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

Yixin Mei, Ruoyu Qian, and Yanhui Wu *

July 23, 2026

Abstract

This paper studies whether social media promotes government accountability in China in the setting of vaccine procurement during 2014–2019. Exploiting an unexpected scandal that induced cross-city variation in citizen complaints about vaccine safety on Weibo, the Chinese leading microblogging platform, we use difference-in-differences estimation instrumented by celebrity fan bases during early Weibo penetration. Stronger Weibo shocks increase rule-following open-bid procurement and reallocate purchases toward reputable suppliers but delay the delivery. Effects are primarily driven by government-critical sentiment instead of specific information and intensify where officials face information deficits and stronger inspection pressure. Beyond vaccine procurement, enhanced information-monitoring capacity magnifies government responsiveness to high-visibility sentiment but not verified petitions. The findings suggest that social media plays a double-edged role in government accountability: its visibility mitigates agency problems through top-down monitoring pressure, but may also trigger welfare-reducing overreactions.

JEL Code: H4, H7, I1, P2

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

*Mei: School of Economics, Huazhong University of Science and Technology, yixinmei@hust.edu.cn, Hongshan District, Wuhan; Qian: Shanghai University of Finance and Economics, qianruoyu@mail.shufe.edu.cn, Yangpu District, Shanghai, China; Wu: University of Hong Kong & CEPR, yanhuiwu@hku.hk, Pok Fu Lam, Hong Kong. This work benefits greatly from comments and suggestions from Oriana Bandiera, Wouter Dessein, Ruben Durante, Jin Li, Ruixue Jia, John Matsusaka, Juanjuan Meng, Joris Mueller, Canice Prendergast, Michael Zheng Song, David Stromberg, and seminar and conference participants at CUHK, HKUST, HKU, NUS, Peking University, Stanford, UConn, UCSD, USC, the China Economics Summer Institute, the JADE-CEPR-TIME Development Economics Conference, and the 2025 Kuznetz Conference on Development and Organization.

1 Introduction

The lack of transparency and truthful communication within government is often perceived as an obstacle to public accountability. In democracies, the media serves as a powerful tool to promote government responsiveness to citizen needs by bringing politician misconduct and policy oversights to public attention (e.g., Besley and Burgess 2002; Strömberg 2004; Ferraz and Finan 2008; Snyder and Strömberg 2010; Besley and Dray 2023).[1] However, in nondemocracies, traditional media typically fails to deliver information conducive to holding officials accountable, primarily due to state control and government interference (Qin, Strömberg, and Wu 2018). With the advent of social media, the public's capacity to generate and disseminate political information has substantially expanded even in authoritarian regimes (Howard 2010; Shirky 2011; Morozov 2011). Can these relatively unrestricted information flows on social media enhance government accountability in nondemocracies? We address this question by investigating the role of social media in local governments' policy implementation in China. We estimate the causal effect of citizen complaints about vaccine safety expressed on social media on Chinese local governments' vaccine procurement and the resulting consequences.

Our empirical setting represents a general political situation of decentralized decision making with centralized control and imperfect monitoring. Local officials implement policies that directly affect citizens. However, they are held accountable to superior authorities (national leaders), who need officials' input to achieve policy goals but also require information to monitor them. This resembles a classic principal-agent-consumers setting, where the uninformed principal can utilize consumer complaints to monitor the actions of more informed agents (Prendergast 2002; 2003). A fundamental insight is that when the principal's objective aligns with those of the consumers, it is optimal to leverage consumer complaints to oversee the agent's behavior.

Thus, even in a nondemocracy, social media has the potential to enhance public accountability in areas where the regime and citizens share interests, such as public health. This echoes the argument that authoritarian leaders can benefit from a freer media for monitoring officials and rectifying policy oversights (Egorov, Guriev, and Sonin 2009; Lorentzen 2014). However, information circulated on social media is often unverified, imprecise, and emotionally charged. Such information may draw disproportionate attention to certain political features, which may generate a salience effect that triggers overreaction (Bordalo, Gennaioli, and Shleifer 2012; 2013). The critical question is whether and how social media provides an effective feedback mechanism to mitigate local officials' agency problems.

Our study on Chinese governments' vaccine procurement provides a unique oppor-

1. For surveys on the political effects of the media, see Dellavigna and Gentzkow (2010); Prat and Strömberg (2013); Strömberg (2015(a)); Strömberg (2015(b)); Dellavigna and La Ferrara (2015).

tunity to examine the effect of social media on government agency problems and policy making. Government procurement plays an important role in the provision of essential goods and services (Best, Hjort, and Szakonyi 2023; Bosio et al. 2022). Ensuring accountability in this process remains a significant societal challenge in many developing countries (Szucs 2024; Gerardino, Litschig, and Pomeranz 2017). In China, vaccines are exclusively administered by the government, and widespread discontent regarding vaccine quality can erode citizens' trust and support for the regime. Thus, vaccine procurement represents a domain where citizen welfare and regime interests are broadly aligned, yet decentralized implementation creates scope for corruption.

Previous studies have documented that Chinese social media is abundant in discussions about social issues, provided that these discussions do not challenge the authority of the Chinese Communist Party (CCP) or incite collective action (King, Pan, and Roberts 2013; King, Pan, and Roberts 2014; Qin, Strömberg, and Wu 2017). Generally, in domains where the issue is not politically sensitive and the regime shares the same concern as the public, information flows on social media are largely unrestricted. We collect evidence showing that public discussions about vaccine safety on Weibo were not subject to censorship until vaccines became a politically sensitive issue in China during the COVID-19 pandemic. For the purpose of this study, we focus on the pre-COVID period.

We assemble a novel panel of over 23,000 vaccine-related procurement items across Chinese cities (prefectures) during the period of 2014–2019. Since 2005, a national regulation has mandated that all vaccines and related services must be procured by governments before being supplied to healthcare providers. However, the specific procurement arrangements are decentralized to city governments. While the regulation requires that normal vaccine procurement be administered through an open-bid (competitive auction) format, many local governments opt for discretionary procurement such as invited bidding and private negotiation. Such deviation is justified on the basis of operational efficiency but is often viewed as an avenue for rent seeking. We use the frequency and share of open-bid procurement conducted by city governments to measure transparency and policy compliance in governments' decision-making.

We link our procurement data to 3.3 million posts containing the keyword "vaccine" published on Sina Weibo (Weibo for short)—the largest Chinese microblogging platform. Using various machine learning techniques (SVM, BERT, and LLM), we extract Weibo users' opinions about vaccine safety and classify "monitoring posts" (citizen complaints raising safety/oversight concerns). These posts are visible to all levels of governments as well as the general public. In fact, many local governments have resorted to acquiring massive information-processing software to monitor public opinion on social media (Beraja et al. 2023). We further classify whether a monitoring post is informative or emotional to separate the information channel from the sentiment channel.

To identify causal effects, we exploit exogenous variation in the intensity of city level social media complaints about vaccine safety generated by unexpected vaccine scandals. We measure the intensity of citizen complaints 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." We then construct a city-specific information shock on Weibo (Weibo shock hereafter) by calculating the change in our intensity measure following a vaccine scandal within a narrow time window. Finally, we interact this cross-sectional Weibo shock with a dummy variable indicating the event timing to construct a difference-in-differences (DID) estimation in a panel of city-month observations.

Our main analysis focuses on a vaccine scandal in March 2016. It involved a drug ring from northern China, which had been illicitly selling defective and expired vaccines for years. The scandal remained concealed until it was unexpectedly leaked to the public by an online media outlet on March 18. Although originating at the local level, the scandal triggered national discussions about vaccine safety on Weibo. Many posts accused governments of procuring vaccines behind the scenes and failing to monitor the circulation of vaccines. However, these posts are mostly emotionally charged without providing precise information targeting specific individuals or entities.

To strengthen the identification of this event study, we instrument the Weibo shock with the regional distribution of the fans of celebrities who entered the Weibo platform during its initial entry period. This historical endowment captures the path-dependent formation of local social media networks and the resulting density of influential users, which serves as a powerful predictor of subsequent Weibo information outbreaks. Horse-race placebo tests demonstrate that the effect is absent for comparable celebrities who were not early platform adopters. A series of validation tests further confirm that our instrument is uncorrelated with important city characteristics and traditional media discussions, with event-study estimates showing no evidence of diverging pre-trends in government responsiveness.

Analysis of the focal event produces two sets of main results regarding (i) procurement formats and (ii) other substantive outcomes. First, in cities experiencing a stronger Weibo shock, local governments significantly increased both the frequency and share of open-bid procurement. Putting the effect in perspective, in a city with an average population of 4 million, if one out of 10,000 citizens tweeted one more post about vaccine safety in the event month, the local government would have increased the share of open-bid procurement by about 10 percentage points (17% of the pre-event mean). Dynamic estimates show an immediate and persistent (8–9 months) increase without pretrends; the result passes the pre-trend test with linear violation (Roth 2022; Bilinski and Hatfield 2020). The estimated effect is robust to the use of IVs and also survives the scrutiny of potential confounding factors, such as the event itself, tightened central regulation, spillovers from other cities, and other information channels (e.g., newspapers and internet search).

Furthermore, we find evidence consistent with the social media effect driven by both information and sentiment posts. When put together, however, the volume of pure sentimental posts appears to be the main driver. The sentiment channel is particularly relevant to the behavior that a local government converts to universal open-bidding in all vaccine procurement, which we interpret as capitulation (discarding private information to comply with the rule) in view of agency theory.

Second, beyond the compliance of procurement formats, the Weibo shock induces (i) reallocation toward large, reputable, and non-local suppliers (reducing local protectionism); (ii) heightened prosecution of vaccine crimes; (iii) intensified government microblogging that pivots from bureaucratic routines to accountability framing; (iv) modest subsequent declines in adult vaccine-preventable disease incidence; and (v) longer procurement durations and contemporary citizen complaints about shortages.

Overall, these results show that citizen complaints on social media enhance policy compliance, which reduces type-I error (low-quality vaccines in the unsafe state) and increases local accountability. However, local governments' responsiveness to sentimental posts, particularly in cities unaffected by the vaccine scandal, suggests government overreaction, which may increase type-II error (wasteful resources in the safe state) as evident by the increased operational frictions (longer duration and shortage).

The above findings align with a top-down-pressure mechanism, in which social media information reduces monitoring costs while social media sentiment undermines the regime's political goal and triggers excessive monitoring. The top-down-pressure mechanism is more pronounced when local governments are more concerned about being found at fault and punished. Consistent with this argument, we find that, associated with the focal event, local governments' responses in shifting to open-bid procurement are stronger when officials are (i) less informed about local public health conditions; (ii) subject to greater inspection pressure; and (iii) have stronger career concerns. We also find that the social media effect is stronger when governments are less trusted by citizens according to surveys. This counts against the benevolent-government argument that a government directly responds to social media information that reveals citizen needs.

We further validate the top-down-pressure mechanism using a placebo vaccine scandal in 2018 that lacked government-critical sentiment, finding muted effects when citizen complaints do not threaten regime stability. Moreover, during the entire period from 2014 to 2019, we observe a positive correlation between the volume of citizen complaints about vaccine safety on Weibo originating from a city and the local government's adoption of open-bid procurement. This correlation is statistically significant only when citizen complaints target government.

Finally, we extend our analysis to all government procurement (2007-2020), showing that enhanced information-monitoring capacity (i.e., ICT adoption) selectively strengthens local governments' responsiveness to highly visible, negative social media sentiment but not to low-visibility verified portal complaints. These results lend further support to the visibility-driven top-down pressure mechanism.

This paper contributes to the emerging study of 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, Messing, and Adamic 2015; Bursztyn et al. 2019; Allcott et al. 2020; Yanagizawa-Drott et al. 2021) and grassroots collective action (e.g., Acemoglu, Hassan, and Tahoun 2018; Enikolopov, Makarin, and Petrova 2020; Qin, Strömberg, and Wu 2024).[2] Several recent papers study the effect of online information flows on government behavior and public accountability in democracies (Gavazza, Nardotto, and Valletti 2019; Petrova, Sen, and Yildirim 2021; Bessone et al. 2022). Our study provides some of the first causal evidence on the effect of social media on government responsiveness and accountability in nondemocracies.[3] Our findings demonstrate a channel through which authoritarian regimes can improve governance when central and citizen interests align—a finding that extends theoretical insights from Egorov, Guriev, and Sonin (2009); Lorentzen (2014).

Closest to our research is Buntaine et al. (2024), who, in a field experiment, find that public appeals about firms' violation of pollution standards to the regulator through social media reduce violations and tilt regulators' focus toward avoiding public unrest. In their study, social media acts as an instrument to transmit verifiable and issue-specific information about policy implementation. While the information channel in our study is consistent with their findings, what we highlight is the effect of social media posts with low precision but strong negative sentiment. This sentiment channel, which is prominent in a natural social media setting, leads to government overreaction which may mitigate the accountability-enhancing effect induced by the information channel.

Finally, our paper advances the study of information control in nondemocracies, which is crucial for understanding the functioning of an authoritarian regime (Guriev and Treisman 2020; Egorov and Sonin 2024). We demonstrate the double-edged nature of social media monitoring in authoritarian contexts. While social media can enhance accountability, a mechanism relying on visibility instead of accurate information can trigger policy rigidity and overreaction—governments may over-comply with regulations even when local conditions do not warrant such responses, leading to inefficiencies such as procurement delays and supply shortages. This insight provides an explanation for why authoritarian regimes may simultaneously embrace and restrict social media.

The remainder of this paper proceeds as follows. Section 2 describes the institutional background. Section 3 sketches a theoretical model to develop testable hypotheses and guide the empirical analysis. Section 4 describes the data and the empirical strategy.

2. See Zhuravskaya, Petrova, and Enikolopov (2020); Aridor et al. (2024) for recent surveys.

3. Several studies provide indirect evidence on the effect of social media on government accountability in nondemocracies. For instance, Enikolopov, Petrova, and Sonin (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.

Section 5 reports the basic results. Sections 6 and 7 present evidence on the proposed mechanisms within and outside the focal event, respectively. Section 8 concludes. Additional tables and figures as well as a full model and some technical details are collected in an online.

2 Background

2.1 Vaccination in China

China stands as the second largest global market for vaccines, surpassed only by the United States. The Chinese government classifies vaccines into two categories. Category-I encompasses 14 vaccines, including DPT and MMR. Vaccination with Category-I vaccines is mandatory and provided free of charge by the government. Category-II comprises common vaccines such as chickenpox, flu, and rabies, as well as some enhanced versions of Category-I vaccines. Vaccination with Category-II vaccines is voluntary, and individual consumers bear the cost. Nationwide, the vaccination rate for Category-I vaccines exceeds 95%, reaching 100% in major cities. Nevertheless, the vaccination rate for Category-II vaccines remains below 20%, even in the developed regions (Fu 2021).

Despite its large size, the Chinese vaccine market is highly fragmented largely due to the involvement of local governments. For manufacturers, entry into a regional market is challenging without collaborating with established distributors and service suppliers who maintain strong ties with local authorities. It is estimated that the prices of Category-II vaccines, which are ultimately borne by consumers, are more than twice the factory prices. The lucrative nature of the vaccine business, combined with extensive government involvement, creates fertile ground for corruption. Over the past two decades, numerous scandals have exposed the circulation of substandard and defective vaccines posing significant health risks to the public.

2.2 Public Procurement

In China, government procurement accounted for 10% of the overall fiscal expenses and 3.3% of the GDP in 2019. Since 2005, the Chinese government has mandated that all vaccines and related products (equipment) and services (storage and transportation) must be procured through government channels. Suppliers are prohibited from selling their products and services directly to health providers. However, the decisions on the quantity, producers, and format of procurement are decentralized to the prefectural Centers for Disease Control and Prevention (CDCs), under the supervision of the corresponding Bureau of Public Health (2021).

The default format is open-bid procurement, wherein the government publicly announces the procurement, and qualified and interested suppliers compete through (typically sealed-bid) auctions. Other procurement formats include invited-bid in which only a small number of suppliers are invited to bid; private negotiation between the government and selected suppliers; and assigned procurement in which a specific supplier is assigned the procurement task.

While open-bid procurement is rule-based and more transparent, the other formats are discretionary and more opaque. The regulatory emphasis on open-bid procurement is intended to reduce corruption. In practice, local governments frequently depart from open bidding, citing various reasons. One common reason is that open-bid procurement is too rigid to meet sudden spikes in demand (e.g., during flu seasons) or operationally inefficient for small-scale purchases. Another justification is the need for specialized and localized services, such as customized storage solutions that can only be provided by specific suppliers. Nevertheless, such departures from the default open-bid format are often viewed as potential avenues for rent-seeking (Fu 2021).

2.3 Monitoring of Policy Oversight

In China, the implementation of national policies is primarily decentralized to local governments with performance-based rewards and top-down monitoring (Ang 2018; Xu 2011). One important tool for top-down monitoring is the so-called "inspection and appraisal" mechanism, in which higher authorities deploy inspection teams to investigate the subordinate officials (Zhou, Ai, and Lian 2012). This top-down inspection aims not only for performance evaluation but also for regime stability because policy failure in one region can cause public distrust of the government as a whole. The inspection can take the form of auditing. It can also be ad hoc action, triggered by large-scale incidents or widespread public sentiment towards certain social problems (Ai 2011). Receiving a low evaluation in an inspection poses a significant threat to the career of an official. The high stakes associated with inspections create substantial top-down pressure on local officials (Zhou 2022).

In ordinary times, local governments' procurement of vaccines follows a regular bureaucratic procedure, subject to the internal audit of the upper-level authority. However, when serious safety problems catch public attention, vaccine procurement can escalate into a political issue and become subject of a wider scope of inspection involving local political leaders.

2.4 Social Media

Among the various social media platforms in China, Sina Weibo is most prominent for public discourse on societal matters. Launched in August 2009, it experienced rapid expansion in the subsequent years and reached its pinnacle with over 500 million users in 2012 (CNNIC 2013). Although Weibo has lost some ground to other social media since 2013, it continues to hold its influence as a platform for public discussion of social issues.

Chinese social media operates under the direct control of the National Office of

Information Control. Local governments exert limited influence over Weibo's censorship practices.[4] China's censorship of social media is widely recognized as strategic: allowing posts that criticize local governments and officials (King, Pan, and Roberts 2013; King, Pan, and Roberts 2014). Qin, Strömberg, and Wu (2017) document that a large volume of posts discussing political issues, such as corruption, strikes, and protests, circulate on Weibo for extended periods of time.

This information-control strategy is particularly applicable to issues where both the central government and the public share a common interest. Vaccination (drug and food safety more broadly) falls into this category as it directly affects people's well-being. The failure to address these issues can lead to public dissatisfaction with the CCP leadership and erode trust in the regime. Citizens often attribute drug and food safety problems to corrupt local officials and unscrupulous companies. Thus, they have a strong incentive to voice their concerns, and the central government allows relatively free discussions.

Posts that contain citizen complaints about social problems (referred to as monitoring posts for convenience) can affect government responsiveness through two basic channels. First, the monitoring posts reveal the consequences of policies and misconduct by local officials, serving as a source of information to reduce the cost of monitoring. Second, even without providing concrete information, these posts, aggregated at a regional level, can reflect widespread public sentiment against specific local governments. This sentiment channel can generate substantial top-down pressure that directly impacts local officials' career prospects.

The potential impact of social media creates substantial incentives for local governments to monitor information flows and gauge public opinion on social media. One common approach is to assign dedicated personnel to search for relevant posts originating from specific localities. Weibo even offers location-based search functions to facilitate the collection of locality-specific information. In addition, there has been increasing government procurement of information technology with the regime's desire to enhance surveillance (Beraja et al. 2023). Since 2010, local governments have made investments to increase their capacity to collect and analyze social media data.

3 Theoretical Framework

To guide the empirical analysis, we build a principal-agent-citizen model in a decentralized system of policy implementation with imperfect monitoring (auditing). The principal can use bottom-up information (e.g., citizen complaints) to monitor the agent, which may in turn suppress the agent's incentive to use local information. Social media provides information that improves monitoring efficiency and sentiment that inflates the

4. There has been anecdotal evidence showing that local government officials bribed Sina Weibo employees to delete unfavorable posts targeting specific governments and officials. However, such interventions are sporadic and have a limited impact on a small number of posts.

principal's political cost of policy failure. We derive conditions under which social media information is welfare enhancing and conduct comparative statics analysis with regard to the impact of social media visibility on local governments' policy choice and the resulting tradeoff between type-I and type-II errors in policy implementation. A full model with formal proofs of the main results can be found in a theory appendix.

3.1 Model

Our model maps the empirical setting with the following primitives and structure.

Players and states. A central government, the principal, delegates a procurement decision to a local official, the agent. The state of the world is risky or safe: $\theta \in \{R, S\}$, with prior $\pi \equiv \Pr(\theta = R) \in (0, 1)$. The agent chooses either rule-following open-bid procurement O or discretionary procurement D.

Payoffs. Open bidding is transparent but rigid. It incurs an operational cost $c > 0$. Discretion is efficient when the state is safe but hazardous when risky. If the agent chooses D in state R and the choice is not corrected by oversight, safety failure occurs with probability $p \in (0, 1]$ and causes social harm $h > 0$. We assume $ph > c$, so that, in the risky state, open bidding is socially preferable to unmonitored discretion. The agent earns a private benefit $b > 0$ from discretion (e.g., rents from flexible contracting, relationship capital, or corruption) and zero from open-bidding.

Information. The agent privately observes a binary signal $\hat{\theta} \in \{r, s\}$, with accuracy $q \in \left(\frac{1}{2}, 1\right)$. After observing his signal, the agent acts on his posterior belief about the state following Bayes's rule, denoted as ρ_r and ρ_s, respectively. Given $q > \frac{1}{2}$, $\rho_r > \rho_s$, meaning that a risky signal raises the agent's assessment of danger more than a safe signal does.

Monitoring. For simplicity, we assume that the principal commits to a safe-harbor scheme: open-bid O is never audited. A choice of D is audited with intensity $m \in [0, 1]$. An audit reveals the state correctly. If it uncovers risky discretion, the principal overturns the procurement format from D to O and penalizes the agent by $C > 0$. Auditing costs $\mu(V)k(m)$, where k is increasing and strictly convex in effort m and $\mu(V)$ scales it with the information environment defined below. Auditing is thus allocatively valuable.[5]

Social media: volume and composition. Let $V \in [0, 1]$ denote total social media exposure in a jurisdiction, and $\sigma \in [0, 1]$ its sentiment composition, measured by the share of content that is emotive rather than informational. Via the information channel, social media exposure lowers monitoring cost through $\mu(V)$. Via the sentiment channel, negative

5. In practice audits often arrive after a bad event. The model can accommodate this through a stationary dynamic reading: detection deters future misallocation rather than correcting the current one. The principal's trade-off then has the same form, with the per-period benefit $p\tilde{h} - c$ replaced by its discounted continuation value $\frac{\delta}{1-\delta}(p\tilde{h} - c)$, where δ is the discount factor. All results below carry through unchanged.

sentiment towards governments undermines public trust of the regime and imposes a harm to the central government. We assume that the sentiment-weighted exposure σV increases the principal's perceived harm: $\tilde{h}(\sigma V) \geq h$.

