation of public procurement as described above. However, the decisions of the quantity and format of procurement are decentralized to prefectural Centers for Disease Control and Prevention (CDCs), under the supervision of the corresponding Bureau of Public Health.

In practice, local governments often deviate from the open-bid format for several reasons. A common justification is that open-bid procurement is too slow to satisfy unexpectedly surging demand (e.g., caused by seasonal flu) or too expensive for small quantities of purchase. Another excuse for not using open-bid is that some services, such as storage, are highly customized and can only be adapted to local conditions by certain suppliers. However, deviation from the open-bid default is usually perceived as a way to increase rent-seeking opportunities—invited-bid and private negotiation are opaque and difficult to monitor. Two agencies problems further increase the possibility of corruption. First, to overcome local protectionism, vaccine suppliers often partner with a regional drug distributor to participate in a procurement. Second, given that vaccine procurement requires professional expertise and experience, a local government typically hires an agency to search for information, screen potential suppliers, and arrange the bid. These agents have an informational advantage over the government and may influence a local government's selection of suppliers in their own interest. Less-transparent procurement is more prone to influence activities.

2.3 Monitoring of Policy Oversight

In China, the implementation of national policies is typically decentralized to local governments, and the central government uses a sequence of metrics to measure the achievement of various policy goals (e.g., GDP growth, environmental protection, and food safety) and decide whether to promote a local official (Ang 2016). This approach grants regional governments some degree of flexibility to adapt a national policy to local conditions. But it may lead to oversight of policies that are difficult to measure and distorted policy implementation (Xu 2011).

To solve the problems in policy implementation, the central government has attempted to collect bottom-up information by allowing citizens to file petitions against local officials, urging local governments to build online forums for citizens to express their grievances, and permitting investigative reports on government misconduct from media outlets. However, these information channels, controlled by local governments (Heurlin 2016; Chen et al. 2016; Qin et al. 2018).

A more effective solution relies on top-down monitoring. One important tool is the so-called "inspection and appraisal" mechanism: the higher authorities send inspection teams down to the subordinate bureaus and evaluate their performance (Zhou et al. 2012). The inspection-and-appraisal process can be periodical and involve an overhaul of a wide range of policy goals. It can also be a one-shot game targeting a specific policy goal, often triggered by large-scale incidents (e.g., coalmine or industrial accidents) or wide-spread public sentiment towards certain social problems (e.g., pollution, food safety). Receiving low evaluation in an inspection will put a local official's career at risk. The process incurs a high cost for the upper-level authorities who have to mobilize considerable resources to collect true information and combat local officials' collusive behavior (Zhou 2017). Another cost is that an inspection may severely disrupt a local government's work flow because officials typically spend weeks, or at times months, preparing for an expected inspection (Ai 2011). Therefore, the central government usually selects only a small subset of local governments for inspection and keeps the process secret in the form of a "sudden attack" (Zhou et al. 2012). Despite the low frequency, the high stakes involved in an inspection create substantial top-down pressure on local officials, particularly those at a lower rank.

2.4 Social Media

Among the many social media in China, Sina Weibo, is the most prominent platform for communicating public issues. It was launched in August 2009, expanded rapidly in 2010, gained its popularity the next year, and reached its peak with over 500 million users in 2012. [5] Since 2013, Weibo has lost some ground to WeChat, a cellphone-based social networking service, but has remained an influential platform for public discussion of social issues.

Unlike traditional media, Chinese social media are directly controlled by the National Office of Information Control under the close supervision of the CCP's Central Propaganda Department. Local governments have very limited influence on the censorship of Weibo. [6] China's censorship of social media is known to be extensive but strategic (King et al. 2013, 2014). Qin et al. (2017; 2021) show that a large amount of posts discussing political issues, such as corruption, strikes, and protests, circulated on Weibo for a long period of time. It seems that the Chinese government does not systematically censor discussion about social problems as

[5] The number of Weibo users surged from 63 million at the end of 2010 to 195 million by mid-2011 and reached over 500 million at the end of 2012 (China Internet Network Information Center 2011, 2012, 2013).

[6] There is anecdotal evidence showing that local government officials bribed employees of Sina Weibo to delete unfavorable posts to specific governments and officials. This type of intervention is occasional and can affect only a small number of posts.

long as they do not challenge the CCP's political doctrines or criticize national leaders.

This information-control strategy is particularly true for issues in which the central government and the public have a common interest. Vaccination (more generally, drug and food safety) is such an issue. It is an issue about the livelihood of ordinary people; failure to address it will generate public dissatisfaction with the CCP leadership and distrust of the regime. From a citizen's perspective, the drug-and-food-safety issues are believed to be caused by corrupt local officials and misbehaved firms. Therefore, citizens have an incentive to voice about them, and the central government will not censor.

The monitoring posts (i.e., posts that discuss a social problem and has accountability implications for governments and firms) can affect government behavior in several channels. First, these posts make the consequence of a policy and local officials' misconduct visible, which will diminish local governments' informational advantages over the central government. Second, these posts, aggregated at a certain regional level, can manifest strong public sentiment against specific local governments. Third, from the perspective of repression, public opinion provides information for top leaders to detect discontent and incipient insurrection in a locality. These channels are all likely to trigger top-down inspection that directly affects a local official's career. Therefore, it has become important for local governments to monitor information flows and gauge public opinion on social media. The easiest way that a local government monitors social media information is to have special personnel search posts about a certain issue or generated from a certain locality. Since 2010, local governments have heavily invested in information technology that enables them to collect and analyze social media data. For instance, up to March 2016—the month of the focal event under our study, more than 200 prefectural governments had purchased mass-opinion monitoring software or AI systems that can be applied to social media, particularly, Weibo. [7]

3 Data and Empirical Strategy

3.1 Data

3.1.1 Vaccine-related Data

Vaccine procurement Our main outcome variables concern the procurement of vaccines and related products by a city government. By regulation, all procurement of vaccines and related products must be publicized on government websites. We collected relevant procurement documents from all national and subnational government websites during the sample period (2014-2019).

For each procurement, there are two corresponding documents: an announcement notice and a deal-finalized notice. From the announcement notice, we extract information on the name of the procuring government entity (typically, a local CDC), the specific item to be procured, the date of the announcement, the procurement format, and the procuring agency (if any). From the final notice, we obtained the names of the participating suppliers and the

[7] This number is based on authors' own calculation according to the procurement records published on government websites.

winner. Note that the public announcement serves as a commitment, and it is rare that a government announces one form of procurement format but implements another form in the final deal. Unfortunately, the price information of specific procured item is largely missing. A significant number of final notices do not contain the transaction value of the procurement. Even for those containing such information, many of them only report the value for the entire procurement combining multiple items (vaccines and services). Therefore, we cannot use price differentials for homogeneous items to measure administrative effectiveness.

In total, we obtain 23,734 item-level observations from 2014 to 2019. We classify each item into one of following three categories: Category-I vaccine, Category-II vaccine, and supplemental material, facilities, and services (supplement for short). We collapsed the item-level data into the prefecture-month level. The key variable is the format of procurement: open-bid versus other less-transparent formats.

Table 1 provides the basic summary statistics of the count and share of open-bid procurement at the city-month level. The average number of open-bid for a city in a month is 0.35 because the count of open-bid procurement contains many zeros. The number of observations reduces to 1,915 when we construct the share of open bid procurement because the share measure is undefinable when there is no procurement in a city-month. Conditional on there being procurement, on average, 60% of vaccine procurement takes the open-bid format. Table A1 reports similar summary statistics by each category of procurement. The share of open-bid is slightly lower for the procurement of Category-II and supplements.

3.1.2 Social Media Data

Our key explanatory variables concern social media coverage of vaccine issues. We collected posts referencing "vaccine" in Chinese published on Weibo during 2014-2019. We obtained the data from a third party specializing in collecting social media data in China. To verify the quality of the data, we scraped Weibo posts referencing the exact same keyword for the period of January-June 2014. The two data sources produced a very similar pool of posts, with our data accounting for approximately 95% of the third-party data. We decided to use the third-party data which are more comprehensive.

We believe that social media discussions about vaccine issues, unless they contain anti-regime content, are not censored by the Chinese government. To verify this, we perform a test by searching the word "vaccine" on a website accessible only outside of mainland China, which gathers Weibo posts that appear on the official Weibo platform only for a short period and are then deleted. We find no posts referencing vaccine on this webisite until 2021 when the vaccination of COVID-19 became a politically sensitive issue. In contrast, when we search other keywords related to top political leaders, collective action events (strikes and protests), or disasters, we can always find posts published as early as 2012. This evidence, together with the eruption of vaccine posts upon sensitive events, supports our view that the flows of information about vaccine issues are largely free on Chinese social media during our sample period.

In total, there are 3,318,732 posts referencing vaccines in our dataset. For each post, we obtained the content, timing of posting, and user location at the prefecture level. The user-location information is obtained based on the self-reported location by users on their Weibo profiles. To verify the accuracy of this location information, we selected a set of users who permitted Weibo to track their locations and extracted their real-time locations. We found that 93% of the self-reported location matched the real-time locations.

We read a randomly-drawn sample of 1000 posts and found that these posts can be roughly classified into three categories. First, many posts indeed talk about vaccine-safety problems such as suspicion of deficient vaccines, complaints about drug safety, and condemnation of corruption and poor government monitoring. Second, some posts just state that a person receives vaccination, tell friends where a vaccine is available, or inquire which vaccine to take. Third, some posts are purely product description including scientific knowledge and company information. We list several examples of these three types of posts in Appendix B. We will use the first type of posts to measure public opinion on vaccine safety, which are precisely the information that a government wants to monitor and has the potential to influence government behavior. For expositional convenience, we call this type of posts "monitoring posts." We will use the aggregation of all three types of posts as a proxy for a region's general attention to vaccine issues.

We use a supervised machine-learning approach to separate the monitoring posts from the others. Specifically, we hired a team of research assistants to manually label 12,000 Weibo posts, randomly drawn from our dataset. We then constructed a training dataset to train a support vector machine (SVM) to classify each post into one of the two categories: monitoring posts and others. After cross-validation, we applied this SVM classifier to the entire dataset. The detailed procedure of our machine-learning approach can be found in Appendix C. As indicated from the confusion matrix and ROC curve (Figure A11 in Appendix C), the accuracy of our classification is reasonably high. Of course, it is possible that we can improve the accuracy by using more-advanced machine learning techniques. However, one important purpose of our classification exercise is to imitate Chinese governments' capacity of gauging public opinion—according to industrial experts, SVM is the most popular approach used in the design of information monitoring software acquired by local governments.

In addition, we apply sentiment analysis to our vaccine-post dataset because Chinese governments also wish to gauge public sentiment. Again, we use one of the most popular sentiment analysis methods for Chinese language; see Appendix C for details. The bottom three rows of Table 1 report the basic summary statistics for the three measures of Weibo discussion about vaccine issues. The average total number of posts referencing vaccine in a city-month is approximately 194 with a substantial variance and a wide range. The mean of the variable Monitoring-Posts (ML), the number of monitoring posts obtained by the machine learning approach, is only 27, less than 1/7 of the mean of the total number. The average number of the posts with a negative sentiment is approximately 70, more than double the mean of the number of monitoring posts. Nevertheless, the two variables, Monitoring-Posts (ML) and Negative-Sentiment-Posts aggregated at the city level, are strongly correlated, as shown in Figure A12. Throughout this paper, we will use Monitoring-Posts (ML) as our baseline measure for Weibo discussion about vaccine safety. Results obtained from using Negative-Sentiment-Posts in lieu of Monitoring-Posts (ML) remain qualitatively the same. [8]

[8] We also use a keyword search approach to identify the posts with monitoring implications by searching rele-

Figure 1 plots the time-series of the above three Weibo variables at the monthly frequency. Most notably, there are several spikes of vaccine posts, which are consequences of vaccine scandals. These spikes would not exist if there were government censorship of public discussions about vaccine issues. The spikes appear so abrupt that they substantially change the information landscape on social media for a short period of time. We will explore the regional variation induced by such a change to identify the social media effect.

3.1.3 Other Media Data

To help assess the relative importance of other information channels, we supplement the Weibo data with two additional types of media data: newspaper coverage of vaccine issues and WeChat discussion about vaccine issues. We obtained all the newspaper articles mentioning "vaccine" from Wisenews, a Hong-Kong-based provider of digital Chinese newspaper archives. During our sample period, Wisenews contains 87 general-interest newspapers across 31 cities in Mainland China. The other media data are from WeChat, the most popular Chinese social media platform, used primarily for private messaging between friends and within a restricted group. WeChat also homes some popular public accounts, which publish articles accessible to followers. We searched for vaccine-related articles published on 59 most well-known public accounts specializing in healthcare. Unfortunately, readership data of the WeChat articles and regional variations of exposure to them are not available.

3.2 Empirical Strategy

In this section, we first describe the background of the focal event that we exploit to draw causal implications. After discussing the identification assumptions, we specify the econometric estimation. Finally, we discuss the use of an instrumental-variable approach to strengthen identification.

3.2.1 Event Setting

Our identification relies on the eruption of public opinion on social media induced by the leak of a vaccine scandal. On March 18, 2016, an online media outlet (The Paper) revealed that a vaccine distributor from Shandong province had sold a large quantity of defective or expired vaccines in 18 provinces for the past six years. The leak of the news to the public immediately kindled heated discussions about vaccine issues on Weibo. As seen in Figure 2, the number of Weibo posts referencing "vaccine" was steadily at a low level (around 1,000) before the date of the news leak, but jumped drastically to 14,765 on March 18, accumulated to the peak of 75,250 on March 22, and then declined quickly after one week. We also plot the number of WeChat articles discussing vaccine issues over the same time period. Although on a much smaller scale, the dynamics of WeChat articles follows the pattern of Weibo posts, which suggests that the outbreak of social media discussion was triggered by the news leak and unexpected to both the public and government.

vant Chinese words. This approach produces many more monitoring posts than the SVM-based approach, primarily because some keywords are commonly used in a wide range of contexts.

In the middle of April 2016, the central government decided to strengthen the regulation regarding vaccine procurement and circulation. First, the procurement of vaccine-related products was required to be reported on the provincial digital platform. Second, the qualification of vaccine distributor was subject to more strict inspection, and a large number of lower-tier distributors were disqualified. However, local governments still kept the decision rights regarding which vaccine products to procure, the format of procurement, and which agency to hire. Although the new regulation was passed in April, it was not implemented in most provinces until the end of 2016.

Within a few months after the news on the scandal, the central government sent inspection teams to investigate the problems of vaccine safety in several regions. After the investigation, 355 individuals involved in the scandal were reported to be arrested; among them, 64 were civil servants from the public health system who were prosecuted for corruption and illegal business operation.

3.2.2 Identification Strategy

In the above setting, the leak of the news on the vaccine scandal generated a strong and unexpected information shock to each region. This regional variation in information shocks is the source of our identification. In Figure 3, we depict the landscape of public opinion on vaccine safety across China, measured by the number of monitoring posts published by Weibo users in each city, from February to May in 2016. In February, there was little public discussion about vaccine issues across regions except for a few central cities which are hubs of Weibo users. The landscape changed sharply in March: discussion about vaccine issues surged with notable regional variations. In the subsequent two months, discussion on vaccine issues on Weibo faded and reverted to the pre-event situation. Our identification exploits the abrupt change in Weibo discussion on vaccine issues in a city during this period: cities experiencing a more substantial change are regarded as receiving stronger treatment.

The key identification assumption is that, absent the region-specific and vaccine-related information shock induced by the event, a local government would not change its behavior in vaccine procurement. The major threats to this assumption are factors that coincide with the timing of the information shock and also affect a local government's procurement behavior. For example, in a region where people pay more attention to vaccine issues, local governments are likely to respond more strongly to the event even in the absence of Weibo. Another confounding factor is other information channels such as newspaper coverage, which may trigger local governments to respond even without information flow on Weibo. Finally, changes in the local social and economic conditions at the timing of the event may also confound the effect of social media.

To address these identification concerns, we use three approaches. First, in our regressions with monthly observations, in addition to city fixed effects and year-month fixed effects, we control for city-level variables changing at a monthly frequency and province-specific time trends. Second, we use an IV estimator to isolate the impact of Weibo. Third, we conduct a series of falsification tests to exclude a number of prominent potential confounders.