One can use machine learning to isolate information posts from sentiment posts. But from the perspective of decision makers, these two channels are often entangled—people express their sentiment in a specific context with implicit information. In this case, we need to assume the curvatures of the above two functional forms to separate the two channels. For example, we can assume that $\tilde{h}(\sigma V)$ is weakly convex to capture the idea that outrage compounds when social media posts go viral and that $\mu(V)$ is decreasing and convex to capture the observation that the first posts reveal genuine local conditions, while later posts are largely retweets that add little new information. With this assumption, when volume rises, the informational return fades while the sentiment force accelerates, so greater social media visibility tilts the system toward the sentiment channel.

3.2 Equilibrium and Results

We focus on the truth-telling equilibrium in which the agent chooses open bidding after a risky signal and discretion after a safe one. Two conditions prevent the agent from choosing this equilibrium strategy. When monitoring is too weak, the agent will take the risk to pursue private benefits. When monitoring is too strong, the agent will lose the incentive to use local information.

The agent will follow his signal only if monitoring lies in an intermediate band, $\underline{m} \equiv \frac{b}{C\rho_r} \leq m \leq \frac{b}{C\rho_s} \equiv \bar{m}$. The lower bound deters rent-seeking discretion when the signal is risky; the upper bound $\bar{m}$ keeps the agent from abandoning efficient discretion when the signal is safe. If pressure exceeds $\bar{m}$, the agent converts to universal open bidding—a situation we call capitulation.

A key discrepancy arises in the assessment of welfare, which defines the accountability issue. Citizens care about the actual harm h while the regime cares about $\tilde{h}(\sigma V)$. We define m^I to be the intensity that maximizes citizen welfare. Two distinct failures can occur. Under-monitoring ($m < m^I$) leaves risky discretion undeterred and sustains safety failures; this is the classic accountability problem. Over-monitoring ($m > m^I$) wastes audit resources and pushes officials toward inefficient rule-following. The information channel attacks under-monitoring by making oversight cheaper while the sentiment channel creates over-monitoring by inflating political harm. Social media can therefore cure one monitoring failure while causing another.

Centering around the above welfare effects of social media, we derive three testable propositions as follows.

Proposition 1 (Visibility and error rates)**.** *There exists a social media visibility threshold $\bar{V}$, below which the agent follows his signal, and when visibility increases, the principal* *monitors more, observed open bidding rises, type-I error (uncorrected safety failure) falls, and type-II error (open bidding in a safe environment) is constant. When social media visibility exceeds $\bar{V}$, the agent discards his private information, capitulates to adopt universal open bidding, and type-II error jumps.*

Proposition 1 presents the basic result that we will test empirically: more social media visibility increases the adoption of open bidding, but it may create a tradeoff between type-I and type-II errors. Note that this media effect is driven by two forces: the principal's increased monitoring effort (below $\bar{V}$) and the agent's capitulation (above $\bar{V}$). The latter force is intensified by the sentiment channel.

Proposition 2 (Information and sentiment channels)**.**

(i) ***Separation case.*** *Both information and sentiment posts increase the share of open-bidding. But the effect of sentiment is greater, and the agent is more likely to capitulate under sentimental pressure.*

(ii) ***Entanglement.*** *Holding the sentiment composition constant, a larger volume of social media posts increases the share of open-bidding.*

Proposition 2 provides two ways to test the sentiment channel empirically. If we assume clean separation of the two channels, the sentiment channel tends to generate a larger response from the agent. However, if the two channels are entangled, one way to test the dominance of the sentiment channel is to look at the effect of the volume holding the observed sentiment-information ratio constant. This is driven by the assumption that sentiment dominates at a high social media volume. In such a case, the accountability gain from social media is mitigated.

The choice of procurement format results from the information asymmetry between the principal and agent and their preference misalignment. These relationships can be affected by the agent's signal quality q and the size of penalty C when the agent's action matches the wrong state. Thus, these two parameters can interact with social media visibility V to affect the agent's procurement choice.

Proposition 3 (Heterogeneous effects)**.**

(i) ***Signal quality*** *The effect of social media visibility on the share of open-bidding is dampened by the agent's signal quality q. This visibility-signal interaction operates through reducing auditing pressure before reaching the capitulation threshold (intensive margin) and increasing the capitulation threshold (extensive margin).*

(ii) ***Penalty*** *The effect of social media visibility on the share of open-bidding is amplified by the penalty C. This visibility-penalty interaction operates only through reducing the capitulation threshold (extensive margin).*

Intuitively, agents with better private information have fewer mistakes to be caught and are more confident to withstand scrutiny. They are less responsive to the pressure created by social media visibility. In contrast, a harsher penalty will make agents more vulnerable to inspection pressure enhanced by social media visibility. Interestingly, the interacting effect between visibility and penalty operates only through the extensive margin. This is because the agent's choice is discrete, and C enters only through the band limits. Within the band, the penalty is a transfer without allocative value. Empirically, this second part of Proposition 3 provides a way to test the overreaction mechanism generated by excessive top-down pressure.

4 Data and Empirical Strategy

4.1 Data

4.1.1 Vaccine-related Data

Vaccine procurement By regulation, all procurement of vaccines and related products must be made public through government websites. We collected relevant procurement documents from both national and subnational government websites from 2010 to 2019.

Each procurement is accompanied by an announcement notice and a deal-closure notice. From the announcement notice, we extracted key information such as the name of the procuring government entity (typically a local CDC), the specific item to be procured, the date of the announcement, and the chosen procurement format. From the deal-closure notice, we obtained the names of the winning companies and, if available, the procurement amount. Unfortunately, the price information for specific procured items is largely missing. We are unable to use price differentials to infer administrative effectiveness and potential corruption.

We obtained a total of 23,035 item-level observations from 2014 to 2019 across 273 cities. These items are classified into three categories: Category-I vaccines, Category-II vaccines, and supplements (supplemental materials, facilities, and services). We aggregate the item-level information at the city-month level.

Table 1 presents the basic summary statistics for the count and share of the rule-following open-bid procurement (as opposed to the discretionary procurement) at the city-month level. The number of observations is limited to 1,630 as we exclude city-month observations without vaccine procurement. On average, open-bidding is used 4.3 times per city-month, accounting for 53% of total vaccine procurement. The share of open-bid procurement is slightly higher for Category-II and supplements (56%), but it is still far below the level if the default format were adopted.

4.1.2 Social Media Data

We collected posts referencing "vaccine" in Chinese published on Weibo from 2014 to 2019. The data was obtained from a specialized third-party provider that collected social media data in China. To ensure data quality, we compared this dataset with posts we scraped from Weibo using the same keyword. The two datasets yielded very similar sets of posts, with our data covering approximately 95% of the third-party data.

We believe that social media discussions about vaccine issues, unless they contain anti-regime content, are largely unimpeded by censorship from the Chinese government. To verify this, we conducted a test by searching for the word "vaccine" on a website (https://freeweibo.com/) exclusively accessible outside mainland China. This website compiled Weibo posts that had briefly appeared on the official Weibo platform before being removed. As shown in Figure A1, we found no posts referencing vaccines on this website until 2021, when the COVID-19 vaccination became a politically sensitive issue. In contrast, when we searched for other keywords related to political leaders, protests, or disasters, we consistently discovered posts dating back as early as 2012.

Our dataset comprises a total of 3,328,889 posts referencing vaccines. For each post, we obtained information on its content, posting time, and the city-level user location. The user-location data was derived from the users' self-reported information on their Weibo profiles.[6]

To separate posts containing citizen complaints about vaccine safety, which we refer to as "monitoring posts," we employ various machine learning approaches, including support vector machine (SVM), sentiment analysis, deep learning (BERT), and large language model (LLM). Below, we briefly describe these approaches and their outcomes, with detailed explanations provided in Appendix C.

Our benchmark approach is SVM, a supervised machine-learning method commonly used in the development of information monitoring software employed by Chinese local governments during our sample period. We hired research assistants to manually label 12,000 Weibo posts randomly selected from our dataset. A part of the labeled data were used to train an SVM classifier, which assigned each post to either the monitoring posts or the non-monitoring posts category. After performing cross-validation to evaluate the classification accuracy (87% overall accuracy indicated in the confusion matrix), we applied the trained SVM classifier to the entire dataset.

Additionally, we employ three alternative algorithmic methods to identify monitoring posts. First, we utilize sentiment analysis to classify the vaccine-related posts with stronger negative sentiment (relative to positive sentiment), based on a commonly used Chinese word dictionary, "HowNet." Second, we apply BERT with fine-turning to auto-

6. To validate its accuracy, we selected a subset of users who had granted Weibo permission to track their real-time locations and compared these to their self-reported locations. Our analysis revealed an over 90% accuracy of match between the self-reported and real-time locations of these users.

matically classify the vaccine-related monitoring posts. Third, we adopt a recent LLM by Alibaba, for which we design a prompt to instruct the pre-trained LLM and fine-tune it to indicate whether a given Weibo post carries monitoring implications. Unsurprisingly, the BERT and LLM approaches outperform the SVM approach in classification accuracy. Nevertheless, we primarily use the monitoring posts classified by SVM to measure the Weibo shock because our classification exercise aims to imitate Chinese local governments' capabilities to gauge public opinion. We use the more sophisticated measures as robustness checks.

Among the SVM-identified monitoring posts, we conduct two further classification exercises to isolate content that can better uncover the mechanisms. First, we identify those "government-critical" posts (targeting governments as opposed to firms or products) by a set of keywords related to government entities and titles of officials. Second, we employ a LLM to separately classify social media posts into (i) whether they are informational and (ii) whether they are sentimental. A post is defined as informational if it contains verifiable or trackable informational cues, including specific locations or time, vaccine manufacturers, brands, batches, distribution, or other factual claims. A post is defined as sentimental if it contains strong negative emotions, such as anger, condemnation, or moral accusations. The two categories are not mutually exclusive. Purely informational posts are defined by excluding sentimental posts from the set of informational posts.

Table 1 reports the summary statistics for the various measures of Weibo discussions about vaccines. On average, there are approximately 173 posts referencing vaccines per city-month, with substantial variance. The mean value of "Monitoring Posts (ML)," which represents the number of monitoring posts identified by SVM, is only 24, less than 1/7 of the mean value of the total number of posts. The means for "Monitoring Posts (DL)" and "Monitoring Posts (LLM)," which measure the number of monitoring posts identified by BERT and LLM respectively, are 30 and 27, close to the mean for the SVM-identified posts. The average number of posts with negative sentiment is approximately 63, more than double the means of the other measures. These four measures of monitoring-posts are strongly correlated. The mean of the "government-critical" posts is 6.77, considerably smaller than the mean of monitoring posts. The bottom three rows of Table 1 report measures of Weibo discussions in the focal event month, during which the mean of "Monitoring Posts(ML)" is 540 and the mean value of informational and sentimental posts identified by LLM are 316 and 381, respectively.

Appendix Figure A2 plots the monthly time-series of the four Weibo variables mentioned above. Notably, two prominent spikes correspond to vaccine scandals. These spikes would likely not have occurred if there had been government censorship of public discussions about vaccine-related issues. The spikes are so sudden that they substantially alter the information landscape on Weibo for a short period. We leverage the regional variations induced by these shifts to identify the social media effect.

4.2 Empirical Strategy

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

4.2.1 Event Setting

The event we study is an outbreak in public opinion on social media triggered by a vaccine scandal. On March 18, 2016, an online media outlet, The Paper, exposed a vaccine distributor from Shandong province for selling defective or expired vaccines worth 88 million US dollars across 18 provinces over a period of six years.[7] The vaccines involved in this scandal were mostly category-II vaccines and related services. This news leak immediately sparked heated discussions about vaccine issues on Weibo. As depicted in Appendix Figure A3, the number of Weibo posts referencing "vaccine" remained relatively low (around 1,000) prior to the news leak. However, it surged to 14,825 on March 18, peaked at 75,506 on March 22, and declined rapidly after one week.

We also plot the number of articles discussing vaccine issues published on WeChat public accounts and newspapers.[8] The dynamics of these articles mirror the pattern observed in Weibo posts, indicating that the information outbreak on Weibo was triggered by the news leak and was unexpected to the public and government.

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. Following the investigations, a total of 355 individuals (including 64 civil servants) involved in the scandal were arrested and prosecuted.[9] At the same time, the central government mandated the procurement of vaccine-related products to be reported on the provincial digital platform. However, local governments still kept the decision rights regarding the choice of vaccine suppliers and procurement formats. While the new regulations were announced in April, they were not implemented in most provinces until the end of 2016.

4.2.2 Econometric Specification

In the aforementioned setting, the release of news about the vaccine scandal created a strong and unforeseen information shock to each city. This regional variation in information shocks serves as the foundation of our identification strategy. Figure A4 illustrates

7. See https://www.thepaper.cn/newsDetail_forward_1445232.

8. WeChat is the most popular Chinese social media platform known for private messaging within a restricted group. We searched for vaccine-related articles published on 59 prominent WeChat public accounts specializing in healthcare. For newspaper data, we retrieved articles mentioning "vaccine" from Wisenews, a digital archive of Chinese newspapers based in Hong Kong. During our sample period, Wisenews covered 87 general-interest newspapers from 31 cities in Mainland China.

9. See https://www.thepaper.cn/newsDetail_forward_2045189.

the landscape of public opinion on vaccine safety in China, measured by the number of monitoring posts published by Weibo users in each city from February to May 2016. In February, public discussions about vaccine issues were minimal; only a few major cities displayed some engagement. However, in March, the landscape experienced a dramatic shift, marked by a surge in vaccine-related discussions and notable regional disparities. Over the subsequent two months, the intensity of discussions on Weibo gradually diminished, returning to pre-event levels. Our identification leverages the abrupt change in Weibo discussions on vaccine issues within a city during this period, with cities exhibiting more pronounced shifts considered to have received a stronger treatment.

Specifically, our estimation applies 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 and t indicates year-month. The dependent variable y_{it} is the number (in logarithms) 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). $WeiboShock_i$ 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 same measure averaged over the three months before the event. This variable captures the intensity of information shock to a city at the timing of the event and is time-invariant. The coefficient of the interaction term β is our interest of estimate.

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, and the share of the secondary industry in GDP.[10] We also include province-specific time trends to control for the provincial-level demand fluctuations and potential procurement guidance from provincial governments. ϵ_{it} is the error term, clustered by city. Clustering the standard error at the province level or two-way clustering at both the city and year-month levels barely changes the statistical significance of the coefficients of the main variables.

The key identification assumption is that, absent the region-specific and vaccine-related information shock induced by the event, local governments would not change their behavior in vaccine procurement. We will examine the dynamic effects to verify the absence of a pretrend. We will also rule out threats to this assumption arising from

10. These variables are sourced from the Chinese City Yearbooks and are available at the year level. 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 and applied a logarithmic transformation.

factors that align with the regional distribution of the Weibo shock during the event and also impact local government procurement practices.

4.2.3 Instrumental Variable

To address the concern that the "treatment" variable $WeiboShock_i$ might be endogenous, we instrument it with an "intention-to-treat" variable: the fan base of celebrities who happened to be earlier Weibo users. Weibo was introduced by Sina Corporation in August 2009, following the Chinese government's ban on Twitter and Facebook as well as the shutdown of several emerging domestic micro-blogging platforms. Initially, Weibo was unknown to most Chinese internet users. In an effort to build its user base, Weibo invited numerous celebrities from various fields to promote the platform. Some celebrities accepted the invitation and became active users, attracting their fans to join Weibo. These early-adopting celebrities and their fans formed a significant portion of the early Weibo user base, which potentially affected the subsequent information outbreaks on the platform. In spirit, this IV is akin to the one used by Enikolopov, Makarin, and Petrova (2020), who utilize the social connections of the founder of the focal social media platform as an instrument for social media penetration in Russia.

Practically, we define early Weibo celebrities as those individuals who had accumulated more than 10,000 followers on Weibo by October 31, 2009. This selection criterion results in a list of 44 celebrities from a variety of sectors. For each celebrity, we measure his or her pre-Weibo popularity in a city using the Google Search Index specific to that celebrity within that city during 2008. Given the ordinal nature of the Google Search Index, we convert it into a cardinal measure of fan base by multiplying it by the total number of followers the celebrity had on Weibo in 2009. Then, we aggregate this city-person measure across all celebrities to capture the influence of Weibo celebrities in each city during the early stage of Weibo's development.

The IV-relevance condition requires that the popularity of the early Weibo celebrities in a city exerts a long-lasting effect on the information outbreak sparked by the 2016 vaccine scandal in that city. A plausible mechanism is that the followers of these celebrities became some of the initial Weibo users, who were more engaged and impactful in shaping social discourse than later users. We provide several pieces of evidence for this mechanism. First, this initial fan base catalyzed early platform adoption. As shown in Table A1, cities with higher 2008 search interest developed significantly higher Weibo penetration rates in the subsequent years during the platform's formative period (2009-2013). Second, cities with a larger initial fan base fostered a significantly higher density of impactful Weibo users, defined as those who had more than 1 million followers by 2012. This suggests that the IV identifies regions that developed strong networks capable of large-scale information dissemination. Third, at the individual level, users who registered in 2009 remained considerably more active and possessed substantially more followers compared to those who joined in later years; see Figure A5. Such a persistent effect of early social media users is also found in Enikolopov, Makarin, and Petrova (2020); Müller and Schwarz (2023); Fujiwara, Müller, and Schwarz (2024).

The exclusion restriction condition necessitates that our measure of city-level celebrity popularity in 2008 does not affect local governments' vaccine procurement other than through the Weibo channel. A primary concern is that the IV might proxy for persistent city characteristics, such as a more ingrained civic culture, higher information demand, or inherent bureaucratic openness.

To isolate the Weibo-specific channel, we perform a series of horse race regressions by including placebo instruments: the 2008 search interest for 2009 Forbes celebrities who were not Weibo users (non-Weibo figures) and 2011 Weibo late-comers. Specifically, we identify a group of 44 comparison celebrities who were equally popular as the aforementioned Weibo celebrities according to the ranking of Forbes 2009 but did not use Weibo during this period. The second placebo group comprises 44 celebrities who were regarded as top Weibo influencers in 2011 but joined Weibo only after 2009. We construct the regional popularity of these two placebo groups of celebrities (named "$Cindex(Forbes)$" and "$Cindex(2011)$", respectively), as we do for the IV.

As a preliminary check, Figure A6 shows that our primary IV strongly predicts the Weibo outbreak during the 2016 event, while the placebo instruments lose their predictive power once the primary IV is controlled for. These results rule out general celebrity interest or civic culture as alternative divers, isolating the early-adopter network effect pertinent to Weibo. Correspondingly, as we will show later, reduced-form estimates confirm that only the primarily IV remains a significant predictor of procurement choice when placebos are included.

We also find that the IV has no significant impact on discussions in traditional print media or general web news (measured by the intensity of web search for the events) during the 2016 scandal, as shown in Table A2. This evidence rules out the possibility that the IV captures a general increase in a city's civic awareness or media attention. Additionally, Table A3 shows that the identifying variation from our IV is not systematically associated with time-varying city covariates, including GDP per capita, fiscal expenditure, and internet and cell-phone penetration.

We next present evidence that supports the exogenous entry of Weibo celebrities. As illustrated in Figure A7, a celebrity's popularity, whether measured by Google search volume or Forbes' exposure index, shows no correlation with the date of their first post on Weibo. Apparently, which celebrities were invited and eventually decided to join Weibo cannot be predicted by their initial popularity. This mitigates the concern that Weibo strategically invited influential celebrities from certain sectors and regions.

Other instrumental variables. To test the robustness of the baseline IV, we develop two complementary instruments. The first augments our baseline IV by incorporating an information shock specific to the focal event, while the second utilizes another source of early Weibo penetration based on the flows of colleague-students across cities.

The augmented IV leverages social media connections between cities established during 2011-2013 when Weibo expanded to its full strength. We define these connections using the number of retweets between city pairs as in Qin, Strömberg, and Wu (2024). We then interact a city's 2008 celebrity popularity with its social media linkage to Jinan (the epicenter of the vaccine scandal). This approach strengthens the baseline instrument by introducing a shock directly tied to the event's location and timing.

The second alternative IV is based on flows of college students from central cities to other cities in 2009. We exploit two facts. First, the penetration of Weibo is highly unbalanced in the initial stage: high penetration in several central cities (Beijing, Shanghai, Guangzhou, and Shenzhen in particular) and low penetration in other cities. Second, college students constitute an important portion of the initial Weibo users. When students from a central city go to colleges in other cities, they are likely to increase the use of Weibo in these destination cities. We use student registration data for college freshmen in 2009 to measure student flows across cities as in Qin, Strömberg, and Wu (2024).

5 Main Results

In this section, we focus on the above event-study setting to test Proposition 1: whether the eruption of Weibo information affected local governments' adoption of open-bidding and the resulting type-I and type-II errors.

5.1 Baseline Estimation

Using specification (1), we estimate the effect of the Weibo shock on vaccine procurement in a panel of Chinese governments that procured at least one vaccine-related product from March 2015 to March 2017. We choose this time window to keep the same period before and after the event and avoid potential overlapping with other events.[11] This selection results in a sample of 124 cities, representing the larger prefectures among the 339 Chinese prefectures. Geographically, these cities are distributed across 31 of the 32 provinces in Mainland China.

Panel A of Table 2 reports the results using the share of open-bid procurement as the dependent variable. The estimation in the first two columns uses data including all vaccine-related procurement. Column (1) corresponds to our baseline specification (1) with controls for city fixed effects and year-month fixed effects. The estimate shows that the Weibo shock increases a local government's share of open-bid procurement by approximately 10 percentage points at the 1% significance level. This implies that one additional monitoring post published per 10,000 people in a city leads to an approximately

11. For instance, in July 2017, the death of a woman after taking rabies vaccines triggered some discussion about vaccine safety.

10 percentage points increase in the share of open-bid procurement, or a 17% increase relative to the pre-event average (0.61).[12] In Column (2), in addition to the existing controls, we add sequence fixed effects. In this way, we reconstruct the panel so that the time series is arranged by the sequence of procurement, which creates a more balanced panel during the months close to the event. The estimated coefficient remains virtually unchanged.

The last four Columns of Table 2 report the Weibo effects by categories—Columns (3)−(4) for Category-I vaccines and Columns (5)−(6) for Category-II vaccines and supplements. We combine Category-II vaccines and supplements into a single outcome variable because they are often procured together, and vaccine safety concerns apply to both of them. Column (5) shows that in cities where one additional monitoring post per 10,000 people is published in a month, the share of open-bid procurement increases by 19.6 percentage points. Such an increase is substantial given the mean share (0.64). The estimated effect barely changes in the specification with sequence fixed effects.

While the effects on the procurement of Category-II vaccines and supplements are sizable and statistically significant at the 1% level, the effects on Category-I vaccines are small and insignificant. Clearly, the positive effect of Weibo on overall procurement presented in the first two columns is attributed to Category-II vaccines and supplements. This makes sense because Category-I vaccines are mandatory and cheap, leaving local governments with little incentive to exercise discretion for rent-seeking. In contrast, Category-II vaccines and supplements are procured with greater discretion and are the major source of safety and corruption concerns.