3.2.3 Econometric Specification

We estimate a panel of 208 cities from March 2015 to March 2017 with the following DID econometric specification:

$$y_{it} = \alpha + \beta WeiboShock_i \times Event_t + X_{it}^{'}\gamma + \lambda_i + \eta_t + \epsilon_{it}, \tag{1}$$

where the subscript i indicates a city government and t indicates year-month. The dependent variable y_{it} is the number (in logarithm) or the share of open bids. The key independent variable is the interaction term $WeiboShock_i \times Event_t$. $Event_t$ is a dummy that equals 1 if an observation is in and after the event month (March 2016) and 0 otherwise. It is not identifiable when time (year-month) fixed effects are included. $WeiboShock_i$ is measured by the difference between the per-capita number of monitoring posts published by users in city i in the event month and the average number of posts published three months before the event. Note that this variable captures the intensity of information shock to a city at the timing of the event and is time invariant. It is not identifiable in the presence of city fixed effects. What we aim to identify is β, the coefficient of the interaction term, which captures the effect of the eruption of vaccine-related monitoring information within a short time period.

In the baseline regressions, we include city-fixed effects, λ_i, and year-month fixed effects, η_t as well as $X_{it}^{'}$, which is a set of time-variant city characteristics including population density, GDP per capita, government expenditure, foreign direct investment, the numbers of internet users, mobile phone users, land-line users, and students, the share of college students among all students, and the share of the secondary industry in GDP. [9] Importantly, we include the number of total posts referencing vaccine in each city in a month. This variable captures the time-variant general attention to vaccine issues in a city, and controlling for it helps purify the effect of information eruption on social media. We also control for province-specific time trends. ϵ_{it} is the error term, clustered in two ways at both the city and year-month levels. [10]

3.2.4 Instrumental Variable

To further single out the effect of the Weibo shock in the specific event, we use Weibo penetration (the number of posts per capita) across cities when Weibo just entered the market in 2009 as an instrument for the variable $WeiboShock_i$ in the baseline regression. This IV idea is similar to Kearney and Levine (2015), who use a broader measure of local area MTV ratings from a pre-period as an IV for the current ratings of a specific MTV reality show (16 and Pregnant) to identify the effect of this particular TV show on teen childbearing. The identifying assumption of this kind of IV in the current setting requires that Weibo penetration in a city in 2009 would be unrelated to vaccine procurement in that city in 2016, except through the outbreak of vaccine posts that have monitoring implications. We next discuss the validity of this assumption.

[9] These city-characteristics variables are available only at the yearly level. Data are obtained from the Chinese City Yearbooks. In the regressions, government expenditure, foreign direct investment, and the numbers of internet users, mobile phone users, land-line users, and students are all weighted by population scale and transformed in logarithm.

[10] Clustering the standard error at the province level barely changes the statistical significance of the coefficients of the main variables of interest.

Weibo was launched by Sina Corporation in August 2009 after China blocked Twitter and Facebook and shut down several domestic micro-blogging services. Initially, Weibo was unknown to most Chinese and only attracted a small number of young and well-educated users who had some experience with micro-blogging. Figure A2 shows that the Weibo users during 2009 to 2010 are mostly young people in their twenties. Therefore, conditional on factors that affected young people's participation in microblogging, such as the numbers of college students and internet users per capita, the early penetration of Weibo into a city is arguably random and unlikely to be related with unobservable regional factors that may affect local governments' procurement of vaccines seven years later. We further verify this conjecture below.

We measure the early Weibo penetration in a city by the per-capita number of Weibo posts (covering all kinds of subjects) published by users in a city in the first three months of Weibo launch (August, September and October of 2009).[11] The left panel of Figure A3 plots the residuals, obtained from regressing our measure of early Weibo penetration on the numbers of college students, internet users, mobile phone users, and land-line users in 2009, against a city's GDP per capita in 2009. The early penetration of Weibo is barely correlated with a city's level of economic development, which is a key determinant of the use of Weibo when it later stabilized. In the right panel of Figure A3, we plot the above residuals in 2009 against the same residuals in 2011 when the use of Weibo experienced the most rapid expansion. The correlation between these two residuals is rather week, supporting our conjecture that the early Weibo penetration, conditional on the four selected variables, is largely random.

The IV-relevance condition requires that the early Weibo penetration is sufficiently correlated with the change in the vaccine monitoring posts associated with the 2016 event. This condition is likely to be satisfied because many early Weibo users, included those accidental users, became opinion leaders on Weibo and actively blogged about social problems. Figure A4 depicts the relationship between the registration time of Weibo users and their total number of posts and followers at the end of 2019. Clearly, users registered in 2009 published many more posts and had substantially more followers than users who registered in later years. We will formally test this IV-relevance condition in the IV estimation.

4 Basic Results

In this section, we focus on the above event-study setting to examine whether the eruption of Weibo information caused local governments' adoption of more-transparent formats for vaccine procurement. We first present the results from the baseline regressions, followed by the results from the IV estimation. Then, we perform a series of robustness checks to verify the basic results. Finally, we report results regarding the effects on other outcomes including local governments' online reaction to vaccine-safety issues, the use of procuring agency, and the supply of vaccines.

[11] The data on the total number of Weibo posts are collected from a database constructed by Weibook. See Qin et al. (2017) for a detailed discussion about the validity of the database.

4.1 Baseline Estimation

Table 2 reports the results from the estimation of specification (1). In Panel A, the dependent variable is the log count of open-bid (or nonopen-bid) procurement. Columns (1) and (2) show that the eruption of monitoring posts had a significant effect on local governments' adoption of open-bid procurement while virtually has no impact on the nonopen-bid procurement. Specifically, in cities where every citizen posts one more monitoring post in a month, the frequency of open-bid will increase by 3.3%. Columns (3) and (4) report the results for the procurement of Category-I (compulsory) vaccines, and the last two columns for Category II vaccines and supplements (equipment and services). We combine these last two categories into one outcome variable because local governments have substantial discretion in procuring Category-II vaccines and supplements, and vaccine safety problems involved both of them. Clearly, the positive Weibo effect is driven only by the Category-II vaccines and supplements. This result is sensible because Category-I vaccines were compulsory with little discretion for local governments to exercise.

Panel B is similarly constructed with the dependant variable being the share of open-bid procurement. As noted before, the number of observations is reduced substantially because observations with zero procurement have to be dropped for calculating the share variable. With this data structure, the panel becomes unbalanced at the city level because some cities are observed only before or after the event during our sample period. To assess the potential bias caused by the unbalanced data, we report the estimated effects controlling for province fixed effects (in the odd-numbered columns) and city fixed effects (in the even-numbered columns), respectively. The control of provincial fixed effects is due to the consideration that there are always observations before and after the event at the provincial level. The first column reports that the Weibo shock, as we measure it, increases a local government's share of open-bid procurement by 11.8%. This coefficient is reduced to 8.4% when the province fixed effects are replaced with the prefecture-fixed effects in the regression (see Column 2). Such a selection effect, however, is primarily driven by Category-I vaccines (Columns 3 and 4). For the procurement of Category-II vaccines and supplements—the outcomes of our main interest, the estimated effects are similar regardless of the level of region fixed effects used in the regression (Columns 5 and 6). Therefore, we will focus on the estimates with control for city-level fixed effects. Column (6) shows that in cities where every citizen publishes one more monitoring post in a month, the share of open-bid increases by 14.4%. This is a notable effect given that the mean is 0.65. [12]

Given the consistency between Panel A and Panel B, we will use the share of open-bid as the outcome variable in our regression analysis in the remaining part of the paper, unless otherwise specified.

Dynamic effects Figure 4 plots the dynamics of the DID estimates (estimated β) around

[12] Admittedly, it is tricky to interpret the magnitude of the Weibo effect. First, the eruption of monitoring posts lasted for a short period, and thus the average number of monitoring posts in a city over the event month is small (around 10,000) relative to the population in a city (typically over one million). If calculated based on the number of posts during the event week, the number of posts in a month will be much larger. Second, the effect of Weibo posts is not necessarily linear. It is likely that the marginal effect is disproportionately large when the number of posts is small. Third, the effect of the Weibo shock persists for a long period, and the accumulated effect is sizeable even if the single-period effect is relatively small.

the event using the share of open-bid in the procurement of Category-II vaccines and supplements as a dependent variable. Specifically, we interact $WeiboShock_i$ with a sequence of time dummies indicating each month before, in and after the month of event. For illustrative convenience, we normalize the effect four months before the event to zero. To see a longer period and more-precise estimate of dynamics, we report the dynamic effects with two-month intervals in Table A2. Figure 4 and Table A2 show that there appear no pretends before the event, and the effect starts to kick in one month after the event and remains significant even four months later after the event. Note that such a speedy reaction is plausible because a local government can immediately announce to use a particular procurement format and start the actual procurement months later.

4.2 IV Estimation

Table 3 reports the IV estimation in which the early-stage Weibo penetration is used to instrument the eruption of monitoring posts in the 2016 event. We exploit the aforementioned quasi-randomness of the IV conditional on a set of predetermined conditions. In particular, we extract the residuals from regressing Weibo penetration during August-October 2009 on the numbers of college students, internet users, mobile phone users, and land-line users in 2009. We then use these residuals to instrument "$WeiboShock_i$" in the estimation of equation (1), following a two-stage-least-squares (2SLS) regression. Note that, if there were no city-fixed effects involved, this IV estimator augmented by a pre-stage is equivalent to a standard 2SLS in which "$WeiboShock_i$" is regressed on the early-stage Weibo penetration together with the four city-level characteristics in 2009. The estimation is slightly more complicated in the presence of city-fixed effects. We therefore implement the following two versions of IV estimation.

First, we use province-fixed effects instead of city-fixed effects in specification (1). In this way, we instrument simultaneously "$WeiboShock_i$" and "$WeiboShock_i \times Event_t$" with "$residuals_i$" and "$residuals_i \times Event_t$." Second, we directly estimate specification (1) with city-fixed effects. In this situation, we can only instrument "$WeiboShock_i \times Event_t$" with "$residuals_i \times Event_t$" because the time-invariant "$residuals_i$" is absorbed by the city-fixed effects. These two versions correspond to the two versions of the baseline estimation reported in Panel B of Table 2. [13]

Table 3 represents the results of our IV estimation.[14] Columns 1 and 3 report the results from the first version of IV implementation (province-fixed effects). The IV estimates are almost identical to their counterparts (Columns 1 and 5) in Table 2, The reported Kleibergen-Paap rk Wald F statistic is over 34, which is considerably larger than the threshold for passing the weak-IV test. Columns 2 and 4 report the results from the second version of IV estimation (city-fixed effects). The Kleibergen-Paap rk Wald F statistic now becomes much larger because the inclusion of city-fixed effects strengthens the IV in the first stage. When the outcome variable is the overall procurement of vaccine and related products, the IV estimate is 0.121 (statistically significant at 1%), which is considerably greater than the baseline estimate (0.084

[13]We also implement a manual version of IV estimation, in which we use the residuals from the pre-stage to instrument "$WeiboShock_i$" in a cross-sectional specification in the first stage. The results are virtually the same.

[14]When the outcome variable is Category-II vaccines and supplements, the number of observations in the IV estimation is smaller than that in the baseline estimation. This is because with the inclusion of city fixed effects, observations whose within-city variation is not strong enough drop automatically in statistical software (e.g., Stata).

and marginally significant). When the outcome includes only Category-II vaccines and supplements, the IV estimate is 0.132, very close to its counterpart (0.144) in the baseline estimation; both estimates are statistically significant at 1%. These results suggest that the endogeneity problem in our baseline estimation is not a serious concern.

4.3 Robustness Checks

Although we have verified the exogeneity of the eruption of monitoring information (see the dynamic effects) and isolated the Weibo channel (see the IV estimation), there may be factors associated with the event that systematically bias our estimation of the Weibo effect. In this section, we present evidence that rules out three sets of prominent confounding factors: (1) event effect, (2) policy effect, and (3) other information channels.

4.3.1 Event and policy effects

One potential confounding factor is the event itself: governments aware of the scandal would respond even without being influenced by public opinion; the effect of Weibo information on government responsiveness is merely a coincidence. To address this confounder, we exploit the fact that not all cities are equally affected by the scandal. As discussed before, the perpetrator of the scandal—the vaccine distributor in Shandong—sold defective vaccines in 18 provinces. Therefore, we include in the baseline regression an interaction term between the event timing dummy and a dummy indicating whether a city was located in one of the 18 affected provinces. As seen from Panel A of Table A3, the coefficient of this newly added interaction term is statistically insignificant while the Weibo effect barely changes.

As noted in Section 3.2.1, the central government tightened up the general regulation of vaccine procurement after the event. If cities where local governments implemented this new regulation more strictly happened to be those experiencing a larger Weibo shock, the estimated effects reported in Table 2 would be spurious. To address this concern, we collected information on each province's actual timing of implementing the new policy from the official websites of provincial Food and Drug Administration.[15] In the data, we observe a large heterogeneity in the timing of implementing this new regulation: a few cities followed the central government's policy immediately after its announcement while many others did it only after half of a year. Nevertheless, the timing of policy implementation has no direct effect on the share of open-bid procurement and does not affect the Weibo effect at all, as can be seen from Panel B of Table A3.

4.3.2 Placebo events

Table 4 reports the results from two placebo events. The first event is caused by the death of a woman bit by a dog even after taking the rabies vaccine in July 2017. Although medical experts clarified that the death was unrelated to the vaccine she received, and no government inspection or regulatory changes were imposed, this event triggered notable discussions about

[15]Out of the 31 provinces in mainland China, Ningxia— an inland province with a relatively low level of economic development— didn't announce the exact timing of implementing the new policy. We drop the observations from this province.

vaccine safety on Weibo. If it is the event and the associated policy change, instead of Weibo information, that have an effect on government procurement, we should expect no Weibo effect to occur around the event time window. However, this is not the case as seen in Panel A—local governments responded to Weibo information induced by such a fake scandal as they did in our main event study.

The second placebo event is another vaccine scandal in November 2017, which involved the distribution of faulty vaccines but did not invoke extensive social media discussions because of the distraction of other social events. [16] We interact a dummy for the event timing with the regional variation in Weibo discussion of vaccines in November 2017 to estimate the Weibo effect using the DID specification as in our baseline estimation. Given the event was only four months after the fake scandal (the dog-bite event), we include both events in the same regression. As shown in Panel B, the lukewarm discussion about vaccine safety on Weibo did not have any significant effect on local governments' vaccine procurement.

4.3.3 Other information channels

Another concern is that a local government may respond to information from other sources, such as traditional media, WeChat, and the internet. In Figure 2, we show that WeChat discussion and newspaper coverage did not precede the outbreak of Weibo posts about vaccine safety. Panel C of Table A3 reports the result of a regression that adds two interaction terms to the baseline DID estimation: the interactions of the 2016 event timing with (1) the change in local newspapers' coverage of vaccine issues and (2) the change in the intensity of search of key words related to vaccines on Baidu, the Chinese equivalent to Google. With such inclusion, the Weibo effect remains unchanged (or even stronger), while the effects of the other two sources are insignificant. This result validates the conjecture that Weibo was the primary informational channel that triggered government response in vaccine procurement.

4.4 Other Outcomes

The above evidence consistently shows that, in the absence of monitoring posts on Weibo, it is unlikely to see the observed changes in the share of open-bid procurement after the event. Given this strong Weibo effect on the format of government procurement, we should expect that the Weibo effect extends to other aspects of the vaccine problem. In this subsection, we investigate whether the outbreak of Weibo information following the event has a significant impact on other outcome variables, such as governments' online interactions with citizens, the hiring of procuring agency, and the supply of vaccines.

4.4.1 Government blogging about vaccine issues

Since 2011 when the use of social media became popular, the Chinese central government has encouraged local governments to actively engage in public communication on social me-

[16]On November 3, 2017, the news that Changsheng Bio-technology, an industry leader in vaccine production, had been producing substandard DPT vaccines was released to the public. During the same period, a controversial murder of a Chinese student in Japan a year ago sparkled substantial social media attention because more details were revealed.

dia. It was reported that, in 2012, there were approximately 50,000 Weibo accounts operated by government offices or individual officials (Sina Weibo 2013). Qin et al. (2017) estimate a much larger presence of government-related users. Fearing that public anger may lead to collective action or trigger top-down inspection, local governments have a strong incentive to sooth online sentiment immediately. Therefore, the eruption of public grievances on Weibo is likely to have an impact on local governments' blogging activities on their Weibo accounts. In particular, we expect to see more discussion about vaccine safety and expression of improving government accountability on the Weibo accounts run by governments in regions experiencing a stronger Weibo shock.

To investigate local governments' blogging activities, we explore the content of governments posts referencing vaccine. We first use a machine-learning approach to identify the posts published by Weibo accounts that represent governments.[17] In addition, we apply the Latent Dirichlet Allocation (LDA) topic modelling approach to the machine-identified government posts. Regardless of the number of topics set in the topic modelling, two topics stand out as closely related to vaccine safety and public accountability while others are mostly related to governments' routine work of disease prevention and regular vaccination. This can be clearly seen in the five-topic model reported in the Appendix C, in which posts that assigned to the first three topics with high probability are referred to as "routine-work posts" while posts assigned to the last two topics with high probability are referred to as "accountability posts." Figure A6 in the Appendix A illustrates that the eruption of government posts upon the shock of the event is driven primarily by the accountability posts.

To formally examine the effect of the Weibo shock on government posts, we run a regression analogous to the baseline DID estimation, except that the observation is at the daily level during a short time window (Feb-April 2016). We construct two timing dummies, with "2016.3.18-2016.3.21" being a dummy for the 4-day period of 18-21 in March 2016 and "After 2016.3.22" being a dummy for the time period from March 22 and onwards. We made this division because March 22 was the first date when the National FDA made an official announcement regarding the scandal. Thus, government posts published in the first time window reflected a government's expressed attitude before the central government's intervention. Table 5 reports the regression results. In cities where citizens' discussion about vaccine safety on Weibo increased more on the date of the event, local governments blogged more posts about vaccine issues, and the content of the posts shifted from routine work to public accountability. Interestingly, after the National FDA's official announcement, local governments responded to the eruption of monitoring posts by talking more about routine work to be aligned with the National FDA's action.

[17] There are two types of Weibo accounts representing governments: government official accounts and private accounts operated by individual officials. The former can be easily identified from the user profile but the latter can only be judged based on post content. Therefore, we use a supervised machine learning approach to identify government users, with the training data labelled by research assistants according to the user profile and post content. Our estimated presence of government posts is approximately 6% of all posts referencing vaccines. This estimation is comparable to Qin et al. (2017). Details of this classification exercise can be found in the online appendix.

4.4.2 Procuring Agencies

One important player in vaccine procurement is the procuring agency who helps match government entities and suppliers, arrange the bidding, and offer consulting services. Chinese governments hire arm's length agencies, instead of procuring officers within the government, in vaccine procurement for two reasons. First, vaccine procurement requires professional knowledge and experience that is not possessed by procurement officers. Second, using an independent agency is a way of curbing corruption as it reduces the direct involvement of government officials. By regulation, these agencies should meet the required level of registered capital and employ enough professionals in the relevant domains in order to obtain the license for being a public procurement agency. In practice, vaccine procurement agencies are mostly companies offering intermediary, legal, and consulting services in the Chinese medical market.

A local government can choose an agency from the same city, outside the city but within the same province, and outside the province. As seen in Table A1, more than 60% of local governments hire agencies from the same cities. Nevertheless, a significant number of governments choose agencies outside the province. By hiring an outside-province agency, a government may expand its choice set and select more suitable suppliers, which however may lead to the risk of losing control. Upon the shock of the focal event, a government concerned about vaccine safety would favor local agencies to gain better control over the quality of the vaccines to be procured. We test this conjecture by using data on the type (by locality) of agency in each procurement.

The first four columns of Table 6 report the regression results regarding the effect of the Weibo shock induced by the event on the share of each type of agencies using the baseline specification. Consistent with our conjecture, the share of local agencies increases with the magnitude of the Weibo shock, suggesting that local governments shifted from outside-province agencies to inside-province agencies.

4.4.3 Vaccine Supply

Public procurement often exhibits a strong home bias in the sense that governments prioritize domestic supplies over foreign supplies (Best et al. 2019) This is also true in the procurement of vaccines in China. The penetration of foreign vaccines produced by famous multinational pharmaceutical firms is generally low in China even though the demand for high-quality vaccines is high. Since vaccine safety is an issue pertinent to domestic manufacturers and service providers, both the demand for and supply (through public procurement) of foreign vaccines would increase in regions experiencing a larger informational shock induced by the event. As shown in the last two columns of Table 6, the estimated effect of the eruption of Weibo posts on the share of foreign vaccines is positive and statistically significant. It may appear puzzling that the reduced home bias is driven by Category-I vaccines instead of Category-II vaccines and supplements. This is because some Category-II vaccines are substitutes for Category-I vaccines and the vaccine-quality problem can be partially solved by procuring more foreign Category-I vaccines.

We also look into the sales of foreign vaccines relative to that of domestic vaccines. Unfortunately, sales data at the regional level are not available. We approximate the sales of a specific vaccine by the quantity of that vaccine approved by health authorities, which is obtained from a large volume of vaccine issuing records. We find that, at the regional (broader than city) level, the share of Category-II vaccines produced by foreign manufacturers that were approved for sales increased after the event, and the increase is larger in regions where the Weibo discussion about vaccine safety following the event was more intense.

5 Further Evidence on Mechanisms

We have shown a body of evidence consistently pointing to the responsiveness of Chinese local governments, in terms of policy implementation in vaccine procurement, upon an information shock caused by the eruption of Weibo posts on vaccine safety. This result rejects the view that, in an authoritarian regime in which strict political control renders the media incapable of influencing government behavior and public accountability. Instead, our empirical findings are consistent with the following two arguments regarding the role of the media in autocracies.

The first is a "benevolent-dictator" argument which posits that local governments, either internalizing the regime's goal or obeying the top leader's command, will respond to public opinion that draws the regime's attention to a social problem or a policy issue. This argument provides a justification for the efficacy of authoritarian ruling: the media in an autocracy plays a similar watchdog role as the media in a democracy, in which a politician responds to public opinion to the extent that voters' preferences are respected. Such an argument corresponds to the so-called Mass Line doctrine that has been advocated by the CCP, in which governments are mobilized to address the concerns of the mass (people) to safeguard regime legitimacy (Dickson 2016).

The second is a "strategic-dictator" argument which contends that, even in an authoritarian regime, some free information flow on the media can help the top leader monitor local officials and improve policy compliance. This argument applies to the current setting because the issue of vaccine safety belongs to the area where the regime and the public share a common interest and public discussion about the issue is not contained. Under this argument, the media effect is limited by the incentive and efficacy of top-down monitoring as part of the authoritarian political process. As described in Section 2.3, top-down monitoring in China, in the form of inspection and appraisal, often involves scrutinizing the implementation of a policy in all aspects, and the format of procurement is a focal point because it is easy to verify.

In this section, we provide evidence to discriminate between these two competing arguments. Such analysis will shed light on the political process that drives the media effect and also help assess the generalizability of our empirical findings. To this end, we first formulate several hypotheses that corroborate one argument but contradict the other. We then test these hypotheses in the same setting as in Section 3.2.1. Finally, we present some evidence beyond the focal event-study setting.

5.1 Hypotheses

In our previous analysis, we show that a local event generates a wave of information shocks through social media, which causes far-reaching government responses across China. The cross-city distribution of such a media effect is rather different under the benevolent-dictator argument and the top-down-monitoring argument. The key distinction lies in whether a government responds directly to citizens' needs or to the perceived top-down pressure. A benevolent government will respond to citizens' needs, and the media informs the government of them. In contrast, a government under the pressure of being monitored responds to public opinion to the extent that public opinion will trigger top-down inspection and that the inspection will impede local officials' career advancement.

In the current setting, a benevolent government will be more responsive when citizens are more concerned about vaccine safety, i.e., the demand for high quality vaccines is greater. As noted in Section 2.1, Category-II vaccines are expensive and paid by consumers themselves; their coverage concentrates on well-developed urban areas, particularly central cities. Therefore, according to the benevolent-dictator argument, the media effect should be stronger in metropolitan areas than in regular cities. Conversely, the top-down-monitoring argument predicts the opposite: the task of policy compliance in vaccine procurement is relatively unimportant for a metropolitan government, which carries out a much wider range of tasks than governments in smaller cities.

Another dimension of city characteristics also helps distinguish the two arguments. Within the Chinese government hierarchy, cities leaders are ranked depending on the administrative level of cities who governments are under their leadership. For instance, Beijing, Shanghai, Tianjin, and Chongqing are four special municipalities directly under the supervision of the central government; they are ranked above provincial capital cities, some of which are in turn ranked above other cities. While more likely to internalize the central government's political goal, a highly-ranked government is less likely to be inspected as it faces fewer upper-lever governments. We frame the following hypotheses in favor of the top-down-monitoring argument.

Hypothesis 1a ***(City size)*** *Relative to governments of other cities, governments of metropolitan cities are less likely to increase the transparency of vaccine procurement in response to public opinion voiced on social media.*

Hypothesis 1b ***(City rank)*** *Relative to lower-ranked governments, higher-ranked governments are less likely to increase the transparency of vaccine procurement in response to public opinion voiced on social media.*

In China, a political leader's career advancement depends on both his ability and loyalty, with the former being manifested in a sequence of performance measures and the latter being subjectively evaluated by the central government (Li and Zhou 2005; Jia et al. 2015; Ang 2016). The importance of the loyalty component increases with a politician's career. In the earlier stage of career, a politician is faced with greater pressure of top-down policy inspection because he has not yet accumulated enough performance measures to signal his ability. This is particularly true for politicians in the first term of his tenure at a new position. In contrast, a politician towards the end of his career or tenure is more likely to internalize the regime's political goal as a way to exhibit his loyalty. We frame another hypothesis favoring the top-down-monitoring argument.

Hypothesis 2 ***(Career concerns)*** *In cities in which political leaders are younger and in the earlier term of tenure, local governments are more likely to increase the transparency of vaccine procurement in response to public opinion voiced on social media.*

Another way to disentangle the two arguments concerns the information asymmetry between the central and local governments. The agency problem within the government hierarchy is largely due to lower-level governments' possession of private information about local conditions. The monitoring role of social media stems from its ability to generate publicly visible information which mitigates the central government's informational disadvantages. Such a monitoring effect is likely to be stronger in a region where preexisting public information is more scarce, because, upon an unexpected information shock, the local government in this region will perceive sharper loss of its informational advantage and a greater probability of top-down inspection. The benevolent-dictator argument does not have such implications. Therefore, we formulate the following hypothesis to be tested.

Hypothesis 3 ***(Information asymmetry)*** *In regions where public information about social problems is generally more scarce, local governments are more likely to increase the transparency of vaccine procurement in response to public opinion voiced on social media.*

5.2 Evidence from event study

We test the above hypotheses in the same event-study setting as in Section 4. In particular, we extend the baseline DID regression (1) to a triple-differences model to estimate a sequence of heterogeneous treatment effects, specified as follows.

$$\begin{aligned} y_{it} = \alpha &+ \theta WeiboShock_i \times Event_t \times Condition_i \\ &+ \beta_1 WeiboShock_i \times Event_t + \beta_2 Condition_i \times Event_t + X'_{it}\gamma + \lambda_i + \eta_t + \epsilon_{it}. \end{aligned} \tag{2}$$

Here, we introduce a new variable $Condition_i$, indicating the predetermined condition related to a city that is used to test a particular hypothesis. For instance, in testing Hypothesis 1a, $Condition_i$ is a dummy to indicate whether a city is in a metropolitan area or not; in testing Hypothesis 1b, $Condition_i$ measures the age or tenure of a political leader in city i. All the other variables are the same as previously defined. The coefficient of the triple-interaction term, θ, is the heterogeneous treatment effect that we are interested in. We include the same set of control variables as in the DID specification and cluster the standard errors in two ways (city and year-month) as before. Table 7 reports the regression results.

Panel A presents the results aimed at testing Hypotheses 1a and 1b. We define a dummy variable, *Metropolitan*, which equals one if a city belongs to the top two tiers in the 5-tier hierarchical classification of Chinese cities.[18] In total, there are 52 cities in the first two tiers.

[18]Endorsed by the Chinese government, the China Business Network(CBN) classifies 337 Chinese cities into five tiers according to the level of economic development, the degree of urbanization and agglomeration, and population. We use the fifth edition of CBN's classification published in 2020.

The average population of these cities is comparable to that of the metropolitan cities world wide, and the GDP per capita of these cities is substantially larger than that of other cities in China. Therefore, it is reasonable to assume that the demand for high-quality vaccines is greater in these cities. We define another dummy variable, $High\ Rank$, to indicate whether a city is among the four special municipalities and 15 sub-provincial cities.[19] The political leaders (Party Secretary and City Chief) of these 19 cities are ranked one or two levels above leaders in regular prefectural cities according to the official administrative level in Chinese government hierarchy. These higher-ranked governments enjoy considerably more autonomy in policy implementation and legislation and thus, are subject to less pressure of top-down intervention than the lower-ranked governments. As reported in Columns (1) and (3), although governments of metropolitan cities responded to the event more than other governments, this responsiveness to the event does not weaken the Weibo effect. The negative and significant coefficient of the triple-interaction term demonstrates that, compared to governments of the regular cities, governments of metropolitan cities were less responsive to the Weibo shock. In fact, the Weibo effect is entirely driven by the regular prefectural governments; metropolitan governments did not respond to the Weibo shock at all, in view of the magnitude of the coefficients and a formal F-test of the sum of the main Weibo effect and the triple-differences effect. Similarly, Colunms (2) and (4) show that relative to lower-ranked cities, higher-ranked cities respond negatively to the Weibo shock. These findings provide support of Hypotheses 1a and 1b, rejecting the benevolent-dictator argument that local governments respond to public opinion by directly internalizing the regime's political goal and thus the concerns of citizens.

Panel B presents the results aimed at testing Hypothesis 2. We define two variables to measure a local political leader's career concerns based on detailed personnel information extracted from the CVs of the mayors of the cities in our sample.[20] The first variable is a binary indicator of a mayor approaching retirement, defined as $Pre-retirement$. According to the promotion rule, a prefectural leader above the cutoff age of 55 will typically serve for one more term before reaching the retirement age of 60 and is thus unlikely to receive further promotion.[21] This age cutoff is also close to the median age (53) of the mayors in our sample. The other variable, $First\ Term\ of\ Tenure$, is an indicator of whether 2016 is within a city mayor's first term. As shown in Columns 3 and 4, the Weibo effect on the procurement of Category-II vaccines and supplements is weaker for local governments headed by leaders above 55 and stronger in their first term of tenure. Such a differential effect is less obvious when the procurement of the compulsory Category-I vaccines is included in the regression (Columns 1 and 2), likely because of limited discretion and rent-seeking opportunities in the procurement of Category-I vaccines. These results confirm Hypothesis 2, supporting the argument that local governments respond to public opinion because of officials' fear of being inspected by upper-level governments.

[19]The 15 sub-provincial cities are Ha'erbin, Changchun, Shenyang, Jinan, Nanjing, Hangzhou, Guangzhou, Wuhan, Chengdu, Xi'an, Dalian, Qingdao, Ningbo, Xiamen, and Shenzhen.

[20]We choose city mayors, instead of party secretary, as political leaders in our analysis because the responsibility of supervising public procurement and vaccine safety is taken primarily by the mayor.

[21]This retirement rule, however, does not apply to the aforementioned four municipalities under the direct control of the central government. Political leaders in these cities have a much longer career. We thus delete these four municipalities from our sample in the regressions.

Panel C is intended to test Hypothesis 3. We use the share of Weibo posts being deleted in a province as a proxy for the degree of informational control faced by the cities in this province. This variable is collected from Bamman et al. (2012), who tracked the deletion of a massive number of Weibo posts across provinces in 2011. Qin et al. (2017) show that this provincial-level measure of censorship is highly correlated with the long-term pro-government bias of local newspapers and the presence of propaganda posts on Weibo between 2009-2013. Thus, more extensive post deletion implies a more strictly controlled media environment and, likely, a greater degree of information asymmetry between the local and central governments. For analytical simplicity, we define a dummy variable, *Delete Post above Mean*, equal to one if the share of posts being deleted is above the mean level across all provinces. As seen in Column 2, the positive and highly significant coefficient of the triple-interaction term shows that the Weibo effect is stronger for governments in a more-controlled informational environment. This finding squares with Hypothesis 3.

5.3 Another event study

From Figure 1, there is another wave of information outbreak on Weibo in July 2018. The information eruption was triggered by a vaccine scandal, in which a public vaccine manufacturer in Jilin province (Changsheng Bio-Technology Co.) were found to provide fake information about its production of a large quantity of rabies vaccine. This event was subsequent to an earlier scandal involving the same firm (described in Footnote 15). The event was reported to the top leaders, who prescribed a thorough investigation of it. Eventually, the firm was delisted and 15 senior managers were sentenced to jail. The event was heatedly discussed on Weibo and also extensively covered by traditional media including the CCP mouthpieces (People's Daily, Xinhua News). Unlike the media coverage of the 2016 event, public opinion this time focused mostly on the issue of firm accountability. This can be clearly seen in Figure A5, in which we show the word clouds of the hot topics generated from the Weibo posts referencing "vaccine", containing monitoring implications and published during the event window. In a situation where governments are directly accountable to the public, a government may take measures to improve firm accountability. In contrast, a government whose behavior is primarily driven by top-down pressure has weak incentive to respond to public opinion that is not aimed at governments and would not provoke intervention from upper-level governments. Perhaps, an exception is those governments in regions where vaccine manufacturers similar to the one involved in the scandal are situated.