Balanced panel Given the high cost of vaccine transportation and the fact that a typical vaccine, if stored appropriately, is valid for 12-24 months after being manufactured, it is not surprising that local governments, particularly those in smaller cities, do not procure vaccines frequently. To address the concern of unbalanced panel data, we estimate the effect on procurement frequency in a balanced panel assigning a value of zero to the city-month with no vaccine procurement. Panel B of Table 2 reports several noteworthy results. First, the mean values of the dependent variable for Category-I vaccines are much smaller than the corresponding means for Category-II vaccines and supplements. This is because Category-I vaccines have a longer expiration time period and are procured less frequently. Second, Columns (1) and (2) show that the Weibo shock significantly increases the frequency of open-bid procurement but does not affect the nonopen-bid procurement. A similar pattern is observed for Category-II and Supplements (Columns

12. Admittedly, interpreting the magnitude of the Weibo effect poses some challenges. First, a single post may be read by many people, leading to a substantial information shock even with a small number of original posts. Second, the effect of Weibo posts is not necessarily linear, and the marginal effect may be disproportionately large when the number of posts is small. Third, the effect of the Weibo shock persists over a long period, and the cumulative effect is sizable even if the effect in a single period is modest.

(5) and (6)). It helps rule out the concern that the Weibo effect on the share of open-bid procurement is driven by the decline of nonopen-bid procurement.

Note that, for Category-I vaccines, the Weibo shock increases the frequencies of both open-bid and nonopen-bid procurement, but only the effect on nonopen-bid procurement is statistically significant. One potential explanation is that open-bid procurement tends to cause delay and temporary supply shortage, and local governments may use the faster nonopen-bid procurement for the safer Category-I vaccines, some of which substitutes Category-II vaccines.

Dynamic effects Figure 1 plots the dynamics of the DID estimates around the event for Category-II vaccines and supplements. Specifically, we interact $WeiboShock_i$ with a sequence of time dummies indicating periods before, during, and after the event. We normalize the effect one period before the event to zero.

Figure 1a reports the dynamic effects on the share of open-bid procurement by calendar month, which corresponds to our baseline specification. To estimate these dynamic effects with greater precision, we choose a two-month interval for each period.[13] For clearer exposition, we omit the estimate for the event month although it is included in the estimation. In the absence of a notable pretrend, the effect becomes evident immediately following the event and remains significant for up to eight months afterward.

Figure 1b reports the dynamics by procurement sequence, corresponding to the specification with sequence fixed effects. As expected, the estimation becomes more precise around the event, and the magnitude of the estimates remains large. Figure 1c reports the dynamics within the balanced panel. The red (blue) dots plot the estimates for the open-bid (nonopen-bid) procurement. The Weibo effects on the frequency of open-bid procurement are significantly positive and precisely estimated within the first six months after the event, while the effects for nonopen-bid procurement are close to zero. The overall pattern is similar to the dynamics of the estimates using the share of open-bid procurement as the outcome variable (Figure 1a). These results provide further reassuring evidence that data unbalancedness is not a significant concern in our setting.

Tests of pretrends The patterns of dynamics from different estimations shown in Figure 1 consistently indicate the absence of any pretrends. To further test the potential violation of parallel trends, we follow the "one step up" approach proposed by Bilinski and Hatfield (2020) and the power test for linear-trend violations suggested by Roth (2022). The former approach shows that adding a linear trend barely affects the estimates of the post-event treatment effects; see Appendix Figure D1. The latter approach shows that the bias caused by linear violation does not invalidate the baseline estimation; see Figure D3. These results demonstrate that our DID estimates are unlikely to be driven by violations of the parallel-trends assumption. Appendix D presents more results with detailed discussions.

13. The monthly dynamics exhibits a similar pattern although the estimates are less precise.

5.2 Other Measures of Weibo Shock

Discrete measure In Table A4, we report regression results using the same specification as in the baseline, but replacing the continuous measure of Weibo shock with a dummy variable that equals one if the magnitude of the Weibo shock exceeds the mean level. The results remain qualitatively consistent with those in Table 2.

LLM-based measure We have constructed two more accurate measures of Weibo shocks based on deep learning (BERT) and LLM, WeiboShock(DL) and WeiboShock(LLM). Figure A8 reports the corresponding DID estimates, with all the Weibo-shock measures being standardized. The first three rows show that the two technologically more advanced measures generate estimates similar to the baseline effects. Furthermore, we regress the outcome variable on two competing interaction terms with the event timing: the baseline Weibo shock and a variable that includes monitoring posts identified by LLM but missed by SVM. This latter variable captures the "measurement error" of the baseline measure from the perspective of researchers with more advanced technology. As shown in the bottom two rows, such measurement error does not increase the share of open-bid procurement. These results suggest that measurement error in the construction of the Weibo shock is unlikely to drive our findings. Rather, they are consistent with the fact that the technology that governments use to gauge public opinion is less-advanced during our sample period.

5.3 IV Estimation

We implement the aforementioned IV approach to validate the estimation of the causal effect of the Weibo shock on vaccine procurement. Specifically, we instrument the variable "$WeiboShock_i$" in specification (1) with our measure of regional fan base of the Weibo celebrities in 2009, referred to as $Cindex$, following a two-stage-least-squares (2SLS) regression.[14] Table 3 presents the results from this celebrity-fans IV estimation with the outcome being the share of open-bid procurement. The IV estimates reported in Column (1) (overall procurement) and Column (2) (category-II vaccines and supplements) are slightly larger than their counterparts in the baseline estimation (Table 2). The Kleibergen-Paap rk Wald F statistic is reasonably large. We further check the adjusted standard error based on the tF critical value (Lee et al. 2022). The adjusted standard error (0.043) is slightly larger than the conventional one (0.035) but the coefficient remains statistically significant.

14. Given the control of city fixed effects, we practically instrument "$WeiboShock_i \times Event_t$" with "$Cindex_i \times Event_t$". Our IV (regional popularity based on Google search) does not cover all the cities in our baseline estimation. To ensure that our IV estimate is comparable with the baseline estimate, we use a LASSO approach to predict the popularity of celebrities for the missing cities, which account for about 30% of the cities in our sample. The R-squared in our LASSO prediction is approximately 80%. Estimation without any predicted values produces similar results despite in a smaller sample. Standard errors calculated by bootstrapping methods are smaller than those reported in this manuscript.

In the remaining part of the table, we focus on the Category-II vaccines and supplements. Column (3) reports the first-stage regression, and Column (4) reports the reduced-form estimation. Columns (5)-(6) present two falsification tests further assessing the validity of the IV. In particular, in the reduced-form regression, we run a horse race between the IV (Weibo celebrities in 2009) and a placebo variable indicating the influence of other types of celebrities, which were previously discussed. Including this placebo variable helps control for the unobservable local conditions correlated with the influence of celebrities in general. Column (5) accounts for the competing effect from the non-Weibo celebrities in 2009 ("$Cindex(Forbes)$") and Column (6) for the competing effect from Weibo-celebrities in 2011 ("$Cindex(2011)$"). Consistent with our previous finding (Figure A6), the coefficients of these two placebo variables are negative and statistically insignificant while the coefficient of our IV ("$Cindex$") remains large and significant.

Figure 1d plots the dynamics of the IV estimates. The pattern is similar to the dynamics of the baseline estimation presented in Figure 1a. Appendix Figure A9 reports the dynamics of a reduced-form event study. The estimates are small and statistically insignificant for all periods leading up to the event. This absence of pre-trends in both the reduced-form and IV event studies indicates that high-IV cities were not on a prior trajectory of institutional improvement.

Appendix Table A5 reports the estimated effects of the Weibo shock on the share of open-bid procurement using the aforementioned two alternative IVs: celebrity-fans augmented with event-specific information shocks and student inflows from the four cities with the highest level of early Weibo penetration. As shown in Columns (1) and (2), compared with the baseline IV, the event-specific augmented IV produces very similar estimates. Columns (3) and (4) report the results when we instrument the 2016 Weibo shock with the student-flows IV. The estimated Weibo effects are slightly larger than the celebrity-fans IV estimates. For comparison, we use the initial Weibo penetration across cities, measured by the per-capita number of all types of posts, as an IV.[15] As shown in Columns (5) and (6), the estimates from this "coarser" IV become slightly smaller. The consistent and similar results derived from different IVs provide compelling evidence that endogeneity in the baseline estimation is not a significant issue.

5.4 Robustness Checks

In Table B1, we present evidence that rules out a number of potential confounding factors: (i) event effect; (ii) policy effect; (iii) administrative reform; (iv) informational spillover; and (v) other information channels. While leaving detailed discussions to Appendix B, we briefly discuss the impact of other informational channels here.

15. In practice, we derive the residuals from regressing the city-level Weibo penetration during its launch period (August-October 2019) on the numbers of college students, internet users, mobile phone users, and land-line users in a city in 2009. We then use these residuals to instrument the Weibo shock in 2016. Standard errors are bootstrapped throughout the entire estimation.

Local governments may respond to information from other sources, such as traditional media, other social media, and the internet. In Appendix Figure A3, we illustrate that WeChat discussion and newspaper coverage did not precede the outbreak of Weibo posts about vaccine safety. Panel B of Table B1 shows that the effect of either newspaper coverage or internet search is small and has little impact on the estimated Weibo effect. The absent impact of newspapers is likely for two reasons. First, after 2014, the Chinese government tightened the control of the media (Qin, Strömberg, and Wu 2024). Second, the newspaper market was severely eroded by social media, and the loss of readership and advertising revenues further crippled the monitoring function of newspapers.

5.5 Evidence on Capitulation

In our theory, one important mechanism that triggers the adoption of open-bid procurement is that local governments capitulate and convert to universal open-bidding. The baseline results presented previously comprise both the monitoring effect and the capitulation effect. Given the limited knowledge about local conditions and the decision process, we examine the capitulation hypothesis using evidence on whether a local government is more likely to choose open-bidding across various procurement items (universal open-bidding) upon the Weibo shock. Specifically, we construct a binary outcome variable that equals one when the share of open-bidding in a city-month is one and zero otherwise. That is, we estimate a linear-probability model. Table A6 reports the results of regressing this dummy outcome variable on the baseline Weibo shock, using the same controls as before.

Columns (1) and (2) use the same sample as in the baseline estimation. Both the OLS and IV estimates are positive and statistically significant. In the next two columns, we include only the city-months where the procurement frequency is greater than one. Despite being smaller, this sample better reflects the idea of universal open-bidding. The effects are small and statistically insignificant.

Columns (5)-(8) repeat the estimation in the first four columns restricting to the sample of category-II vaccine and supplements. The effects of the Weibo shock on the measure of capitulation are sizable and statistically significant in the sample with all procurement (Column 5) and in the sample with multiple procurement (Column 7). The corresponding IV estimates produce even stronger effects. Overall, the evidence supports the existence of capitulation.

5.6 Type-I and Type-II Errors

As argued in the theory, the choice of procurement format results in a tradeoff between type-I error and type-II error. In this subsection, we present evidence on these aspects to assess the social consequences of the Weibo shock.

5.6.1 Tradeoff between Quality and Operational Efficiency

Local protectionism One important distinction between open-bid procurement and the less-transparent formats is the larger pool of participants involved in open bidding. This helps alleviate the problems of local protectionism, the phenomenon where regional governments offer preferential treatment to firms operating within their administered cities. Local protectionism results in a widespread supply of lower-quality products and services provided by small local firms in China (Tang, Wang, and Wu 2025).

Measuring vaccine quality is challenging due to the limited availability and accuracy of price information. Therefore, we focus on the effect of the Weibo shock on the composition of vaccine suppliers. Table 4 reports the regression results for different binary classifications of producers: large vs. small firms and non-local vs. local firms.[16] The econometric specification follows the baseline model, with the dependent variable being the share of procurement frequency by firm type.

Panel A shows that cities experiencing a stronger Weibo shock are more likely to procure vaccines from reputable large firms. The effect is sizeable and statistically significant for both open-bid and nonopen-bid procurement. Panel B reveals that a stronger Weibo shock led city governments to procure vaccines from nonlocal firms more frequently. Interestingly, the effect is primarily driven by nonopen bid procurement. This is likely because local protectionism is more pertinent to nonopen procurement.

Operational efficiency One potentially detrimental consequence of open-bid procurement is the protracted duration due to its requirement for extensive participation and thorough review. This may lead to a delay in the supply of vaccines to the market and potentially even temporary shortages. Table A7 reveals a substantial positive effect of Weibo on the duration of open-bid procurement at the 1% significance level.[17]

The prolonged duration of vaccine procurement coincides with anecdotal reports on the shortage of vaccine supply. Complaints filed with local leaders' offices in certain cities reveal a shift in the proportion of complaints related to vaccine shortages.[18] Prior to the vaccine scandal in 2016, the share of such complaints stood at 12.3%, increasing to 35.1% thereafter. To further investigate this issue, we selected a subset of Weibo posts that

16. Local firms, in this context, refer to companies whose registration address is located within the administrative regions of the procuring government. These firms encompass both vaccine manufacturers and distributors. Large firms are defined based on their size, specifically those whose sales exceed the median; all the listed companies fall under the category of large firms.

17. Because the deal-closure information is not accurately recorded, we measure the duration by the number of days between the announcement date and the expiration date of the announcement notice, which usually coincides with the first meeting with potential bidders. The average duration before the event is approximately 20 days.

18. The data of citizen complaints are sourced from the Local Leader Message Board (LLMB), which is a government platform for citizens to submit their complaints and requests to top officials in a city. Using "vaccine" as a keyword, we identify 2,677 vaccine-related complaints from 2013 to 2020. Each complaint is manually reviewed to assess whether it pertains to a supply shortage. Approximately 25% of the complaints are classified as related to vaccine shortages.

contained keywords related to vaccine shortages. The average number of posts expressing concerns about shortages rose from approximately 50 before the event to nearly double that number afterward. Notably, a spike in Weibo posts discussing vaccine shortages was observed in July 2016, a quarter after the scandal.

5.6.2 Vaccine-related Infectious Diseases

In addition to impacting the quality of vaccines available in the market, the Weibo shock may also influence citizens' decisions regarding the vaccination of discretionary Category-II vaccines. Unfortunately, reliable data on the vaccination rates of Category-II vaccines is not accessible to researchers. We turn to the incidence of infectious diseases for which vaccination has a preventive effect.

Specifically, we estimate the effect of the Weibo shock on the incidences of infectious diseases in a region using the baseline DID specification. Unfortunately, data on vaccination rates or infectious diseases at the city-month level are unavailable. We have to use province-monthly data on infectious diseases in this part of the analysis.[19] As shown in Table A8, the Weibo shock reduces the number of cases (per 100,000 people) of vaccine-preventable diseases by 0.068 (0.034/0.5) standard deviations. This effect is predominantly driven by the most prevalent adult infectious diseases such as influenza, tuberculosis, and hepatitis, which are associated with Category-II vaccines. Note that the effect on diseases affecting children is negligible, likely because these diseases are mostly related to Category-I vaccines. The monthly dynamics of our DID estimates (Figure A11) reveal no pretrend, with the effect becoming significant and stabilized nine months after the event and remaining persistent afterwards.

5.7 Evidence on Government Blogging and Monitoring

5.7.1 Government Blogging about Vaccine Issues

Since 2011 when social media gained popularity in China, the central government has encouraged local governments to engage with the public through platforms like Weibo. In 2012, there were approximately 50,000 Weibo accounts operated by government offices or individual officials (Sina Weibo 2013). These provide a channel for the central government to observe and monitor local officials' activities, particularly in response to crises. Therefore, the eruption of public grievances on Weibo is likely to affect local governments' blogging activities on their Weibo accounts.

To investigate local governments' blogging activities, we use a machine-learning approach to identify the posts published by Weibo accounts representing governments; see Appendix C for details. Our estimated presence of government posts is approximately 4.81% of all posts referencing vaccines, comparable to the estimate by Qin, Strömberg, and Wu (2017). But these posts account for only a tiny share (0.7%) of the monitoring

19. These data are obtained from numerous official documents that we collected from various sources.

posts. Then, we apply the Latent Dirichlet Allocation (LDA) topic modeling approach to the machine-identified government posts, from which we distinguish topics related to vaccine safety and public accountability from those topics about government routine work. As shown in Table A9, in cities where discussions about vaccine safety on Weibo increased more on the event date, local governments responded by publishing a greater number of posts on vaccine-related issues. Additionally, the content of these posts shifted from routine work to matters of public accountability. Interestingly, after the National FDA's official announcement, local governments responded to the eruption of monitoring posts by refocusing more on routine work to align with the National FDA's guidance.

5.7.2 Legal Inspection of Vaccine Crimes

In response to the vaccine scandal, China's Supreme People's Procuratorate mandated that procuratorial agencies nationwide pursue related criminal cases. This directive establishes the prosecution of illegal vaccine activities as a measure of government accountability. Because Chinese courts are integrated into the political system, they are susceptible to media influence, making the surge of public discourse on Weibo a potential driver of judicial outcomes.

To investigate this, we collected 399 vaccine-related criminal cases during 2014-2019 via the China Judgments Online Database.[20] These cases primarily involve bribery, abuse of power, and illegal business operations. We manually coded 128 cases as being directly related to the scandal and identified 213 cases that involved government officials. From this case-level data, we constructed a city-year panel aggregating the number and proportion of cases, criminals, and implicated officials, distinguishing between those related and unrelated to the scandal.

Using a city-year panel from 2016 to 2019, we estimate the effect of the Weibo shock on these criminal prosecutions. The OLS results in Table A10 show a positive and significant correlation between the Weibo shock and scandal-related outcomes, both in absolute numbers (Columns (1)-(3)) and as a share of total cases (Columns (4)-(6)). Conversely, we find no significant relationship for non-scandal-related outcomes.

6 Evidence on Mechanisms

We have shown evidence consistent with the basic theoretical result presented in Section 3. We now test the mechanisms, centering around Propositions 2 and 3.

6.1 Information and Sentiment Channels

One important observation of the Chinese social media data is that citizen complaints on Weibo often contain little information about specific local conditions and entities. Instead, Weibo posts are highly sentimental, especially in regions where the volume is

20. For detailed description of the database, see Chen, Chen, and Yang (2025); Liu et al. (2022).

large. We decompose the Weibo shock into the shock of informative posts and the shock of sentimental posts. Specifically, we aggregate the LLM-classified informational and sentimental posts to construct city-level variables $WeiboShock(Info = 0, Senti = 1)$ and $WeiboShock(Info = 1, Senti = 0)$ analogous to $WeiboShock$.

Panel A of Table 5 reports the estimates of the effects of the informative and sentimental shocks, using the baseline DID specification. Columns (1) and (2) show that both channels independently and significantly increase the share of open-bidding. Column (3) reports the result when putting together the two channels in the same regression. The sentiment effect is positive, sizable and statistically significant while the information effect is negative and statistically insignificant. In Column (4), we test the second part of Proposition 2, which indicates that fixing the composition of sentiment and information posts, the volume of the Weibo shock has a positive effect on open-bidding under the assumption that sentiment posts dominate at a high volume. The result supports this hypothesis. The same pattern is found in Columns (5)-(8) when the outcome variable concerns Category-II vaccines and supplement.

Panel B of Table 5 reports the same regression results using the capitulation measure as previously defined: a binary variable equal to one when the share of open-bidding is one and zero otherwise. The results are consistent with those in Panel A, providing further support to the hypothesis that, in addition to the information channel, the sentiment channel exerts a substantial effect on government responsiveness. Its dominance suggests the possibility of government overreaction, which offsets the informational effect.

6.2 Heterogeneous Effects

Proposition 3 states, under the top-down-monitoring pressure, how social media visibility interacts with local governments' private information and the size of penalty to influence their choice of procurement formats. We test these implications within the same event-study framework. We extend the baseline DID regression (1) to a triple-differences model to estimate a series of heterogeneous treatment effects 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 a predetermined city condition used to test a specific hypothesis. All other variables are defined as in the previous specification. The coefficient attributed to the triple-interaction term, θ, captures the heterogeneous treatment effects of interest. We include the same set of control variables as in specification (1) and continue to cluster standard errors at the city level. Table 6 reports the regression results.

Information asymmetry We assume that an official who is more informed of local conditions are more likely to draw a high-quality signal. Thus, we construct two measures to assess how officials are informed of local conditions: (i) local officials' career background and (ii) the informational environment of the region. For the first measure, we gather the resumes of the mayors (the chief mayor and the one who is responsible for public health) and calculate the number of years they spent working in the public health sector or as a local officer before 2016. To measure the informational environment in a region, we use the proportion of Weibo posts deleted in a province as a proxy for the scarcity of local information available to officials. This variable is sourced from Bamman, O'Connor, and Smith (2012), who tracked the deletion of a substantial number of Weibo posts across provinces in 2011. A higher rate of post deletion indicates a more tightly controlled media environment and likely less-informed officials.

As shown in Columns (1)-(4) of Table 6, in cities where officials are less informed of local conditions, the social media shock has a more pronounced impact on local governments' use of open-bid procurement.

Inspection penalty We first consider local politicians' perceived probability of being inspected due to exogenous events. Since 2013, the Chinese central government has launched a series of anti-corruption campaigns aimed at dismantling patronage networks within a region. In these campaigns, cities that have faced frequent corruption allegations tend to attract increased scrutiny from top-down inspections (Jiang, Shao, and Zhang 2022). Therefore, local politicians in areas that had recently experienced more of these anti-corruption actions were likely to be more sensitive to negative shocks. Using data on anti-corruption allegations at the provincial level collected from the Procuratorial Yearbooks of China, we find that the Weibo shock has a stronger effect in cities where anti-corruption allegations were more frequent right before the event. This is evident by the sizable and significant coefficient of the triple-difference term in Column (6).

We next consider local officials' career concerns. Generally, the stake of being inspected and punished diminishes over an official's career cycle. Moreover, younger officials or those in the earlier tenure are more susceptible to top-down inspections, because they have not yet amassed sufficient performance metrics and loyalty credits (Jia, Kudamatsu, and Seim 2015; Ang 2018). Based on detailed personnel information in the CVs of the mayors in our sample cities, we construct a variable, $Pre-retirement$, which is a binary indicator for city leaders whose age is 54 and older, considering that the mandatory retirement age is 60 and a term of city leaders is five years.[21] As shown in Column (8), the Weibo effect on the procurement of Category-II vaccines and supplements is more pronounced for officials with stronger career concerns.

As argued in Proposition 3, the visibility-penalty interaction operates exclusively through the capitulation margin and thus provides support to the top-down-pressure

21. We also construct another variable, $First\ Term\ of\ Tenure$, which indicates whether 2016 fell within the mayor's initial term in office. The result is similar.

mechanism. Yet, one may argue that our results can be explained by local politicians' direct responses to citizens' need in their constituents as in democracies. In nondemocracies, citizens have limited power to constrain politicians' behavior, and the direct response relies, to a large extent, on a benevolent government. Below, we provide evidence that helps limit this explanation.

Public trust To measure how benevolent a local government is, we use responses to questions regarding citizen trust in local officials from surveys conducted in the Chinese General Social Survey (CGSS).[22] Based on the proportion of survey respondents who express a positive attitude towards local officials, we construct a dummy variable for public trust in local officials, which equals one if the relevant proportion in a city is above the median for all sample cities. According to the benevolent government explanation, a government that citizens trust more should be more responsive to citizen complaints. However, as shown in Columns (9)-(10) of Table 6, in cities with lower levels of public trust, the effect of the Weibo shock on procurement transparency is stronger. Interestingly, the coefficient of the interaction term $Low\ Trust \times Event$ is negative. This is plausible because officials who are less trusted by the public are more likely to ignore the impact of the event if it did not spark citizen complaints.