We apply the event-study approach used in the previous sections to the event in 2018, with the Weibo shock being defined as the difference between the per-capita number of monitoring posts published by users in a city in the event month (July 2018) and the average number of posts published three months before the event. Table 8 presents the regression results using the baseline specification (1). We choose the time window to be between November 2017 and July 2019. The starting month is not extended to an earlier period to avoid overlapping with a precedent information outbreak (recall the placebo event discussed in Section 4.2.2). The ending month is chosen for practical reasons to ensure enough observations in the panel regression with city fixed effects. Columns (1) and (3) report a positive coefficient regarding the effect of the Weibo shock on the share of open-bid procurement. However, the standard errors are substantial, rendering the coefficient statistically insignificant. Columns (2) and (4) show different government responses in regions with Category-II vaccine manufacturers and those without. This differential effect is large and statistically significant for the procurement of category-II vaccines and supplements. The overall finding suggests that public opinion diffusing over social media helps the effect of an event on government responsiveness break geographical boundaries to the extent that local governments face similar problems. This is consistent with the top-down-pressure mechanism revealed in the 2016 event study.

5.4 Evidence beyond event study

So far, we have exploited information eruption triggered by sudden events to examine the effect of Weibo on government procurement of vaccines. Does the Weibo effect go beyond these extraordinary information shocks? Figure A8 plots the time series of the share of open-bid procurement of vaccines during our entire sample period. The considerable fluctuations suggest that local governments do not institutionalize their vaccine-procurement practices despite their increased policy compliance during the periods of information outbreak. Rather, they sometimes comply with the procurement regulation more strictly while deviating from it at other times. The event-study results suggest that public opinions on vaccine issues voiced on social media can constitute some pressures on local governments' behavior in vaccine procurement. On the other hand, the effect of social media may be absent given the weak voice of public opinion in the regular time. We investigate this inquiry outside of the two event time window.

Table 9 reports the results from OLS regressions of the per-capita numbers of Weibo posts (referencing vaccine and containing negative sentiment) in the past three months, respectively, on the current period of vaccine procurement from January 2014 to December 2019, excluding the Weibo-information outbreak periods of March-August 2016 and July-December 2018. The regressions control for prefecture and time fixed effects as well as time-variant prefectural characteristics. The first three columns show a modest, albeit statistically significant, correlation between Weibo discussions about vaccine issues in a city two months ago and local governments' adoption of the open-bid format for vaccine procurement.

Given our conjecture that, for Weibo to have an effect on government behavior, a government needs to be able to perceive the information shock on social media and is sufficiently sensitive to public opinion, we split the sample based on a predetermined condition—cities which had announced procurement of information-monitoring technology by 2012 and those had not. As discussed in Section 2.4, Chinese local governments invested heavily in information technology to gauge grassroots information on social media. Thus, earlier investment in such technology indicates that a government has greater capacity to monitor information flow on Weibo. We choose whether a government had made procurement announcements by the end of 2012 to split the sample because it usually took at least six months for the procurement to be completed and it took some time for a local government to utilize information technology. The middle two columns of Table 9 show that the positive correlation between the lagged Weibo information flow and the share of open-bid procurement only appears in the sample of cities that had adopted information-monitoring technology before the sample period.

Furthermore, to examine the the role of local governments' sensitivity to social media information, we split the sample by the intensity of strikes in a city before 2012. In China where collective action is regarded as a potential threat to the regime, how to handle social unrest is an important political task faced by local governments. Qin et al. (2021) show that information flows on Weibo increase both the incidence and geographical spread of strikes across Chinese cities. Therefore, in a city with appearance of strikes, the local government is likely to be more sensitive to public opinion on social media. The last two columns of Table 9 confirm this conjecture.

Although having no causal implications, the above correlations suggest that local governments are generally responsive to public opinion voiced on social media. However, they also send an alarming message regarding the effect of social media on government accountability. If an outbreak of Weibo posts had a persistent effect on local governments' compliance with the central government's policy, the effect of social media information on governments' policy compliance would be rather limited after the first information outbreak. The persistent correlations implied in Table 9 suggest that local governments' strategic response: using a more-transparent procurement format under the pressure of public opinion but switching back to the less-transparent formats when the pressure of public opinion fades. This is consistent with the periodical outbreaks of public grievances about vaccine safety and the recurrence of vaccine scandals in 2018.

6 Conclusion

While enlarging political participation, social media may also enhance the power of authoritarian governments who hold a grip of the new information technology. What are the political impact and policy effect of such a combination between enlarged political participation and enhanced government control? We answer this question in the context of China. Theoretically, social media can help authoritarian leaders' utilize bottom-up information to improve policy implementation and bureaucratic control, particularly in areas where the regime and the public share common interests. Empirically, numerical anecdotal evidence suggests that Chinese governments rapidly respond to social media discussion, whereas there is no lack of accusation of government failure to address citizen grievances furiously voiced on social media. Our study discerns the conditions under which public opinion on social media has an impact on government behavior by systematically examining how information outbreaks on social media affects local governments' implementation of vaccine-procurement policies in China.

Our main finding is that, although social media discussions provide only noisy and coarse information, it significantly improves local governments' compliance with the central government's policy that requires more transparent and competitive procedures in public procurement. To the extent that such policy compliance gains better control over misconduct (e.g., hiring more local procuring agencies) and improves vaccine safety (e.g., opening more opportunities to multinational vaccine producers), social media can enhance public accountability in an authoritarian regime. Thus, our finding rejects a popular view positing that, in an authoritarian regime, social media under strict government control is barely consequential for government behavior and public accountability.

Moreover, we show that the top-down political process in an authoritarian system is the key driver of the social media effect. We demonstrate three findings. First, the social media effect is weaker in metropolitan areas and provincial capital cities, where political leaders face less top-down monitoring pressure while the demand for high-quality vaccines is larger. Second, the social media effect is stronger in cities where political leaders are younger or in their early term of tenure. Third, the effect is stronger in regions where the informational environment is generally more strictly controlled and the problem of information asymmetry between the central and local governments is more severe. These findings lend support to a "fear" mechanism: public opinion propagated on social media triggers local officials' fear of top-down inspection which puts their careers at risk and thus forces them to address public grievances and citizens' needs. Consistent with theoretical studies of authoritarian politics (e.g., Xu 2011; Gehlbach et al. 2016; Egorov and Sonin 2020), the insight derived from our empirical results can be generalized to other authoritarian settings.

It may appear that, regardless of the specific mechanism, as long as governments are responsive to the opinion and concerns of citizens within their administrations, the effect of social media in our study is observationally comparable to the effect of a free media in democracies. Such a comparison, however, is delusive. A vote-maximizing politician in a democracy responds to media information that communicates citizens' needs. This response is likely to be localized, but the effect is directed to respect citizens' true preferences and solve real problems, because unsatisfied citizens will continue to produce information that holds politicians accountable. In contrast, a promotion-oriented politician in an autocracy responds to media information that generates the pressure of top-down monitoring. Such a response can go beyond a particular administration. This is precisely the power of social media which spreads local information to a national scale and triggers a wave of responses even in regions where citizens' true concerns are not warranted. The media effect is not only limited by the need of responding to potential top-down intervention but can also increase the risk of policy rigidity. In this study, we show that local governments were responsive to social media sentiment induced by the dog-bite accident. Local governments' overreaction under enormous top-down pressure can lead to the failure of local adaption in policy implementation, as evident in China's extremely strict lockdown policy in the combat against COVID. In this sense, free information flows on social media, even if they do not impose a threat to the regime, are a double-edged sword to an authoritarian regime.

References

Acemoglu, D., T. A. Hassan, and A. Tahoun (2018). The Power of the Street: Evidence from Egypt's Arab Spring. *The Review of Financial Studies 31*(1), 1–42.

Ai, Y. (2011). *Appraisal Inspection and Copying Strategies in the Chinese Bureaucracy: An Organizational Analysis (in Chinese)*. Society.

Allcott, H., L. Braghieri, S. Eichmeyer, and M. Gentzkow (2020). The Welfare Effects of Social Media. *American Economic Review 110*(3), 629–676.

Anderson, S. E., M. T. Buntaine, M. Liu, and B. Zhang (2019). Non-Governmental Monitoring of Local Governments Increases Compliance with Central Mandates: A National-Scale Field Experiment in China. *American Journal of Political Science 63*(3), 626–643.

Ang, Y. Y. (2016). *How China Escaped the Poverty Trap, Ithaca*. NY: Cornell University Press.

Bakshy, E., S. Messing, and L. A. Adamic (2015). Exposure to Ideologically Diverse News and Opinion on Facebook. *Science 348*(6239), 1130–1132.

Bamman, D., B. O'Connor, and N. Smith (2012). Censorship and Deletion Practices in Chinese Social Media. *First Monday 17*.

Benkler, Y., R. Faris, and H. Roberts (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford University Press.

Besley, T. and R. Burgess (2002). The Political Economy of Government Responsiveness: Theory and Evidence from India. *Quarterly Journal of Economics 117*(4), 1415–1451.

Besley, T. and S. Dray (2022). The Political Economy of Lockdown: Does Free Media Matter? *CEPR Working Paper*.

Bessone, P., F. Campante, C. Ferraz, and P. C. Souza (2020). Internet, Social Media, and the Behavior of Politicians: Evidence from Facebook in Brazil. *Working Paper*, 68.

Best, M. C., J. Hjort, and D. Szakonyi. Individuals and Organizations as Sources of State Effectiveness, and Consequences for Policy Design. pp. 80.

Bosio, E., S. Djankov, E. Glaeser, and A. Shleifer (2022). Public Procurement in Law and Practice. *American Economic Review 112*(4), 1091–1117.

Buntaine, M., M. Greenstone, G. He, M. Liu, S. Wang, and B. Zhang (2022). Citizen Participation and Government Accountability: National-Scale Experimental Evidence from Pollution Appeals in China. *Working Paper*.

Buntaine, M. T., B. Zhang, and P. Hunnicutt (2021). Citizen Monitoring of Waterways Decreases Pollution in China by Supporting Government Action and Oversight. *Proceedings of the National Academy of Sciences 118*(29).

Bursztyn, L., G. Egorov, R. Enikolopov, and M. Petrova (2019). Social Media and Xenophobia: Evidence from Russia.NBER Work. Pap. 26567. *NBER Working Paper*.

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.

CNNIC (2011). The 28th Statistical Report on Internet Development in China. Technical report, China Internet Network Information Center, Beijing.

CNNIC (2012). The 30th Statistical Report on Internet Development in China. Technical report, China Internet Network Information Center, Beijing.

CNNIC (2013). The 32nd Statistical Report on Internet Development in China. Technical report, China Internet Network Information Center, Beijing.

Dellavigna, S. and M. Gentzkow (2010). Persuasion: Empirical Evidence. *Annual Reviews 2*, 643–669.

Dellavigna, S. and E. La Ferrara (2015). Economic and Social Impacts of the Media. In *Handbook of Media Economics*, Volume 1, pp. 723–768. North-Holland.

Dickson, B. J. (2016). *The Dictator's Dilemma: The Chinese Communist Party's Strategy for Survival*. Oxford University Press.

Edmond, C. (2013). Information Manipulation, Coordination, and Regime Change*. *The Review of Economic Studies 80*(4), 1422–1458.

Egorov, G., S. Guriev, and K. Sonin (2009). Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data. *American Political Science Review 103*(4), 645–668.

Egorov, G. and K. Sonin (2020). The Political Economics of Non-democracy. *Working Paper*.

Enikolopov, R., A. Makarin, and M. Petrova (2020). Social Media and Protest Participation: Evidence From Russia. *Econometrica 88*(4), 1479–1514.

Enikolopov, R., M. Petrova, and K. Sonin (2018). Social Media and Corruption. *American Economic Journal: Applied Economics 10*(1), 150–174.

Ferraz, C. and F. Finan (2008). Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes *. *Quarterly Journal of Economics 123*(2), 703–745.

Frost and Sullivan (2022). Global and China Vaccine Industry Report. Technical report.

Fu, C. (2021). *Vaccine Procurement and Distribution in China (in Chinese)*. Shanghai Science and Technology Press.

Gavazza, A., M. Nardotto, and T. Valletti (2019). Internet and Politics: Evidence from U.K. Local Elections and Local Government Policies. *Review of Economic Studies 86*(2092-135), 52.

Gehlbach, S., K. Sonin, and M. W. Svolik (2016). Formal Models of Nondemocratic Politics. *Annual Review of Political Science 19*(1), 565–584.

Gerardino, M. P., S. Litschig, and D. Pomeranz (2017). Distortion by Audit: Evidence from Public Procurement. *NBER Working Paper*, 23978.

Heurlin, C. (2016). *Responsive Authoritarianism in China*. Cambridge University Press.

Howard, P. (2010). *The Digital Origins of Dictatorship and Democracy: Information Technology and Political Islam*. Oxford University Press.

Jia, R., M. Kudamatsu, and D. Seim (2015). Political Selection in China: The Complementary Roles of Connections and Performance. *Journal of the European Economic Association 13*(4), 631–668.

Kearney, M. S. and P. B. Levine (2015). Media Influences on Social Outcomes: The Impact of MTV's 16 and Pregnant on Teen Childbearing. *American Economic Review 105*(12), 3597–3632.

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.

King, G., J. Pan, and M. E. Roberts (2014). Reverse-engineering Censorship in China: Randomized Experimentation and Participant Observation. *Science 345*(6199), 1251722.

Li, H. and L.-A. Zhou (2005). Political Turnover and Economic Performance: The Incentive Role of Personnel Control in China. *Journal of Public Economics*, 1743–1762.

Lin, H. (2022). *Petition against the Police: Conflicts between Police and Civilians in China, 2003-2012*. Singapore: Palgrave Macmillan.

Lorentzen, P. (2014). China's Strategic Censorship. *American Journal of Political Science 58*(2), 402–414.

Margetts, E., J. Peter, H. Scott, and Y. Taha (2016). *Political Turbulence: How Social Media Shape Collective Action*. Princeton University Press.

Maskin, E., Y. Qian, and C. Xu (2000). Incentives, Information, and Organizational Form. *Review of Economics Studies* (67), 359–378.

Morozov, E. (2011). *The Net Delusion: The Dark Side of Internet Freedom*. New York: PublicAffairs.

Petrova, M., A. Sen, and P. Yildirim (2021). Social Media and Political Contributions: The Impact of New Technology on Political Competition. *Management Science*.

Prat, A. and D. Strömberg (2013). The Political Economy of Mass Media. *Advances in Economics and Econometrics: Tenth World Congress*.

Qian, Y., G. Roland, and C. Xu (2006). Coordination and Experimentation in M-Form and U-Form Organizations. *Journal of Political Economy 114*(2), 366–402.

Qin, B. (2013). Chinese Microblogs and Drug Quality. *Working Paper*.

Qin, B., D. Strömberg, and Y. Wu (2017). Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda. *Journal of Economic Perspectives 31*(1), 117–140.

Qin, B., D. Strömberg, and Y. Wu (2018). Media Bias in China. *American Economic Review 108*(9), 2442–2476.

Qin, B., D. Strömberg, and Y. Wu (2021). Social Media and Collective Action in China. *Working Paper*.

Snyder, J. M. J. and D. Strömberg (2010). Press Coverage and Political Accountability. *Journal of Political Economy 118*(2), 355–408.

Strömberg, D. (2004). Radio's Impact on Public Spending. *The Quarterly Journal of Economics 119*(1), 189–221.

Strömberg, D. (2015a). Media and Politics. *Annual Review of Economics 7*(1), 173–205.

Strömberg, D. (2015b). Press Coverage and Political Accountability: Theory and Evidence. In *Handbook of Media Economics*, Volume 1, pp. 595–622.

Sunstein, C. (2018). *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.

Szucs, F. Discretion and Corruption in Public Procurement. pp. 41.

Xu, C. (2011). The Fundamental Institutions of China's Reforms and Development. *Journal of Economic Literature 49*(4), 1076–1151.

Yanagizawa-Drott, D., G. Russo, M. Petrova, and R. Enikolopov (2021). Echo chambers: Does online network structure affect political polarization? *Working Paper*.

Zhou, X. (2017). *The Institutional Logic of Governance in China: An Organizational Approach (in Chinese)*. SDX and Harvard-Yuncheng Academic Library.

Zhou, X., Y. Ai, and H. Lian (2012). The Limit of Bureaucratic Power in organizations: The case of the Chinese Bureaucracy. In *Research in the Sociology of Organizations*, Volume 34, pp. 81–111. Emerald Group Publishing Limited, Bingley.

Zhuravskaya, E., M. Petrova, and R. Enikolopov (2020). Political Effects of the Internet and Social Media. *Annual Review of Economics 12*(1), 415–438.

1 Figures

Figure 1: Time Trend of Weibo Posts

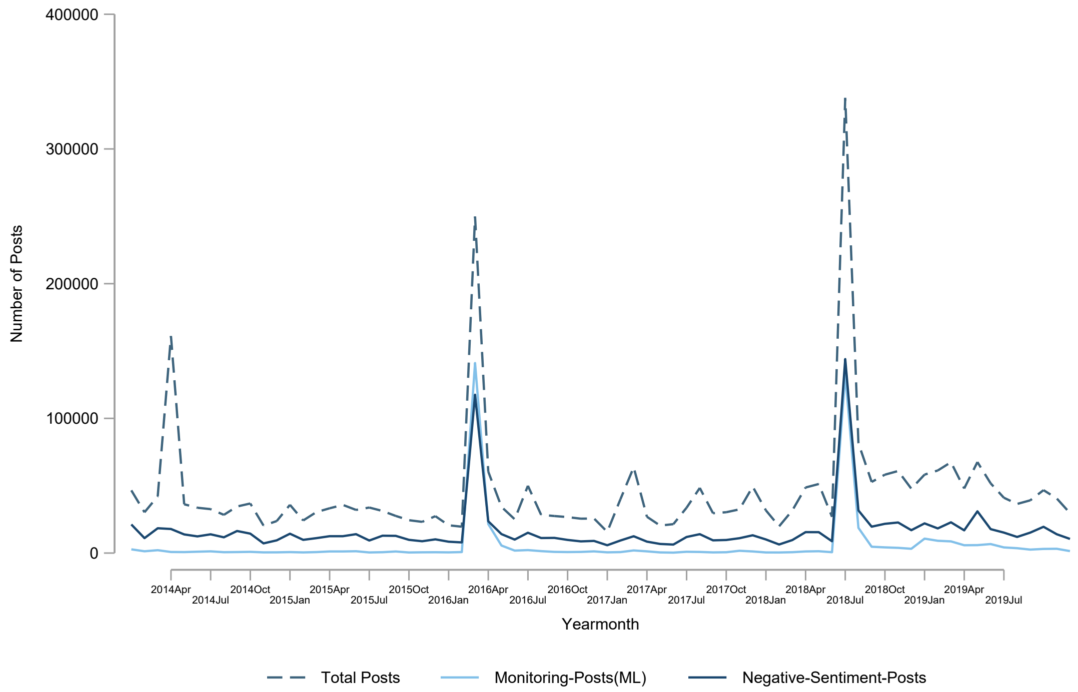

Figure 2: Information Eruption During the 2016 Event

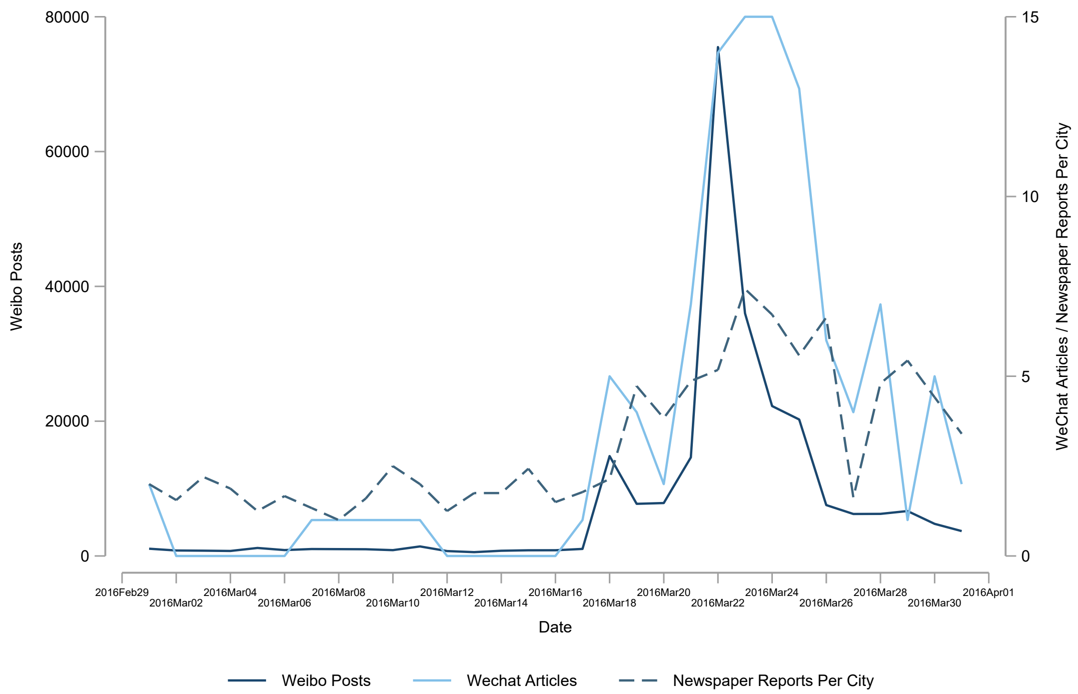

Figure 3: Landscape of Monitoring Posts within 4 Months

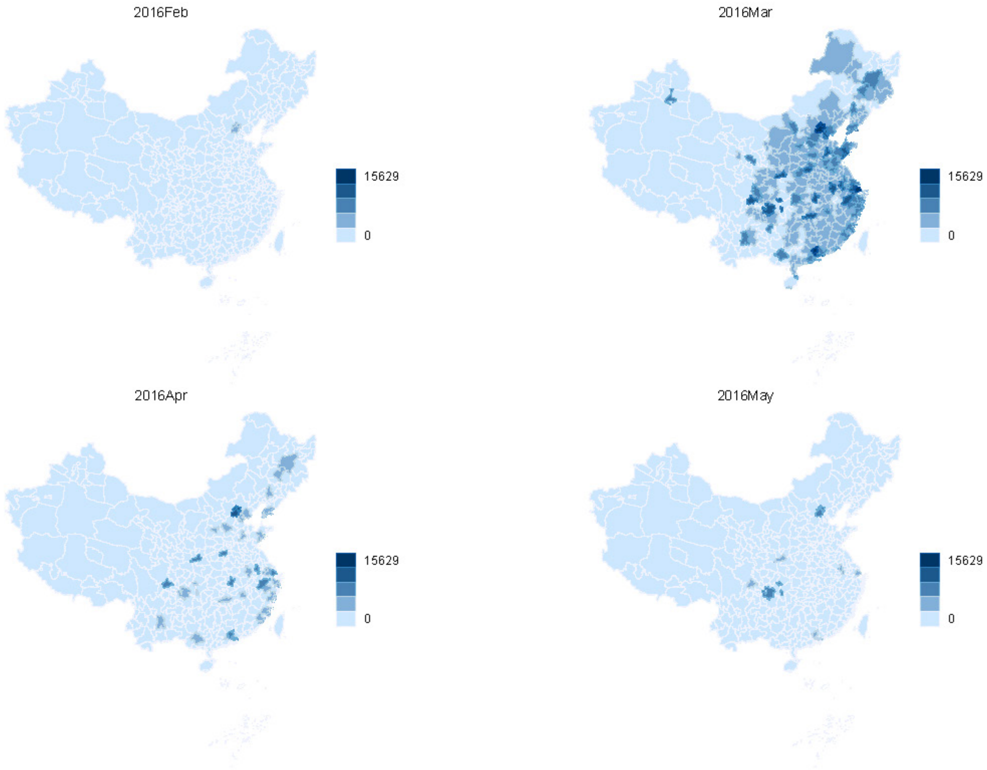

Figure 4: Dynamic Effect of the DID Estimation

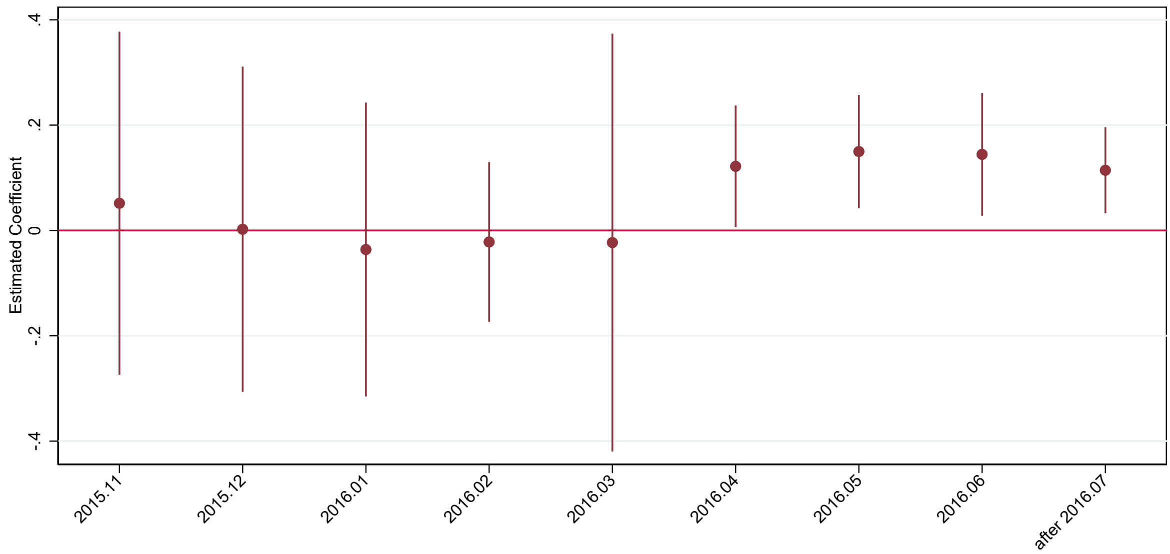

2 Tables

Table 1: Summary Statistics for Main Variables

	count	mean	sd	min	max
Procurement Variables					
Number of procured items	26416	0.48	4.56	0	346
Number of open-bid procurement	26416	0.35	4.14	0	346
Share of open-bid procurement	1915	0.60	0.47	0	1
Weibo Variables					
Total posts	14976	194.18	890.36	2	55422
Monitoring-Posts(ML)	14976	27.11	258.52	0	15629
Negative-Sentiment-Posts	14976	69.75	304.78	0	18138

Note: The unit of observation is city-month. The time frame is from January 2014 to December 2019. "Total posts" is the count of posts referencing the Chinese word of vaccine. "Monitoring-Posts(ML)" is the count of posts referencing vaccine and having monitoring implications identified by a supervised machine learning approach. "Negative-Sentiment-Posts" is the count of posts referencing vaccine and containing a negative sentiment identified by sentiment analysis. "Number of procured items" is the count of total procured items. "Number of open-bid procurement" is the count of open-bid procured items and "Share of open-bid procurement" is the share of open-bid procured items.

Table 2: Baseline Results of DID Estimation

Panel A: Frequency of procured items in different formats

	Overall		Category I		Category II and Supplement	
	(1)	(2)	(3)	(4)	(5)	(6)
	log(open)	log(nonopen)	log(open)	log(nonopen)	log(open)	log(nonopen)
WeiboShock × Event	0.033*	0.003	0.015	0.004	0.026**	−0.003
	(0.018)	(0.009)	(0.014)	(0.006)	(0.012)	(0.011)
Observations	5150	5150	5150	5150	5150	5150
DV Mean	0.097	0.042	0.042	0.015	0.062	0.030
Adjusted R^2	0.126	0.096	0.122	0.096	0.071	0.065
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Prefecture FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Provincial Time Trend	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Share of open-bid procurement

	Overall		Category I		Category II and Supplement	
	(1)	(2)	(3)	(4)	(5)	(6)
	open share	open share	open share	open share	open share	open share
WeiboShock × Event	0.118***	0.084*	0.044	−0.037	0.148***	0.144***
	(0.027)	(0.044)	(0.085)	(0.072)	(0.031)	(0.034)
Observations	431	431	179	179	315	315
DV Mean	0.662	0.662	0.667	0.667	0.652	0.652
Adjusted R^2	0.198	0.222	0.062	0.144	0.187	0.230
Regional FE	Province	Prefecture	Province	Prefecture	Province	Prefecture
Prefectural Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Provincial Time Trend	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from March 2015 to March 2017. "Event" is a dummy that equals 1 if an observation is in and after the event month (March 2016) and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. Controls are time-variant city characteristics including (log) population density, (log) GDP per capita, (log) government expenditure per capita, (log) foreign direct investment per capita, (log) number of internet users per capita, (log) number of mobile phone users per capita, (log) number of land-line users per capita, (log) number of students per capita, the share of college students among all students, and the share of the secondary industry in GDP. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Results of IV Estimation

	Overall		Category II and Supplement	
	(1)	(2)	(3)	(4)
WeiboShock × Event	0.120***	0.121***	0.150***	0.132***
	(0.026)	(0.027)	(0.032)	(0.024)
WeiboShock	−0.064		−0.015	
	(0.052)		(0.047)	
Observations	431	383	313	265
Kleibergen-Paap rk Wald F statistic	34.526	243.336	34.937	370.527
Regional FE	Province	Prefecture	Province	Prefecture
Other Full Baseline Controls	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from March 2015 to March 2017. "Event" is a dummy that equals 1 if an observation is in and after the event month (March 2016) and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. "WeiboShock" and its interaction term with "Event" are instrumented by residuals and its interaction term with "Event" in column 1 and 3, where residuals are obtained from regressing Weibo penetration on (log) numbers of college students per capita,(log) number of internet users per capita, (log) number of mobile phone users per capita, (log) number of land-line users per capita in 2009. The interaction term of "WeiboShock" and "Event" is instrumented by the interaction term of residuals and "Event" in column 2 and 4, where residuals are similarly obtained. Other Full baseline controls include time-variant city characteristics as well as year-month fixed effect, and provincial time trend. The number of mobile phone users is excluded in 2SLS estimation due to concerns on colinearity. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Weibo Effect in the Fake and Silent Scandals

Panel A: Fake Scandal (October 2016-October 2017)

	Overall	Category-II and Suplement
	(1)	(2)
WeiboShock-2017.07 × Event-2017.07	16.484	22.296***
	(12.407)	(6.963)
Observations	268	195
Adjusted R^2	0.208	0.308
Baseline Full Controls	Yes	Yes

Panel B: Silent Scandal (October 2016-June 2018)

	Overall	Category-II and Suplement
	(1)	(2)
WeiboShock-2017.11 × Event-2017.11	−2.207	0.839
	(3.648)	(2.660)
WeiboShock-2017.07 × Event-2017.07	10.556	10.039
	(6.410)	(8.403)
Observations	401	298
Adjusted R^2	0.189	0.263
Baseline Full Controls	Yes	Yes

Note: Observations are at the city-month level. The time window of the regressions in Panel A and Panel B are from October 2016 to October 2017 and from October 2016 to June 2018 respectively. The variable "Event-2017.07" is a dummy that equals 1 if an observation is in and after the event month (July 2017) and 0 otherwise. "Event-2017.11" is similarly defined. "WeiboShock-2017.07" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. "WeiboShock-2017.11" is similarly defined. Full baseline controls include time-variant city characteristics as well as year-month fixed effect, prefecture fixed effect, and provincial time trend. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Weibo Effect on Government Blogging by Topics

		Routine-work Posts			Accoutability Posts	
	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	Topic1	Topic2	Topic3	Topic4	Topic5
WeiboShock(Daily) × 2016.3.18-2016.3.21	3.615***	−0.110	0.146	0.230	2.339***	3.282***
	(0.652)	(0.113)	(0.199)	(0.292)	(0.815)	(0.689)
WeiboShock(Daily) × After 2016.3.22	1.272***	0.261*	0.164*	0.569**	0.308**	1.087***
	(0.347)	(0.136)	(0.092)	(0.253)	(0.120)	(0.303)
Observations	12627	12627	12627	12627	12627	12627
Adjusted R^2	0.704	0.401	0.406	0.476	0.575	0.695
Prefecture FE	Yes	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-date level. The time window of the regression is from February 2016 to April 2016. "WeiboShock(Daily)" is measured by the difference between the per-capita number of monitoring posts published by users in a city on the day 2016.3.18 and the number of posts published one day before. "2016.3.18-2016.3.21" is a dummy that equals 1 if an observation is between 2016.3.18 and 2016.3.21 and 0 otherwise."After 2016.3.22" is a dummy that equals 1 if an observation is on and after 2016.3.22. Control variables include prefecture fixed effect and date fixed effect. Standard errors, clustered in two ways (city and date), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Weibo Effect on Procuring Agency and Winner