7 Further Evidence

We provide further evidence on this mechanism beyond the focal event. First, we exploit a 2018 vaccine scandal with limited government-critical sentiment as a placebo. Second, we examine correlations between fluctuations in monitoring posts and vaccine procurement transparency outside event windows. Finally, we move beyond vaccine procurement, leveraging variation in the timing of local adoption of information and communication technology (ICT) (2007–2020) across all government procurement to examine whether enhanced monitoring capacity strengthens government responsiveness to highly visible social media sentiment or to low-visibility verified complaints.

7.1 Placebo Event

In our argument, the top-down pressure arises from the regime's political goal to maintain stability and public trust, which can be undermined by the prevalent government-critical sentiment on social media. In contrast, social media discussions that do not express government-critical sentiment, even if they contain information regarding vaccine safety, are less likely to propel local governments to respond. To test this conjecture, we

22. We use CGSS data from 2010 and 2012 due to data availability. The survey includes the following question: "How would you rate your level of trust in local officials?" Based on the answer in terms of scale, we create a binary indicator for trust, which equals one if a surveyee's chosen trust level exceeds a neutral position (i.e., "Somewhat Trustworthy" or "Completely Trustworthy"). We also utilize data from the 2012 and 2014 waves of the China Family Panel Studies (CFPS) to conduct the test. The result remains consistent.

leverage another surge of information on Weibo.

In July 2018, a public vaccine manufacturer, Changsheng Bio-Technology Co., located in Jilin province, was discovered to have provided false information about the production of a significant quantity of rabies vaccine. This event sparked intense discussions on Weibo and received extensive coverage from traditional media outlets. However, public opinion during this time primarily centered on the issue of corporate accountability. In terms of the number of related posts, the ratio of content targeting government to content targeting firms is 3.84 in 2016 and 0.68 in 2018. In terms of word frequencies, the ratio is 7.19 in 2016 and 0.45 in 2018.

We employ the same event-study approach to analyze this new context, with the Weibo shock being defined as the difference between the per-capita number of monitoring posts in a city during the event month (July 2018) and the average number of posts from three months prior to the event. Table A11 shows that, across all specifications, the effects of the Weibo shock on the share of open-bid procurement are minimal and statistically insignificant. Note that the average share of open-bid procurement before 2018 is comparable to that before the focal event in 2016. It is not the case that the lack of government responsiveness is due to the share of open-bid procurement having reached a sufficiently high level with no further room for adjustment. Consistent with this primary finding, we observe that the Weibo shock has a limited impact on other outcome variables. These findings suggest that when public sentiment does not generate top-down political pressure, the influence of media is subdued.

7.2 Evidence beyond Event Study

Figure A10 plots the time series depicting the share of open-bid procurement of vaccines throughout the entire period of 2014-2019. The noticeable fluctuations suggest that local governments do not consistently institutionalize their vaccine procurement practices, even though they exhibit increased policy compliance during periods of information outbreaks. Instead, their adherence to procurement regulations appears to vary over time, with periods of strict compliance followed by deviations from these regulations. We investigate whether this variability in local governments' adoption of open-bid procurement is associated with fluctuations in information flows on Weibo.

Table A12 reports the OLS estimation regressing the share of open-bid procurement for all categories of vaccines in the present month on the per-capita numbers of Weibo posts referencing vaccine and containing negative sentiment. In the first two columns, the regressions are run within the entire sample from January 2014 to December 2019, while in the last two columns, we exclude the periods of Weibo-information outbreaks triggered by the two aforementioned events, namely, March to August 2016 and July to December 2018. The regressions control for city and time fixed effects, provincial time trends, as well as time-varying city characteristics.

Columns (1) and (3) show a significant positive correlation between the number of posts referencing vaccine with negative sentiment and the share of open-bid procurement, with a two-month lag. In Columns (2) and (4), we separate two types of vaccine posts: those targeting the government and those targeting firms. The correlation between Weibo opinion and government procurement formats is driven by the government-targeted posts, but not by the firm-targeted posts.

7.3 Evidence beyond Vaccine Procurement

We move beyond vaccine procurement to examine the proposed top-down mechanism. Because social media data on a range of social issues are not available over the full period for which reliable procurement data exist, we employ an indirect approach that exploits variation in local governments' capacity to gauge public opinion.

As discussed in Section 2.4, Chinese local governments made substantial investments in ICT to expand their capacity to gather and analyze citizen opinions—especially those voiced on social media. The adoption of ICT can produce two different effects depending on the mechanism through which media matters. If the information mechanism dominates, ICT increases a government's awareness of and responsiveness to citizen needs. However, under the top-down-pressure mechanism, ICT would disproportionately heighten responses to opinions more likely to trigger top-down inspection.

We exploit the timing of each local government's first ICT adoption and estimate its effect using a city-year panel from 2007 to 2020, covering all government procurement. The dependent variable is the share of open bids among all procured items, our proxy for rule-following policy implementation as before. We retrieve the timing of ICT adoption from our procurement dataset, since all ICT items are mandated to be acquired through formal procurement. The decision to adopt ICT is decentralized to city governments and depends largely on local fiscal capacity. Most governments adopt between 2010 and 2016, generating substantial cross-city timing variation. All regressions use OLS controlling for city characteristics and city and year fixed effects.

Most importantly, we interact the ICT adoption dummy with two measures of citizen complaints. The first is the intensity of citizen complaints submitted to the aforementioned Local Leader Message Board (LLMB), which measures verified information without public visibility. The second measure draws on the geographic distribution of followers of right-leaning (i.e., government-critical) key opinion leaders (KOLs). These KOLs frequently comment on social issues with strong emotion and effectively generate government-critical sentiment. The per-capita number of their followers in a city proxies the potential for the spread of negative sentiment. We construct a "KOL index" using the city-level intensity of Baidu (Chinese equivalent to Google) searches in 2009—the early period of Weibo expansion—for the most prominent right-leaning KOLs.

Table A13 reports the regression results. The first two columns present the effect of ICT adoption for overall government procurement. Column (1) shows no statistically significant average effect on procurement transparency. With interactions added (Column (2)), the coefficient on the KOL index is sizable and significant at the 1% level, while the coefficient on official-portal complaints is not. This pattern suggests that ICT adoption enhances responsiveness to public opinion that is potentially negative and highly visible.

To further test this mechanism, we estimate effects by procurement category (remaining columns of Table A13). The interaction between ICT adoption and verified complaints is small and insignificant in most categories except legal matters and agriculture—the two areas with the most frequent petitions. In contrast, the interaction between ICT adoption and social media sentiment (KOL index) is generally sizeable and significant, with agriculture and government administrative matters being the only exceptions—categories that typically attract limited public attention.

8 Conclusion

The rise of social media has provided authoritarian leaders with greater control over bottom-up information flows. Does this redistribution of information enhance government accountability and thus the competency of authoritarian regimes? Motivated by this question, this paper demonstrates that, in China, social media acts as a channel for increasing local governments' responsiveness and accountability, particularly in areas where the central government and the public share common interests, such as vaccine procurement. We find that intensified citizen complaints regarding vaccine safety on social media in a city cause its local government to improve policy compliance by increasing the share of more-transparent (open-bid) vaccine procurement. This, in turn, leads to more frequent procurement from large reputable providers rather than small local providers, likely resulting in better health outcomes.

Moreover, we find that the social media effect on local governments' behavior is primarily driven by high volumes of posts containing low-precision but emotionally charged complaints. The effect is particularly pronounced for a group of relatively "vulnerable" governments, whose top officials are less informed about local conditions, less trusted by local citizens, subject to greater top-down inspection pressure, and driven by stronger career concerns.

Our findings align with a theoretical framework in which the principal (authoritarian leader) leverages citizen complaints from social media to monitor agents (local officials). In this relationship, beyond the informational value of citizen complaints, the high visibility of social media information amplifies the cost of ignoring citizen complaints about policy oversights for the authoritarian leader. As a result, top leaders have a strong incentive to pressure local governments to address public grievances and citizens' needs. This insight, consistent with theoretical studies on authoritarian politics (e.g., Xu 2011;

Gehlbach, Sonin, and Svolik 2016; Egorov and Sonin 2024), is likely applicable to other authoritarian regimes. Moreover, the result that social media offers a channel to escalate a local problem to a political salient signal is also relevant to democratic settings.

It might seem that the government's responsiveness to public opinion and citizen concerns in our study mirrors the effect of a free media in democracies. However, the analogy is limited. In a democracy, citizens (voters) act as the principal and can remove non-compliant officials when adequately informed. In an authoritarian regime, citizens resemble customers with little direct power. Citizen complaints become operational only when the principal's preferences align with theirs, and observed "responsiveness" is mediated by potential top-down intervention. This structure creates heightened risks of policy rigidity: local governments under strong upward pressure may overreact, undermining flexible adaptation. One of the most striking examples is China's exceptionally strict lockdown policies during the COVID-19 pandemic. Thus, free information flows on social media, even when non-threatening to the regime, operate as a double-edged sword in an authoritarian context.

References

Acemoglu, Daron, Tarek A. Hassan, and Ahmed Tahoun. 2018. "The Power of the Street: Evidence from Egypt's Arab Spring." *The Review of Financial Studies* 31 (1): 1–42.

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

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020. "The Welfare Effects of Social Media." *American Economic Review* 110 (3): 629–676.

Ang, Yuen Yuen. 2018. *How China Escaped the Poverty Trap.* Cornell University Press.

Aridor, Guy, Rafael Jiménez-Durán, Ro'ee Levy, and Lena Song. 2024. "The Economics of Social Media." *Journal of Economic Literature* 62 (4): 1422–1474.

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. "Exposure to Ideologically Diverse News and Opinion on Facebook." *Science* 348 (6239): 1130–1132.

Bamman, David, Brendan O'Connor, and Noah Smith. 2012. "Censorship and Deletion Practices in Chinese Social Media." *First Monday* 17.

Beraja, Martin, Andrew Kao, David Y. Yang, and Yuchtman Noam. 2023. "AI-tocracy." *Quarterly Journal of Economics* 138 (3): 1349–1402.

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

Besley, Timothy, and Sacha Dray. 2023. "The Political Economy of Lockdown: Does Free Media Matter?" *European Journal of Political Economy* 78:102361.

Bessone, Pedro, Filipe Campante, Claudio Ferraz, and Pedro CL Souza. 2022. "Social Media and the Behavior of Politicians: Evidence from Facebook in Brazil." *Working Paper.*

Best, Michael Carlos, Jonas Hjort, and David Szakonyi. 2023. "Individuals and Organizations as Sources of State Effectiveness." *American Economic Review* 113 (8): 2121–67.

Bilinski, Alyssa, and Laura A. Hatfield. 2020. "Nothing to See Here? Non-inferiority Approaches to Parallel Trends and Other Model Assumptions." *Working Paper.*

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *The Quarterly Journal of Economics* 127 (3): 1243–1285.

———. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121 (5): 803–843.

Bosio, Erica, Simeon Djankov, Edward Glaeser, and Andrei Shleifer. 2022. "Public Procurement in Law and Practice." *American Economic Review* 112 (4): 1091–1117.

Buntaine, Mark, Michael Greenstone, Guojun He, Mengdi Liu, Shaoda Wang, and Bing Zhang. 2024. "Does the Squeaky Wheel Get More Grease? The Direct and Indirect Effects of Citizen Participation on Environmental Governance in China." *American Economic Review* 114 (3).

Bursztyn, Leonardo, Georgy Egorov, Ruben Enikolopov, and Maria Petrova. 2019. "Social Media and Xenophobia: Evidence from Russia." *NBER Working Paper.*

Chen, Heng, Yuyu Chen, and Qingxu Yang. 2025. "Women in the Courtroom: Technology and Justice." *Review of Economic Studies, forthcoming.*

CNNIC. 2013. *The 32nd Statistical Report on Internet Development in China.* Technical report. Beijing: China Internet Network Information Center.

Dellavigna, Stefano, and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Reviews* 2:643–669.

Dellavigna, Stefano, and Eliana La Ferrara. 2015. "Economic and Social Impacts of the Media." In *Handbook of Media Economics, Chapter19,* 1:723–768. North-Holland.

Dumais, Susan, John Platt, David Heckerman, and Mehran Sahami. 1998. "Inductive Learning Algorithms and Representations for Text Categorization." *Proceedings of the Seventh International Conference on Information and Knowledge Management,* 148–155.

Egorov, Georgy, Sergei Guriev, and Konstantin 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, Georgy, and Konstantin Sonin. 2024. "The Political Economics of Non-democracy." *Journal of Economic Literature* 62 (2): 594–636.

Enikolopov, Ruben, Alexey Makarin, and Maria Petrova. 2020. "Social Media and Protest Participation: Evidence From Russia." *Econometrica* 88 (4): 1479–1514.

Enikolopov, Ruben, Maria Petrova, and Konstantin Sonin. 2018. "Social Media and Corruption." *American Economic Journal: Applied Economics* 10 (1): 150–174.

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

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

Fujiwara, Thomas, Karsten Müller, and Carlo Schwarz. 2024. "The Effect of Social Media on Elections: Evidence from The United States." *Journal of the European Economic Association* 22, no. 3 (June): 1495–1539.

Gavazza, Alessandro, Mattia Nardotto, and Tommaso Valletti. 2019. "Internet and Politics: Evidence from U.K. Local Elections and Local Government Policies." *Review of Economic Studies* 86 (5): 2092–2135.

Gehlbach, Scott, Konstantin Sonin, and Milan W. Svolik. 2016. "Formal Models of Nondemocratic Politics." *Annual Review of Political Science* 19 (1): 565–584.

Gerardino, Maria Paula, Stephan Litschig, and Dina Pomeranz. 2017. "Distortion by Audit: Evidence from Public Procurement." *NBER Working Paper,* 23978.

Guriev, Sergei, and Daniel Treisman. 2020. "A Theory of Informational Autocracy." *Journal of Public Economics* 186:104158.

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

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

Jiang, Junyan, Zijie Shao, and Zhiyuan Zhang. 2022. "The Price of Probity: Anticorruption and Adverse Selection in the Chinese Bureaucracy." *British Journal of Political Science* 52 (1): 41–64.

Joachims, Thorsten. 1998. *Making Large-Scale SVM Learning Practical.* Technical report.

King, Gary, Jeniffer Pan, and Margaret E. Roberts. 2014. "Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345 (6199): 1251722.

King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 326–343.

Lee, David S., Justin McCrary, Marcelo J. Moreira, and Jack Porter. 2022. "Valid T-Ratio Inference for IV." *American Economic Review* 112 (10): 3260–3290.

Liu, Ernest, Yi Lu, Wenwei Peng, and Shaoda Wang. 2022. "Judicial Independence, Local Protectionism, and Economic Integration: Evidence from China." *Working Paper.*

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

Morozov, Evgeny. 2011. *The Net Delusion: The Dark Side of Internet Freedom.* PublicAffairs.

Müller, Karsten, and Carlo Schwarz. 2023. "From Hashtag to Hate Crime: Twitter and Antiminority Sentiment." *American Economic Journal: Applied Economics* 15 (3): 270–312.

Petrova, Maria, Ananya Sen, and Pinar Yildirim. 2021. "Social Media and Political Contributions: The Impact of New Technology on Political Competition." *Management Science* 67 (5): 2997–3021.

Prat, Andrea, and David Strömberg. 2013. "The Political Economy of Mass Media." *Advances in Economics and Econometrics* 2:115.

Prendergast, Canice. 2002. "Consumers and Agency Problems." *The Economic Journal* 112 (478): 34–51.

———. 2003. "The Limits of Bureaucratic Efficiency." *Journal of Political Economy* 111 (5): 929–958.

Qin, Bei, David Strömberg, and Yanhui Wu. 2017. "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda." *Journal of Economic Perspectives* 31 (1): 117–140.

———. 2018. "Media Bias in China." *American Economic Review* 108 (9): 2442–2476.

———. 2024. "Social Media and Collective Action in China." *Econometrica* 92 (6): 1993–2026.

Roth, Jonathan. 2022. "Pre-Test with Caution: Event-study Estimates After Testing for Parallel Trends." *American Economic Journal: Insight* 4 (3): 305–322.

Sebastiani, Fabrizio. 2002. "Machine Learning in Automated Text Categorization." *ACM Computing Surveys* 34 (1): 1–47.

Shirky, Clay. 2011. "The Political Power of Social Media: Technology, the Public Sphere, and Political Change." *Foreign Affairs* 90 (1): 28–41.

Snyder, James M Jr., and David Strömberg. 2010. "Press Coverage and Political Accountability." *Journal of Political Economy* 118 (2): 355–408.

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

———. 2015(b). "Media and Politics." *Annual Review of Economics* 7 (1): 173–205.

———. 2015(a). "Media Coverage and Political Accountability." In *Handbook of Media Economics, Chapter13,* 1:595–622. Elsevier.

Szucs, Ferenc. 2024. "Discretion and Corruption in Public Procurement." *Journal of the European Economic Association* 22 (1): 117–160.

Tang, Wei, Yuan Wang, and Jiameng Wu. 2025. "Local Favoritism in China's Public Procurement: Information Frictions or Incentive Distortion?" *Journal of Urban Economics* 145:103716.

Xu, Chenggang. 2011. "The Fundamental Institutions of China's Reforms and Development." *Journal of Economic Literature* 49:1076–1151.

Yanagizawa-Drott, David, Gianluca Russo, Maria Petrova, and Ruben Enikolopov. 2021. "Echo Chambers: Does Online Network Structure Affect Political Polarization?" *Working Paper.*

Zhou, Xueguang. 2022. *The Logic of Governance in China: An Organizational Approach.* Cambridge University Press.

Zhou, Xueguang, Yun Ai, and Hong Lian. 2012. "The Limit of Bureaucratic Power in Organizations: The Case of the Chinese Bureaucracy." In *Research in the Sociology of Organizations,* 34:81–111. Emerald Group Publishing Limited.

Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov. 2020. "Political Effects of the Internet and Social Media." *Annual Review of Economics* 12 (1): 415–438.

Figures and Tables

Figure 1: Dynamic Effects of the DID Estimation

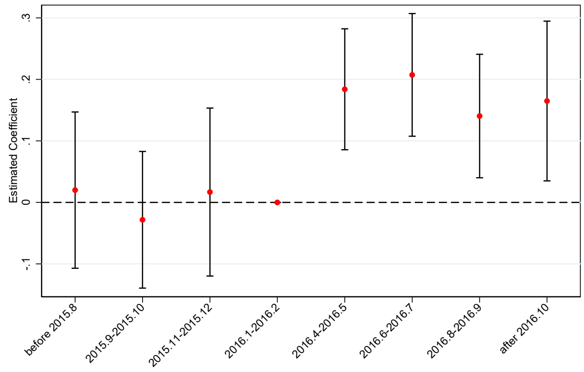

(a) open share by procurement month

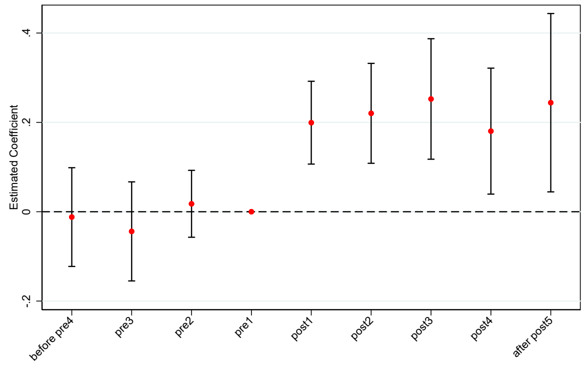

(b) open share by procurement sequence

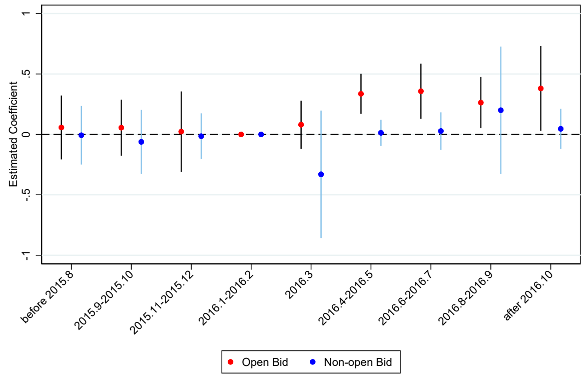

(c) frequency by procurement month

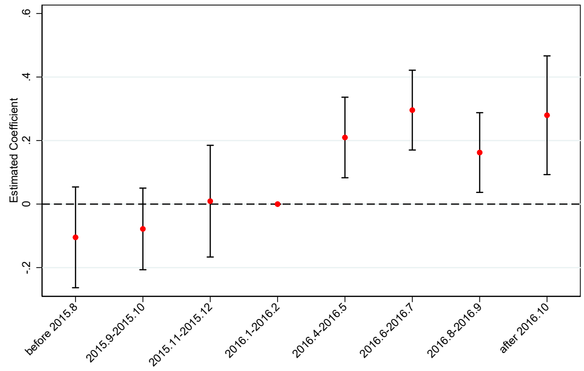

(d) open share by procurement month, 2SLS

Note: This figure plots the dynamic effects of the Weibo shock. Panel (a) shows the dynamics by calendar month. Panel (b) shows the dynamics by procurement sequence. Panel (c) shows the dynamics by calendar month with the outcome variable being the frequency of procurement, where we assign zero to city-month observations without the occurrence of procurement. Panel (d) shows the dynamics by calendar month using IV estimation. Note that in Panels (a) and (d), although the figure does not display the estimate for the event month for clearer exposition, this month is included in the estimation.

↑ *Return to Main Text: p.21*

Table 1: Summary Statistics of Main Variables

	count	p50	mean	sd	min	max
Procurement Variables						
Number of procured items	1630	2.000	6.640	13.849	1	200
Number of open-bid procurement	1630	1.000	4.292	12.225	0	200
Share of open-bid procurement	1630	0.667	0.530	0.473	0	1
Weibo Variables /1,000						
Total posts	17064	42.000	173.280	835.604	0	55422
Monitoring Posts(ML)	17064	2.000	24.148	242.368	0	15629
Monitoring Posts(DL)	17064	2.000	30.838	323.665	0	22599
Monitoring Posts(LLM)	17064	6.000	27.968	120.247	0	6804
Negative Sentiment Posts(ML)	17064	16.000	62.528	286.221	0	18138
Government-critical Monitoring Posts(ML)	17064	0.000	6.770	81.625	0	5834
Weibo Variables During Event Shock /1,000						
Monitoring Posts(ML)	237	186.000	540.899	1322.202	4	15629
Information Posts(LLM)	237	113.000	316.527	791.260	3	9816
Negative Sentiment Posts(LLM)	237	135.000	381.481	838.033	2	8622

Note: The unit of observation is city-month. The time frame is from January 2014 to December 2019. "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. "Total posts" is the count of posts referencing "vaccine" in Chinese. "Monitoring Posts(ML)" is the count of vaccine posts containing monitoring implications identified by the Support Vector Machines approach. "Monitoring Posts(DL)" is the count of vaccine posts containing monitoring implications identified by the deep learning (BERT) approach. "Monitoring Posts(LLM)" is the count of vaccine posts containing monitoring implications identified by large language models. "Negative Sentiment Posts" is the count of vaccine posts carrying negative sentiment identified by automated sentiment analysis. "Government-critical Monitoring Posts(ML)" is the count of vaccine posts that contain monitoring implications and also target government entities identified by a supervised machine learning approach.