	Overall		Category-II and Supplement		Overall	Category-II and Supplement
	(1)	(2)	(3)	(4)	(5)	(6)
	province/all	prefecture/all	province/all	prefecture/all	foreign/all	foreign/all
WeiboShock × Event	0.031	0.052***	0.026	0.067***	0.018**	0.002
	(0.020)	(0.018)	(0.022)	(0.019)	(0.008)	(0.005)
Observations	409	409	313	313	438	302
Adjusted R^2	0.440	0.413	0.390	0.366	0.017	0.052
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from March 2015 to March 2017. "Event" is a dummy that equals 1 if an observation is in and after the event month (March 2016) and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. Full baseline controls include time-variant city characteristics as well as year-month fixed effect, prefecture fixed effect, and provincial time trend. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Evidence on Mechanisms: Heterogeneous Treatment Effects

Panel A: City Size and Rank

	Overall		Category II and Supplement	
	(1)	(2)	(3)	(4)
WeiboShock × Event	0.616***	0.275***	0.392*	0.453***
	(0.141)	(0.093)	(0.222)	(0.133)
Metropolitan × Event	0.372		0.362	
	(0.261)		(0.325)	
WeiboShock × Event × Metropolitan	−0.536***		−0.270	
	(0.157)		(0.218)	
High Rank × Event		0.314		0.092
		(0.292)		(0.331)
WeiboShock × Event × High Rank		−0.219*		−0.289*
		(0.110)		(0.140)
Observations	431	431	315	315
Adjusted R^2	0.232	0.232	0.223	0.250
Full Baseline Controls	Yes	Yes	Yes	Yes

Panel B: Career Concerns

	Overall		Category II and Supplement	
	(1)	(2)	(3)	(4)
WeiboShock × Event	0.097	−0.111	0.204**	−0.128
	(0.089)	(0.078)	(0.076)	(0.139)
Pre-retirement × Event	0.681*		1.271**	
	(0.396)		(0.610)	
WeiboShock × Event × Pre-retirement	−0.212*		−0.409**	
	(0.107)		(0.155)	
First Term of Tenure × Event		−0.305		−0.674**
		(0.273)		(0.271)
WeiboShock × Event × First Term of Tenure		0.204**		0.299**
		(0.090)		(0.127)
Observations	396	396	283	283
Adjusted R^2	0.226	0.226	0.190	0.192
Full Baseline Controls	Yes	Yes	Yes	Yes

Panel C: Information Asymmetry

	Overall	Category II and Supplement
	(1)	(2)
WeiboShock × Event	0.068	0.126***
	(0.048)	(0.036)
Delete Post above Mean × Event	−0.494*	−1.105**
	(0.258)	(0.463)
WeiboShock × Event × Delete Post above Mean	0.277*	0.710**
	(0.158)	(0.288)
Observations	431	315
Adjusted R^2	0.224	0.232
Full Baseline Controls	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from March 2015 to March 2017. "Event" is a dummy that equals 1 if an observation is in and after the event month (March 2016) and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. "Metropolitan" is a dummy the equals 1 if the city belongs to the top two city tiers and 0 otherwise. "High Rank" is a dummy equals 1 if the city belongs to four special municipalities or 15 sub-provincial cities and 0 otherwise. "Pre-retirement" is a dummy that equals 1 if the age of the city mayor is above 55 and 0 otherwise. "First Term of Tenure" is a dummy that equals 1 if the city mayor is in his/her first 4 tenure years and 0 otherwise. "Delete Post above Mean" is a dummy that equals 1 if the number of posts being deleted is above the mean level of all provinces. Full baseline controls include time-variant city characteristics as well as year-month fixed effect, prefecture fixed effect, and provincial time trend. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Weibo Effect during the Vaccine Scandal in 2018

	Overall		Category II and Supplement	
	(1) openshare	(2) openshare	(3) openshare	(4) openshare
WeiboShock × Event	0.010	−0.029	0.011	−0.078
	(0.033)	(0.067)	(0.042)	(0.046)
Category II Company × Event		−0.167		−0.231
		(0.194)		(0.152)
WeiboShock × Event × Category II Company		0.055		0.109**
		(0.065)		(0.049)
Observations	476	476	416	416
Adjusted R^2	0.163	0.160	0.225	0.227
Prefectural Controls	Yes	Yes	Yes	Yes
Prefecture FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Provincial Time Trend	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from November 2017 to July 2019. "Event" is a dummy that equals 1 if an observation is in and after the event month (July 2018) and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. "Category II Company" is a dummy that equals 1 for cities where there existed companies producing or selling Category-II vaccines from July 2017 to June 2018. Controls are time-variant city characteristics including (log) population density, (log) GDP per capita, (log) government expenditure per capita, (log) foreign direct investment per capita, (log) number of internet users per capita, (log) number of mobile phone users per capita, (log) number of land-line users per capita, (log) number of students per capita, the share of college students among all students, and the share of the secondary industry in GDP. Standard errors, clustered in two ways (city and month), are in parentheses.

* p < 0.1, ** p < 0.05, *** p < 0.01.

Table 9: Correlation between Weibo Posts and Government Procurement

	Overall			Adoption of Surveillance Technology before 2012		Appearance of Strikes before 2012	
	(1) log(open)	(2) log(nonopen)	(3) openshare	(4) openshare	(5) openshare	(6) openshare	(7) openshare
L1.#posts pca	−0.053	0.045	−0.035	0.205	0.039	−0.342	0.035
	(0.062)	(0.052)	(0.051)	(0.439)	(0.058)	(0.775)	(0.055)
L2.#posts pca	0.102**	−0.071*	0.271***	−0.214	0.361***	−0.252	0.289***
	(0.043)	(0.041)	(0.042)	(0.266)	(0.061)	(0.694)	(0.057)
L3.#posts pca	−0.019	−0.004	0.012	−0.281	0.028	0.192	−0.059
	(0.085)	(0.042)	(0.087)	(0.358)	(0.100)	(0.963)	(0.107)
Observations	11742	11742	1025	465	560	327	697
Adjusted R^2	0.121	0.078	0.180	0.243	0.114	0.238	0.109
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prefecture FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Provincial Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from January 2014 to December 2019, excluding two scandal periods, from March 2016 to August 2016 and from July 2018 to December 2018. "L1.#posts pca" is one-period lag of the per-capita number of posts referring vaccine and containing a negative sentiment identified by sentiment analysis. "L2.#posts pca" and "L3.#posts pca" are similarly defined but to measure two-period and three-period lags of the per-capita number of posts. "Adoption of Surveillance Technology before 2012" indicates whether a city had announced procurement of information surveillance technology by 2012. "Appearance of Strikes before 2012" indicates whether there had been strikes in a city by 2012. We control for the lagged counts of government procured items (up to three periods). Other control variables include (log) population density, (log) GDP per capita, (log) government expenditure per capita, (log) foreign direct investment per capita, (log) number of internet users per capita, (log) number of mobile phone users per capita, (log) number of land-line users per capita, (log) number of students per capita, the share of college students among all students, and the share of the secondary industry in GDP. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendices

A Additional Empirical Results

Figure A1: Time Trend of Wisenews and WeChat

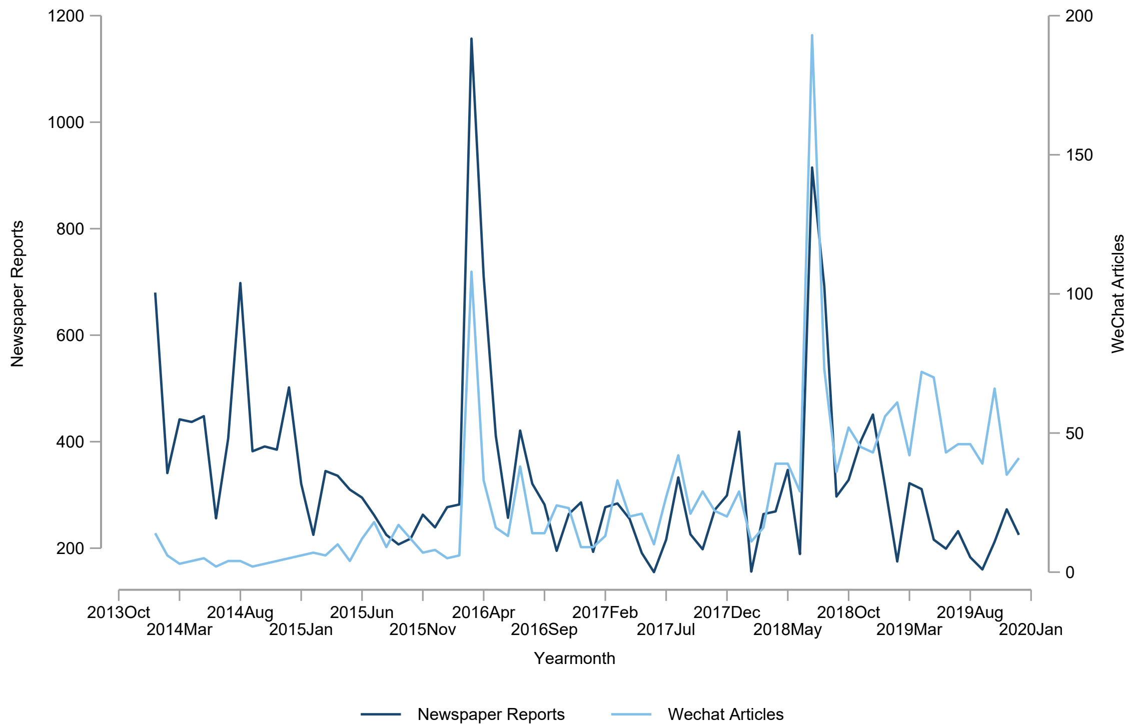

Figure A2: Age Distribution of Early Weibo Users

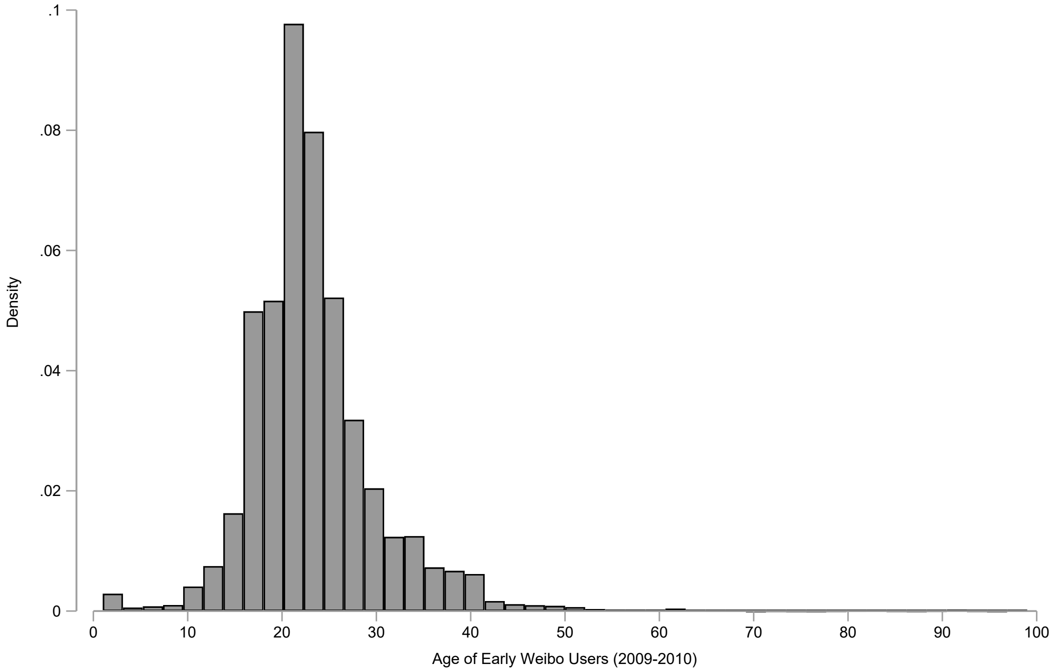

Figure A3: Randomness of Weibo Penetration

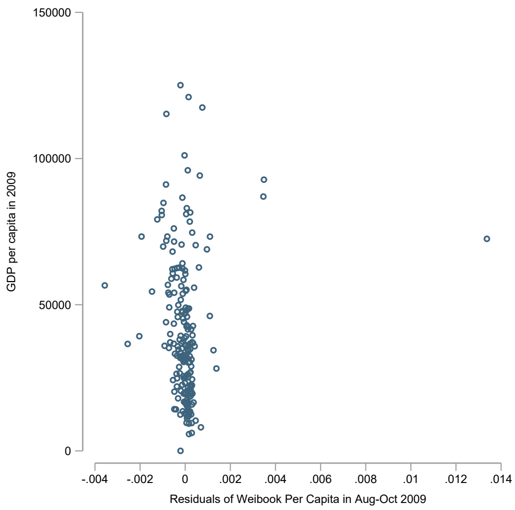

(a)

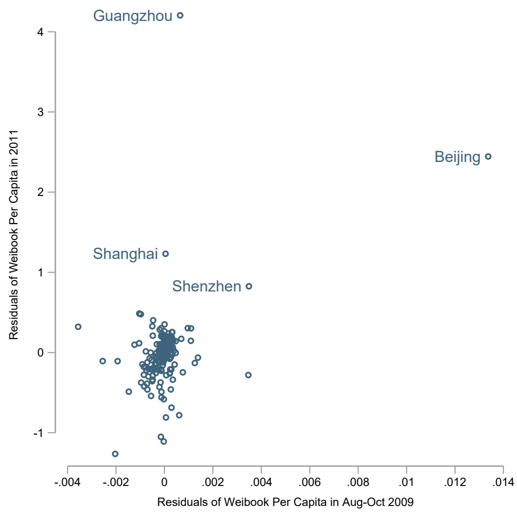

(b)

Figure A4: KOLs and Registration Time

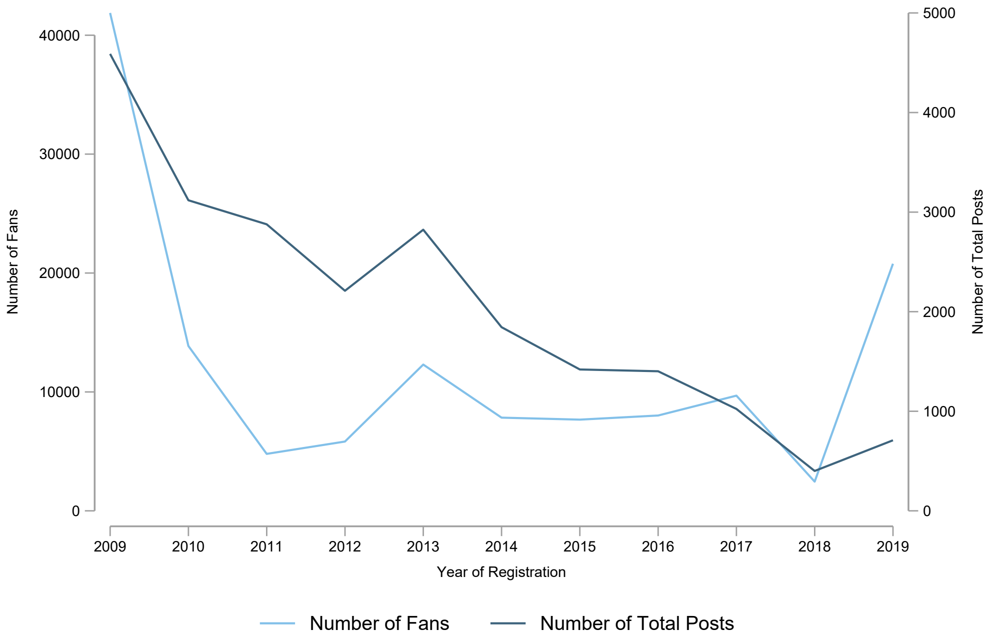

Figure A5: Word Cloud

(a) Government Posts in 2016

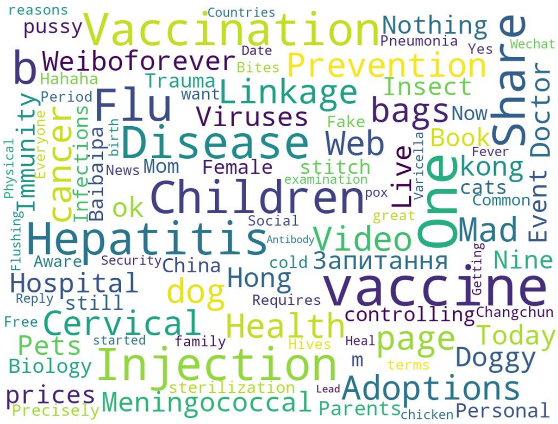

(b) Citizen Posts in 2016

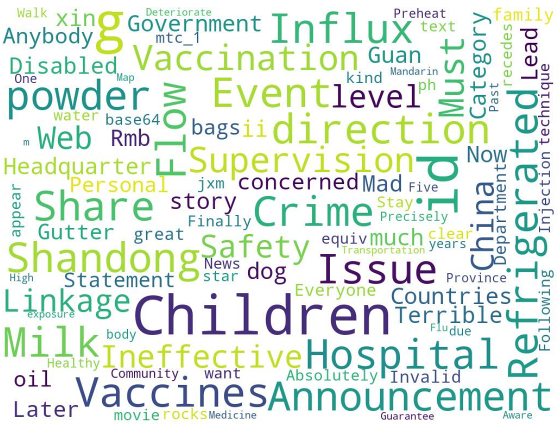

(c) Monitoring Posts in 2016

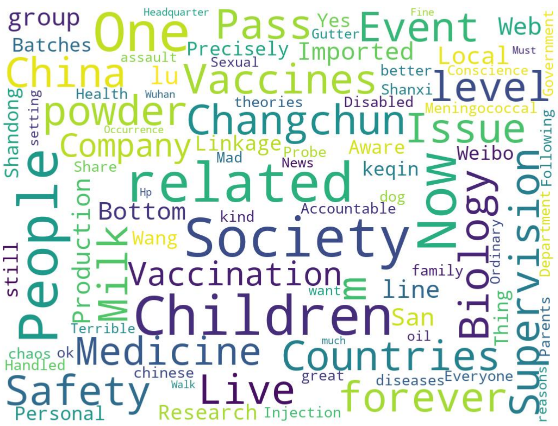

(d) Monitoring Posts in 2018

Figure A6: Time Trend of Posts for Different Topics

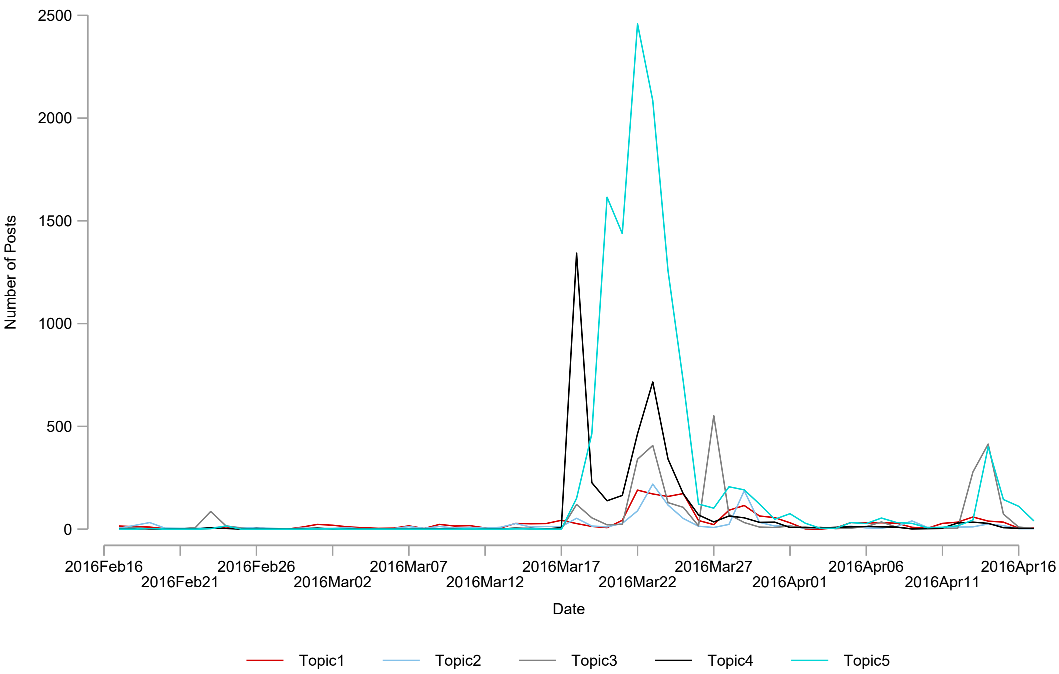

Figure A7: Information Eruption During the 2018 Event

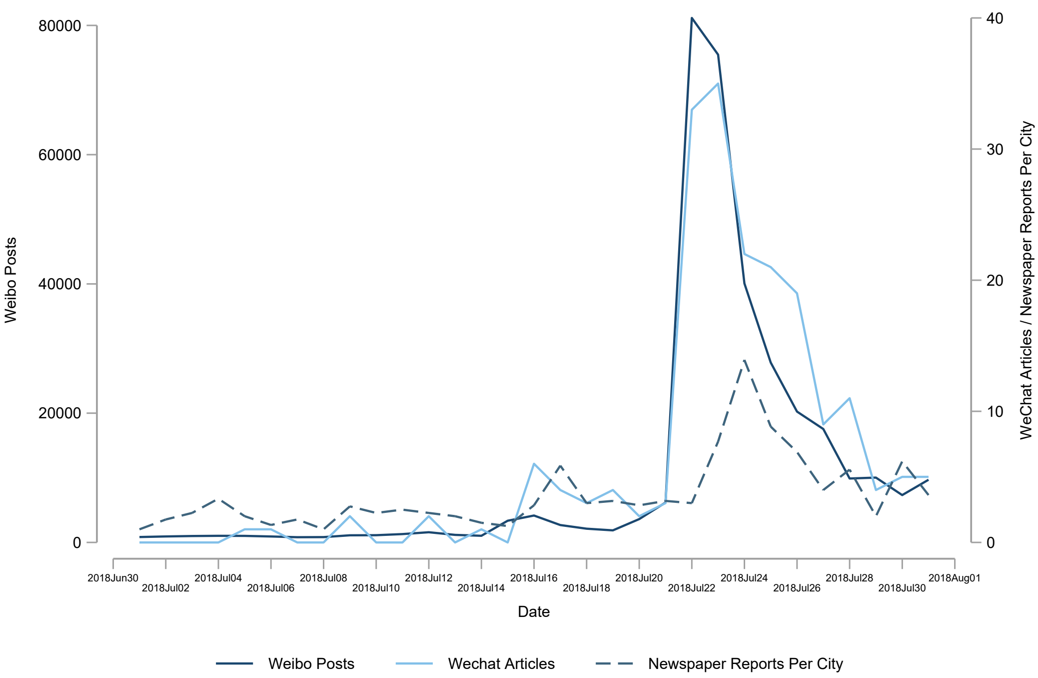

Figure A8: Time Trend of Open Share

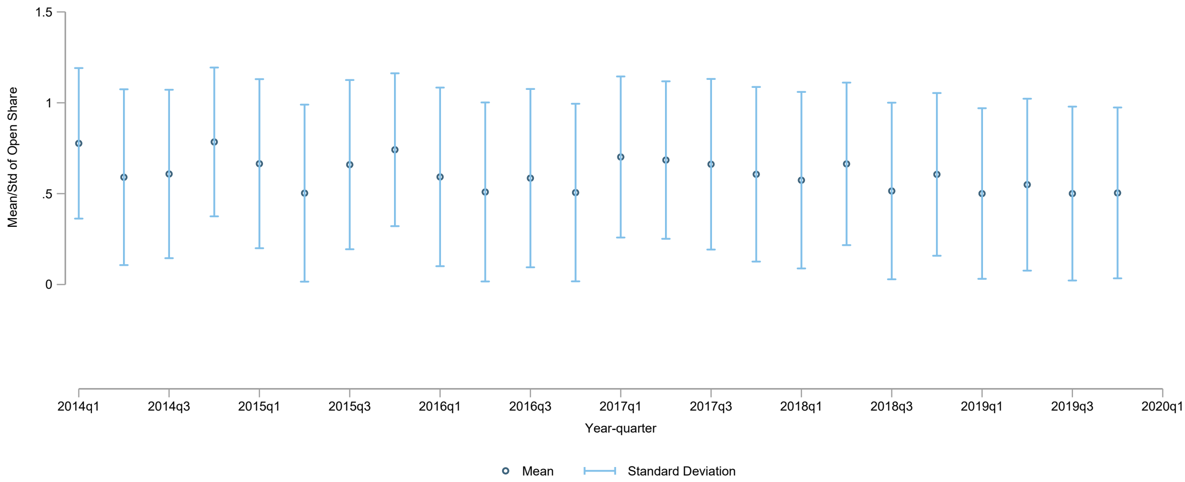

Table A1: Summary Statistics by Category

	count	mean	sd	min	max
Category I					
Number of open-bid procurement	26416	0.13	1.47	0.00	72.00
Share of open-bid procurement	653	0.69	0.45	0.00	1.00
Share of foreign winner	820	0.05	0.18	0.00	1.00
Share of local agency (within-prefecture)	747	0.82	0.38	0.00	1.00
Share of local agency (within-province)	747	0.89	0.30	0.00	1.00
Category II					
Number of open-bid procurement	26416	0.18	3.70	0.00	346.00
Share of open-bid procurement	640	0.65	0.47	0.00	1.00
share of foreign winner	671	0.03	0.10	0.00	1.00
Share of local agency (within-prefecture)	544	0.77	0.42	0.00	1.00
Share of local agency (within-province)	544	0.87	0.34	0.00	1.00
Supplements					
Number of open-bid procurement	26416	0.05	0.75	0.00	70.00
Share of open-bid procurement	985	0.54	0.48	0.00	1.00
Share of foreign winner	1163	0.00	0.03	0.00	1.00
Share of local agency (within-prefecture)	845	0.72	0.44	0.00	1.00
Share of local agency (within-province)	845	0.87	0.33	0.00	1.00

Note: The unit of observation is city-month. The time frame is from January 2014 to December 2019. For each category, "Category I", "Category II" or "Supplements", the listed variables are similarly defined using the corresponding sub-category sample. "Number of open-bid procurement" is the count of open-bid procured items and "Share of open-bid procurement" is the share of open-bid procured items. "Share of foreign winner" is the share of procuring winners that are foreign companies. "Share of local agency (within-prefecture)/(within-province)" is the share of local procuring agencies defined by comparing the province/prefecture of a procuring entity and that of the corresponding agency.

Table A2: Dynamics Effect of DID Estimation

	Overall	Category II and Supplement
	(1)	(2)
WeiboShock × 2015.06-2015.07	−0.016	0.017
	(0.029)	(0.037)
WeiboShock × 2015.08-2015.09	−0.049	−0.041
	(0.033)	(0.025)
WeiboShock × 2015.10-2015.11	−0.037	−0.045
	(0.044)	(0.037)
WeiboShock × 2015.12-2016.01	0.005	0.006
	(0.051)	(0.054)
WeiboShock × 2016.02-2016.03	−0.003	−0.016
	(0.075)	(0.084)
WeiboShock × 2016.04-2016.05	0.067	0.137***
	(0.071)	(0.045)
WeiboShock × 2016.06-2016.07	0.064	0.128***
	(0.058)	(0.039)
WeiboShock × 2016.08-2016.09	0.055	0.091**
	(0.073)	(0.034)
WeiboShock × after 2016.10	0.048	0.096*
	(0.082)	(0.049)
Observations	431	315
Adjusted R^2	0.215	0.207
Full Baseline Controls	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from March 2015 to March 2017. The time period variables are a set of dummy that equals 1 if an observation is within the time period as the name of each time variable indicates and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. Full baseline controls include time-variant city characteristics as well as year-month fixed effect, prefecture fixed effect, and provincial time trend. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Robustness Checks

Panel A: Event Effect

	Overall		Category II and Supplement	
	(1)	(2)	(3)	(4)
WeiboShock × Event		0.087*		0.133***
		(0.042)		(0.040)
Province Involved × Event	0.029	−0.079	0.466*	0.196
	(0.224)	(0.238)	(0.235)	(0.272)
Observations	431	431	315	315
Adjusted R^2	0.208	0.219	0.197	0.226
Full Baseline Controls	Yes	Yes	Yes	Yes

Panel B: Policy Effect

	Overall		Category II and Supplement	
	(1)	(2)	(3)	(4)
WeiboShock × Event		0.081*		0.146***
		(0.044)		(0.038)
WeiboShock × Vaccine Policy	0.021	0.007	0.029	−0.003
	(0.022)	(0.027)	(0.031)	(0.022)
Observations	431	431	315	315
Adjusted R^2	0.209	0.219	0.187	0.224
Full Baseline Controls	Yes	Yes	Yes	Yes

Panel C: Other Information Channel

	Overall		Category-II and Supplements	
	(1)	(2)	(3)	(4)
WeiboShock × Event	0.175***	0.094*	0.167*	0.129***
	(0.051)	(0.054)	(0.086)	(0.042)
Newspaper Shock × Event	−0.006**		−0.001	
	(0.003)		(0.005)	
Search Index Shock × Event		−0.060		0.132
		(0.187)		(0.213)
Observations	431	431	315	315
Adjusted R^2	0.234	0.219	0.224	0.225
Full Baseline Controls	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The time window of the regression is from March 2015 to March 2017. "Event" is a dummy that equals 1 if an observation is in and after the event month (March 2016) and 0 otherwise. "WeiboShock" is measured by the difference between the per-capita number of monitoring posts published by users in a city in the event month and the average number of posts published three months before the event. "Newspaper Shock" is a similar measure of information shock using the number of newspaper reports. "Search Index Shock" is a similar measure of information shock using per capita Baidu search index. "Province Involved" is a dummy that equals 1 if the province is supplied with the problematic vaccines. "Vaccine Policy" is a province-varied dummy that equals 1 if the month is in and after the implementation time of the vaccine policy in each province. Full baseline controls include time-variant city characteristics as well as year-month fixed effect, prefecture fixed effect, and provincial time trend. Standard errors, clustered in two ways (city and month), are in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B Content Description of Weibo Posts

This section provides several examples of Weibo posts in each of the following three categories.

B.1 Vaccine Safety and Public Grievance

Example 1 *"Poor children! Their health is stolen. Those people who sold defective vaccines should go to hell! It is absolutely necessary to solve the problem of vaccine deficiency. How many times has this kind of event happened? These firms were fined before but came back to the market again. How can this happen? What exactly did the regulatory agency do? Shouldn't the government reflect on this? Are they really responsible? Do they feel shameful when paid so much by taxpayers?"*

Example 2 *"It is unbelievable that 500-million-Yuan-worth defective vaccines were circulated in 18 provinces. These criminal activities should be investigated thoroughly! The government should also make it public where these vaccines went, how many remain in the market, which batch of vaccines are problematic. Please forward my post if you still have conscience. We cannot turn a blind eye to this event. Save children!"*

Example 3 *"I beg the government to tell us which community hospitals have used the defective vaccines. We need to know whom have been given these vaccines and how many are still circulated. This information must be transparent! Otherwise, no one can protect the health of our children in the future!"*

B.2 Personal Experience

Example 1 *"After taking encephalitis vaccine, my baby started to have a fever in the evening. Physical cooling did not work. Finally, I fed her with ibuprofen. But her temperature rose again at around 12 o'clock. Is it because my baby did not take enough breast milk to resist the reaction?"*

Example 2 *"I went to hospital today. I was told that there has been no supply of polio sugar pills for three months. The same for chickenpox vaccine. I plan to go to hospital again next Monday."*

Example 3 *"I also want to play with cats and dogs! I just got the rabies vaccine and I don't need to be scared about being bitten within half a year. So happy."*

B.3 Medicare/Product Information

Example 1 *"Yesterday, the First People's Hospital of the city received more than 500 patients every hour. Half of them got a cold or flu. Experts remind that Type A3 influenza is a common seasonal flu. Citizens should take the flu vaccine as soon as possible."*

Example 2 *"Here is a headline article: 'Vaccination: why does induration appear right after the injection? How to deal with it?' Please click this link for more information."*

Example 3 *"Mothers all know that every baby is vaccinated with various vaccines from birth. Which vaccines are required? How to get free vaccines? Moms, please check here for more information."*

C Details on Textual Analysis

In this section, we describe the detailed process of textual analysis, which generates two important inputs for our empirical study: (1) monitoring posts—the Weibo posts that reflect public grievances and have the potential to affect government behavior—and (2) government posts—those published by governments.