↑ *Return to Main Text: p.12*

Table 2: Baseline Results of DID Estimation

Panel A: Share of open-bid procurement

	Overall		Category I		Category II and Supplement	
	(1) open share	(2) open share	(3) open share	(4) open share	(5) open share	(6) open share
WeiboShock × Event	0.099***	0.099***	−0.077	−0.042	0.196***	0.198***
	(0.029)	(0.032)	(0.053)	(0.065)	(0.035)	(0.040)
Observations	471	471	181	181	367	367
DV Mean	0.606	0.606	0.583	0.583	0.641	0.641
R^2	0.499	0.544	0.655	0.739	0.644	0.673
City Controls	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	Prefecture	Prefecture	Prefecture	Prefecture	Prefecture	Prefecture
Period FE	Year-month	YM+Sequence	Year-month	YM+Sequence	Year-month	YM+Sequence
Provincial Trend	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Frequency of procurement in balanced panels

	Overall		Category-I		Category-II and Supplement	
	(1) open	(2) nonopen	(3) open	(4) nonopen	(5) open	(6) nonopen
WeiboShock × Event	0.407***	0.151	0.105	0.082**	0.302***	0.069
	(0.148)	(0.152)	(0.070)	(0.035)	(0.114)	(0.140)
Observations	3100	3100	3100	3100	3100	3100
DV Mean	0.853	0.323	0.201	0.070	0.652	0.254
R^2	0.095	0.105	0.171	0.135	0.077	0.090
City Controls	Yes	Yes	Yes	Yes	Yes	Yes
City 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

Note: Observations are at the city-month level. The regression time window is from March 2015 to March 2017. The sample in Panel A includes only observations where procurement occurs. The sample in Panel B includes observations in all city months with zero being assigned to city-month without procurement. "Event" is a dummy for observations in and after the event month (March 2016). "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. The time-variant city characteristics 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 (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.19*

Table 3: Results of IV Estimation

	2SLS		First Stage	Reduced Form		
	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	Cat.II & Suppl.	WeiboShock × Event	Cat.II & Suppl.	Cat.II & Suppl.	Cat.II & Suppl.
WeiboShock × Event	0.104***	0.238***				
	(0.035)	(0.055)				
Cindex × Event			40.481***	9.625***	16.453*	19.199***
			(10.408)	(3.380)	(8.911)	(6.641)
Cindex(Forbes) × Event					−9.857	
					(10.240)	
Cindex(2011) × Event						−14.621
						(9.725)
Observations	471	367	367	367	367	367
K-P rk Wald F statistic	22.172	15.127				
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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. The interaction term "$WeiboShock_i \times Event_t$" is instrumented by the interaction term $Cindex_i \times Event_t$, where "Cindex" is measured by the city-level popularity of Weibo celebrities named in 2009. "Cindex(Forbes)" is measured by the city-level popularity of comparable celebrities named by Forbes in 2009. "Cindex(2011)" is measured by the city-level popularity of Weibo celebrities named in 2011. Full baseline controls include time-variant city characteristics as well as year-month fixed effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.22*

Table 4: Weibo Effect on Vaccine Procurement by Company Size and Locality

Panel A: Share in Terms of Procurement Frequency, Large Company

	Overall	Open Bid	Non-open Bid
	(1)	(2)	(3)
WeiboShock × Event	0.068**	0.118**	0.161***
	(0.029)	(0.049)	(0.049)
Observations	302	142	199
DV Mean	0.628	0.570	0.683
R^2	0.601	0.776	0.805
Full Baseline Controls	Yes	Yes	Yes

Panel B: Share in Terms of Procurement Frequency, Non-local Company

	Overall	Open Bid	Non-open Bid
	(1)	(2)	(3)
WeiboShock × Event	0.009	−0.033	0.105***
	(0.028)	(0.056)	(0.036)
Observations	302	142	199
DV Mean	0.840	0.795	0.867
R^2	0.521	0.680	0.655
Full Baseline Controls	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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 effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.25*

Table 5: Weibo Effect Through Information and Sentiment Channels

Panel A: share of open-bidding

	(1) Overall	(2) Overall	(3) Overall	(4) Overall	(5) Cat.II & Suppl.	(6) Cat.II & Suppl.	(7) Cat.II & Suppl.	(8) Cat.II & Suppl.
WeiboShock(Info=0,Senti=1) × Event	0.325***		0.335**		0.636***		0.540***	
	(0.093)		(0.137)		(0.111)		(0.180)	
WeiboShock(Info=1,Senti=0) × Event		0.305***	−0.015			0.607***	0.133	
		(0.109)	(0.108)			(0.172)	(0.146)	
WeiboShock(Info=1 or Senti=1) × Event				0.111***				0.225***
				(0.033)				(0.038)
Share of Sentiment Posts × Event				1.747				2.459
				(1.926)				(2.689)
Observations	471	471	471	471	367	367	367	367
DV Mean	0.606	0.606	0.606	0.606	0.641	0.641	0.641	0.641
R^2	0.500	0.496	0.500	0.500	0.645	0.638	0.645	0.646
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: I(share of open-bidding=1)

	(1) Overall	(2) Overall	(3) Overall	(4) Overall	(5) Cat.II & Suppl.	(6) Cat.II & Suppl.	(7) Cat.II & Suppl.	(8) Cat.II & Suppl.
WeiboShock(Info=0,Senti=1) × Event	0.264**		0.445***		0.612***		0.751***	
	(0.103)		(0.135)		(0.131)		(0.188)	
WeiboShock(Info=1,Senti=0) × Event		0.159	−0.265**			0.467**	−0.192	
		(0.130)	(0.115)			(0.224)	(0.161)	
WeiboShock(Info=1 or Senti=1) × Event				0.079**				0.206***
				(0.036)				(0.047)
Share of Sentiment Posts × Event				2.559				3.854
				(2.144)				(3.052)
Observations	471	471	471	471	367	367	367	367
DV Mean	0.564	0.564	0.564	0.564	0.618	0.618	0.618	0.618
R^2	0.478	0.473	0.479	0.478	0.589	0.577	0.590	0.590
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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(Info=0,Senti=1)" is measured by the difference between the per-capita number of non-informative but sentimental posts published by users in a city in the event month and the average number of posts published three months before the event. "WeiboShock(Info=1,Senti=0)" is similary defined based on informative but non-sentimental posts. "WeiboShock(Info=1 or Senti=1)" is similarly defined based on either informative or sentimental posts. "Share of Senti" is measured by the share of sentimental posts over "WeiboShock(Info=1 or Senti=1)". Full baseline controls include time-variant city characteristics as well as year-month fixed effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.28*

Table 6: Evidence on Mechanisms: Heterogeneous Treatment Effects

	Official Experience		Information Environment		Inspection		Career Concerns		Public Trust (CGSS)	
	(1) Overall	(2) Cat.II & Suppl.	(3) Overall	(4) Cat.II & Suppl.	(5) Overall	(6) Cat.II & Suppl.	(7) Overall	(8) Cat.II & Suppl.	(9) Overall	(10) Cat.II & Suppl.
WeiboShock × Event	0.045	0.142***	0.092***	0.194***	0.102***	0.139***	0.157***	0.336***	0.141***	0.134***
	(0.051)	(0.051)	(0.029)	(0.038)	(0.036)	(0.037)	(0.042)	(0.059)	(0.045)	(0.047)
Lack Local Experience × Event	−0.130	−0.137								
	(0.204)	(0.232)								
WeiboShock × Event × Lack Local Experience	0.121**	0.139**								
	(0.058)	(0.065)								
Deleted Posts × Event			−0.353	−0.479*						
			(0.265)	(0.282)						
WeiboShock × Event × Deleted Posts			0.168	0.474***						
			(0.222)	(0.160)						
Inspection × Event					−0.081	−0.707***				
					(0.290)	(0.265)				
WeiboShock × Event × Inspection					−0.005	0.165***				
					(0.065)	(0.057)				
Pre-retirement × Event							0.023	0.293		
							(0.205)	(0.252)		
WeiboShock × Event × Pre-retirement							−0.078	−0.201***		
							(0.049)	(0.068)		
Low Trust × Event									−0.564*	−0.856***
									(0.314)	(0.254)
WeiboShock × Event × Low Trust									0.015	0.106**
									(0.047)	(0.045)
Observations	471	367	471	367	471	367	471	367	229	175
DV Mean	0.606	0.606	0.606	0.641	0.641	0.641	0.641	0.641	0.641	0.641
R^2	0.505	0.651	0.500	0.652	0.499	0.652	0.502	0.654	0.543	0.736
F Test(row1+trippleD=0)	0.000	0.000	0.251	0.000	0.075	0.000	0.022	0.006	0.002	0.000
Other Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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. "Lack Local Experience" is a dummy that equals 1 if the number of years that the mayors in a city worked in the local administration is below the median level for the mayors in all cities. "Deleted Posts" is a dummy that equals 1 if the proportion of posts being deleted is above the mean level of all provinces. "Low Trust" is a dummy that equals 1 if the proportion of individuals who expressed distrust in local officials in a city during the 2010 and 2012 CGSS surveys exceeds the median level among all cities. "Inspection" is a dummy that equals 1 if the number of anti-corruption allegations in a province within three years before the event is above the median level among all provinces. "Pre-retirement" is a dummy that equals 1 if the age of the city mayor is above 54 and 0 otherwise. Full baseline controls include time-variant city characteristics as well as year-month fixed effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text:* p.28

Online Appendices (Not for Publication)

A Additional Empirical Results

Figure A1: Deletion of Weibo Posts by Subject

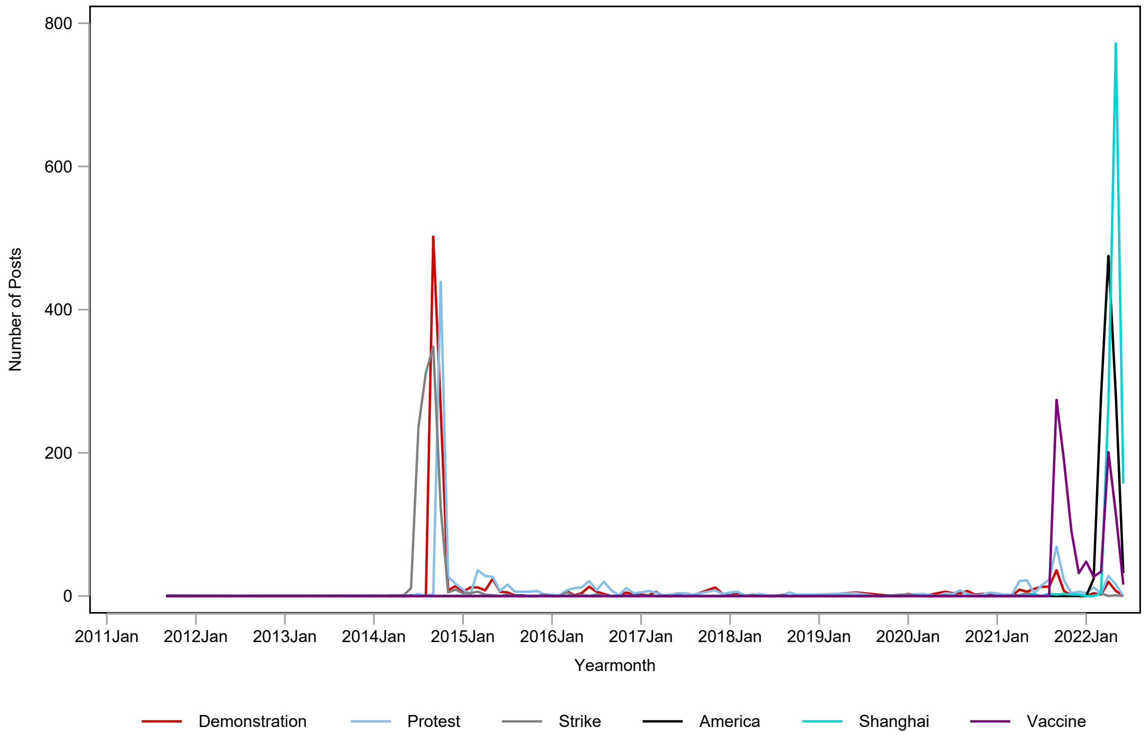

Note: This figure plots the numbers of posts referring different subject keywords that were deleted from Weibo over time. The deletion count is based on the number of posts republished on www.freeweibo.com, a website that collects and republishes posts deleted from the official Weibo website.

↑ *Return to Main Text: p.13*

Figure A2: Time Trends of Weibo Posts on Vaccines

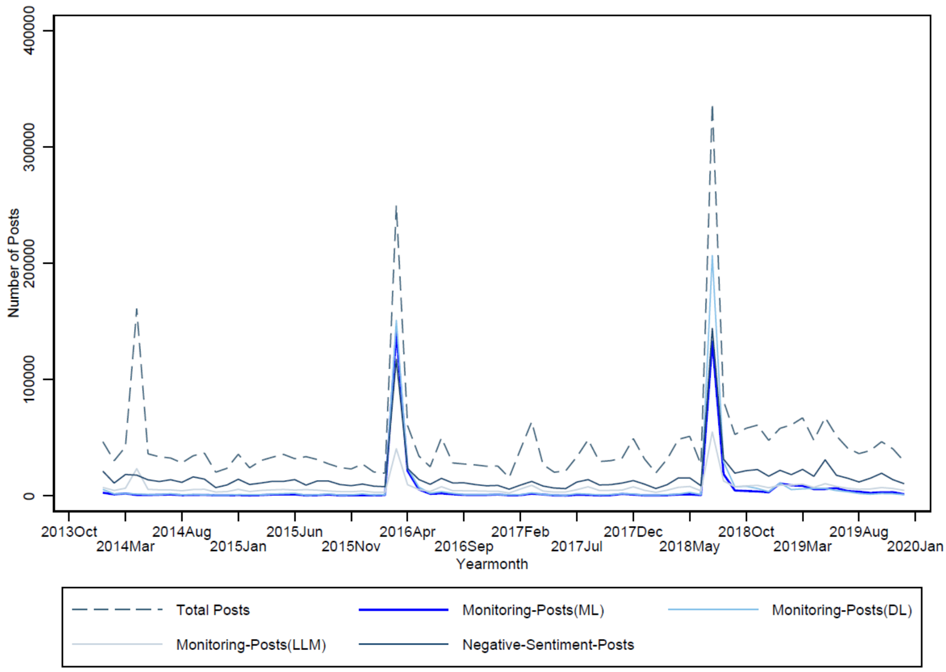

Note: This figure plots the time series of the numbers of all vaccine-related Weibo posts, posts with negative sentiment, "monitoring" posts identified by three different approaches: machine-learning (SVM), deep-learning (BERT), and a large language model (Qwen1.5:14b), respectively.

↑ *Return to Main Text: p.14*

Figure A3: Information Eruption during the 2016 Event

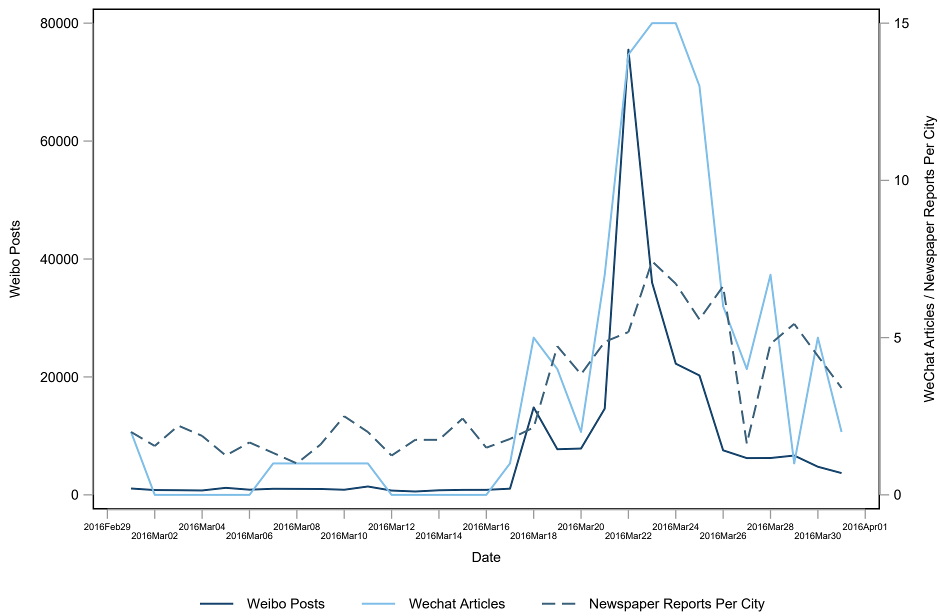

Note: This figure plots the numbers of Weibo posts, WeChat articles and newspaper reports on vaccines at the daily frequency during the event month (March 2016).

↑ *Return to Main Text: p.24*

Figure A4: Changes in the Landscape of Monitoring Posts in 2016

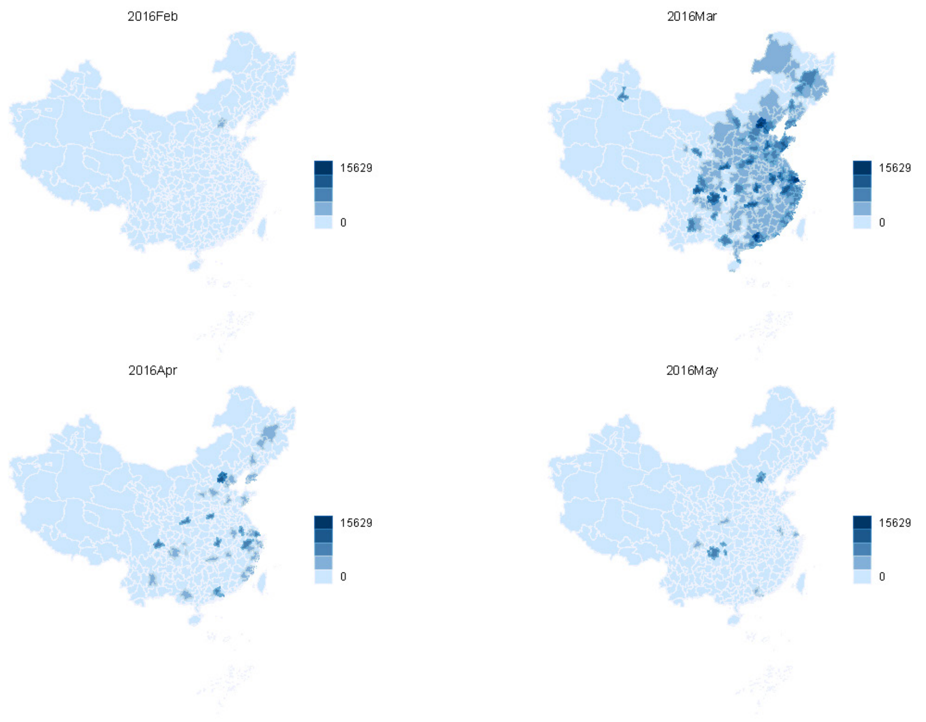

Note: This figure depicts the regional variations in the number of monitoring posts identified by the machine-learning approach (SVM) for each month from February 2016 to May 2016.

↑ *Return to Main Text: p.15*

Figure A5: Registration Time and Influence of Active Weibo Users

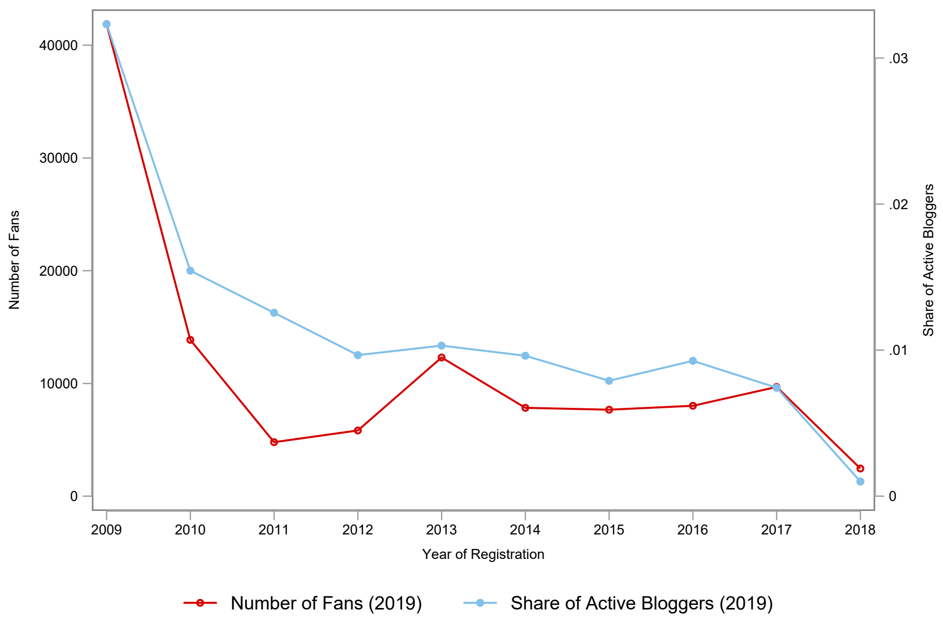

Note: This figure plots the average number of followers of Weibo users and the share of active bloggers against the registration time. Active bloggers are defined as those whose total number of posts falls within the top 1% among all Weibo users.

↑ *Return to Main Text: p.18*

Figure A6: Correlation between IV and the 2016 Weibo Shock

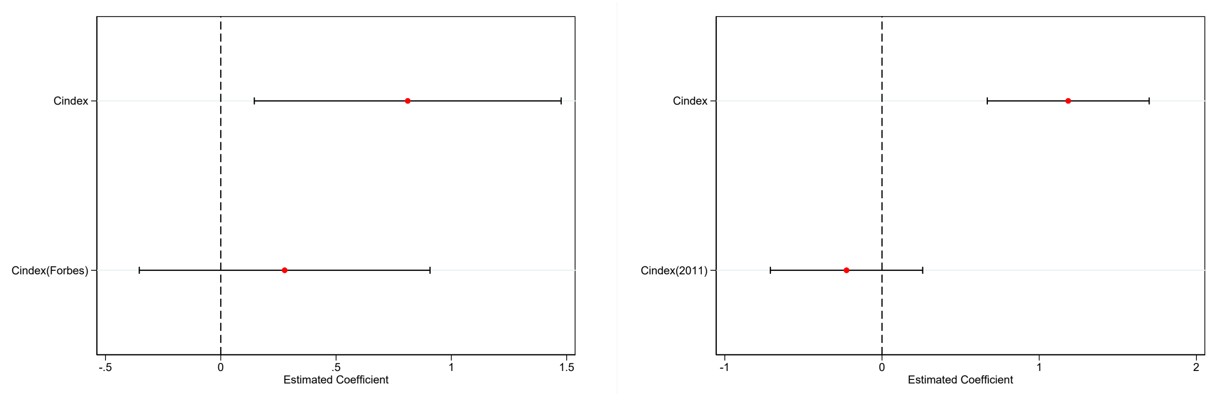

(a) Placebo: Forbes Celebrities 2009

(b) Placebo: Weibo Celebrities 2011

Note: This figure plots the estimated coefficient by regressing the Weibo shock in 2016 on the actual instrument ("Cindex") and the placebo instruments ("Cindex(Forbes)" or "Cindex(2011)") at the city level. "Cindex" is measured by the city-level popularity of Weibo celebrities named in 2009. "Cindex(Forbes)" is measured by the city-level popularity of comparable celebrities named by Forbes in 2009 but did not join Weibo during that year. "Cindex(2011)" is measured by the city-level popularity of Weibo celebrities named in 2011 and joining Weibo after 2010. The regression controls for city characteristics in 2008 include the (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, and the share of the secondary industry in GDP.

↑ *Return to Main Text: p.23*

Figure A7: Celebrity Popularity and the Date of the First Post

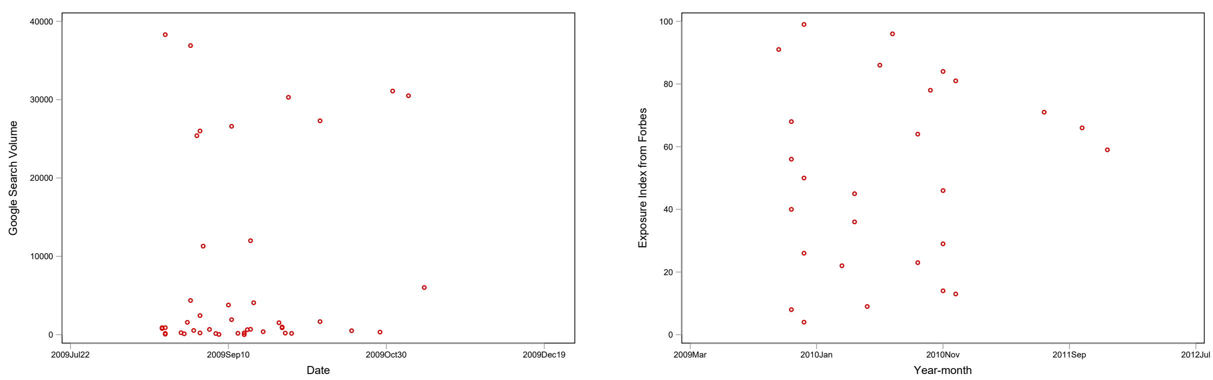

(a) 2009 Weibo celebrity

(b) 2012 Weibo celebrity

Note: This figure depicts the relationship between the popularity of Weibo celebrities and the date of their first post. Panel (a) plots the Google search volume in 2008 for a Weibo celebrity who registered by the end of 2009 against the date when they published their first post. Panel (b) plots the exposure index according to the Forbes 2010 ranking for a Weibo celebrity who appeared on the Weibo 2012 Celebrity List against the date when they published their first post. The date of the first post is tracked by the historical records from each celebrity's Weibo account.

↑ *Return to Main Text: p.18*

Figure A8: Results of DID Estimation with Advanced Measures of Weibo Shock

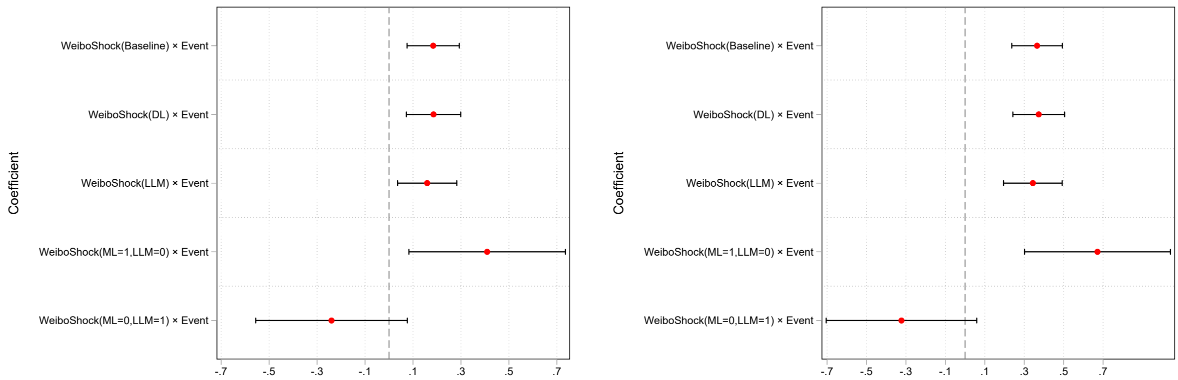

(a) Overall (b) Category II and Supplement

Note: This figure plots the estimates of the Weibo effect on the share of open-bidding. Observations are at the city-month level. The regression time window 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(DL)" is measured by the difference between the per-capita number of deel-learning-classified 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(LLM)" is similarly defined but using large language model. "WeiboShock(ML=1,LLM=0)" and "WeiboShock(ML=0,LLM=1)" are similarly defined but measures exclusive information shocks identified by SVM or LLM, as indicated in parentheses. The time-variant city characteristics 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 (in parentheses) are clustered by city.

↑ *Return to Main Text: p.22*

Figure A9: Dynamic Effects of the IV Estimation: Reduced Form

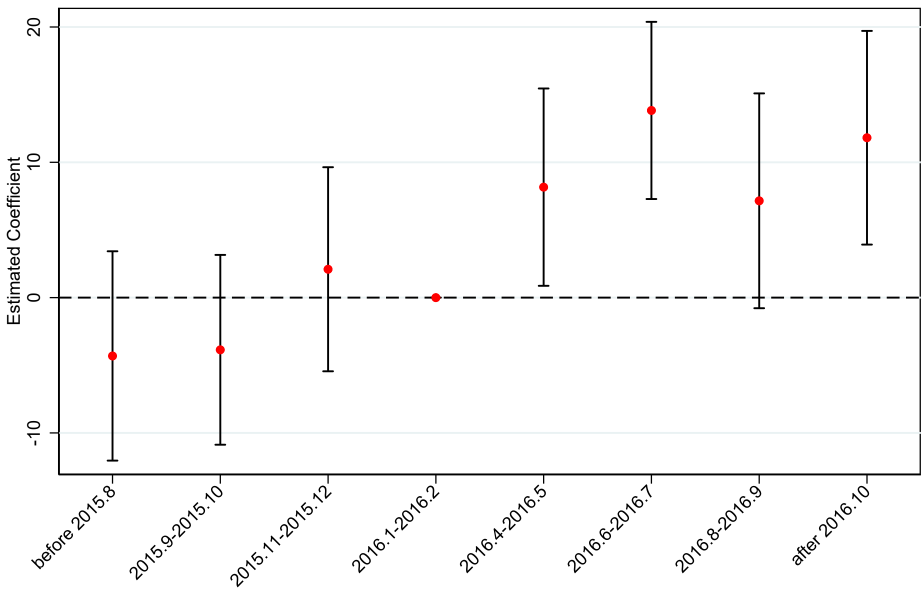

Note: This figure plots the dynamic effects of the Weibo shock using IV as independent variables. The outcome variable being the frequency of procurement, where we assign zero to city-month observations without the occurrence of procurement. Note that the figure does not display the estimate for the event month for clearer exposition, this month is included in the estimation.

↑ *Return to Main Text: p.23*

Figure A10: Share of Open-bid Procurement over Time

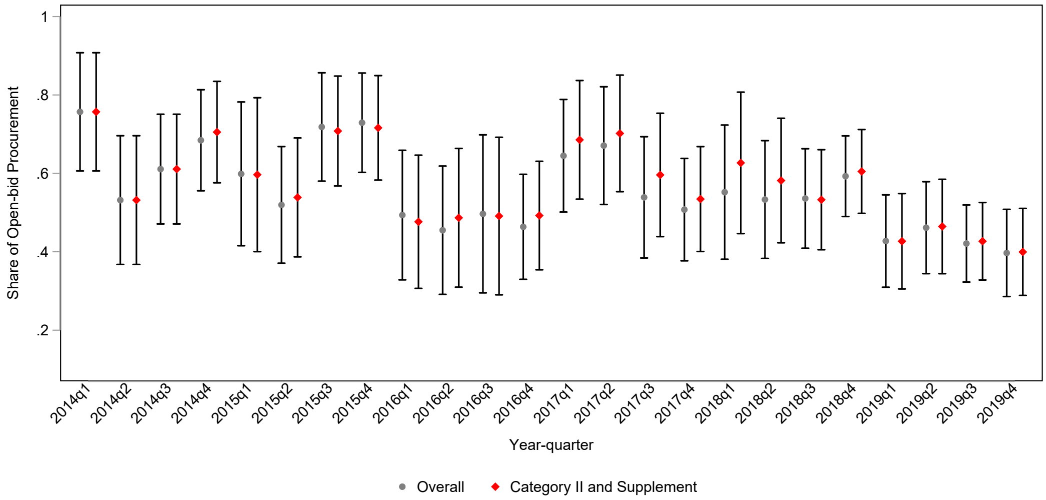

Note: This figure plots the quarterly time trend of the share of open-bid procurement at the national level between January 2014 and December 2019, showing the mean and standard deviation in each quarter.

↑ *Return to Main Text: p.31*

Figure A11: Effects on the Incidence of Infectious Diseases

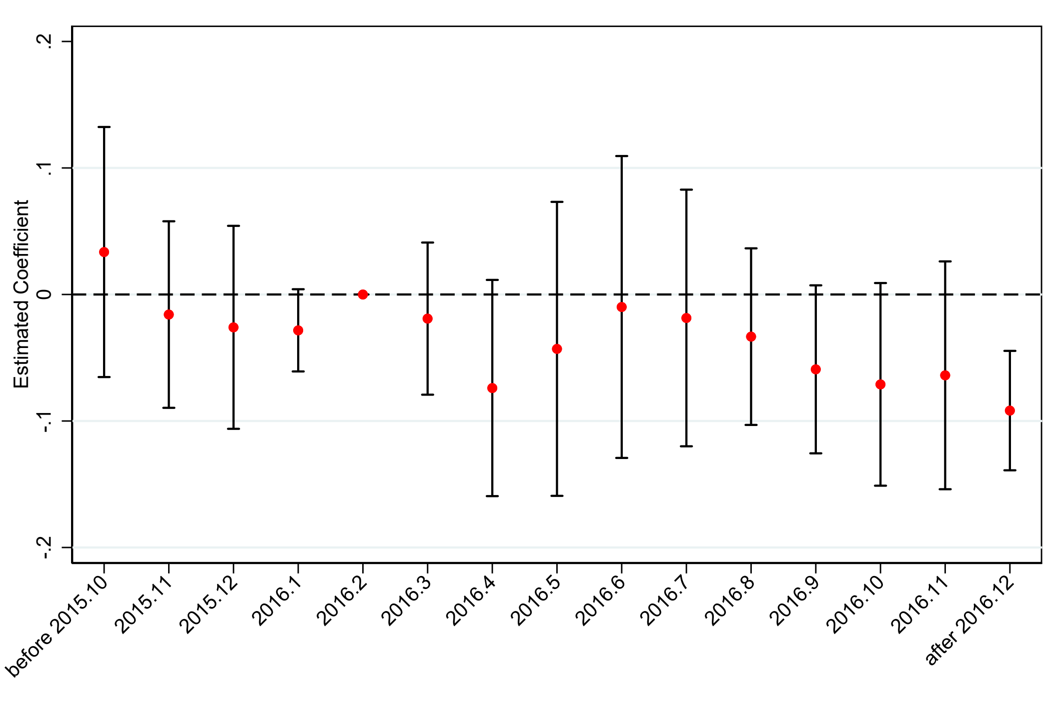

Note: This figure plots the dynamic estimates of the Weibo effect on the logarithm of the total number of infectious disease incidents. The observations are at the province-month level. The regression time window is from March 2015 to March 2018. Each dot represents the estimated coefficient of the interaction term between "WeiboShock" and a set of dummies that equal 1 if an observation falls within the specific time period and 0 otherwise.

↑ *Return to Main Text: p.26*

Table A1: Correlation between IV and Weibo Expansion

	Weibo penetration				Weibo KOL density
	(1)	(2)	(3)	(4)	(5)
	2010	2011	2012	2013	Followers
Cindex (std)	0.533***	2.514***	4.791***	5.109***	0.098***
	(0.162)	(0.721)	(1.248)	(1.320)	(0.023)
Observations	206	206	206	206	73
R^2	0.703	0.801	0.866	0.853	0.859
City Controls	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city level. This table shows the correlation between our instrument and two measures of Weibo activity: early Weibo penetration and the density of impactful users. In columns (1)–(4), the dependent variable is Weibo penetration, defined as the per-capita number of Weibo posts in each city-year from 2010 to 2013. In columns (5), the dependent variable is the density of impactful Weibo users in 2012, measured by the per-capita number of users with more than 1 million followers. The instrument is standardized here to facilitate coefficient interpretation. All regressions control for city-level characteristics in the corresponding year and include province fixed effects. Standard errors are clustered at the province level.

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

↑ *Return to Main Text: p.17*

Table A2: Exclusion Restrictions (Only through Weibo platform)

	(1)	(2)	(3)
	Weibo	Newspaper	Online Search
Cindex	33.478***	−69.213	−3.494
	(10.141)	(119.867)	(2.564)
Observations	124	24	124
R^2	0.872	0.318	0.667
City Controls	Yes	Yes	Yes

Note: Observations are at the city level. This table reports the impact of our IV on information outbreaks across three different information channels: Weibo, newspapers, and the Internet. For each channel, the dependent variable is constructed as the difference between two components: (i) the value of the corresponding metric in the event month, and (ii) the average value of the same metric over the three months preceding the event. Specifically, the three metrics are the per-capita number of monitoring posts on Weibo, the number of vaccine-related news articles per media outlet, and the Baidu search intensity for vaccine-related keywords. All regressions control for city-level characteristics in 2016. Standard errors are clustered at the province level.

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

↑ *Return to Main Text: p.18*

Table A3: Exclusion Restrictions

	Exclusion Restrictions							
	(1) %2ind/GDP	(2) Log govexp	(3) Log GDPpc	(4) Log pop den	(5) Log fdipc	(6) Log internet	(7) Log mobile	(8) Log land-line
Cindex × Event	−5.038	0.260	−0.381	0.040	−0.242	−0.000	−0.000	−0.000
	(4.376)	(0.257)	(0.367)	(0.059)	(0.191)	(0.000)	(0.000)	(0.000)
Observations	2975	2975	2975	2975	2975	2975	2975	2975
City FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Provincial Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window is from March 2015 to March 2017. The sample includes only observations where procurement occurs. The dependent variable in each column is the share of the secondary industry in GDP, (log) government expenditure per capita, (log) GDP per capita, (log) population density, (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, respectively. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.18*

Table A4: Effect with Discrete Measure of Weibo Shock

	Overall		Category-I		Category-II and Supplement	
	(1) open share	(2) open share	(3) open share	(4) open share	(5) open share	(6) open share
WeiboShockDmy × Event	0.311***	0.342*	0.404**	−0.339	0.334***	0.485*
	(0.089)	(0.175)	(0.162)	(0.346)	(0.100)	(0.252)
WeiboShockDmy	−0.116		−0.307		−0.079	
	(0.093)		(0.193)		(0.100)	
Event	−0.337***		−0.477***		−0.321***	
	(0.068)		(0.140)		(0.074)	
Observations	471	471	181	181	367	367
DV Mean	0.606	0.606	0.583	0.583	0.641	0.641
R^2	0.088	0.495	0.138	0.655	0.107	0.625
City Controls	Yes	Yes	Yes	Yes	Yes	Yes
City FE	No	Yes	No	Yes	No	Yes
Year-Month FE	No	Yes	No	Yes	No	Yes
Provincial Time Trend	No	Yes	No	Yes	No	Yes

Note: Observations are at the city-month level. The regression time window 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. "WeiboShockDmy" is a dummy variable indicating that the Weibo Shock (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) is above the mean level among all cities. The time-variant city characteristics include the (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 (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.22*

Table A5: Results of Alternative IV Estimation

	Augmented Celebrity IV		Student Inflow from 4 cities (2009)		All Posts (2009.Aug-Oct)	
	(1)	(2)	(3)	(4)	(5)	(6)
	Overall	Cat.II & Suppl.	Overall	Cat.II & Suppl.	Overall	Cat.II & Suppl.
WeiboShock × Event	0.110***	0.213***	0.115***	0.278***	0.095	0.188**
	(0.034)	(0.050)	(0.043)	(0.057)	(0.075)	(0.082)
Observations	456	359	466	362	471	367
K-P rk Wald F statistic	18.885	13.996	35.007	31.431	167.714	256.251
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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. The interaction term "$WeiboShock_i \times Event_t$" is instrumented by the interaction term between the IV variable and the event dummy. In columns (1)-(2), the IV variable is the interaction term between the popularity of Weibo celebrities named in 2009 and the social media connections to the city of the scandal. In columns (3)-(4), the IV variable is the number of college freshmen who registered at a college in a particular city but came from four cities with the highest level of Weibo penetration in 2009 ("Beijing", "Shanghai", "Guangzhou", and "Shenzhen"). In columns (5)-(6), the IV variable is the residuals obtained from regressing Weibo penetration in Aug.-Oct. 2009 on the (log) numbers of college students per capita, (log) number of internet users per capita, (log) number of mobile phone users per capita, and (log) number of land-line users per capita in 2009. Full baseline controls include time-variant city characteristics as well as year-month fixed effects, city fixed effects, and provincial time trends. In the first four columns, standard errors (in parentheses) are clustered by city. In the last two columns, we bootstrap the standard errors of the estimates throughout the entire estimation.

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

↑ *Return to Main Text: p.23*

Table A6: Weibo Effect on Capitulation

	Overall		Overall(units>1)		Cat.II & Suppl.		Cat.II & Suppl.(units>1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
WeiboShock × Event	0.072**	0.076*	0.035	0.037	0.177***	0.259***	0.112*	0.273*
	(0.036)	(0.039)	(0.045)	(0.053)	(0.049)	(0.065)	(0.065)	(0.159)
Observations	471	471	290	290	367	367	219	219
DV Mean	0.564	0.564	0.564	0.564	0.618	0.618	0.618	0.618
R^2	0.476	0.050	0.637	0.079	0.586	0.075	0.715	0.105
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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 effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.24*

Table A7: Effect on the Duration of Procurement

	Overall	Open Bid	Non-open Bid
	(1)	(2)	(3)
	log(days)	log(days)	log(days)
WeiboShock × Event	0.324***	0.362***	0.123
	(0.079)	(0.097)	(0.094)
Observations	471	471	471
DV Mean	2.498	1.900	0.717
R^2	0.545	0.502	0.516
Full Baseline Controls	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window is from March 2015 to March 2017. The duration of open-bid procurement is measured by the number of days between the announcement date of procurement and the expiration date of the announcement notice, which usually coincides with the first meeting with potential bidders. "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 effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text:* p.25

Table A8: Effect on the Incidence of Infectious Diseases

	Overall	Prevalent Disease	Child Disease
	(1)	(2)	(3)
WeiboShock × Event	−0.034	−0.080***	0.026
	(0.020)	(0.021)	(0.016)
Observations	1110	1110	1110
DV Mean	5.553	4.910	4.514
R^2	0.787	0.886	0.835
Province Controls	Yes	Yes	Yes
Province FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Provincial Trend	Yes	Yes	Yes

Note: Observations are at the province-month level. The regression time window is from March 2015 to March 2018. "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 province in the event month and the average number of posts published three months before the event. Full baseline controls include time-variant province characteristics as well as year-month fixed effects, province fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by province.

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

↑ *Return to Main Text:* p.26

Table A9: 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	3.389***	−0.108**	0.187	0.228	2.348***	3.147***
	(0.692)	(0.049)	(0.150)	(0.191)	(0.456)	(0.526)
WeiboShock(Daily) × After 2016.3.22	1.214***	0.244**	0.211**	0.590***	0.304***	1.113***
	(0.280)	(0.101)	(0.092)	(0.121)	(0.106)	(0.290)
Observations	6059	6059	6059	6059	6059	6059
DV Mean	0.123	0.057	0.042	0.026	0.012	0.005
R^2	0.731	0.455	0.465	0.531	0.625	0.730
City FE	Yes	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-date level. The regression time window is within a short time frame 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. This division is based on the fact that March 22 marked the date when the National FDA issued an official announcement regarding the scandal. Therefore, government posts published during the first time window reflect a local government's expressed attitude prior to the central government's intervention. Note that the government posts only account for a tiny share (0.7%) of the monitoring posts, and the concern of reverse causality is mitigated. Control variables include city fixed effects and date fixed effects. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.27*

Table A10: Results of Criminal Cases

	Numbers: scandal-related			Shares: scandal-related			Numbers: Non-scandal-related		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Cases	Crim.	Crim. gov	% cases	% crim.	% crim. gov	Cases	Crim.	Crim. gov
WeiboShock	0.124**	0.183**	0.047**	0.047***	0.047***	0.022**	0.017	0.014	0.016
	(0.049)	(0.077)	(0.023)	(0.015)	(0.015)	(0.009)	(0.040)	(0.050)	(0.030)
Observations	924	924	924	924	924	924	924	924	924
R^2	0.196	0.176	0.141	0.220	0.219	0.128	0.124	0.126	0.134
City Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Provincial Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-year level. The regression time window is from 2016 to 2019. "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. The dependent variables in columns (1)-(3) are the number of cases, criminals and implicated officials related to the scandal, while columns (4) -(6) show the proportion of the outcomes related to the scandal. Columns (7)–(9) contain the dependent variables for non-scandal-related outcomes. Standard errors (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.27*

Table A11: Effect of Weibo Shock during the Vaccine Scandal in 2018

	Overall		Category-I		Category-II and Supplement	
	(1) open share	(2) open share	(3) open share	(4) open share	(5) open share	(6) open share
WeiboShock × Event	–0014	0.006	0.068	0.036	–0028	–0011
	(0.021)	(0.029)	(0.124)	(0.110)	(0.022)	(0.023)
Observations	564	564	166	166	484	484
DV Mean	0.512	0.512	0.439	0.439	0.568	0.568
R^2	0.210	0.453	0.494	0.520	0.266	0.545
Regional FE	Province	Prefecture	Province	Prefecture	Province	Prefecture
City 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 regression time window is from November 2017 to July 2019. The starting month is not extended to an earlier period to avoid overlapping with a precedent mild information outbreak. The ending month is chosen for practical reasons to ensure enough observations in the panel regression with city fixed effects. "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. The time-variant city characteristics include the (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 (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.31*

Table A12: Correlation between Weibo Posts and Government Procurement

	Entire Sample		Excluding Periods of Information Outbreaks	
	(1)	(2)	(3)	(4)
	share of open-bid	share of open-bid	share of open-bid	share of open-bid
L1.#posts pca	0.024		−0061	
	(0.015)		(0.066)	
L2.#posts pca	0.052**		0.253***	
	(0.025)		(0.048)	
L3.#posts pca	−0018		0.102	
	(0.027)		(0.093)	
L1.#govt-specific posts pca		−0002		−0125
		(0.020)		(0.134)
L2.#govt-specific posts pca		0.037**		0.208
		(0.016)		(0.142)
L3.#govt-specific posts pca		0.022		0.170*
		(0.028)		(0.086)
L1.#firm-specific posts pca		0.012		0.281
		(0.017)		(0.192)
L2.#firm-specific posts pca		−0019		−0241
		(0.014)		(0.197)
L3.#firm-specific posts pca		−0028*		−0192
		(0.015)		(0.127)
Observations	1523	1523	1237	1237
R^2	0.316	0.317	0.347	0.347
City Controls	Yes	Yes	Yes	Yes
City 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 regression time window is from January 2014 to December 2019, with all observations included in the first two columns (entire sample) and with the Weibo-information outbreak periods of March-August 2016 and July-December 2018 being excluded from the sample in the last two columns. "posts.pca" is the per capita number of posts that referencing vaccine and containing a negative sentiment. "L1.", "L2." and "L3." indicate one-, two-, three- period lags, respectively. "#govt-specific posts pca" is the per-capita number of negative-sentiment posts referencing vaccine and specifically targeting government. "#firm-specific posts pca" is the per-capita number of negative-sentiment posts referencing vaccine and specifically targeting firms. The time-variant city characteristics include the (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 (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.31*

Table A13: Effect of Information Monitoring Technology on Procurement Transparency

	All		Legal	Infrastructure	Agriculture	Administrative	Education	Health
	(1) open share	(2) open share	(3) open share	(4) open share	(5) open share	(6) open share	(7) open share	(8) open share
ICT Adoption	−0.007	−0.046*	−0.005	−0.028	0.002	−0.036	−0.042	−0.027
	(0.015)	(0.024)	(0.035)	(0.024)	(0.030)	(0.026)	(0.031)	(0.022)
ICT Adoption × Complaints		0.013	0.045*	0.013	0.065**	0.009	0.011	0.030
		(0.020)	(0.027)	(0.018)	(0.026)	(0.021)	(0.026)	(0.024)
ICT Adoption × KOL-index		0.295***	0.297**	0.248***	−0.041	0.154	0.305***	0.260***
		(0.087)	(0.131)	(0.088)	(0.087)	(0.105)	(0.093)	(0.082)
Observations	3600	2236	1682	2045	1783	2061	2054	2096
R^2	0.435	0.434	0.398	0.382	0.394	0.417	0.379	0.404
City Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-year level. The dependent variable is the share of open-bid procurement. The regression time window is from 2007 to 2020. "ICT Adoption" is a dummy that equals 1 if a city government has ever procured ICT systems to monitor information flows."Complaints" is measured by the number of complaining messages filed with local leader's offices on the portal of the Local Government Leaders Forum. "KOL-index" is measured by the aggregated Baidu search index of early right-leaning Weibo KOLs in 2009. The time-variant city characteristics 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 (in parentheses) are clustered by city.