C.1 Monitoring Posts

We use two common machine-learning approaches to identify the monitoring posts. An important purpose of our textual analysis is to imitate the function of the massive-information-monitoring software or AI systems acquired by local governments. According to our interviews with industrial experts, sentiment analysis and the support-vector-machine (SVM) type of supervised machine learning technique are the two most popular methods used by social media information-monitoring software.[1]

C.1.1 Sentiment Analysis

We apply sentiment analysis to all Weibo posts referencing "vaccine" during our sample period (2014-2018). Following the common practice, we adopted HowNet, a pre-trained dictionary that is widely used in natural language processing in the Chinese context, to identify sentiment words. [2] We then used an intuitively simple sentiment word-count method for sentiment calculation. The sentiment polarity of an input text is calculated by the count difference between negative words and positive words. A larger sentiment score for a post is associated with stronger negative sentiment.

C.1.2 Classification

We used a binary classification method to separate monitoring posts from other posts. To begin with, we drew a random sample of 12,000 posts from our Weibo data published in 2016—the year of the focal scandal in our study.[3] We read 1000 Weibo posts and labelled a post as 1 if a post contains monitoring content; otherwise we labelled the post as 0. With the experience from this pilot sample, we came up with a detailed instruction on the labeling task as shown in Figure A9. Two research assistants were then hired to independently label the remaining 11,000 Weibo posts in the random sample according to the instruction. Approximately 87% of the posts were assigned the same label by the RAs. For those inconsistent labelling, we tried

[1] We could have used more-advanced techniques, such as neural networks and reinforcement learning, to improve the accuracy of classifying monitoring posts. But these techniques are less widely used by massive-information monitoring software during our sample period.

[2] Other Chinese dictionaries available for sentiment analysis include the Chinese Sentiment Word Weight Table from BoYuan of Tsinghua University, the Chinese Emotional Vocabulary developed by Dalian University of Technology Information Research Laboratory, and the National Taiwan University Sentiment Dictionary.

[3] We restricted our selection of posts to 2016 to avoid the issue of paucity of monitoring posts in the non-scandal period. After all, the variation in the monitoring post we exploit in our empirical analysis comes from the event period within 2016.

to resolve the inconsistencies by careful reading. Eventually, the labelled dataset contains 4633 monitoring posts, accounting for 38.61% of the entire sample.

Next, we divided the labelled dataset into 10 folds, using seven folds as the training data and three folds for testing. The pre-processing of the data involves segmentation and removal of all alphabets, numbers, punctuation, and other meaningless characters. After vectorizing the corpus into a word frequency matrix, we applied various machine-learning algorithms to train the classifier, including Support Vector Machine(SVM), Naïve Bayes, Logistics Regression, Decision Tree, Random Forest, and AdaBoost. As shown in Figure A10, SVM outperforms other models in terms of both recall and precision. This advantage of SVM is also found in other classification tasks (Dumais, Platt, Heckerman, and Sahami 1998; Joachims 1998; Sebastiani 2002) and is particularly true in the classification of social media posts whose length is short.

Therefore, we adopted the SVM as our baseline algorithm to classify the monitoring posts. Figure A11a presents the confusion matrix, and Figure A11b reports the Receiver Operating Characteristic (ROC) curve and the Precision-Recall curve. Both figures show that the SVM classifier produces good fit and a high level of classification accuracy. Finally, we apply the trained SVM classifier to predict the labelling of posts in our entire dataset. These predicted monitoring posts are used to compute the regional variation in the Weibo shock in our empirical analysis. As shown in Figure A12, at the city level, this measure of monitoring posts is highly correlated with the negative-sentiment posts based on our sentiment analysis.

C.2 Government Posts

One straightforward way to identify government posts is through user information. Official government accounts on Weibo are typically operated by users whose profiles clearly indicate their identity such as a certain government division. However, identifying government posts by user names is likely to underestimate the appearance of posts published by governments because some governments also open accounts under private user names or require some government employees to post official messages. Therefore, we also use a machine-learning approach to identify government posts.

C.2.1 Classification

We used the same supervised machine learning approach to classify government posts as we did in the classification of monitoring posts. The key is to construct an accurately labelled sample for the training of a classifier. We used two ways to label a sample of 12,000 posts selected from our vaccine dataset. The first way relies on the user information—a post is labelled as a government post if it is published by a user whose name clearly indicates its government identity. The second way is manual coding based on the content of a post—whether a post is talking about messages sent and announcements made by a government. Two RAs were hired to independently label the selected posts, and their labelling achieved a consistency rate of 94%. In the final dataset we used for machine learning, 22.84% (3287 posts) were labelled as government posts.

Again, we divided the labelled dataset into 10 folds, using seven folds as training data and three folds for testing. We also used SVM to classify government posts. The classification accuracy is very high.

C.2.2 Topic Modeling

In Section 4.4.1 of the paper, we examine governments' blogging activities in response to the information outbreak. Such examination requires knowledge about the content of government posts. We applied the Latent Dirichlet Allocation (LDA) topic modelling method to the government posts obtained from the above textual classification exercise. As a user-friendly unsupervised machine-learning method, LDA is widely used in textual analysis in the Chinese context.

In the baseline version of LDA, we set the number of topics to be five. Given that the content of government posts is simple and rather concentrated, this choice of a relatively small number of topics has the advantage of being easy to interpret. Figure A13 lists the distribution of top 10 words for each topic. Across all topics, the first line is the Chinese words, the second line is the corresponding English translation, and the third line reports the proportion that each word contributes to the topic. Topic 1 is dominated by words such as *work, carry out, preventive vaccination*, indicating this topic is largely about government routine work. The topic words in Topics 2 and 3 are mostly names of specific vaccines and terms used in vaccination. Topics 4 and 5 are rather different. The dominant topic words in Topic 4 include *problem, children (kids), parents, and flow*, indicating that the topic is related to vaccine safety. Topic 5 is clearly about government reaction to vaccine problems, as shown by words such as *illegal (business), legal cases, FDA, Bureau of General Administration, and announcement*. Based on this reading of topic words, we regard Topics 1-3 as being related to government routine work and Topics 4-5 as being related to vaccine safety and government accountability. Since a post is associated with a probability distribution of topic weight, we go one step further to assign a post to a topic by the highest topic-weight. For example, we define a post belongs to Topic 1 if Topic 1 carries the highest weight in the post. After this post-topic assignment, we name posts under Topics 4 and 5 as "accountability posts" and those under Topics 1-3 as "routine-work posts." These measures are used for the regression analysis in Section 4.4.1.

In other versions of LDA, we varied the number of topics. The division between "routine-work" topics and "accountability" topics is rather clear, even when the number of topics is set to 10 or a larger number.

Figure A9: Labeling Instruction

Instructions for Labeling Work
General Information: In March 2016, a vaccine scandal, "Shandong Case", happened and there were a lot of online discussions about vaccines. We constructed a large dataset of Weibo posts containing the keyword "Vaccine" in Chinese characters while at the current stage we don't have a clear picture of the detailed topics and contents that Weibo users talked about. Therefore, we plan to look into the variable **"zhengwen"**, which records the content of all posts. Specifically, we will carry out a task that classifies the posts into subgroups. To begin with, we need to generate a label indicating the category of each piece of post.
Rules for Classification 1: 1. Label 0 represents posts without monitoring power, including some neutral posts promoting scientific knowledge or objective description about vaccination. 2. Label 1 represents monitoring posts expressing personal negative emotions to vaccine issues on security, validity, quality, and so on. It can include emotions like despair, furiousness, umbrage, panic, caricature, complaint, inquiry and etc. These monitoring posts are not necessarily targeted at government/officials/suspects/media actions. 3. vaccine cases mentioned in other social issues should also be considered as relevant posts and classified accordingly. 4. We do not explicitly focus on a specific vaccine scandal. Instead, we will find out all general posts with monitoring power to vaccine issues. 5. For some personal complaint that is quite ambiguous to judge, we will take into account the time of the posts. For example, it is more likely to be label posts as 1 when the posts are complaining of vaccine issues shortly after 2016-03-18.
Rules for Classification 2: 1. Label 0 represents posts from citizens. 2. Label 1 represents posts from governments/officials. 3. We will make the judgment mainly based on the content and tone of posts. If the posts describe or respond to vaccine issues officially and formally, then they should be classified to category 1. 4. When common people forward posts from governments/officials, it should be classified as category 1. That is, we will label posts according to the content of the original post.

Figure A10: Comparison of Model Performance

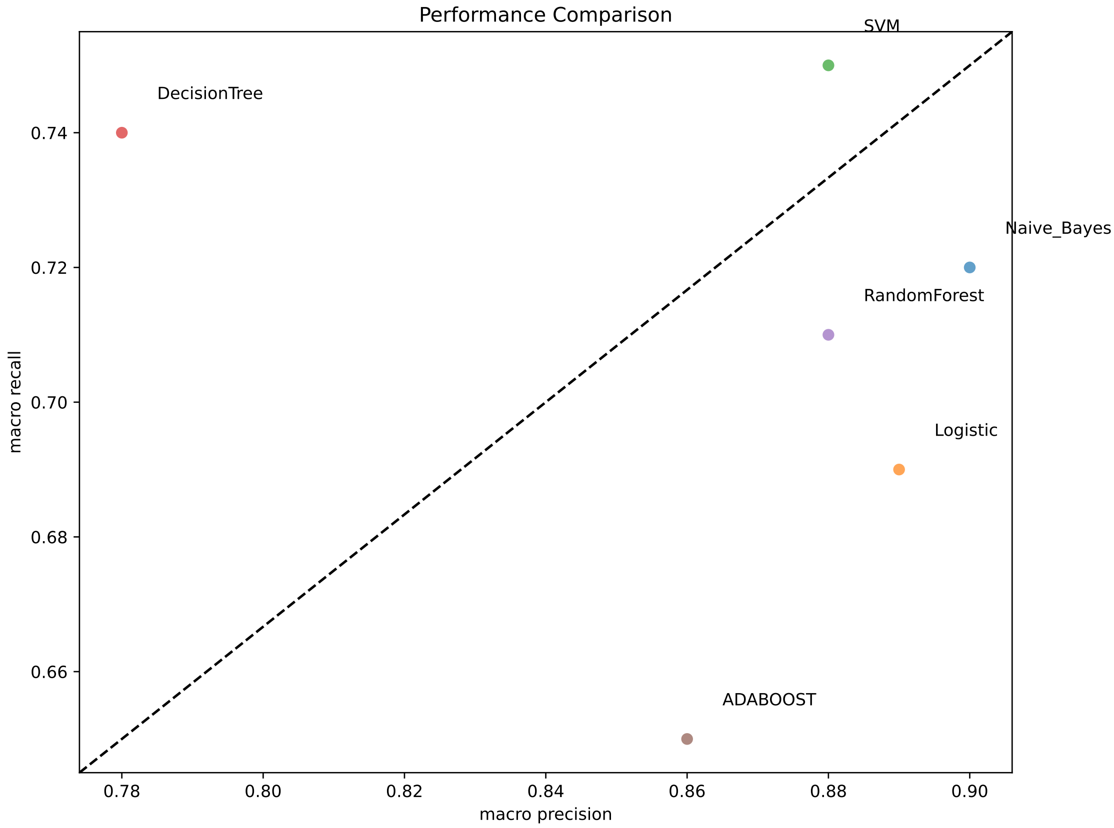

Figure A11: Performance Evaluation for Trained SVM Model

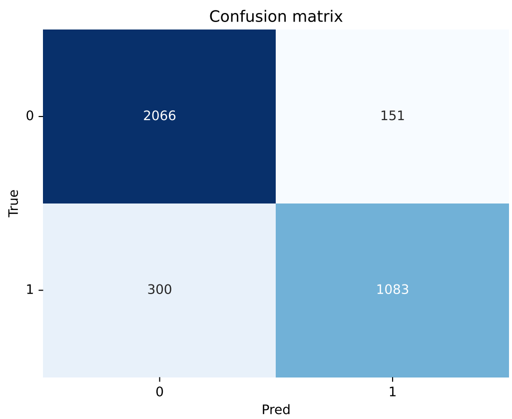

(a) Confusion Matrix

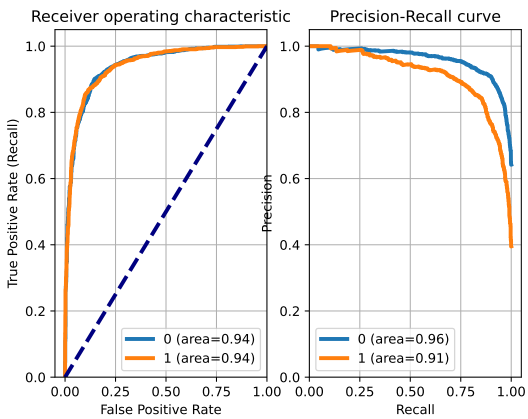

(b) ROC Curve and PR Curve

Figure A12: Correlation of Monitoring-Posts(ML) and Negative-Sentiment-Posts

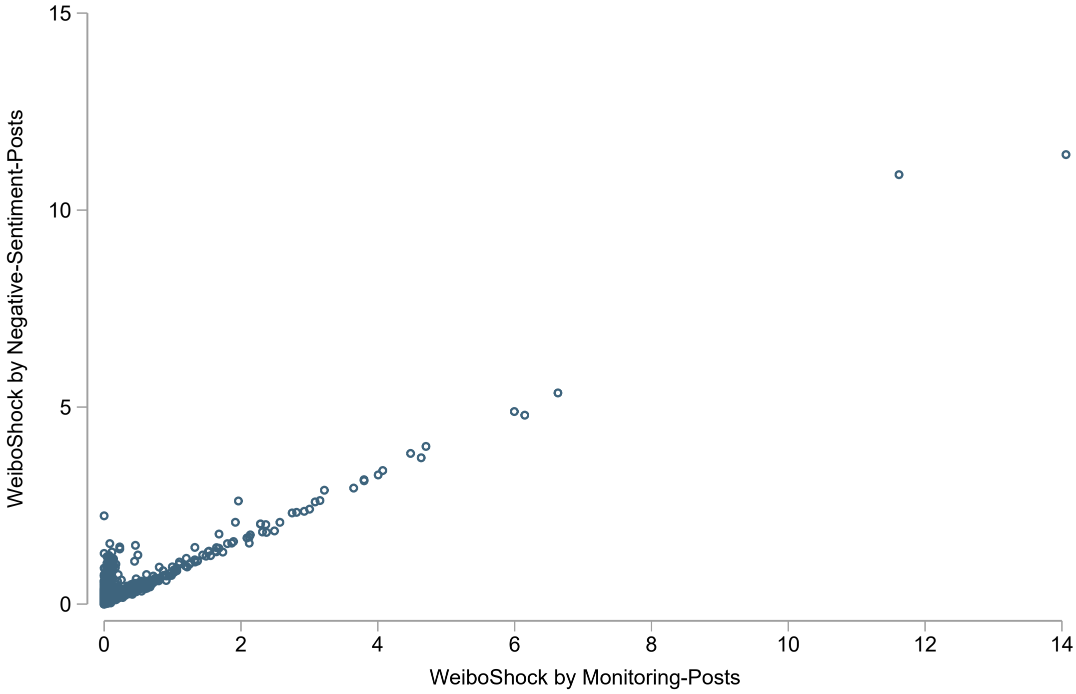

Figure A13: Words Distribution of Each Topic

Topic	Top 10 Words									
	接种	工作	免疫	开展	进行	检查	流感	预防接种	免费	疾病
1	vaccination	work	immune	carry out	proceed	inspection	flu	**preventive vaccination**	free	disease
	0.039	0.015	0.012	0.011	0.010	0.010	0.010	0.009	0.008	0.008
	接种	狂犬病	预防	麻疹	补种	注射	咬伤	病毒	反应	出现
2	vaccination	rabies	prevent	measles	revaccinate	injection	bite	virus	react	appear
	0.022	0.018	0.014	0.013	0.009	0.008	0.008	0.007	0.007	0.007
	接种	活疫苗	预防	脊灰	国务院	脊髓灰质炎	国家	链接	网页	减毒
3	vaccination	valid vaccines	prevent	polio	state department	poliomyelitis	country	link	website	attenuation
	0.020	0.014	0.013	0.013	0.010	0.009	0.009	0.009	0.009	0.008
	接种	问题	疾控中心	儿童	预防接种	家长	孩子	手足口	门诊	流入
4	vaccination	problem	CDCs	children	preventive vaccination	parents	kids	hand-foot-mouth	outpatient service	flow
	0.052	0.016	0.016	0.014	0.013	0.011	0.011	0.010	0.010	0.009
	山东	非法经营	非法	公布	涉案	问题	食药监	总局	案件	药品
5	Shandong	illegal business	illegal	announce	involve the case	problem	FDA	general administration	law case	drug
	0.037	0.029	0.022	0.018	0.016	0.014	0.013	0.013	0.011	0.010