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

↑ *Return to Main Text: p.32*

B Robustness Checks

In this section, we presents details of several robustness checks for the baseline estimation.

Event effect The vaccine scandal, once revealed to the public, was likely to provide information about the state of the world, and local governments might have responded even without being influenced by public opinion. To address this, we use the fact that the vaccine distributor responsible for the scandal sold defective vaccines in 18 provinces. Therefore, we include an interaction term between the event timing indicator and a dummy indicating whether a city is located in one of the affected provinces. The results, as shown in Column (1) to Column (4) of Panel A of Table B1, reveal that the coefficient of this newly added interaction term is statistically insignificant, while the estimated Weibo effect remains barely changed. Note that this does not imply that the event has no effect on vaccine procurement. Rather, it means that the event itself does not confound the social media effect, which suggests a mechanism other than pure information provision as we will investigate in the next section.

Policy effect As noted in Section 4.2.1, subsequent to the event, the central government made some policy adjustments aiming to enforce stricter regulations regarding vaccine procurement. If cities with stricter implementation of the new regulations coincidentally experienced a stronger Weibo shock, our estimated effects could be spurious. To address this concern, we consider the timing of the implementation of the new regulations, which varied significantly across provinces.[23] Most cities did not implement the central government's new regulation until six months after the event. As demonstrated in Panel A of Table B1 (Columns (5)-(8)), the interaction between the Weibo shock and the timing of policy implementation has a weak effect on the share of open-bid procurement but barely changes the estimate of the interaction term "$WeiboShock \times Event$." This suggests that the Weibo shock may influence local governments' implementation of the central government's ad hoc policy adjustment, which in turn affects the adoption of the procurement format.

Scope of the effect The social-media effect could potentially extend to other public procurement activities if it led local governments to enhance overall administrative transparency. To test this, we examine the effect of the Weibo shock on the procurement of two alternative types of products: (1) health-related products excluding vaccines and (2) non-health products. Figure B1 plots the dynamic effects of the Weibo shock on the share of open-bid procurement for them. We observe no significant effects for either product type, suggesting that the Weibo effect is limited in scope. It also helps alleviate the concern that the Chinese government may have engineered an administrative reform

23. This information was collected from the official websites of provincial Food and Drug Administration.

coinciding with the timing of the Weibo shock, which could have triggered greater social media reactions for more compliant city governments.

Other information channels Local governments may respond to information from other sources, such as traditional media, other social media, and the internet. In Figure A3, we illustrate that WeChat discussion and newspaper coverage did not precede the outbreak of Weibo posts about vaccine safety. Panel B of Table B1 examines the impact of other informational channels on the estimated Weibo effect. Columns (1) and (2) report the results from the baseline DID regression with the "$WeiboShock$" variable being replaced by the "$Newspaper\ Shock$" variable, which measures the increase of the number of news articles covering vaccine issues per media outlet in a city in the event month relative to its average level in the three months prior to the event. The newspaper effects are small and statistically insignificant. In Columns (3) and (4), we run a horse race between the Weibo and newspaper effects by including both interaction terms in the regression. The coefficients of the newspaper effect become even smaller than their counterparts in Columns (1) and (2) while the coefficients of the Weibo effect remain almost the same as those in our baseline estimation.

The above results suggest that newspaper coverage did not affect government behavior in vaccine procurement.[24] This is likely for two reasons. First, after 2014, the Chinese government tightened the control of the media (Qin, Strömberg, and Wu 2024). Second, the newspaper market was severely eroded by social media, and the loss of readership and advertising revenues further crippled the monitoring function of newspapers.

Another potential information source is news portals on the internet. We add to our baseline regression a term indicating the interaction of the event timing with the change in the search intensity for vaccines-related keywords on Baidu, the Chinese equivalent to Google. Columns (5) and (6) of Panel B of Table B1 show that the effect of this additional term is negligible and hardly affects the estimate of the Weibo effect.

Information spillover We have shown that the government of a locality responds to social media opinions originated from that locality. This makes sense provided that local governments care more about the opinions of citizens in their administered regions and have the ability to gauge local information by using modern information technology. Nevertheless, information rapidly spreads over social media, and public opinions in one region may also affect the responses of governments in other regions. To examine this potential information-spillover effect, we construct two confounding variables, "$WeiboShock(prov.)$" and "$WeiboShock(capital)$," which are measured by the number

24. A caveat regarding these regression results is that the newspaper sample is considerably smaller than the Weibo sample because digital archives are available only for newspapers in large cities. To assess the potential bias due to such sample selection, we look into the manner of news coverage by newspapers in cities experiencing the strongest and weakest Weibo shocks within the newspaper-sample cities. Regardless of the magnitude of the Weibo shock, all newspapers cover vaccine issues using similar (official) narratives, and the intensity of coverage, measured by the number of reports, is close. Therefore, we suspect that newspapers from smaller cities cover the issue in a similar way.

of monitoring posts originating from the entire province (excluding the procuring city itself) and from the provincial capital city, respectively. We include these two competing variables in our baseline regression. Columns (7) and (8) of Panel B show that the inclusion of these two newly added interaction terms, if anything, strengthens the effect of the Weibo shock from the procuring city. These results suggest that information spillover across regions is not a major concern for our estimate of the Weibo effect.

Table B1: Robustness Checks

Panel A: Event Effect and Policy Effect

	Event Effect				Policy Effect			
	(1) Overall	(2) Cat.II & Suppl.	(3) Overall	(4) Cat.II & Suppl.	(5) Overall	(6) Cat.II & Suppl.	(7) Overall	(8) Cat.II & Suppl.
WeiboShock × Event			0.100***	0.196***			0.102***	0.193***
			(0.028)	(0.037)			(0.030)	(0.036)
Province Involved × Event	0.013	0.187	–0.061	–0.001				
	(0.218)	(0.223)	(0.211)	(0.239)				
WeiboShock × Vaccine Policy					0.036**	0.041**	0.040**	0.032
					(0.017)	(0.020)	(0.017)	(0.024)
Observations	471	367	471	367	471	367	471	367
DV Mean	0.606	0.641	0.606	0.641	0.606	0.641	0.606	0.641
R^2	0.489	0.617	0.499	0.644	0.491	0.619	0.502	0.646
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Other Information Channels and Spillover Effect

	Newspaper				Online Search		Spillover Effect	
	(1) Overall	(2) Cat.II & Suppl.	(3) Overall	(4) Cat.II & Suppl.	(5) Overall	(6) Cat.II & Suppl.	(7) Overall	(8) Cat.II & Suppl.
WeiboShock × Event			0.127**	0.172*	0.210**	0.296***	0.163***	0.309***
			(0.054)	(0.090)	(0.086)	(0.057)	(0.060)	(0.070)
Newspaper Shock × Event	0.019	0.035	0.008	0.021				
	(0.024)	(0.038)	(0.022)	(0.041)				
Search Index Shock × Event					–0.042	0.014		
					(0.299)	(0.244)		
WeiboShock(prov.) × Event							–0.079	0.340*
							(0.198)	(0.182)
WeiboShock(capital) × Event							–0.054	–0.300*
							(0.143)	(0.161)
Observations	180	126	180	126	336	264	426	329
DV Mean	0.606	0.641	0.606	0.641	0.606	0.641	0.606	0.641
R^2	0.421	0.657	0.441	0.676	0.608	0.747	0.521	0.671
Full Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Observations are at the city-month level. The regression time window 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. In Panel A, "Province Involved" is a dummy that equals 1 if a city is located in a province that was affected by the company involved in the vaccine scandal. "Vaccine Policy" is a set of dummy variables measuring the timing of the implementation of the new vaccine policy in each province. In Panel B, "Newspaper Shock" (or "Search Index Shock") is defined using the number of newspaper reports on vaccine-related issues (or the per-capita intensity of Baidu search for the keyword vaccine) in lieu of the number of relevant Weibo posts. "WeiboShock(prov.)" for a city is the Weibo Shock measure at the provincial level excluding the shock to that city. "WeiboShock(capital)" for a city is the Weibo Shock from the capital city of the province where the city is located. In both panels, full baseline controls include time-variant city characteristics as well as year-month fixed effects, city fixed effects, and provincial time trends. Standard errors (in parentheses) are clustered by city.

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

Figure B1: Robustness Check: General Procurement Practices

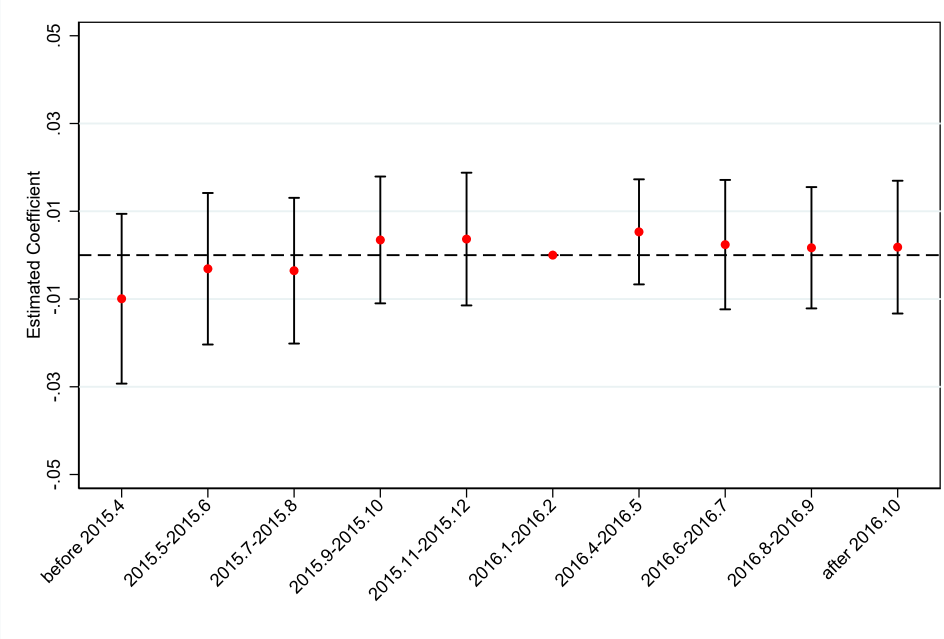

(a) Non-health products

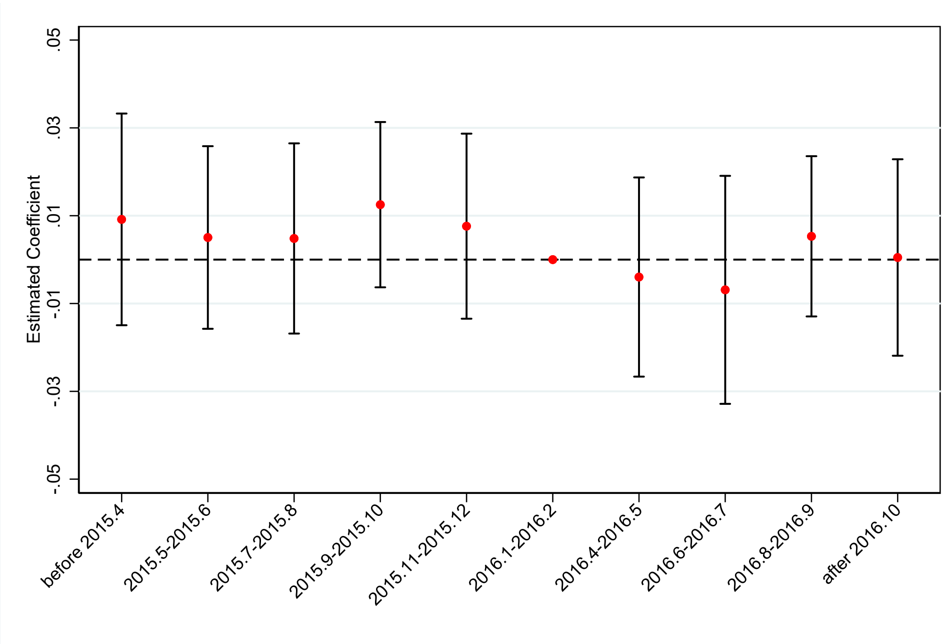

(b) Non-vaccine health products

Note: This figure plots the dynamics of the estimates of the Weibo effect on the share of open-bid procurement of non-health products (Panel (a)) and non-vaccine health products (Panel (b)), using the same specification as in our baseline estimation. The observations are at the city-month level, and the regression time window is from March 2015 to March 2017. Each dot represents the estimated coefficient of the interaction term between "WeiboShock" and a set of dummies that equals 1 if an observation falls within the specific time period and 0 otherwise.

C Details on Textual Analysis

C.1 Monitoring Posts

C.1.1 Sentiment Analysis

We apply sentiment analysis to all Weibo posts referencing "vaccine" within our sample period (2014-2019). Specifically, we employ HowNet, a pre-trained dictionary widely used in natural language processing in the Chinese context, to identify sentiment words. For sentiment calculation, we use a straightforward approach known as sentiment word count. The sentiment polarity of a given text is determined by the difference in the count of negative words and positive words. A higher sentiment score in a post indicates a stronger negative sentiment.

C.1.2 Support Vector Machine (SVM)

We apply a class of supervised machine-learning classification methods to distinguish monitoring posts from other posts. Initially, we obtain a random sample of 12,000 posts from our Weibo data, specifically from the year 2016, which was the focal year in our study.[25] A small sample of 1,000 Weibo posts was manually labeled to identify monitoring content, with a label of 1 being assigned to posts with monitoring content and 0 to others. Based on the experience gained from this pilot sample, we developed a comprehensive instruction for the labeling task. Subsequently, two research assistants were hired to independently label the remaining 11,000 Weibo posts in the random sample following the provided instructions. Approximately 87% of the posts received the same label from both research assistants. In cases where inconsistencies arose, we resolved them through careful reading and discussion. Ultimately, the labeled dataset consisted of 4,633 monitoring posts, which accounted for 38.61% of the entire sample.

We divide the labeled dataset into 10 folds, with 7 folds used for training and 3 folds for testing purposes. Pre-processing of the data involves segmenting the text and removing alphabets, numbers, punctuation, and other irrelevant characters. The corpus is then vectorized into a word frequency matrix. We apply various machine-learning algorithms, including Support Vector Machine (SVM), Naïve Bayes, Logistic Regression, Decision Tree, Random Forest, and AdaBoost, to train the classifier. As depicted in Figure C4, SVM demonstrates superior performance in terms of both recall and precision among the supervised machine learning approaches. This advantage of SVM has been observed in other classification tasks as well, particularly in the classification of short social media posts (Dumais et al. 1998; Joachims 1998; Sebastiani 2002).

Hence, we decide to use SVM as the baseline algorithm for classifying the monitoring posts. Figure C5a presents the confusion matrix, while Figure C5b displays the

25. We concentrate our post selection on 2016 to address the potential scarcity of monitoring posts during non-scandal periods.

Receiver Operating Characteristic (ROC) curve and the Precision-Recall curve. Both figures demonstrate that the SVM classifier exhibits a good fit and achieves a reasonably high level of classification accuracy. Subsequently, we apply the trained SVM classifier to predict the labels of posts in our entire dataset. These predicted monitoring posts are used to calculate the regional variation in the Weibo shock for our empirical analysis. As shown in Figure C8, at the city level, this measure of monitoring posts is strongly correlated with the negative-sentiment posts identified through our sentiment analysis.

C.1.3 Deep Learning (BERT)

We further employ the popular deep-learning model, Bidirectional Encoder Representations from Transformers (BERT), to classify the Weibo posts referencing vaccine into monitoring posts and others. To make the results comparable with those obtained by the above SVM method, we also divide the labeled dataset into the training and testing set with a ratio of 7:3. Figure C6a presents the corresponding confusion matrix and Figure C6b provides the Receiver Operating Characteristic (ROC) curve and the Precision-Recall curve for the evaluation of the BERT model. The classification accuracy with BERT is around 91%, outperforming the SVM classifier.

C.1.4 Large Language Model (LLM)

Finally, we employ large language models for posts classification. For the entire vaccine-related posts, we employ a GPT-type of large language model, "Qwen1.5:14b," developed by Alibaba for Chinese contexts, to classify the monitoring posts. The prompt for the task performance is designed specific to the vaccine scandal as shown in Figure C1:

Figure C1: Prompt for identifying monitoring implication in Weibo comments

Prompt: Monitoring Implication Classifier

In 2016, an illegal vaccine operation case occurred in Shandong Province. The illegally sold vaccines caused extremely serious social harm, triggering public dissatisfaction with government regulation. Below is a Weibo comment related to vaccines. Please determine whether this comment has a "monitoring implication" (where "monitoring implication" refers to a post expressing criticism of insufficient government regulation, hoping the government will take timely measures to crack down on businesses selling illegal vaccines and protect public life and health): `<text>`.

Output requirements: Please return a standardized JSON string in the following format:

- `{"supervisory nature": "yes"}`
- or `{"supervisory nature": "no"}`

To enhance the model performance, we adopt the LLaMA-Factory framework and use LoRA as the training algorithm to fine-tune the initial model. The fine-tuning leads to a high accuracy exceeding 93%, as shown in Figure C7.

To decompose Weibo posts into informational and sentimental content during the two event shocks (i.e., from December 2015 to March 2016, and from April 2018 to July 2018), we further employ a Gemini-type large language model to classify vaccine-related posts, using the prompt shown in Figure C2.

Figure C2: Prompt for informativeness and sentiment classification of Weibo posts

Prompt: Weibo Text Informativeness and Sentiment Classifier

Role: You are a Weibo (Chinese microblog) text classifier. You need to simultaneously judge two types of information: (1) whether the text contains verifiable/traceable information clues; (2) whether the text expresses strong negative sentiment, i.e., whether it is sentimental.

Core output definitions

- `informative`: `1` if clear verifiable facts/clues exist (judge cautiously); `0` if mainly emotional venting, moralizing, inference, or repost without verifiable content.
- `sentimental`: `1` if strong negative emotion is expressed (anger, fear, despair, condemnation, abuse, accusation); `0` if tone is neutral, factual, or only mildly expressive.

Informative determination rules: Judge `informative=1` if any of the following conditions is met:

1. The text explicitly cites/paraphrases specific content from an institutional or official source (department, official, original media text, announcement, link, etc.), not merely emotional expression;
2. The text provides clear traceable elements: region (province/city, etc.) + event type, or a specific time point, or vaccine manufacturer/brand/batch/distribution channel/handling action;
3. The text contains a clear factual statement: what happened, who did what, and what the specific reference is (rather than merely "fear/condemnation/anger");
4. The text raises a specific and verifiable inquiry (e.g., "domestic or imported" or "is the pentavalent vaccine only available as an imported product"), i.e., a question with a clear verification target.

Otherwise, judge `informative=0`.

Sentimental determination rules: Judge `sentimental=1` if any of the following conditions is met:

1. The text clearly expresses strong negative emotion, such as anger, fear, disgust, despair, hatred, breakdown, or intense anxiety;
2. The text uses strong condemnation, verbal abuse, cursing, sarcasm, or moral judgment;
3. The text uses noticeably emotional punctuation or tone, such as excessive exclamation marks, rhetorical questions, or extreme wording;
4. The primary purpose of the text is to vent negative emotion, rather than to state facts

neutrally or ask an ordinary question.

Otherwise, judge `sentimental=0`

Output requirements: Please return a standardized JSON string in the following format:

- `{"informative": 0 or 1}`
- `{"sentimental": 0 or 1}`

C.1.5 Monitoring Posts Targeting Governments and Firms

Among the monitoring posts identified by SVM, we further isolate a subset of monitoring posts that target government entities and those targeting firms or products using keyword matching. Specifically, the following two groups of keywords are used to identify Weibo posts targeting governments and firms, respectively:

Government-targeted posts: "Supervision", "Drug Administration", "General Administration", "Government", "State", "Central", "Department", "Official", "Disease Control", "Planning Commission", "Ministry", "State Council", "Leadership", "Investigation Team", "Ministry of Health", "Inspection", "Procuratorate".

Firm-targeted posts: "Enterprise", "Company", "Manufacturer", "Vendor", "Domestic", "Merchant", "State-owned Enterprise", "Pharmaceutical Company", "Unscrupulous Merchant", "Private Enterprise", "Manufacturing", "Production".

C.2 Posts Published by Government Weibo Users

C.2.1 Classification

We employ the same supervised machine-learning approach (SVM) to classify government posts as we did for the monitoring posts. We use two methods to label a selective sample of 12,000 posts from our vaccine dataset. The first method relies on user information, where a post is labeled as a government post if it is published by a user with a clear indication of government identity in their user profiles. The second method involves manual coding based on the content of the post, determining whether it discusses messages and announcements from the government. Two research assistants independently labeled the selected posts, achieving a high consistency rate of 94%. In the final dataset used for machine learning, 22.84% (3287 posts) were identified as government posts. This disproportionately large share of government posts is due to our intentional selection of the labeling sample, in which government posts are heavily overrepresented. In our entire data referencing vaccines, government posts account for only a small fraction. Using a randomly selected sample will lead to substantial asymmetries between the two classes in our classification exercise.

As before, we divide the labeled dataset into 10 folds, using seven folds for training and three folds for testing. SVM is employed once again for the classification of government posts, yielding a very high accuracy. Our estimated presence of government posts is approximately 4.81%, , which is comparable to Qin, Strömberg, and Wu (2017).

C.2.2 Topic Modeling

In Section 5.7.1 of the paper, we conduct an analysis of governments' blogging activities in response to the information outbreak. In order to gain a deeper understanding of the content of government posts, we use the Latent Dirichlet Allocation (LDA) topic modeling method.

In the baseline version of LDA, we opt for five topics to analyze the content of government posts. This choice is motivated by the simplicity and concentration of the posts, allowing for easier interpretation. Figure C9 presents the distribution of the top 10 words for each topic. The first line displays the Chinese words, the second line provides the English translation, and the third line shows the proportion of each word contributing to the topic. Topic 1 primarily revolves around words such as *work, carry out, preventive vaccination*, indicating a focus on government routine work. Topics 2 and 3 mainly consist of specific vaccine names and vaccination-related terms. On the other hand, Topics 4 and 5 exhibit distinct characteristics. Topic 4 includes words such as *problem, children (kids), parents, and flow*, suggesting a connection to vaccine safety. Topic 5 clearly relates to the government's response to vaccine issues, as evidenced by words like *illegal (business), legal cases, FDA, Bureau of General Administration, and announcement*. Based on the interpretation of the topic words, we classify Topics 1-3 as pertaining to government routine work and Topics 4-5 as relating to vaccine safety and government accountability. Since each post is associated with a probability distribution of topic weights, we assign a post to a specific topic based on the highest topic weight. For example, if Topic 1 carries the highest weight in a post, we classify the post as belonging to Topic 1. Following this post-topic assignment, we label posts under Topics 4 and 5 as "accountability posts" and those under Topics 1-3 as "routine-work posts."

In alternative versions of LDA, we explore different numbers of topics. Despite varying the number of topics to 10 or more, the distinction between "routine-work" and "accountability" topics remains clear.

References

Dumais, Susan, John Platt, David Heckerman, and Mehran Sahami. 1998. "Inductive Learning Algorithms and Representations for Text Categorization." *Proceedings of the Seventh International Conference on Information and Knowledge Management,* 148–155.

Joachims, Thorsten. 1998. *Making Large-Scale SVM Learning Practical.* Technical report.

Qin, Bei, David Strömberg, and Yanhui Wu. 2017. "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda." *Journal of Economic Perspectives* 31 (1): 117–140.

Sebastiani, Fabrizio. 2002. "Machine Learning in Automated Text Categorization." *ACM Computing Surveys* 34 (1): 1–47.

Figure C3: Machine Learning Procedure

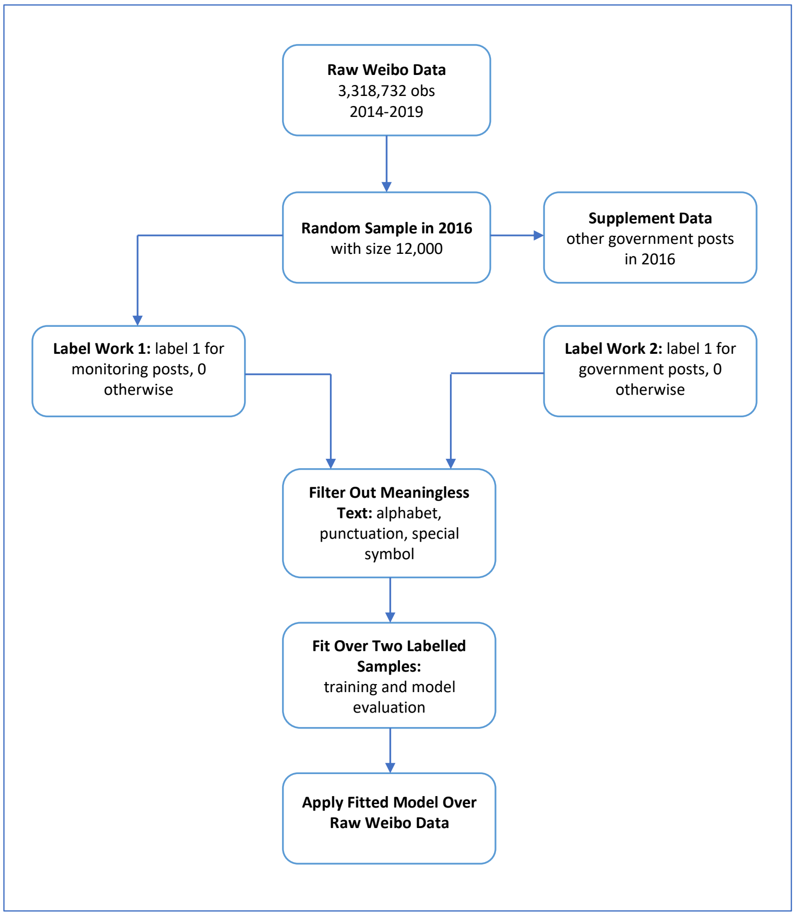

Figure C4: Comparison of Model Performance

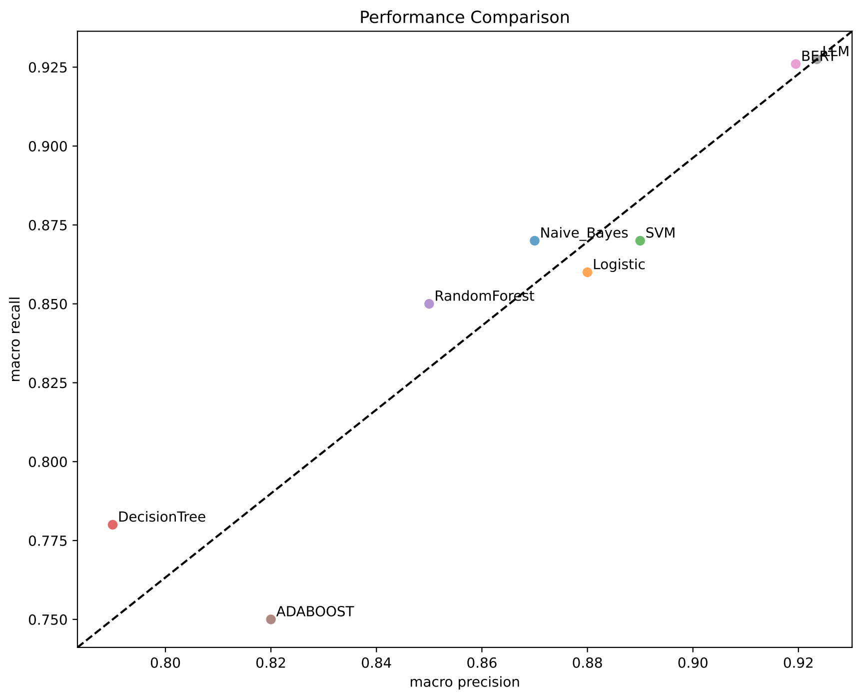

Figure C5: Performance Evaluation for Trained SVM Model

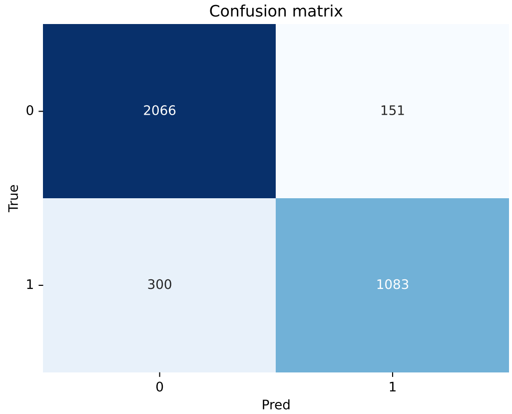

(a) Confusion Matrix (SVM)

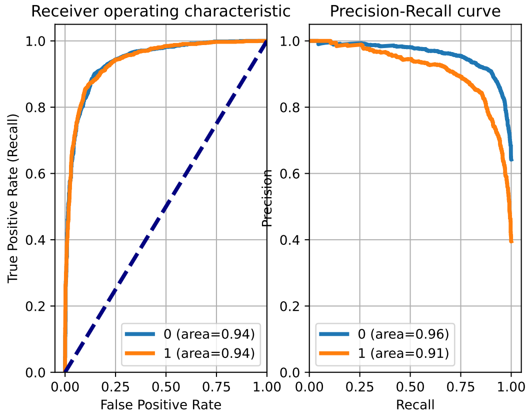

(b) ROC Curve and PR Curve (SVM)

Figure C6: Performance Evaluation for Trained BERT Model

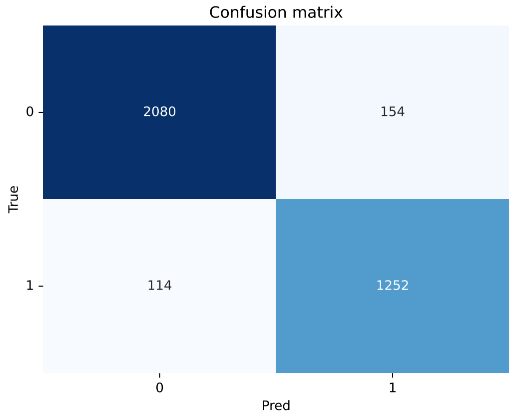

(a) Confusion Matrix (BERT)

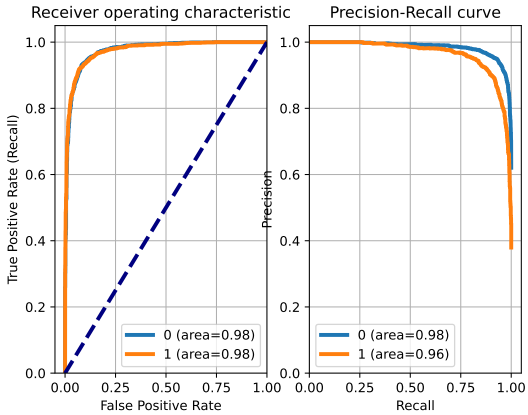

(b) ROC Curve and PR Curve (BERT)

Figure C7: Performance Evaluation for Large Language Model

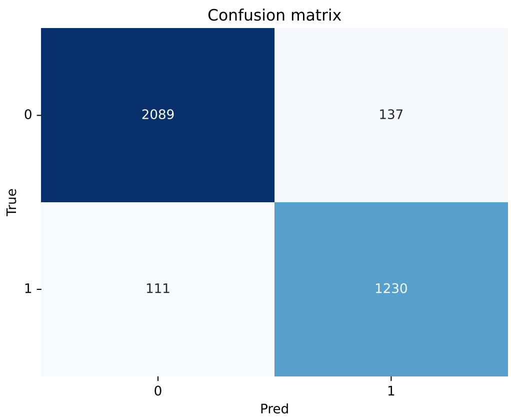

Figure C8: Correlation with Monitoring-Posts under Different Algorithms

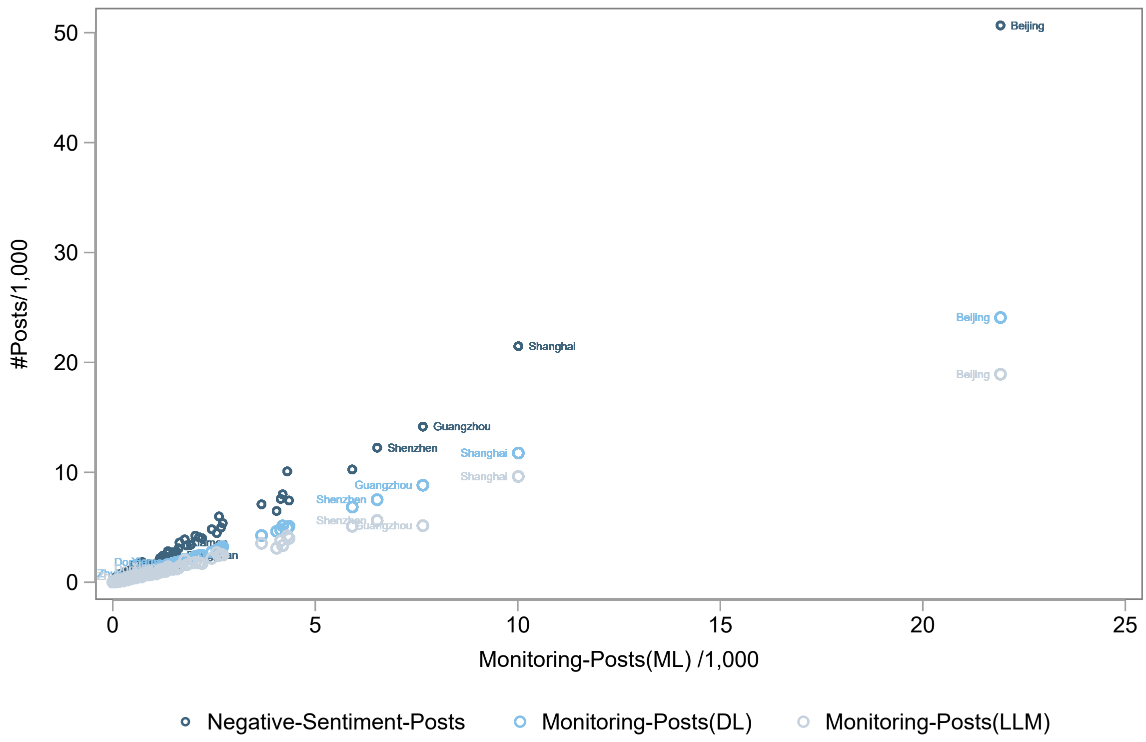

Figure C9: Words Distribution of Each Topic

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

D Pretrend Test

In this section, we describe how we test the parallel-trend assumptions, following several recent practices to further assess the validity of our DID estimation presented in the paper.

Bilinski-Hatfield one-step-up approach We assess whether there are different trends among sample cities, following the "one step up" approach proposed by Bilinski and Hatfield (2020). Specifically, we first estimate the period-wise post-treatment coefficients when the parallel-trend assumption holds for the pre-treatment coefficients. Applying the following specification (D.1), named as the "Reduced Model," we replicate the baseline estimation except that we interact the Weibo shock with dummy indicators for each post-event period.

$$y_{it} = \beta_0 + \sum_{j=T_0}^{T} \beta_j WeiboShock_i \times 1(t=j) + \lambda_i + \gamma_t + X_{it}^{'} + \epsilon_{it}, \tag{D.1}$$

Next, we modify the estimation by including a linear trend difference among cities as the following "Expanded Model" (specification (D.2)). If the parallel trend assumption holds, then the coefficient of the linear trend in (D.2) should be statistically insignificant. Alternatively, we can view (D.2) as the true model and (D.1) as a model with omitted-variable bias. If this bias is negligible, the magnitudes of the coefficients of interest should be similar in the two models.

$$y_{it} = \beta_0 + \sum_{j=T_0}^{T} \beta_j WeiboShock_i \times 1(t=j) + \theta WeiboShock_i \times t + \lambda_i + \gamma_t + X_{it}^{'} + \epsilon_{it}, \tag{D.2}$$

Figure D1 plots the estimation results with and without the linear trends. The averaged post-treatment coefficients from the two estimations are statistically indifferent with a P-value equal to 0.21, suggesting that the parallel-trends assumption holds.[26] On the right-hand side of the Figure D1, we plot the individual post-treatment coefficients for each period. The coefficients of interest from the reduced and the expanded specifications demonstrate similar post-event trends.

Roth linear violation approach We follow Roth (2022) to evaluate the effect of differential linear trends among cities on the estimated bias. We adopt a linear trend such that conventional pre-trends tests——reject the null whenever any of the pre-event coefficients are statistically significant at the 5% level——would have 50% to detect the pre-trends. In Figure D2, we plot the unconditional bias in green and the bias conditional on passing the pre-trend test in red. The baseline treatment effects and 95% confidence intervals are also presented in the figure for comparison. All the values in the figures

26. Statistically, it is equal to testing the null hypothesis $H_0 : \beta - \beta' = 0$ and testing $H_0 : \theta = 0$.

are normalized by the standard error of the estimated treatment effect. The results in Figure D2 indicate that the magnitudes of the conditional and unconditional biases are relatively small, and the treatment effects cannot be primarily driven by the linear violation of pre-trends, either using procurement months as periods or using procurement sequences as periods. The bias caused by the linear violation, relative to the treatment effects, appears negligible in the first and second periods after the event.

We also visualize the influence of linear violation by extrapolating the pre-treatment trends and presenting the counterfactual post-treatment effects. The extrapolated trends (in red), as well as the corresponding coefficients (in dashed blue) under such linear violation, are shown in Figure D3. The slope of the linear trends is 0.04 based on the power of 50% to detect pre-trends. The resulting plots show that the economically plausible violation of parallel trends cannot fully explain the pattern in the conventional event study. Consistent with Figure D2, the social media effects, when accounting for the linear trend, primarily come from the first two periods (four months) after the event.

References

Bilinski, Alyssa, and Laura A. Hatfield. 2020. "Nothing to See Here? Non-inferiority Approaches to Parallel Trends and Other Model Assumptions." *Working Paper.*

Roth, Jonathan. 2022. "Pre-Test with Caution: Event-study Estimates After Testing for Parallel Trends." *American Economic Journal: Insight* 4 (3): 305–322.

Figure D1: Post-event Estimated Coefficient with Pretrend

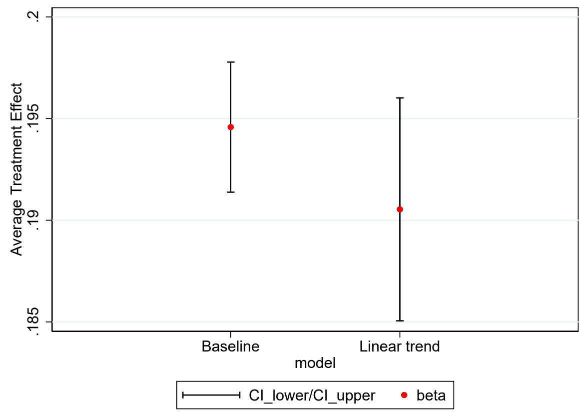

(a) average effect

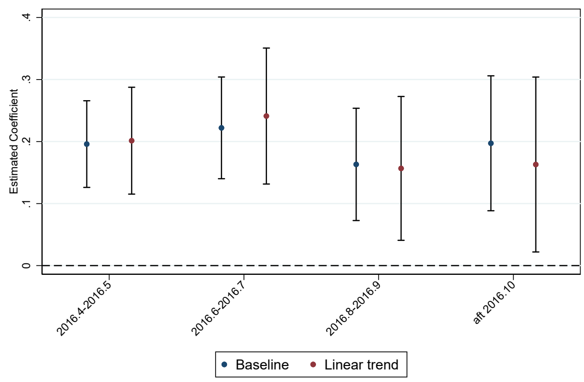

(b) period by period

Figure D2: Evaluation of Estimation Bias

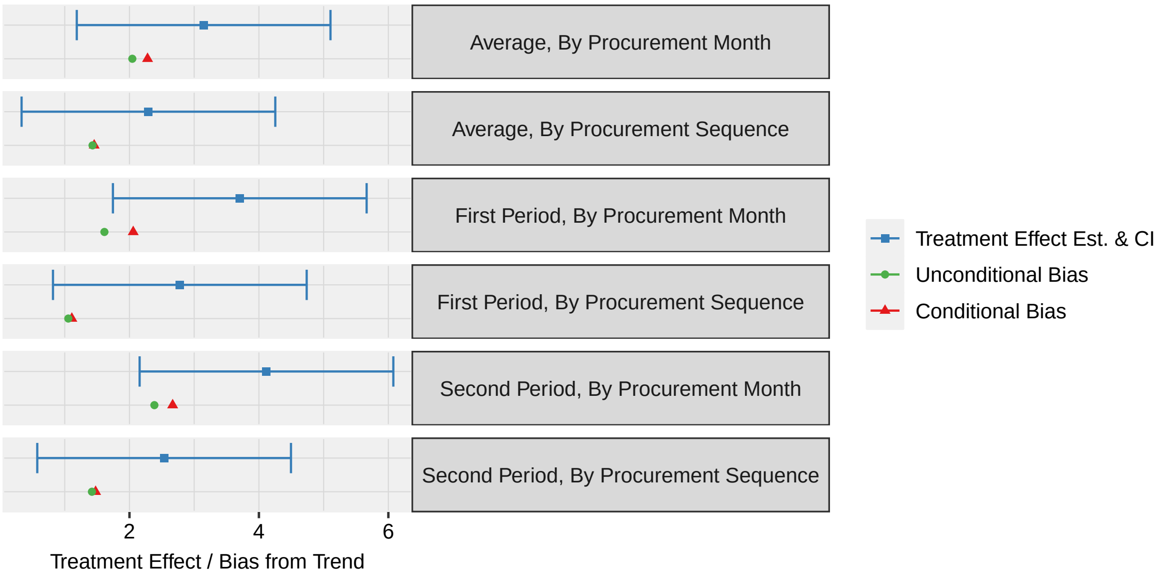

Figure D3: Dynamic Effects of the DID Estimation: Linear Violation

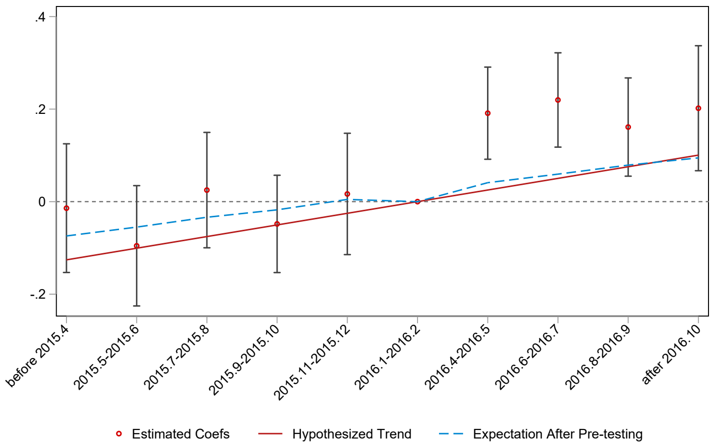

(a) by procurement month

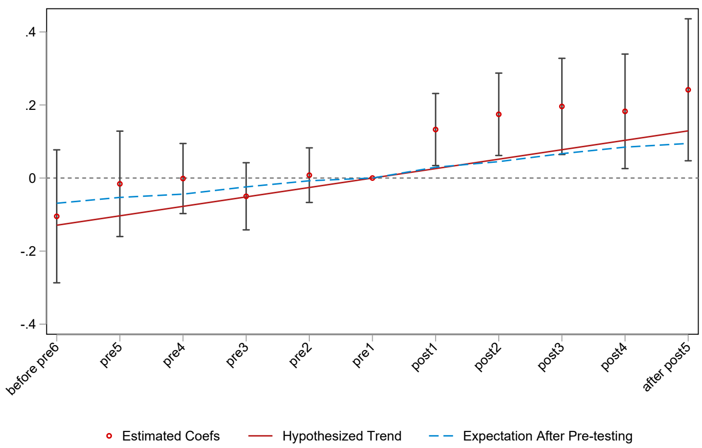

(b) by procurement sequence